USING LARGE-SCALE ADMINISTRATIVE TAX DATA TO ASSESS HOUSING AND RETIREMENT POLICIES

A Dissertation
submitted to the Faculty of the
Graduate School of Arts and Sciences
of Georgetown University
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy
in Economics

By

Elena Derby, M.A.

Washington, DC
April 22, 2020

USING LARGE-SCALE ADMINISTRATIVE TAX DATA TO ASSESS HOUSING AND RETIREMENT POLICIES

Elena Derby, M.A.

Thesis Advisor: Arik Levinson, Ph.D.

ABSTRACT

In this dissertation I use administrative tax data to evaluate two policies in the fields of low income housing and retirement. In particular, I construct novel data sets to assess whether growing up in housing constructed under the Low Income Housing Tax Credit leads to higher educational attainment and earnings later in life, and to measure the effect of automatic enrollment in firm-sponsored retirement plans on individual and household savings. On the topic of housing, I find that each additional year spent in LIHTC housing as a child is associated with an average 3.9 percent increase in the likelihood of attending a higher education program for four years or more, and a 5.2 percent increase in future earnings. I conclude that the reason I see a positive LIHTC effect is likely that the housing subsidy provides families with a more stable living situation and with more disposable income. On the topic of retirement, I estimate an increase in plan participation of 32.4 percentage points due to automatic enrollment, and a more modest increase in average saving of 0.9 percentage points, with a negative nudge towards lower default rates for employees who opted in prior to the policy change. However, I also find that automatic enrollment leads to a 36.1 percent increase in withdrawals taken from other retirement accounts, and that this completely offsets the increase in contributions from firm-sponsored retirement plans for employees in the bottom wage quintile, although the net effects remain positive for employees with higher wages.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: DOES GROWING UP IN TAX-SUBSIDIZED HOUSING LEAD TO HIGHER EARNINGS AND EDUCATIONAL ATTAINMENT?

The Low Income Housing Tax Credit (LIHTC) is the largest federal subsidy for the construction of low income housing in the United States. Since its establishment through the Tax Reform Act of 1986, the LIHTC has helped finance the construction and renovation of over three million units, and currently costs over $9 billion per year in forgone tax revenue (JCT, 2018; HUD, 2019). Yet despite the size and importance of the LIHTC we still know relatively little about the people who reside in LIHTC buildings; nor do we know whether access to subsidized housing funded by the LIHTC improves people's lives in measurable ways. The main reason for this is the lack of data available on LIHTC residents and their outcomes. The tax credits are issued directly to developers, and these developers are not required to track or report information on the tenants residing in their buildings.

In this paper I addresses this problem by using a newly created data set of families residing in qualifying LIHTC properties to evaluate the long-run effects of growing up in LIHTC housing. Specifically, I use administrative tax records to estimate whether individuals who grow up in LIHTC housing have higher wages as adults than they would otherwise, and whether they are more likely to enroll in higher education programs. Using administrative tax records, I create a database of families with children under the age of 18 who lived in LIHTC housing between 1999 and 2012. I identify families living in LIHTC housing during these years using the publicly available addresses of LIHTC buildings, matched with parents' addresses listed on their information returns (such as W-2 and 1099 forms), and tax returns. I match parents and children based on the parents' and the children's ages, and whether the children are listed as dependents on their parents' tax returns. I then use the children's Social Security numbers to find their adult wages in 2018, when they are between ages 24 and 36. I also observe how many years they enroll in higher education programs using 1098-T tuition statements, which are filed by American colleges and universities.

In order to estimate the effects of living in LIHTC housing, I exploit variation in the number of years that individuals spend living in these buildings as children. This "dosage" measure of exposure to LIHTC housing helps address potential biases that may arise from comparing the outcomes of individuals whose families made the decision to move into LIHTC housing to those who did not live in subsidized low income housing during the same time period (or who lived in public housing). By comparing individuals whose families all decided to move into LIHTC housing when they were different ages, I eliminate potential bias from omitted unobserved variables such as parents' motivation to secure affordable housing for their families. However, there may still exist some bias associated with circumstances leading families to leave LIHTC housing, and the timing of families moving into LIHTC housing. To address these issues I use the stratified sampling procedure outlined below.

The first step in my stratified sampling approach is to estimate the effect of spending one additional year in LIHTC housing as a weighted average across groups of individuals who left LIHTC housing at different ages. This should eliminate any bias arising from both negative and positive circumstances that could cause families to leave LIHTC housing, such as eviction or moving in with a family member. This strategy also helps deal with differences in the effect of spending the same amount of time in LIHTC housing at different ages. For example, an individual who lives in LIHTC housing from ages 5 to 10 may experience a different effect than someone who lives in LIHTC housing from ages 13 to 18, despite spending the same number of years living in a LIHTC building. Instead, all variation comes from the age of entry into LIHTC housing.

Second, I further limit my sample population to include only individuals who move into LIHTC housing the same year or one year after the building is placed in service. This helps deal with bias that may arise from families strategically moving into LIHTC housing when their child is a specific age. For example, a family may decide to move into a better school district when their child is 11 years old and about to enter middle school. Using this stratified sampling method, even if parents target their move into a LIHTC building based on the age of their child, it will be purely coincidental that they enter into a building that was put into commission that same year, or the year prior. This approach assumes that parents who want to move into subsidized housing when their child is a particular age do not also favor moving into a newly constructed building. I assume that parents are concerned with moving into a subsidized unit, possibly timing the move according to the age of their child, and they do not care whether the building they move into was built in the last two years (as opposed to buildings that were constructed in the last three years or more).

Using this strategy, I find that spending a longer amount of time growing up in LIHTC housing has a positive and statistically significant effect on both earnings and education. Under my preferred specification I find that *for every additional year* spent in LIHTC housing as a child, individuals are 3.9 percent more likely to enroll in a higher education program for four or more years, and are 3.7 percent more likely to enroll in two or more years of higher education. Additionally, individuals earn approximately 5.2 percent more as adults for every additional year spent in LIHTC housing. The cumulative effect is large. For example, I find that individuals who stay in LIHTC housing for seven years are 28.0 percent more likely to attend a college, university, or trade school for four or more years than those who live in LIHTC housing for just one year. Those who stay for seven years also earn 24.7 percent more on average in income than their counterparts who stay for one year.

In the following sections I describe these findings in greater detail and explore possible mechanisms that may be driving these results. In particular, I examine differences in the effect of living in LIHTC housing located in neighborhoods with different characteristics, including varying poverty levels, racial composition, median household incomes, high school graduation rates, and measures of opportunity constructed by Chetty, et al. (2018). I also look at heterogeneous effects based on the level of housing security a family has prior to entering into LIHTC housing – measured by the number of times that an individual changes addresses prior to moving into a LIHTC building – and heterogeneous effects based on household income. I find that although there are differences in the "LIHTC effect" based on both neighborhood and family characteristics, the latter appears to play a more important role in explaining my results.

1.1 Background

Congress passed the Low Income Housing Tax Credit (LIHTC) as part of the 1986 Tax Reform Act, touting the credit as an incentive to encourage private investment in the development and rehabilitation of low income housing in the United States (Keightley, 2019). During this time period the number of households living in LIHTC units increased from approximately 400,000 in 1993 to roughly 2.6 million households in 2016, nearly equaling the number of households served by all other public housing programs combined (Kingsley, 2017). Yet despite the size and importance of the LIHTC program, we know relatively little about how the tax credit affects the low income renters it intends to serve.

Unlike other federal low income housing subsidies, the LIHTC is issued directly to developers – as opposed to individual renters or homeowners – and is used solely to help finance the construction or renovation of low income housing. The Internal Revenue Service (IRS) and Department of Treasury administer the tax credit at the federal level, but selection of projects eligible for the credit, and administration and oversight of the program, are the responsibilities of state and city governments. Each year the federal government allocates tax credits to each state's Housing Finance Agency (HFA) based on the state's population – or the city's HFA if it is a "home rule" city – with the minimum credit level set at $3.1 million as of fiscal year 2018 (Novogradac, 2018). The HFAs then award tax credits to developers based on federal guidelines, and on criteria that each state or city sets in its Qualified Action Plan (QAP) (Ellen, et al., 2015).

There are two types of Low Income Housing Tax Credits: nine percent credits, which are capped based on state population, and four percent credits, which are uncapped and typically used to help fund renovation projects, or new construction projects financed with tax-exempt bonds (Keightley, 2019). Each credit is awarded for a period of 10 years, and designated LIHTC units must remain affordable to low income renters – as defined by federal regulations – for a period of 15 years, with additional rent restrictions for a three to 15 year period afterward (Black, 2014; Baum-Snow and Marion, 2009). The recipient of the credit can claim a percentage (either nine or four percent depending on the type of credit) of the building's qualified basis each year. However, the applicable rates are typically lower, as they fluctuate with interest rates. In practice, the total subsidy is equivalent to approximately 30 percent of the initial value of the project's qualified basis in the case of the four percent credit, and 70 percent of the qualified basis in the case of the nine percent credit (Keightley, 2019). The term "qualified basis" has a lengthy legal definition, but is essentially understood to be the total cost of construction or renovation that will ultimately benefit low income residents.

Developers apply to receive credits by proposing plans directly to state HFAs, who approve a set number of plans each year to receive nine percent credits (with total allocated credits falling below the annual limit), and an unrestricted number of plans to receive four percent credits. Since approval is decided at the state level, there are considerable differences between states when it comes to the location of LIHTC housing, which are driven by the states' Qualified Action Plans (QAPs). Some state QAPs promote economic and racial diversification, so proposals for LIHTC buildings in more affluent or mixed-income neighborhoods are given priority. Other state QAPs place greater emphasis on economic development of impoverished communities, prioritizing proposals for buildings in low income neighborhoods (Ellen, et al., 2015). In general, although LIHTC developments are typically located in neighborhoods with higher poverty rates compared to national averages, they are also located in relatively more affluent neighborhoods compared to project-based public housing (McClure, 2006). A larger percentage of LIHTC buildings are also built in suburban neighborhoods, compared to public housing projects.

Developers (or more commonly investors) cannot claim the LIHTC on their federal returns until their project is completed and occupied by tenants (Keightley, 2019). In order to fully qualify for the credit, developers must prove that at least 20 percent of units in their building are occupied by families or individuals with incomes lower than 50 percent of the area median income (AMI), or at least 40 percent of the units are occupied by families or individuals with incomes lower than 60 percent of the AMI, with slightly lower or higher income thresholds depending on family size (Keightley, 2019). This is commonly referred to as the 20-50 test, or the 40-60 test. Developers must also prove that the rent level for these units is set no higher than 30 percent of the designated income threshold for the building (also adjusted based on family size). For example, in a building that qualifies based on the 20-50 test, rent for a 4-person family can be set no higher than 30 percent of 50 percent of the area median income (or 15 percent of the AMI).

Since the LIHTC cannot be claimed until buildings are occupied, developers typically sell these tax credits to third party investors – usually banks and other financial institutions – for equity to fund the construction or renovation of their buildings (Keightley, 2019). These partnerships between developers and investors are sometimes brokered by syndicators, who charge an additional fee for processing the transactions (GAO, 2019). On top of benefiting from a generous tax credit, banks in particular have an incentive to purchase Low Income Housing Tax Credits in order to meet federal requirements for local investment under the Community Reinvestment Act of 1977 (GAO, 2019).

One of the largest differences between the LIHTC and other federally funded housing programs is that income limits are set according to characteristics of the building, unit, and area; not according to the income of the tenant (O'Regan, et al., 2013). Thus, unlike with housing vouchers (for example), families living in LIHTC housing often pay more than 30 percent of their household income towards rent, a rate that is commonly considered the "rent burden" threshold. Yet a 2012 study by the Furman Center at New York University found that over 45 percent of LIHTC residents are considered extremely low income, meaning that their income is less than 30 percent of the area median income (O'Regan, et al., 2013; Hollar, 2014). These people can afford to live in LIHTC housing in large part because they often also receive additional housing subsidies – both directly through housing vouchers and indirectly through building- or unit-specific subsidies.

The Furman Center found that over one-third of LIHTC households receive some other form of rental assistance, including Section 8 housing vouchers (O'Regan, et al., 2013). This is especially true of extremely low income tenants, 70 percent of whom receive some other type of rent subsidy in addition to lower LIHTC rents (NYU, 2012). One of the main reasons why a lot of Section 8 voucher recipients reside in LIHTC buildings is that voucher holders face considerable housing discrimination in unregulated markets, especially in high or middle income neighborhoods (Emple, 2014). Voucher holders often report being ignored by landlords, or turned away when applying for apartments that are not designated as low income units. In contrast, LIHTC buildings need these tenants in order to qualify for the tax credits and are less likely to turn them away. However, it is also the case that LIHTC developments themselves receive a number of additional federal and local subsidies, and in return these buildings charge their tenants lower rents (O'Regan, et al., 2013).

This overlap between the LIHTC and other housing subsidies does make it difficult to separate the effect of living in LIHTC housing from that of receiving a Section 8 voucher, for example. However, these programs can be considered complimentary in many ways, since many low income renters could not afford to live in LIHTC buildings without other housing subsidies, and many voucher recipients are not able to find housing in non-LIHTC buildings. My identification strategy does help isolate the effect of LIHTC housing by restricting my sample to families that move into new LIHTC buildings, essentially measuring the effect of new subsidized housing being introduced to a given area. However, in general I interpret my results as the effect of a combination of housing subsidies, since the estimate would likely change in the absence of having extremely low income renters comprising almost half of all LIHTC residents.

1.1.1 Housing Policy and Economic Mobility

There are several reasons why growing up in LIHTC housing might lead to higher educational attainment and higher earnings later in life. First, if LIHTC housing is built in a more affluent neighborhood where low income families could not otherwise afford to rent an apartment or house, then the children growing up in these buildings might benefit from attending better, well-funded schools, with access to resources and social networks that they would not otherwise have. Previous studies on peer effects show that organizational norms and structures at the high school level – as well as social expectations around college attendance – can have a large effect on the probability of individual students attending college, especially for low income students (Sokatch, 2006; Roderick, et al., 2011; Falk and Ichino, 2006).

The most famous experiment testing this theory of location effects on adult outcomes is the US Department of Housing and Urban Development's (HUD's) decade-long Moving to Opportunity (MTO) program, which relocated families living in high poverty neighborhoods to low poverty neighborhoods through a restricted voucher system (Ludwig, et al., 2013; Leventhal and Brooks-Gunn, 2003; Sanbonmatsu, et al., 2011). The MTO experiment aimed to fix what HUD perceived as a problem with the Section 8 voucher program. In theory, families who received these vouchers had unlimited choice when it came to which neighborhood they could move to. However, the majority of tenants who received Section 8 vouchers remained in high poverty neighborhoods, both because of personal preference – choosing to remain in a familiar neighborhood where they felt comfortable – and as a result of discriminatory housing practices, with landlords refusing to accept vouchers as a form of rent payment.

As an alternative, MTO required random voucher recipients to move into low poverty neighborhoods to test the theory that housing policy could go beyond simply providing people with a cheaper place to live, and could help improve the lives of tenants by moving them into areas with better opportunities for education and employment. HUD created local programs in each of the test cities (Baltimore, Boston, Chicago, Los Angeles, and New York) and entered into agreements with local non-profit organizations to provide counseling and assistance to participating families in order to help them navigate different housing markets (Gennetian, et al.,2011). In addition, they worked directly with landlords to encourage their participation in the program and reduce the incidence of discrimination against voucher holders.

The initial results of the MTO experiment were mixed. The data did show improved mental and physical health outcomes for children who moved to low poverty neighborhoods, such as lower obesity and diabetes rates (Ludwig, et al., 2013; Leventhal and Brooks-Gunn, 2003; Sanbonmatsu, et al., 2011). The study also concluded that outcomes for girls who moved into low poverty neighborhoods were generally better than they were for boys. However, the experiment revealed no significant effect of moving to a low poverty neighborhood on students' test scores, or on their future employment and wages (Ludwig, et al., 2013). Thus, the study concluded that moving to a high opportunity neighborhood had no discernible impact on individuals' long-run outcomes.

However, a few years after publication of the initial MTO findings, Chetty, Hendren, and Katz (2016) re-evaluated the MTO data and found that moving to a low poverty neighborhood actually did lead to improved outcomes, particularly for individuals who moved into these areas before age 13. The reason the initial MTO study missed this pattern is that HUD collected its final round of outcomes data in 2008, when most individuals who participated in the program were not yet old enough to enter the labor market. Furthermore, the participating individuals who had entered the labor market by 2008 had moved into low income neighborhoods at an older age. So the original conclusions of MTO were based solely on the outcomes of children who had low "exposure" to low poverty neighborhoods, and who had moved at an age at which the negative disruption in their schooling and personal life may have outweighed the benefits of moving to a more prosperous area.

In general, Chetty, et al. (2016) find that the neighborhood in which a person grows up has a large effect on their outcomes later in life, and that "high opportunity" neighborhoods are not randomly scattered across the country. They are clustered in very specific geographic areas, with significant differences even across adjacent Census tracts (Chetty, et al.,2016). The more time children from low income families spent in high opportunity neighborhoods – neighborhoods where individuals eventually ended up earning more than other people in their age group on average – the more likely they were to end up in a higher income bracket than their counterparts in low opportunity neighborhoods.

Another reason why living in LIHTC housing may lead to higher earnings and educational achievement may be that the construction of LIHTC housing has a revitalizing effect on low income neighborhoods themselves. In their recent paper, Rebecca Diamond and Tim McQuade (2019) find that the construction of LIHTC housing in low income areas leads to an increase in housing prices, a decrease in crime rates, and greater racial and economic diversification in these neighborhoods. Baum-Snow and Marion (2009) also find that the construction of LIHTC buildings leads to an increase in property values in declining areas. In the same way that building LIHTC housing in a high income neighborhood can lead to better resources and opportunities for children growing up in these buildings, improvements to low income neighborhoods may have similar positive effects on schools and available resources in these communities. Thus, children growing up in LIHTC buildings constructed in high poverty areas may benefit from a general improvement in the neighborhood where they are constructed, because the buildings themselves bring about changes that affect childhood development.

There are several other possible reasons why moving into subsidized housing might be beneficial to low income families, and why it might lead to better outcomes later in life for the children who move in at younger ages. Spending less income on rent may give parents more disposable income to spend on other resources for their children. They might be able to afford more nutritious food, or pay for after school programs and tutors, all of which would benefit their children in the long run. Entry into LIHTC housing may also alleviate homelessness for some families, or provide a more stable living situation, especially if the family was moving around a lot prior to securing low income housing in a LIHTC building. LIHTC housing may also provide a safer environment for families than their previous living situation.

This paper examines three of these possible explanations for why growing up in LIHTC housing may lead to better outcomes: location, housing stability, and household income. To do this I decompose the effect by both neighborhood and household characteristics to determine if the effect differs between children growing up in different neighborhoods or in families facing different financial and housing circumstances. The results from these two analyses suggest that the positive estimated effect of growing up in LIHTC housing is more the result of improved housing stability and higher net household earnings than of moving to a better neighborhood. However, as I explain in greater detail in the following sections, I cannot rule out the importance of location since most families do not move to a higher opportunity neighborhood when they move into LIHTC housing.

1.2 Data

Most of my data come from the population of tax and information returns in the United States, including tax returns filed by individuals, and information returns issued by employers to employees and contractors, such as W-2 and 1099 forms. Using these data I am able to track families' addresses over time, link parents with their children (listed as dependents on tax returns), find earnings information for both the parents and the adult children, and ascertain children's educational achievement using 1098-T tuition statements. These records are limited in some ways. I only have access to data starting in 1999, and I am not able to track families for the oldest cohorts for more than a few years. Some data are also missing from administrative tax records. In particular I find that when it comes to tax returns (as opposed to information returns) the tax years associated with a lot of addresses are missing, which makes it difficult to determine which years each tax filer lived at each address. This seems to be particularly common in earlier years as there is a jump in the number of addresses listed on tax returns filed after 2003. As such, I rely heavily on information returns to locate parents prior to this year.

My second main source of data is the US Department of Housing and Urban Development (HUD) database for buildings constructed under the Low Income Housing Tax Credit (HUD, 2019). HUD provides a wide range of information on each building, including the building's address and zip code, the year it was placed in service, the total number of units in each building, the number of low income units in the building, the income threshold rule (20-50 or 40-60), the Census tract number (in both 2000 and 2010) where the building is located, county and state identifiers, and cost information related to the tax credit issued for the building. There are some problems with these data as well, particularly with missing or partially available addresses (for example, a few buildings provide just a zip code). There are some missing data in other fields as well, like the number of LIHTC units in the building compared with the overall number of units. As such, I do exclude a small number (about 6 percent) of buildings due to missing or incomplete data.

In order to identify families living in LIHTC housing, I match addresses of LIHTC buildings with addresses listed either on parents' information returns or their individual tax returns. I provide a detailed explanation of the matching process I use to identify families in **Appendix A**. The fully merged data set includes 540,839 individuals born between 1982 and 1994 whom I observe living in LIHTC housing between 2000 and 2012. The population includes only dependents that I identify as living in LIHTC units based on their parents' income and on the income thresholds unique to each building. I exclude anyone I observe living in LIHTC housing in 1999 because I cannot tell how long they lived in the building prior to that year. As such, all of the individuals in my data moved into LIHTC housing in 2000 or later.

I observe these individuals moving into LIHTC housing between ages 6 and 18, and approximately one third of them remain in LIHTC housing through age 18. They are between ages 24 and 36 in 2018, when I observe their annual earnings. **Figure 1** below shows the distribution of ages and amounts of time spent in LIHTC housing in the full data set. As shown in the graphs, there are a greater number of individuals in younger cohorts simply because the data only go as far back as 1999, so I observe the families of individuals born in earlier years for a shorter period of time. Since I observe older individuals for a shorter amount of time, the number of years spent in LIHTC housing also tends to be skewed towards the lower end.

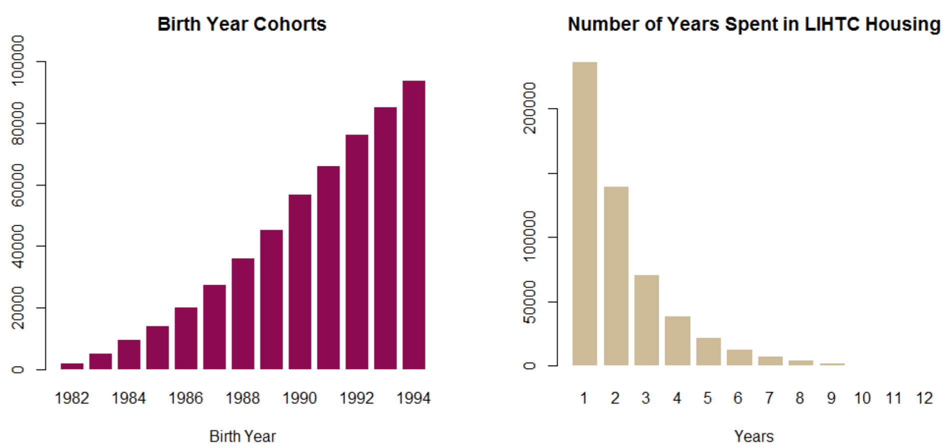

Figure 1: Summary Statistics

Table 1 below provides summary statistics for each variable in the data including LIHTC years (the number of years spent growing up in LIHTC housing), household income (in 2018 US dollars, adjusted for inflation), birth year, gender, family size, the parent's age at birth, total moves (the number of times an individual moved over the period of time I observe them), and number of units in each LIHTC building. The table includes mean, median, and standard deviation values. A full description of each variable is provided in **Appendix B**. One thing to note in the table is that the standard deviation of household income is quite high, but this is reflective of large variations in area median income (and consequently large variation in the rent threshold) between different counties or cities.

Table 1: Summary Statistics: Low Income Housing Tax Credit Data

Variable	**Median**	**Mean**	**Standard Deviation**
LIHTC Years	2.0	2.3	1.7
Birth Year	1991	1991	2.8
Household Income	\$22,630	\$25,999	\$55,111
Family Size	2.0	1.8	0.9
Parent's Age at Birth	24	25.4	8.1
Total Moves	4.0	4.0	2.6
Number of Units in Building	108	134	103.6
Percent Female		49.9%	

To further examine the mechanisms that may be underlying the effect of growing up in LIHTC housing I also use Census tract data from the US Census Bureau's 2000 Census of Population and Housing to estimate differences in the effect for individuals growing up in LIHTC housing located in neighborhoods with varying characteristics (Census, 2000). I use these data to analyze differences in the estimated effect for buildings constructed in neighborhoods with varying levels of poverty, high school graduation rates, racial composition, and median household incomes. In the following section I provide further details on how I divide my sample into different quantiles based on neighborhood characteristics in order to conduct this analysis.

In addition, I use data from Chetty, et al.'s (2018) Opportunity Atlas (`www.opportunityatlas.org/`) to look at how the LIHTC effect varies among neighborhoods with different opportunity measures. The opportunity measure I use is the fraction of children who grew up in a specific area whose household income in 2014/2015 (when they are in their mid-30s) is in the top 20 percent of the national income distribution for children born in the same year. It is a measure of the average differences in the outcomes of individuals who grow up in different neighborhoods (Chetty, et al., 2018). In other words, the opportunity measure tells us how likely it is that a child who grows up in a particular area will be in a relatively higher or lower income percentile than their peers growing up in other neighborhoods. This provides a more dynamic measure of poverty than more traditional measures that show a snapshot of neighborhood conditions.

Looking at the raw data, there does appear to be a small positive correlation between the number of years spent in LIHTC housing ("LIHTC Years") and the likelihood of enrolling in a higher education program, without controlling for other variables. However, there does not appear to be as strong of a relationship between earnings and number of years spent in LIHTC housing. The first graph in **Figure 2** below shows the percent of individuals who enrolled in four or two years of post-secondary education, according to the number of years each group spent in LIHTC housing. The second graph shows 2018 earnings at the 25^{th}, 50^{th} and 75^{th} percentiles by group. These raw patterns in the data do not control for several factors that may introduce some bias in the effect of growing up in LIHTC housing on adult outcomes. In the following section I further describe my approach to estimating this effect, which turns out to be larger and more statistically significant than the graphs in **Figure 2** might suggest.

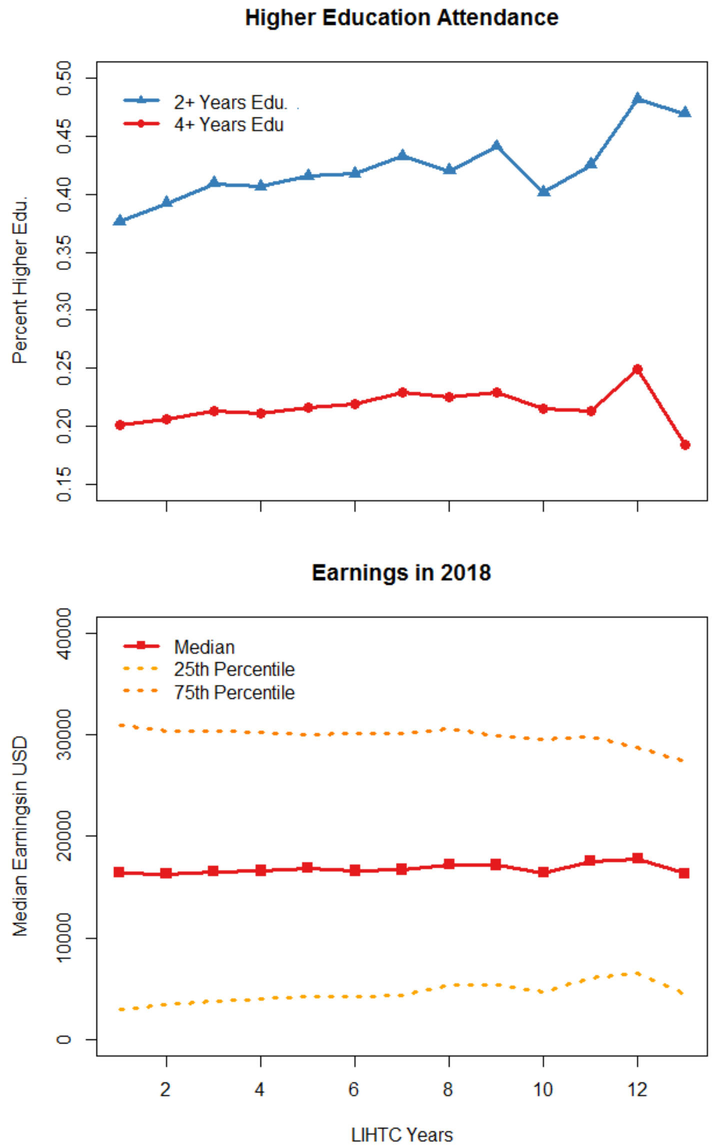

Figure 2: Educational Attainment and Earnings by LIHTC Years

1.3 Empirical Specification

The focus of this paper is estimating the effect of growing up in LIHTC housing on educational achievement and adult earnings. As such, it may seem natural to compare children growing up in LIHTC housing with similar individuals growing up in non-LIHTC housing. However, this comparison may not be appropriate if parents who secure housing in a LIHTC building are in some way inherently different from parents who do not live in LIHTC housing. In particular, it would not be possible to control for unobservable characteristics like parental ambition, which might affect both the parent's ability and willingness to find a rent-subsidized unit in a LIHTC building, and would also likely affect their children's future outcomes. If parents who rent units in LIHTC buildings are more proactive about securing housing for their families, then they may be proactive in other ways that would help their children succeed, such as enrolling them in free after school programs.

Rather than compare individuals who grow up in LIHTC housing with those who do not, I instead compare individuals who spend different amounts of time in LIHTC housing as children with each other. Instead of estimating the overall effect of growing up in LIHTC housing I measure the effect of spending *one additional year* growing up in LIHTC housing. My main variable of interest, which I call "LIHTC Years" can be considered a kind of dosage or exposure variable, where individuals who spend a longer amount of time living in LIHTC housing as children are considered to have greater exposure or a larger dose of LIHTC housing than their counterparts who spend just one or two years living in a LIHTC building.

I estimate the effect using parametric regressions of adult earnings and educational attainment on the number of years spent growing up in LIHTC housing, controlling for a number of individual, family, and location characteristics. I estimate the effect of growing up in LIHTC housing, $\hat{\theta}$, using the following regression, where $y_{i,b,z}$ is the outcome variable (either educational attainment or adult earnings) for individual i with birth year b growing up in LIHTC housing located in zip code z, regressed on LIHTC years $h_{i,b,z}$, and a vector of control variables $X_{i,b,z}$, with birth year and zip code fixed effects, γ_b and δ_z. Standard errors are robust and clustered at the zip code level.

$$y_{i,b,z} = \theta h_{i,b,z} + \beta X_{i,b,z} + \gamma_b + \delta_z + \varepsilon_{i,b,z}$$

I use a logistic regression for both of my binary education outcome variables (the incidence of attending four or more years of higher education and the incidence of attending two or more years of higher education), and I use an ordinary least squares regression for my earnings outcome variable. I control for a number of individual, household and building characteristics including household income (for all years I am able to observe the family before the child turns 18 years old, from 1999 to 2012), gender, parents' ages relative to their children, area median income, family size, the number of units in each building, and parents' filing status (e.g. married or single). Each variable is explained in greater detail in **Appendix B**. One notable control missing from my estimates is race. Race and ethnicity are not reported to the Internal Revenue Service or the Social Security Administration.

1.3.1 Stratified Sampling

Even after controlling for the variables listed above, there is reason to believe that the estimated LIHTC effect using the full population, $\hat{\theta}_{All}$, may suffer from omitted variable bias on unobserved characteristics. Thus, I employ a stratified sampling approach to deal with the potential sources of bias explained below, and to establish a more convincing level of causal inference for my estimates.

One source of potential bias has to do with the circumstances that may cause a family to leave LIHTC housing. Approximately two thirds of the dependents in my sample leave LIHTC housing before age 18. This could be for positive reasons I cannot observe, like the family moving into a better neighborhood with a relative, or for negative reasons, like eviction. If there is a high incidence of eviction among families who spend just one or two years in LIHTC housing, for example, then my estimates will be positively biased. The positive effect of living in LIHTC housing for a longer period of time will be amplified because the reasons underlying the family leaving LIHTC housing may further disadvantage the children whose families spend a short amount of time in LIHTC housing due to eviction.

Thus, in order to eliminate potential bias coming from families leaving LIHTC housing, I divide my sample into groups based on the last age at which I observe each child in a LIHTC building. I then estimate the effect of spending an additional year in LIHTC housing on each of these sub-groups, stratified by age of exit, and then combine the total effect using weights based on the percent of individuals represented in each group. The regression specification is defined as follows:

$$\hat{\theta}_S = \sum_{k=8}^{18} \hat{\theta}_k w_k$$

$$s.t.\ E[y_{i,b,z} \mid i \in \mathbf{S}_k] = E[\theta_k h_{i,b,z} + \beta X_{i,b,z} + \gamma_b + \delta_z \mid i \in \mathbf{S}_k]$$

$$w_k = \frac{\sum_i I\{i \in \mathbf{S}_k\}}{\sum_i I\{i\}}$$

Here, $\mathbf{S}_k$ represents "stayers", i.e. the group of individuals who stay in LIHTC housing up until age k. All other variables remain as before. By undertaking this procedure, all of the variation in my estimate now comes from the age of entry into LIHTC housing, and compares individuals only to others who exit LIHTC housing at the same age. For example, among individuals who stay in LIHTC housing through age 18, all individuals who spend one year in LIHTC housing enter at age 18, and all individuals who spend eight years in LIHTC housing enter at age 11. The equation assumes a linear effect, meaning the single-year effect of moving in one year earlier is the same from age 10 to 11 as it is from age 17 to 18. I will relax this assumption in the following section.

A second source of potential bias is families timing their move into a different neighborhood or a different rent-subsidized building when their child is a particular age. For example, parents may look for housing in a neighborhood with a good middle school district when their oldest child is about to leave elementary school (around age 10 or 11). Although I do not see any evidence that families are doing this on average – there appears to be a uniform distribution of age of entry within each birth cohort – parents who are generally more ambitious may plan their move into LIHTC building at a "better" age than other parents. There may be less disruption in a child's life, for example, if they move at the same time they were supposed to change schools anyway.

For this reason, I further limit the sample of individuals in my third regression to children who move into a LIHTC building the same year or one year after that building is placed in service. In this series of "New Building" regressions I am still using the stratified sampling procedure based on age of exit. Note that in this specification the leaving age begins at 11 years rather than eight due a low number of individuals leaving new buildings at younger ages in the data. The regression is specified as follows:

$$\hat{\theta}_{NB} = \sum_{k=11}^{18} \hat{\theta}_k w_k$$

$$s.t.\ E[y_{i,b,z} \,|\, i \in \mathbf{S}_k, i \in \mathbf{NB}] = E[\theta_k h_{i,b,z} + \beta X_{i,b,z} + \gamma_b + \delta_z \,|\, i \in \mathbf{S}_k, i \in \mathbf{NB}]$$

$$w_k = \frac{\sum_i I\{i \in \mathbf{S}_k\}}{\sum_i I\{i\}}$$

In the regression, **NB** is the set of individuals who move into a LIHTC building the same year or one year after the building is placed in service. If there are parents who are planning their move based on the age of their child, then it will be purely coincidental that they move into a building that was recently placed in service. Some bias may still exist if parents are planning to move into a new building the same year that their child turns a specific age, but it seems unlikely that such targeting would be successful since there is some uncertainty about when buildings are placed in service, and families would more likely search for a new apartment based on availability and location. Even families with a preference for new buildings would probably not differentiate much between a building constructed three or four years ago and a building constructed one year ago.

By limiting my sample to only families that I observe moving into a new building I also take advantage of the exogenous event of new affordable housing appearing in a given family's local area. The location data I have for parents show that most of these families (85 percent) move less than 10 miles from their previous location when they move into a LIHTC building, and over half (55 percent) do not move out of their current zip code. This is especially true for families in the "New Building" sub-sample. 58 percent of individuals in this subset do not change zip codes when moving into LIHTC housing, and 88 percent travel less than 10 miles to move into LIHTC housing (95 percent travel less than 20 miles). Moreover, there is little correlation between the distance a family moves and the affluence of their new neighborhood (measured by things like poverty level and median income). **Figures 3 and 4** below show the distribution of distances that families travel when they move into LIHTC housing, for the full data set and for the"New Building" sub-sample, respectively. The distributions are graphed by continuous miles and by discrete categories, and have large spikes at zero. This suggests that families move into low income housing as it becomes available in their area.

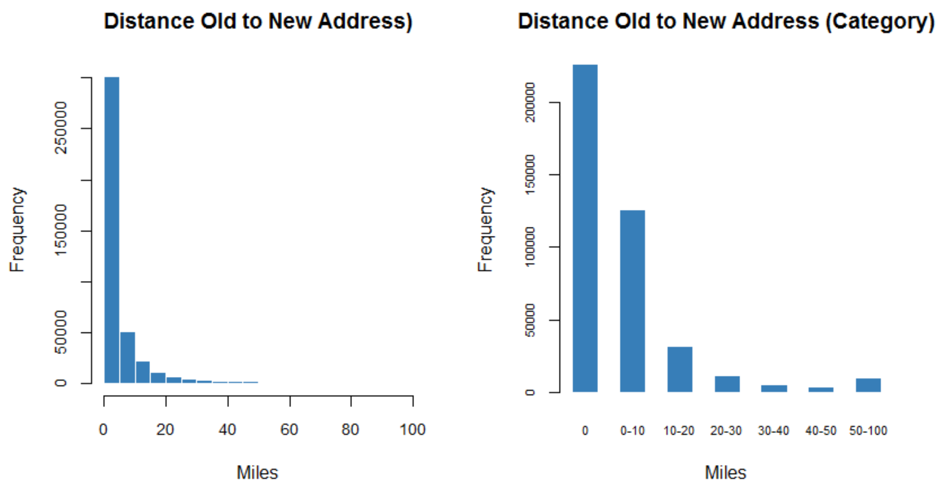

Figure 3: Distance between Old and New Addresses (Full Sample)

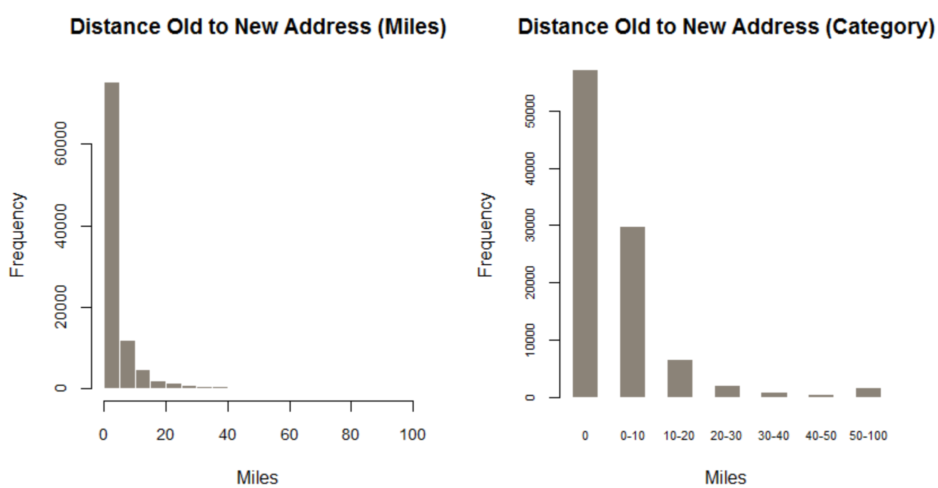

Figure 4: Distance between Old and New Addresses (New Building)

1.4 Results

1.4.1 Stratified Sampling Regression Results

As shown in **Table 2** below, the estimated effect of spending one additional year in LIHTC housing on the probability of attending four or more years of post-secondary education is positive, significant, and large. The estimates presented in **Table 2** are the relative odds ratios (minus one). When running the regression on all LIHTC residents, I find that individuals are 3.9 percent more likely to enroll in a four-year program for every additional year they spend in LIHTC housing. When I use the stratified sampling procedure outlined above (a weighted estimate for each sub-sample of "stayers"), the estimate increases significantly to 6.8 percent.

The difference in estimates suggests that there is some negative bias on $\hat{\theta}_{All}$ coming from families leaving LIHTC housing at various ages. However, when limiting the sample further to include only families who move into new buildings, the estimated effect goes back down to 3.9 percent, indicating that there may also be positive bias coming from families timing their move into LIHTC housing based on the age of their child.

Most of the control variables also have statistically significant estimated effects on adult outcomes. Higher household income and older parental age at birth are both associated with a higher probability of college attendance; as is higher area median income, although this effect disappears when using stratified sampling. This may be because AMI is defined at the county level, and the zip code fixed effect is absorbing most of the effect (though not all since I calculate area median income as an average over all years spent in LIHTC housing).

I also find that males are significantly less likely to attend a four-year program than females. This result is in line with the original Moving to Opportunity (MTO) study, which found that girls who moved to low opportunity neighborhoods fared better than boys (Ludwig, et al., 2013; Clampet-Lundquist, et al., 2006; Kling, et al., 2007). Girls who participated in the experiment did 13.6 percent better across a number of educational measures − including test scores and graduation rates − than girls in the control group. In comparison, boys who moved to low poverty neighborhoods performed worse than their counterparts in the control group, although the differences were not statistically significant (Kling, et al., 2007). The finding is also in line with gender comparisons of college enrollment rates for low income students: low income boys are generally less likely to enroll in a four-year university program than low income girls (Semuels, 2017).

As shown in **Table 3** below, the estimated effect of growing up in LIHTC housing on attending two or more years of post-secondary education is similar. According to my preferred specification, for every additional year spent growing up in LIHTC housing, individuals are 3.7 percent more likely to attend two or more years of higher education. Differences in the estimate across regressions and the estimated effect of each of the control variables also follow the same patterns as in **Table 2**, with a slightly larger effect measured in the first regression on the full sample. It is important to note that since the results I am presenting in each table are log odds, the results should be understood as a percent increase over the baseline education rates of 23.2 percent (four years) and 42.6 percent (two years). For example, someone who spends 8 years in LIHTC housing would be approximately 8.3 percentage points more likely to attend university than someone who spends one year in LIHTC housing (a 3.9 percent increase in enrollment for eight years over a baseline rate of 23.2 percent). Similarly, someone who spends 8 years in LIHTC housing would be approximately 14.4 percentage points more likely to enroll in two or more years of higher education compared to someone who spent one year in LIHTC housing.

I am concerned that the positive estimated effect of LIHTC years might simply be explained as the effect of remaining in any given place for a longer period of time (regardless of whether the family is living in a LIHTC building). In order to check whether this is the case, I run the same regressions including the variable "Total Moves", defined as the total number of times that a family moves over the entire time period that I observe them. As shown in **Tables 26 through 28** in **Appendix C**, including this variable in my regressions does not significantly change my estimates.

Table 2: Regression Results: 4+ Years Higher Education

	4+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.039***	0.068***	0.039**
	(0.004)	(0.011)	(0.018)
Male	-0.537***	-0.540***	-0.540***
	(0.007)	(0.011)	(0.018)
Log Household Income	0.228***	0.254***	0.234***
	(0.013)	(0.025)	(0.039)
Parent's Age at Birth	0.008***	0.009***	0.009***
	(0.001)	(0.002)	(0.003)
Log Area Median Income	1.481***	0.111	-0.021
	(0.359)	(0.249)	(0.316)
Family Size	-0.102***	-0.121***	-0.133***
	(0.007)	(0.012)	(0.020)
Log Units in Building	0.033**	0.047*	0.030
	(0.014)	(0.024)	(0.052)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.200	0.207	0.232

Note: *p<0.10, **p<0.05, ***p<0.01

Table 3: Regression Results: 2+ Years Higher Education

	2+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.050***	0.078***	0.037**
	(0.004)	(0.009)	(0.015)
Male	-0.534***	-0.539***	-0.547***
	(0.005)	(0.009)	(0.015)
Log Household Income	0.181***	0.206***	0.191***
	(0.010)	(0.019)	(0.031)
Parent's Age at Birth	0.003***	0.004***	0.005**
	(0.001)	(0.001)	(0.002)
Log Area Median Income	1.829***	0.257	0.330
	(0.337)	(0.236)	(0.391)
Family Size	-0.087***	-0.108***	-0.121***
	(0.006)	(0.010)	(0.017)
Log Units in Building	0.033***	0.044**	0.003
	(0.011)	(0.020)	(0.043)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.376	0.391	0.426

Note: $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$

Table 4 below shows results from the same three regressions – the full sample and two stratified sampling regressions – estimating the effect of growing up in LIHTC housing on adult earnings, as observed in 2018. As detailed in the table, the estimated effect of spending one additional year growing up in LIHTC housing on adult earnings is positive and statistically significant at the five percent level in all three regressions, which is not surprising given the positive estimated effect on education (on average, people with greater educational attainment have higher earnings). When looking at the full population, for every additional year spent in LIHTC housing individuals earn approximately 3.6 percent more in income.

This estimate is once again higher when running the regression using the "Stayers" stratified sampling approach, which estimates that on average individuals earn 5.7 percent more for every additional year spent in LIHTC housing. The estimated effect for families moving into new buildings, $\hat{\theta}_{NB}$ is actually close to this estimate: for every additional year spent in LIHTC housing individuals earn approximately 5.2 percent more in income as adults. Since the baseline wage level (the mean wage among individuals who spend one year in LIHTC housing) is $24,127, the estimated effect equates to an increase in earnings of approximately $1,255 as a result of one additional year spent in LIHTC housing. The cumulative effect of spending 8 years in LIHTC housing (compared to one year) is equivalent to an increase in earnings of about $12,066.

There are similar patterns in the estimated effect of each control variable as in the regressions on education: males earn less on average compared to females, household income has a positive estimated effect on earnings, and family size has a negative estimated effect on earnings. Variables like area median income, the age of the parent at birth, and the number of units in the LIHTC building do not seem to have as much of an effect on wages as they do on education outcomes. This may be because the individuals I observe are between ages 24 and 36 in 2018, an age range in which the gap in earnings between college- and non-college-educated workers tends to be smaller. The relationship may be stronger if I were to compare incomes at a later age, which I am unable to do at present time due to data restrictions (data are not available before 1999).

Since the estimated effect of spending an additional year in LIHTC housing on earnings is similar to the effect on education, I further test whether the education is the likely mechanism through which spending additional time in LIHTC housing may affect earnings. In **Table 5** below I test this theory by including the incidence of attending two or more years of education as a control variable in my earnings regressions. In doing so I find that the effect is lower across the board, and is no longer statistically significant at the five percent level when I use the stratified sampling approach. Thus, it appears that the effect of spending an additional year in LIHTC housing on future earnings works primarily through the effect on education: individuals who spend a longer amount of time growing up in a LIHTC building are more likely to enroll in two or more years of higher education, and as a result are more likely to earn higher wages as adults. For this reason, when I decompose the effect based on neighborhood and family characteristics in the following section I focus solely on the effect of LIHTC housing on educational outcomes.

In sum, spending more time in LIHTC housing does appear to be strongly correlated with both higher educational achievement and higher earnings. Even after controlling for a number of variables including location, family, and individual characteristics, and using stratified sampling techniques to control for unobserved correlates, the estimated effect, $\hat{\theta}$ is positive, significant, and large. Furthermore, when I include education as a control variable in my earnings regression it absorbs most of the estimated LIHTC effect, suggesting that the primary mechanism through which growing up in LIHTC housing affects earnings is through higher education.

Table 4: Regression Results: Earnings in 2018

	Log Adult Earnings (in 2018)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.036***	0.057***	0.052**
	(0.003)	(0.014)	(0.025)
Male	-0.430***	-0.447***	-0.525***
	(0.021)	(0.039)	(0.062)
Log Household Income	0.192***	0.204***	0.196***
	(0.007)	(0.024)	(0.041)
Age of Parent at Birth	-0.002***	-0.002	-0.001
	(0.001)	(0.002)	(0.004)
Log Area Median Income	0.663***	-0.321	0.726
	(0.112)	(0.418)	(0.786)
Family Size	-0.101***	-0.106***	-0.107***
	(0.007)	(0.020)	(0.035)
Log Units in Building	0.016	0.017	-0.056
	(0.013)	(0.034)	(0.080)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	102,476
Baseline Mean Wage	$23,007	$22,505	$24,127

Note: *p<0.10, **p<0.05, ***p<0.01

Table 5: Regression Results: Earnings in 2018 (Controlling for Education)

	Log Adult Earnings (in 2018)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.014***	0.024*	0.029
	(0.003)	(0.014)	(0.018)
2+ Years Education	2.019***	2.039***	1.610***
	(0.020)	(0.037)	(0.046)
Male	-0.083***	-0.097**	-0.126**
	(0.020)	(0.039)	(0.049)
Log Household Income	0.125***	0.131***	0.101***
	(0.007)	(0.022)	(0.030)
Age of Parent at Birth	-0.003***	-0.004**	-0.001
	(0.001)	(0.002)	(0.003)
Log Area Median Income	0.194*	-0.226	0.446
	(0.103)	(0.403)	(0.586)
Family Size	-0.060***	-0.054***	-0.036
	(0.006)	(0.019)	(0.026)
Log Units in Building	0.002	-0.002	-0.034
	(0.011)	(0.032)	(0.059)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	102,476
Baseline Mean Wage	$23,007	$22,505	$24,127

Note: *p<0.10, **p<0.05, ***p<0.01

1.4.2 LIHTC Fixed Effects

In the regression results presented above, I make assumptions about the linear or log-linear effect $\hat{\theta}$. In other words, I assume that the effect of spending one additional year in LIHTC housing is the same at any age. For example, the effect of spending one more year in LIHTC housing from ages six to seven is the same as spending an additional year in LIHTC housing from ages 16 to 17. However, this assumption may not be correct, and the effect may be heterogeneous in age of entry. It may be the case that the effect of spending one more year in LIHTC housing at age six has a much greater effect on adult outcomes than spending an additional year in the same housing at age 16.

In order to test the linearity assumption I make in the previous section, I transform the variable "LIHTC years" into a fixed effect to separately estimate the effect of spending different amounts of time in LIHTC housing, relative to one year. In place of $h_{i,b,z}$, I regress outcomes on a vector of 11 binary variables, $H_{i,b,z}$, one for every possible number of years spent in LIHTC housing, from 2 to 12 years. I use the same stratified sampling procedure with this new regression specification. The results from these regressions are presented in **Figures 5 through 7** below. **Figure 5** graphs the estimated difference in the probability of attending a higher education program for four or more years, for individuals who spend different amounts of time growing up in LIHTC housing. Each point represents the estimated effect and the tails are 95 percent confidence intervals. The effect for each group (categorized by LIHTC years) is the estimated percentage change compared to spending one year growing up in LIHTC housing. **Figure 6** graphs the same results for two or more years of education, and **Figure 7** graphs the percent difference in earnings for each group relative to those who spend one year in LIHTC housing as children.

Although the pattern in estimated differences is not perfectly linear, there is little evidence of a non-linear increase in gains from spending one additional year growing up in LIHTC housing at different ages. Rather, the effect is cumulative: the benefit of moving in at an earlier age is that one can potentially spend a longer amount of time living in LIHTC housing, with each additional year associated with a similar increase in future earnings and probability of enrolling in higher education. There does appear to be a small increase in the effect when comparing individuals who spend six versus seven years in LIHTC housing, and this pattern is stronger when you look at individuals who remain in LIHTC housing through age 18, which is equivalent to moving in at age 12 rather than age 13. **Figure 37** in **Appendix D** graphs the results of this analysis, comparing only individuals who stay in LIHTC housing through age 18. As shown in the figure, those who move in prior to age 13 see a jump in the estimated effect of spending one additional year in LIHTC housing compared to those who move in at age 13 or later. This is consistent with previous analysis of the MTO experiment, with greater positive effects for individuals who move into a low poverty neighborhood before age 13 (Chetty, et al., 2018).

Overall, this analysis suggests that the estimated gains from spending an additional year growing up in LIHTC housing do not vary by age, aside from a small increase in the effect if an individual moves in before high school and remains in LIHTC housing through age 18. As shown in **Figure 7**, the confidence intervals for the effect of LIHTC on earnings are similar, although there is a lot more noise in the estimate, resulting in larger standard errors. As mentioned earlier, this is likely because the individuals in my data are between ages 24 and 36 in 2018, a time at which differences in salaries between college- and non-college-educated workers are smaller and more varied than they are later in life. Regardless, there is also no sign of heterogeneous age effects in these results.

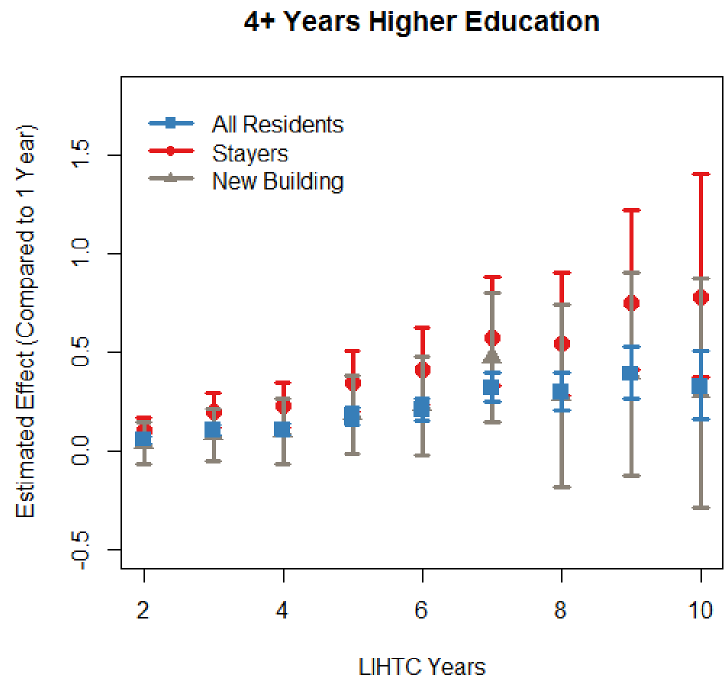

Figure 5: LIHTC Fixed Effects Results: 4+ Years Education

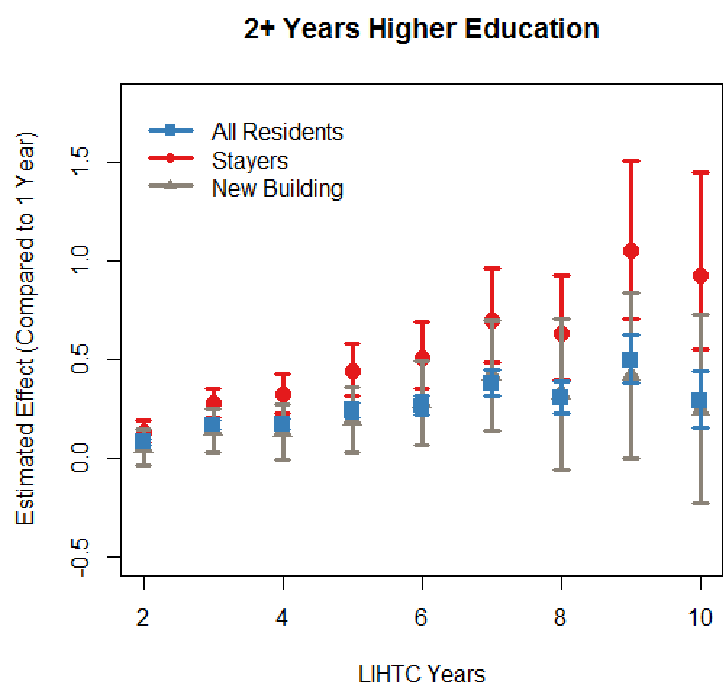

Figure 6: LIHTC Fixed Effects Results: 2+ Years Education

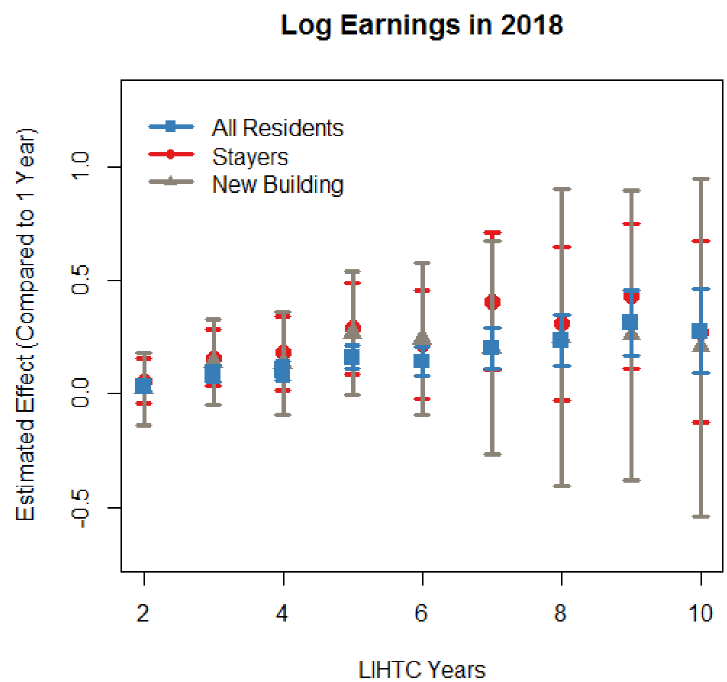

Figure 7: LIHTC Fixed Effects Results: Earnings

All three graphs also illustrate the large cumulative estimated effect of growing up in LIHTC housing. According to the "New Building" fixed effect regression results, a person who spends eight years in LIHTC housing as a child is on average 28.0 percent more likely to enroll in four or more years of higher education than someone who spends only one year in LIHTC housing. Similarly, those who spend eight years in LIHTC housing earn 24.7 percent more in income on average than individuals live in LIHTC housing for one year.

There are several possible explanations for the observed effect. Perhaps families are moving into better neighborhoods with higher performing schools, or parents are paying less in rent, which allows them to invest more income in their children's education. It is also possible that the availability of affordable housing simply provides stability for families that previously moved around a lot. In the following section I attempt to answer the question of why spending an additional year in LIHTC housing might affect educational attainment by looking at three possible sources of heterogeneity: building location, housing stability, and household income. First, I look at differences in the LIHTC effect for buildings in neighborhoods with difference characteristics. Second, I look at differences in the effect for families that vary in how many times they change addresses prior to entering LIHTC housing. And third, I look at differences in the effect based on parents' earnings.

1.5 Heterogeneous Effects

1.5.1 Neighborhood-Level Effects

Many characteristics of a neighborhood can affect a child's trajectory in life. Measurable demographic variables like racial composition, median wealth, poverty rates, and education rates can all affect the opportunities that a child has for greater educational attainment or higher wages. Thus, as an extension of my analysis of the effects of the LIHTC on individuals' long-run outcomes, I want to know if the effect of growing up in LIHTC housing is more positive in certain neighborhoods compared to others. In the following analysis I focus on the effect of LIHTC housing on education since this appears to be the driving mechanism behind the effect on earnings.

I first look at differences in the LIHTC effect across neighborhoods with different characteristics. I focus on five observable sources of variation at the Census tract level: poverty rates, median household income (in 1990 and in 2012), high school graduation rates among adults 25 and older in the neighborhood, racial composition (the percent of the population that is white), and opportunity measures from Chetty, et al. (2018), as described in the previous section. I also look at differences across zip codes with varying public high school graduation rates as a measure of secondary education, although there are some problems with these data, as I discuss below.

To measure the differences in the estimated effect, $\hat{\theta}$, across neighborhoods with varying characteristics I match each LIHTC building with location data using the Federal Information Processing Standards (FIPS) codes provided by HUD – also known as Census tract codes – and then divide the full data set into 10 evenly divided groups based on quantiles of each neighborhood characteristic. I then run 10 separate regressions for each group and graph the predicted effects, $\hat{\Theta}$, to compare results based on each characteristic. **Table 29** in **Appendix E** provides the cutoff points for each quantile of each variable.

Figure 8 below shows the regression results for the estimated effect $\hat{\theta}$ by neighborhood characteristic on four years of higher education, and **Figure 9** shows differences in the effect on two or more years of education. For the first three characteristics – poverty level, median household income (in 1990 and 2012), and high school graduation rate (of adults over 25) – there do appear to be trends in the direction one would expect, but the differences are not statistically significant. The estimated LIHTC effect decreases slightly as the poverty level rises, and it increases as household income and high school graduation rates increase. However, the confidence intervals are overlapping even when comparing the lowest and highest estimates.

On the other hand, there are significant differences in the estimated effect when it comes to racial composition and opportunity measures. As shown in **Figure 8**, the estimated effect of spending an additional year in LIHTC housing on the probability of enrolling in four or more years of higher education is 1.6 percent in a neighborhood that is 0-11.3 percent white. In comparison, the estimated effect for individuals who move into a LIHTC building in a neighborhood that is 70.6-79.0 percent white is 3.7 percent.

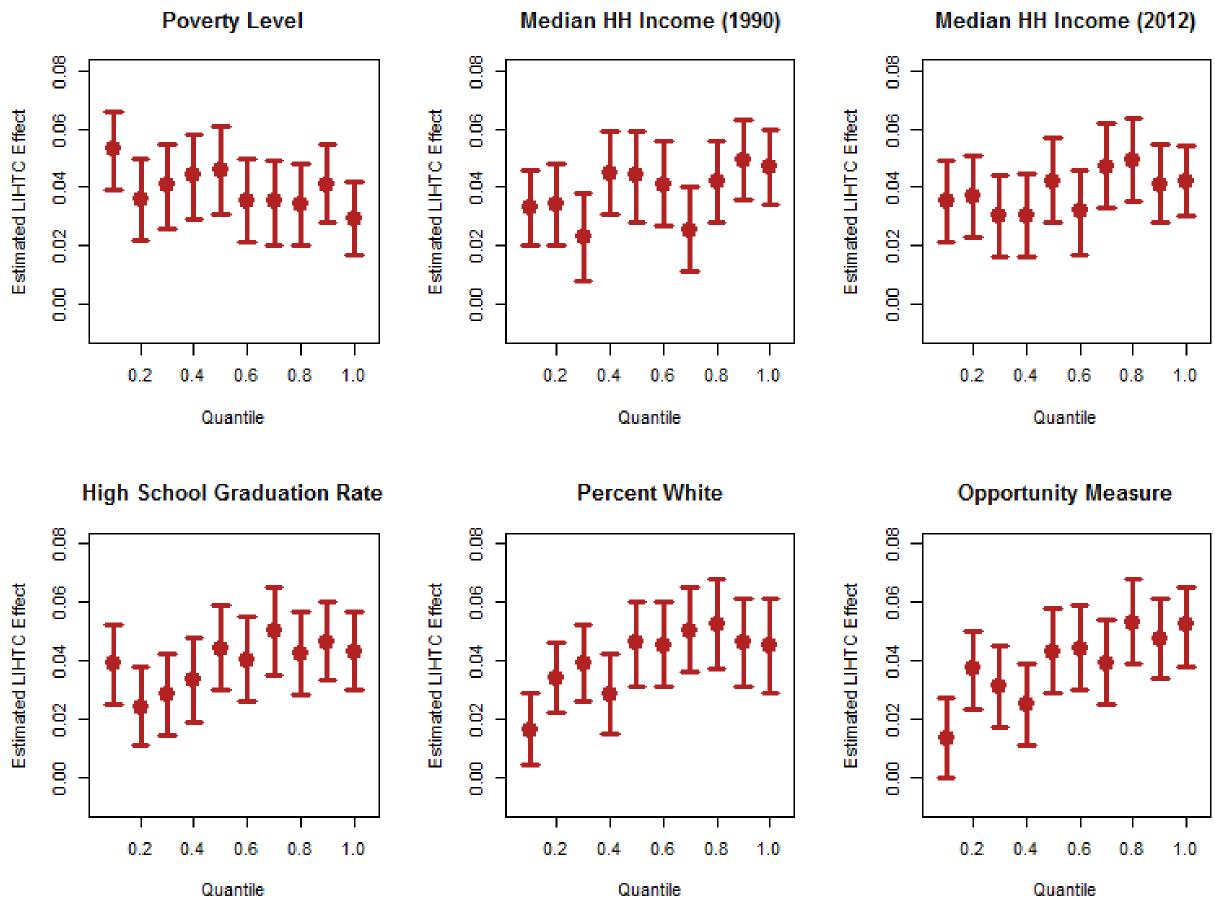

Figure 8: Neighborhood Effects, 4+ Years Higher Education

However, because there is a strong correlation between a neighborhood's racial composition and the racial identities of the individuals that live in that neighborhood, I cannot draw strong conclusions about differences in the LIHTC effect between neighborhoods of varying racial makeup. By construction, people of color are more likely to live in neighborhoods that have fewer white people, and white people are more likely to live in neighborhoods that have fewer people of color. So differences in the LIHTC effect may actually be capturing differences in the benefit that white people get from moving into LIHTC housing for a longer period of time compared to other racial groups. The ability to include race in these regressions would be very helpful in determining whether this difference in effect persists when controlling for the race of each individual.

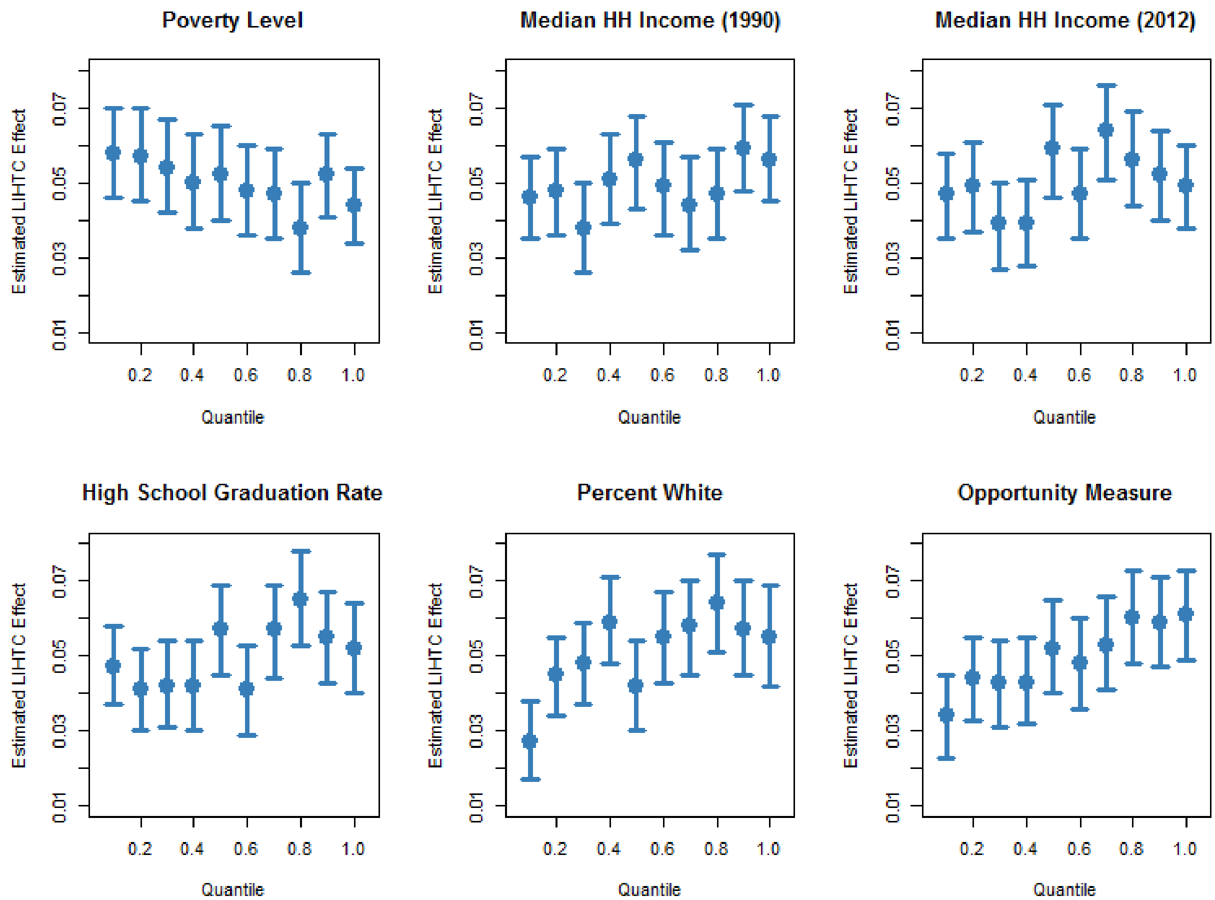

Figure 9: Neighborhood Effects, 2+ Years Higher Education

One thing to note when looking at differences in effect across neighborhoods with varying racial composition is that even the lowest estimated LIHTC effect of 1.6 percent is still positive and significant. Moreover, the estimated effect for the second lowest quantile (neighborhoods that are 11.3-28.3 percent white) is 3.4 percent, which is quite high and close to the estimated average effect for the entire population. This suggests that those who spend more time in LIHTC housing as children have better outcomes regardless of what neighborhood they live in. An individual who lives in a neighborhood that is mostly white may benefit slightly more – either because of the neighborhood, because of the individual's race, or a combination of both factors – but spending additional time growing up in LIHTC housing still appears to be generally beneficial across all neighborhoods.

I also see significant differences between neighborhoods with varying opportunity measures. The estimated effect for individuals who grow up in LIHTC housing in neighborhoods with an opportunity measure of 0 to 0.04 is 1.3 percent. In comparison, the estimated effect for individuals growing up in a neighborhood with an opportunity index of 0.22 to 0.27 is 5.2 percent. Once again, we cannot necessarily draw conclusions about these differences as neighborhood opportunity measures and racial composition are highly correlated. Moreover, there is a small correlation between the number of buildings that are constructed in a given area and the opportunity measure, although the correlation is negative at -0.16.

Finally, I look at differences in the estimated effect between zip codes that have varying rates of graduation at local public schools. This analysis is a lot noisier than the other comparisons. One reason for this is that I measure high school graduation rates at the zip code level due to lack of data available on public school graduation rates at the Census tract level. I also do not know what percent of students in each zip code attend public school (versus private or religious schools). Finally, the data − from the US Department of Education, and the California Department of Education − are more sparse than the Census data I use in my other analyses, so I am only able to break up the data into five groups. The estimated heterogeneous effects based on public school graduation rates (on four and two years of higher education, respectively) are shown in **Figure 10** below.

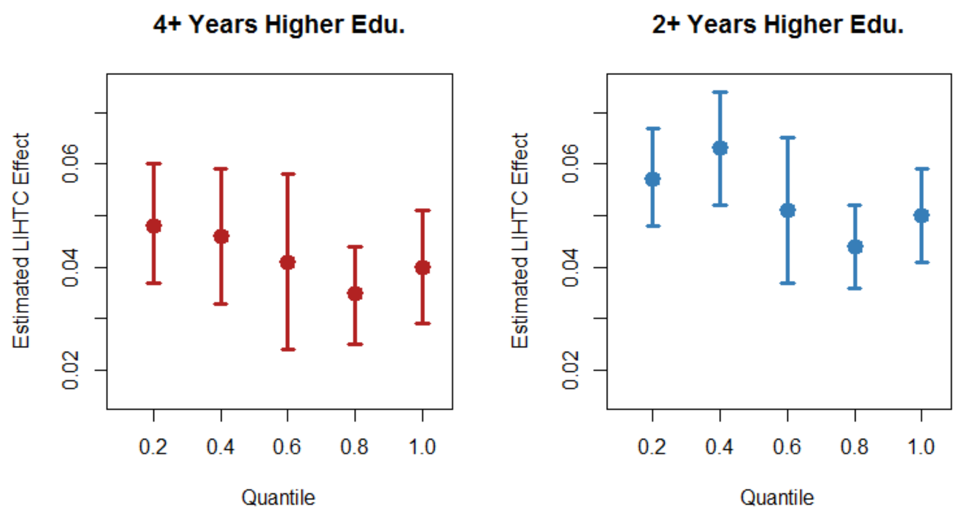

Figure 10: Neighborhood Effects: Public High School Graduation Rates

As shown in the figure, the trend in differences between zip codes with different public school graduation rates appears to be negative, suggesting that the effect of growing up in LIHTC housing is lower for individuals who live in zip codes with better performing schools. However, the differences are not significant, and because of measurement issues when it comes to high school performance, the data most likely do not paint a full picture of how children benefit from enrolling in schools with higher graduation rates, so no conclusions should be drawn from this analysis. Furthermore, the differences in high school graduation rates reported by zip code are quite small when comparing the top three quantiles. Those in the third quantile live in a zip code where the public high school graduation rate is between 87 and 95 percent. In comparison, those in the top quantile live in a district with graduation rates between 97 and 100 percent. Thus, one would not expect to see large differences in the effect given the similarity of schools in each group.

In general, I find that while the location of LIHTC housing appears to matter somewhat when it comes to the size of the LIHTC effect, differences between neighborhoods do not fully explain the positive estimated effect of growing up in LIHTC housing on adult outcomes. Regardless of where the LIHTC housing is located, it appears that individuals benefit from spending additional time living in these buildings as children. Even when I estimate statistically significant differences in the effect between neighborhoods with different racial demographics and opportunity scores, I cannot rule out the possibility that these differences are attributable to individual rather than neighborhood characteristics (in particular the race of the individual).

Furthermore, as I mention earlier, most people do not move very far away from where they are living when they move into LIHTC housing. As shown in **Figure 4** above, most of the families in my data stay in the same zip code or move to a zip code less than 10 miles away when they move into a LIHTC building. Moreover, the families who do change locations do not seem to move to a "higher opportunity" area on average.

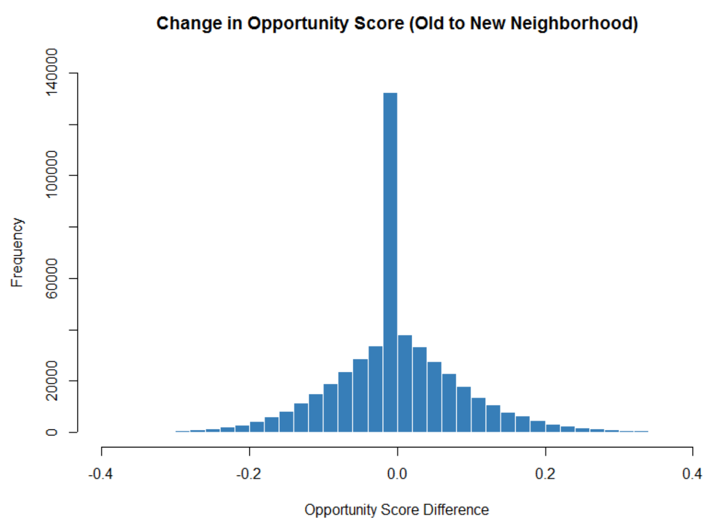

Figure 11: Differences in Opportunity Measure from Old to New Neighborhood

As shown in **Figure 11** above, when we graph the difference in opportunity scores between old and new neighborhoods – the new neighborhood being where the LIHTC building is located – there appears to be a normal distribution with a large spike at zero, which represents individuals who do not change zip codes. This trend is in line with the original Moving to Opportunity study, which found that people did not generally move to higher opportunity neighborhoods when given the choice (Ludwig, et al., 2013; Sanbonmatsu, et al., 2011). Thus, moving to a "better" neighborhood does not seem to be a very good explanation for why individuals might be better off when they spend more time in LIHTC housing as children.

1.5.2 Family-Level Effects

Another possible explanation for why individuals might have better outcomes the longer they spend in LIHTC housing as children is that subsidized housing provides a more stable living situation for families who are suffering from a lack of housing security. Housing insecurity can cause families to move from one living situation to another, disrupting the lives of the children, particularly when it comes to their schooling. When affordable housing becomes available in a family's area − a new LIHTC building is constructed or a subsidized unit becomes available to rent − this can provide a more stable living situation for the family, as they are more likely to be able to afford that housing for a longer period of time.

In order to test to what extent improved housing security is a driver of the estimated LIHTC effect I divide my full sample into five groups based on the number of times that they change addresses *before* entering into LIHTC housing. This provides a relatively good measure of the level of housing stability the family has prior to securing low income housing in a LIHTC building. Since the number of pre-LIHTC moves is a discrete variable, I group individuals into five categories based on the number of times the family changes addresses: one category each for zero to three moves and one category for four or more moves. I then run the same regression for each group. Since taxpayers sometimes receive information returns at more than one address, I find the most commonly reported address for parents of each child in each year before they move into LIHTC housing, and then sum up the number of moves based on these addresses.

As shown in **Figure 12** below, I find strong evidence of heterogeneous effects in housing stability. I estimate that for a person who did not change addresses prior to entering into LIHTC housing, one additional year spent in LIHTC housing is associated with a 2.1 percent increase in the likelihood of attending four or more years of higher education. In comparison, the estimated effect for someone who moved three times is 6.7 percent, and for someone who moved four or more times the estimated effect is 11.7 percent. Unlike in the neighborhood effects analysis based on opportunity scores, the correlation between the number of pre-LIHTC moves and neighborhood racial composition is low (about 0.036), which suggests that differences in housing stability cannot be attributed to the race of the individual.

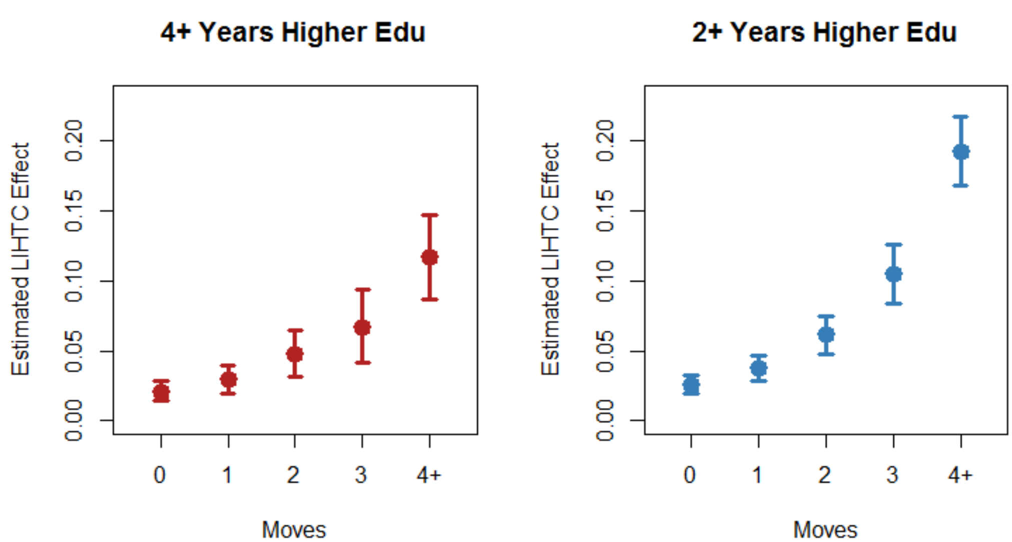

Figure 12: Heterogeneous Effects: Housing Stability

As with the analysis of heterogeneous neighborhood effects, it is encouraging that the estimated effect for the lowest group here is still positive and significant, albeit using the full sample, without implementing the stratified sampling procedures. Nevertheless, the large effect that an additional year in LIHTC housing seems to have on individuals who move around a lot prior to entering into subsidized housing suggests that this is likely one of the main drivers of the overall effect.

Another possible explanation for the positive estimated LIHTC effect is that families have more financial resources to invest in their children. In order to test this theory I once again split my population into ten groups based on average household income (calculated as an average for all years I observe the child from ages 6 to 18). As shown in **Figure 13** below, I see that there are strong patterns in the LIHTC effect based on household income as well. The estimated effect on attending four or more years of higher education for a family in the bottom 10 percent of earnings is 6.8 percent. In comparison, the estimated effect for a family in the top 10 percent of earnings is 1.3 percent. The difference in the effect between these two groups is statistically significant a the one percent level. As shown in **Figure 13**, the same pattern also is evident in regressions estimating the effect on the incidence of attending two or more years of higher education.

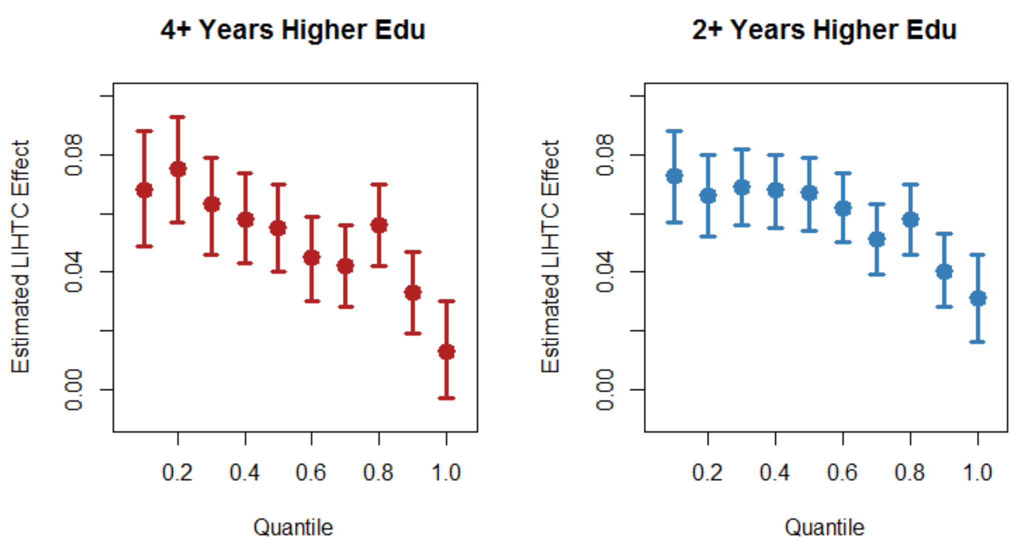

Figure 13: Heterogeneous Effects: Household Income

These results suggest that the underlying reasons why spending more time in LIHTC housing might be beneficial to children can be explained by differences at the household level. When families have access to rent subsidies they have better housing security and more disposable income, both of which appear to have positive effects on the educational outcomes of their children (and subsequently on their earnings as adults). Although there are still positive and statistically significant effects for families with higher earnings and fewer moves prior to entering into a LIHTC building, the high estimates that we see for those with lower earnings and housing security are likely driving up the overall estimated LIHTC effect.

1.6 Conclusion

In sum, I find that growing up in LIHTC housing has a large positive effect on both education and earnings, with an estimated 3.9 percent increase in the probability of attending four or more years of higher education, and a 5.2 percent increase in earnings for every additional year spent living in a LIHTC property. I find that the effect is somewhat heterogeneous in age, with a greater increase in the probability of enrolling in higher education for those who move in before age 13 (comparing individuals who remain in LIHTC through age 18). I also find that the primary mechanism through which LIHTC housing affects earnings is through the effect on higher education.

In general, I find that changes in location do not provide a good explanation for the positive estimated effect of spending an additional year in LIHTC housing. Although there is evidence in the data and in the previous literature to suggest that where people grow up has a large impact on their future outcomes, it is not clear whether location is a driving mechanism behind the LIHTC effect. The only statistically significant differences in the estimated LIHTC effect based on location are between individuals residing in neighborhoods with different racial composition, or in neighborhoods with different opportunity scores. However, there is reason to believe that these differences can be attributed to individual characteristics, such as race.

Moreover, the findings from my analysis of neighborhood effects are inconclusive because a majority of individuals in my data do not move far away from the neighborhood they already lived in prior to entering LIHTC housing, and the ones who do move are not relocating to "better" neighborhoods on average (places with lower poverty rates or higher opportunity scores). It may be that growing up in a more affluent neighborhood has a large effect on individuals' future outcomes, but the estimated LIHTC effect in this study cannot be attributed to changes in neighborhood since most of the families who move into LIHTC housing do not change locations.

On the other hand, there is strong evidence that the LIHTC effect is in large part a measure of the effect of housing stability and household income. The more times an individual changes addresses prior to moving into a LIHTC building, the more positive the effect of LIHTC housing is on their outcomes. Similarly, the lower the family's income, the greater the effect on education. Based on these findings, it seems that one of the main benefits of constructing subsidized housing is that it provides low income families with a more stable living situation, and with more disposable income to invest in their children. Generally, I find that families who are in a worse financial or housing situation prior to entering into LIHTC housing are the ones whose children benefit the most. We see the largest effect of LIHTC housing on education – and subsequently earnings – for children whose families earn less and move around more.

There are still some unanswered questions when it comes to understanding why growing up in LIHTC housing has a positive effect on future education and earnings. Even among individuals who do not change addresses prior to moving into a LIHTC building, and among those with higher household incomes, the estimated LIHTC effect is positive and significant, suggesting that there are other underlying mechanisms that can help explain these results. Further research is also needed to understand the ways in which the LIHTC may benefit other groups of renters, expanding the analysis beyond the outcomes of children. The data and methods used in this paper will hopefully help inform these types of analyses in the future.

CHAPTER 2: SAVINGS RESPONSES TO AUTO-ENROLLMENT: EVIDENCE FROM A LARGE PANEL OF WORKER-EMPLOYER LINKED DATA[1]

Over the past decade it has become increasingly common for private companies to adopt automatic enrollment as a feature of their retirement plans, and the practice has been encouraged by federal and state governments. Several papers examining individual companies that adopted automatic enrollment have also provided evidence supporting the effectiveness of this "nudge," finding large positive effects on employee participation and contributions (Falk & Karamcheva, 2019; Beshears, 2009; Nessmith, Utkus, and Young, 2007; Thaler & Benartzi, 2004; Choi et al., 2004; Madrian & Shea, 2001). The positive estimated effects of the policy are particularly stark when compared to the relatively small effects of other interventions like increasing employer match rates on 401(k) participation and contribution rates (Engelhardt & Kumar, 2007; Duflo, et al., 2006; Even & Macpherson, 2005; Munnel, et al., 2001; Kusko, et al. 1998; Papke & Poterba, 1995).

In this paper, we study whether employees who are automatically enrolled in retirement plans are more likely to stay enrolled than employees who have to opt in to their company-sponsored retirement plan. We also estimate differences in the average percent of wages that individuals from each group defer towards a retirement plan. Our paper departs from previous research that typically examines auto-enrollment adoption in a single firm or small number of firms. Instead, we analyze the effects of automatic enrollment using employer-employee linked data from 279 large US-based firms that adopted the policy between 2010 and 2016, with information extracted from detailed deferred contribution plan documents and individual tax and information returns. With this expanded data set we are able to use heterogeneity at both the firm and individual level to

[1] This chapter is co-authored with Jacob Mortenson from The Joint Committee on Taxation

evaluate how the effect varies based on different firm and employee characteristics.

Previous studies examining the effects of automatic enrollment date back to Madrian and Shea (2001), who were the first to estimate the effects of the policy on participation in deferred compensation plans. In their seminal paper, they find that 401(k) participation among employees with three to 15 months' tenure at a large health services corporation more than doubled after the firm began automatically enrolling employees in the company-sponsored retirement plan.[2] They found that enrollment among new employees increased from 37 percent enrollment to 86 percent. Controlling for gender, race, age, compensation, and months of tenure, they estimate an increase in employee participation in retirement plans of 50.1 percentage points as a result of the policy change.

They also find significant changes in the distribution of contribution rates before and after the introduction of automatic enrollment, with 76 percent of participants enrolling in the default rate of three percent after the policy change, compared to just 12 percent of employees before the change. They estimate that the regression-adjusted average difference in saving rates before and after automatic enrollment, controlling for other correlates, is 2.2 percentage points (Madrian & Shea, 2001). They explain the low increase as the result of a larger number of employees enrolling at the low default rate of 3 percent. It is important to note that these findings are based on employees at one firm, which differs from the national workforce in at least two important ways: the workers are mostly female (78 percent), and are located primarily in the East and Midwest of the United States.

Choi, et al. (2004) expand this analysis to two additional firms. In their analysis, they use data on employees from three companies that adopted automatic enrollment policies in the same

[2]Throughout the paper we use 401(k) – the most common type of employer-sponsored defined contribution plan – as short hand for employer-sponsored defined contribution plans.

year, including the same company studied by Madrian and Shea. Two of the firms in their sample implemented a three percent auto-enrollment policy, and the third adopted a two percent default auto-enroll rate. Among individuals with five months tenure, they find that the difference in enrollment between employees hired before and after implementation of automatic enrollment falls between 60 and 70 percentage points, although the difference decreases with tenure (the rate is closer to 40 percent for those with two years' service). However, in their measurement of the effect of automatic enrollment on participation in retirement plans, they do not control for potentially correlated variables like employee wages, which is an important omission if firms appear to be adopting automatic enrollment during a time of expansion, as we see with a few large firms in our data.

More recently, Falk and Karamcheva (2019) find positive changes in the saving behavior of federal workers after a government policy change to automatically enroll federal employees into the government's Thrift Savings Plan (TSP) at a default rate of three percent. They find that a considerably higher percent of workers participated in the plan after the policy change, with enrollment jumping from 60 percent to 97 percent. Overall, they estimate a 19 percentage point increase in participation due to automatic enrollment (Falk & Karamcheva, 2019). When looking at the effect on employee contributions, they note the same pattern of higher participation rates coupled with lower average contributions due to the relatively low auto-enroll default rate adopted by the federal government. They find that average contribution rates increase by just 0.6 percentage points after implementation, and they estimate larger average increases for policies that adopt a higher enrollment rate.

Our analysis contributes to the existing body of research on automatic enrollment in three ways. First, we use a larger and more diverse sample of firms in our baseline estimate of the effect of the policy, controlling for potential correlates like wages and employer matching policies. The companies in our sample range in size from 200 to over 65,000 employees, and adopt different default rates, ranging from one to six percent. The firms also come from a wide range of industries including construction, retail, wholesale trade, transportation, real estate, communications, scientific services, manufacturing, finance, insurance, professional services, technology, administrative support, education, health care, entertainment, and food services. Clustering firms by characteristics like auto-enrollment policy, size, and industry allows us to decompose the estimated effect so that we can better understand how automatic enrollment affects workers in different sectors of the economy.

Second, we analyze whether setting a relatively low default rate has a negative "nudge" effect on individuals who are already inclined to save. Using W-2 data we are able to identify which employees in the treatment and control groups opted in to company-sponsored plans in the years prior to joining one of the firms in our sample, and then estimate the effect of automatic enrollment on their saving levels. By looking at the effect of automatically enrolling these "savers" in plans with different default rates we are able to determine if setting a lower default rate leads to lower than average savings among workers who may have otherwise chosen a higher contribution rate.

Third, we analyze the distributional effects of automatic enrollment among employees by separately analyzing responses at different percentiles of the wage distribution within firms. Determining whether there are differences in the distributional effects of the policy is particularly relevant as automatic enrollment has gained notable traction among firms attempting to qualify for exemptions from deferral nondiscrimination testing requirements (IRS, 2020). If not exempt, firms are required to test their traditional 401(k) plans each year to ensure that the contributions made by and for "rank-and-file" (non-highly-compensated) employees are proportional to contributions to the plan made by and for highly-compensated employees such as owners and managers (IRS, 2020).

We estimate that automatically enrolling employees in a company-sponsored retirement plan increases participation by approximately 32.5 percentage points. In our sample, the mean saving rate among all employees is 28 percent, which approximately equates to a 116 percent increase in enrollment, up to 60.5 percent, controlling for other correlates like wages, gender, age, and firm. We also estimate that automatic enrollment is associated with an average increase in the individual saving rate of 0.9 percentage points, a relatively low increase that reflects both the large positive effect on participation and relatively low default auto-enroll rates adopted by firms. Among individuals who defer a positive amount of wages to a company-sponsored plan in the years before and after joining one of the firms in our sample, we find that automatic enrollment is actually associated with a 0.9 percentage point *decrease* in the level of savings, suggesting that although the overall effect of the policy is positive, setting a relatively low default rate acts as a negative nudge on individuals who would likely opt in to the plan at higher rates.

We further find that the effect of automatic enrollment on employee saving is heterogeneous in the default rate that firms adopt. Firms that automatically enroll their employees at lower default rates see both a higher increase in participation and a lower increase in the average saving rate. Part of the reason for this trend is the endogenous nature of default rates, since there are both firm and employee characteristics that are highly correlated with the default rate that each company chooses to adopt. However, when we look at the effect on savers – i.e. individuals who contributed to a firm-sponsored plan under their previous employer, before they were hired at one of the auto-enroll companies in our data – there is a clear pattern in the negative nudge effect: savers who are automatically enrolled at lower rates see a larger decrease in average contributions relative to their counterparts in the control group.

In addition, we see that the effects of automatic enrollment vary with wage level relative to other employees within the same firm, although differences in the absolute effect are not statistically significant. Employees at the lower end of the wage distribution see a greater increase in their participation rate relative to those at the top end of the wage distribution. However, differences in the effect between employees in the middle of the distribution are small and not significant. Conversely, employees at the lower end of the wage distribution do not see as large of an increase in their average saving rate due to automatic enrollment, even when looking at groups of employees auto-enrolled at the same default rate, although the differences are once again small and not significant.

We further try to explain the difference in automatic enrollment that we see compared to the previous literature by looking at differences in the effect among firms with different characteristics − specifically, firms that vary in size, mean wage, and industry. We find that that there are statistically significant differences in the effect of automatic enrollment when comparing firms across all three measures, which helps explain the difference between our baseline estimates and those in previous papers. By considerably expanding our sample of firms we are able to see that the effect of adopting automatic enrollment can vary considerably based on both workers and employers, thus providing a broader understanding of how the policy might affect the economy at large if implemented universally.

2.1 Data and Summary Statistics

The retirement plan-level data we use come primarily from Form 5500 filing attachments downloaded from the Department of Labor's (DOL) Employee Benefits Security Administration (EBSA) database. We begin by identifying companies that adopted automatic enrollment at some point between 2010 and 2016 using the Department of Labor's "Bulletin" files (which are cleaned by DOL) and we cross-reference these data with Form 5500 attachments downloaded from the EBSA database. These attachments provide details on each defined contribution (DC) plan, including eligibility requirements, employer matching rates, vesting schedules, and automatic enrollment rates. We use a text reading algorithm to identify which Form 5500 attachments mention automatic enrollment, and catalogue these by firm and year. A more detailed explanation of this process is available in **Appendix F**.

Once we have downloaded and processed the Form 5500 attachments, we link plan data to W-2 filings from 2010 to 2016. Importantly, we incorporate known parent-subsidiary relationships as part of this linkage. This process is also described in greater detail in **Appendix F**. After matching the two data sets, we use a text-reading algorithm to identify the exact year – and often the exact date – when each company began automatically enrolling employees in their 401(k) plan (or plans). In order to ensure the accuracy of our treatment variable, we restrict our sample of firms to those that explicitly mention a start date of automatic enrollment in one of their Form 5500 attachments.

Using W-2 filings we are able to identify employees that were hired one year before or one year after each firm adopted automatic enrollment. Individuals hired in the year before each firm adopts automatic enrollment – and must opt in to a retirement plan – are designated as our "control" population and employees who were hired one year after each firm adopted automatic enrollment – and were thus automatically enrolled in a retirement plan – are our "treated" populations. In order to ensure sufficient within-firm variation among employees we only include companies that hire 100 new employees both in the year before and the year after adopting automatic enrollment. This restriction reduces our sample considerably to only a few hundred firms. We also restrict our data to employees whose wages are greater than an employee working full time at minimum wage for three months. We do this to eliminate short-term employees and those who may not yet be eligible to participate in a retirement plan.

In the Form 5500 attachments, firms often include service requirements for eligibility to participate in the company-sponsored retirement plan. Employees must often wait for a period ranging from one day to one year before they become eligible to enroll. According to the Profit Sharing/401k Council of America (PSCA), in 2010 approximately 57.6 percent of firms had at least a 3 month waiting period before employees became eligible to enroll in 401(k) plans, and we also observe that roughly 60 percent of eligible firms in our data have a three month minimum service requirement (Gelber, 2011; PSCA, 2010). This complicates our measurement of treatment since we do not want to identify new employees as "treated" if they have only worked at the firm a short period of time, and are not yet eligible to enroll (or be automatically enrolled) in the company's retirement plan.

As a consequence, we use the Form 5500 DC plan attachments – which provide details on employee eligibility requirements and waiting periods – to further restrict our sample to firms that have less than a six-month service requirement for participation in the firm's retirement plan. For employees with less than a three-month waiting period, we calculate the percent of wages deferred to a retirement plan the same year they are hired. For individuals with a three- to six-month waiting period, we calculate the percent of wages deferred in the year after they join the firm, to allow for enough time for them to become eligible to contribute to the retirement plan. The hiring date for employees in the control group at companies that have a three- to six-month eligibility period is actually two years prior to adoption of automatic enrollment, to ensure we are still calculating their contributions before the company adopts the policy.

We also exclude firms from our sample that do not allow all workers (both part- and full-time) to participate in their plan. Finally, if a company has more than one retirement plan, or if the company has subsidiaries with their own retirement plans, we exclude the firm if there are conflicting auto-enrollment start dates across different plan documents. In general, we try to ensure that if a firm shows up in our sample as adopting automatic enrollment in year y, that most, if not all, of the employees at that firm are eligible to enroll in the firm's 401(k) plan within three months from when we observe their contributions. In general, we want to ensure that we do not observe low enrollment rates due to employee exclusions or long service requirements for eligibility.

Our full data set includes 390,733 new employees – hired either one year before or after each firm adopts automatic enrollment – from 279 unique firms. 56 percent of the people in our sample are "treated" employees, meaning they were hired at one of the firms in our sample after the company adopted automatic enrollment. Conversely, we refer to the employees hired in the year *before* each firm adopted automatic enrollment as our "control" population. The firms range in size from 222 to 65,699 employees. The median firm size in our sample is 1,104, although we have a much larger population of employees in our sample from a handful of firms with over 10,000 employees. The firms represent a broad range of industries including construction, retail, wholesale trade, transportation, real estate, communications, scientific services, manufacturing, finance, insurance, professional services, technology, administrative support, education, health care, entertainment, and food services. All firms in our data adopted automatic enrollment for defined contribution (DC) plans between 2010 and 2016. The most common auto-enrollment default rate in our sample is three percent (46 percent of firms implement this rate), although a majority of the firms (54 percent) adopt a different default rate, ranging from one to six percent.

The raw data reveal a sharp increase in employee contributions to DC plans after companies adopt automatic enrollment. **Figure 14** below shows the percent of newly hired employees who contributed a positive, non-zero percent of wages to their company's 401(k) plan in the years before and after each company adopted automatic enrollment. As shown in the graph, the median enrollment rate at the firm level jumps from approximately 20 to 80 percent after the policy changes go into effect. The distribution of enrollment rates across the full sample also changes from a right-skewed distribution clustered at 20 percent before auto-enrollment to a left-skewed distribution with 401(k) participation rates clustering around 80 percent after adoption of automatic enrollment.

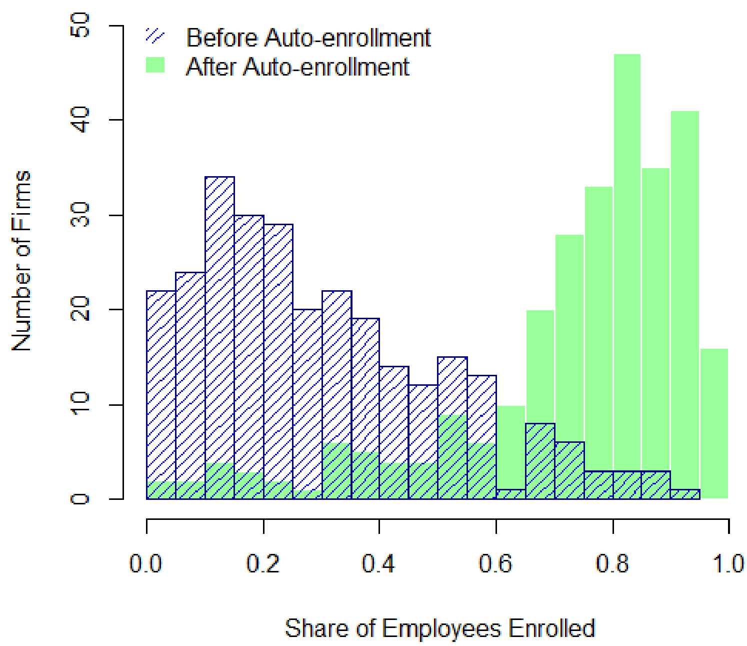

Figure 14: DC Plan Enrollment Before and After Automatic Enrollment

At the individual level, there are similar increases in employee contributions, although these are tempered by changes in the distribution of saving rates among plan participants. When we compare new employee participation rates before and after automatic enrollment, we can see that among those who participate in a firm-sponsored plan there is bunching at the default enrollment rate set by the firm among employees hired after automatic enrollment goes into effect. **Figure 15** below shows the distributions of employee saving rates before and after each firm adopted automatic enrollment, grouped by default rates from one to six percent.

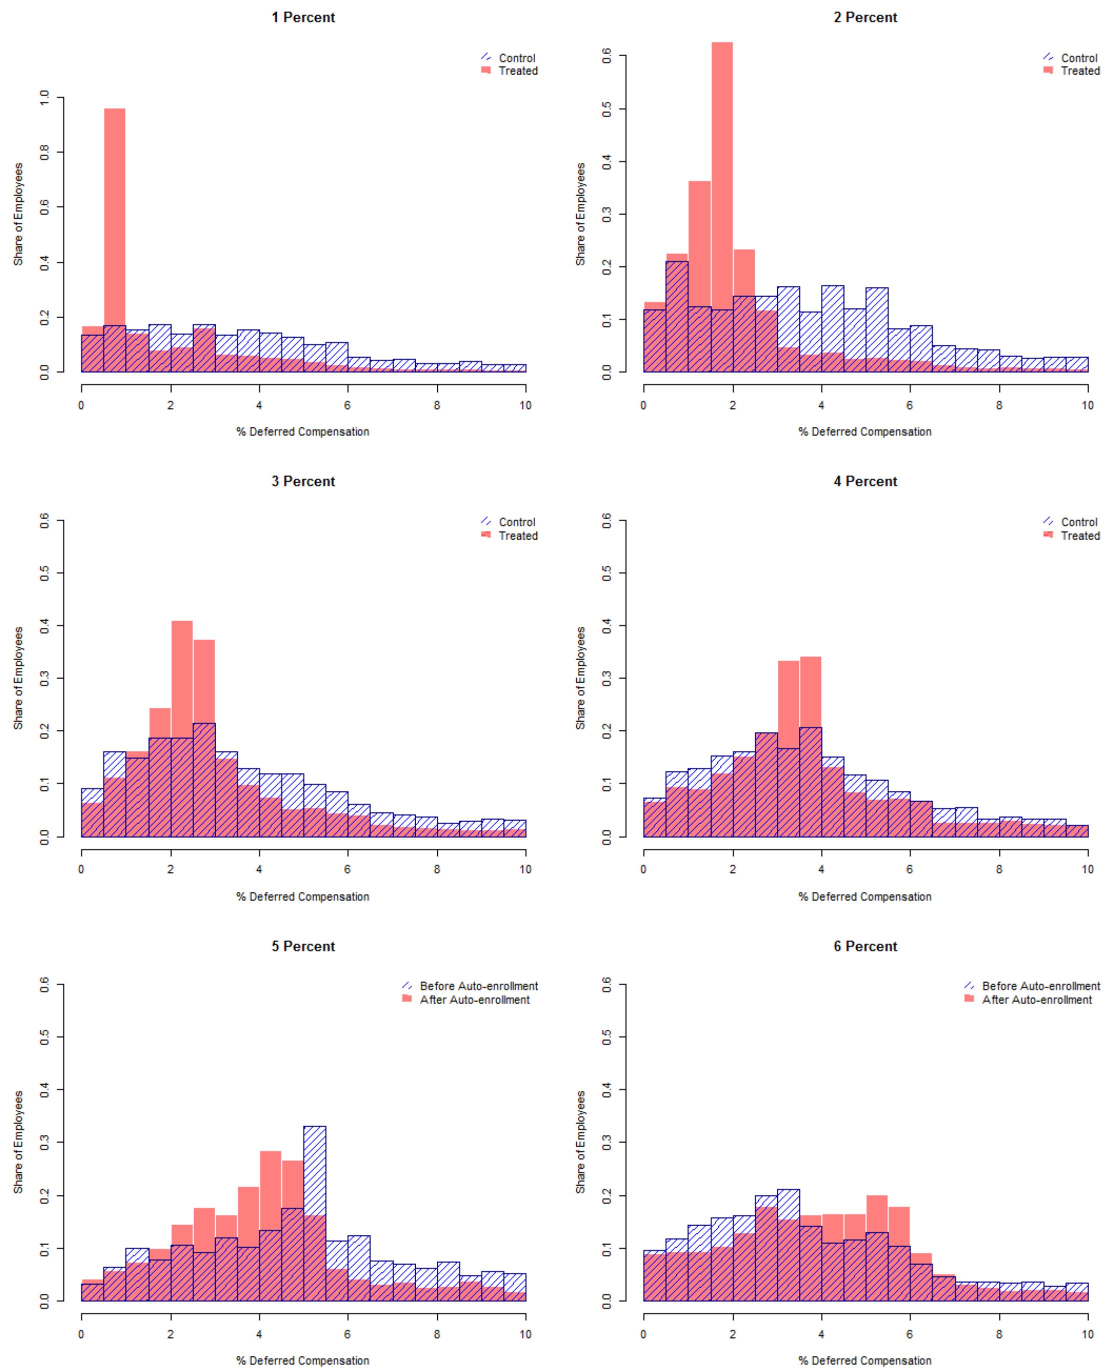

Figure 15: Employee Saving Rates Before and After Automatic Enrollment

The figure only includes only active participants in a company-sponsored retirement plan (individuals who are deferring a positive, non-zero amount of wages). As we can see, the percent of employees contributing at the default rate increases sharply in every group, and the percent of employees contributing at other rates decreases across the board. There is also evidence among firms with three to six percent default rates that the median contribution rate to the company-sponsored retirement plan is close to the default rate selected by the firm. As we discuss in greater detail below, this suggests that the default auto-enroll rate selected by each firm is endogenous to the saving behavior – and other characteristics – of its employees.

Based on the raw data, there does appear to be evidence that automatic enrollment increases employee participation in 401(k) plans, and that employees who save are more likely to contribute at the default rate than they would otherwise. However, there are other changes occurring over the same time period that may account for part of the perceived effect. In particular, we observe a statistically significant increase in wages between the control and treated groups, even after adjusting for inflation. There were also changes in employer discretionary matching policies at some firms in the same year that these firms adopt automatic enrollment.

These changes may explain at least part of the difference in enrollment since both variables are correlated with the probability of saving. Thus, in order to estimate the true magnitude of effect of automatic enrollment on employee saving, we use the estimation procedure outlined in the following section, controlling for correlated firm- and individual-level characteristics. We then decompose the effect based on characteristics at the firm and individual levels in order to determine how the policy may vary depending on employee and employer type.

2.2 Empirical Specification

Our primary estimation approach compares the contribution rates of employees hired in the year before each firm adopts automatic enrollment with the savings of individuals hired the year after the policy goes into effect, including control variables for firm and individual-level characteristics that may also affect individual saving behavior. The estimation procedure follows the specification below. Standard errors in all regressions are robust and clustered at the firm level.

$$Y_{i,f} = \alpha + \phi T_{i,f} + \beta \mathbf{X}_{i,f} + \delta_f + \varepsilon_{i,f}$$

In the equation, $Y_{i,f}$ represents either the individual employee's incidence of saving (a binary variable indicating whether or not the employee is deferring a positive, non-zero percent of wages to a firm-sponsored plan), or their contribution rate to the firm-sponsored retirement plan (measured as percent of wages). $T_{i,f}$ is a binary treatment variable that indicates whether the employee began working at the firm before or after the firm adopted automatic enrollment. The control variables, $\mathbf{X}_{i,f}$, include plan-specific variables like employer discretionary matching rates and individual characteristics like gender, wage level, and age. We also include a variable to indicate whether the firm had an eligibility requirement of more than three (but only up to six) months. Finally, δ_f is a firm fixed effect, which controls for variation coming from firm-level characteristics like the size of the firm, distribution of wages, and industry type. We do not include a year fixed effect because the combination of the firm fixed effect, the eligibility indicator, and the treatment variable perfectly predict the year of employment.

We include two measures of employer discretionary matching policies in our regressions:

"maximum match level" and "matching rate." The matching level is the maximum percent of employee contributions that the firm will match (at some rate), and the matching rate is the rate at which firms match employee savings at the maximum level. For example, if a firm matches 50 percent of employee contributions up to six percent, then employees at that firm would have a maximum match level of 0.06, and a matching rate of 0.5. Similarly, if a firm matches 100 percent of contributions up to three percent and 50 percent of contributions between three and five percent, then the maximum match level for employees at that firm is 0.05 and the matching rate is 0.8.

We include both matching variables because some firms change their matching policy around the same time that they adopt automatic enrollment. There are also some policies that set a fixed dollar maximum matching level, which equates to a different percent of wages depending on the salary of each employee. However, it is important to keep in mind that most of the firms in our sample to not change their discretionary matching policies when they adopt automatic enrollment, and do not have fixed dollar limits on contributions. Thus, all of the variation in our estimates on matching comes from a few firms that meet one of these two criteria.

Identification relies on the assumption that employees who join a firm that has implemented automatic enrollment are not more prone to save than employees who join a firm that does not automatically enrolled its employees in a DC plan. We generally assume that other factors – such as job responsibilities, wages, health care, and other benefits – are the deciding factors that determine whether an employee accepts a position at a firm, and that employees put little to no emphasis on whether the firm will automatically enroll them in a retirement plan when they are making their decision. The assumption we need for identification is actually weaker: we assume that even if employees *do* take automatic enrollment into consideration, that employees who are more likely to save are also not more likely to accept a job based on the firm having an auto-enrollment policy.

We cross-reference the results of these regressions with a difference-in-differences analysis, comparing changes in saving between the two groups of employees (those automatically enrolled and those not) before and after they are hired at each firm. For this portion of our analysis we use the following specification:

$$Y_{i,f} = \alpha + \gamma H_{i,f} + \phi T_{i,f} + \eta H_{i,f} \times T_{i,f} + \beta \mathbf{X}_{i,f} + \delta_f + \varepsilon_{i,f}$$

In the equation above, $H_{i,f}$ is a binary variable indicating whether the saving rate is measured before or after the individual is hired by firm f (one of the firms in our sample that adopts automatic enrollment between 2010 and 2016), and the coefficient of interest is η, the estimated effect of the interaction term $H_{i,f} \times T_{i,f}$. This is the estimated effect of automatic enrollment on the treated population after they are hired by company f. We use this specification only for robustness as there are some issues with comparing saving rates before and after hire for each employee. We do not know any information about the policies of the company each employee works for before joining company f. So, if the prior employer also happens to automatically enroll its employees in the company-sponsored retirement plan, or if they have an eligibility requirement of more than six months, we will not be able to control for these things in our estimates. Still, as explained in the following section, we do find that the results are nevertheless quite close to those of our preferred specification.

We check to see if our treated and control populations differ in meaningful ways by comparing individual- and firm-level statistics for the two groups, including wages, firm size, gender, and age. In order to do this comparison, we run a series of t-tests to evaluate whether the differences in each characteristic are statistically significant at the five percent level. The results of this analysis are presented in **Table 13** below. As shown in the table, the differences in mean age and percent female between the two groups are not statistically significant. However, the treated group does have slightly higher wages on average, even after adjusting for inflation. This suggests that firms may adopt automatic enrollment during a period of growth.

Table 6: Balance Test: Control and Treated Population Differences

Individual-Level	**Mean Control**	**Mean Treated**	**95% CI (Treated-Control)**	
Wages	$33,889	$35,342	$868	$2,038
Age	38	38	-0.1	0.0
Percent Female	50.7%	50.6%	-0.004%	0.002%
Auto-Enrolment Rate	3.2%	3.2%	-0.00%	0.01%
Firm-Level	**Mean Control**	**Mean Treated**	**95% CI (Treated-Control)**	
Firm Size	2,635	3,117	-1,552	589
Mean Wage	$49,978	$50,127	-$5,271	$5,570
Median Wage	$37,520	$38,761	-$3,051	$5,534

When looking at differences between the treated and control populations at the firm level (also in **Table 13**), we generally do not see statistically significant differences in firm size or wages (at either the mean or the median) before and after automatic enrollment. The statistically significant differences in wage growth that we see at the individual level are thus likely explained by greater-than-average growth of a few large firms in our data. The individual-level tests are essentially weighted comparisons of firms, and not necessarily reflective of a general trend among companies adopting automatic enrollment.

Still, the growth in wages at large firms in our sample may affect our estimates if they are not taken into account. Simply looking at a comparison of averages may overstate the effect of the policy if looking at one firm or a handful of large companies. We attempt to control for this by regressing on potential correlates like employee wages, and employer matching policies. We also include a firm fixed effect in our regression to control for other firm-specific differences like distribution of wages and industry type.

2.3 Results

2.3.1 Employee Participation in Retirement Plans

We begin by estimating the effect of automatic enrollment on the percent of employees enrolled in 401(k) plans, which we refer to as the "likelihood of saving." We estimate the effect using the specification outlined in the previous section, where the dependent variable is an indicator for whether or not the employee deferred a positive amount of wages. We regress on three different samples based on employee tenure: employees who work at the firm for at least three months, employees who receive at least two W-2 forms from the company (we call this "one-year tenure" for simplicity), and employees who work at the firm long enough to receive three W-2s from the firm (we call this "two-year tenure" for simplicity).

All the regressions are conducted using data from the first year or second year the employee is at the firm (depending on the length of the firm's service requirement for eligibility), so these estimates do not measure how participation changes with tenure. Rather, they are comparisons of how automatic enrollment affects different groups of employees based on whether they continue working at that company for a shorter or longer period of time. For all three groups we run logistic regressions with robust standard errors clustered at the firm level.

The results are provided in **Table 7** below. The coefficients reported in the table are marginal effects evaluated at the means. We use the group of employees who remain at the firm for at least one year as our baseline estimate for the effect of automatic enrollment on the likelihood of saving. The reason for doing so is that by restricting our sample to one-year tenure employees, we exclude workers who are likely to opt out of automatic enrollment for reasons unrelated to saving, such as seasonal or short-term employees. Furthermore, using this specification ensures that all employees in our sample work at the same firm for at least two consecutive tax years, instead of just employees who have a three- to six-month eligibility period, who by construction are observed working at one of the firms in our sample for at least two consecutive tax years (since their savings are calculated in the year after hire).

Based on this specification and controlling for a number of correlates listed in the table below, we estimate that automatic enrollment increases employee participation in 401(k) plans by approximately 32.4 percentage points. The mean enrollment rate among employees in the control group with one year of tenure is 28 percent. Thus, our estimate is equivalent to a 116 percent increase in participation due to automatic enrollment. This measure is somewhat lower than the estimates of Madrian and Shea (2001) and Choi, et al, (2004), although it is also larger than the 19 percentage point effect estimated by Falk and Karamcheva (2019).

Part of the explanation for these differences is the inclusion of a more diverse group of firms and workers in our data. The estimates in Madrian and Shea are based on a single large health services corporation with over 30,000 workers. The same firm is one of three used in Choi, et al.'s

Table 7: Automatic Enrollment Effect on Likelihood of Saving (By Tenure Type)

	Likelihood of Saving (Logit Marginal Effects)		
	All Employees	**1-Year Tenure**	**2-Year Tenure**
Treated	0.266***	0.324***	0.426***
	(0.014)	(0.015)	(0.024)
Female	0.035***	0.038***	0.038***
	(0.008)	(0.009)	(0.005)
Age	-0.000	-0.000	0.000
	(0.001)	(0.001)	(0.000)
Age Squared	-0.000	-0.000	-0.000*
	(0.000)	(0.000)	(0.000)
Log Wages	0.129***	0.143***	0.154***
	(0.016)	(0.017)	(0.007)
Max Match Level	0.133	0.147	0.341**
	(0.165)	(0.165)	(0.113)
Matching Rate	0.257***	0.288***	0.208*
	(0.069)	(0.087)	(0.086)
3-6 Month Wait	0.202	0.528	1.209*
	(0.697)	(0.770)	(0.525)
Firm Fixed Effect	✓	✓	✓
Mean Control	0.210	0.280	0.343
	(0.407)	(0.449)	(0.475)
Observations	390,733	270,589	109,910
Number of Firms	279	279	279
Pseudo R^2	0.438	0.365	0.345

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

analysis of the policy. In order to test how our estimate might compare to these two studies we look at the effect of automatic enrollment in a similar group of three firms in our data. Specifically, we estimate the effect among three large corporations in the field of health services with employee populations of over 40,000, 17,000, and 11,000 respectively. Using our preferred specification, we estimate that automatic enrollment leads to a 45.1 percentage point increase in participation at these three firms. This is much closer to Marian and Shea's regression-adjusted estimate of 50.1 percentage points than our baseline estimate.

Conversely, the estimated effect measured in Falk and Karamcheva is likely lower than ours due to the relatively high enrollment rate of federal workers prior to implementation of automatic enrollment. The median participation rate in firm-sponsored retirement plans in our sample is close to 20 percent prior to adoption of automatic enrollment. In comparison, 60 percent of federal workers in Falk and Karamcheva's study were enrolled in the federal TSP plan prior to the policy change (Falk & Karamcheva, 2019). As we discuss further in the following sections, other differences in enrollment, firm size, wage level, and industry can cause the estimated effect to vary significantly.

When we cross-reference the results from our main specification with our difference-in-differences analysis we find that the estimated effect of automatic enrollment on employee participation is similar. As shown in **Table 30** in **Appendix G**, we find that automatic enrollment increases 401(k) participation by an average of 36.6 percentage points for employees with at least one year of tenure − only a 1.2 percentage point difference from the coefficient estimated using our preferred specification. The estimated effect on employees with two years of tenure is also quite close at 41.9 percentage points (compared to 42.6 percentage points in our preferred specification).

We also find that there are statistically significant effects on saving due to gender, wages, and discretionary matching policies. We find that women are more likely to save by an average of 3.8 percentage points, and there is a large increase in the likelihood of saving as wages rise − about a 14 percentage point increase for every one percent increase in wages. There is no estimated effect of age on probability of saving, although this may be due to the inclusion of wages in our regression, which are highly correlated with age. Finally, we do not see evidence that the matching level set by firms has a significant effect on 401(k) participation. The estimated effect is positive but not statistically significant at the five percent level. However, we do see that the rate at which employers match contributions has a large estimated effect of 52.8 percentage points. However, as mentioned earlier, this effect is driven by differences between employees at a small number of firms.

We are further interested in finding out whether adopting different auto-enrollment policies has varying effects on 401(k) participation. We begin trying to answer this question by running the same regression (with individuals that have one year of tenure) on groups of employees working at firms that adopt different automatic enrollment default rates. **Figure 16** below graphs the estimated effect of automatic enrollment on participation by default enrollment rate. As we can see in the figure, the increase in enrollment varies from 27.6 percentage points to 44.5 percentage points depending on the level of the default rate.

Interestingly, the effect is not monotonically increasing or decreasing based on the auto-enroll rate, although there are statistically significant differences between groups with each type of policy. For employees automatically enrolled at a rate of one or two percent, there is an associated increase in participation of 40 to 44 percentage points. In comparison, employees automatically enrolled at three to six percent have a much lower estimated increase in enrollment, between 29 and 35 percentage points. We provide the full results of these regressions in **Table 32** in **Appendix G**.

One might assume that the differences we observe in the effect based on default auto-enroll rates has to do with the way that employees react to the rate that each firm sets. For example, employees may be more likely to remain enrolled at a default rate of one percent than at a rate of six percent because the first policy represents a smaller decrease in their net take home pay. However, part of the reason we see this trend is that employees working at firms that adopt a five or six percent enrollment rate are much more likely to participate in the firm-sponsored retirement plan before the firm adopts automatic enrollment.

As shown in **Figure 16** below, when we measure the effect of automatic enrollment as a relative percent change, rather than in percentage points, we see a more stark trend: the change in enrollment relative to the previous level is much higher in firms with one to four percent default rates because they have lower enrollment rates prior to implementation of the policy. Note that the size of the circles in **Figure 16** represents the number of firms that adopt that rate.

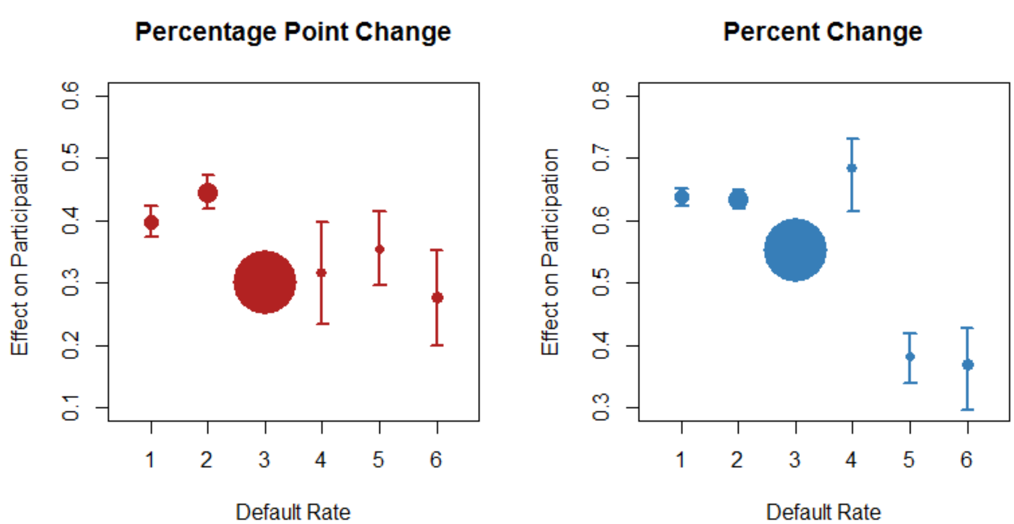

Figure 16: Effect of Automatic Enrollment on Employee Participation (by Default Rate)

When we compare employees at firms adopting each level of auto-enrollment, we notice other differences as well. As shown in **Table 8** below, firms with lower auto-enroll default rates tend to hire employees with lower average wages when compared to companies that opt for a higher rate of five to six percent. Companies with automatic enrollment rates in the middle (three or four percent) also have a larger number of employees on average, compared to firms that adopt either a low (one to two percent) or a high (five to six percent) default rate. Interestingly, companies with high rates also seem to have greater differences between employees at the 75^{th} and 25^{th} percentiles of the wage distribution. The inter-quartile range for wages at firms that adopt a five or six percent default rate is about 50 percent higher than that of firms with lower default rates.

Table 8: Firm Characteristics by Automatic Enrollment Default Rate

	Mean Wage	Mean Firm Size	Mean Wage IQR	# Firms
1% Auto-enroll	$45,680	1,521	$33,705	28
2% Auto-enroll	$44,665	1,989	$30,558	40
3% Auto-enroll	$50,039	3,115	$38,094	128
4% Auto-enroll	$42,855	3,893	$34,203	21
5% Auto-enroll	$58,301	1,139	$50,906	19
6% Auto-enroll	$56,560	2,094	$45,057	24

These results suggest that the default rate that companies decide to adopt is endogenous to the characteristics of their employees. It appears that companies tend to adopt auto-enrollment policies with higher default rates when their employees have higher average incomes and are more likely to contribute to 401(k) plans. In comparison, companies tend to adopt lower rates when their employees have lower average incomes. Larger firms are also more likely to adopt a middle-range default rate, possibly to accommodate a larger number of employees that may have different responses to automatic enrollment at higher or lower levels. We also have much smaller samples of firms that adopt an automatic enrollment rate different from three percent. The number of firms adopting other policies ranges from 19 to 40, compared to 128 firms that adopt a rate of three percent. This makes it difficult for us to draw broad conclusions about differences in the effect based on the default rate.

Things are further complicated by the regulations guiding exemption of plans from nondiscrimination testing. For a plan to be considered a Qualified Automatic Contribution Arrangement (QACA) − and qualify for exemption − firms must automatically enroll all employees at a default rate of at least three percent (IRS, 2020). So, firms that adopt policies of one or two percent default enrollment likely differ from those that implement enrollment rates of three percent or higher. In particular, firms that adopt a default rate of three percent or higher are more likely to have greater within-firm variation in wages, and larger gaps in wages between highly compensated and rank-and-file employees.

We will take into consideration these differences between employees grouped by firm default rates when we decompose the estimated effect of automatic enrollment in the following section. We will look specifically at differences in the effect among employees with earnings at percentiles of the wage distribution within their respective firms, and at differences across firms of varying size and with different average wages and industry types. This will give us a deeper understanding of how the effect of automatic enrollment varies between different types of employers and workers.

2.3.2 Deferred Compensation Percent

Next, we estimate the effect of automatic enrollment on how much individuals choose to save, measured as a percent of wages (we call this variable "deferred compensation percent"). As before, we begin by estimating the overall effect of automatic enrollment on savings, without considering variation in default rates. As shown in **Table 9** below, we find that automatically enrolling employees in a 401(k) plan leads to an average increase in deferred contributions of 0.93 percentage points, equivalent to approximately a 63 percent increase in savings (over a baseline saving rate of 1.48 percent).

We expect that most of this increase is coming from the effect of automatic enrollment on non-savers: if the policy causes an increase in participation, as we find previously, then this will have a positive effect on aggregate contributions. However, the increase in savings will be tempered if the average contribution rate among savers prior to the policy change is higher than the default rates set by firms. We define savers in this context as individuals who contributed to a firm-sponsored plan at their previous job, before they were hired at one of the companies in our sample.

Once again, when we cross-reference the estimated effect on savings from our main specification with the results from the difference-in-differences analysis we find similar results. Using this second method we find that automatic enrollment increases employee savings by an estimated 0.96 percentage points (compared to 0.93 percentage points in our main specification) for employees with one year of tenures, and by 1.23 percentage points (compared to 1.26 percentage points) for employees with two years of tenure. The full results from the difference-in-differences analysis is presented in **Table 31** in **Appendix G**. Despite the fact that we know relatively little about the firms that each employee works for prior to joining the companies that we include in our data, we find that the estimates using either specification are nevertheless quite close.

These findings are in line with previous studies looking at the effect of automatic enrollment on the overall level of saving. Once again, our estimate falls between that of Madrian and Shea (2001), who estimate an overall increase in savings of 2.2 percentage points, and Falk and Karamcheva (2019), who see a much lower increase of 0.6 percentage points (likely due to the propensity of federal workers to save at higher rates in the absence of auto-enrollment). We also find, as Falk and Karamcheva do, that the effect on savings increases with the automatic enrollment rate, as one would expect. As shown in **Figure 17** below, enrolling employees at a rate of three percent increases savings by an average of 0.9 percentage points, whereas automatically enrolling employees at five percent increases average savings by 1.44 percentage points.

Our analysis of the effect on saving levels goes one step further, testing whether the low increase is saving is simply a compositional effect − with more non-savers enrolling at the relatively lower default rate while savers continue to contribute at the same levels they did previously − or if there is also a negative nudge affecting savers (those who contributed to a DC plan at a prior job), prompting them to remain enrolled at a lower default rate rather than increase their deferred contributions to the same percent of wages they were saving under their previous employer. In order to evaluate this we estimate the effect of auto-enrollment on the saving levels of employees who contributed to a firm-sponsored retirement plan both before and after joining one of the auto-enroll firms in our sample. This should give us a better idea of whether automatic enrollment has some negative effects on saving at the same time that it increases participation among non-savers. The results of this analysis are presented in **Table 10** below.

When we look at the results among savers with one year of tenure, we find that there is actually an average *decrease* in the savings of 401(k) participants of approximately 0.9 percentage points due to automatic enrollment. This is equivalent to a 14 percent decrease in savings from a baseline rate of 6.34 percent among savers in the control group. Thus, we can conclude that the total average effect of 0.9 does not simply reflect an increase in the percent of employees contributing at a low default rate (relative to average saving levels among employees who opt in). This estimate also reflects a negative nudge on employees who saved at higher levels prior to being automatically enrolled in a firm-sponsored plan. As we see in **Figure 17** below, when we compare the effect between savers at firms that adopt different default rates, this negative nudge only disappears − meaning the estimate is still negative but no longer statistically significant − when firms institute a default enrollment rate closer to the average saving among savers of five to six percent.

Our estimates of the effect of each control variable on average saving rates − also provided in **Tables 9 and 10** below − follow the same general trends that we see in our previous results in **Table 7** (the effects on participation). Controlling for all other variables, we find that women save more on average (by approximately 0.20 percentage points), and that there is a small increase in percent of wages that people defer as they get older, with employees increasing their contributions by an average of 0.02 percentage points per year (although the estimates are not statistically significant).

Similarly, employees with higher salaries save more on average (by about one percentage point for every one percent increase in wages), and we see statistically significant differences in the firm fixed effect, as we expect. Although the effect of employer matching policies (both the matching rate and the match level) are still positive, we do not see a significant effect of these policies on employee saving rates at the five percent level. We also do not see statistically significant effects of working at a firm with a longer eligibility requirement of three to six months, although we continue to include this variable as it is directly correlated with the tax year in which we observe wages and 401(k) contributions.

Table 9: Automatic Enrollment Effect on Saving Rate (By Tenure Type)

	% Deferred Compensation (OLS)		
	All Employees	**1-Year Tenure**	**2-Year Tenure**
Treated	0.69***	0.93***	1.26***
	(0.13)	(0.10)	(0.15)
Female	0.20***	0.22***	0.23***
	(0.05)	(0.05)	(0.05)
Age	0.02**	0.02***	0.03***
	(0.01)	(0.01)	(0.00)
Age Squared	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)
Log Wages	1.00***	1.05***	1.03***
	(0.06)	(0.05)	(0.05)
Max Match Level	0.08	0.16	0.10
	(1.10)	(0.85)	(1.41)
Matching Rate	0.51	0.45	0.49
	(0.66)	(0.68)	(0.53)
3-6 Month Wait	6.76	1.32	0.35
	(6.95)	(6.11)	(6.90)
Firm Fixed Effect	✓	✓	✓
Mean Control	1.11	1.48	1.88
	(3.24)	(3.65)	(4.03)
Observations	390,733	270,589	109,910
Number of Firms	279	279	279
Adjusted R^2	0.251	0.244	0.230

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

Table 10: Automatic Enrollment Effect on Saving Rate Among Savers (By Tenure Type)

	% Deferred Compensation (OLS)		
	All Employees	**1-Year Tenure**	**2-Year Tenure**
Treated	-0.98***	-0.90***	-1.14***
	(0.51)	(0.15)	(0.23)
Female	0.42***	0.46***	0.38**
	(0.09)	(0.10)	(0.12)
Age	0.08***	0.08***	0.08***
	(0.01)	(0.01)	(0.01)
Age Squared	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)
Log Wages	0.59***	0.60***	0.42***
	(0.22)	(0.23)	(0.28)
Max Match Level	0.28	0.68	-0.39
	(1.56)	(1.84)	(2.50)
Matching Rate	-1.22	-1.17	-0.27
	(0.96)	(0.97)	(0.61)
3-6 Month Wait	2.86	1.30	4.58
	(9.73)	(9.65)	(14.88)
Firm Fixed Effect	✓	✓	✓
Mean Control	6.31	6.34	6.46
	(5.31)	(5.27)	(5.24)
Observations	48,468	43,755	19,482
Number of Firms	279	279	279
Adjusted R^2	0.131	0.131	0.126
Note:			*p<0.05, **p<0.01, ***p<0.001

As before, we run the same regression for each group of individuals based on the default auto-enrollment rates set by firms in order to see how total savings rise and fall based on the policy each firm adopts. **Figure 17** below shows a much clearer pattern between the default contribution rate set by the firm and the intensive margin of saving. As the default rate increases, so does the effect on the rate of deferred contributions. Setting the default rate at one percent is associated with an average increase in saving of 0.6 percentage points. In contrast, setting the default rate at six percent is associated with an increase in saving of 1.5 percentage points: more than double that of the increase associated with the one percent policy. Since the baseline saving rate of each group differs (the mean among those in the control group is 1.2 percent and 2.3 percent, respectively), the relative differences are slightly smaller, and for the most part not significant. Relative differences are presented in second graph in **Figure 17** and the full results of these regressions are presented in **Table 33** in **Appendix H**.

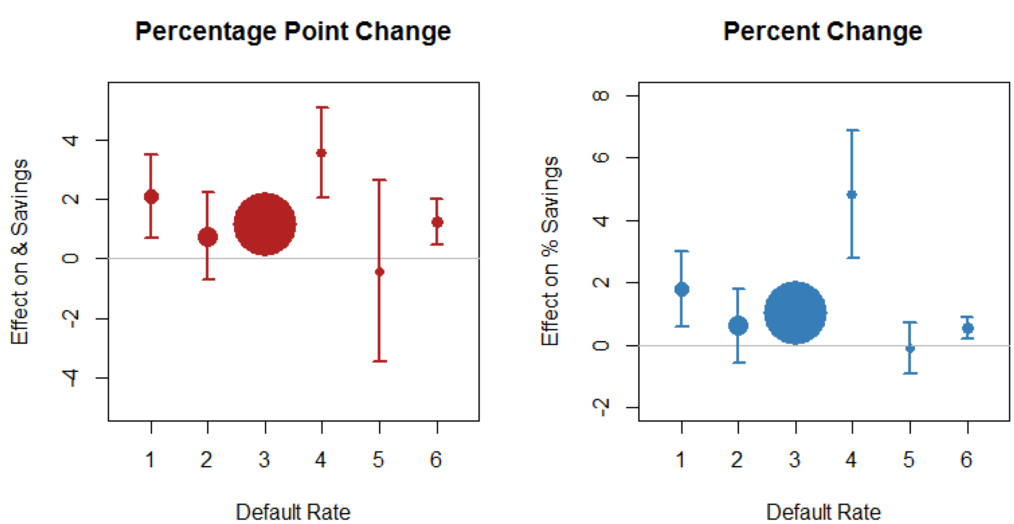

Figure 17: Automatic Enrollment Effect on Percent Saved

When we look at differences in the effect of auto-enrollment on saving among savers, we get a better idea of how setting the default rate at a lower level offsets increases in participation due to the adoption of automatic enrollment. When comparing groups enrolled at different default rates, we see that the effect is negative and statistically significant at the five percent level for all employees automatically enrolled at less than five percent. Participants automatically enrolled at one percent have an estimated decrease in saving of 2.3 percentage points, whereas plan participants enrolled at six percent have a much smaller decrease close to zero.

This indicates that although employees are a lot more likely to participate after adoption of automatic enrollment, they are also more likely to remain enrolled at the default rate, which is often lower than the rate people select when they opt *in* to a retirement plan. Once again, the relative differences are smaller given the higher average levels of pre-automatic enrollment savings in firms with higher default rates. However, it is important to note that differences in saving are not as pronounced among savers. Savers in the control group actually contribute to retirement plans at similar rates across firms that implement different defaults, with average contributions ranging from 5.6 to 6.7 percent. Full results of this analysis are presented in **Table 34** in **Appendix H**.

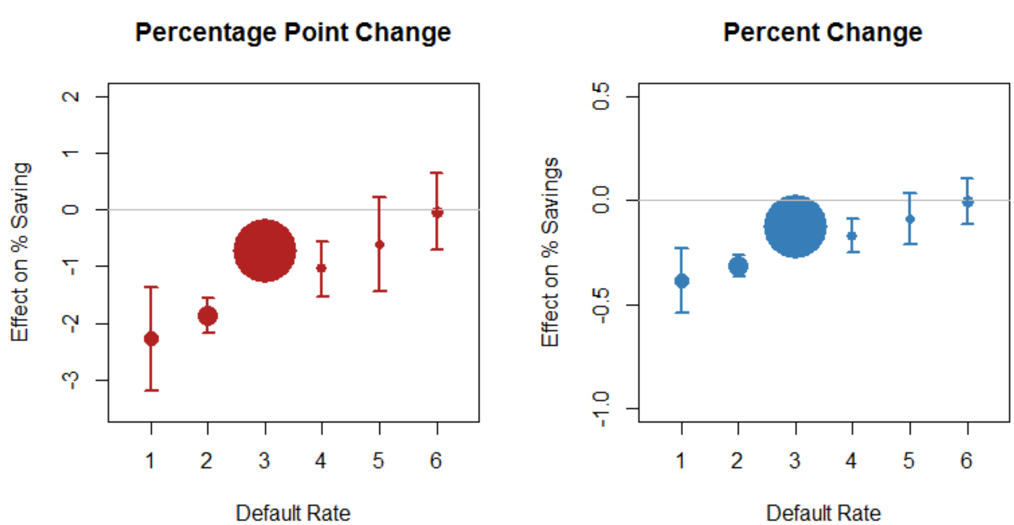

Figure 18: Automatic Enrollment Effect on Percent Saved (Among Savers)

In the following sections we look more closely at differences in the effect between individuals in different percentiles of wages, both within and across firms, to get a better idea of how the policy affects employees at different parts of the wage distribution. We also explore differences in the effect between firms with different characteristics including size, average wage level, and industry.

2.4 Heterogeneous Effects: Employees

2.4.1 Within-Firm Wage Effects

In addition to average treatment effects, we explore heterogeneous responses within and across firms by earnings. As explained in detail below, when we compare differences in the estimated effect between individuals sorted by their wage level within each firm, we find that both the propensity to save and level of saving decrease in relative wages. We see larger differences in the propensity to save, although the effect on saving level also varies between some groups (with differences statistically significant at the five percent level). However, the difference in the effect is only evident when comparing increases relative to the baseline means for each group. In other words, we do not see much variation in the effect of automatic enrollment when we look at the percentage point increase across groups, but the increases for individuals at the bottom of the wage distribution represent larger relative changes since these employees have lower participation and saving levels to begin with.

We first compare the outcomes of individuals at different income levels within firms to see if there are distributional differences in the effect between people in the top and bottom percentiles of wages within each firm. We do this by calculating the percentile at which each individual falls in the estimated distribution of income for their given firm and tax year, and then sort individuals into five categories based on this percentile value. Since all of the employees in our data are new employees, we only observe partial-year wages for most of them in the first year that they are hired (the year in which we calculate their saving level as a percent of wages). Thus, in order to determine where each employee falls within their respective firm's wage distribution, we compare their wages in the year after we observe them joining the firm. In order to further ensure that we are not using partial-year wages in this measure, we also restrict our sample to employees who are still working for the same firm two tax years after their year of hire.

As such, our sample in this analysis is limited to employees with two years of tenure at the same firm, which amounts to about one third of our total population. We are still comparing average savings in the year of hire (or the year after hire if the firm has a three- to six-month eligibility period), but we are restricting that sample to individuals with three consecutive W-2 returns at the same firm in order to get a better measure of where they fall in the wage distribution. As shown in **Table 7** above, the average estimated effect for this group of employees is a bit higher than the estimated effect for one-year tenure employees. We find that for two-year tenure employees automatic enrollment increases 401(k) participation by 42.6 percentage points. It is also important to keep in mind that the average wages of new employees tend to be lower than average wages at each firm, because new hires tend to be younger and less experienced. For example, the wages of employees in the third quintile of within-firm wage distributions fall between the 32nd and 44th percentiles of wages at their firms. **Table 11** below provides the cutoff points for each quintile of workers based on where they fall in the within-firm wage distribution.

Figures 19 and 20 below graph differences in the effect of automatic enrollment on 401(k) participation and saving levels (respectively) between employees at different quintiles in the wage distribution within firms. As before, we include two graphs in each figure: one for the absolute change in enrollment (in percentage points) and the other for the relative change (in percent).

Table 11: Quantile Cutoff Points Based on Within-Firm Wage Distributions

	0%	20%	40%	60%	80%	100%
Percentile	0.00	0.22	0.32	0.44	0.60	1.00

At first glance there do not appear to be large differences in the effect of automatic enrollment on 401(k) participation when looking at employees at different percentiles of the within-firm wage distribution. Although there is a pattern in the effect, which decreases as the percentile bin increases, the differences are not statistically significant at the five percent level. However, when we look at the relative change in enrollment, we do see variation in the effect due to differences in the baseline enrollment rates of each group. Employees in the bottom quintile of wages see an estimated increase in enrollment of 42.3 percentage points, which equates to a 150 percent increase in enrollment (from a baseline participation rate of 28.2 percent). In contrast, employees in the top quintile of wages see a similar estimated increase in enrollment of 36.9 percentage points, but this is equivalent to a much lower relative increase in enrollment of 82 percent (from their base participation rate of 44.7 percent). The full results comparing these three groups are presented in **Table 35** in **Appendix H**.

As we can see in **Figure 20** below, there also do not appear to be statistically significant differences in the effect of automatic enrollment on the level of saving when comparing individuals at different quintiles of the within-firm wage distribution. However, as with participation, there do appear to be differences in the relative effect since individuals in the top two wage quintiles save at higher rates prior to the policy change. The relative increase in saving among those in the lowest quintile is 84 percent versus 49 percent in the highest quintile. Interestingly, there do not appear to be large differences in the estimated effect among those in the first three quintiles. This is mainly due to the fact that the baseline saving rates of the three groups only differ by about 0.1 percentage point. The full results of each regression are presented in **Table 36** in **Appendix I**.

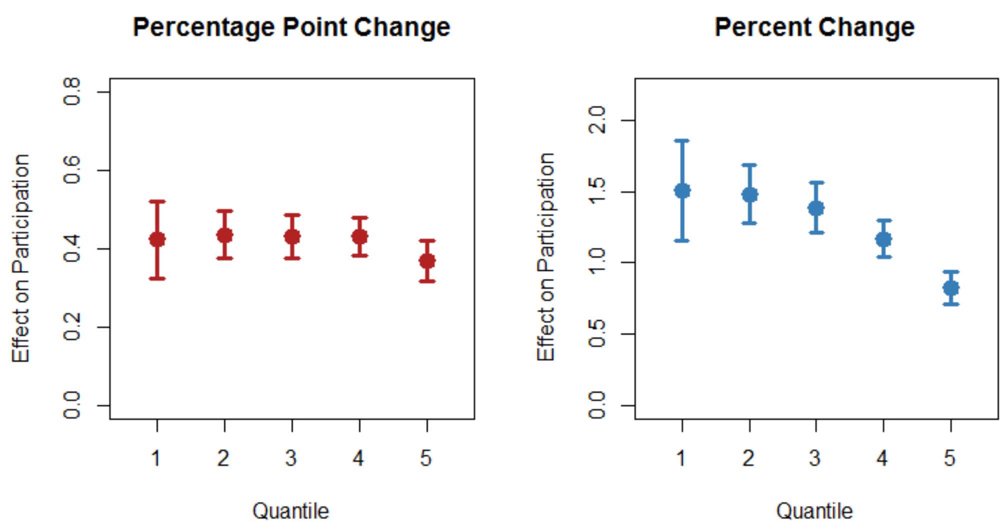

Figure 19: Effect of Automatic Enrollment on Participation by Wage Level (Within Firms)

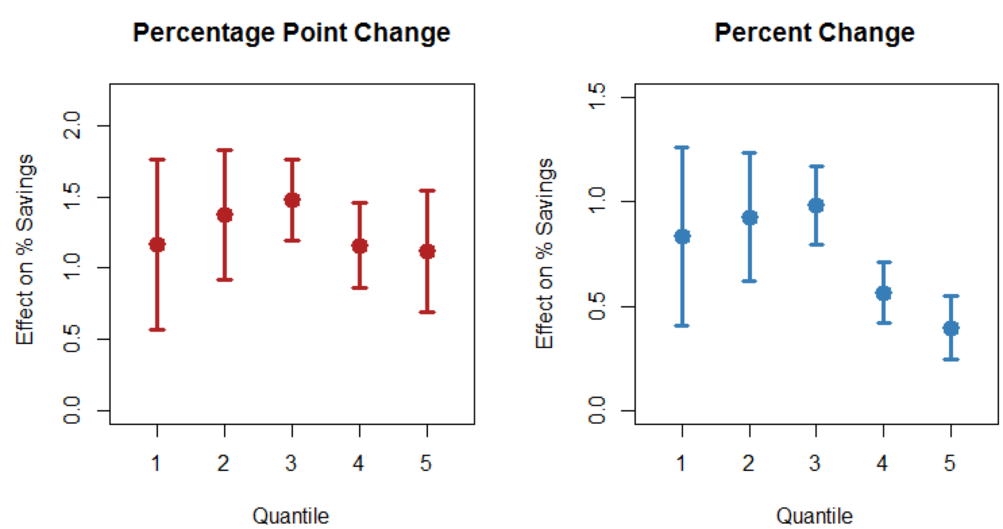

Figure 20: Effect of Automatic Enrollment on Percent Saved by Wage Level (Within Firms)

In general, we do see downward trends suggesting that the effect of automatic enrollment on both participation and saving rates is higher for individuals with lower incomes. However, this conclusion is limited to relative effects. Since individuals with lower wages are less likely to contribute to a 401(k) plan before automatic enrollment, and contribute a lower percent of wages to retirement plans on average, a similar percentage point increase in saving constitutes a larger relative change for these employees when compared to their counterparts at the top of the within-firm wage distribution. However, if we are concerned with whether automatic enrollment decreases inequality between workers at different income levels, then what matters more is the absolute change in participation and contributions. In this case, we do not find that automatic enrollment closes the gap in saving between workers at the top and bottom of the wage distribution. Instead, all employees see a similar increase in participation/savings.

2.5 Heterogeneous Effects: Firms

In order to further decompose the effect of automatic enrollment − to better understand how the policy affects different groups of employees − we look at the differences of the effect based on varying firm characteristics. We look specifically at differences in the effect based on three types of firm-level variation: number of employees (firm size), mean wage among employees, and industry type. To conduct the first two parts of this analysis (heterogeneous effects on firm size and mean wage) we group firms into quintiles based on each characteristic, and then run regressions on the full population of employees in each group of firms. Since our analysis is based on grouping firms in equally populated bins based on number of workers or mean wage, our sample size for each group of employees varies from 15,474 to 176,960 for firm size, and from 32,424 to 83,666 for mean wage. **Table 12** below lists cutoff points for each grouping by size and average wage. The cutoff points are based on characteristics of the firm before adoption of automatic enrollment.

Table 12: Quantile Cutoff Points: Firm Level

	0%	20%	40%	60%	80%	100%
Firm Size	222	546	897	1,414	2,661	58,188
Average Wage	\$9,605	\$23,853	\$34,791	\$50,356	\$71,439	\$239,918

2.5.1 Firm Size Effects

When we estimate differences in the effect of automatic enrollment on 401(k) participation by firm size, we find that there are no significant differences in the effect among employees in the bottom four quintiles. However, the effect is lower for individuals in the largest firms (those with more than 2,661 employees). Both the absolute and relative differences in the effect are small but statistically significant compared to groups of employees at other firms. The estimated percentage point increase in enrollment for individuals working at large firms is 28.1 percentage points, compared to 39.4 percentage points for workers in the fourth quintile based on firm size (1,414 to 2,660 employees). The results are graphed in **Figure 21** below. Although these estimates are limited to workers observed with one year of tenure, part of the reason for this difference could be that large firms employ a larger number of part-time and seasonal workers. Although we only include firms in our data that automatically enroll all employees, these individuals may still be more prone to opting out of a company-sponsored plan if they do not intent to stay at that company for a long time. Full results from these regressions are presented in **Table 37** in **Appendix J**.

Similarly, we do not see statistically significant differences in the effect of automatic enrollment on saving rates when comparing employees at firms of different sizes. **Figure 22** below graphs the estimated effect by group based on firm size. As we can see in the figure, the effect is still slightly lower for those in the top quintile, but the differences are not statistically significant at the five percent level. The rest of the estimates are quite close together for both absolute and relative measures of the effect.

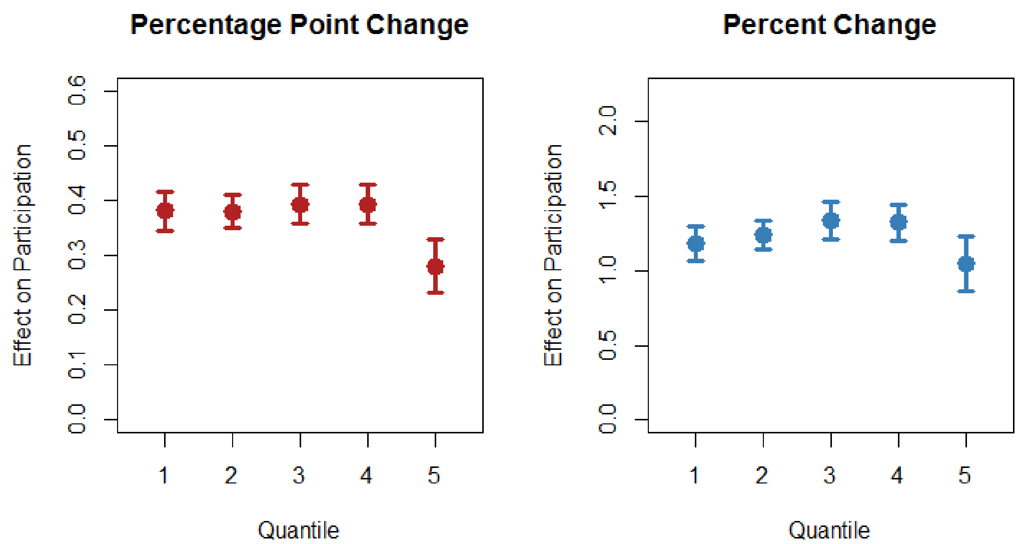

Figure 21: Effect of Automatic Enrollment on Employee Participation by Firm Size

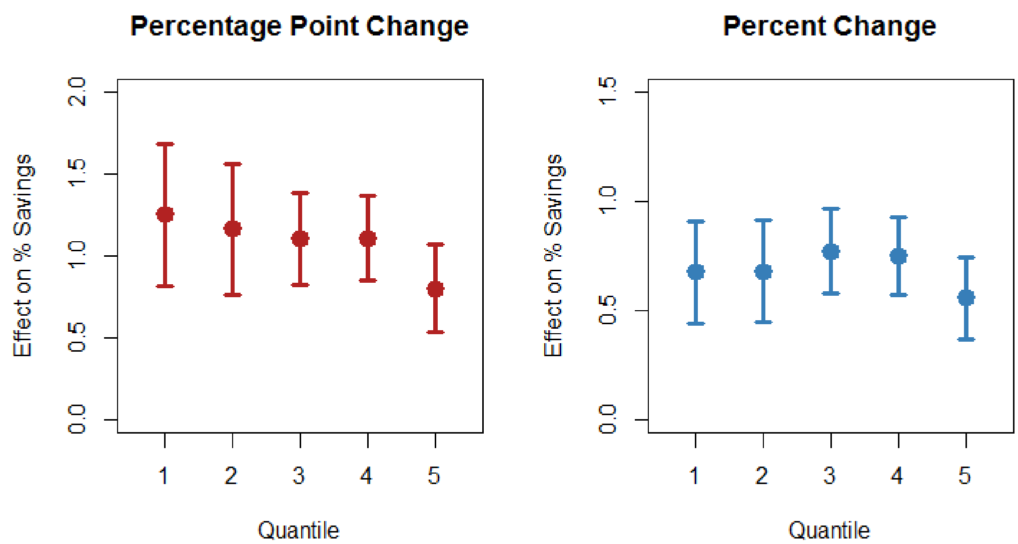

Figure 22: Effect of Automatic Enrollment on Employee Saving Rates by Firm Size

Since we know that there is variation in the effect on saving based on default enrollment rates, and that the rate firms adopt is correlated with the size of the firm, we want to make sure that these results do not reflect some underlying pattern in the policy itself. However, when we compare average rates across firms and across individuals in each group, we see that the mean enrollment rate does not vary much and is close to three percent for all groups (ranging from 2.83 to 3.38 percent). This provides further evidence that the size of the firm does not seem to have much of an impact on the size of the effect.

2.5.2 Mean Wage (Firm-Level) Effects

Next, we look at how the effect of automatic enrollment varies between firms with different mean wages. Once again, we group the firms into five quantiles based on the mean wage among new employees at each firm. This allows us to determine if the differences we see in the effect on participation and saving – and the differences between our baseline estimate and those of other studies – are due to variation in the average type of employee working at each firm. One might expect, for example, that firms with relatively higher mean wages may be comprised of more highly skilled workers, which may cause variation in the effect of automatic enrollment on both participation and saving rates.

Figure 23 below graphs differences in the estimated effect of automatic enrollment on 401(k) participation based on mean firm wages. As we can see in the figure, the absolute effect on participation appears fairly uniform, but the relative effect is much larger for firms with lower average wages. In other words, you see close to the same percentage point increase in participation across firms with different mean wages, but this represents a much larger increase for employees at low-wage firms, given that their auto-enrollment rates are much lower prior to the policy change.

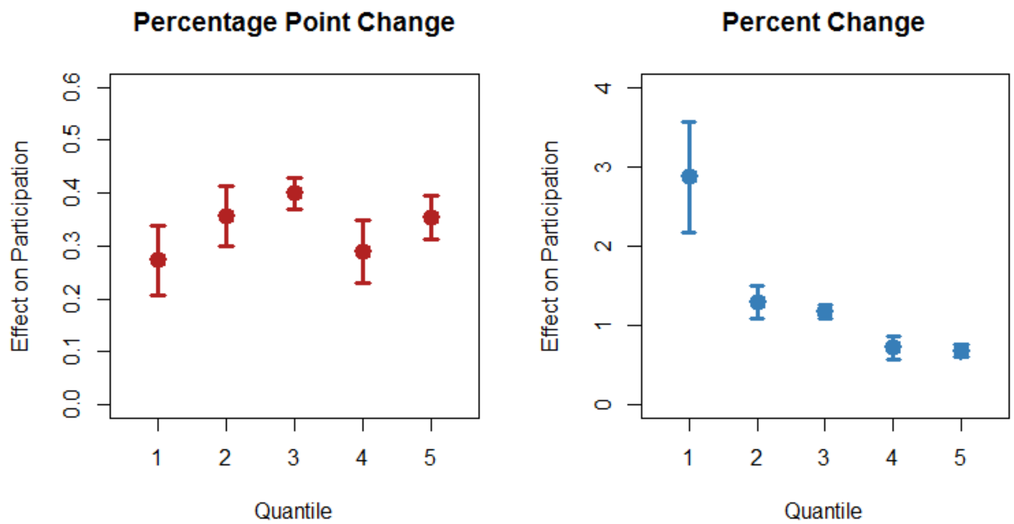

Figure 23: Effect of Automatic Enrollment on Employee Participation by Average Wage (Firm Level)

As shown in the figure, enrollment among employees with an average wage under $23,853 is only 9.5 percent pre-automatic enrollment, compared to 51.7 percent at firms with average wages over $71,439. Those in the top quintile actually see a slightly larger percentage point increase in enrollment – 35.3 percent versus 27.3 percent for the bottom quintile – but the large difference in baseline participation rates is equivalent to a large difference in relative changes in participation, with enrollment at firms with lower average wages nearly tripling as a result of the policy change.

We see a similar trend when comparing the effect of automatic enrollment on the level of saving across employees at firms with different average wages. **Figure 24** below graphs differences in the effect by quantile of mean wages at the firm level. As we can see, the absolute effect on saving is generally the same among employees at firms with varying average wages, although employees at firms in the first quintile do have a lower estimated effect. However, due to large differences in the mean saving rates among employees prior to automatic enrollment, we get a similar distribution of relative effects to those in **Figure 23** above. As discussed previously, this is because employees of firms in the bottom quintile of wages have a much lower baseline saving rate of 0.4 percent relative to firms with higher average wages. As such, employees working at firms in the bottom quintile have an estimated relative increase in savings of 160 percent due to automatic enrollment. In comparison, those in the top quintile (with a baseline saving rate of 3.4 percent) see a relative increase in savings of just 26 percent.

Despite large differences in the effect of automatic enrollment relative to baseline levels of participation and savings, the absolute difference in automatic enrollment appears to be similar across groups clustered by wage and population at the firm level. This tells us that although employees working at relatively smaller firms with lower average wages see a large percent increase in their participation (and through that effect a larger relative increase in saving), that the effect of automatic enrollment at the level of participation across different types of firms – based on these two characteristics – is essentially the same. Full results for all of these regressions are provided in **Tables 44 and 45** in **Appendix J**.

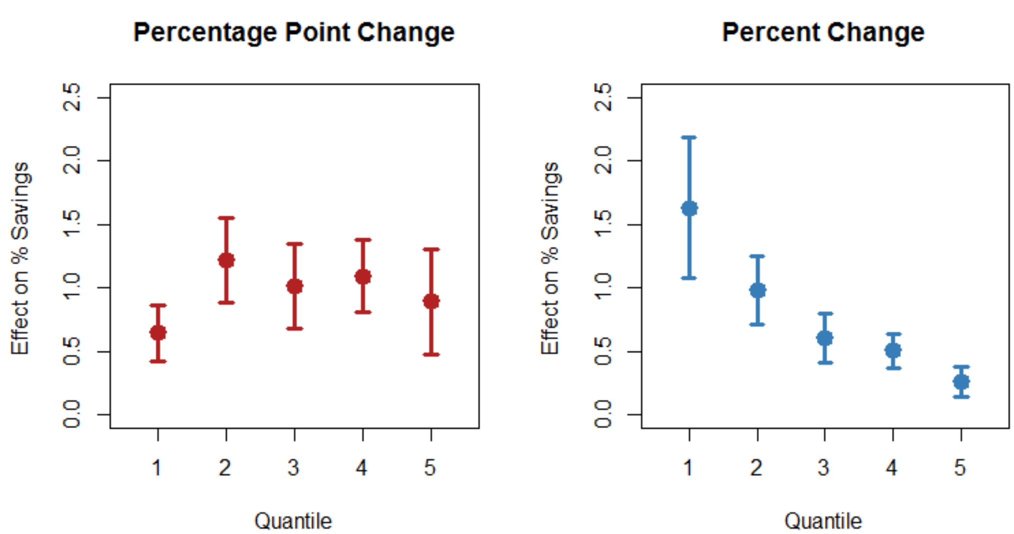

Figure 24: Effect of Automatic Enrollment on Employee Saving Rates by Average Wage (Firm Level)

However, as we will explain in greater detail below, we do start to see differences in the level of the effect (in percentage points) when we estimate how automatic enrollment affects employees working at firms in different industries. These differences help to explain why we end up with lower or higher baseline estimates than those in previous studies of the effects of auto-enrollment.

2.5.3 Industry Effects

Finally, we examine variation in the effect of automatic enrollment based on industry type. To conduct this analysis, we group firms by two-digit North American Industry Classification System (NAICS) codes and run the same regression on each sub-group. A full description of each NAICS code is provided in **Table 41** in **Appendix K**. We use the broadest possible classification (two-digit NAICS codes) to ensure that the sample size in each group is sufficiently large enough to draw conclusions abou the effect of automatic enrollment on employees working at each group of firms. The total number of firms and employees in each group are provided in **Table 42** in **Appendix K** with the estimated effects from each regression.

Figure 25 below graphs the effect of automatic enrollment on 401(k) participation for employees working in different industries. As shown in the figure, we find that the effect on enrollment varies widely across industries, from an increase of 20 percentage points (entertainment) to nearly 50 percentage points (education). However, it is difficult to tell why we find greater effects in some industries compared to others. For example, we see that high-skill industries like education and health care have some of the largest increases in participation due to auto-enrollment, but we see similar increases in participation after auto-enrollment among employees working in relative low-skill industries like manufacturing. In comparison, industries like retail have a relatively low effect – as we might expect to see in an industry with more part-time and seasonal workers –

but it is similar to the effect on employees in firm management, which presumably has a greater percentage of salaried employees.

We find even larger differences in the effect of automatic enrollment on participation relative to baseline (pre-auto-enrollment) rates among each control group, ranging from about 80 percent (news and information) to over 1,500 percent (entertainment, which is not pictured in **Figure 25** for purposes of scale). In particular, there seems to be a large relative increase in participation among employees working in the sectors of retail, real estate, administrative services and waste management, education, and health care. In some cases − for example, health care − this is due to a relatively higher effect on the level of participation, and in others − like retail − it is due to relatively low enrollment rates prior to the policy change.

We further find that in industries like manufacturing and professional services the relatively high baseline enrollment rate equates to a lower relative increase in participation resulting from to automatic enrollment. Again, these are industries with similar effects but very different types of workers. Interestingly, the largest relative increase in in saving is among employees of the entertainment industry − comprised primarily of casino workers in our data − who have an extremely low participation rate of 1.3 percent prior to the policy change.

However, this is same group of employees also has the lowest estimated percentage point increase in retirement plan participation due to automatic enrollment. Similarly, automatic enrollment has a relative low effect on 401(k) participation of workers in the retail sector (27 percent), but have one of the highest relative increases due to low enrollment rates among employees in the control group.

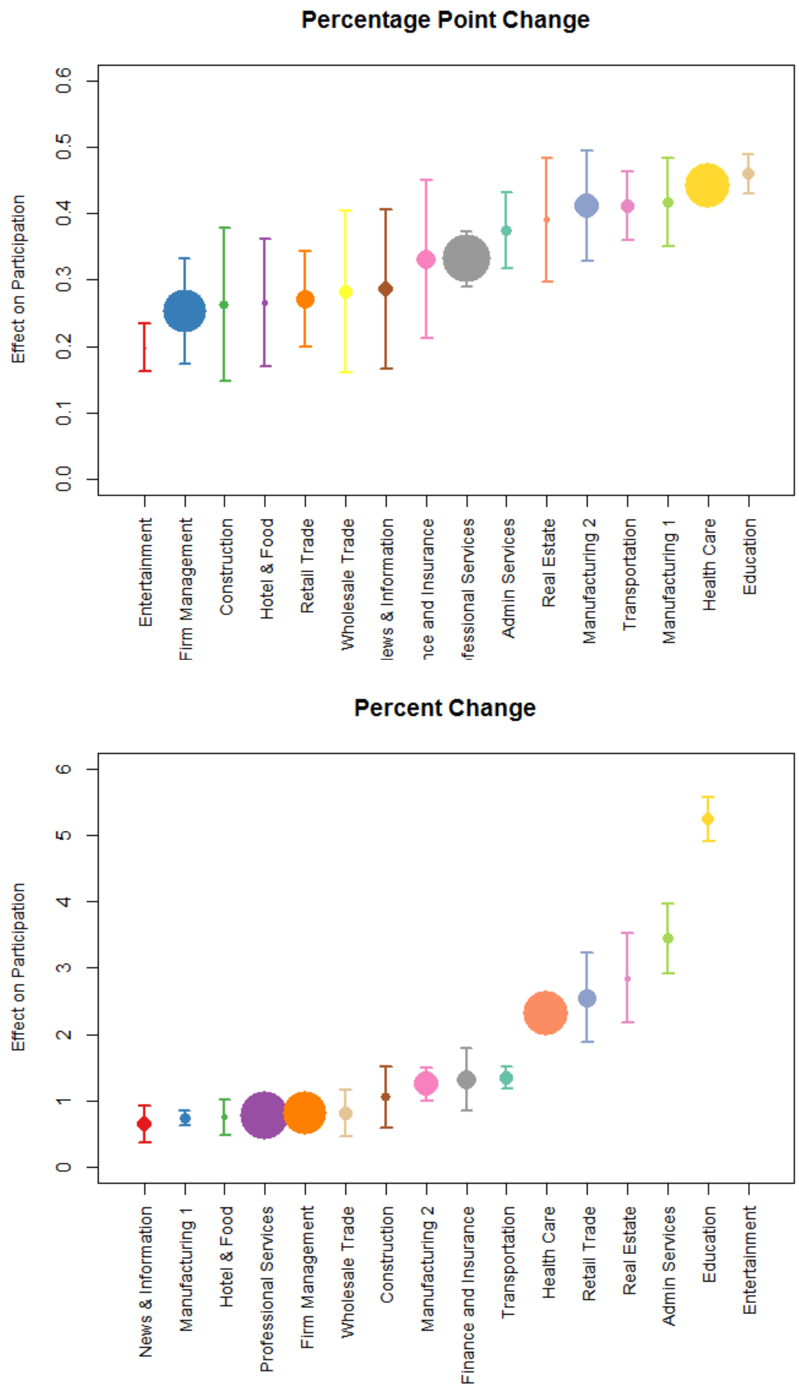

Figure 25: Effect of Automatic Enrollment on Employee Participation by Industry

We also see large differences between groups of employees clustered by industry when we look at the effect of automatic enrollment on average saving rates. **Figure 26** below graphs these estimated effect of automatic enrollment on each group of employees clustered by industry. As shown in the figure, the effect ranges from an increase of approximately 0.4 percentage points (retail and admin services) to nearly 2.5 percentage points (real estate). We see the largest increase in the level of savings due to automatic enrollment among employees in the industries of real estate and wholesale trade. Part of the reason we see this large increase in savings among those in the real estate sector is that they have a relatively high average default enrollment rate of 5.7. However, despite experiencing a similar effect on savings, the average enrollment rate of employees in the wholesale trade industry is below average, at just 2.7 percent. This group of employees also has a relatively low increase in participation due to automatic enrollment, so it is unclear why their increase in savings is among the highest compared to other industries.

We find that employees in health care, education, and manufacturing also have relatively high saving rates − likely because of their relatively higher increases in participation. Those with the lowest average increases in saving work in the fields of retail, entertainment, and administrative services. With the exception of admin services, saving patterns for these groups do seem to follow trends in the effect of auto-enrollment on participation. Those who work in administrative services are also enrolled at the lowest average default rate among any group at 2.1 percent.

When looking at relative changes in the saving rate we see similar results, with relative increases ranging from about 0.3 percent (professional services and manufacturing 1) up to 4.8 percent (real estate). Note that the largest relative increase in savings is actually among employees in entertainment (due to their extremely low levels of enrollment pre-policy change) but this is again excluded from the graph for purposes of scale.

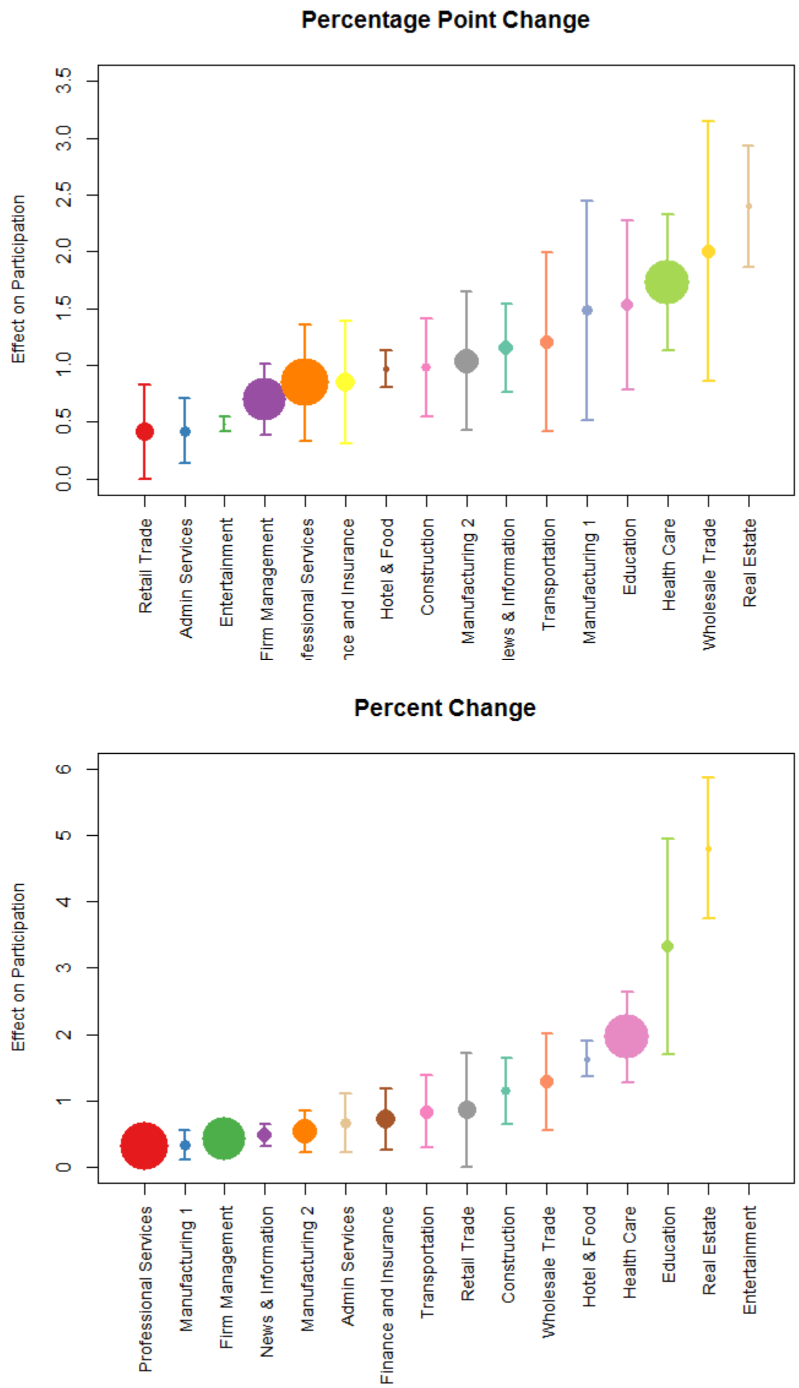

Figure 26: Effect of Automatic Enrollment on Employee Saving by Industry

The highest relative increases are once again in industries with large relative changes in participation like education and health care, and in companies with high auto-enrollment default rates like real estate (which also saw a relatively high increase in participation after the policy change). We once again see the lowest relative increases in the sectors of manufacturing and professional services, which reflects their relatively low increases in participation rates due to high baseline levels of pre-policy-change enrollment.

In general, we find that differences in the effects of automatic enrollment by industry are the best explanation for the discrepancy between our baseline estimates compared to the estimated effects in previous papers. If one evaluates the policy in a small number of firms clustered in particular industries, this is likely to result in only a partial estimate of how automatic enrollment might affect 401(k) participation and saving rates on a larger scale. This is why we get similar results to previous papers like Madrian and Shea (2001) and Choi, et al. (2004), when we restrict our sample to large firms in the health services sector. Similarly, we fine that sectors with relatively high enrollment rates in the control population like news/information and professional services − similar to federal workers in tFalk and Karamcheva (2019) − see relatively lower effects compared to our baseline estimates.

2.6 Conclusion

In conclusion, we find that automatic enrollment increases participation in company-sponsored retirement plans by approximately 32.4 percentage points and increases average savings by 0.9 percentage points. We find that this relatively low increase in average savings is both compositional − in that automatic enrollment has a positive effect on the participation of non-savers, enrolling them at a relatively low rate − and reflective of a negative nudge on savers. We estimate that when firms implement automatic enrollment at a default rate of less than five percent, this decreases the average saving of individuals who previously opted in to a company-sponsored retirement plan (at an average rate of around six percent).

We generally find differences in the effect of automatic enrollment on participation and saving rates when looking at firms that adopt different default enrollment rates. However, although we find that the effect on saving rates in particular is greater for employees that are automatically enrolled at higher rates, we see that the default rates adopted by each firm are endogenous to the characteristics of the employees that work there (particularly the average wage level and saving rate of employees at the firm). Nevertheless, we see that there is considerable bunching at the enrollment rate set by employers, and that the negative nudge is not present in firms that set their rate at five or six percent.

Our findings differ from previous estimates of the effects of automatic enrollment due to the greater diversity of firms and workers that we include in our data. We find that the effects of the policy on 401(k) participation and contribution levels vary considerably across different types of employees and firms. In particular, we find large differences in the effect of automatic enrollment across firms clustered by industry type, with no clear pattern as to why workers in some industries are more likely to participate in a company-sponsored employer plan after implementation of the policy (at least not when it comes to the size of the firm or average wage of employees at each firm).

These differences in the effect across industries best explain why we see a lower effect than that of Madrian and Shea (2001) and Choi, et al. (2004). When we restrict our analysis to companies that more closely resemble the firm(s) in these studies we obtain similar results. However, analyzing the effect among employees working at this select group of firms provides us with only a partial understanding of how workers might respond to the policy change in the economy at large. When we analyze how automatic enrollment affects employees in different sectors, we see wide variation in the effect on 401(k) participation ranging from 20 to 50 percentage points. These differences also help explain why the estimates in these studies vary significantly from the lower estimated effect in Falk and Karamcheva (2019).

Additional work is still needed to evaluate how automatic enrollment would effect other sectors of the economy. Since restrictions surrounding our data only allow us to look at relatively large firms, we cannot evaluate what effect the policy might have on smaller companies with less than 200 employees, for example. Further work is also needed to evaluate if the policy may have secondary effects on other forms of saving, such as contributions to individual retirement accounts and withdrawals from existing saving accounts.

CHAPTER 3: ARE THE POSITIVE EFFECTS OF AUTOMATIC ENROLLMENT OFFSET BY CHANGES IN OTHER FORMS OF SAVING?[3]

Over the past decade it has become increasingly common for private companies to adopt automatic enrollment as a feature of their retirement plans, a practice encouraged by federal and state governments. Several papers examining individual companies that adopted automatic enrollment have provided evidence supporting the effectiveness of this "nudge," finding that changing the decision to save from one of opting in to opting out has large positive effects on employee participation and contributions (Derby & Mortenson, 2020; Falk & Karamcheva, 2019; Beshears, 2009; Nessmith, Utkus, and Young, 2007; Thaler & Benartzi, 2004; Choi et al., 2004; Madrian & Shea, 2001). The positive estimated effects of the policy are particularly stark when compared to the relatively small effects of other interventions like increasing discretionary employer matching rates (Engelhardt & Kumar, 2007; Duflo, et al., 2006; Even & Macpherson, 2005; Munnel, et al., 2001; Kusko, et al. 1998; Papke & Poterba, 1995).

However, little is known about whether automatic enrollment increases net saving for individuals and households. In particular, we do not know if the associated decrease in take-home pay that comes with automatic enrollment in company-sponsored retirement plans is offset by a decrease in the existing savings of employees or their spouses. For example, cash-constrained individuals who are automatically enrolled in employer plans, which are often accompanied by an employer-match, might choose to remain enrolled in the plan and withdraw funds from retirement accounts to cover expenses. Alternatively, individuals might stop contributing − or contribute less − to individual retirement accounts (IRAs). The decrease in take-home pay may also lead to changes in saving at

[3]This chapter is co-authored with Jacob Mortenson from The Joint Committee on Taxation

the household level, with spouses reducing their level of contributions to a retirement account to compensate for the lower net wages of their partners.

Analyzing how automatic enrollment affects saving both within and outside firms is also important for determining whether the policy has distributional effects. This question is particularly relevant as automatic enrollment has gained notable traction among firms attempting to qualify for exemptions from nondiscrimination testing requirements (IRS, 2020). If not exempt, firms are required to test their traditional 401(k) plans each year to ensure that the contributions made by and for "rank-and-file" (non-highly-compensated) employees are proportional to contributions to the plan made by and for highly-compensated employees such as owners and managers (IRS, 2020). If higher-wage employees disproportionately increase their net savings in response to automatic enrollment, then the policy exacerbates inequalities rather than reduces them.

In order to help answer these questions, we construct an original data set of employer-employee-household linked data comprised of 390,733 employees at 279 large, US-based firms that adopted automatic enrollment between 2010 and 2016. We use a difference-in-differences (DiD) approach to estimate changes in the saving behavior of employees who were hired before and after their firms adopted automatic enrollment. In our data, employees hired one year after each firm adopted automatic enrollment (and their spouses) are the "treated" population, and those hired one or two years before the policy change (depending on firm-level service requirements for eligibility) are the "control" population. In order to estimate the effect of automatic enrollment on other forms of saving – aside from employee contributions to company-sponsored plans – we compare changes in contributions to and withdrawals from savings accounts between the two groups for both employees and their spouses.

At the individual (employee) level we find that automatic enrollment does have significant spillover effects on the saving behavior of individuals who are automatically enrolled in company-sponsored retirement plans. In our previous paper we find that automatic enrollment is associated with a 32.4 percentage point increase in 401(k) enrollment and a more modest 0.9 percentage point increase in contributions to employer-sponsored plans, measured as a percent of wages (Derby & Mortenson, 2020). In this paper we find that automatic enrollment is also associated with a 5.1 percentage point increase in the percent of individuals who withdraw funds from existing retirement accounts, and a 36.1 percent increase in the level of withdrawals from those accounts. Conversely, we do not find statistically significant effects on individuals' contributions to individual retirement accounts (IRAs). We further find that withdrawal of funds from retirement accounts reduces the increase in saving resulting from automatic enrollment.

The negative effect on withdrawals is especially large for employees in the bottom quintile of wages within our sample (those with annual wages under $17,700), who are 10.7 percentage points more likely to withdraw funds from existing retirement accounts as a result of being automatically enrolled. The corresponding increase in the average amount of withdrawals is 80.8 percent, equivalent to roughly $2,000. When we look at net savings we see that automatic enrollment has essentially no effect on the savings of individuals in this group, with the estimated increase in withdrawals completely offsetting the increase in contributions to firm-sponsored plans. Interestingly, we also see that employees in the lowest quintile of wages are more likely to make IRA contributions due to automatic enrollment, but the effect is small enough to be considered negligible, especially when compared to the estimated effect on withdrawals for the same group.

In comparison, we estimate that automatic enrollment is associated with a much smaller 5.5 percent increase in withdrawals among employees in the top quintile of the wage distribution, and the estimate is not statistically significant at the five percent level. For all employees with wages higher than the 20^{th} percentile we do see a positive overall effect of the policy on net savings, but this effect is also increasing in the wage level. We estimate that employees at the highest wage levels have disproportionately positive responses to automatic enrollment, compared to those at the bottom of the wage distribution. The average increase in net savings among those in the top 20 percent of earnings is $2,896.

When we look at the saving behavior of spouses of employees automatically enrolled in employer-sponsored plans, we find that automatic enrollment has no significant effect. Although we see a small increase in contributions to company-sponsored retirement plans among spouses in the treated group relative to those in the control group, when we control for other factors like spousal income, gender, age, and firm characteristics, those differences disappear and are no longer significant at the five percent level. We also do not see significant increases in spousal contributions to IRAs in the treated group relative to spouses of employees in the control group; nor do we see statistically significant differences in the withdrawals from retirement accounts between spouses in each group. Thus, we conclude that the spillover effects of automatic enrollment on other forms of saving are restricted to effects on saving at the individual level, rather than at the household level, and that these effects are largely negative.

3.1 Data and Summary Statistics

The retirement plan-level data we use come primarily from Form 5500 filing attachments downloaded from the Department of Labor's (DOL) Employee Benefits Security Administration (EBSA) database. We begin by identifying companies that adopted automatic enrollment at some point between 2010 and 2016 using the Department of Labor's "Bulletin" files (which are cleaned by DOL) and we cross-reference these data with Form 5500 attachments downloaded from the EBSA database. These attachments provide details on each defined contribution (DC) plan, including eligibility requirements, employer matching rates, vesting schedules, and automatic enrollment rates. We use a text reading algorithm to identify which Form 5500 attachments mention automatic enrollment, and catalogue these by firm and tax year. We provide a more detailed explanation of this process in **Appendix F**.

Once we have downloaded and processed the Form 5500 attachments, we link plan data to W-2 filings from 2010 to 2016. Importantly, we incorporate known parent-subsidiary relationships as part of this linkage. This process is also described in greater detail in **Appendix F**. After matching the two data sets, we use another text-reading algorithm to identify the exact year − and often the exact date − when each company begins automatically enrolling employees in their 401(k) plan(s). In order to ensure the accuracy of our treatment variable, we restrict our sample of firms to those that explicitly mention a start date of automatic enrollment in one of their Form 5500 attachments.

Using W-2 filings we are able to identify employees that were hired one year before or one year after each firm adopted automatic enrollment. Individuals hired in the year before each firm adopted automatic enrollment − who must opt in to a retirement plan − are designated as our "control" population. employees hired one year after each firm adopted automatic enrollment − automatically enrolled in a retirement plan − are our "treated" population. In order to ensure sufficient within-firm variation among employees we only include companies that hire 100 new employees both in the year before and the year after adopting automatic enrollment. This restriction reduces our sample considerably to only a few hundred firms. We also restrict our data to employees whose wages are greater than an employee working full time at minimum wage for one quarter of the year. We do this to exclude short-term employees from our data, as well as employees who are not yet be eligible to participate in a retirement plan.

In the Form 5500 plan details, firms often include service requirements for eligibility to participate in the company-sponsored retirement plan. Employees must often wait for a period ranging from one day to one year before they become eligible to enroll. According to the Profit Sharing/401k Council of America (PSCA), in 2010 approximately 57.6 percent of firms had at least a three month waiting period before employees became eligible to enroll in 401(k) plans (Gelber, 2011; PSCA, 2010). We also observe that roughly 60 percent of eligible firms in our data have a three month minimum service requirement. This complicates our measurement of treatment since we do not want to identify new employees as "treated" if they have only worked at the firm a short period of time, and are not yet eligible to enroll (or be automatically enrolled) in the company's retirement plan.

As a consequence, we use the Form 5500 DC plan attachments – which provide details on employee eligibility requirements and waiting periods – to further restrict our sample to firms that have less than a six-month service requirement for participation in the firm's retirement plan. For employees with less than a three-month waiting period, we calculate the percent of wages deferred to a retirement plan the same year they are hired. For individuals with a three- to six-month waiting period, we calculate the percent of wages deferred in the year after they join the firm, to allow enough time for them to become eligible to contribute to the retirement plan. The hiring date for employees in the control group at companies that have a three- to six-month eligibility period is actually two years prior to adoption of automatic enrollment, to ensure we are still calculating their contributions before the company adopts the policy. The hiring date for treated workers is still one year after each firm adopts automatic enrollment, and we measure their contributions the following year.

We also exclude firms from our sample that do not allow all workers (both part- and full-time) to participate in their plan. Finally, if a company has more than one retirement plan, or if the company has subsidiaries with their own retirement plans, we exclude the firm if there are conflicting auto-enrollment start dates across different plan documents. In general, we try to ensure that if a firm shows up in our sample as adopting automatic enrollment in year t, that most, if not all, of the employees at that firm are eligible to enroll in the firm's 401(k) plan within three months from when we observe their contributions. In general, we want to ensure that we do not observe low enrollment rates due to employee exclusions or long service requirements for eligibility.

Our full data set includes 390,733 new employees – hired either one year before or after each firm adopted automatic enrollment – from 279 unique firms. Once we have identified this population we use household identifiers to pull wage and deferred contribution data for both individuals and their spouses from W-2, 1099-R, and 5498 filings. These data include contributions of both individuals and their spouses to individual retirement accounts (IRAs) – an aggregated measure of contributions to Roth IRAs, SIMPLE IRAs, and Simplified Employee Pension (SEP) IRAs – and withdrawals from retirement accounts (excluding any transfers from a retirement account to a Roth IRA).

3.1.1 Summary Statistics: Employee Saving Responses

When we look at the raw data we do see some differences in the saving behavior of employees in the control and treated groups after they are hired at auto-enrollment firms. As we find previously, there are large changes in the distribution of enrollment rates in company-sponsored plans among treated employees relative to those in the control group. As shown in **Figure 27** below, the distributions of enrollment rates prior to individuals being hired at one of the firms in our sample are almost identical when comparing the treated and control populations. However, after employees are hired, we see that treated employees (those automatically enrolled in a firm-sponsored plan upon hire) are more likely to defer a nonzero percent of wages to a 401(k) plan, and more likely to enroll at levels between one and six percent (the default rates) than their counterparts in the control group (hired at the firm before the policy change).

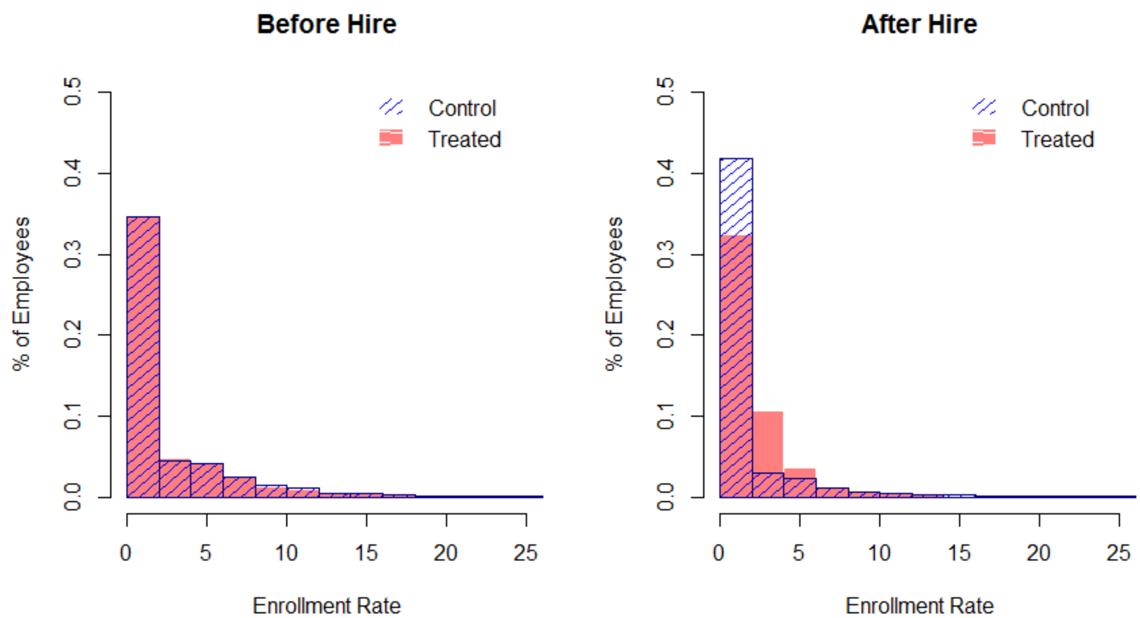

Figure 27: Effect of Automatic Enrollment on Employee Saving at Firm

These patterns in the raw data are consistent with the effects of automatic enrollment on employee contributions to company-sponsored plans that we estimate in our previous paper, which uses the same sample of employees and firms. Using these data, we find that automatic enrollment is associated with a 32.4 percentage point increase in employee participation in company-sponsored retirement plans, and an average increase in savings (measured as a percent of wages) of 0.9 percentage points (Derby & Mortenson, 2020). We also see considerable heaping around the default enrollment rate chosen by each firm, which is typically lower than the average saving rate of employees who opt in to the plan. This results in a negative nudge on "savers" (employees who previously opted in to an firm-sponsored plan), countering the positive effect of auto-enrollment on participation.

These findings illustrate the direct effects of automatic enrollment on employees. However, we also see effects of the policy on other forms of saving, beyond changes to enrollment in company-sponsored plans. As shown in **Figure 28** below, we also see a relative increase in the withdrawals taken by employees in the treated group relative to those in the control group. The first graph in **Figure 28** shows the percent of employees in each group that withdraw funds from an existing savings account before and after they are hired at an auto-enrollment firm. As we can see, there is almost no difference in the two groups before hire, and in both groups there is an increase in the percent of employees that take withdrawals after hire. This may be because employees are withdrawing funds to cover a gap in wages during a job transition.

However, the increase in the percent of workers taking withdrawals is much higher for individuals in the treated group relative to the control group. This suggests that despite the positive effect automatic enrollment has on employee participation in and contributions to firm-sponsored plans, ignoring withdrawals (from the employer-sponsored plans and other forms of retirement savings) might overstate the net effect.

The second graph in **Figure 28** shows the percent of employees in each group − treated and control − who contribute to an IRA before and after they are hired by one of the firms in our sample. As we can see, there is almost no change in the percent that individuals contribute to IRAs before and after hire (for either the control or treated employees). Thus, it does not appear that the increase in withdrawals is counterbalanced by an increase in other forms of saving aside from increases in deferred compensation to firm-sponsored plans.

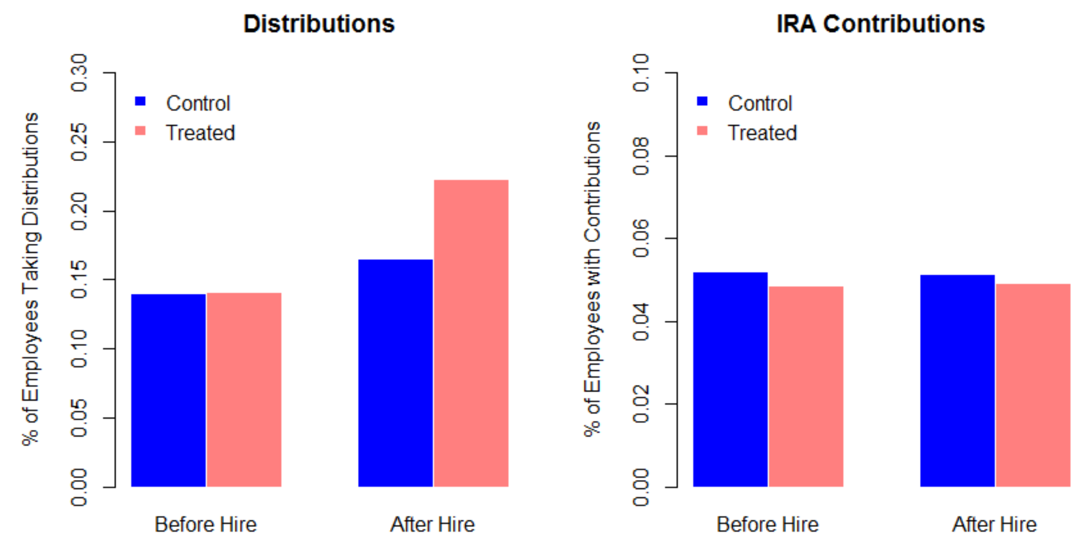

Figure 28: Effect of Automatic Enrollment on Withdrawals and IRA Contributions

When we look more closely at changes in the level of withdrawals we find further evidence of differences between the treated and control groups before and after being hired at an auto-enrollment firm. We see differences between the two groups in the amount that employees withdraw from retirement savings accounts, in addition to the propensity of taking a withdrawal. As shown in the first graph of **Figure 29** below, in the year before they are hired at one of the firms in our sample, employees in both the treated and control groups who withdraw funds from a retirement account withdraw roughly the same amounts. Note that the amounts are calculated as the natural log of withdrawals taken each year. There are slightly more employees in the treated group that withdraw funds from retirement accounts before hire – partially because there are slightly more employees in that group – but the distributions are almost identical.

However, as we can see in the second graph, which plots log withdrawals taken after hire at one of the firms in our sample, although both groups increase their level of withdrawals, the increase in withdrawals among employees in the treated (auto-enrolled) group by far exceeds the increase in withdrawals among those in the control group. Although the median level of withdrawals is now lower for the treated group, there are a greater number of treated employees taking withdrawals above zero relative to those in the control group, and withdrawals among treated employees exceed the number of withdrawals taken by those in the control group at every level. The differences are also quite large, with thousands more employees in the treated group withdrawing between $400 and $3,000 from retirement accounts (roughly equal to log withdrawals between six and eight).

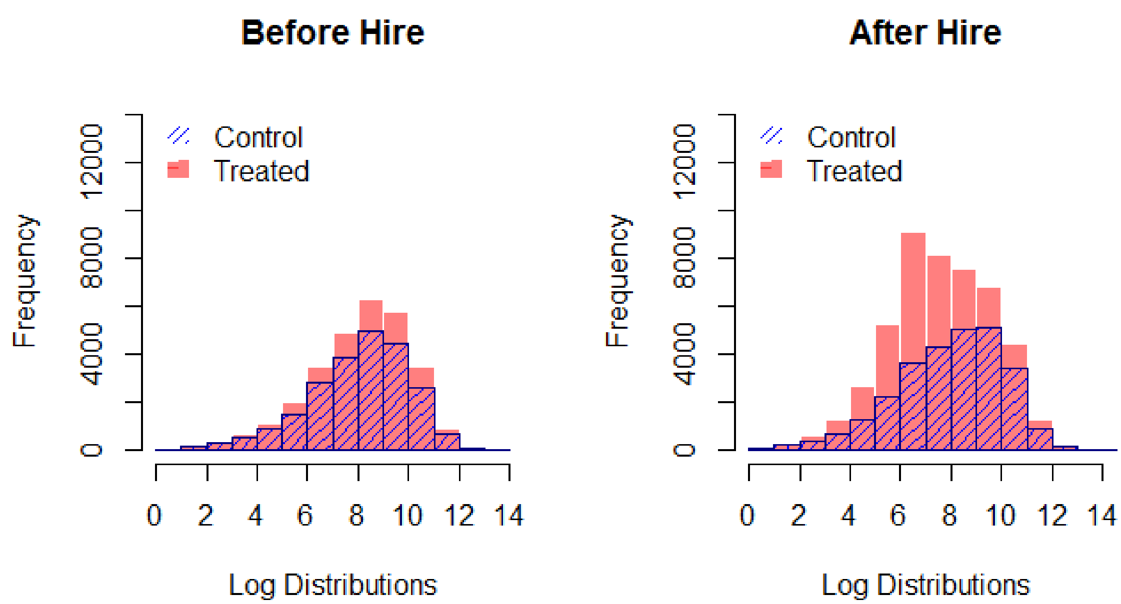

Figure 29: Effect of Automatic Enrollment on Withdrawal Levels

Conversely, when we graph IRA contributions made by employees in the treated and control groups – provided in **Figure 30** below – we do not see a meaningful difference in the change in IRA contributions between the two groups before and after hire. The histograms for each group in the first and second graph are almost identical, with a slight increase in IRA contributions for the treated group. Here again, we see that employees in the treated group are slightly over-represented, so their IRA contributions are higher across the board.

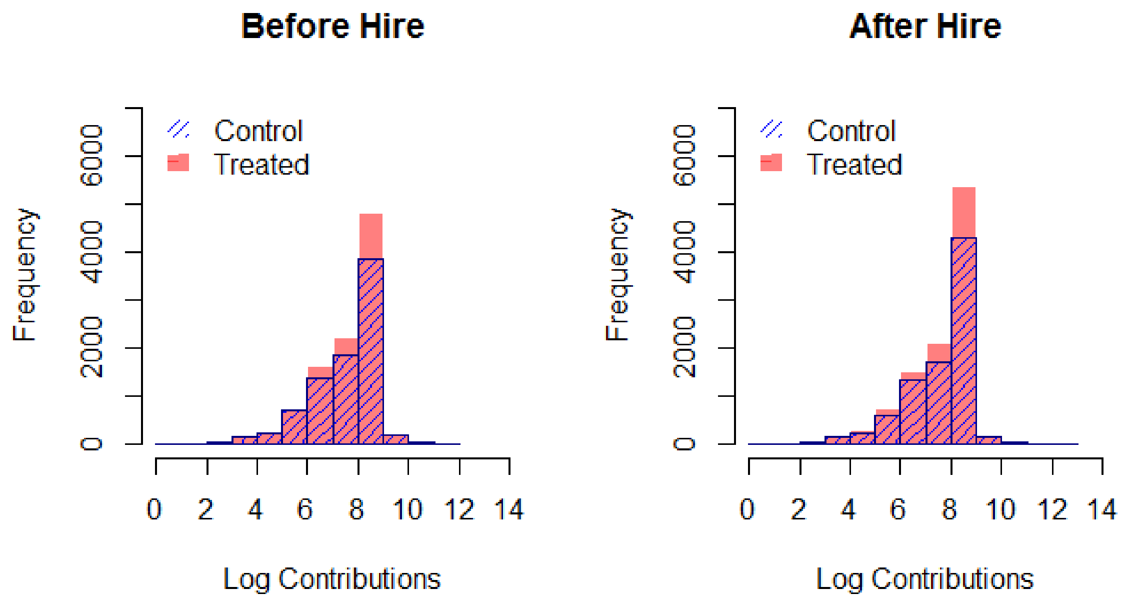

Figure 30: Effect of Automatic Enrollment on IRA Contribution Levels

This suggests that the secondary effects of automatic enrollment on saving at the individual level are mainly negative. When we look at patterns in the raw data it appears that employees who are automatically enrolled in a company-sponsored plan are more likely to withdraw savings from retirement accounts and are not more likely to increase contributions to IRAs. However, we are not controlling for potential correlates in this analysis, such as changes in other income. In Section 3.2 we investigate this pattern further to determine if the differences are statistically significant when controlling for other factors.

3.1.2 Summary Statistics: Spousal Saving Responses

We further examine saving responses to automatic enrollment at the household level by looking at the saving behavior of employees' spouses. We might expect to see either positive or negative spousal responses to automatic enrollment. For example, if an employee is automatically enrolled in a firm-sponsored plan, and thus has a reduction in take home pay that exceeds spending, the employee's spouse might respond to this decrease in household net wages by reducing their own contributions to firm-sponsored savings plans or withdrawing funds from existing savings accounts to cover expenses. Conversely, if an employee who previously did not save is automatically enrolled in a firm-sponsored plan and does not face budget constraints as a result, this may have a positive learning effect at the household level. The employee may also encourage their spouse to begin deferring wages – or defer a larger percent of their wages – to their own firm-sponsored plan (or contribute to an IRA).

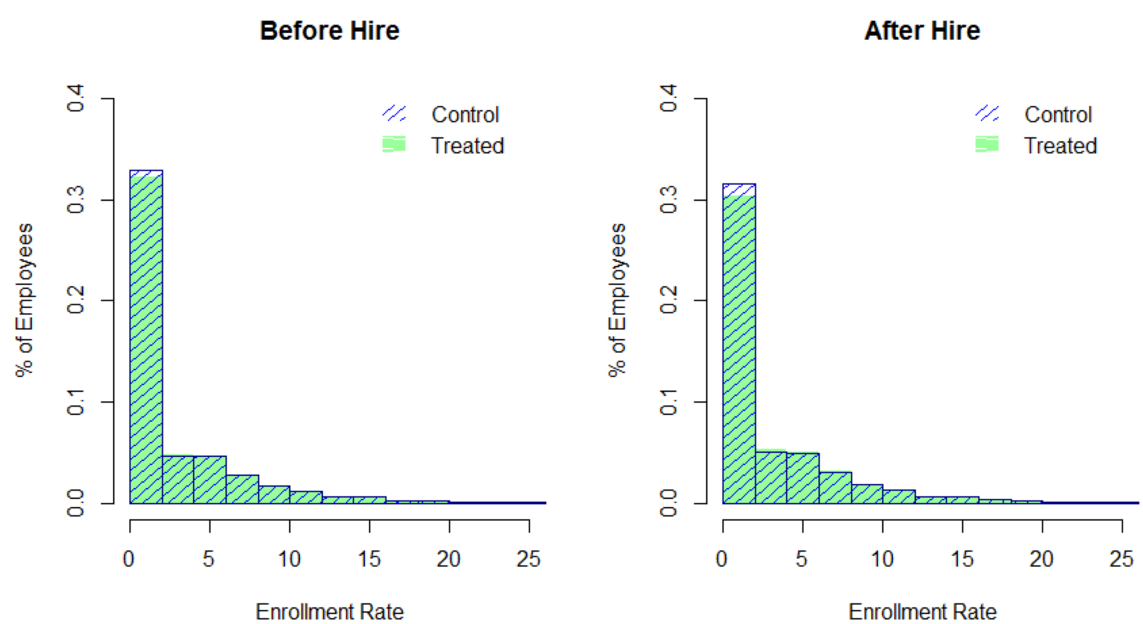

Figure 31: Effect of Automatic Enrollment on Spousal Saving at Firm

As shown in the **Figure 31** above, when we look at the effect of automatic enrollment on spousal contributions to company-sponsored plans (at the company where the spouse works, which are not included in our sample of auto-enrollment firms), we see small differences in saving between the treated and control groups. The first graph in **Figure 31** shows the distribution in enrollment of each group of employees' spouses before their partners are hired at auto-enrollment firms, and the second graph shows spousal enrollment after the employees are hired at auto-enrollment firms. Presumably the companies where the spouses work did not adopt automatic enrollment during the same period, although we do not know whether this is the case for all firms.

As shown in **Figure 31**, when we compare the distribution of spousal saving rates between the treated and control groups in the years before and after hire, we do see a reduction in the number of spouses with zero contributions in the treated group compared to the control group, although the difference is very small. While there do not appear to be large differences in the distribution of contributions themselves, the decrease in spouses enrolled at a rate of zero suggests that there may be a small effect of auto-enrollment on spousal participation in 401(k) plans at their respective firms.

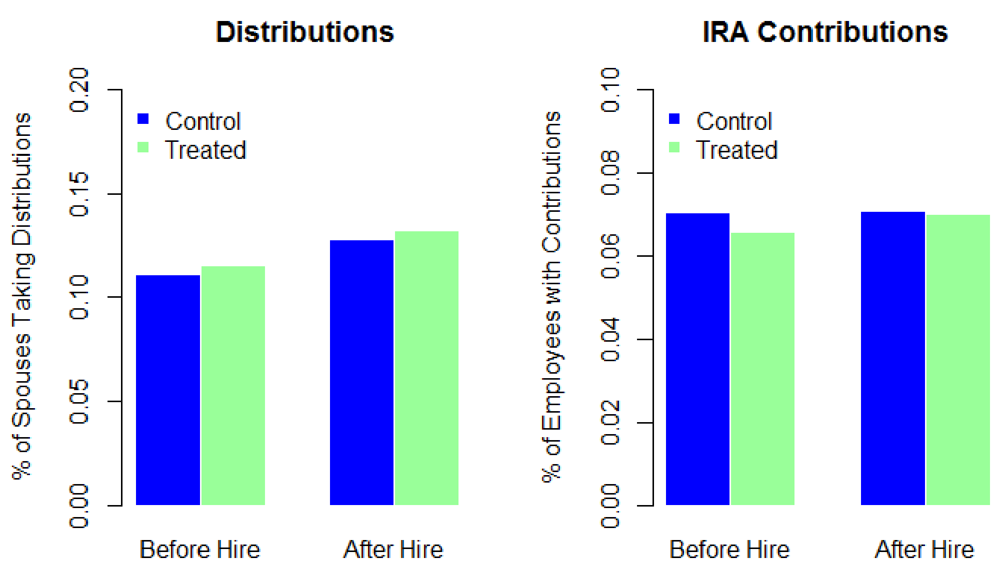

Figure 32: Effect of Automatic Enrollment on Spousal Withdrawals and IRA Contributions

We also see negligible saving responses at the household level when it comes to spousal IRA contributions and withdrawals from retirement savings accounts. When we look at spousal contributions to IRA accounts we see a small relative increase in the treated group relative to the control group before and after hire. The second graph in **Figure 32** graphs the percent of employees' spouses contributing to an IRA in the treated and control groups before and after employees are hired at an auto-enrollment firm.

As we can see in the graph, relatively fewer spouses in the treated group contribute to an IRA in the year before hire compared to spouses in the control group, but this difference disappears after employees are hired at an auto-enrollment firm. The percent of spouses who withdraw funds from an existing retirement account, shown in the first graph **Figure 32** below, tells a similar story. There is a small difference between the baseline rates of withdrawals before hiring, and this difference remains after hiring, suggesting no net change resulting from automatic enrollment.

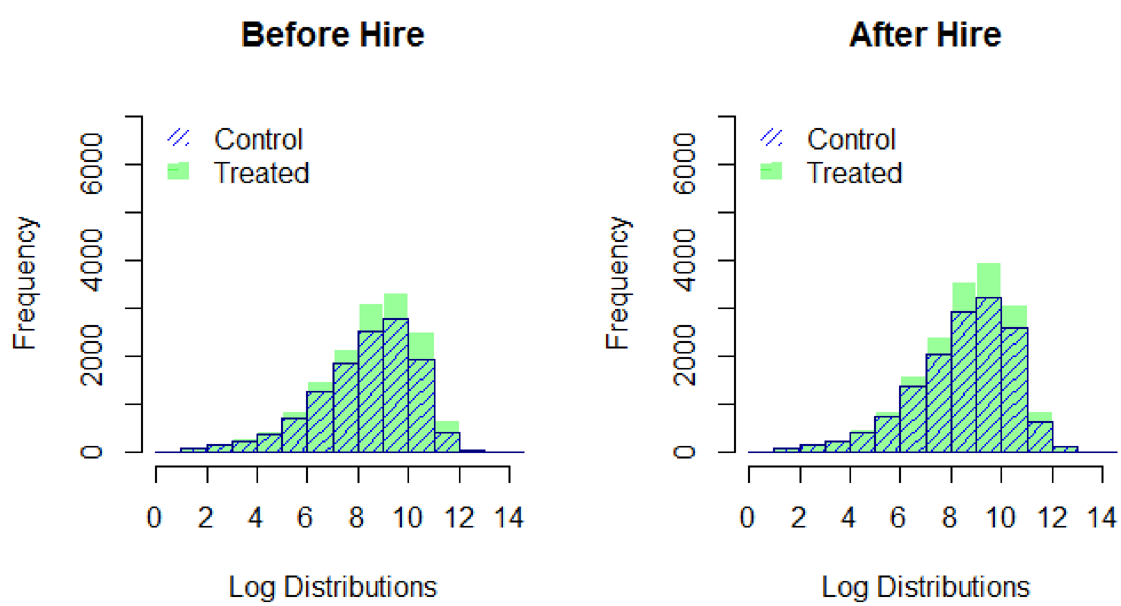

Figure 33: Effect of Automatic Enrollment on Spousal Withdrawal Levels

The pattern is similar when examining the levels of savings by spouses. The distribution of spousal withdrawals from retirement accounts – shown in **Figure 33** above – reveal only minor changes between the two groups before and after hire. It does appear that there is a slight increase in spousal withdrawals from existing accounts in both groups, and that there may be a slightly larger increase for spouses of employees in the treated group, but these differences are quite small relative to what we see when graphing employee withdrawals.

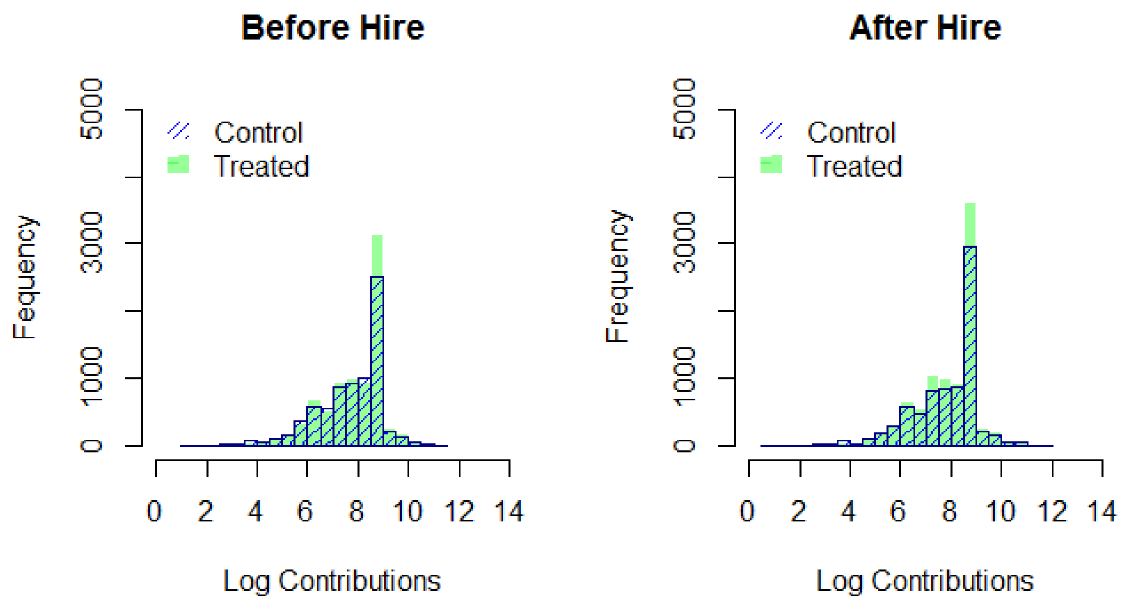

Figure 34: Effect of Automatic Enrollment on Spousal IRA Contribution Levels

When we look at the level of IRA contributions of employees' spouses before and after employees are hired at auto-enrollment firms, we see that automatic enrollment appears to have a small but positive effect on spousal contributions. As shown in the first graph of **Figure 34** above, spouses in the treatment group contribute slightly more to begin with (although this can be explained by the slight over-representation of employees in the treated group), but the distributions are largely overlapping. However, in the year after hire there is a small increase in the IRA contributions of spouses in the treated group relative to those in the control group, particularly around the median.

Thus, although automatic enrollment appears to be associated with withdrawals from savings accounts, there is also an associated increase in spousal saving both in company-sponsored retirement plans and individual retirement accounts. Given the relative increase that we see in withdrawals taken by the employees themselves, these differences raise questions about the net secondary effects of the policy. In the following sections we will examine each of these effects separately, controlling for other correlates like wage level, gender, age, and firm-level characteristics in order to determine the overall effect of the policy. We will then estimate the combined secondary saving effect of automatic enrollment, or the combination of other forms of saving net of withdrawals from existing accounts.

3.2 Empirical Specification

We use a difference-in-differences approach to estimate the effect of automatic enrollment on other forms of saving, utilizing differences in saving behavior over time and between individuals who are hired before and after each firm in our sample adopts automatic enrollment. The regression specification is defined as follows:

$$Y_{i,f} = \alpha + \gamma H_{i,f} + \psi T_{i,f} + \eta\left(H_{i,f} \times T_{i,f}\right) + \beta \mathbf{X}_{i,f} + \delta_f + \varepsilon_{i,f} \tag{1}$$

In equation (1), $H_{i,f}$ is a binary variable indicating whether the dependent variable is calculated before or after the individual is hired by firm f (the firm that adopts automatic enrollment either one year before or one year after the employee is hired), and $T_{i,f}$ is a binary treatment variable that indicates whether the employee began working at the firm before or after the firm adopted automatic enrollment. The coefficient of interest is η, the estimated effect of the interaction term $H_{i,f} \times T_{i,f}$. In other words, η represents the estimated effect of automatic enrollment on the withdrawals from or contributions to existing savings accounts after employees are hired by company f.

For regressions on saving at the individual level, the control variables $\mathbf{X}_{i,f}$ include employer discretionary matching rates (at the firm-level) and the individual's gender, wage level, and age. We also include a variable to indicate whether the firm had an eligibility requirement of more than three (but only up to six) months. Finally, δ_f is a firm fixed effect, which controls for variation coming from time-invariant firm-level characteristics such as industry type. We do not include a year fixed effect because the firm fixed effect, the eligibility indicator, and the treatment variable perfectly predict the year of employment. $Y_{i,f}$ is either a binary variable indicating whether the individual made a retirement savings contribution or withdrawal, or the natural log of the level of contributions to/withdrawals from retirement accounts.

We include two measures of employer discretionary matching policies in our regressions: "maximum match level" and 'matching rate." The matching level is the maximum percent that employees will receive discretionary employer matching contributions up to, and the matching rate is the rate at which firms match employee savings at the maximum level. For example, if a firm matches 50 percent of employee contributions up to six percent, then employees at that firm would have a maximum match level of 0.06, and a matching rate of 0.5. Similarly, if a firm matches 100 percent of contributions up to three percent and 50 percent of contributions between three and five percent, then the maximum match level for employees at that firm is 0.05 and the matching rate is 0.8.

We include both matching variables because some firms change their matching policies around the same time that they adopt automatic enrollment. There are also some policies that set a fixed dollar maximum matching level, which equates to a different percent of wages depending on the salary of each employee. However, it is important to keep in mind that most of the firms in our sample do not change their discretionary matching policies when they adopt automatic enrollment, and do not have fixed dollar limits on contributions. Thus, all of the variation in our estimates on matching comes from a few firms that meet one of these two criteria.

Withdrawals and IRA contributions are both calculated in the year after employees are automatically enrolled, or become eligible to enroll, in a firm-sponsored retirement plan. Since we cannot observe what month each employee begins working at an auto-enrollment firm, we want to ensure that we are capturing the effect on withdrawals and IRA contributions *after* employees are automatically enrolled and/or eligible to enroll in a 401(k) plan. If an employee is hired halfway through the year, for example, the effect on withdrawals may be diluted by differences in their saving behavior in the months prior to that date.

We use the same approach outlined above to measure the effects of automatic enrollment on spousal contributions to IRAs and withdrawals from existing accounts. The only difference in the estimation procedure for these two spousal variables is that we use control variables that pertain to the spouse rather than the employee. The spousal controls are the wage level, age, gender, and tax year and North American Industry Classification System (NAICS) code fixed effects. We use these instead of a firm fixed effect because the spouses in our sample are spread across a much larger number of firms − very few work at the same firm − and we include the tax year fixed effect to control for time-dependent variables, which is not necessary in our individual-level specification due to the high correlation between the firm fixed effect and tax year (as explained above).

The nature of spousal wages requires a different estimation approach than with the other dependent variables. Some spouses in our data are not working, which may introduce selection bias in estimates of the saving responses of spouses. We do not believe that this is the case when it comes to IRA contributions and withdrawals from existing accounts because the correlation between the decision to contribute/withdraw and the decision to work at a firm with a company-sponsored plan is not as strong. In an attempt to address any possible sample selection, we use a Heckman two-step estimator to correct for the bias of not observing the saving level of spouses who do not earn a wage. The specification is defined as follows:

$$Y_{2i,t,n} = \alpha + \gamma H_{i,t,n} + \psi T_{i,t,n} + \eta(H_{i,t,n} \times T_{i,t,n}) + \beta \mathbf{X}_{i,t,n} + \delta_t + \zeta_n + \sigma_{12}\lambda(\hat{\theta}\mathbf{Z}_{i,t,n}) + \nu_{i,t,n} \qquad (2)$$

where $\hat{\theta}$ is obtained using a first-step probit regression of Y_1 on $\mathbf{Z}$, where Y_1 is the incidence of working at a firm that issues a W-2 statement, and $\mathbf{Z}$ includes the same treatment, hiring, and interaction variables as in equation (2) and the same control variables used in the regressions on spousal IRA contributions and withdrawals, including NAICS code and tax year fixed effects (δ_t and ζ_n, respectively). $\mathbf{Z}$ also includes the exclusion restriction of employee wages, which we posit is correlated with a spouse's decision to work, but not directly correlated with the amount that the spouse chooses to contribute to a firm-sponsored plan if they choose do work. The reasoning behind this assumption is that if one person in the couple earns enough money to cover all joint expenses, then there is less of a need for the spouse to work (and vice versa if each spouse does not earn enough to cover expenses alone). However, the wage level of one's spouse would be less directly related to the decision of how much to contribute to a 401(k) plan, and in the data we see that this correlation is in fact quite low at 0.06.

In equation (2), $\lambda(\hat{\theta}\mathbf{Z}) = \phi(\hat{\theta}\mathbf{Z})/\Phi(\hat{\theta}\mathbf{Z})$ is the estimated inverse Mills Ratio from the first-stage regression and σ_{12} is the correlation between the two error terms in the first and second stage regressions. Including this term in the second stage equation addresses the sample selection issue assuming the error term in the stage equation is a linear function of the error in the first stage equation. Similarly, if we assume the shape of the Mills Ratio is approximately linear, the system of equation is identified without the exclusion restriction (Heckman, 1979).

In the second stage regression, Y_2 is the percent of wages that the spouse of employee i contributes to the company-sponsored retirement plan where they work. We do not use the Heckman correction to estimate the incidence of spouses deferring a positive amount of wages to a 401(k) because the functional form is not appropriate to use with a binary dependent variable. However, for the sake of comparison, we run the same difference-in-differences regression used for our other dependent variables (ordinary least squares for the level of saving and logistic for the incidence of saving), using data from only the spouses who we observe working, to see how automatic enrollment affects this sub-sample of individuals. These estimates are included solely for comparison purposes as some selection bias is introduced with the exclusion of spouses who do not earn a wage.

Spousal contributions to their own employer-sponsored plans are also measured in the year after employees begin working at the one of the firms in our sample. As explained above, the reason for this is that we cannot observe what month each employee begins working at an auto-enrollment firm, and we want to make sure that we are calculating spouses' saving responses *after* employees are automatically enrolled and/or eligible to enroll in a 401(k) plan.

3.2.1 Balance Tests

We check to see if our treated and control populations differ in meaningful ways by comparing individual- and firm-level statistics for the two groups, including wages, firm size, gender, and age. In order to do this comparison we run a series of t-tests to evaluate whether the differences in each characteristic are statistically significant at the five percent level. The results of this analysis are presented in **Table 13** below. As shown in the table, the differences in mean age and percent female between the two groups are not statistically significant. However, both employees and spouses in the treated group do have slightly higher wages on average, even after adjusting for inflation. This suggests that firms may adopt automatic enrollment during a period of firm-level or macroeconomic growth.

When looking at differences between the treated and control populations at the firm level (also in **Table 13**), we generally do not see statistically significant differences in firm size or wages before and after automatic enrollment (at either the mean or the median). The differences in wage growth that we see at the individual/spousal level are thus likely explained by the expansion of a few large firms in our data. Essentially, the individual-level tests are the outcome of a weighted sample of firms, and not necessarily reflective of a general trend among companies adopting automatic enrollment.

Still, the growth in wages at some firms in our data (and at the firms of some spouses) may affect our estimates of the effects of automatic enrollment if they are not taken into account. Simply looking at a comparison of averages may overstate the effect of the policy if looking at one firm or a handful of large companies. We attempt to control for this by regressing on potential correlates like employee wages and employer matching policies. We also include a firm fixed effect in our employee-level regressions to control for time-invariant firm-specific differences, such as industry type. As explained above, for household-level (spousal) regressions we control for both industry type and tax year, which we find are the firm-level characteristics most highly correlated with saving levels (Derby & Mortenson, 2020).

Table 13: Balance Test: Control and Treated Population Differences

Individual-Level	**Mean Control**	**Mean Treated**	**95% CI (Treated-Control)**	
Wages	$33,889	$35,342	$868	$2,038
Age	38	38	-0.1	0.0
Percent Female	50.7%	50.6%	-0.004%	0.002%
Auto-Enrolment Rate	3.2%	3.2%	-0.00%	0.01%
Spouse/HH-Level	**Mean Control**	**Mean Treated**	**95% CI (Treated-Control)**	
Wages	$52,162.06	$55,241.83	$2,344.03	$3,815.51
Age	42.06	41.03	0.924	1.126
Percent Female	0.512	0.510	-0.002	0.006
Firm-Level	**Mean Control**	**Mean Treated**	**95% CI (Treated-Control)**	
Firm Size	2,635	3,117	-1,552	589
Mean Wage	$49,978	$50,127	-$5,271	$5,570
Median Wage	$37,520	$38,761	-$3,051	$5,534

3.3 Results

3.3.1 Employee Withdrawals

We first examine the secondary effects of automatic enrollment on employee withdrawals. As described in the previous section, we run two different types of regressions for each dependent variable: ordinary least squares (OLS) for log withdrawals and a logistic maximum likelihood regression for the incidence of taking withdrawals. Standard errors in both regressions are robust and clustered at the firm level. In general, we find that the effects of the policy are mainly negative: with increased withdrawals in the treated group relative to the control group, and no difference in IRA contributions.

As detailed in the first column of **Table 14**, the coefficient on the interaction term between being hired and treated is 0.051, indicating that automatic enrollment is associated with a 5.1 percentage point increase in the probability of positive withdrawals. The baseline rate is 16.5 percent (third row from the bottom of the table), which means the probability of taking a withdrawal increases by roughly 31 percent. The coefficients associated with the regression of log-level withdrawals are presented in column 2. The coefficient on the interaction term is 0.361, indicating that auto-enrollment increases the level of withdrawals by 36.1 percent. This is equivalent to an increase of $737 relative to individuals in the control group, who on average withdraw $2,041 from retirement savings accounts after they are hired at one of the firms in our sample. Both estimates are statistically significant at the 0.1 percent level. These results are consistent with the raw data in **Figure 28** in the previous section, showing an increase in total withdrawals taken by employees in the treated group relative to those in the control group after being hired at an auto-enrollment firm.

We also estimate significant effects of gender, age, and wages on both the incidence of taking withdrawals and the average amount funds withdrawn. We see that women are on average less likely to take withdrawals, employees are more likely to withdraw funds as they get older (and as their total savings increase), and those with higher wages are less likely to take withdrawals. Notably, we also see a 3.2 percentage point increase in the probability of taking withdrawals and a 23.7 percent increase in the amount of funds withdrawn from existing accounts as a result of being hired at an auto-enrollment firm (across both the treated and control populations).

Table 14: Automatic Enrollment Effect on Withdrawals from Retirement Accounts

	Incidence (Logit ME)	**Log Withdrawals (OLS)**
Hired	0.032***	0.237*
	(0.009)	(0.100)
Treated	-0.003	-0.013
	(0.003)	(0.025)
Hired×Treated	0.051***	0.361***
	(0.011)	(0.063)
Female	-0.011*	-0.160***
	(0.005)	(0.044)
Age	0.006***	0.064***
	(0.000)	(0.004)
Age Squared	0.000	0.000
	(0.000)	(0.000)
Log Wages	-0.013*	-0.093*
	(0.006)	(0.037)
Max Match Level	0.133	0.770
	(0.085)	(0.559)
Matching Rate	0.030	0.246
	(0.052)	(0.364)
3-6 Month Wait	-0.664	-8.764
	(0.468)	(5.478)
Firm Fixed Effect	✓	✓
Baseline	0.165	$2,040.97
	(0.001)	($28.83)
Observations	365,071	365,071
Number of Firms	279	279

Note: *p<0.05, **p<0.01, ***p<0.001

We are concerned that there may an issue of censoring in our data, with some employees taking zero withdrawals simply because they do not have sufficient savings to withdraw funds from. In order to test whether this is the case, we restrict our sample only to employees who we observe contributing to a firm-sponsored plan – or to an individual retirement account – in the year before or two years before we observe them in our data (effectively, two to three years before they are hired at auto-enrollment firm). We then run the same regressions as above to see if there is any difference in the estimated effect.

When we restrict our population to savers (individuals who contributed to a retirement plan at their previous job), we see that the effect of automatic enrollment on withdrawals is similar to the result presented in **Table 14** above. Individuals who are automatically enrolled in a 401(k) plan are 5.5 percentage points more likely to take withdrawals than their counterparts in the control group, which is only a 0.4 percentage point difference from our estimated effect using the full population. Similarly, automatically enrolled individuals withdraw 33.5 percent more in savings from existing accounts than those in the control group, a 2.6 percent difference from the full population estimate. The results of this analysis are presented in **Table 43** in **Appendix L**.

Next, we explore whether the estimated effect of automatic enrollment on withdrawals varies across the wage distribution. To answer this question, we divide our population of employees into five groups based on their earnings (calculated as aggregate wages in the year after becoming eligible to enroll in a retirement plan), and run the same regressions on each sub-group stratified by wage. As shown in **Figure 35** below, we see that there are statistically significant differences in the effect based on wage level. Treated employees in the lowest wage quintile are 10.7 percentage points more likely to take withdrawals compared to their counterparts in the control group, and their withdrawals are on average 80.8 percent greater than those in the control group, a much larger effect than we estimate for the full population. In comparison, employees in the top quintile of wages essentially see no effect on withdrawals due to automatic enrollment, with estimates in both regressions close to zero and not significant at the five percent level.

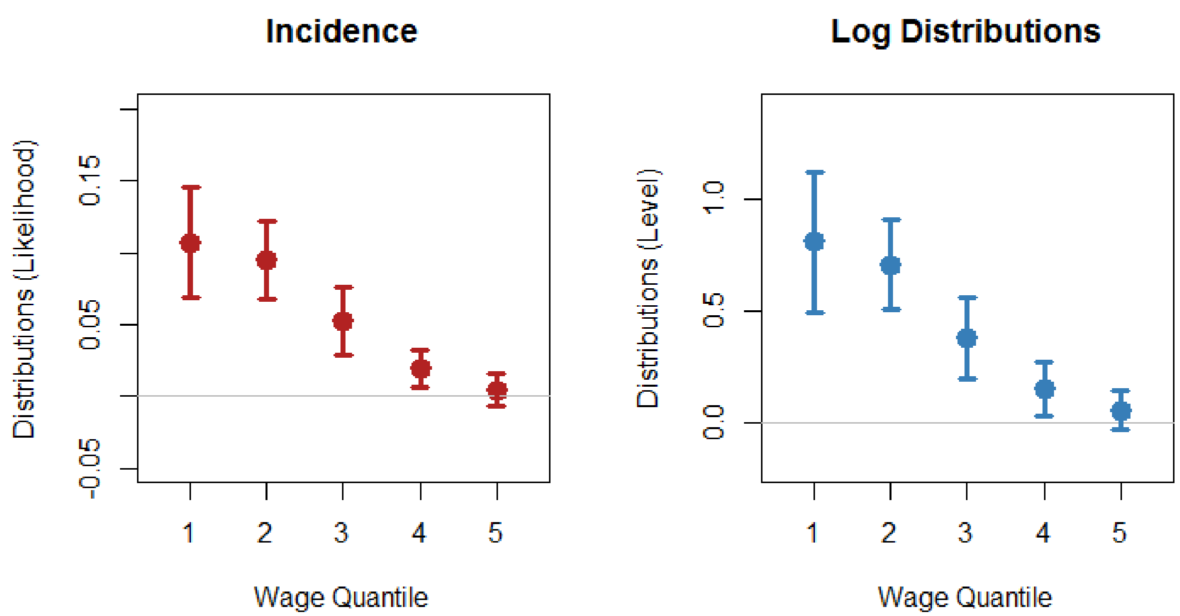

Figure 35: Effect of Automatic Enrollment on Employee Withdrawals by Wage Level

These results are consistent with what we might expect based on the way that automatic enrollment affects people who earn different wages. Low wage workers automatically enrolled in defined contribution plans are less likely to be able to adjust to a decrease in take-home pay, and are therefore more likely to need additional funds to make payments if their wages are being reduced by automatic enrollment in a company-sponsored plan. By remaining enrolled in a 401(k) plan, these individuals may also be taking advantage of employer-matching policies. In particular, if the employer match rate exceeds the penalty for early withdrawals (10 percent for those under age 59.5), the employee can increase their compensation by $(\gamma - 10\%) * contribution$, where γ is the employer match rate, by making contributions to their retirement account and withdrawing the full value of that contribution after the employer match is made. Withdrawing funds from existing accounts may also be less time-consuming, or require less paperwork, than opting out of a firm-sponsored plan.

On the other hand, those at the top end of the wage distribution are more interested in the tax-deferral benefits of saving in tax-preferred accounts, and are less likely to need the full value of their compensation to fund contemporaneous consumption. As we see in **Figure 35** above, the effect of automatic enrollment on withdrawals is monotonically decreasing in the wage distribution, with each group less affected by the policy. The full results from these regressions by wage level are provided in **Tables 44 and 45** in **Appendix M**.

3.3.2 Employee IRA Contributions

Consistent with the raw comparison of means above, we do not see much of an effect of automatic enrollment on employee contributions to IRAs. As shown in **Table 15** below, we estimate that automatic enrollment is associated with a 0.1 percentage point increase in the probability of contributing to an IRA, and a 0.6 percent increase in the level of contributions to IRAs. Both estimates are small and neither is statistically significant at the five percent level. Furthermore, unlike in our analysis of the effects on withdrawals, we also do not see any significant effect of being hired at one of the firms in our sample on the incidence or level of IRA contributions. In fact, we find that only age and wage level have a positive effect on contributions. Note that in our analysis of the effect of automatic enrollment on employee IRA contributions we do not think that there is any issue of censoring, as we do in our analysis of withdrawals, seeing as employees only need some positive level of earnings to be able to contribute to an IRA (which all employees in our sample do, by construction).

Table 15: Effect of Automatic Enrollment on Contributions to an IRA Accounts

	Incidence (Logit ME)	**Log Contributions (OLS)**
Hired	0.000	0.007
	(0.001)	(0.007)
Treated	-0.002	-0.011
	(0.001)	(0.008)
Hired×Treated	0.001	0.006
	(0.001)	(0.008)
Female	-0.000	-0.011
	(0.003)	(0.022)
Age	0.001***	0.006***
	(0.000)	(0.001)
Age Squared	-0.000	-0.000
	(0.000)	(0.000)
Log Wages	0.013***	0.123***
	(0.001)	(0.013)
Max Match Level	0.033	0.181
	(0.018)	(0.174)
Matching Rate	-0.000	-0.104
	(0.016)	(0.136)
3-6 Month Wait	0.150	0.181
	(0.184)	(1.297)
Firm Fixed Effect	✓	✓
Mean Control	0.051	$181.86
	(0.001)	($3.58)
Observations	365,071	365,071
Number of Firms	279	279

Note: *p<0.05, **p<0.01, ***p<0.001

We are also interested to find out if there are differences in the effect of automatic enrollment on employee contributions based on wages. In order to answer this question we once again group employees based on wage level (relative to each other) and run the same set of regressions on each group. As shown in **Figure 36** below, we actually do see a statistically significant effect for individuals in the bottom quintile of wages, but not for any other group of employees above the 20^{th} percentile of earnings. For this first group we find that automatic enrollment is associated with an 0.5 percentage point increase in the probability of contributing to an IRA and a 3.2 percent increase in the level of contributions to IRAs. The effects are small but significant at the five percent level. Full results from these regressions are provided in **Tables 46 and 47** in **Appendix M**.

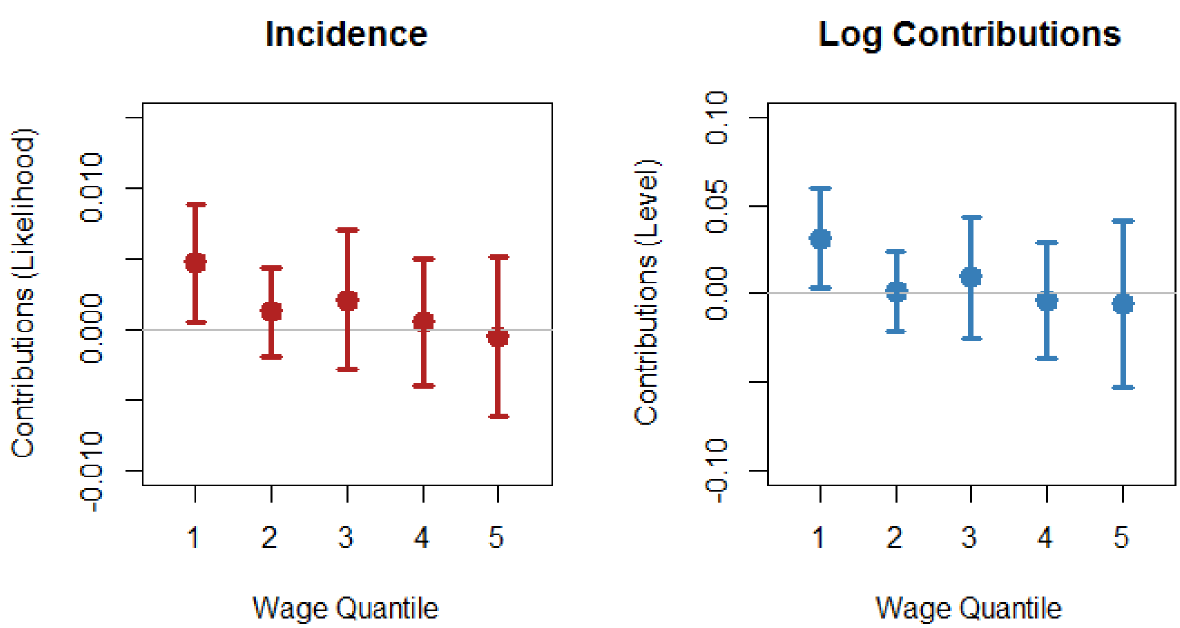

Figure 36: Effect of Automatic Enrollment on Employee IRA Contributions by Wage Level

In sum, although we find that the secondary effects of automatic enrollment on saving at the individual level are largely negative – with larger increases in withdrawals taken by treated employees, particularly for those in the bottom 20 percent of the wage distribution – we also find that there is a small positive effect on saving for those in the lowest quintile of earnings. However, it is important to keep in mind that due to relatively low levels of saving among employees in the bottom quintile of wages, the estimated increase of 3.2 percent in IRA contributions is equivalent to an average increase of just $2.45, which is a negligible amount, especially compared to the increase in withdrawals due to automatic enrollment for this group, which we estimate is equivalent to an average increase of $2,007 (from a baseline withdrawal amount of $2,484). Therefore, although it is statistically significant among employees with lower wages, we consider the effect of automatic enrollment on IRA contributions to be basically zero.

3.3.3 Spousal Contributions to Firm-Sponsored Retirement Plans

We next estimate the effects of automatic enrollment at the household level, beginning with the effect on spousal contributions to firm-sponsored retirement plans. When we run the same logistic and OLS regressions on the incidence of spouses contributing to a company-sponsored plan, and the savings rate (percent of wages deferred), we do not see any effect of automatic enrollment on spousal saving. As shown in **Table 16** below, the effect of automatically enrolling employees on both the incidence of spousal saving and the spousal savings rate is negative but not statistically significant at the five percent level. Standard errors presented in the table are robust and clustered at the NAICS code level.

As we discuss in the previous section, we do think that there may be some bias present in these estimates due to the exclusion of spouses who do not receive a W-2, so we run another regression on the level of spousal contributions using the Heckman two-step procedure from equation (2), including employee earnings as an exclusion restriction. The third column in **Table 16** presents the results, including bootstrapped standard errors. As shown in the table, even after correcting for sample selection due to incidence of spouses working, we still do not see any effect of automatic

Table 16: Effect of Automatic Enrollment on Spousal Contributions to 401(k) Plans

	Incidence (Logit ME)	**Savings Rate (OLS)**	**Savings Rate (Heckman)**
Hired	0.026***	0.293***	0.293***
	(0.003)	(0.031)	(0.032)
Treated	0.021***	0.246***	0.246***
	(0.003)	(0.027)	(0.032)
Hired×Treated	-0.005	-0.070	-0.069
	(0.004)	(0.044)	(0.044)
Female	0.042***	0.178*	0.179***
	(0.009)	(0.072)	(0.017)
Age	0.004***	0.019*	0.019***
	(0.001)	(0.009)	(0.005)
Age Squared	-0.000**	0.000**	0.000***
	(0.000)	(0.000)	(0.000)
Log Wages	0.221***	0.992***	0.992***
	(0.003)	(0.062)	(0.007)
Tax Year FE	✓	✓	✓
NAICS Code FE	✓	✓	✓
Mean Control	0.450	2.788	2.788
	(0.498)	(5.200)	(5.200)
Observations	179,895	179,895	201,624

Note: *p<0.05, **p<0.01, ***p<0.001

enrollment on spousal contributions to firm-sponsored retirement plans. The estimated effect is once again negative but not significant at the five percent level. In fact, the results are quite close to those obtained using the more simple OLS regression (column 2), suggesting that selection bias may not actually be much of an issue in estimating the true effect of the policy on spousal savings.

3.3.4 Spousal Withdrawals and IRA Contributions

We further estimate the effects of automatic enrollment on spousal withdrawals from existing savings accounts and contributions to individual retirement accounts (IRAs). Here again, we see generally null effects of automatic enrollment on the saving behavior of spouses. We estimate that spouses of individuals who are automatically enrolled are more likely to take withdrawals from retirement accounts after hire, and are less likely to contribute to IRA accounts, but all of these estimates are small (or effectively zero) and not statistically significant at the five percent level. For each set of regressions we apply the same difference-in-differences approach using a logistic regression on the incidence of taking withdrawals or making IRA contributions, and an OLS regression on log withdrawals or log contributions. All standard errors are robust and clustered at the NAICS code level (the industry code of the spouse's employer).

When looking at differences in the level of spousal contributions to IRAs between the two groups, we get similar results. As shown in **Table 18** below, the estimated effects are negative, but they are small (rounding up to zero), and not statistically significant at the five percent level. The effects of each control variable on spousal contributions are also similar to what we see in the employee-level regressions. As shown in **Table 17**, we see similar patterns in the estimated effects of gender, age, and wages on the level and probability of taking withdrawals for spouses as we do for employees. Female spouses are about three percent less likely to take withdrawals, and spouses are more likely to take withdrawals as they get older (with a 0.5 percentage point increase in the likelihood of taking withdrawals with every additional year). We also see in these results − as we do in the analysis of employee withdrawals − that spouses in both groups are more likely to withdraw funds from retirement accounts when their partner changes jobs.

Table 17: Effect of Automatic Enrollment on Spousal Withdrawals

	Incidence (Logit ME)	**Log withdrawals (OLS)**
Hired	0.007**	0.008**
	(0.003)	(0.002)
Treated	0.001	0.001
	(0.002)	(0.002)
Hired×Treated	0.002	0.002
	(0.002)	(0.002)
Female	-0.028***	-0.030***
	(0.005)	(0.005)
Age	-0.005***	-0.016***
	(0.001)	(0.001)
Age Squared	0.000***	0.000***
	(0.000)	(0.000)
Log Wages	0.002	0.001
	(0.001)	(0.001)
Tax Year FE	✓	✓
NAICS Code FE	✓	✓
Mean Control	0.129	$2,375.97
	(0.001)	($67.89)
Observations	231,609	231,624
Note:	*p<0.05, **p<0.01, ***p<0.001	

Table 18: Effect of Automatic Enrollment on Spousal IRA Contributions

	Incidence (Logit ME)	**Log Contributions (OLS)**
Hired	0.009***	0.066***
	(0.002)	(0.013)
Treated	0.006***	0.046***
	(0.001)	(0.011)
Hired×Treated	-0.000	-0.000
	(0.002)	(0.016)
Female	-0.005***	-0.053***
	(0.00164)	(0.0161)
Age	0.005***	0.032***
	(0.002)	(0.011)
Age Squared	-0.000*	-0.000*
	(0.000)	(0.000)
Log Wages	0.006***	0.034***
	(0.001)	(0.003)
Tax Year FE	✓	✓
NAICS Code FE	✓	✓
Mean Control	0.072	$293.82
	(0.001)	($5.28)
Observations	231,602	231,624

Note: *p<0.05, **p<0.01, ***p<0.001

3.3.5 Net Employee Savings

Finally, we want to find out if the positive estimated effect of automatic enrollment on employee contributions is offset by the negative effect of the policy on employee withdrawals from retirement accounts. We previously estimate that employee contributions, measured as a percent of wages, increase by an average of 0.9 percentage points after employees are automatically enrolled in a firm-sponsored plan, equivalent to a 60.8 percent increase in contributions relative to those in the control group (Derby & Mortenson, 2020). The average level of deferred contributions among employees in the control group in our data is $1,345, which equates to an average increase in savings of $823 due to automatic enrollment. In comparison, we estimate that withdrawals increase by an average increase of $737 as a result of the policy, almost completely offsetting the increase in saving, with an estimated average net increase of only $86.

In order to further test whether the increase in withdrawals offsets the effect on savings at the firm level, we use the same procedure as before − our difference-in-differences specification for employees in equation (1) − to estimate the effect of automatic enrollment net savings, calculated as 401(k) contributions minus withdrawals from retirement accounts. For simplicity, we limit our analysis to just the combination of these two amounts, since we do not find that automatic enrollment affects other forms of saving at either the employee or household level. We also estimate the effect by wage level to see if there are differences in the effect of automatic enrollment on net saving for employees at different percentiles of the wage distribution. Note that we restrict our analysis to employees with one year of tenure since that is the preferred specification that we use to measure the effect of the policy in our previous paper (Derby & Mortenson, 2020).

Table 19 below presents the results from these regressions. We find that net savings increase

Table 19: Automatic Enrollment Effect on Net Savings by Wage Quintile

	Log Net Savings by Wage (OLS)					
	All	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
Hired	0.741**	-0.698	0.751***	1.276***	1.236***	1.025***
	(0.225)	(0.518)	(0.174)	(0.247)	(0.257)	(0.213)
Treated	0.099	0.220	0.187*	0.129	-0.094	0.297
	(0.089)	(0.114)	(0.088)	(0.124)	(0.156)	(0.160)
Hired×Treated	0.864***	-0.319	0.632*	1.137***	1.437***	0.884***
	(0.186)	(0.196)	(0.271)	(0.215)	(0.228)	(0.164)
Female	0.107	0.320***	0.369***	0.392***	0.476***	0.393***
	(0.055)	(0.043)	(0.056)	(0.058)	(0.070)	(0.091)
Age	-0.039***	-0.074***	-0.040***	-0.045***	-0.044***	-0.043***
	(0.005)	(0.004)	(0.004)	(0.004)	(0.006)	(0.008)
Age Squared	-0.000	-0.000	-0.000	-0.000*	-0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	1.133***	0.130*	0.022	0.174**	0.256***	0.695***
	(0.064)	(0.061)	(0.071)	(0.062)	(0.052)	(0.059)
Max Match Level	0.806	-0.762	-1.553	2.004	1.935	0.662
	(1.224)	(1.002)	(1.103)	(1.158)	(1.037)	(1.261)
Matching Rate	0.251	0.0821	1.853***	1.196	0.558	-0.037
	(0.975)	(0.553)	(0.610)	(0.689)	(0.552)	(0.677)
3-6 Month Wait	1.925	11.190	9.412	11.201*	1.840	-6.096
	(6.811)	(10.240)	(6.073)	(5.592)	(7.521)	(11.060)
Firm Fixed Effect	✓	✓	✓	✓	✓	✓
Mean Control	$84.12	- $2,206.06	- $705.20	- $434.11	$343.94	$3,276.56
	($33.94)	($76.26)	($45.68)	($52.72)	($68.41)	($114.33)
Observations	257,002	42,482	53,463	54,858	52,431	53,768

Note: *p<0.05, **p<0.01, ***p<0.001

by 86.4 percent. However, this increase is over a much lower baseline net savings level of \$84. Thus our estimated average increase in savings due to automatic enrollment comes out to about \$73, which is close to the difference in our separate estimates of the effect on contributions to 401(k) plans and withdrawals. Furthermore, we find that among those in the lowest wage quintile the estimated effect of automatic enrollment on net savings is negative and no longer significant. It appears that for those with the lowest wages, the increase in contributions to company-sponsored plans due to automatic enrollment is entirely offset by the increase in withdrawals they take from retirement accounts. It is important to note that average net savings are negative for employees in the bottom three wage quintiles, meaning that on average, employees in the bottom 60 percent of wages are withdrawing more funds than they contribute to a 401(k) plan.

In comparison, we find that automatic enrollment has a positive and statistically significant effect on net savings for employees in the top 60 percent of wages. The effect is particularly large for those in the top quintile of the wage distribution. As shown in **Table 19**, we find that net savings increase by 88.4 percent due to automatic enrollment for employees in the top 20 percent of wages, equivalent to an increase in savings of \$2,896. Thus it appears that the policy primarily benefits employees with higher wages.

3.4 Conclusion

We find that the conventional wisdom of the savings benefits of automatic enrollment requires an important caveat. While automatic enrollment does have large positive effects on the participation rates and contribution levels of employees to firm-sponsored retirement plans, those effects are mitigated by individuals withdrawing funds from existing savings accounts. We estimate that automatic enrollment is associated with a 5.1 percentage point (31 percent) increase in the percent of workers withdrawing savings from retirement accounts, and a 36.1 percent increase ($737) in the level of funds withdrawn from those accounts. However, we do not find that employees reduce their IRA contributions in response to automatic enrollment, nor do we find that spousal contributions to retirement savings are affected.

When we estimate the effect of automatic enrollment on net saving we see that the associated increase in withdrawals from retirement accounts almost entirely offsets the increase in deferred contributions to firm-sponsored plans. We estimate that the average increase in net savings across our full sample of employees is still positive, but relatively small, equating to an average increase of just $73 per year. We also find that the benefit is not evenly distributed among employees when we allow the effect to vary by wage quintile. For those in the bottom 20 percent of earnings we find that automatic enrollment has effectively no impact on net savings. Once you take withdrawals into consideration, the estimated effect on saving for this group is negative, and no longer statistically significant. Conversely, we find that automatic enrollment has a large effect on employees in the top 20 percent of wages, who see an average increase in net savings of $2,896 as a result of the policy.

These findings suggest automatic enrollment does not increase saving at the bottom of the wage distribution, and likely exacerbates the inequities of tax-subsidies for retirement savings. Rather than build up savings for lower-income workers, the policy appears to primarily benefit those at the top of the wage distribution, though it is possible low wage individuals are simply taking distributions after receiving the employer match (which does not increase retirement savings, but does increase compensation). This brings into question the premise of requiring firms to implement automatic enrollment as a prerequisite for exemption from nondiscrimination testing requirements. If the objective of this policy is to mitigate the flow of tax subsidies for retirement savings to top earners, it does not appear achieve this goal.

The spillover effects we document highlight the need to more holistically measure individual retirement saving, and suggest previous work that primarily rely only upon data available to firms are missing important information. Simply enrolling employees in a company-sponsored plan may not be enough to get them to save more for retirement if the associated decrease in take-home pay causes them to take withdrawals from retirement accounts. Employees under age 59.5 generally also have to pay a 10 percent penalty on early withdrawals from retirement savings accounts, and if the employer is not offering a match (or the match is below 10 percent) the individual could be worse off due to automatic enrollment.

APPENDIX A: DATA MATCHING PROCEDURE FOR IDENTIFYING FAMILIES IN LIHTC HOUSING

In order to identify families living in LIHTC housing between 1999 and 2012, I match parents' addresses with the publicly available addresses of LIHTC buildings placed in service prior to 2012. I match addresses based on year, zip code, street number, and a unique word or phrase from the street name. For example, if the street address of a LIHTC building constructed in 2002 is "333 Garden Park Road", the algorithm pulls any W-2 form filed in 2002 or later with the same zip code, with the street number "333", and with a street name including the phrase "Garden Park". Common words like "Road" are ignored to ensure that the algorithm also matches street names with abbreviated words like "Garden Park Rd". Numerical street names like "13th Street" are matched using both the numerical form of the street address ("13th") and the alpha-numerical version ("thirteenth"). In order to ensure that these unique street name words or phrases result in accurate matches, I check each LIHTC address individually (38,722 in total), making sure to use phrases that are neither too broad (resulting in too many matches) nor too narrow. Using the same example, I make sure I match street names based on the phrase "Garden Park" rather than "Garden", as the phrase "Garden Park" would provide a better match that would not likely be abbreviated or altered on tax or information returns.

This algorithm is designed to match the addresses of existing LIHTC buildings (placed in service before the date I observe people living in that building) with the addresses of the *parents* of individuals in my sample. A "parent" is defined as any person who claims a child (a person under the age of 18) as a dependent on any tax return from 1999 to 2012. I later restrict this sample to include only families who have children born between 1982 and 1994. If a child is claimed by more than one filer I consider both people to be parents of that child. A very small percentage (less than 1 percent) of children in my sample are claimed as dependents by more than two people. In this case, I select the two people who claim the child as a dependent for the longest period of time and assign those people as the child's parents. I also compare the ages of the parents to that of the individuals they are claiming as dependents. If the difference between the ages of the parent and child is less than 13 or greater than 75, I drop the individual from my sample. This restricts my sample to children between the ages of 6 and 18 who are living with their parents or grandparents (or another relative such as an aunt).

Once I have identified parents living in LIHTC housing, I use their tax identification numbers (TINs) to extract data on their birth dates, addresses, earnings, marital status, and the TINs of children they list as a dependents from 1999 to 2012. This information comes primarily from the Social Security Administration (SSA) database, W-2 form, 1099-MISC form, and 1040 individual tax returns. Since not all units in LIHTC buildings are subsidized, using data from HUD I calculate the LIHTC income ceilings for each building in each year and keep only households that I observe with income limits below the threshold. If I observe two people claiming a child as a dependent at the same address in the same year, I use the sum of both parents' annual earnings as my measure of household income. I also drop anyone in my sample that I observe living in LIHTC housing in 1999 (the first year I am able to observe my data) because I cannot observe how many additional years they may have spent living in the same building prior to 1999. Thus, I limit my sample to families who moved into LIHTC housing in 2000 or later.

I use dependents' TINs to find their birth dates and gender (SSA database), earnings in 2018 (W-2, 1099-MISC, and F-1040 forms), and higher education attendance (1098-T form). I use information return 1098-T to determine how many years each dependent is enrolled in post-secondary education programs. 1098-T forms are tuition statements filed on behalf of students by eligible education institutions in the United States, including all accredited colleges, universities, and vocational schools. I count every year that I observe each dependent receiving a 1098-T tuition statement as a year they are enrolled in higher education, regardless of whether the statement has a positive tuition amount listed (there are circumstances under which a student would still receive a 1098-T tuition form even if they did not pay tuition that year). I then create indicator variables for dependents I observe with two or more years of higher education, and with four or more years of higher education.

All post-secondary education institutions who qualify for federal financial aid under Title IV of the Higher Education Act of 1965 are required to file the 1098-T tuition form, with a few exceptions. Institutions are not required to file for students who are "nonresident aliens", for students whose "qualified tuition and related expenses are entirely waived or paid entirely with scholarships", or for students whose "qualified tuition and related expenses are covered by a formal billing arrangement between an institution and the student's employer or a governmental entity" (The Internal Revenue Service, 2018). This does introduce some measurement error into my data, as there are likely individuals in my sample who receive a full scholarship every year they are in college or university. However, although institutions are not required to file 1098-T forms for students if their expenses are paid entirely with scholarships many schools do still issue these statements with zero tuition amounts for enrolled students (I see over a million 1098-T tuition statements in my data with zero or missing tuition amounts). Furthermore, data on the 1098-T form are not self-reported to the IRS by individuals, and do include students who receive partial scholarships, or a mix of scholarships and loans.

APPENDIX B: REGRESSION VARIABLE DETAILS

4+ Years of Education. My first outcome variable is a binary indicator equal to 1 if individual i is enrolled in a higher education program for four or more years, and 0 otherwise. As explained above, I determine the number of years that an individual is enrolled in a higher education program by summing up the number of years that I observe them receiving a 1098-T tuition statement, including statements filed with zero tuition payments or missing tuition amount, since there are circumstances under which a school would file a tuition statement for as student who did not pay tuition that semester. For the youngest cohort, this is a measure of the number of years up to age 24 that the individual is enrolled in a higher education program. Although some individuals may take longer to complete a four year program, or delay entry into university to serve in the military, most individuals complete four year university degrees by this age. My sample does also include older cohorts − up to age 36 in 2018 − who are more likely to have completed a four year program, if they enrolled in one.

2+ Years of Education. This outcome variable is also a binary indicator, equal to 1 if individual i is enrolled in a higher education program for two or more years, and 0 otherwise. This measure of education would include individuals who completed a 2 year associate's degree, or partially completed a four-year college or university program. It is calculated in the same way I calculate the outcome variable for four years of education.

Log Earnings in 2018. My final outcome variable is the log of gross earnings that individual i earned in 2018. I use the larger amount of two possible calculations of earnings: either the sum of all pre-tax earnings listed on information returns (W-2 and 1099 forms), or the total gross income listed on the F1040 individual tax return. If an individual does not have any income reported on information or tax returns, I interpolate earnings using 2017 returns, adjusted for inflation. Some of the individuals in my sample (about 16 percent) are missing income information or have a reported income of zero in 2017 and 2018. Since they are at an age in 2018 when they could plausibly have zero earnings, I allow earnings to equal zero for these individuals.

Male. This control variable is a binary gender indicator. The variable is equal to 1 if individual i is male, and 0 otherwise. The determination of gender is made according to data from the Social Security Administration database.

Log Household Income. Household income is the most important variable that I control for, since it directly impacts whether families qualify for low income housing, and can also affect children's outcomes later in life. I control for parental income by taking the average of the parents' combined gross earnings − or just a single parent's earnings if only one person claims the child as a dependent − over the full period of time that I observe the family between 1999 and 2012, up until the year the child is 18 years old. If two separate filers claim the same individual as a dependent then I use the sum of those two people's gross earnings as my measure of household income, whether or not the couple are identified as spouses. All earnings are converted to 2018 dollars to adjust for inflation. Since my data start in 1999 I only have a one or two years of data for some families.

Household income/earnings (I use the two terms interchangeably in this paper) are gross earnings that are reported either on Form 1040 individual tax returns, or on W-2 and 1099-MISC information returns filed by employers. If there are missing data for parents in some years I impute the parents' earnings for those years using other income information I have for the parents when the child is between 6 and 18 years old (with different windows of time depending on the birth year of the child). After this imputation only 0.56 percent of dependents have parents with zero or missing household income, and I allow household income to simply remain at zero for these families.

Since household income is such an important variable I try two other specifications to ensure I am controlling for household income in the best way possible. First, I run a separate regression controlling for log household income for every observed year from 1999 to 2012. This ensures that I capture both average income effects and income fluctuation over time. It also controls for household income beyond age 18 for all but one of the dependent cohorts. Aside from the estimated coefficient on household income, changing this specification does not seem to have much of an effect on my results. I also seem to lose some predictive power from household income as there seems to be a much less significant effect on the control variable once children are older than 18. This causes the coefficient on household income in later years (2005 to 2012) to become smaller and close to zero. Thus, I do not use this specification in my final regression.

I also run my regressions with log mean household income (as before) as well as the standard deviation of household income, and the ratio between the maximum and minimum income observed for each family to control for volatility in parents' earnings from year to year. Again, my regression results remain the same, save for the predicted household income effect. So for the ease of interpretation I do not include the standard deviation and min/max ratio of household income in my final regression results. The results from both of these regressions, run with the alternate household income specifications are found in **Tables 20 through 25** below.

Parent's Age at Time of Birth. I use parental age at time of birth to control for differences in children's outcomes as adults stemming from variation in the ages of their parents. I specifically use the mother's age at time of birth, if available. About 65 percent of the individuals in my sample do have a woman of an appropriate age who claims them as a dependent. This does not necessarily mean that the rest of the individuals live only with a father or grandfather. Instead, the reason for this is probably that men are more likely to be the primary tax filer, and mothers who do not work are not likely to report income. If an individual does not have a female parent claiming them as a dependent then I use the male parent's age at the time of birth instead.

Log Area Median Income. I include gross area median income (AMI) as a control variable since it directly affects LIHTC income limits and could also affect dependents' outcomes later in life. AMI is reported by the US Department of Housing and Urban Development annually at the county level. Since AMI changes from year to year, I use the following approach to calculate the control variable. First, I find the AMI for the county corresponding to the LIHTC building that each dependent lives in each year (converted into 2018 dollars to adjust for inflation). I then take the average AMI across all years I observe the dependent living in that building. For dependents who live in more than one LIHTC building, I use the average AMI across county and year (if the buildings are located in different counties). I then take the log of this average to estimate the effect of AMI on education and earnings in terms of percentage changes.

Family Size. Since the number of children a family has can affect the type of unit and building that the family moves into, I include family size (in terms of siblings) in my regression as well. Family size may also affect outcome variables like education and earnings since families with a greater number of children might have less income to spend on each individual child when they are growing up (less income to invest in education, for example). I measure family size by finding the number of unique tax identification numbers (TINs) up to age 18 claimed as dependents by each parent or set of parents.

Log Units in Building. I use the reported number of units in each LIHTC building (provided by HUD) as a control for variation in outcomes that may arise from differences in the type of

LIHTC housing that families move in to. There may be differences, for example, between moving into a duplex with two units and a large apartment complex with hundreds of units.

Filing Status. I include parents' filing status as a control variable as this can influence both the timing of when a family moves into LIHTC housing and dependents' outcomes later in life. I use dummy variables to indicate if parents file in one of six categories: Single, Married Filing Jointly, Married Filing Separately, Head of Household, Qualifying Widow(er) and Spouse not Filing (Other). This variable does not determine marital status exactly, as one can file as Head of Household whether single or married, but most filers fall into one of the first three categories, which provide a clearer marital status designation.

Total Moves. This is the total number of times that a dependent moves to a different zip code before age 18. I do not include this variable in my main regressions but I do include it to check for estimate robustness in **Tables 26 through 28** in **Appendix C**. I include the variable in these regressions to control for variation coming from moves to a different area in a city, or to a different city or state. These moves could contribute both to the timing of when families move into a LIHTC building, and may also affect children's development and subsequent outcomes later in life. One concern with controlling for this variable is that housing stability is a reason why moving into LIHTC housing may lead to better outcomes later in life. Thus, including this control variable may take away some of the explanatory power of my estimate $\hat{\theta}$. However, I also do not want my estimated LIHTC effect to be solely a measure of housing stability, since this would not tell me anything about whether growing up in LIHTC housing specifically leads to better outcomes. If that is the case then I am simply measuring the effect of living in one place for a longer period of time. Howver, as shown in **Appendix C**, when I include total moves in my regression, the estimated $\hat{\theta}$ does not change significantly.

Birth Year. I used a birth year fixed effect to control for differences between individuals of different ages, who enroll in higher education and enter the workforce in different years. Including this fixed effect also helps control for the fact that I cannot observe where individuals are living prior to 1999, so I have data for earlier childhood years for younger cohorts, and less so for the older ones.

Zip Code. I use a zip code fixed effect to account for differences in outcomes due to the location of LIHTC housing. This helps control for differences in housing built in areas that are more rural versus urban, are located in different cities, or even located in parts of the same city that have very different characteristics.

Table 20: Alternate Household Income Results (Part 1): 4+ Years Higher Education

	4+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.039***	0.068***	0.038**
	(0.004)	(0.011)	(0.018)
Male	-0.537***	-0.540***	-0.540***
	(0.007)	(0.011)	(0.018)
Log Household Income	0.203***	0.233***	0.205***
	(0.016)	(0.031)	(0.048)
Log Standard Deviation, HH Inc	0.019***	0.017	0.023
	(0.007)	(0.013)	(0.022)
Log Min/Max Ratio, HH Inc	-0.007*	-0.004	-0.019
	(0.006)	(0.011)	(0.018)
Age of Parent at Birth	0.008***	0.009***	0.009***
	(0.001)	(0.002)	(0.003)
Log Area Median Income	1.482***	0.113	-0.025
	(0.360)	(0.250)	(0.315)
Family Size	-0.103***	-0.122***	-0.134***
	(0.007)	(0.012)	(0.020)
Log Units in Building	0.033**	0.047*	0.030
	(0.014)	(0.024)	(0.052)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.200	0.207	0.232

Note: *p<0.10, **p<0.05, ***p<0.01

Table 21: Alternate Household Income Results (Part 1): 2+ Years Higher Education

	2+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.050***	0.078***	0.036**
	(0.004)	(0.009)	(0.015)
Male	-0.534***	-0.539***	-0.547***
	(0.005)	(0.009)	(0.015)
Log Household Income	0.152***	0.182***	0.164***
	(0.012)	(0.024)	(0.038)
Log Standard Deviation, HH Inc	0.023***	0.021*	0.022
	(0.006)	(0.011)	(0.019)
Log Min/Max Ratio, HH Inc	-0.009*	-0.006	-0.022
	(0.005)	(0.009)	(0.015)
Age of Parent at Birth	0.003***	0.005***	0.005**
	(0.001)	(0.001)	(0.002)
Log Area Median Income	1.831***	0.257	0.328
	(0.337)	(0.236)	(0.391)
Family Size	-0.088***	-0.109***	-0.122***
	(0.006)	(0.010)	(0.017)
Log Units in Building	0.033***	0.044**	0.003
	(0.011)	(0.020)	(0.043)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.376	0.391	0.426

Note: *p<0.10, **p<0.05, ***p<0.01

Table 22: Alternate Household Income Results (Part 1): Earnings

	Log Adult Earnings (in 2018)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.034***	0.055***	0.051**
	(0.003)	(0.014)	(0.025)
Male	-0.430***	-0.447***	-0.525***
	(0.021)	(0.039)	(0.062)
Log Household Income	0.119***	0.130***	0.138***
	(0.009)	0.031)	(0.053)
Log Standard Deviation, HH Inc	0.071***	0.071***	0.056*
	(0.006)	(0.019)	(0.033)
Log Min/Max Ratio, HH Inc	-0.038***	-0.036**	-0.040
	(0.005)	(0.015)	(0.026)
Age of Parent at Birth	-0.002**	-0.001	0.001
	(0.001)	(0.002)	(0.004)
Log Area Median Income	0.659***	0.316	0.712
	(0.112)	(0.418)	(0.785)
Family Size	-0.104***	-0.109***	-0.110***
	(0.007)	(0.020)	(0.035)
Log Units in Building	0.016	0.017	-0.056
	(0.013)	(0.034)	(0.080)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Mean Wage	$23,007	$22,505	$24,127

Note: *p<0.10, **p<0.05, ***p<0.01

Table 23: Alternate Household Income Results (Part 2): 4+ Years Higher Education

	4+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.040***	0.068***	0.038**
	(0.004)	(0.011)	(0.018)
Male	-0.538***	-0.541***	-0.541***
	(0.007)	(0.011)	(0.018)
Age of Parent at Birth	0.009***	0.010***	0.010***
	(0.001)	(0.002)	(0.003)
Log Area Median Income	1.312***	0.089	-0.044
	(0.336)	(0.246)	(0.310)
Family Size	-0.103***	-0.122***	-0.134***
	(0.007)	(0.012)	(0.020)
Log Units in Building	0.032**	0.046*	0.029
	(0.014)	(0.024)	(0.052)
Log Household Income 1999-2012	✓	✓	✓
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.200	0.207	0.232
Note:	*p<0.10, **p<0.05, ***p<0.01		

Table 24: Alternate Household Income Results (Part 2): 2+ Years Higher Education

	2+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.050***	0.078***	0.037**
	(0.004)	(0.009)	(0.015)
Male	-0.535***	-0.540***	-0.548***
	(0.005)	(0.009)	(0.015)
Age of Parent at Birth	0.004***	0.005***	0.006**
	(0.001)	(0.001)	(0.002)
Log Area Median Income	1.679***	0.235	0.320
	(0.320)	(0.233)	(0.390)
Family Size	-0.088***	-0.108***	-0.121***
	(0.006)	(0.010)	(0.017)
Log Units in Building	0.032***	0.043**	0.002
	(0.011)	(0.020)	(0.043)
Log Household Income 1999-2012	✓	✓	✓
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.376	0.391	0.426

Note: *p<0.10, **p<0.05, ***p<0.01

Table 25: Alternate Household Income Results (Part 2): Earnings

	Log Earnings (in 2018)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.036***	0.057***	0.051**
	(0.003)	(0.014)	(0.025)
Male	-0.430***	-0.447***	-0.526***
	(0.021)	(0.039)	(0.062)
Age of Parent at Birth	-0.002**	-0.002	0.002
	(0.001)	(0.002)	(0.004)
Log Area Median Income	0.612***	0.295	0.680
	(0.112)	(0.418)	(0.785)
Family Size	-0.101***	-0.105***	-0.108***
	(0.007)	(0.020)	(0.035)
Log Units in Building	0.015	0.016	-0.055
	(0.013)	(0.034)	(0.080)
Log Household Income 1999-2012	✓	✓	✓
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Mean Wage	$23,007	$22,505	$24,127

Note: *p<0.10, **p<0.05, ***p<0.01

APPENDIX C: LIHTC EFFECT REGRESSION RESULTS CONTROLLING FOR TOTAL MOVES

Table 26: Regression Results Controlling for Total Moves: 4+ Years Higher Education

	4+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.034***	0.063***	0.031**
	(0.004)	(0.011)	(0.014)
Male	-0.537***	-0.540***	-0.475***
	(0.007)	(0.011)	(0.015)
Log Household Income	0.230***	0.254***	0.183***
	(0.013)	(0.024)	(0.032)
Parent's Age at Birth	0.006***	0.008***	0.005**
	(0.001)	(0.002)	(0.002)
Log Area Median Income	1.513***	0.122	-0.030
	(0.365)	(0.251)	(0.263)
Family Size	-0.098***	-0.118***	-0.105***
	(0.007)	(0.012)	(0.171)
Total Moves	-0.028***	-0.026***	-0.032***
	(0.003)	(0.005)	(0.008)
Log Units in Building	0.034**	0.049***	0.041
	(0.014)	(0.024)	(0.044)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.200	0.207	0.232

Note: *p<0.10, **p<0.05, ***p<0.01

Table 27: Regression Results Controlling for Total Moves: 2+ Years Higher Education

	2+ Years Higher Education (Odds Ratios)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.047***	0.075***	0.032**
	(0.004)	(0.009)	(0.015)
Male	-0.535***	-0.539***	-0.547***
	(0.005)	(0.009)	(0.015)
Log Household Income	0.182***	0.207***	0.193***
	(0.010)	(0.019)	(0.031)
Parent's Age at Birth	0.002***	0.004***	0.004
	(0.001)	(0.001)	(0.002)
Log Area Median Income	1.846***	0.265	0.336
	(0.339)	(0.237)	(0.392)
Family Size	-0.085***	-0.106***	-0.118***
	(0.006)	(0.010)	(0.017)
Total Moves	-0.015***	-0.013***	-0.023***
	(0.002)	(0.004)	(0.007)
Log Units in Building	0.033***	0.044**	0.003
	(0.011)	(0.020)	(0.043)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	108,825
Baseline Enroll Rate	0.376	0.391	0.426

Note: *p<0.10, **p<0.05, ***p<0.01

Table 28: Regression Results Controlling for Total Moves: Earnings in 2018

	Log Adult Earnings (in 2018)		
	All LIHTC	Stayers	New Building
LIHTC Years	0.042***	0.063***	0.055**
	(0.003)	(0.015)	(0.025)
Male	-0.430***	-0.447***	-0.525***
	(0.021)	(0.039)	(0.062)
Log Household Income	0.189***	0.201***	0.195***
	(0.007)	(0.024)	(0.041)
Age of Parent at Birth	-0.000	0.001	0.002
	(0.001)	(0.002)	(0.004)
Log Area Median Income	0.652***	0.302	0.729
	(0.112)	(0.418)	(0.786)
Family Size	-0.106***	-0.111***	-0.110***
	(0.007)	(0.020)	(0.035)
Total Moves	-0.033***	0.036***	0.019
	(0.004)	(0.008)	(0.014)
Log Units in Building	0.016	0.015	-0.057
	(0.013)	(0.034)	(0.080)
Filing status fixed effects	✓	✓	✓
Birth year fixed effects	✓	✓	✓
Zip code fixed effects	✓	✓	✓
Observations	540,839	531,088	102,476
Baseline Mean Wage	$23,007	$22,505	$24,127

Note: *p<0.10, **p<0.05, ***p<0.01

APPENDIX D: LIHTC FIXED EFFECT RESULTS FOR INDIVIDUALS WHO REMAIN IN LIHTC HOUSING THROUGH AGE 18

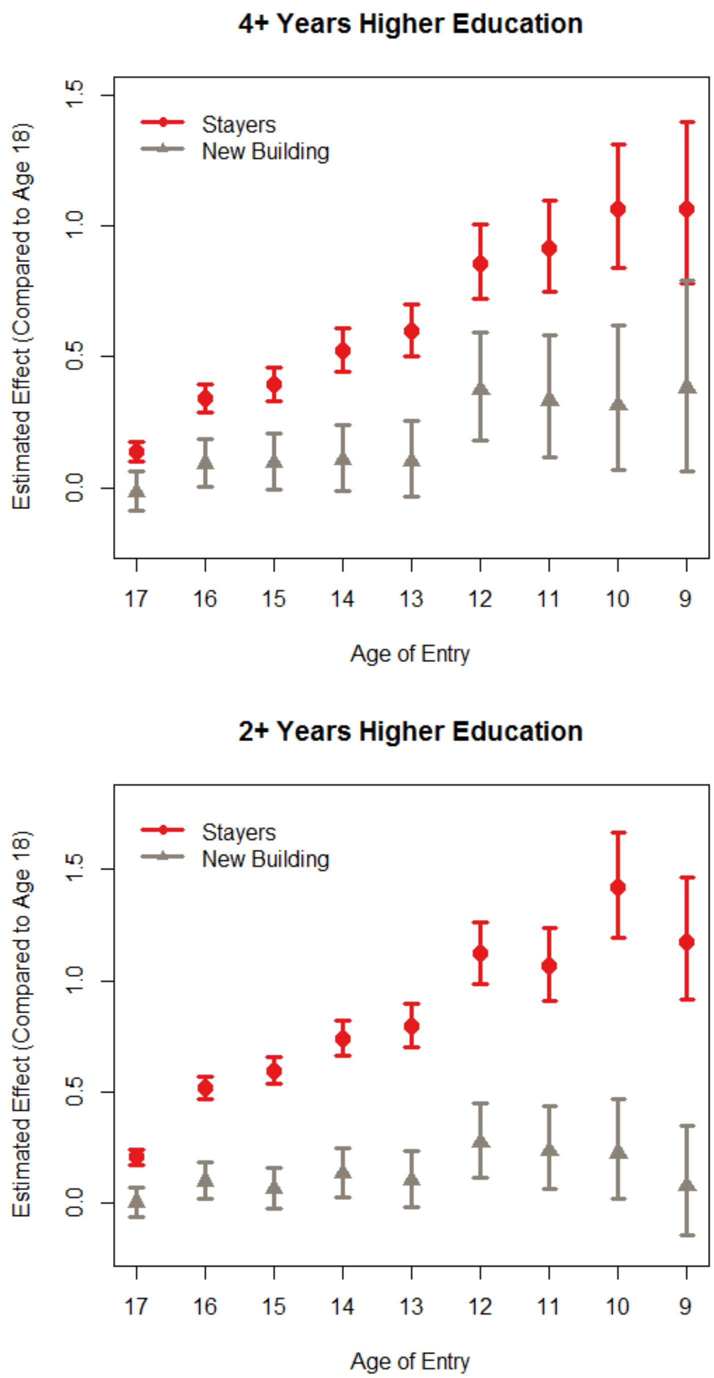

Figure 37: LIHTC Fixed Effects Results for Stayers through Age 18 (Education)

APPENDIX E: HETEROGENEOUS NEIGHBORHOOD EFFECTS QUANTILE CUTOFFS

Table 29: Neighborhood Characteristics: Quantile Cutoff Points

Quantile	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**
Poverty Rate	5.5	8.4	11.3	14.5	17.9	22.2	27.1	32.3	39.7
Med HH Inc (1990)*	21.7	29.1	33.9	38.5	43.3	48.6	53.8	60.3	72.8
Med HH Inc (2012)*	21.6	26.1	30.1	34.2	37.6	42.1	47.7	54.7	66.8
Percent High School Grad	7.5	12.4	16.7	21.1	25.2	30.2	35.3	41.4	48.8
Percent White	11.3	28.3	43.5	53.1	62.0	70.6	79.0	85.6	92.2
Opportunity Measure	0.04	0.06	0.08	0.10	0.12	0.15	0.18	0.22	0.27
Public School Grad		44.00		87.06		95.00		97.00	

* In thousands of US dollars, adjusted for inflation

APPENDIX F: AUTOMATIC ENROLLMENT DATA CONSTRUCTION

This appendix describes the process used to construct the tax data and pension data underlying our analyses.

Tax Data: Form W-2

The employee-level wage and retirement contribution data used in this paper are drawn from the population of Form W-2 filings from 2010 to 2016. The process of extracting these data involves three steps. First, each EIN is classified as belonging to a partnership, corporation, non-profit, government, sole proprietor, or some other type of organization. For this project, we are only interested in private entities, and drop W-2 EINs belonging to state, local, or federal governments. If the EIN belongs to an overarching "parent" company, the EIN of the parent company is also recorded. Second, all W-2s that are not deemed to belong to a government are aggregated by W-2 EIN. As part of the aggregation, we count the number of W-2s, the number of W-2s with deferred compensation, and the amount of Medicare wages. To limit the effect of part time workers in our sum total, we drop all W-2s with less than one quarter of full time work at the federal minimum wage. To reduce computational resources, we drop all W-2 EINs with fewer than 10 W-2s. The final step is to further aggregate the data by parent EIN and year, again counting the number of W-2s, the share with deferred compensation, and the amount of wages.

Pension Data: Form 5500

The plan-level information on pension details used in this paper are drawn from the population of Form 5500 filings from 2010 to 2016. We use the Department of Labor's (DOL) "Bulletin" files, which are cleaned by DOL (cite). We limit our data to single employer defined contribution plan filings with the latest plan year ending date. We do not include plans for which the accrual of benefits has been suspended or plans that have zero participants at the start of the year. We then cross reference this list of filings with the Form 5500 Annual Report from the Department of Labor's Employee Benefits Security Administration. The reports are a set of indices used for bulk downloading of Form 5500 images and attachments in PDF format. (cite) Since we are ultimately interested in using Form 5500 attachments to identify details about the plans' auto-enrollment policies, we limit the list of plans to those that have attachments available for download.

Linking Wage and Pension Data

After gathering the W-2 data and augmented Form 5500 data, we merge the two sets. First, we aggregate all w2 data – the number of employees, the number of employees with deferred compensation, total compensation, and total deferred compensation – at the parent EIN level. Then we merge the two data sets by parent company EIN, and tax year. The table below provides the weighted match rate against all W-2 data in each year, which ranges from 60.02% to 63.53%. The match rate is weighted by the number of employees with deferred contributions, as a percentage of the total number of employees with deferred contributions in each year.

Year	**2010**	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**
Match rate	69.80%	70.18%	69.99%	69.49%	68.94%	68.25%	67.21%

The data set is then further refined to include only companies that are observed in 2 or more consecutive years, so that it is possible to observe a change in their auto-enrollment policy over time. Only 3.2% of plans in the merged data set are excluded as a result of this step.

Identifying Adoption of Auto-Enrollment in Pension Plans

After merging the W-2 data with the augmented Form 5500 data and excluding plans that do not have attachments available for download, we identify companies that adopted auto-enrollment using the following procedure.

First, we use data from the Form 5500 Bulletin files to create an "auto-enrollment flag" for each year from 2010 to 2016. All plans in the Form 5500 data that have a type pension benefit code containing the character string "2S" are flagged as having auto-enrollment. Although this flag is not very accurate when it comes to identifying the exact year in which a company adopted auto-enrollment, it is accurate at identifying companies that adopted automatic enrollment *at some point*, if the company is flagged as having auto-enrollment in one of those six years.

Second, we use Form 5500 attachments, downloaded from the Bulk Form 5500 Image/Attachments Requests site, to identify which year each plan/company adopted auto-enrollment. We use a text-reading algorithm that searches each attachment for specific key words and phrases related to auto-enrollment. Once we have identified whether each plan mentions auto-enrollment, we compare the results to the flag we created using the F5500 Bulletin files. There is a

63.81% match rate between the auto-enrollment flag produced by the text-reading algorithm and the auto-enrollment flag provided in the Bulletin files.

In order to identify the exact year in which each company adopts auto-enrollment, the earliest mention of auto-enrollment in the attachments is compared with the first time that the plan is flagged as having auto-enrollment in the Bulletin files. We also analyze the text of each plan attachment to extract the start date of auto-enrollment, if it is explicitly mentioned in any of the plan attachments for each company. 3,221 companies mention the specific date of auto-enrollment in one of their plan attachments.

APPENDIX G: DIFFERENCE-IN-DIFFERENCES RESULTS

Table 30: Automatic Enrollment Effect on Likelihood of Saving (Difference in Differences)

	Likelihood of Saving (Logit Marginal Effects)		
	All Employees	**1-Year Tenure**	**2-Year Tenure**
Hired	-0.188*	-0.065	-0.017
	(0.090)	(0.039)	(0.024)
Treated	0.001	0.013	0.031
	(0.015)	(0.010)	(0.020)
Hired × Treated	0.304***	0.336***	0.419***
	(0.025)	(0.017)	(0.032)
Female	0.028***	0.024***	0.019***
	(0.005)	(0.005)	(0.004)
Age	0.003***	0.003***	0.003***
	(0.001)	(0.000)	(0.000)
Age Squared	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
Log Wages	0.115***	0.115***	0.120***
	(0.006)	(0.007)	(0.005)
Max Match Level	0.263	0.237	0.309***
	(0.164)	(0.132)	(0.083)
Matching Rate	0.160***	0.135***	0.0907*
	(0.053)	(0.044)	(0.038)
3-6 Month Wait	-0.458	-0.759	-0.267
	(0.580)	(0.440)	(0.460)
Firm Fixed Effect	✓	✓	✓
Mean Control	0.210	0.280	0.343
	(0.407)	(0.449)	(0.475)
Observations	390,733	270,589	109,910
Number of Firms	279	279	279
Pseudo R^2	0.438	0.365	0.345

Note: *p<0.05, **p<0.01, ***p<0.001

Table 31: Effect of Automatic Enrollment on Percent Saved (Difference in Differences)

	Deferred Compensation Percent (OLS)		
	All Employees	**1-Year Tenure**	**2-Year Tenure**
Hired	-1.13*	-0.42	-0.157
	(0.56)	(0.23)	(0.11)
Treated	-0.10	0.02	0.03
	(0.09)	(0.07)	(0.11)
Hired × Treated	0.87***	0.96***	1.23***
	(0.10)	(0.07)	(0.17)
Female	0.12***	0.10*	0.13***
	(0.05)	(0.05)	(0.05)
Age	0.04***	0.05***	0.05***
	(0.00)	(0.01)	(0.00)
Age Squared	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
Log Wages	0.12***	0.12***	0.12***
	(0.01)	(0.01)	(0.01)
Max Match Level	0.47	0.56	0.44
	(0.67)	(0.62)	(0.92)
Matching Rate	-0.17	-0.24	-0.10
	(0.58)	(0.55)	(0.50)
3-6 Month Wait	-5.59	-10.07	-10.02
	(6.01)	(6.31)	(6.42)
Firm Fixed Effect	✓	✓	✓
Mean Control	1.11	1.48	1.88
	(3.24)	(3.65)	(4.03)
Observations	390,733	270,589	109,910
Number of Firms	279	279	279
Adjusted R^2	0.125	0.144	0.155

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

APPENDIX H: AUTOMATIC ENROLLMENT EFFECT BY DEFAULT RATE

Table 32: Effect of Automatic Enrollment on Participation in Retirement Plans by Default Enrollment Rate

	Likelihood of Saving by Default Rate (Logit ME)					
	1%	**2%**	**3%**	**4%**	**5%**	**6%**
Treated	0.398***	0.445***	0.300***	0.315***	0.355***	0.276***
	(0.012)	(0.014)	(0.023)	(0.042)	(0.030)	(0.039)
Female	0.019**	0.035***	0.046**	0.003	0.049***	0.041***
	(0.007)	(0.010)	(0.015)	(0.009)	(0.009)	(0.012)
Age	0.000	-0.001	-0.001	0.001	0.003***	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Age Squared	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.113***	0.136***	0.150***	0.106***	0.182***	0.154***
	(0.014)	(0.017)	(0.028)	(0.015)	(0.016)	(0.012)
Max Match Level	-0.775	-3.128	0.087	4.913***	-0.174	0.263
	(5.701)	(2.281)	(0.175)	(1.152)	(1.081)	(0.312)
Matching Rate	0.001	-0.179	0.301	0.089	-0.343	0.456***
	(0.765)	(0.123)	(0.169)	(0.116)	(0.278)	(0.116)
3-6 Month Wait	1.020	1.482	1.204	-0.389	-1.213	-0.899
	(1.437)	(0.805)	(1.111)	(1.201)	(1.451)	(1.635)
Firm Fixed Effect	✓	✓	✓	✓	✓	✓
Mean Control	0.225	0.257	0.242	0.147	0.580	0.474
	(0.3418)	(0.437)	(0.429)	(0.354)	(0.494)	(0.499)
Observations	16,436	24,822	145,480	39,740	10,903	18,501
Number of Firms	28	40	128	21	19	24
Pseudo R^2	0.440	0.400	0.345	0.453	0.283	0.257

Note: *p<0.05, **p<0.01, ***p<0.001

Table 33: Effect of Automatic Enrollment on Percent Saved, by Default Enrollment Rate

	% Deferred Compensation by Default Rate (OLS)					
	1%	**2%**	**3%**	**4%**	**5%**	**6%**
Treated	0.60***	0.79***	0.87***	1.06*	1.44***	1.50***
	(0.20)	(0.17)	(0.12)	(0.46)	(0.22)	(0.18)
Female	-0.10	0.01	0.28***	-0.08	0.55**	0.46**
	(0.07)	(0.12)	(0.05)	(0.06)	(0.18)	(0.14)
Age	0.02**	0.02***	0.02*	0.02***	0.07**	0.03**
	(0.01)	(0.01)	(0.01)	(0.00)	(0.02)	(0.01)
Age Squared	-0.00	0.00	-0.00	-0.00*	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log Wages	0.84***	0.91***	1.13***	0.82***	1.045***	1.01***
	(0.11)	(0.09)	(0.08)	(0.08)	(0.07)	(0.11)
Max Match Level	-3.73	-19.93*	1.21	32.11	3.92	-3.80***
	(70.68)	(9.01)	(0.80)	(18.21)	(4.84)	(0.85)
Matching Rate	0.94	-2.19***	0.63	0.84	-1.70*	0.87*
	(6.12)	(0.32)	(1.63)	(1.17)	(0.74)	(0.34)
3-6 Month Wait	12.72	-4.03	1.74	51.91	-36.10	-16.77
	(11.72)	(8.66)	(8.47)	(25.41)	(23.70)	(12.00)
Firm Fixed Effect	✓	✓	✓	✓	✓	✓
Mean Control	1.16	1.25	1.13	0.74	3.79	2.30
	(3.64)	(3.32)	(3.12)	(2.70)	(5.08)	(4.08)
Observations	16,436	24,822	145,480	39,740	10,903	18,501
Number of Firms	28	40	128	21	19	24
Adjusted R^2	0.152	0.157	0.201	0.267	0.167	0.205

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

Table 34: Effect of Automatic Enrollment on Percent Saved Among Savers, by Default Enrollment Rate

	% Deferred Compensation by Default Rate (OLS)					
	1%	**2%**	**3%**	**4%**	**5%**	**6%**
Treated	-2.27***	-1.87***	-0.73**	-1.04***	-0.62	-0.03
	(0.46)	(0.16)	(0.25)	(0.25)	(0.42)	(0.35)
Female	-0.04	-0.10	0.53***	0.10	0.61*	0.96***
	(0.22)	(0.24)	(0.12)	(0.13)	(0.28)	(0.16)
Age	0.05***	0.07***	0.08***	0.07***	0.09***	0.07***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Age Squared	-0.00	-0.00	-0.00	-0.00**	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log Wages	0.79***	0.83***	0.71***	0.78***	-0.31	0.51
	(0.18)	(0.15)	(0.07)	(0.15)	(0.20)	(0.25)
Max Match Level	55.80	20.47	2.32	11.10	43.89	-5.02
	(45.14)	(30.59)	(1.21)	(23.48)	(70.48)	(3.18)
Matching Rate	28.71*	-0.80	-1.02	-0.34	-4.16	-1.02
	(13.15)	(1.51)	(2.43)	(0.84)	(6.36)	(0.80)
3-6 Month Wait	38.29	2.74	-0.06	34.12*	-8.33	-19.55
	(25.13)	(22.73)	(12.58)	(13.54)	(19.14)	(17.30)
Firm Fixed Effect	✓	✓	✓	✓	✓	✓
Mean Control	5.97	5.91	5.64	6.07	6.76	6.12
	(6.42)	(5.14)	(5.11)	(5.97)	(4.63)	(4.90)
Observations	2,214	3,789	22,717	4,182	4,024	3,644
Number of Firms	28	40	128	21	19	24
Adjusted R^2	0.113	0.106	0.105	0.173	0.081	0.107

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

APPENDIX I: AUTO-ENROLLMENT HETEROGENEOUS WAGE EFFECTS

Table 35: Automatic Enrollment Effect on 401(k) Participation by Wage Level

	Likelihood of Saving by Wage Quintile (Logit ME)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Treated	0.423***	0.435***	0.430***	0.432***	0.369***
	(0.050)	(0.031)	(0.028)	(0.025)	(0.026)
Female	0.040***	0.032***	0.028***	0.028***	0.045***
	(0.009)	(0.008)	(0.006)	(0.006)	(0.007)
Age	0.001*	0.001*	0.001	0.001*	-0.000
	(0.001)	(0.000)	(0.000)	(0.000)	(0.001)
Age Squared	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.149***	0.162***	0.141***	0.153***	0.149***
	(0.012)	(0.022)	(0.018)	(0.013)	(0.014)
Max Match Level	0.396*	0.308	-1.080*	-1.266*	0.172
	(0.197)	(0.391)	(0.462)	(0.564)	(0.373)
Matching Rate	0.416**	0.139	0.175	-0.198*	-0.070
	(0.148)	(0.203)	(0.153)	(0.090)	(0.115)
3-6 Month Wait	0.473	0.869	0.742	1.132	0.849
	(0.874)	(0.752)	(0.858)	(0.766)	(1.372)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	0.282	0.294	0.310	0.370	0.447
	(0.450)	(0.455)	(0.463)	(0.483)	(0.497)
Observations	21,248	21,090	21,236	21,345	21,818
Number of Firms	279	279	279	279	279
Pseudo R^2	0.335	0.361	0.368	0.335	0.342

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

Table 36: Automatic Enrollment Effect on Percent Saved by Wage Level

	% Deferred Compensation by Wage Quintile (OLS)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Treated	1.17***	1.38***	1.48***	1.16***	1.12***
	(0.30)	(0.23)	(0.14)	(0.15)	(0.22)
Female	0.35***	0.18**	0.16*	0.30***	0.40***
	(0.10)	(0.07)	(0.06)	(0.09)	(0.12)
Age	0.03***	0.03***	0.02***	0.03***	0.03***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Age Squared	0.00	-0.00	0.00	0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log Wages	0.93***	0.79***	0.69***	0.92***	0.60*
	(0.10)	(0.10)	(0.11)	(0.09)	(0.25)
Max Match Level	1.58	-1.23	-4.84	-1.80	-14.97*
	(1.10)	(2.54)	(4.54)	(4.46)	(6.06)
Matching Rate	1.90**	1.00	0.90	-0.81	-0.69
	(0.58)	(0.52)	(1.62)	(0.50)	(0.69)
3-6 Month Wait	-2.36	3.00	-7.37	-5.79	8.32
	(6.35)	(11.91)	(8.29)	(12.00)	(14.78)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	1.40	1.49	1.51	2.05	2.82
	(3.51)	(3.86)	(3.53)	(4.33)	(4.56)
Observations	21,982	21,982	21,982	21,982	21,982
Number of Firms	279	279	279	279	279
Adjusted R^2	0.414	0.246	0.235	0.217	0.214

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

APPENDIX J: AUTO-ENROLLMENT HETEROGENEOUS FIRM EFFECTS

Table 37: Automatic Enrollment Effect on 401(k) Participation by Firm Size

	Likelihood of Saving by Firm Size (Logit ME)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Treated	0.382***	0.381***	0.394***	0.394***	0.281***
	(0.018)	(0.015)	(0.019)	(0.018)	(0.025)
Female	0.027**	0.042***	0.054***	0.033***	0.038**
	(0.009)	(0.010)	(0.011)	(0.007)	(0.013)
Age	0.001	0.001	0.000	0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Age Squared	0.000	-0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.139***	0.144***	0.155***	0.148***	0.141***
	(0.012)	(0.010)	(0.012)	(0.012)	(0.024)
Max Match Level	0.184*	-1.669**	-0.051	0.167	0.328
	(0.080)	(0.516)	(0.501)	(0.107)	(0.315)
Matching Rate	0.895**	-0.348*	0.065	0.136	0.376***
	(0.343)	(0.167)	(0.172)	(0.137)	(0.090)
3-6 Month Wait	-0.469	1.406	-0.553	0.363	0.372
	(0.814)	(1.089)	(0.887)	(1.037)	(1.195)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	0.323	0.307	0.295	0.298	0.268
	(0.468)	(0.461)	(0.456)	(0.458)	(0.443)
Observations	15,474	17,001	21,356	39,798	176,960
Number of Firms	56	57	54	56	56
Pseudo R^2	0.377	0.338	0.322	0.326	0.289

Note: *p<0.05, **p<0.01, ***p<0.001

Table 38: Automatic Enrollment Effect on Percent Saved by Firm Size

	% Deferred Compensation by Firm Size (OLS)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Treated	1.25***	1.17***	1.11***	1.11***	0.81***
	(0.22)	(0.20)	(0.14)	(0.13)	(0.14)
Female	0.19*	0.31***	0.22**	0.15**	0.23**
	(0.08)	(0.09)	(0.07)	(0.05)	(0.08)
Age	0.03***	0.03***	0.02***	0.03***	0.02**
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)
Age Squared	0.00	-0.00	-0.00	-0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log Wages	1.05***	0.95***	1.11***	1.08***	1.05***
	(0.07)	(0.07)	(0.09)	(0.06)	(0.08)
Max Match Level	1.39**	-7.94	-0.98	2.15***	-0.11
	(0.45)	(7.48)	(3.00)	(0.57)	(1.81)
Matching Rate	4.01*	-3.56*	-1.79	0.62*	1.01
	(1.54)	(1.58)	(1.08)	(0.24)	(0.89)
3-6 Month Wait	-2.12	13.29	8.24	1.55	-9.66
	(15.31)	(11.73)	(13.13)	(7.13)	(7.26)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	1.85	1.71	1.43	1.48	1.44
	(4.46)	(4.39)	(3.76)	(3.41)	(3.52)
Observations	15,474	17,001	21,356	39,798	176,960
Number of Firms	56	57	54	56	56
Adjusted R^2	0.141	0.235	0.202	0.182	0.232

Note: *p<0.05, **p<0.01, ***p<0.001

Table 39: Automatic Enrollment Effect on 401(k) Participation by Mean Wage (Firm Level)

	Likelihood of Saving by Firm Wage (Logit ME)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Treated	0.273***	0.357***	0.400***	0.289***	0.353***
	(0.034)	(0.029)	(0.015)	(0.030)	(0.021)
Female	0.033	0.036***	0.029***	0.039***	0.057***
	(0.020)	(0.010)	(0.006)	(0.012)	(0.007)
Age	-0.002	0.001	0.001*	0.002	0.003*
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
Age Squared	-0.000	-0.000	-0.000	0.000	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.157***	0.112***	0.135***	0.162***	0.172***
	(0.017)	(0.025)	(0.010)	(0.011)	(0.009)
Max Match Level	0.810	-0.423	-0.426	0.145	1.350***
	(0.623)	(0.439)	(0.353)	(0.237)	(0.239)
Matching Rate	0.121	0.429**	0.648***	0.193	0.199
	(0.107)	(0.162)	(0.162)	(0.178)	(0.111)
3-6 Month Wait	3.850	0.955	1.094	-2.369	-3.080*
	(3.076)	(0.708)	(0.883)	(1.237)	(1.278)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	0.095	0.275	0.341	0.401	0.517
	(0.294)	(0.447)	(0.474)	(0.490)	(0.500)
Observations	83,666	82,272	35,124	37,103	32,424
Number of Firms	56	56	55	56	56
Pseudo R^2	0.421	0.276	0.315	0.233	0.288

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

Table 40: Automatic Enrollment Effect on Percent Saved by Mean Wage (Firm Level)

	% Deferred Compensation by Firm Wage (OLS)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Treated	0.64***	1.22***	1.01***	1.09***	0.89***
	(0.11)	(0.17)	(0.17)	(0.15)	(0.21)
Female	0.11	0.20*	0.14*	0.41***	0.41***
	(0.10)	(0.08)	(0.06)	(0.11)	(0.10)
Age	0.00	0.04***	0.03***	0.05***	0.08***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.02)
Age Squared	-0.00	0.00***	0.00	0.00*	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log Wages	0.98***	0.96***	1.06***	1.10***	1.19***
	(0.17)	(0.04)	(0.06)	(0.06)	(0.08)
Max Match Level	1.16	0.74	-1.22	-0.49	7.34
	(0.99)	(1.25)	(1.52)	(0.86)	(4.37)
Matching Rate	-0.17	-0.41	3.65*	0.05	-0.32
	(0.22)	(1.77)	(1.38)	(0.63)	(0.96)
3-6 Month Wait	17.86	-22.84***	-14.44	-20.69**	-46.40
	(22.77)	(5.21)	(13.02)	(7.12)	(23.99)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	0.40	1.24	1.68	2.17	3.40
	(1.96)	(3.23)	(3.92)	(4.31)	(5.00)
Observations	83,666	82,272	35,124	37,103	32,424
Number of Firms	56	56	55	56	56
Adjusted R^2	0.201	0.149	0.165	0.168	0.204

Note: *p<0.05, **p<0.01, ***p<0.001

APPENDIX K: AUTO-ENROLLMENT REGRESSION RESULTS BY NAICS CODE

Table 41: NAICS Code Descriptions for Industry Sub-Samples

Code	**Category**	**Description**
23	Construction	Construction of buildings or engineering projects
32	Manufacturing 1	Wood, paper, plastic, and chemical manufacturing
33	Manufacturing 2	Metal, machinery, computer, and electronics manufacturing
42	Wholesale Trade	Wholesaling merchandise, without transformation, and rendering services incidental to the sale of merchandise
44	Retail Trade	Retailing merchandise, without transformation, and rendering services incidental to the sale of merchandise
48	Transportation	Providing transportation of passengers and cargo, warehousing and storage for goods, scenic and sightseeing transportation
51	News & Information	Producing and distributing information and cultural products, as well as data or communications
52	Finance & Insurance	Financial transactions (transactions involving the creation, liquidation, or change in ownership of financial assets)
53	Real Estate	Renting, leasing, or otherwise allowing the use of tangible or intangible assets
54	Professional Services	Professional, scientific, and technical services in the fields of law, engineering, computer programming, and others
55	Firm Management	Establishments that administer, oversee, and manage enterprises (strategic or organizational planning and decision making roles)
56	Admin Services	Administrative Support and Waste Management: routine support activities for the day-to-day operations of other organizations
61	Education	Instruction and training provided by specialized establishments, such as schools, colleges, universities, and training centers
62	Health Care	Providing health care and social assistance for individuals
71	Entertainment	Operating facilities or providing services to meet varied cultural, entertainment, and recreational interests of their patrons
72	Hotel & Food	Providing customers with lodging and/or preparing meals, snacks, and beverages for immediate consumption

Table 42: Effects of Automatic Enrollment on Participation (Logit) and Savings (OLS) by Industry

Industry	**# Firms**	**Observations**	**Effect (Logit ME)**	**Effect (OLS)**
Construction	8	3,807	0.26 (0.06)	0.98 (0.22)
Manufacturing 1	9	6,097	0.42 (0.03)	1.48 (0.49)
Manufacturing 2	19	12,418	0.41 (0.03)	1.04 (0.49)
Wholesale Trade	11	7,190	0.28 (0.06)	2.00 (0.58)
Retail Trade	10	6,334	0.27 (0.04)	0.41 (0.21)
Transportation	11	6,967	0.41 (0.03)	1.12 (0.40)
News & Information	12	11,066	0.29 (0.06)	1.15 (0.20)
Finance & Insurance	16	21,397	0.33 (0.06)	0.85 (0.28)
Real Estate	5	2,921	0.39 (0.05)	2.40 (0.27)
Professional Services	40	18,935	0.33 (0.02)	0.85 (0.26)
Firm Management	36	103,038	0.25 (0.04)	0.70 (0.16)
Admin Services	9	5,599	0.38 (0.03)	0.42 (0.14)
Education	10	5,521	0.46 (0.01)	1.53 (0.38)
Health Care	37	22,987	0.44 (0.01)	1.73 (0.31)
Entertainment	3	20,309	0.20 (0.02)	0.48 (0.03)
Hotel & Food	5	4,611	0.27 (0.05)	0.97 (0.08)

APPENDIX L: EFFECTS OF AUTOMATIC ENROLLMENT ON WITHDRAWALS (SAVERS)

Table 43: Effect of Automatic Enrollment on Employee Withdrawals (among Savers)

	Incidence (Logit)	**Savings Rate (OLS)**
Hired	0.054***	0.376***
	(0.005)	(0.039)
Treated	-0.016*	-0.085
	(0.007)	(0.047)
Hired×Treated	0.055***	0.335***
	(0.010)	(0.080)
Female	-0.012**	-0.142***
	(0.004)	(0.029)
Age	0.005***	0.059***
	(0.000)	(0.003)
Age Squared	-0.000	0.000
	(0.000)	(0.000)
Log Wages	-0.041***	-0.263***
	(0.004)	(0.026)
Max Match Level	0.022	0.247
	(0.103)	(0.803)
Matching Rate	-0.020	0.053
	(0.079)	(0.568)
3-6 Month Wait	0.348	-3.309
	(0.440)	(3.838)
Firm Fixed Effect	✓	✓
Mean Control	0.209	$2,682.17
	(0.002)	($68.17)
Observations	291,399	291,399
Number of Firms	279	279

Note: *p<0.05, **p<0.01, ***p<0.001

APPENDIX M: HETEROGENEOUS WAGE EFFECTS EMPLOYEE SAVINGS

Table 44: Automatic Enrollment Effect on the Incidence of Employee Withdrawals by Wage Level

	Incidence by Wage Quintile (Logit ME)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Hired	0.072***	0.021*	0.009	0.015***	0.019***
	(0.019)	(0.009)	(0.007)	(0.005)	(0.005)
Treated	-0.017***	-0.015***	-0.004	0.007	-0.006
	(0.006)	(0.005)	(0.005)	(0.006)	(0.005)
Hired×Treated	0.107***	0.095***	0.052***	0.019***	0.004
	(0.019)	(0.014)	(0.012)	(0.007)	(0.006)
Female	-0.018***	-0.012*	-0.008	-0.006	-0.007
	(0.004)	(0.004)	(0.004)	(0.006)	(0.004)
Age	0.007***	0.005***	0.005***	0.006***	0.006***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age Squared	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.022***	0.010	-0.009	-0.011***	-0.019***
	(0.006)	(0.007)	(0.005)	(0.004)	(0.003)
Max Match Level	0.141	0.202*	0.042	0.104	0.049
	(0.123)	(0.094)	(0.088)	(0.094)	(0.070)
Matching Rate	0.060	-0.009	-0.034	-0.001	0.032
	(0.073)	(0.039)	(0.045)	(0.038)	(0.034)
3-6 Month Wait	-1.014	-0.183	-0.297	-0.397	-0.111
	(0.746)	(0.484)	(0.432)	(0.497)	(0.517)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	0.162	0.148	0.165	0.167	0.162
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	42,397	53,436	54,852	52,416	53,751

Note: *p<0.05, **p<0.01, ***p<0.001

Table 45: Automatic Enrollment Effect on the Level of Employee Withdrawals by Wage Level

	Log Withdrawals by Wage Quintile (OLS)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Hired	0.470*	0.107	0.051	0.109*	0.155***
	(0.213)	(0.073)	(0.053)	(0.046)	(0.039)
Treated	-0.121***	-0.093*	-0.025	0.065	-0.056
	(0.045)	(0.036)	(0.041)	(0.046)	(0.041)
Hired×Treated	0.808***	0.704***	0.379***	0.151*	0.055
	(0.160)	(0.102)	(0.092)	(0.061)	(0.045)
Female	-0.269***	-0.145***	-0.094***	-0.089	-0.103***
	(0.037)	(0.034)	(0.031)	(0.046)	(0.038)
Age	0.080***	0.050***	0.056***	0.058***	0.065***
	(0.004)	(0.003)	(0.002)	(0.003)	(0.004)
Age Squared	0.000	0.000	0.000*	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.153***	0.123*	-0.033	-0.061*	-0.177***
	(0.0560)	(0.049)	(0.0290)	(0.0254)	(0.0221)
Max Match Level	0.503	1.190	0.227	0.561	0.359
	(0.549)	(0.814)	(0.495)	(0.624)	(0.613)
Matching Rate	0.368	-0.121	-0.253	-0.045	0.180
	(0.632)	(0.278)	(0.314)	(0.231)	(0.288)
3-6 Month Wait	-10.460	-3.580	-8.733*	-7.734	-2.318
	(9.655)	(5.118)	(3.733)	(4.862)	(5.688)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	$2,484.03	$1,195.68	$1,480.02	$1,861.48	$3,177.00
	($62.91)	($42.19)	($48.66)	($56.84)	($97.98)
Observations	42,397	53,436	54,852	52,416	53,751

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

Table 46: Automatic Enrollment Effect on the Incidence of Employee IRA Contributions by Wage Level

	Incidence by Wage Quintile (Logit ME)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Hired	-0.003**	-0.000	0.001	0.005***	-0.004
	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)
Treated	-0.004*	-0.002	-0.003	-0.002	-0.001
	(0.002)	(0.001)	(0.002)	(0.003)	(0.003)
Hired×Treated	0.005*	0.001	0.002	0.005	-0.001
	(0.002)	(0.002)	(0.003)	(0.002)	(0.003)
Female	0.001	-0.000	0.003	0.003	-0.006*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Age	0.001***	0.001***	0.001***	0.000***	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age Squared	0.000	-0.000	-0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	0.003*	-0.001	0.001	0.005*	0.005*
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Max Match Level	0.001	-0.038	-0.026	0.052	-0.004
	(0.012)	(0.020)	(0.025)	(0.031)	(0.044)
Matching Rate	0.024	0.011	-0.013	-0.006	-0.034
	(0.018)	(0.006)	(0.010)	(0.013)	(0.028)
3-6 Month Wait	0.140	0.069	-0.069	-0.514	
	(0.161)	(0.144)	(0.169)	(0.321)	(0.437)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	0.022	0.026	0.041	0.070	0.100
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Observations	42,482	53,463	54,858	52,432	53,767

Note: *p<0.05, **p<0.01, ***p<0.001

Table 47: Automatic Enrollment Effect on the Level of Employee IRA Contributions by Wage Level

	Log Contributions by Wage Quintile (OLS)				
	Q1	**Q2**	**Q3**	**Q4**	**Q5**
Hired	-0.025*	0.004	0.016	0.046***	-0.029
	(0.013)	(0.008)	(0.012)	(0.014)	(0.022)
Treated	-0.036**	-0.013	-0.021	-0.012	-0.001
	(0.014)	(0.008)	(0.014)	(0.020)	(0.027)
Hired×Treated	0.032*	0.002	0.010	-0.004	-0.006
	(0.015)	(0.012)	(0.018)	(0.017)	(0.024)
Female	-0.001	-0.007	0.016	0.022	-0.050
	(0.013)	(0.012)	(0.015)	(0.017)	(0.026)
Age	0.008***	0.007***	0.006***	0.004***	-0.001
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Age Squared	0.000	-0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log Wages	-0.002	0.013	0.040**	0.059**	
	(0.014)	(0.008)	(0.011)	(0.015)	(0.020)
Max Match Level	-0.076	-0.332	-0.220	0.300	0.062
	(0.165)	(0.272)	(0.225)	(0.272)	(0.417)
Matching Rate	0.140	0.059	-0.085	-0.047	-0.288
	(0.088)	(0.042)	(0.073)	(0.114)	(0.241)
3-6 Month Wait	-3.328	-0.058	-0.484	-0.124	-3.365
	(2.087)	(1.418)	(1.435)	(2.510)	(3.411)
Firm Fixed Effect	✓	✓	✓	✓	✓
Mean Control	$76.47	$76.91	$108.64	$222.00	$436.08
	($8.45)	($7.15)	($4.14)	($6.51)	($11.66)
Observations	42,482	53,463	54,858	52,432	53,767

Note: $^{*}p<0.05$, $^{**}p<0.01$, $^{***}p<0.001$

REFERENCES

Alana Semuels. Poor Girls Are Leaving Their Brothers Behind. *The Atlantic*, 27 November 2017.

Baum-Snow, Nathaniel, and Justin Marion. The Effects of Low Income Housing Tax Credit Developments on Neighborhoods. *Journal of Public Economics*, 96, 2009.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States. In *Social Security Policy in a Changing Environment*, pages 167–195. University of Chicago Press, 2008.

Black, David. Low-Income Housing Tax Credits: Affordable Housing Investment Opportunities for Banks. *Office of the Comptroller of the Currency*, 2014.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter. The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility. *National Bureau of Economic Research*, 2018.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, 106(4):855–902, 2016.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick. For Better or for Worse: Default Effects and 401 (k) Savings Behavior. In *Perspectives on the Economics of Aging*, pages 81–126. University of Chicago Press, 2004.

Clampet-Lundquist, Susan, Kathryn Edin, Kathryn, Jeffrey R. Kling and Duncan, Greg. Moving At-Risk Teenagers out of High-Risk Neighborhoods: Why Girls Fare Better than Boys. *American Journal of Sociology*, 116(4), 2006.

Derby, Elena C. and Jake Mortenson. Savings Responses to Auto-Enrollment: Evidence from a Large Panel of Worker-Employer Linked Data. *Joint Committee on Taxation Working Paper*, 2020.

Diamond, Rebecca and Tim McQuade. Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development. *Journal of Political Economy*, 127 (3), 2019.

Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez. Saving Incentives for Low-and Middle-Income Families: Evidence from a Field Experiment with H&R Block. *The Quarterly Journal of Economics*, 121(4):1311–1346, 2006.

Ellen, Ingrid Gould, Keren Horn, Yiwen Kuai, Roman Pazuniak, and Michael David Williams. Effect of QAP Incentives on the Location of LIHTC Properties. *U.S. Department of Housing and Urban Development*, 2015.

Emple, Hannah. Donald Sterling and Other Landlords Openly Discriminate Against Low-Income Renters. *The Atlantic*, 22 May 2014.

Engelhardt, Gary V. and Anil Kumar. Employer Matching and 401 (k) Saving: Evidence from the Health and Retirement Study. *Journal of Public Economics*, 91(10):1920–1943, 2007.

Even, William E. and David A. Macpherson. The Effects of Employer Matching in 401 (k) Plans. *Industrial Relations: A Journal of Economy and Society*, 44(3):525–549, 2005.

Falk, Armin and Andrea Ichino. Clean Evidence on Peer Effects. *Journal of Labor Economics*, 24 (1):39–57, 2006.

Falk, Justin, and Nadia Karamcheva. The Effect of the Employer Match and Defaults on Federal Workers' Savings Behavior in the Thrift Savings Plan. *Congressional Budget Office*, 2019.

Gelber, Alexander M. How do 401(k)s Affect Saving? Evidence from changes in 401(k) Eligibility. *American Economic Journal: Economic Policy*, 3:103–112, 2011.

Gennetian, Lisa A., Lisa Sanbonmatsu, and Jens Ludwig. An Overview of Moving to Opportunity: A Random Assignment Housing Mobility Study in Five US Cities. *Neighborhood and Life Chances: How Place Matters in Modern America*, pages 163–178, 2011.

Heckman, James. Sample Selection Bias as a Specification Error. *Econometrica*, 47(1):153–161, 1979.

Hollar, Michael. Understanding whom the LIHTC program serves: Tenants in LIHTC units as of December 31, 2012. *The United States Department of Housing and Urban Development*, 2014.

Keightley, Mark P and Jeffrey M. Stupak. An Introduction to the Low-Income Housing Tax Credit. *Congressional Research Service*, 2019.

Kingsley, G Thomas. Trends in Housing Problems and Federal Housing Assistance. *Urban Institute*, 2017.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz. Experimental Analysis of Neighborhood Effects. *Econometrica*, 75(1), 2007.

Kusko, Andrea, James M. Poterba, David Wilcox, Olivia Mitchell, and Sylvester Schieber. *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement*. University of Pennsylvania Press, 1998.

Leventhal, Tama and Jeanne Brooks-Gunn. Moving to Opportunity: An Experimental Study of Neighborhood Effects on Mental Health. *American Journal of Public Health*, 93(9):1576–1582, 2003.

Ludwig, Jens, Greg J. Duncan, Lisa A. Gennetian, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, and Lisa Sanbonmatsu. Long-Term Neighborhood Effects on Low-Income Families: Evidence from Moving to Opportunity. *American Economic Review*, 103(3):226–31, 2013.

Madrian, Brigitte C. and Dennis F. Shea. The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187, 2001.

McClure, Kirk. The Low-Income Housing Tax Credit Program Goes Mainstream and Moves to the Suburbs. *Housing Policy Debate*, 17(3), 2006.

Munnell, Alicia H., Annika Sunden, and Catherine Taylor. What Determines 401 (k) Participation and Contributions. *Social Security Bulletin*, 64:64, 2001.

Nessmith, William E., Stephen P. Utkus, and Jean A. Young. Measuring the Effectiveness of Automatic Enrollment. *Vanguard Center for Retirement Research*, 2007.

Novogradac, Michael. All States to See LIHTC, Bond Caps Increases. *Novogradac Affordable Housing Resource Center*, 4 May 2018.

O'Regan, Katherine M. and Keren M. Horn. What Can We Learn about the Low-Income Housing Tax Credit Program by Looking at the Tenants? *Housing Policy Debate*, 23(3):597–613, 2013.

Papke, Leslie E. and James M. Poterba. Survey Evidence on Employer Match Rates and Employee Saving Behavior in 401(k) Plans. *Economics Letters*, 49(3):313–317, 1995.

Profit Sharing/401k Council of America. 401(k) and Profit Sharing Plan Eligibility Survey. `https://www.psca.org/uploads/pdf/research/2010/2010_Eligibility_Survey_Report.pdf`, 2010.

Roderick, Melissa, Vanessa Coca, and Jenny Nagaoka. Potholes on the Road to College: High School Effects in Shaping Urban Students' Participation in College Application, Four-year College Enrollment, and College Match. *Sociology of Education*, 84(3):178–211, 2011.

Sanbonmatsu, Lisa, Lawrence F. Katz, Jens Ludwig, Lisa A. Gennetian, Greg J. Duncan, Ronald C. Kessler, Emma K. Adam, Thomas McDade, and Stacy T. Lindau. Moving to Opportunity for Fair Housing Demonstration Program: Final Impacts Evaluation. *US Dept of Housing and Urban Development*, 2011.

Sokatch, Andrew. Peer Influences on the College-Going Decisions of Low Socioeconomic Status Urban Youth. *Education and Urban Society*, 39(1):128–146, 2006.

Thaler, Richard H., and Shlomo Benartzi. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Savings. *Journal of Political Economy*, 112(1):164–187, 2004.

The Furman Center for Real Estate and Urban Policy, New York University. What Can We Learn about the Low-Income Housing Tax Credit Program by Looking at the Tenants? `https://furmancenter.org/files/publications/LIHTC_Final_Policy_Brief_v2.pdf`, 2012.

The Internal Revenue Service. Instructions for Forms 1098-E and 1098-T. https://www.irs.gov/pub/irs-pdf/i1098et_18.pdf, 2018.

The Internal Revenue Service. Retirement Topics - Automatic Enrollment. https://www.irs.gov/retirement-plans/plan-participant-\employee/retirement-topics-automatic-enrollment, 2020.

The Joint Committee on Taxation. Estimates of Federal Tax Expenditures for Fiscal Years 2018-2022. *Prepared for the House Committee on Ways and Means and the Senate Committee on Finance*, JCX-81-18, 4 October 2018.

The United States Census Bureau. 2000 Census of Population and Housing: Population and Housing Unit Counts. https://www.census.gov/topics/research/guidance/planning-databases.2000.html, 2000.

The United States Department of Housing and Urban Development, Office of Policy Development and Research. Low Income Housing Tax Credits. https://www.huduser.gov/portal/datasets/lihtc.html, 24 May 2019.

The United States Government Accountability Office. Low-Income Housing Tax Credit: The Role of Syndicators. GAO-17-285R, 16 February 2017.