Essays in Financial Economics

by

Allison Cole

B.A. Boston University (2013)
M.A. Boston University (2013)
S.M. Massachusetts Institute of Technology (2020)

Submitted to the Department of Management
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN MANAGEMENT

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Authored by: Allison Cole
MIT Sloan School of Management
May 5, 2023

Certified by: Jonathan A. Parker
Professor of MIT Sloan School of Management, Thesis Supervisor

Accepted by: Eric So
Sloan Distinguished Professor of Financial Economics
Professor, Accounting and Finance
Faculty Chair, MIT Sloan PhD Program

Essays in Financial Economics

by

Allison Cole

Submitted to the Department of Management
on May 5, 2023, in partial fulfillment of the
requirements for the degree of
DOCTOR OF PHILOSOPHY IN MANAGEMENT

Abstract

In Chapter 1, joint work with Bledi Taski, we pose the question: how do workers value retirement benefits relative to wages and what impact do these benefits have on firm hiring? We find that dollars paid in employer contributions to 401(k) plans have nearly double the effect on a firm's recruiting success than dollars paid in wages. However, the effect is driven primarily by high-income and higher-age occupations. We use two novel instruments to identify the results: 1) IRS mandated non-discrimination testing of retirement plans and 2) corporate policies of national wage setting. We then develop and estimate an on-the-job search model which shows that the average worker requires only a 0.25 percentage point increase in employer contribution dollars to offset a 1% decrease in wages. Again, retirement valuations are positively correlated with salary. We confirm the channel in an online survey setting: participants are willing to give up total pay to get a higher employer match to get a non-matching employer-sponsored 401(k). The results imply that 80% of firms could improve their probability of a job offer being accepted by increasing 401(k) contributions.

Chapter 2, joint work with Jonathan Parker, Antoinette Schoar, and Duncan Simester, documents the share of investable wealth that middle-class U.S. investors hold in the stock market over their working lives. This share rises modestly early in life and falls significantly as people approach retirement. Prior to 2000, the average investor held less of their investable wealth in the stock market and did not adjust this share over their working life. These changes in portfolio allocation were accelerated by the Pension Protection Act (PPA) of 2006, which allowed employers to adopt target date funds (TDFs) as default options in retirement saving plans. Young retail investors who start at an employer shortly after it adopts TDFs have higher equity shares than those who start at that same employer shortly before the change in defaults. Older investors rebalance more to safe assets. We also study retirement contribution rates over the life-cycle and find that average retirement saving rates increase steadily over the working life. In contrast to what we find for investment in the stock market, contribution rates have been stable over time and across cohorts and were not increased by the PPA.

In Chapter 3, I use administrative data on very small businesses (median 5 employees) to measure the effects of the Paycheck Protection Program (PPP). Firms that applied for PPP increased employment by 7.5% relative to similar firms that did not apply. The positive effects on employment occur primarily in industries which were less affected by COVID-19: industries with more employees that are able to work remotely, those that have fewer hourly workers and essential businesses. Novel data on hiring shows that PPP worked as intended by preserving employment matches. My estimates imply a cost of approximately $270,000 per job-year at small firms.

Thesis Supervisor: Jonathan A. Parker
Title: Robert C. Merton (1970) Professor of Financial Economics

Thesis Errata Sheet

Author Allison Cole

Primary Dept. Management

Degree PhD Graduation date 06/2023
mm/yyyy

Thesis title

Essays in Financial Economics

Brief description of errata sheet

The changes adjust citations.

Number of pages 5 (11 maximum, including this page)

▶ **Author**: I request that the attached errata sheet be added to my thesis. I have attached a copy prepared as prescribed by the current *Specifications for Thesis Preparation*.

Signature of author Allison Cole Digitally signed by Allison Cole Date: 2024.03.01 15:00:48 -07'00' Date 3/1/2024

▶ **Thesis Supervisor** or **Dept. Chair**: I approve the attached errata sheet and recommend its addition to the student's thesis.

Signature Jonathan A. Parker Digitally signed by Jonathan A. Parker Date: 2024.03.02 15:45:54 -05'00' Date 3/2/2024

Name Jonathan A. Parker ☑ Thesis supervisor ☐ Dept. Chair

▶ **Vice Chancellor or their designee**:

I approve the attached errata sheet and direct the MIT Libraries to insert it into all electronic and/or print copies of the student's thesis held by the MIT Libraries.

Signature Ian A. Waitz Date 3/13/24

Name ~~Steven R. Lerman~~ Ian A. Waitz

1 Page 26, Line 14

Original text:

In order to understand how these valuations of retirement contributions affect where people choose to work, in the third part of the paper we develop and estimate an on-the-job search model. Similar to Burdett and Mortensen (1998), we assume workers make binary choices over two jobs that offer different wages and different idiosyncratic firm-worker specific match values.

Corrected Text:

In order to understand how these valuations of retirement contributions affect where people choose to work, in the third part of the paper we estimate an on-the-job search model. In the spirit of Burdett and Mortensen (1998), following the recent models of Sorkin (2018), Bonhomme and Jolivet (2009), and most closely Lehmann (2023), we assume workers make binary choices over two jobs that offer different wages and different idiosyncratic firm-worker specific match values.

2 Page 30, Line 23

Original text:

Also related is a series of papers using structural estimation of on-the-job search models to estimate compensating differentials. Burdett and Mortensen (1998) develops the originating on-the-job random search model; this model generates wage dispersion but does not address non-wage valuations. Others have since expanded upon the model to include worker valuation of non-wage benefits.

Corrected Text:

Also related is a series of papers using structural estimation of on-the-job search models to estimate compensating differentials. Burdett and Mortensen (1998) develops the originating on-the-job random search model; this model generates wage dispersion but does not address non-wage valuations. Sorkin (2018), Bonhomme and Jolivet (2009), Hall and Mueller (2018), Sullivan and To (2014), and others have since expanded upon the model to include worker valuation of non-wage benefits. We follow the estimation method of Becker (2011), Hall and Mueller (2018), and most closely, Lehmann (2023), which have developed tractable methods to estimate worker preference parameters via maximum likelihood by observing job transitions.

3 Page 65, Line 5

Original text:

Motivated by these facts, in this section, we develop a random on-the-job search model, similar to Burdett and Mortensen (1998) and Sorkin (2018), in which workers value retirement benefits (on both the intensive and the extensive margin.) and the other non-wage portions of compensation separately.

Corrected Text:

Motivated by these facts, in this section, we develop a model in the spirit of Burdett and Mortensen (1998) and Sorkin (2018) in which workers value wage and non-wage benefits separately, with a setup as in Lehmann (2023). We extend the model so that workers value different types of non-wage benefits separately and have taste for retirement benefits on both the intensive and the extensive margin.

4 Page 70, Line 2

Original text:

The main estimating equation from the model, (1.10), is a likelihood function and can be estimated by standard maximum likelihood techniques.

Corrected Text:

The main estimating equation from the model, (1.10), is a likelihood function and can be estimated by standard maximum likelihood techniques, as in Lehmann (2023), Hall and Mueller (2018), and Becker (2011).

5 Page 79, Line 20

Original text:

As a validation exercise of the model, we estimate firm-level amenity valuations, net of the estimated retirement valuations and show that the amenity valuations do not change around NDT failure for the subset of firms who failed in our estimation sample.

Corrected Text:

As a validation exercise of the model, we replicate the procedure in Lehmann (2023) to estimate firm-level amenity valuations. As an extension, we then net out the estimated retirement valuations to measure the residual, non-wage, non-retirement amenity valuations. We then show that these residual amenity valuations do not change around NDT failure for the subset of firms who failed in our estimation sample.

References

Becker, Dan. 2011. "Non-Wage Job Characteristics and the Case of the Missing Margin." *SSRN Electronic Journal*.

Bonhomme, Stéphane, and Grégory Jolivet. 2009. "The pervasive absence of compensating differentials." *Journal of Applied Econometrics*, 24(5): 763–795.

Burdett, Kenneth, and Dale Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review*, 39(2): 257–73.

Hall, Robert E., and Andreas I. Mueller. 2018. "Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values." *Journal of Political Economy*, 126(4): 1594–1637.

Lehmann, Tobias. 2023. "Non-Wage Job Values and Implications for Inequality." Working Paper.

Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *The Quarterly Journal of Economics*, 133(3): 1331–1393.

Sullivan, Paul, and Ted To. 2014. "Search and Nonwage Job Characteristics." *The Journal of Human Resources*, 49(2): 472–507.

Acknowledgments

I first want to express deep gratitude for the invaluable guidance and support of my dissertation committee: Taha Choukhmane, Jonathan Parker, Antoinette Schoar and Anna Stansbury. They have been generous with their time and advice throughout the course of completing my thesis. Their insights and guidance pushed my work to be more precise, more convincing, and more ambitious.

My heartfelt thanks especially go to my dissertation committee chair, Jonathan Parker. As both my thesis chair and my advisor since day one of the PhD, he has truly guided me through this six-year journey. Jonathan helped me to navigate through numerous research ideas and pivots with patience and support. He also provided a sympathetic ear to many challenges throughout the years, research-related and otherwise.

I am also grateful to be part of the vibrant community at MIT. Christopher Palmer, Larry Schmidt, David Thesmar, and Emil Verner gave helpful feedback over the years. I want to thank Adrien Verdelhan for his dedication to the Finance PhD program and efforts in recruiting so many wonderful candidates to be my colleagues. I benefited tremendously from conversations with my classmates, Pierre Jaffard, Yura Olshanskiy and Albert Shin. Mazi Kazemi, Olivia Kim, and Pari Sastry shared their recent experience to help guide me through the job market. Maya Bidanda has been an invaluable friend.

For helpful comments, on the various parts of this dissertation, I also thank Jesse Bricker, Sylvain Catherine, Isil Erel, Joe Hazell, Giuseppe Moscarini, Jim Poterba, Josh Rauh, John Sabelhaus, Eva Xu, and participants at various seminars.

Crucially, I thank my friends and family. I consider myself uniquely lucky to have a wonderful group of friends outside of academic life who always supported me, even when my concerns about “identification", “exclusion restrictions", or the “incidental parameters problem" seemed far-fetched. Thank you Bridget, Brittany, Meredith, Rachel and Tressa. I also had the support of my four-legged friends, Egg, Fig and Flash who brought me so much joy over the years. Thank you to Spencer for believing in me and reminding me not to take myself too seriously. My sister, Katie, has always been my biggest cheerleader. Lastly, I never could’ve done this without the unconditional love and support of my parents, Francis and Tracy. I dedicate this dissertation to them for teaching me that I could accomplish anything I set my mind to.

Dislcaimer: All authors of Chapter 2 except Simester served as unpaid consultants for the financial services company that provided data for that chapter. Simester served as a paid consultant of the company during this time.

Contents

List of Figures

List of Tables

Chapter 1

Worker Valuation of Retirement Benefits

With Bledi Taska

1.1 Introduction

From 2010-2019, employers in the United States spent an aggregate 1.3 trillion dollars on contributions into defined contribution (DC) retirement plans such as 401(k) and 403(b) accounts. Defined contribution plans are important determinants of household disposable income and consumption over the life-cycle, and the wealth accumulated in these accounts is an increasingly important source of household funds for retirement. According to the 2016 Survey of Consumer Finance, the median household holds 76 percent of their total (non-real estate) wealth in DC retirement accounts. Given the importance of DC plans for households savings, wealth accumulation, and firm costs, a large body of research has focused on how employees save and invest in DC retirement plans.[1]

There is less understanding of how employees value the plan as a job feature and how DC retirement plans affect equilibrium labor market behavior. Several studies have shown that many 401(k) participants do not take full advantage of employer matches (Mitchell et al. (2007), Engelhardt and Kumar (2007), Choukhmane et al. (2022)). On the other hand, there is conventional wisdom amongst human resource professionals that retirement benefits are an essential tool to attract talent.[2] Moreover, policymakers and regulators value these plans substantially, as revealed by the increasingly large tax expenditure that these plans represent.

In this paper we measure both how workers value a dollar of retirement contributions, relative to a dollar of wages, and the impact of retirement contributions on labor market flows. First, we use a revealed-preference approach and instruments for exogenous wage and benefit changes to measure how wages and retirement benefits affect a firm's recruiting success. We use data on worker transitions across firms to show that – for the average job – one dollar increase in employer contributions to DC retirement accounts has nearly twice the effect on recruiting success as a one dollar increase in wages. Second, we design and conduct an online survey experiment to show that, consistent with the results in part one, most workers are willing to take a lower-paying job when it offers either higher employer contributions to its DC plan or when it offers a DC plan at all. Third, we develop and estimate on on-the-job search model in which we confirm that – at the equilibrium – the majority of workers are willing to give up some total compensation to get a larger share of compensation as a retirement benefit. We also show in a counterfactual exercise that 80 percent of firms in our sample could have improved their

[1]For these various issues, see Gomes et al. (2018), Autor et al. (2020), Sialm et al. (2015), Parker et al. (2022), Carroll (2000), Bernheim et al. (2015), Choukhmane (2019), Choi et al. (2004).

[2]See Wasick (2016), Weber (2022), Whalen and Tergensen (2022), Miller-Merrell (2013)

recruiting outcomes if they increased retirement contributions, consistent both with the steady increase in benefits over the period studied and also with regulatory constraints on benefit plan equality that make changes to retirement plans costly. In the following paragraphs, we lay out our data, methods, and findings in more detail.

In the first part of the paper, we measure the effect of wages and retirement benefits on firm recruiting success. To do so, we construct a novel data set which merges 1) 30 million online job postings from Lightcast (formerly Burning Glass Technologies), 2) 83 million online resumes, also from Lightcast, and 3) detailed financials on retirement plans from regulatory filing Form 5500 for every U.S. company with a retirement plan. Together, these data link information typically only available through proprietary or administrative access: information on a firm's hiring success (inferred from the resumes, which enable us to see whether a posted vacancy is filled), a job's posted wage (from the online job postings), and the retirement benefits available to workers in each job (from the regulatory filings). The final estimation sample represents over half a million worker transitions to 24,000 firms and 150,000 occupation by CBSA groups from 2010-2019.

Ideally, one could regress job choices on a measure of wages, retirement, and other benefits and estimate directly worker preferences for each type of compensation. However, both wage setting and benefit policy suffer from an endogeneity concern: the way that firms set wages and retirement contributions is likely to be correlated with unobservable firm characteristics that may also be correlated with how attractive the firm is to job-seeking workers. [3] We address this concern using two novel instruments.

First, we use IRS mandated non-discrimination testing (NDT) as an instrument for firm retirement policy. Each year, firms that offer DC retirement plans must undergo this

[3] See for example, Sockin (2021), Sorkin (2018), Maestas et al. (2018).

test to show that their plan does not disproportionately favor "highly-compensated employees." We confirm that, following failure, plans increase their effective contribution rate by approximately half a percentage point relative to non-failers who are similar on observables.The identifying assumption is that, conditional on observables, NDT failure is orthogonal to changes in the firm's recruiting ability over time. Consistent with this assumption, we show that firms that fail NDT do not adjust their wages, healthcare benefits, or the composition of their targeted hires in job postings following failure.

The second instrument is the corporate policy of national wage setting. This policy induces exogenous variation in wages at the occupation and geographical level. First documented by Hazell et al. (2022), approximately 30-40 percent of firms follow a policy of setting occupational wages nationally. That is, rather than tailoring wages to local labor market conditions, they offer the same wage to all workers within an occupation across the country. Firms that follow this policy have 1-3 percent higher wages on average.[4] For the firms that follow this policy, we construct an instrument that measures how much an occupation by CBSA specific wage is pushed up, relative to the CBSA and occupation average, due the firm being a national wage setter. The identifying assumption is that national wage setters' other job features that attract workers are not affected by year to year deviations in the wage offered due to national wage setting. Consistent with this assumption, national wage setting is uncorrelated with most observable firm characteristics. Moreover, healthcare and retirement benefits are not significantly correlated with national wage setting.

The instrumented results show that within a firm and occupation by CBSA job market, on average, an extra dollar of retirement contributions has twice the effect on the

[4]The pattern of higher wages holds at the firm level, controlling for industry fixed effects, and at the occupation by CBSA level, controlling for industry by occupation by CBSA fixed effects.

likelihood of a firm successfully filling a position than an extra dollar of wage. Specifically, We find that a one percent increase in wages increases recruiting success by 1.4 percent while an equivalent dollar increase in employer retirement contributions increases recruiting success by 2.7 percent. This large value of benefits is driven primarily by workers in high-income and higher-age occupations. An equivalent dollar increase in retirement contributions has nearly three times the effect on the recruiting success of a firm when looking only at the higher income and older occupations. For lower age occupations, retirement dollars have one-third the effect of wages. Occupations which are female-dominated have a similar sensitivity to retirement relative to wages as those that are male-dominated.

Following this first portion of the paper that infers workers' valuation of retirement benefits by revealed preference, in the second part of the paper, we use on-line survey experiments to directly measure workers' reported choices and infer from these their valuation of retirement benefits. We find that the stated preferences imply valuations that are consistent with those implied by revealed preferences. We recruit 1,600 online survey participants and measure their willingness to pay for varying levels of retirement benefits. To measure willingness to pay, we show participants side-by-side job offers which are identical other than the level of wages and retirement benefits offered.[5] The discrete choice framework allows us to flexibly estimate the distribution of willingness to pay measures using maximum likelihood estimation. We also design the survey to show choices that are strictly dominated, thus we can measure and correct for inattention. We test six retirement related conditions in total, which measure the willingness to pay both for the availability of a 401(k) plan at the job (the extensive margin) and for different levels of employer matching contributions (the intensive margin).[6]

[5]The experimental design is similar to Mas and Pallais (2017).

[6]We also test one condition that elicits willingness to pay for remote working capability, in order to

The survey results show that the majority of workers, 50 to 80 percent, depending on the condition, will choose a job that offers better retirement benefits, even when that job pays lower total compensation, *inclusive of the match* and *net of tax differences*. In the conditions that test the intensive margin, or how large of a dollar match the company offered (versus a plan with no match), participants were willing to give up approximately half a percent in total compensation for each percentage point increase in the employer match. The implied willingness to pay is 1.5 percent of total compensation to get a 401(k) with a 3 percent match. In the condition that tested the extensive margin, or whether or not the job offered a 401(k) plan at all, participants were willing to pay about 3.4 percent of total compensation to get a plan, even when it offered no match. These results are consistent with the finding of our first main analysis of actual job flows which implies that most job-seekers value retirement dollars roughly two times as much as they value wage dollars.

In order to understand how these valuations of retirement contributions affect where people choose to work, in the third part of the paper we develop and estimate an on-the-job search model. Similar to Burdett and Mortensen (1998), we assume workers make binary choices over two jobs that offer different wages and different idiosyncratic firm-worker specific match values. As in Sorkin (2018), Bonhomme and Jolivet (2009) and Hall and Mueller (2018) compensation includes non-wage benefits, and we further allow non-wage benefits to be valued differently from wages on both the intensive and extensive margins. Workers have indirect utility over wages, retirement, healthcare, other amenities, and their idiosyncratic match values. We use the search model to estimate retirement valuations (for a subset of workers) from their revealed preference and then show how retirement benefits affect where people choose to work, relative to wages and other bene-

---

compare with other estimates of worker valuation of flexibility, such as Mas and Pallais (2017) and Wiswall and Zafar (2018).

fits. The model is identified from the net flows of workers between firms and measurable wage, healthcare and retirement differentials between firm pairs.

The average worker in the estimation sample is willing to give up one percent of wage (550 dollars on average) to get just a 0.25 percentage point increase employer in retirement contributions (110 dollars on average). Strikingly, about 75 percent of the distribution of workers is willing to give up some of total compensation to get a higher retirement benefit. The remaining 25 percent, who also tend to be lower income, need larger compensating differentials; total compensation must increase if the wage decreases for those workers. On the extensive margin, the average worker is willing to give up about 2 percent of wages to get a 401(k) plan, which is close in magnitude to the survey estimates. This estimate is also increasing in income, though 90% of workers place a positive value on the availability of the plan.

The model differs from the revealed preference results in two distinct ways, yet produces remarkably similar results to the revealed preference estimation. First, the model is estimated on a specific subset of job-switchers whose compensating differentials are primarily driven by firm-specific characteristics. Specifically, we estimate the model only on transitions within occupation, CBSA, and industry. This eliminates drivers of job change due to career or location changes. Second, the model estimation allows us to directly estimate the weight workers place on retirement and other non-retirement, non-healthcare benefits *separately*. While the instrumental variables estimates are at the firm-level, the model is estimated from observable compensation differences *between firm pairs*. Combining information on hundreds of firm pairs within an industry by occupation group, the average weight placed on each part of compensation as well as the residual or "amenity" difference between firms for workers in that group are each separately identified. We also validate the model by using the NDT instrument. We show that firms that fail NDT change their retirement following the failure, but that the average difference in amenity valuations in the industry by occupation group does not change. Hence, the estimated amenities term is not picking up changes in retirement contributions.

The results have implications for both firm compensation setting policy and regulation of DC retirement plans. First, our results indicate that 80 percent of firms in the estimation sample could have improved their average recruiting success (across all occupations) if they shifted compensation from wages to employer contributions to DC plans. We obtain these results by conducting a counterfactual exercise in which one firm a time increases either wages or retirement contribution dollars by one percent and then calculate their new unconditional probability of having an offer accepted based on the corresponding estimated valuation weights for each type of worker they are trying to hire. This exercise also shows that all firms that do not offer retirement plans (30 percent in-sample) could improve their recruiting success if they offered a 401(k) plan, holding everything else constant. Ignoring for now the regulatory and set-up costs associated with increasing retirement contributions, dollar for dollar, increases in retirement contributions have two to three times the effect on recruiting success as wages, depending on the estimate used. In other words, it would take a two to three percent increase in wage dollars to induce that same change in recruiting success as a one percent change in employer contribution dollars.[7] However, the regulatory constraint of non-discrimination testing makes increasing retirement contribution prohibitively costly for most firms.

Changes in firm retirement policy would have disparate effect on workers across the income distribution. Because higher-income workers place a higher value on retire-

[7]Due to the wide distribution of retirement valuations, this varies based on the type of worker a firm is targeting. If a firm in particular wants to hire in an occupation with lower retirement valuations, it would be better off increasing wages. But given the distribution of workers at most firms, increasing retirement contributions has a larger effect on average.

ment contributions, more generous retirement offerings provide larger gains to higher income workers while doing relatively little for workers on the lower end of the income distribution. Due to the structure of NDT, retirement benefits cannot vary across workers; firms must offer the same policy of contributions to everyone. Hence NDT places a binding constraint on any firm that employs workers with differing valuations of retirement. Firms are unable to tailor these benefits to worker preferences and must cater to the majority, which, holding the income distribution constant, shapes compensation in a way that tends to favor higher income workers. In ongoing work, we use the model and framework to assess the welfare implications of non-discrimination testing on worker valuations.

Contribution: This paper has three main contributions. First, we document new empirical facts about worker valuation of defined contribution retirement plans. Relatively little academic literature has studied how DC retirement plans contribute to worker labor market decisions. We show that 1) workers value retirement plans as a job feature, despite the existing evidence that retirement plans are underused by savers 2) DC retirement plans and contributions are a significant driver of compensating differentials between firms. Our findings suggest that workers value retirement plans above and beyond just the dollars paid. Qualitative survey responses suggest that 401(k)s and matching provide value as both a signal of firm quality and a commitment device, but more research is needed to fully understand why workers have strong preferences for the plan at the outset of a job. We also introduce a novel instrument in non-discrimination testing, which shows that firms change their plan design due to regulatory constraints. The variation induced by this testing has many potential applications across both corporate and household finance research agendas. Second, we expand on the class of search models in which workers have values over job features by adding taste for retirement benefits on both the extensive and intensive margin. Moreover, we construct a novel data set that links employers and employees in order to estimate the model. Third, we document how DC plan regulations favor firm compensation policy that disproportionately benefits higher income workers. Few papers have examined why firms structure DC plans the way they do; this paper shows that regulatory constraints have a significant effect on plan design. The following paragraphs discuss in more detail how the paper relates and adds to the existing literature.

Relation to Literature: A large literature has studied the importance of compensating differentials in labor markets. Originating with Rosen (1986) household finance and labor economics have long been interested in understanding what job features make up for differences in wages. Miller (2004) and Sheiner (1999) showed that healthcare benefits are typically passed of into lower wages and valued by workers by as much 10 percent of wages. Simon and Kaestner (2003) shows that offering pensions does not crowd out wages. More recently, several papers have shown that workers place a high value on non-wage, non-retirement and non-health benefits, such as remote-work, working conditions, or job flexibility (Maestas et al. (2018), Mas and Pallais (2017), Wiswall and Zafar (2018)). There is mixed evidence as to whether or not firms with higher wages offer better (Sockin (2021), Becker (2011)) or worse (Lamadon et al. (2022)) amenities. Most studies conclude that non-wage job features make up a large part of job valuation, explaining as much as half of the variance of job valuations (Taber and Vejlin (2020), Sorkin (2018)). Moreover, non-wage characteristics are thought to contribute more to inequality both within and between firms (Kristal et al. (2020), Azar et al. (2022), Ouimet and Tate (2022)).

Also related is a series of papers using structural estimation of on-the-job search models to estimate compensating differentials. Burdett and Mortensen (1998) develops the originating on-the-job random search model; this model generates wage dispersion but does not address non-wage valuations. Others have since expanded upon the model to include worker valuation of non-wage benefits. Bonhomme and Jolivet (2009) shows that workers have strong preferences for non-wage job features, such as job security. Sorkin (2018) find that non-wage features account for over half the firm component of the variance of earnings. Several other papers have shown that that non-wage compensation significantly contributes to worker valuation differences across jobs (Becker (2011), Sullivan and To (2014), Hall and Mueller (2018)).

We contribute to this strand of the literature in two ways. First, we build a data set that connects wages, retirement, and healthcare benefits. Most previous studies lack detailed data on benefits, especially retirement, and thus are only able to measure the non-wage portion of valuations in aggregate. Our merged data set of posted wages and regulatory filing with benefit financials allows us to separately identify the direct effect of retirement alone on worker valuations. Second, we focus specifically on retirement. Studies that have used data with information on benefits have focused primarily on healthcare or general amenities, not retirement. We build on this class of models by adding explicitly taste for retirement on the extensive and the intensive margin.

A smaller literature has studied the effect of retirement plans on worker mobility, primarily focusing on defined benefit, or DB (pension), plans, which differ significantly from DC plans. As most employers shifted from DB to DC plans in the 2000s and 2010s, there was concern over the loss of DBs leading to higher turnover (Johnson (2013)), because DBs typically required much longer vesting periods than DCs to receive full benefits. This fear was largely shown to be unfounded (Goldhaber et al. (2017), Gustman and Steinmeier (2002), Gustman et al. (1994), Goda et al. (2017)). A few studies have looked more generally at how benefits correlate with turnover and have found a positive association (Johnson (2013), Bennett et al. (1993), Lee et al. (2006)).

We add to this literature by updating the findings on mobility for the modern retirement landscape. Currently, 60 percent of workers have access to a DC while only 25 percent have access to a DB, thus making the findings about DBs effect on mobility less relevant for the modern worker. While we focus on recruitment, rather than turnover, our findings shows that variation between DCs and whether or not a DC is offered are important drivers of job valuation.

Very few papers in finance and economics have analyzed how or why employers design retirement plans. Two exceptions are Bubb et al. (2015) and Bubb and Warren (2020). These papers show that, theoretically, employers design plans to take advantage of the myopia of participants by offering generous matching that they know won't be taken advantage of. This paper complements their findings and together, the findings explain the puzzle of why workers value the plan but often do not use it. Workers value DC plans when choosing jobs perhaps because they plan to use them, but just as employers anticipate, many workers don't end up using them, thus saving the firm costs. Arnoud et al. (2021) documents the current landscape of plan design in the U.S., but it is beyond the scope of that paper to analyze the drivers of plan design. Fadlon et al. (2016) uses a tax reform in Denmark to show that employers adjust their contributions to be consistent with worker preferences.

This paper offers new insights into employer's motivations and incentives for DC plan design. While we cannot measure the mechanism directly, we offer empirical evidence that supports two pieces of motivation. First, at least a subset of (higher income) workers highly values DC retirement plans, so offering such benefits can help firms to more effectively recruit those workers. Second, non-discrimination testing limits a firm's ability to cater plans to individual worker preferences. Thus firms must choose a plan that they think will appeal to either the largest cross-section of workers or the workers they most want to attract. Only one other paper, to our knowledge, has examined the effects of NDT on firm compensation. Ouimet and Tate (2022) show that firms with more high-wage workers also tend to offer higher benefits, which has spill-over effect on lower-income employees. This is consistent with our finding that higher-income workers place higher-value on retirement benefits. This paper shows directly that high-wage workers value retirement benefits more than low-wage workers, which complements the findings in Ouimet and Tate (2022).

The paper proceeds as follows. In section 1.2, we outline the research design of our instrumental variables approach and the data used for this approach. Next, in Section 1.3, we describe the instrumental variables results. In Section 1.4 we describe the survey design and the sample of participants. In Section 1.5, we describe the survey results. Section 1.6 describes the on-the-job search model. Section 1.7 describes the implication of the results for firm policy and worker valuations. Section 1.8 concludes.

1.2 Instrumental Variables Approach

In this section, we detail our first method for estimating the impact of retirement benefits on recruiting success and worker's valuations of retirement benefits. In Section 1.2.1, we detail the empirical specification of our instrumental variables approach. In sections 1.2.1 and 1.2.1, we describe the two instruments that we use to induce exogenous variation in retirement and wage setting policy: non-discrimination testing and the national wage setting, respectively. In section 1.2.2, we describe our data and its sample representativeness.

1.2.1 Method

It is well documented that workers place value on many different parts of compensation, other than just wages (Sorkin (2018), Mas and Pallais (2017), Wiswall and Zafar (2018), Bonhomme and Jolivet (2009), Taber and Vejlin (2020)). Imagine a simple indirect utility function from working at firm j for worker i:

$$V_{i,j} = \alpha w_{i,j,} + \beta r_j + \gamma h_j + \delta a_j + \epsilon_{i,j} \tag{1.1}$$

Worker i values the wage, $w_{i,j,}$, the retirement, r_j, the healthcare h_j, other benefits or amenities, a_j and there is an additional firm worker specific match component that can affect valuation. Note that the wage can be worker specific, but benefits cannot.[8] Normalize $\alpha = 1$ so that all other terms in (1.12) are in wage-equivalent units.

The primary objective of this paper is to measure β, the worker's sensitivity to the retirement contributions, relative to his sensitivity to wages. There are several empirical challenges to estimating this. First, one does not observe directly worker valuations of benefits and wages. Second, data on wages, retirement, healthcare, and other amenities is not readily available and is difficult to collect. Third, the way firms set their compensation and benefit policies is likely to be correlated with each other *and* correlated with unobservable characteristics that also increase worker valuations. In other words, wages and retirement benefits are endogenous.

To deal with the first issue, we use a firm-level measure of recruiting success to infer worker valuations by revealed preference. Although we cannot measure worker valuations directly, we can observe where people choose to work. Comparing the firms chosen

[8] This restriction mimics the equality regulation on benefit plans, such as non-discrimination testing.

to other choices with different compensation bundles reveals which components of compensation workers place higher value on. We construct the recruiting success measure by comparing job postings to the resume data and thus see when and if a posted job is filled.

To deal with the second issue, we have built a data set that contains detailed information on the first three terms in equation (1.12): wages retirement, and healthcare. We describe the data in more detail in Section 1.2.2. This dataset allows us to measure the three largest dollar parts of compensation, but not amenities. Lack of data on amenities contributes to the third challenge described above, an issue we address in the following paragraphs.

If one has perfect data, including information on firm amenities, the third challenge described above is mitigated. For example, firms that offer better healthcare may also offer better parental leave, both of which add value for workers. If one could observe and control for the parental leave, then one could estimate the worker's true valuation of healthcare alone. However most firm-level data does not indicate how good a firm's parental leave policy is. The parental leave example is one that deals with an amenity that is, at least in theory, measurable. However, there are other amenities that would not be measurable even if perfect data did exist. For example, a firm that offers better retirement benefits may also have more financially savvy employees, from which there is a positive spillover to other employees who work there. This type of amenity is unmeasurable to the econometrician. In sum, the fact that our data does not have information on non-retirement and non-wage amenities, and that some amenities are unmeasurable, means that the endogeneity issues remains.

We address this in two ways. First, our specification includes a firm by occupation by CBSA fixed effect. Thus, we compare firms to themselves in prior years, within the same CBSA and occupation (or market) as their wages or benefits in that market change. This helps reduce bias in the estimates due to endogeneity so long as we believe that changes in within firm amenities over time are smaller than differences in between firm amenities.[9] The empirical specification is thus:

$$HireSuccess_{j,t,l} = \alpha w_{j,t,l} + \beta r_{j,t} + \gamma h_{j,t} + X_{j,t} + \delta_{j,l} + \delta_t + \epsilon_{j,t,l} \tag{1.2}$$

Where HireSuccess is a dummy variable equal to one if the firm successfully filled a job in that market, measured at the firm, occupation, CBSA and year level. $w_{j,t,l}$ is the posted wage at firm j in year t for occupation and CBSA (market) l. $r_{j,t}$ is the employer contribution rate offered by firm j in year t; note that this does not vary by occupation as this policy must be constant within firm. Similarly $h_{j,t}$ is the healthcare benefits offered by firm j in year t and does not vary across occupations. $\delta_{j,l}$ is a firm by market fixed effect and δ_t is a year fixed effect. $X_{j,t}$ is an additional control for time varying firm characteristics not captured by the firm fixed effect. The specification thus measures how much hiring success changes within a firm and occupation by CBSA market as wages and retirement and healthcare change, controlling for yearly trends in the hire success rate and other observed firm characteristics.

The fixed effects are however, not sufficient to fully address the endogeneity concern. It is possible that firms change other unobservable amenities over time. To further address the issue, we use two instrumental variables to estimate (1.2). The first, nondiscrimination testing, has a direct effect on retirement contributions. The second, national wage setting, has a direct affect on wages. we describe the instruments in detail in

[9]This assumption is supported by other work, such as Ouimet and Tate (2022) and Kristal et al. (2020).

the following subsections.

Non-discrimination Testing of Defined Contribution Plans

Each year, about 60 percent of firms with retirement plans must undergo IRS mandated non-discrimination testing (NDT).[10] The purpose of the test is for employers to show that their plan does not disproportionately favor "highly compensated employees" or HCEs. As of 2022, these are employees who make over 135,000 dollars per year.[11] There are various steps to the test, but the main objective is to show that HCEs do not have a significantly higher contribution rate than non-HCEs, inclusive of the employer match.[12]

When a plan fails, there are two options for correction. First, they can give more contributions to their non-HCES to raise that group's effective contribution rate. Second, they can take contributions back from HCEs and distribute them as income, which is now taxable. Either method presents costs for the firms. The first method is costly in dollar terms: the firm must make payouts to some or all of their employees who make less than $130,000. The second method, while not financially costly, presents a significant administrative burden and likely disgruntles HCEs who now have more taxable income than expected. If firms do not correct the failure within 3 months of the end of the filing year, they must pay 10% excise tax on the corrective distribution amount. If they don't

[10]Firms that choose a safe harbour contribution schedule are exempt from testing. Arnoud et al. (2021) estimate that about 40 percent of all firms choose safe harbor plans. The three available safe harbor provisions are: 1) Non elective safe harbor: the employer contributes 3% of salary to all employees which is immediately vested, regardless of how much the employee contributes to the plan; 2) Basic safe harbor match: the employer matches 100% of the first 3% of the employee's contribution and 50% of the next 2%; 3) Enhanced safe harbor match: the employer matches 100% of the first 4% of each employee's contribution.

[11]See https://www.irs.gov/retirement-plans/plan-participant-employee/definitions for details on the definition of an HCE.

[12]See Appendix Table A.1 for a concrete example of how NDT is implemented.

correct it within a year, they must enter an IRS corrective program and are at risk of losing their qualified status as a plan.

Despite the cost, we find that failure is not uncommon. Based on corrective distributions paid, which is observable in the regulatory filings we use for our retirement plan financial data, roughly 5% of all firms (or 8.5% of the firms that must test) fail each year (Figure 1-1). About 10-12% of firms pay some corrective distributions each year.[13] This can include small penalties for other plan mistakes, so we assume that a firm failed NDT testing only if its corrective distributions per person are in the top 10% of the distribution that year. This is likely a conservative definition of NDT failure, as the majority of corrective distributions paid are due to NDT failure. Figure 1-1b shows that about 10% of firms in our main estimation sample fail NDT each year.

Due to the cost and administrative hassle of failure, it is natural to think that most firms that fail want to avoid failing again. Indeed, we observe that less than 10% of firms that fail fail more than once. While we do not directly observe if a plan chooses a safe harbor provision that would exempt them from future testing, our conversations with benefit administrators indicate that the most common policy change after failure is to either elect into safe harbor or introduce auto-enrollment (if the firm did not already use auto-enrollment). More often that not, this results in the firm having a more generous contribution overall. All three safe harbors have an effective contribution rate of 3-4% (see footnote 10), which is typically an improvement over what non-safe harbor plans offer. We drop any failing firms that had an effective contribution rate of higher than 4% prior to the NDT failure in order to ensure that the monotonicity condition of the

[13] Appendix Figure A.1 shows that while large firm ($>$ 100 employees) are more likely to pay some corrective distributions, they are not more likely to fail NDT. The median dollar amount in corrective distributions paid, conditional on paying some is around $400 per person in our sample period. Larger firms pay higher dollar amounts per person.

instrumental variable is met (Imbens and Angrist, 1994). Without this restriction, it could be the case that higher income employees who were previously maxing out on a high employer match would receive a lower match after the failure if the firm switched to a safe harbor. This restriction applies only to about 10% of the firms that failed.

Indeed, we find that retirement plans in our main estimation sample that fail NDT in a given year increase their ratio of contributions by 2.5% , and their contribution rate by approximately .5% in the three years immediately following the failure. Figure 1-2 show the parallel trends comparing failers to non-failers, controlling for year by industry fixed effects, log number of employees, and log dollars of assets in the retirement plans. Table I shows the corresponding regressions. The specification compares those that failed to those that did not fail, starting three years prior to failure and ending three years after failure.[14] For the control group, the comparative year zero is taken as the median year the firm appears in the sample. Robust standard errors are clustered at the firm level. Note that we do not find a significant effect of NDT failure on autoenrollment for firms in our sample, indicating that these firms do not typically use autoenrollment as a way to remediate NDT failure.

Over the same time period, NDT-failers do not significantly change wages. The event study plot is shown in Figure 1-3. They also do not significantly change how many jobs they post, their experience requirements for a job, or their spending on healthcare (Table II). We do observe that NDT-failers slightly increase their number of new hires after failure; this is consistent with the finding that workers find better retirement benefits to be an attractive feature.

[14]In unreported robustness checks, we find the results are similar using two alternate controls groups. First, we limit the control group to be only firms that do not fail, but have a high probability of failing based on a predictive regression of failing on firm characteristics. Second, we limit the control group to firms that will fail in the future, but compare them to failing firms in the years before the control group actually failed.

A few notable differences between firms that fail and those that do not are present. In particular, firms that fail have ex-ante lower contribution rates. This is consistent with the fact that they ultimately fail the NDT. Moreover, these firms tend to have higher salaries, more job postings, and more new hires. Controlling for industry and size reduces these differences, but does not eliminate them entirely. Hence, there is reason to believe that NDT failing firms are somehow different from firms that don't fail, in a way that might by correlated with recruiting outcomes. This is further motivation for using a firm fixed effect in our estimating equation. Because these firms appear different on several dimensions, comparing them only to themselves in prior years is likely to reduce the endogeneity concern that unobservable firm characteristics are correlated with the firms' attractiveness to workers. The identifying variation from the NDT instrument represents *time series variation* within firm when they switch from non-failing to failing. Thus, the instrumental variable estimates capture the firm's change in recruiting success in a specific occupation by CBSA market that is due only to the retirement plan changing at that firm.

The exclusion restriction is that, conditional on observables, NDT failure is orthogonal to the firm's recruiting ability. With a firm and occupation by CBSA fixed effect, this means that the NDT failure does not affect other firm characteristics (like amenities) that may also contribute to recruiting outcomes. While we don't directly observe other amenities, we do observe that healthcare and wages do not change around NDT failure, implying that the firms at least do not adjust on those margins. In addition, we find that hiring slightly increases following NDT failure at failing firms, thus there is no indication that these firms try to decrease hiring after failing the test. Moreover NDT is a relatively unknown institutional procedure that potential employees are not likely to know about and it should only affect job choice by how it changes retirement benefits.

National Wage Setting

The second instrument we use to induce exogenous variation in wages is corporate policy of national wage setting. First documented by Hazell et al. (2022), approximately 30-40% of occupation have their wages set nationally by the firm, rather than tailoring wages to local labor market conditions. For example, if a firm employs one accountant in New York City and one in Santa Fe, it pays the two employees the same salary, despite the differing labor market conditions between the two cities. Firms that follows this policy sometimes do so for only select occupations or sometimes they do so for the majority of their workforce. Figure 1-4 shows that in our estimation sample, about 20-25 percent of firms predominately set wages nationally, meaning they do so for at least 75% of their occupations. Around 15% of occupations have the wage set nationally.[15]

National wage setters also pay a wage premium: on average, nationally identical jobs pay 1-3% percent more than other comparable jobs within their markets. Table III shows that this wage premium holds both at the firm level, controlling for industry by year fixed effects, and at the market level, controlling for CBSA by occupation by industry and year fixed effects.

We construct an instrument in which national wage setting is interacted with the difference between an occupation and CBSA specific wage and the predicted occupation and CBSA specific wage for non-national wage setters in that CBSA. The instrument is:

$$Instrument = \begin{cases} [ln(\widehat{Sal_{j,t,l}}) - ln(Sal_{t,l})] & \text{If Wage set Nationally} \\ 0 \quad \text{If Wage not Set Nationally} \end{cases}$$

[15]Note that the incidence of national wage setting declines significantly in 2018 and 2019 as Lightcast significantly expanded its coverage of postings,

where $Sal_{j,t,l}$ is the salary for firm j for occupation and CBSA l in year t, $\widehat{Sal_{t,l}}$ is the predicted wage for that same occupation and CBSA in year t, estimated only for non-national wage setters.[16] The instrument thus measures a firm's difference from the predicted value for non-national wage setters in that same occupation and geography. It is a measure of how much the wage deviates from the predicted value due to national wage setting. Appendix Figure A.2 shows the median value of the instrument across various geographical and firm characteristics (including only occupations for which the wage is set nationally). Wages are pushed up the most by national wage setting in lower-population and lower cost of living areas and at at larger firms.

National wage setters differ from other firms on several dimension. First, by definition, they must have multiple establishments and thus tend to be larger firms. However, when compared to other multi-establishment firms, they are actually slightly smaller by employment size. Correspondingly, they also have fewer job postings and hire fewer new employees on average when compared to other multi-establishment firms. Table IV documents these results.

Most importantly for our setting, however, national wage setters do not differ significantly on measures of turnover, retirement contributions, or healthcare benefits (see Table IV). This indicates that nationally wage setting firms do not offer substantially different benefits or amenities than other firms. Hazell et al. (2022) provide further evidence that national wage setting firms appear similar to non-national wage setting firms, and that the decision to set wages nationally is typically related to organizational structure and concentration, but not benefits.

The identifying variation of this instrument comes from time series variation within

[16]We estimate the predicted wage by regressing the average posted salary on CBSA, 6-digit SOC code occupation, and year fixed effects, including only jobs for which the wage is not set nationally.

firm and market. The instrument measures how much a firm's wage is being pushed up or down in a given year relative to the local average for that occupation and CBSA due to being a national wage setter. A firm may increase wages in some CBSA and occupation for reasons unrelated to local conditions if they decide to increase the national wage. With firm and market fixed effects, it measures how much wages are being pushed up or down in a given market due to being a national wage setter, relative to the same firm and market in previous years. So long as firms that set wages nationally do not change their other benefits or amenities (at the occupation and CBSA level) when the nationally set wage changes, then the instrument captures changes in the firm's recruiting success that are due only to the changing wage.

How the Instruments Work Together

The effects that we identify in our instrumental variables specification apply only to the sample of treated firms (Imbens and Angrist (1994)). We estimate the effect of wages and the employer contribution rate on recruiting success for firms that 1) have recently failed a non-discrimination test and 2) set wages nationally.

Table V shows summary statistics for firms that are affected by each instrumental variable separately, the two IVs together, and the full sample. Firms that are affected by either NDT Failure or being a National Wage Setter tend to be larger than the average (or median) firm, both in terms of plan assets and number of employees. As expected, firms that fail NDT test tend to have lower contribution rates and levels of employer contributions. Firms to which both IVs apply are by definition, multi-establishment firms, versus only 41 percent of firms in the full sample. In the sample affected by both instrumental variables, there are 588 firms with 55,786 unique jobs and 66,938 transitions from 2010-2019.

Figure 1-5 shows the geographic, industry, and occupational composition of the firms affected by the instrumental variables and the full sample. Each group has broad representation across geographies and sectors. In the main results, we include only firms that both have DC plans and that have greater than two establishments.

1.2.2 Data

This paper uses a panel data set of posted wages at the firm by occupation by CBSA level, individual (worker) level resume data with detailed job information, and firm-level retirement plan financials. We aggregate each source to the yearly level from 2010-2019 and merge all three sources to combine information on wages, new worker transitions, and retirement plans at the firm-level. The resume and posted wage data are from Lightcast (formerly Burning Glass Technologies) and the retirement plan financial data are from regulatory filing Form-5500. In the following subsections, we describe each data source in detail. In Section 1.2.2, We describe how we construct the estimation sample and compare the sample to the average U.S. firm.

Lightcast

Job Postings Data The Lightcast data on posted wages contains the near universe of online job postings. The postings are collected from over 40,000 distinct sources including company websites and online job boards, with no more than 5% of vacancies from any one source (Hazell and Taska (2020), Schubert et al. (2020)). Azar et al. (2020) shows that in 2016, the Lightcast job-posting database captured around 85% of all job vacancies, including offline jobs. So, while Lightcast likely omits job-postings in certain occupations where offline or informal postings are more common, it does capture the majority of posted jobs. Schubert et al. (2020) finds that particularly underrepresented occupations include low-wage food service jobs, cleaners, home health aides, laborer and cashiers. Thus, our estimates should be interpreted while keeping in mind that some occupations, particularly low-income ones, are underrepresented.

The main data we extract from the job postings is the posted wage. This is available for about 20% of postings, which equates to around 40 million postings from 2010-2019 with non-missing posted wages. The wage can appear either as single number or a range; when it is a range, we take the median as the posted wage. The data also include pay frequency, i.e. whether pay is hourly or annual, the type of salary (base or bonus pay), thus we can aggregate all wages up to an annual level. Appendix Table A.2 shows that the posted wages in Lightcast match well the wages in the Occupational Employment Statistics (OES).[17]

The final sample of posted wages also only includes jobs for which the SOC code, industry, and location information are available. We then collapse wages to the 5-digit SOC code, CBSA, year, and firm level. The final posted wage data set, prior to matching with the other data sources, has posted wages for over 8 million jobs at 1.2 million distinct firms. 437 out of 459 possible 5-digit SOC codes and 929 out of and 939 possible CBSAs are represented. The average annual salary is $50,247, the median $29,205. Appendix Table A.3 shows summary statistics of this sample.

One may be concerned about how well the postings data with wage information represents all jobs in the United States. Hazell and Taska (2020) show that the postings data, limited to postings with wages and job information are largely representative of the

[17] We collapse wages at the 5-digit SOC code by CBSA by year level and regress the OES salary (or hourly wage) on the Lightcast salary (or hourly wage), using the within-occupation and CBSA, medians, means, and within-occupation quantiles.

population of U.S. employment. Compared to data from the Bureau of Labor Statistics (BLS) and the Occupation Employment Statistics (OES), the data match the regional and occupation distribution of actual posted jobs well. Compared to Dun & Bradstreet data, the Lightcast data also represents the population of establishments well, based on industry classification and establishment age. we discuss the representativeness of the final estimation sample in more detail in Section 1.2.2.

Resume Data The next data set we use to construct the final estimation sample is a collection of resumes taken from online sources, also constructed by Lightcast. Resumes were sourced from a variety of Lightcast partners, including recruitment and staffing agencies, workforce agencies, job boards and social media. The resumes form a longitudinal data set, since we observe all jobs that an individual lists on their resume.[18]

In total, the data represent 83 million unique resumes with non-missing current job info. In 2010, the Lightcast resumes capture 26% of the total workforce; in 2019, this figure increases to 35% (see Appendix Figure A.3). Appendix Table A.4 shows high-level summary statistics for the sample. Across the 83 million resumes, there are 106 million transitions to new jobs, 65 million of which are to a new firm. Each resume has an average of 2.5 jobs represented with 1.6 transitions to a new job. The median job length is three years, compared to about four years in the BLS. The mean span of years observed on the resume is 15 years.

While the data represents a large percentage of the total U.S. workforce, it is not completely representative of the average U.S. worker. Appendix Figure A.4a shows that the average worker in Lightcast is younger than the average worker in the U.S., with about 75% of workers in Lightcast being under the age of 45.[19] The average worker

[18] A job here means a firm by occupation pair.

[19] We impute age by using education information on the resume. If the resume gives a year of high school

in Lightcast also has a higher education level than the average worker in the BLS data (Appendix Figure A.4b). About half of all the Lightcast resumes have non-missing education information. Using only that information, around 75% of workers in Lightcast have a college degree or more, versus only about 38% in the BLS data. However, if we assume that missing education info indicates that the worker did not receive education beyond college, then roughly 33% have greater than a college education, which matches the BLS much more closely. Some occupations are over-represented in Lightcast: management, business and finance, computer and mathematical, engineering, and arts and design (Appendix Figure A.4c). Others are underrepresented, such as office and administrative support, sales, food preparation, and healthcare. The Lightcast data matches the geographical distribution of all workers in the BLS at the state level quite well (Appendix Figure A.4d). In general, while the resume data captures a significant portion of the labor force, our results should be interpreted with the caveat they they apply to a younger, higher educated sample which works in typically higher-income occupations.

Form 5500

Our final data source is regulatory filing Form 5500, which contains detailed financial information on retirement and health plans for all U.S. firms. This is publicly available data, published by the Department of Labor.[20] Form 5500 is required yearly of all retirement plans which have qualified status under the Employee Retirement Income Security Act (ERISA). The form contains detailed information about their benefit plans, including what type of plan it is, financial information about inflows and outflows, how the funds

or college graduation, We assume the individual was 18 at the time of high school graduation and 23 for college graduation and use that year to calculate the worker's current age.

[20]https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/public-disclosure/foia/form-5500-datasets

are invested, the number of participants covered in the plan, and some information about plan features.

There are two version of Form-5500, one for large plans (those with greater than 100 participants) and one for small plans (those with less than 100 participants). The version required for large plans is significantly more detailed than the version required for small plans. However, small plans make up approximately three-quarters of all plans. Thus, we elect to include small plans in our final merge across all data sets in order to preserve sample size. The form for small plans provides enough information to 1) back out the effective employer contribution rate and 2) know whether or not the firm offers a healthcare plan and thus is sufficient for our analysis.

The administrative data do not include information about the specific default contribution rate or the structure of employer contributions offered in the plan. However, Schedule H of the Form 5500 for large plans and Schedule SF for small plans gives the amount in dollars that the employer contributes to the plan each year. It also gives the amount in dollars that participants contribute to the plan each year.

Form-5500 also contains details about healthcare plans. The main Form 5500 for large plans and Schedule SF for small plans have indicators as to whether or not the firm has a health plan. Schedule A of Form 5500 has specific information about health plans, including the insurance carrier and dollars paid by the plans on claims and total plan expenses, but this is available only for large plans.

Appendix Table A.5 shows summary statitics for all DC plans in Form 5500 from 2010-2019. Appendix Figure A.5 shows how the firms with retirement plans in Form-5500 compare to all U.S. firms, as measured by the Bureau of Labor Statistics. Roughly two-thirds of firms offer retirement plans, thus Form 5500 is representative only of those firms. Appendix Figure A.5a shows the distribution of industry (2-digit NAICS code) in each data set by number of firms, Firms with retirement plans are more likely to be in the Healthcare or Professional/Scientific/Technical industry than other U.S. firms. They are also less likely to be in Construction, Trade, and Service Industries. Appendix Figure A.5b compares firms in Form 5500 to the BLS by employment size. Firms with retirement plans are larger, with firms with more than 20 employees representing about one-third of all firms amongst firms with retirement, versus less than 10% in the BLS. Moreover, only about 20% of firms have 1-4 employees in Form 5500, versus about 61% in the BLS.

Variable Construction and Estimation Sample

Our final data set is a merged panel of posted wages, new-hire transitions and retirement plan information from the three combined data sets. We first merge the wage data with the resume data at the firm by occupation by CBSA level. That is, from the 65 million between-firm transitions with full job information in the posted resumes, We match the new workplace of the person changing jobs with the 8 million job-level posted wages, at the firm by occupation by CBSA level. Then, we merge in the retirement plan information from Form 5500 at the firm-level. All three data sources are matched through fuzzy merging on firm name, as the firm name may vary slightly between the three data sources.

The final estimation sample represents over half a million worker transitions to 24,000 firms in 486,000 unique CBSA by occupation jobs. Figure 1-6 shows how the main estimation sample compares to all firms in the U.S. The sample over-represents firms in the Professional/Scientific/Technical industry, as well as firms in Finance and Insurance, Information, and Healthcare. This is expected given the representativeness of both firms that have retirement plans (see Figure A.5a) and firms that are represented in the Lightcast data (see Hershbein and Kahn (2018) and Schubert et al. (2020)).

By size, the matched sample over-represents large firms, with about 18% of sample firms having over 500 employees, compared to less than 1% of firms in the BLS. This is driven by both the distribution of firms that have retirement plans, as these firms skew larger (see Figure A.5b) and due to the matching process, which is more likely to pick-up larger firms.

At the occupation level, we compare the distribution of all workers by occupation to the distribution of transitions in the Lightcast resume data. New workers in the sample over-represent those in Management, Business/Finance, Computer and Engineering occupations. This is consistent with the distribution of all resumes in Lightcast shown in Appendix Figure A.4c.

In Section 1.2.2 we showed that the posted wage in Lightcast closely matches the wage in the BLS's Occupational Employment Statistics. Table VI compares the distribution of wages in the OES with that of the matched sample. The means and median wage in the sample are slightly lower than that in the BLS. But the 90th percentile is significantly higher; this is expected given the skew toward higher income occupations and industries. However, the distribution of wages in the sample matches that of the OES reasonably well.

Table VII shows the main summary statistics for the matched sample. In the following paragraphs, we describe how we construct the variables needed for our analysis.

Employer Contribution Rate: Form 5500 does not directly disclose the employer's matching formula or contribution rate. Instead, we use the combined Form 5500 with the Lightcast wages to calculate an “effective" employer contribution rate. Schedule H of Form 5500 (or Schedule SF for small plans), gives the total dollars that the employer contributes to the plan each year. Then, from the Lightcast job postings data, we calculate the average wage at the firm. From the average wage and the number of employees at the firm (from Form 5500), we calculate total wages paid. Then, we divide the employer contribution by total wages paid to infer an effective employer contribution rate. While this measure does not capture the exact matching formula, it does serve as a measure of the generosity of the employer's plans. The average effective contribution rate at firms in our sample is 5.12%. This aligns well with the descriptive evidence in Arnoud et al. (2021), which finds that the majority of plans offer at least 3% and 40% of plans offer matching contributions up to 6% of salary.

There are two alternative measures of plan generosity used in robustness checks. First is the employer contributions in dollar per person, which is directly from Form 5500. Second is the employer's ratio of all contributions, relative to contributions from the employer and the participants summed together, which is also directly from Form 5500. All three variables give similar results.

Hiring Success: Hiring success at the firm, occupation, and CBSA level is the main outcome variable for our instrumental variables specification. We measure this by comparing the resume data with the postings data. When we see a job postings in the postings data, We can then see if it is filled within the year in the resume data. The average hire success in our matched sample is 11%. The average percentage of employees captured in the resume data, using Form 5500 as the true number of employees, is 42%. Given that the resume data do not capture all employees, its is likely that our estimates of hire success is downward biased for many firms.

1.3 Effect of Wages and Retirement on Firm Recruiting

In this section we describe the results of the instrumental variable approach to measuring the effect of wages and retirement contributions on firm hiring success outlined in Section 1.2. The empirical specification is shown in equation (1.2). The sample studied in this section is the 24,000 firms and 486,000 unique firm by CBSA by occupation jobs for which we matched data on posted wages, new worker flows, and retirement plans. Summary statistics for the sample are shown in Table VII.

1.3.1 Instrumental Variable Results

Table VIII shows the results from the main estimating equation in (1.2) on the full sample of firms. We find that a one percentage point increase in the employer contribution rate increases the likelihood of filling a position in each CBSA by occupation in which a firm recruits by 2.7%. A 1% increase in wages has an effect of half the magnitude: 1.4%. Note that these unit increases for the independent variable are roughly equivalent in dollar terms, as a one percentage point increase in the employer contribution rate has the same effect as a 1% increase in salary.[21,22] The average hire success in the sample is 11%, so these are large effects. The estimation sample is limited to those that both have at least two establishments and have a DC plan, as those are the firms to which the instruments are applicable. The specification has firm by CBSA by occupation fixed effects and controls for log employment size by year, log assets in the DC plan by year and whether or not the firm has a healthcare plan.

[21] For example, with a salary of $100,000 going from a 2% to a 3% employer matching rate means going from $2,000 to $3,000 in employer contributions, or a $1,000 increase. This is equivalent to a 1% increase in the $100,000 salary.

[22] We say roughly equivalent because these calculation do not account for taxes or non-linear matching formulas.

Comparing the OLS and IV results (columns (3) and (4) of Table VIII) suggests that there is a significant correlation between retirement and unobserved characteristics that affect hiring success and that the OLS coefficients are indeed biased. There is a very large difference between the two coefficients on the employer contribution rate. However, there is not a large difference between the coefficients on log salary. Consistent with the discussion of endogeneity in Section 1.2, the OLS coefficient on the employer contribution rate fails to pick up employee valuation of retirement contributions.

These results suggest that, by revealed preference, workers value a dollar of retirement contributions roughly double what they value a dollar of wages. In other words, the average worker in this sample values a dollar of employer contributions to his DC account nearly twice as much as an extra dollar of annual salary. The results applies to this sample of workers, who are higher income and working in selected occupations and industries. Moreover, the effects are measured for those who went to work at firms affected by both national wage setting and failure of non-discrimination tests.

The specification has strong first-stages, with an F-statistic of 136. As outlined in Sections 1.2.1-1.2.1, national wage setting has a positive effect on salary and NDT failure has a positive effect on the employer contribution rate. Within firm, employer contribution rates increase by .33 percentage points following NDT failure. With the firm-fixed effect, national wage setting does appear to have an effect on contribution rates (unlike in the difference-in-difference results shown in Section 1.2.1, in which national wage setting was not correlated with contribution rates). However, the negative effect of national wage setting is smaller in magnitude than the effect of NDT, so the net effect is still positive. The national wage setting instrument measures how much the firm's wage is being pushed up due to national wage setting. Thus, it is natural that employer contributions might decrease within firm in years when there is more upward wage pressure in order to even out costs.

Heterogeneity by Income

Table IX shows that the effect of the employer contribution rate is driven primarily by high-income occupations. In the high-income sample (column (8)), recruiting success increases by 6.03% for a one percentage point increase in employer contribution rate, relative to a 2.03% increase for a 1% increase in salary. In contrast, the low-income occupations (column (4)) are roughly equally effected by similar dollar increases in wages and employer contribution rate; hiring success increases by 1.87% and 1.60%, respectively. These results are consistent with standard lifecycle saving and consumption theory that higher income individuals tend to save more, and would therefore value these contributions more highly (Gourinchas and Parker (2002), Parker et al. (2022), Carroll (2000)).

These estimates come after re-estimating (1.2) on sub-samples of occupations, partitioned by salary levels. We take the median annual salary of all industry by occupation groups and designate low-income occupations as those below the median ($44,508) and high-income occupations as those above the median.[23] Thus, the partitioning is based on occupational salaries, not individual job salaries, so that we still compare jobs within the same occupation.

As before, both instruments have a strong first stage in each of these subsamples. The F-statistics are 140.67 for the low-income occupations and 23.00 for the high-income occupations. The national wage setting instrument has a positive association with salary and employer contribution rates increase following NDT failure in both subsamples.

[23] After creating these subsamples, the low-income group has a median (mean) annual salary of 31,000 ($33,484). The high-income group has a median (mean) annual salary of $63,585 ($70,470).

Heterogeneity by Age

Table X shows that the powerful effect of employer contribution rates on recruiting success is also increasing in age.[24] In high-age occupations, a one percentage point increase in employer contribution rates increases recruitment success by 7.2%. A 1% increase in wage increases recruiting success by only 2.36%. On the other hand, recruiting success in low-age occupations is much more affected by wages than by retirement. A 1% increase in wages improves recruiting outcomes by 2.05%; a one percentage point increase in contribution rates has no statistically significant effect on recruiting outcomes in this sample. This finding is consistent with lifecycle theory of saving and consumption (Gourinchas and Parker (2002), Scholz et al. (2006)) in that older people would be expected to save more and thus would value retirement contributions more highly. Again, both instruments have a strong first stage in the expected direction in each of the subsamples.

Heterogeneity by Gender

Table XI shows that the effect of retirement and wages on recruiting success in similar across occupations with differing gender compositions. In both male-dominated and female dominated occupations, a one percentage point increase in the employer contribution rate increases hiring success by between 3% and 4%. A one percent increase in annual salary increases hiring success by between 1.4% and 2% for both groups.[25]

[24]We partition the sample by median age in the occupation, as measured by the BLS's Occupational Employment Statistics. Occupations with an median age of greater (less) than 42 are "older" ("younger") occupations. The highest-age occupations are: motor vehicle operators (other than taxi drivers, chauffeurs, truck drivers and bus drivers), crossing guards and agricultural managers. The lowest-age occupations are: lifeguards, restaurant hosts and hostesses, and counter attendants at cafeterias, concession shops and coffee shops.

[25]We divide occupations into two groups: majority female or majority male, as measured by the BLS's Occupational Employment Statistics. The most female occupation are: skincare specialists, preschool and

In sum, for the average firm in our estimation sample, a one percentage point increase in contribution rates has nearly twice the effect on recruiting success as a 1% increase in wages. By revealed preference, this suggests that job-switchers in this sample value retirement contributions nearly twice as much as wages. The results are identified by comparing within firm, CBSA and occupation across years when wages and benefits change due to national wage setting and failure of NDT, respectively.

1.4 Eliciting Willingness to Pay in a Survey

In this section, we outline our second method of measuring employees valuations of retirement benefits: an online experimental survey that directly measures valuations by eliciting willingness to pay (WTP) for various levels of retirement benefits. We first discuss the method of eliciting WTP and then the sample of participants who completed the survey experiment.

1.4.1 Method

In each survey question, we show participants two similar job offers in which only the wage and the retirement vary slightly. Showing different combinations across many participants allows us to estimate the complete distribution of willingness to pay for the retirement benefits, similar to Mas and Pallais (2017).

In each question, the participant was shown two job offers. One of these was always

kindergarten teachers, and executive administrative assistants. The most male occupations are: cement masons, concrete finishers, and terrazzo workers, extraction workers, and electrical power line installers and repairers.

a "baseline" job offer that has no retirement, which is stated explicitly. Then, the salary was varied in downward increments (randomly) across participants for the second offer, while adding a retirement benefit. Figure 1-7 shows two example questions. Note that participants were told explicitly what the difference in take home pay would be between the two jobs. They were also told, prior to seeing each condition, that the two jobs they were choosing between were exactly the same, other than what was observed in the table. Vacation, healthcare, and remote work were always exactly the same between the two choices in which retirement varied.

We tested 5 conditions in total[26]:

1. Willingness to pay for a 401(k) with a 3% match versus no 401(k)
2. Willingness to pay for a 401(k) with a 5% match versus no 401(k)
3. Willingness to pay for a 401(k) with no match versus no 401(k)
4. Willingness to pay for a 401(k) with a 3% match versus a 401(k) with no match
5. Willingness to pay for working remotely for 2 days a week, versus no days of remote work

The first two conditions simultaneously measure the willingness to pay for both the intensive and the extensive margin of retirement benefits. The third measures only the extensive margin. The fourth measures only the intensive margin. The last measures willingness to pay for a non-monetary amenity, remote work, in order to compare the estimates from our sample to other estimates in the literature.

[26] Appendix Figure A.6 shows example questions for each condition

We follow the procedure in Mas and Pallais (2017), using a discrete choice framework to estimate willingness to pay. Imagine individual i is shown two jobs with wage difference $w_1 - w_0 = \Delta_w$ where job 1 offers the better retirement (R_i) than job 0. Then her willingness to pay (WTP_i) if she is fully attentive for the (better) retirement benefit is:

$$P_{\Delta w} = \Pr(WTP_i > -\Delta w) \tag{1.3}$$

However, some survey participants are likely to be inattentive. Those participants are equally likely to choose either job. Imagine 2α percent of participants are inattentive; then α of them will choose a dominated option by chance. Then, the probability that that an individual chooses the job with the better retirement benefit is:

$$Pr(R_i = 1|\Delta w) = P_{\Delta w}(1-\alpha) + (1 - P_{\Delta w})\alpha \tag{1.4}$$

$$= F(b\Delta w + c; \mu, \sigma)(1-\alpha) + \alpha \tag{1.5}$$

where μ is the population mean willingness to pay, while σ is the population standard deviation. Equation (1.5) is a mixture model that can be estimated by maximum likelihood. We assume the distribution follows a logistic distribution and bootstrap the standard error. Although the fraction of inattentive participants, α, is identified in (1.5), we estimate it directly through an attention check. A fraction of participants in each condition view a dominated condition - that is the job offers higher pay and the better retirement benefit. We use the fraction which do not choose this job to estimate α directly. We find that only a small number (2.5%) of participants are inattentive.

Testing each of the conditions across hundreds of participants means that we can flexibly estimate the full distribution of the willingness to pay in the population of survey participants. We thus can estimate willingness to pay for higher retirement benefits both on the intensive and extensive margin.

1.4.2 Data and Sample

For our survey sample, we recruited 1,600 online participants via Prolific. Prolific is an online platform on which participants are paid to take part in survey research.[27] We limited our sample to only those who live in the U.S., speak English as a first language, and were currently working or looking for work. We also balanced the sample equally between non-college graduates and those with a college degree or more. Figure 1-8 shows summary statistics for the survey participants. The average age is 34 (median 32). The sample is equally balanced on gender. The majority (75%) of participants are white.

1.5 Survey Results

In Section 1.3, we showed that, by revealed preference, workers place nearly double the value on retirement contributions than on wages when selecting between two otherwise similar jobs. In this section, we show that workers place about 1.5 times the value on employer contributions to 401(k)s than on wages when comparing between similar jobs. We measure this by eliciting valuations of retirement benefits directly in an online experimental survey setting. The method of eliciting willingness to pay and the sample of participants are described in Section 1.4.

[27] Prolific is similar to MTurk, which has been more commonly used in Economics and Finance studies. However, several studies from other fields have shown that Prolific produces higher data quality than MTurk (Peer et al. (2017), Péer et al. (2021)).

First, we start by measuring the value of additional retirement benefits on the intensive margin by varying the dollar value given in contributions by the employer conditional on having a retirement plan (a 401(k)). These results are most directly comparable to the revealed preference results of Section 1.3. Specifically, we asked participants if they would prefer a job with a 401(k) with no match or a 401(k) with a 3% match (see Figure 1-7a). The wage was higher in the condition with no match and varied in downward increments randomly across participants for the job offering the 3% match.[28]

The first finding is that a majority of participants (80%), exhibit some willingness to pay for the 3% match. This is shown in the black line in Figure 1-9, with corresponding summary statistics in Table XII, Column (1). The plot shows the fraction of the participants who chose the job with the 3% match against the total compensation gap.[29] 80% chose the job with the 3% match when the total compensation difference was $0, indicating that only 20% of participants place no value on the employer match. Moreover, 55% of participants chose the job with the 3% match even when it paid less overall.

Focusing now on the participants to the left of the black line, who saw conditions in which the job with the match paid less overall, their implied willingness to pay in total compensation is $748, or 1.5% of compensation offered in the higher-paying job. Net of taxes, the willingness to pay is $258.[30] In terms of wages, rather than total compensation, the willingness to pay is $2,226, or 4.3% of the annual salary offered by the job without the match.

Next, we show that the distribution estimated by maximum likelihood and corrected

[28]In some conditions, the job with the 3% match offered a higher wage in order to test for inattention.

[29]The total compensation gap includes the dollar value of the employer match.

[30]The tax calculation accounts for the difference in taxes paid because the job with the 3% match offers a lower wage and the match dollars are not taxable at the time they are paid out. It does not account for future taxation of the employer contributions.

for inattention implies an even higher willingness to pay: \$909 or 1.8% of total compensation. The top panel of Table XIII shows the results. Participants at the 25th percentile of the willingness to pay distribution are willing to give up only \$60 in annual pay to get the 3% match. The top 25% of workers are willing to give up at least \$1757, or 3% of total compensation to get the employer match. Hence, there is significant heterogeneity across workers in the valuation of this benefit; but the majority of workers are willing to pay for the 3% match, relative to a 401(k) with no match. These results come from estimating the inattention rate using the procedure described above in Section 1.4; the inattention-corrected shares are plotted in the red dotted line in 1-9a, along with the inattention corrected maximum likelihood estimates, plotted in a blue dashed line.[31]

Relative to the revealed preference results in Section 1.3, the relevant comparison here is how many dollars in wage participants are willing to give up for each dollar in match. The willingness to pay estimates suggest that workers will give up about 1.4% of wages for each one percentage point increase in the match. This is less than the near two to one trade-off estimated in the instrumental variables specification on the full sample, but directionally consistent in that it implies a higher valuation of employer contribution dollars than wage dollars when comparing amongst similar jobs.

Next, we show the results for the condition that tested for the WTP for the extensive margin of retirement benefits, or whether or not the job offers a 401(k) at all. In this condition, participants were given the choice between a job with a 401(k) that has no match or a job with no 401(k) (see Figure 1-7b).

Starting with the raw share of participants that chose the job with the 401(k), 80% chose the 401(k) job when it paid only \$500 less, implying that less than 20% of participant

[31] Note that the inattention correction changes the distribution only slightly, as we found very few participants to be inattentive in the sample.

have no willingness to pay for the 401(k).[32] Figure 1-9b shows the results; corresponding summary statistic are in Column (2) of Table XII. Moreover, 49% chose the job with the 401(k), even when it paid less. When the compensation premium is less than $2000, this figure increases to 60%.

This implies that the majority of participants have some willingness to pay for the 401(k), even when it offers no match. Focusing on the participants to the left of the black line, who saw conditions in which the job with the 401(k) paid less overall, their implied willingness to pay for the 401(k) is $1,775 or 3.4% of total pay. Note that the willingness to pay in wages and total compensation are the same in the condition, as there is no added dollar value from a match. Net of taxes, the willingness to pay is $1,345.[33]

As in the intensive margin condition, the estimation of the distribution by maximum likelihood results in an even higher mean estimated willingness to pay - $2,268 or 4.4% of total pay. The bottom panel of Table XIII shows the estimated distribution from the maximum likelihood procedure. The 25th percentile is $823 or 1.5% of total pay. The 75th percentile is $3,711, meaning that 25% of participants are willing to pay at least 7.2% of total pay to get a 401(k).

The test for the extensive margin is not directly comparable to the revealed preference results from Section 1.3, as the specification in that analysis included only firms with a DC plan. However, the survey results show that, when accounting for the match dollars, job-seekers are willing to pay even more just to get a 401(k), even when it offer no match, than they are to get an employer match, conditional on having a plan. The average

[32] There is no condition that had an exact $0 difference in compensation in this case, but the closest is that in which the total compensation difference is only $500

[33] In this condition, the inattention correction (red dashed line) does very little to change the estimates as inattention was not prevalent in the sample. The attention-corrected maximum likelihood estimates (blue dashed line) smooths out the distribution and slightly alters the tails, but does not significantly change the median.

willingness to pay for the 401(k) alone is about three times the willingness to pay for the 3% match (4.4% of total pay versus 1.5% of total pay). This suggests that workers value DC retirement plans beyond the dollars offered and that the plan itself provides value to employees even when it does not offer any dollar matching from the employer.

In Appendix Figure A.6 and Tables A.6-A.7,we show the results for the remaining conditions that we tested. Two conditions tested simultaneously for the intensive and extensive margin of benefits, offering participants a choice of a job with a 401(k) with 3% or 5% match versus a job with no 401(k). The final condition tested the willingness to pay for the ability to work remotely for two days per week, versus no remote work option. The willingness to pay for the extensive and intensive margin simultaneously aligns with the results that test each condition separately; the WTP for both is higher than the WTP for each separately. comparing the 3% conditions, the estimated average willingness to pay for a 401(k) and a 3% match separately is around \$3,200 versus \$3,600 for both the 401(k) and the 3% match simultaneously.. For the remote work condition, We find an average willingness to pay of \$2,935 in annual salary for two days of remote work per week. This aligns well with the finding from Mas and Pallais (2017) who estimate and average willingness to pay of \$2,533 in annual salary.[34]

Reasons for Valuing Retirement: In a sub-sample of participants ($N = 600$) who were tested for the extensive margin condition, we asked why they chose the selected job. One-half were given multiple choice options and one-half were given a text box in which they could write freely. For the multiple choice options, participants could choose between the following:

1. Chosen job has a higher wage

[34]Mas and Pallais (2017) find an average willingness to pay of \$1.33 per hour for the option to work from home, which we scale up to an annual salary assuming full time work of 1,920 hours per year.

2. Chosen job has higher total compensation

3. Chosen job has a retirement plan

4. Chosen job is a better job

For those choosing between a 401(k) or no 401(k) when the job with the 401(k) paid less, the overwhelming majority (91%) say they chose it because of retirement plan, as opposed to only 5% who chose is because it's a "better job".[35] This indicates that participants value the retirement plan itself for some reasons besides the dollar value, and not that they see it as signal of employer quality. In the open text-box responses, participants who chose the 401(k) job when it paid less almost unanimously said that they chose the job because of the 401(k) or the job having better benefits. 4% mentioned the tax advantages and 5% mentioned the importance of saving for the future. A few selected answers below capture the qualitative nature of the majority of the responses:

- "Because it had a retirement plan, even though it didn't have an employer match."
- "$83 lower income per month is nothing compared to the long-term benefit of having money set aside for retirement. "
- "The company sponsored 401(k) instead on $1000 annually seems like a good deal."
- "Retirement benefits are always good."

[35]The remaining 4% erroneously said they chose the job because it had higher total compensation, when it fact did not. They may have considered the 401(k) to be part of compensation, but there was no dollar value associated with it. See Appendix Figure A.7.

1.6 On-the-job Search Model with Retirement Benefits

In Sections 1.3 and 1.5, we showed across two distinct empirical settings that job-seekers place approximately 1.4-2 times the value on retirement contributions than on wages when comparing between two otherwise similar jobs. Moreover, the survey results show that most workers also exhibit willingness to pay for having a DC retirement plan at all. Motivated by these facts, in this section, we develop a random on-the-job search model, similar to Burdett and Mortensen (1998) and Sorkin (2018), in which workers value retirement benefits (on both the intensive and the extensive margin.) and the other non-wage portions of compensation separately. The model allows us to both directly estimate worker valuations and show the effect of retirement policy on labor market outcomes. We first describe the model setup and how it is estimated in our data. Then we describe the results and a validation exercise using the non-discrimination testing exercise.

1.6.1 Model Setup

The model is an on-the-job random search model, in the category of Burdett and Mortensen (1998), The model is partial equilibrium in the sense that wages and firm behavior are exogenous. Employers post contracts in which they are willing to pay workers a (exogenous) wage premium, proportional to the worker's skills and match value with the firm and exogenous benefits. Workers make binary choices over job offers, based on their valuation of the wage, benefits, and idiosyncratic features offered by the job. We focus only on transitions within industry, occupation and CBSA, to highlight job changes that are directly related to firm differences, rather than career or location changes.

Firms: There are J firms in the economy. Each firm employs workers in L_j unique occupation by industry by CBSA markets (henceforth, markets), indexed by $l_j = 1_j, ..., L_j$. Each firm is characterized by the tuple: $\mathbf{J_j}, r_j, \mathbb{1}_{r_j}, \mathbb{1}_{h_j}, \mathbf{g_j}, \mathbf{f_j}, \mathbf{a_j}$ with

- $\mathbf{J_j}$: a $1xL_j$ vector in which each element is the log wage premium paid by firm j to all workers in a market l equally
- r_j: the firm's employer contribution rate, a constant within firm
- $\mathbb{1}_{r_j}$ an indicator equal to one if the firm offers a retirement plan
- $\mathbb{1}_{h_j}$: an indicator equal to one if the firm offers a healthcare plan
- $\mathbf{g_j}$: $1xL_j$ vector in which each element is the number of employees in market l at firm j. Denote $\sum_{l=1}^{L} g_j = G_j$ where G_j is the total number of employees working at firm j.
- $\mathbf{f_j}$ a $(J-1)xL_j$ vector of firm j's recruiting intensities, where each element $f_{j,k,l}$ is the intensity with which firm j makes offer to employees of firm k in market l.
- $\mathbf{a_j}$: a $1xL_j$ vector in which each element is the non-wage, non-healthcare, non-retirement amenities offered by the firm to workers in market l.

To fix notation consider an example firm, Amazon, in the retail trade industry. Denote Amazon as firm 1. Amazon operates in the Retail Trade industry and employs workers in many occupations and cities. Denote marketing managers in San Jose as occupation $l_1 = 1_1$, marketing managers in New York City as $l_1 = 2_1$ and human resource managers in San Jose as $l_1 = 3_1$. The $\mathbf{J_1}$, $\mathbf{g_1}$ and $\mathbf{a_1}$ vectors thus contain the corresponding values to each of the markets for Amazon's workers (log wage premium, number of employees, and amenities, respectively). $\mathbb{1}_{h_j}$, $\mathbb{1}_{r_j}$, r_j are constants for all employees. Each row of $\mathbf{f_1}$ contains all the recruiting intensities of Amazon for the corresponding market in the

Retail Trade Industry, For example, the first row of $\mathbf{f_1}$ contains the recruiting intensity of Amazon toward marketing managers in San Jose at all other firms in the Retail Trade industry, etc.

Workers: M workers are characterized by $m_{i,l}$ which encompasses their skill-level, labor market experience and other factors for which they will be compensated equally by all employers while working in market (occupation, industry and CBSA) l. A worker's indirect utility from working at firm j is a linear combination of his log wage, $w_{i,j,l}$ plus the value of having a health plan, the log dollars of retirement benefits $r_{j,l}$, the value of having a retirement plan, the log-dollar value he places on the amenities at firm j, denoted $ln(\bar{a_{i,j}})$, and his idiosyncratic valuation for working at j:

$$V_{i,j,l} = \gamma_{i,l} ln(w_{i,j,l}) + (1-\gamma_{i,l}) ln(w_{i,j,l}(1+r_j)) + \beta_{i,l}\mathbb{1}_{r_j} + \alpha_{i,l}\mathbb{1}_{h_j} + ln(\bar{a_{i,j}}) + \epsilon_{i,j,l} \tag{1.6}$$

$$= \gamma_{i,l} ln(w_{i,j,l}) + (1-\gamma_{i,l}) ln(w_{i,j,l}) + (1-\gamma_{i,l}) ln(1+r_j) + \beta_{i,l}\mathbb{1}_{r_j} + \alpha_{i,l}\mathbb{1}_{h_j} + ln(\bar{a_{i,j}}) + \epsilon_{i,j,l}$$

$$= ln(w_{i,j,l}) + (1-\gamma_{i,l}) ln(1+r_j) + \beta_{i,l}\mathbb{1}_{r_j} + \alpha_{i,l}\mathbb{1}_{h_j} + ln(\bar{a_{i,j}}) + \epsilon_{i,j,l}$$

Note that the wage is individual specific, due to the individual's skill level, but the retirement is not: firms must pay the same retirement (as a fraction of salary) to all workers. $\gamma_{i,l}$ is the weight worker i in market l places on wages, (1-$\gamma_{i,l}$) is the weight placed on total dollar compensation (wages + employer retirement contributions) by worker i when working in market l. $\beta_{i,l}$, and $\alpha_{i,l}$ are the weights placed on the firm having a retirement plan and healthcare dollars, respectively. In the modeling framework, weights are individual and market specific. That is, workers can value the distinct parts of compensation differently than other workers and individual workers may value compensation components differently when they are working in different markets. Assume $\epsilon_{i,j,l} \sim N(0, \sigma_l^2)$.

Note that the distribution of the idiosyncratic match-value is market specific. This log-additive form of indirect utility is supported by findings in Maestas et al. (2018) and Mas and Pallais (2017).

Search and Transitions: Employed workers receive job offers sequentially (one at a time) from other employers randomly. Offers are received at an exogenous rate, λ.[36] When another offer is received, workers make a binary choice over the two jobs. Firms offer:

$$ln(w_{i,j,,l}) = m_{i,l} + \tilde{\eta}_{j,l} + \varphi_{i,j,l} \tag{1.7}$$

where $\varphi_{i,j,l}$ is a random draw from a mean zero distribution and $\tilde{\eta}_{j,l} = \eta_{j,l} - E[\varphi_{i,j,l}|$ Offer Accepted] is the pay premium j offers to workers in market l adjusted for the fact that those with a higher match value are more likely to be accepted. By offering $\tilde{\eta}_{j,l}$, the firm ensures that the actual average log-wage premium paid to workers is $\eta_{j,l}$.

If a worker is employed at firm j and receives an offer from firm k, he makes a binary choice over the two jobs. He will leave his current job if $V_{i,k,l} > V_{i,j,l}$ which occurs with

[36]Note that λ depends on the $\mathbf{f_j}$s of other firms in the market, but the functional form is not crucial for the subsequent analysis.

probability:

$$
\begin{aligned}
&P(V_{i,k,l} > V_{i,j,l}) \\
&= P(\gamma_{i,l} ln(w_{i,k,l}) + (1-\gamma_{i,l}) ln(w_{i,k,l}(1+r_k)) + \beta_{i,l}\mathbb{1}_{r_k} + \alpha_{i,l}\mathbb{1}_{h_k} + ln(a_{i,k}) + \epsilon_{i,k,l} \qquad (1.8)\\
&> \gamma_{i,l} ln(w_{i,j,l}) + (1-\gamma_{i,l}) ln(w_{i,j,l}(1+r_j)) + \beta_{i,,l}\mathbb{1}_{r_j} + \alpha_{i,l}\mathbb{1}_{h_j} + ln(a_{i,j}) + \epsilon_{i,j,l})\\
&= \Phi[\gamma_{i,l}(ln(w_{i,k,l}) - ln(w_{i,j,l})) + (1-\gamma_{i,l})(ln(w_{i,k,l}(1+r_k)) - ln(w_{i,j,l}(1+r_j)))\\
&+ \beta_{i,l}(\mathbb{1}_{r_k} - \mathbb{1}_{r_j}) + \alpha_l(\mathbb{1}_{h_k} - \mathbb{1}_{h_j}) + (ln(a_{i,k}) - ln(a_{i,j}))] \qquad (1.9)
\end{aligned}
$$

where Φ is the normal CDF $\sim N(0, 2\sigma_l^2)$.

Let $\Omega_l = ([j, k, \Delta ln(w_l)]_1, ..., [j, k, \Delta ln(w_l)]_{Sl})$ be the set of all S employer-to-employer transitions within market l. The joint likelihood of observing all such transitions, conditional on offers being made, is:

$$
\begin{aligned}
\mathbb{L}_l &= \Pi_{s=1}^{S\in l} \Phi[V_{i,k,l} - V_{i,j,l}] \qquad (1.10)\\
&= \Pi_{s=1}^{S\in l} \Phi[\gamma_{i,l}(ln(w_{i,k,l}) - ln(w_{i,j,l})) + (1-\gamma_{i,l})(ln(w_{i,k,l})(1+r_k)) - ln(w_{i,j,l}(1+r_j)))\\
&+ \beta_{i,l}(\mathbb{1}_{r_k} - \mathbb{1}_{r_j}) + \alpha_l(\mathbb{1}_{h_k} - \mathbb{1}_{h_j}) + (ln(a_{i,k}) - ln(a_{i,j}))] \qquad (1.11)
\end{aligned}
$$

In words, the likelihood of the given transitions occurring is the product of the likelihood of each individual transition.

1.6.2 Estimation of Random Search Model

The main estimating equation from the model, (1.10), is a likelihood function and can be estimated by standard maximum likelihood techniques. With an ideal data set which contains all information about wages, healthcare, retirement, amenities and all offers from outside employers, one could estimate each of the weights and the amenities term in (1.10). However, two challenges prevent us from directly estimating these parameters. First, our data set does not have information on amenities. Second, we do not observe rejected offers, only accepted ones. In the following paragraphs, we detail how we deal with each of these issues in the model estimation.

First, to address the lack of data on amenities, we move from estimating individual worker level weights to estimating weights at the occupation by industry level. That is, instead of the indirect utility function in (1.6), we have[37]:

$$V_{i,j,l} = \gamma_l ln(w_{j,l}) + (1-\gamma_l) ln(w_{j,l}(1+r_j)) + \beta_l \mathbb{1}_{r_j} + \alpha_l \mathbb{1}_{h_j} + ln(\bar{a}_j) + \epsilon_{i,j,l} \tag{1.12}$$

Estimating $\gamma_{i,l}$, $\beta_{i,l}$ and $ln(a_{i,j})$ from the likelihood function in (1.10) is not possible because the parameters are not identified. Even with data on transitions across many firm-to-firm pairs, the individual worker's weights cannot be pinned down without data on multiple transitions for the same worker. However, when the weights and the amenities term are averaged across all workers in an occupation by industry group, the parameters γ_l, β_l, and the mean of $ln(a_j) - ln(a_k)$ for all firm to firm pairs in S are clearly identified. This is because the measured retirement, wage and healthcare differentials across firms, plus an unobserved difference in amenities must explain the observed probabil-

[37]Note also that the i subscript on wages has been dropped, as we only observe occupation wages, not worker-specific wages.

ity of workers moving between firms. Using the observed transition probabilities for all workers moving between firms in the industry by occupation group gives sufficient degrees of freedom and variation to estimate the average value of the parameters for all workers.

The variation that identifies the weight workers place on retirement benefits and the amenities term is driven by the net flows of workers between employers. The estimation results in three main outcomes of interest: 1) a unique weight on the intensive margin of retirement contributions for each occupation by industry group, 2) a unique weight on the extensive margin of having a DC retirement plan for each occupation by industry group, 3) a residual term that corresponds to the amenities term in equation (1.10). This term explains other job features that are not captured by wages, healthcare, retirement, occupation, industry, or CBSA differences but that drive transitions between employers. In this estimation method, the estimated amenities term represents an average difference in amenities across all firm-to-firm pairs in that industry by occupation group.

To address the second issue, that we do not observe rejected offers, we borrow from Sorkin (2018) and Bonhomme and Jolivet (2009) and make an assumption about offer intensities at the firm by market-level. We assume that firms make offers to unemployed workers in that market at the same rate as they do to employed workers in that market and that unemployed workers never reject offers, that is $f_{k \to j,l} = f^{NE}_{j,l}\ \forall j \neq k \in J$. Thus, we use observed transitions out of unemployment to measure offer intensity ($f_{j,l}$).[38] This allows us to convert the conditional (on receiving an offer) probabilities of transition that we observe to unconditional probabilities of transition.

Normalizing $\gamma = 1$ so that all all other terms in the worker's valuation function are

[38] We measure unemployment from gaps in employment on an individual's resume. If a person has a gap of at least 6 months between two jobs on their resume, we assume they were unemployed.

in log-wage equivalent units, the empirical counterpart of (1.10) is :

$$
\begin{aligned}
\mathbb{L}_{n,l} &= \Pi_{s=1}^{S\in l}\Phi[V_{i,j,l} - V_{i,k,l}]^{\frac{1}{f_{j,l}^{NE}}\frac{1}{g_k}} && (1.13)\\
&= \Pi_{s=1}^{S\in l}\Phi[((ln(w_{k,l}) - ln(w_{j,l})) + \hat{\gamma}_l(ln(w_{k,l}(1+r_k)) - ln(w_{j,l}(1+r_j))) \\
&\quad + \hat{\beta}_l(\mathbb{1}_{r_k} - \mathbb{1}_{r_j}) + \hat{\alpha}_l(\mathbb{1}_{h_k} - \mathbb{1}_{h_j}) + \Delta\widehat{(ln(a_l)})]^{\frac{1}{f_{j,l}^{NE}}\frac{1}{g_k}} && (1.14)
\end{aligned}
$$

This is simply the likelihood function, weighted by both the inverse of the joining firm's offer intensity and the leaving firm's size. These estimation weights account for flows observed in the data that are due to firm size and recruiting efforts, but not the valuation of job qualities. This likelihood is estimated separately for each industry and occupation to get distinct weights on retirement benefits by worker type. Table XIV details the definitions of each component of the model.

Another step in the estimation of the search model is to limit the estimation sample to job-transitions that occur within industry, within occupation, and within CBSA. The motivation for doing so is two-fold. First, this choice eliminates variation in job-choice that comes from career or location switches. In this model framework, those differences would be picked up as "amenities." However, the preferences that drive these kinds of switches may not be directly comparable to wage and benefit levels. So, eliminating this variation makes the estimates of the different component valuations more directly comparable. Second, both the revealed preference results in Section 1.3 and the survey results in Section 1.5 relied on job-switchers who chose between very similar jobs. In the instrumental variables setting, the effect is measured within a firm, CBSA, and occupation, in which the wage does not vary drastically over time. In the survey setting, the wage difference between the jobs offered was only about $2,000, on average and the other benefits (besides retirement), were always exactly the same. Hence, these methods are measuring the willingness to pay for retirement benefits when choosing amongst jobs with relatively similar wage levels and other benefits and job features. The choice to only keep this specific type of job-switcher creates a similar setting, in which the choice to move is not driven by large differences in the wage or other job characteristics that are specific to an occupation and CBSA.

As a final step before estimating the model, we also discount the dollar value of retirement contributions at the firm level by the participation rate in the plan. The participation rate can be calculated directly form Form 5500, by dividing the number of actively contributing employees by the number of eligible participants. This step is necessary because if we assumed that all employees participated in the plan, then that would lead to an over-valuation of benefits. Rather, we assume that a smaller number of employees (only the participating ones) are getting the benefit. The average participation rate in DC plans in the estimation sample is 75%, which is lightly higher than but close to the BLS estimate of 68%.[39]

Summary statistics for the estimation sample are shown in Table XV. This sample represents about 35,000 transitions to 9,500 firms, relative to the half a million transitions to 24,000 firms in the instrumental variables results.[40] Moreover, this sample represents a higher-income group of workers - the mean salary is $56,000, versus $49,554 in the main sample; the median is $50,000 versus $41,000 in the main sample. Firms in the estimation sample are also significantly larger: the average (median) number of employees is 3,800 (856), compared to 400 (66) in the main sample. Table XVI shows the input values from the sample corresponding to the terms in equation (1.14).

[39]More details on: https://www.bls.gov/ncs/

[40]See Appendix Table A.9 to see how transitions change when each of the criteria are added.

The main parameters of interest are $\hat{\gamma}_l = \frac{1-\gamma_l}{\gamma_l}$ and $\hat{\beta}_l$, which are measured at the industry by occupation level. $\hat{\gamma}_l$ is the weight on log total dollar compensation, relative to a weight of one on log wages only. $\hat{\beta}_l$ is the weight on having a retirement plan relative to a weight of one on log wages. While we could directly estimate $\hat{\alpha}_l$, the weight on having a health plan, we choose instead to calibrate it in order to preserve degrees of freedom and focus on estimating retirement-related parameters.

To understand the interpretation of the coefficients, note that the coefficient $\hat{\gamma}_l = \frac{1-\gamma_l}{\gamma_l}$ is the ratio of the weight on log total pay to the weight on log wages. Consider two jobs at firms j and k that offer an agent the same indirect utility. For simplicity assume both firms have retirement plans, healthcare plans, and the same valuation of amenities by workers. It must be that:

$$ln(w_{k,l}) + \hat{\gamma}_l(ln(w_{k,l}(1+r_k))) = ln(w_{j,l}) + \hat{\gamma}_l(ln(w_{j,l}(1+r_j))) \tag{1.15}$$

$$\implies \hat{\gamma}_l = \frac{ln(w_{k,l}) - ln(w_{j,l})}{ln(w_{j,l}(1+r_j)) - ln(w_{k,l}(1+r_k))} \tag{1.16}$$

$$= \frac{\text{-\% change in wage}}{\text{\% change in total compensation}} \tag{1.17}$$

$\hat{\gamma}_l$ thus measures what percentage change in total compensation is necessary to compensate a worker for a 1% decrease in wage (or vice versa). Note the following ranges of interest for $\hat{\gamma}_l$:

- $\hat{\gamma}_l < -1$: workers place a higher weight on total compensation than on wages. A 1% decrease in wage must be compensated by some smaller increase in retirement. The worker can get the same utility at a lower level of total compensation, so long as retirement has increased.

- $\hat{\gamma}_l > 0$: workers place some weight on wages, but less than on total compensation.

A 1% decrease in wages must be compensated by a larger increase in total compensation, meaning retirement must increase[41]

1.6.3 Estimates of Retirement Valuations

Table XVII shows the main results following the estimation of (1.14), The mean weight on wages is -2.83 and the corresponding mean weight on total compensation is 3.83. This implies an elasticity of wages to total pay of -0.74, meaning that a 1% decrease in wages can be compensated for by an increase in retirement that results in a -0.74% decrease in total compensation. The median is similar to the mean. Industry by occupation groups at the 10th percentile of the wage valuation distribution have an elasticity of -0.88, meaning that they require a smaller compensating differential in retirement contributions than those at the mean. At the 90th percentile of the distribution, workers have $\gamma_l = .83$. This means that workers place a higher value on wages than on total pay. Approximately 25% of the industry by occupation groups has $\gamma_l > 0$, and thus places a higher weight on only the wage portion of compensation versus the total compensation (See Appendix Figure A.8). The remaining 75% would be willing to give up some total compensation to get a higher employer contribution to their DC plan.

The estimates of β_l, the weight placed on the extensive margin of having a plan, show an average valuation of .02, meaning that the average occupation by industry group of

[41]Special cases:

- When γ_l , the non-normalized weight on wages is greater than one, then $-1 \leq \hat{\gamma}_l \leq 0$: this means that workers care about wages more than total compensation and a decrease in wages can never be compensated for by an increase in total pay.
- When γ_l, the non-normalized weight on wages, is equal to zero, then workers only care about total compensation. Any decrease in wages must be compensated for exactly (dollar for dollar) in total compensation.

workers is willing to give up 2% of total pay to get a DC plan. The 10th percentile of the distribution is slightly above 0, but still positive, indicating that most workers are willing to pay for this benefit. The 90th percentile of the distribution of β_l estimates value the DC plan as 4% of wages. These estimates align well with the survey estimated value of 3.4% of wages to get a 401(k).

The results also show that workers in higher income industry by occupation groups place a higher value on DC retirement plans. Figure 1-10 shows the relationship between retirement valuations and salary. Figure 1-10a shows a binscatter of the weight on total pay $(1-\gamma_l)$ versus the average salary in the occupation by industry group. There is a strong positive correlation, with a coefficient of .38 and p-value of .009. Figure 1-10b shows a binscatter of the weight on having a DC plan (β_l) versus the average salary in the occupation by industry group. Again, there is a positive correlation, with a coefficient of .002 and p-value of .026, though the slope is much smaller than in the case of the intensive margin valuations. The willingness to pay for having the plan applies to almost all workers, while a significant portion of the distribution of workers (25%) has no willingness to pay for higher dollar contributions to the plan.

Compensating Differentials: In the following paragraphs, we detail what the model estimates imply about compensating differentials in retirement contribution dollars. Table XVIII shows the results from an exercise which supposes that wages were reduced by 1%. What compensating differential in retirement contribution would be required to give worker the same valuation as before the wage reduction, holding all other benefits constant? We split the sample by income groups to show how the valuation of retirement contributions varies with the income distribution.

The first column of Table XVIII shows the results for the 10th-25th percentile income group in the estimation sample. The average salary in these occupation by industry groups is $41,618. A 1% loss in wage is equivalent to a loss of about $414. For workers to get the same valuation from retirement contribution after this wage reduction, they need retirement contributions to increase by $2,240. Thus, total compensation must increase by $1,825 or 4.1%. This groups places a higher value on wages than on total pay, and thus needs large compensating differential to be made indifferent between the two compensation bundles. The increase in retirement needed is equivalent to the employer increasing their contribution rate by 5.48 percentage points (i.e. the employer goes from offering a 2% contribution rate to a 7.48% contribution rate).

Moving to the second column, we show the results for the middle 10% of the income distribution. This group has an average salary of $57,000, so a 1% decrease in wages is equivalent to a lost of about $567. To be made indifferent, this group only needs retirement to increase by $108, meaning thy will accept a total pay cut of $460, or 0.75%. This group of workers requires the employer contribution rate to increase by only 0.28 percentage points (i.e. the employer goes from offering a 2% contribution rate to a 2.28% contribution rate). This group values total pay more than wages, so they will accept a wage cut to get a slightly higher retirement contribution.

Similar to the middle income group, the high income group also values total pay more than wages. A 1% pay cut is equivalent to a loss of around $740 for this group. They require only a $130 increase in retirement contributions to get the same valuation as before the pay cut. So, this group will accept a total pay cut of about .77% so long as the employer contribution rate increases by .25 percentage points.

The instrumental variable estimates, the survey results, and the results from this structural model all show that most workers are willing to take less total compensation in return for a greater share of compensation as retirement benefits. The instrumental variables estimates and the structural model imply quite similar willingness to pay for higher income households. The instrumental variable results indicated by revealed preference that the average work values a dollar of retirement contribution nearly twice as much as a dollar of wage, while high-income workers value it nearly three times a much as a dollar of wage. In the estimated search model, the trade-off is 3-4 dollars of retirement for one dollar of wage amongst higher income workers; thus the estimates are fairly comparable in magnitude. Recall that the estimation sample for the model captures more high-income workers and occupations, thus it is expected that the results match more closely to the high-income only results from the instrumental variables analysis.

In the survey of hypothetical job choices, workers were willing to give up about 1.4% of wages for each one percentage point increase in the match - this is less than the trade-off estimated in the models. In terms of the extensive margin, survey participants were willing to give up 3.4% of wages to get a 401(k), versus about 2% on average here. Around 80% of participants had some willingness to pay for the 401(k) plan in this survey; about 90% of workers do according to the model estimates. Survey participants were 50% college graduates, compared to 75% in the Lightcast resume data. Survey participants had a median age of 32 versus about 40 in the Lightcast resume data. These differences could explain the higher valuation on retirement values found in the model estimation, as older and higher income workers seem to value retirement more.

1.6.4 Model Validation

While the total pay and amenity terms are separately identified in (1.14), one may still be concerned that the weights pick up variation from the amenities term and vice versa.

That is, it could be that the preference for retirement is reflective of the fact that firms with better retirement also have better amenities. To address this, we use non-discrimination testing to show that firm-level amenity valuations do not change around NDT testing.

To complete this exercise, we must estimate *firm-level* valuations, rather than industry by occupation average weights. In the main modeling framework, we elect to estimate weights at the occupation by industry level for two main reasons. First, we are interested in estimating a worker-level quantity, not a firm-level one. Ideally, we could estimate these quantities at the individual level, but as discussed above, this is not possible in our data. Hence, the occupation by industry averages are a way to estimate a proxy for the desired parameter (worker-level weights on retirement) that is possible in our data. Second, we have data on retirement contributions and plans. Most papers that estimate firm-level valuations (Sorkin (2018), Lehmann (2022), Bonhomme and Jolivet (2009)) do not have such data, and thus estimate the valuation of all benefits at the firm-level. We have data on benefits (retirement and healthcare), and thus can estimate worker valuations for each component separately, using variation across workers in the same industry by occupation group.

This estimation method results in an industry by occupation average amenity difference, denoted $\Delta\widehat{ln(a_l)}$ above, which represents the average amenity difference between each firm pair in the industry-by-occupation group. This does not reveal firm-level amenity valuation. As a validation exercise of the model, we estimate firm-level amenity valuations, net of the estimated retirement valuations and show that the amenity valuations do not change around NDT failure for the subset of firms who failed in our estimation sample. The likelihood function to estimate firm-level valuations is:

$$\mathbb{L}_{n,l} = \Pi_{s=1}^{S \in n,l} \Phi[(\gamma[ln(w_{j,l}) - ln(w_{k,l})] + ln(\bar{a_{j,l}}) - ln(\bar{a_{k,l}}]]^{\frac{1}{f_{j,l}^{NE}} \frac{1}{g_k}} \tag{1.18}$$

Note that in this method, we no longer use any data on retirement or healthcare. We estimate firm-level valuations for all amenities, inclusive of retirement, healthcare, and any other benefits. As before, we normalize $\gamma = 1$ so that the amenities term is in log-wage equivalent units. Each firm's valuation is estimated against a base-firm, that is $\bar{a_{k,l}}$ is normalized to zero for some firm k, which we select to be the firm with the most transitions in the estimation sample.

We have shown in Sections 1.2.1 and 1.3 that the shock of NDT failure induces changes in the measured retirement contributions. Thus, NDT failure can be used to test the model by showing that the firm-level amenity term, net of implied retirement valuations, does not change around NDT failure. That is, $[ln(\bar{a}_j) - ln(\bar{a}_k)] - (1 - \gamma_l)(w_{j,l} r_j)$, should not change following a firm's NDT failure.

Figure 1-11 shows the results of a difference and difference regression for this estimated quantity around NDT failure. There is no significant different in amenity valuation, net of the implied valuation on retirement contributions, for NDT failing firms after NDT failure. In the model estimation sample, around 800 firms fail NDT at some point from 2010-2019. The regression controls for year by industry fixed effects, log number of employees, and log dollars of assets in the retirement plan. Hence, firm-level amenity valuations do not change around NDT failure. The results of this test indicate that the estimated weights on retirement contributions in Section 1.6.3 indeed are measuring the valuation of retirement, not other amenities.

1.7 Firm Policy Implications

In this section, we assess whether or not firms could improve recruiting outcomes with changes in compensation structure in partial equilibrium under the constraint that retirement benefits are common across workers and firms hire workers across the income distribution who value retirement benefits heterogeneously. We consider the distribution of workers in our data – not just new hires – and measure the differential impact on worker well-being across the income distribution.

1.7.1 Setup

The search model estimates yield a value for γ_l in each industry by occupation group, which reveals how much workers value a 1% increase in retirement contributions relative to a 1% increase in salary.

Let $\Omega_{j,n,l} = ([j, 1, \Delta ln(w_l)], ..., [j, k, \Delta ln(w_l)]_J)$ be the set of all possible transitions of workers in market l to firm j from all other firms in the same industry employing workers in the same market. When looking at all such transitions, there is an average probability that workers who are offered jobs at firm j will accept the offer. The probability that j offers a greater value than a given firm, k is:

$$\begin{aligned} Pr(V_{i,j,l} > V_{i,k,l}) = \Phi[&\gamma_{i,l}(ln(w_{k,l}) - ln(w_{j,l})) + (1-\gamma_{i,l})(ln(w_{k,l}(1+r_k)) - ln(w_{j,l}(1+r_j))) \\ &+ \beta_{i,l}(\mathbb{1}_{r_k} - \mathbb{1}_{r_j}) + \alpha_l(\mathbb{1}_{h_k} - \mathbb{1}_{h_j}) + (ln(a_{i,k}) - ln(a_{i,j}))] \end{aligned} \tag{1.19}$$

Denote this probability as $\phi_{j,k,l}$. Thus the expected probability that any offer from firm j to a worker in market l is accepted (unconditional of the worker's current employer) is:

$$\phi_{j,l} = \frac{\sum_{k=1}^{J_n} \phi_{j,k,l} g_{k,l}}{\sum_{k=1}^{J_n} g_{k,l}} \tag{1.20}$$

which is the weighted average of the probability of transitions from each possible leaving firm, weighted by the leaving firms' sizes. This means that firm j will hire

$$g_{j,l,new} = \phi_{j,l} \times f_{j,l} \times (g_l - g_{j,l}) \tag{1.21}$$

new workers in market l in a given time period. This is the probability of acceptance, times the offer intensity of firm j to all other firms in occupation l times the number of workers in occupation l at other firms.

Consider the case where firm j increases its wages in occupation l by 1%. This also implies that $r_{j,l}$ increases by 1%, assuming that retirement is a percentage of compensation. Holding wages and retirement at all other firms constant, this necessarily increases $\phi_{j,l}$ and $g_{j,l.new}$. The magnitude of the increase depends on the magnitude of the wage and retirement differential between j and all other firms employing workers in occupation l, the size of those competitor firms, the size of γ_l, and the standard deviation of the individual to firm match component, σ_l^2. Denote the new number of new workers for firm j in market l as $\hat{g}_{j,l,new} = A g_{j,l,new}$ where A is some constant > 1.

Now consider the case where firm j raises its employer contribution by one percentage point. Holding wages and retirement at all other firms constant, this also necessarily increases $\phi_{j,l}$ and $g_{j,l,new}$. Again, the magnitude of the increase depends on the other terms in (1.19) and (1.20). Denote this new number of new workers as $\tilde{g}_{j,l,new} = Bg_{j,l,new}$ where B is some constant > 1.

The total net cost of increasing wages by 1% in occupation l is exactly:

$$(w_{j,l}(1 + r_j)) \times .01 \times g_{j,l,new} \tag{1.22}$$

.

The net cost for new workers of increasing retirement to occupation l by 1% is:

$$w_{j,l}(r_j + .01) \times (g_{j,l} + g_{j,l,new}) \tag{1.23}$$

However, recall that due to NDT regulations, the firm must increase the employer contribution rate for all occupations if it does so for one occupation. The total cost, inclusive of increasing contributions of existing employees, is:

$$\sum_{l=1}^{L_j} w_{j,l}(r_{j,l} + .01) \times (g_{j,l} + g_{j,l,new}) \tag{1.24}$$

When increasing retirement contributions, the firm must do so for all employees and thus pays a cost that is summed across all occupations and all existing and new employees. If the firm increases wages, the cost is limited to new workers in that occupation. For a given firm that wants to increase its hiring efficiency in occupation l, it has two options each of which have different costs. It can increase wages at a cost of

$$\frac{(w_{j,l}(1+r_j)) \times g_{j,l,new}}{Ag_{j,l,new}} \tag{1.25}$$

per new worker in market l. Or it can increase retirement at cost:

$$\frac{\sum_{l=1}^{L_j} w_{j,l}(r_{j,l}+.01) \times (g_{j,l}+g_{j,l,new})}{Bg_{j,l,new}} \tag{1.26}$$

per new worker in market l. This trade-off will vary by occupation and depending on the firm's set of competitors and its worker composition.

1.7.2 Effect on Firm Recruiting

Table XIX shows the results from the counterfactual exercise of firms increasing their wages by 1% or their contribution rate by one percentage point.[42] The exercise is partial equilibrium; we assume that each firm increases its wage, one at a time, that other firms do not respond, and that firms do not change their recruiting intensity. The top panel shows the results for a 1% increase in wage. On average, firms would improve their recruiting success (the likelihood of a job offer being accepted) by 0.16% if they increased wages by 1% in all occupations. The average cost of doing so per one new worker is about \$543. Spreading the total cost for all of the net new workers hired over all employees results in a cost of just 6 cents per employee to increase wages by 1% across the board for

[42]Note that these unit increases for the wage and retirement are roughly equivalent in dollar terms, as a one percentage point increase in the employer contribution rate has the same effect as a 1% increase in salary. For example, with a salary of \$100,000 going from a 2% to a 3% employer matching rate means going from \$2,000 to \$3,000 in employer contributions, or a \$1,000 increase. This is equivalent to a 1% increase in the \$100,000 salary.

new hires. So, the increase has relatively small effect on recruiting outcomes but it is also not very costly.

The bottom panel show the effect on recruiting success and cost of increasing the employer contribution rate by one percentage point. First, we explain the costs per new worker, which is more directly comparable to the exercise of increasing wages, as those costs apply only to new hires. A one percentage point increase in the employer contribution rate leads to an average .41% increase in the recruiting success rate. This is about 2.5 times the effect of the equivalent dollar increase in wages. The cost per one new hire is roughly similar to the wage exercise: $503. The cost is slightly lower for increasing retirement contributions because of tax considerations; firms do not have to pay payroll taxes on the retirement portion of compensation. The net cost of the new hires spread across all employees is 10 cents per worker. This is slightly higher than the wage cost per worker because increasing retirement contributions gets the firms more new workers. Thus, on a per new worker basis, increasing retirement contributions has about 2.5 times the effect on recruiting success as increasing wages at a slightly lower cost.

However, the retirement exercise explained above ignores the constraint of non-discrimination testing. As described in Section 1.2.1, firms must offer equitable retirement contributions across all workers. Thus, increasing retirement contributions only for new workers is not an option. The last two rows of Table XIX shows the cost of increasing the employer contribution rate by one percentage point for all workers, including existing workers. The cost per one new worker increases nearly 17 fold to $8,282. Spread across all existing workers, the cost of hiring the new workers brought on by this change is about $35 per worker. So, while firms might like to use retirement contributions to attract workers, as doing so has a larger effect on success, regulations make it prohibitively costly.[43]

The top panel of Appendix Table A.10 shows the equivalent increase in wages that would be required to get the same effect on recruiting success as the one percentage point increase in retirement. To increase the probability of offer acceptance by .41%, the average firm needs to increase wages by 2.7%, which has a net cost of $1,480 per new worker, or 17 cents per existing worker. Thus, if one considers the cost of increasing retirement contributions for all workers at the firm (not just the costs per new worker), increasing wages to get the same effect is still significantly cheaper, about one-eighth the cost, than increasing retirement per one new worker hired.

The bottom panel of Appendix Table A.10 shows the effect of increasing wages at roughly the same cost as increasing retirement, inclusive of the costs for existing workers. For the increase in wages to cost roughly $8,200 per new worker, the firms would need to increase wages for new hires by 14%. This will results in nearly a 2% increase in new hires, at a net cost of just 92 cents per existing worker.

There are three main takeaways from this exercise. First, retirement contributions, dollar for dollar have a about 2.5 times the effect on the recruiting success of new hires. If considering only the cost per new hire, increasing retirement contributions is a much more effective way to recruit in this sample of job-switchers. Second, the equity regulations on retirement plans make increasing retirement contributions for new hires extremely costly. When considering the cost of increasing contribution for all workers, rather than just new hires, the cost per new worker increases nearly 17 fold. Thus, although firms may want to use DC retirement contributions as a recruiting tool, doing so

[43]It is likely also challenging for firms to increase wages only for new hires. This is not for regulatory reasons, but rather because of workplace organization and bargaining dynamics (Grigsby et al. (2021), Galuscak et al. (2012)). We abstract from this issue here, but note that our estimates on the cost of increasing wages are a lower-bound.

may be prohibitively costly. Lastly, increases in wages and retirement have a relatively small effect on recruiting success, even in this specific sample of switchers.

1.7.3 Effect on Worker Valuations

The previous section discussed the effect of changing compensation policy on firm outcomes and concluded that while increasing retirement contributions is more effective for most firms, increasing wages is often cheaper, due to regulatory constraints. But what effect would these changes have on worker job valuations? This section discusses the effects of changing the different types of worker compensation on worker valuations.

This counterfactual exercise involves increasing either wage by 1% or the employer contribution rate by one percentage point, one firm at a time. Then, we calculate the new valuation a worker would get from that job, from the worker's indirect utility function (1.6). Thus, this valuation change represents a potential valuation change if the worker were to move to the firm that changed its policy. The valuation increase should thus be thought of as an increase to the worker's outside option, or valuation at potential employers, not a valuation increase they get immediately at their current job.[44] As above, the exercise is partial equilibrium in the sense that only one firm changes its policy at a time, other firms do not react, and firms change nothing else about their hiring or compensation.

Table XX shows the effect on worker valuations from either changing the employer contribution rate by one percentage point or increasing the wage by 1%, split by how much the group values retirement contributions. Starting with the top 10th percentile, or

[44]Note that changes to the retirement contribution would impact workers at their current job because of the NDT regulation. We abstract from that in this exercise.

those who value retirement the most and whose average income is \$70,000, we see that a one percentage point change in the contribution rate of a potential employer increases those workers' potential valuations by 43%. In contrast, a 1% increase in the wages of a potential employer increase their potential valuation by only 4%. For the middle 10th percentile of retirement valuations, whose average income is \$67,000, the effects are 11% and 3%, respectively. So, for the majority of the distribution, the retirement increase has a much larger (3-10 times) effect on outside options. However, for the bottom 10th percentile of retirement valuers, who are also lower income (average=\$57,000), the retirement increase increases outside options by only .69%. In contrast, the 1% wage increase increase outside options by over 10 times as much: 8%.

This exercise demonstrates how the potential valuation differences between workers are due to the heterogeneous preferences for retirement. The workers who place the most value on the employer contributions get large positive effects to their outside options if firms offer higher contribution rates; this thus helps firms to recruit these workers more so than an increase in wages would. However, because the employer contribution rate must change for everybody, due to NDT regulations, this has a negative spillover on lower-income workers. Rather than getting a wage increase, which would provide them with a greater valuation increase, they would be offered higher retirement contributions, which do little to affect their valuations.

1.8 Conclusion

Defined contribution (DC) retirement accounts, such as 401(k) and 403(b) accounts, and employer contributions to these accounts are an increasingly important part of both household wealth accumulation and corporate labor costs. In this paper, we study both how workers value these plans as well their effect on labor market outcomes. We show across three distinct methods that most workers exhibit willingness to pay for both the intensive and extensive margin of DC retirement plans.

Our first finding, focusing on plausibly exogenous variation in firm-level retirement plans and posted wages, is that the average worker values a dollar of retirement contributions nearly twice as much as a dollar of wage when choosing between similar jobs. This ratio increases to 3:1 when focusing on high-income or older workers. The variation in wages and retirement plans is driven by changes induced by non-discrimination testing and firms that set common wages nationally for each given job type.

Second, we design and conduct an online experimental survey that uses hypothetical choices to measure how much potential workers are willing to give up in total compensation in order to get a 401(k) plan and to get higher employer contributions. Workers in the survey are willing to give up 3.4% of total pay to get a job with a 401(k) and 1.4% of wages for each one percentage point increase in the employer match. The majority of participants exhibit willingness to pay for both the intensive and extensive margin of the DC plans.

Lastly, we build and estimate an on-the-job search model which estimates worker valuations and allows us to conduct counterfactual analyses on firm compensation policy. Consistent with the instrumental variable results and the direct estimation of valuations in the survey, we show that the majority of workers (75%) are willing to give up some total pay to get a higher employer contribution. The valuation of retirement contributions increases with income, as in the revealed preference results. Moreover, 90% of workers are willing to give up some of total pay just get a 401(k). The results suggest increasing retirement contributions has a much larger (2.5 times) effect on firm recruiting success than increasing wages. However, doing so disproportionately benefits higher income workers who place a higher value on these contributions.

In future work, using the framework and data developed here, we plan to study the impact that non-discrimination testing has on worker welfare. Given the higher valuation placed on retirement benefits for higher-income workers, firms are incentivized to design plans to attract these workers. This could have spill-over effects on the wages of lower-income workers who don't value the retirement. Another natural follow-up is to study how retirement benefits affect retention. Many DC retirement contributions are vested and require workers to stay at the firm for some amount of time to reap the full benefits of employer matching. This paper focused on hiring, but the employer contributions and vesting schedules likely have an affect on retention; understanding this would improve our understanding of how DC plans affect equilibrium labor market outcomes. We would also like to study further the mechanism that leads to the high valuations of retirement benefits that we find. There are three prevailing theories to explain the high valuation 1) the signaling power of retirement benefits 2) retirement accounts are valuable as a commitment device 3) the valuation can be explained by a combination of taxes and discount rates. Combining both empirical methods and experiments, we plan to look for evidence of these theories in future work.

Figures

Figure 1-1: NDT Failure over Time

(a) % of DC Plans with a Corrective Distribution or NDT Failure

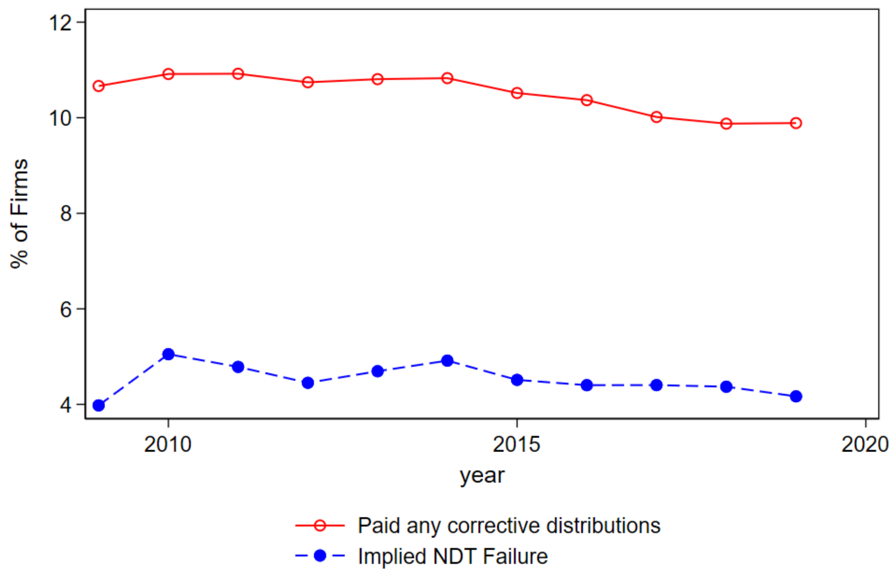

(b) Estimation Sample: % of DC Plans with a Corrective Distribution or NDT Failure

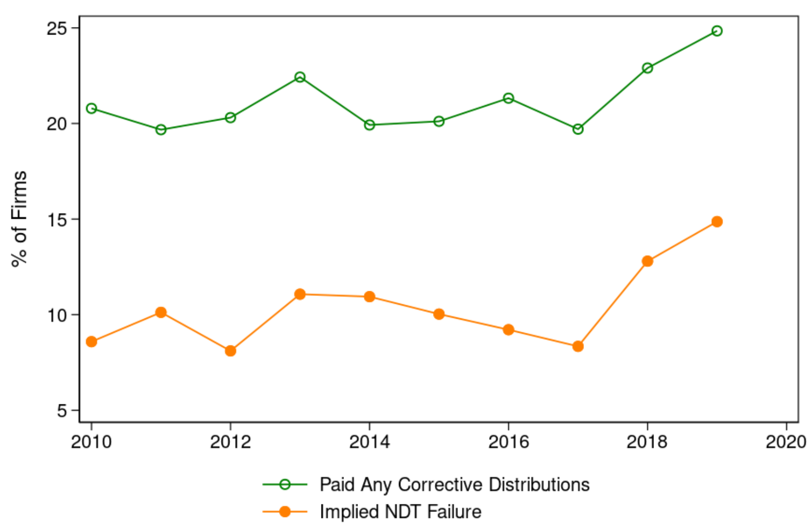

Notes: Authors' calculations from the Form-5500 data, 2010-2019. Includes only DC plans. The top panel shows the entire universe of form-5500 filers. The bottom panel shows only firms in our estimation sample.

Figure 1-2: Employer Retirement Contributions Before and After NDT Failure: Event Study Plots

(a) Employer's Ratio of all Contributions

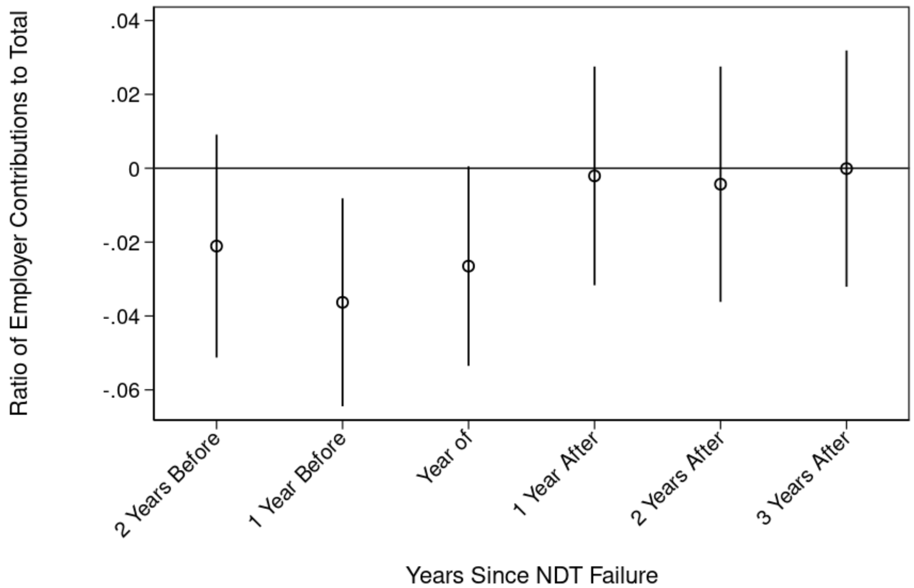

(b) Employer Contribution Rate

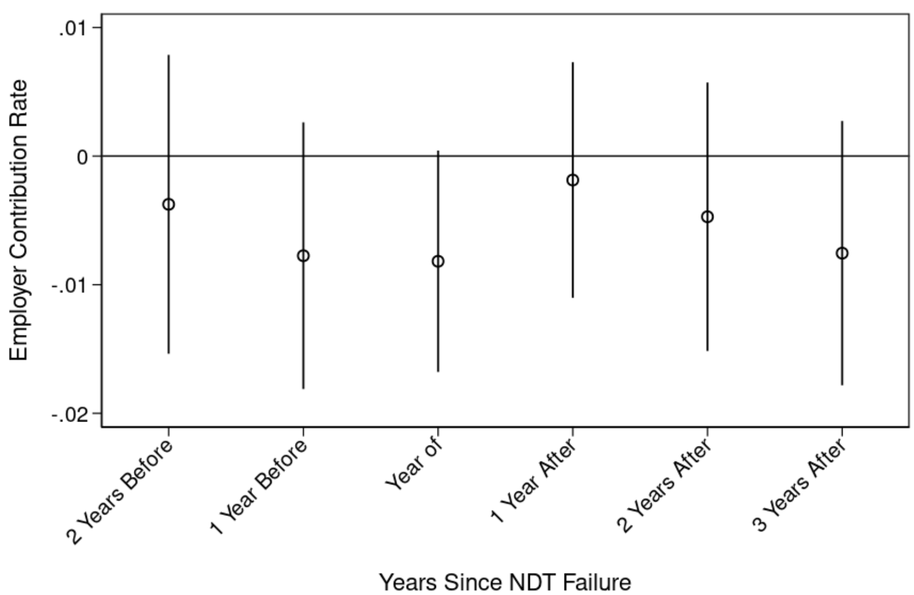

Notes: Control group is firms that do not fail NDT with the median year taken as year zero. Treated firms are those that fail in year zero. Regressions include industry by year fixed effects and controls for log number of employees and log dollars in assets in the retirement plan.Robust standard errors are clustered at the firm level. Confidence intervals are at the 95% significance level.

Figure 1-3: Average Annual Salary Before and After NDT Failure

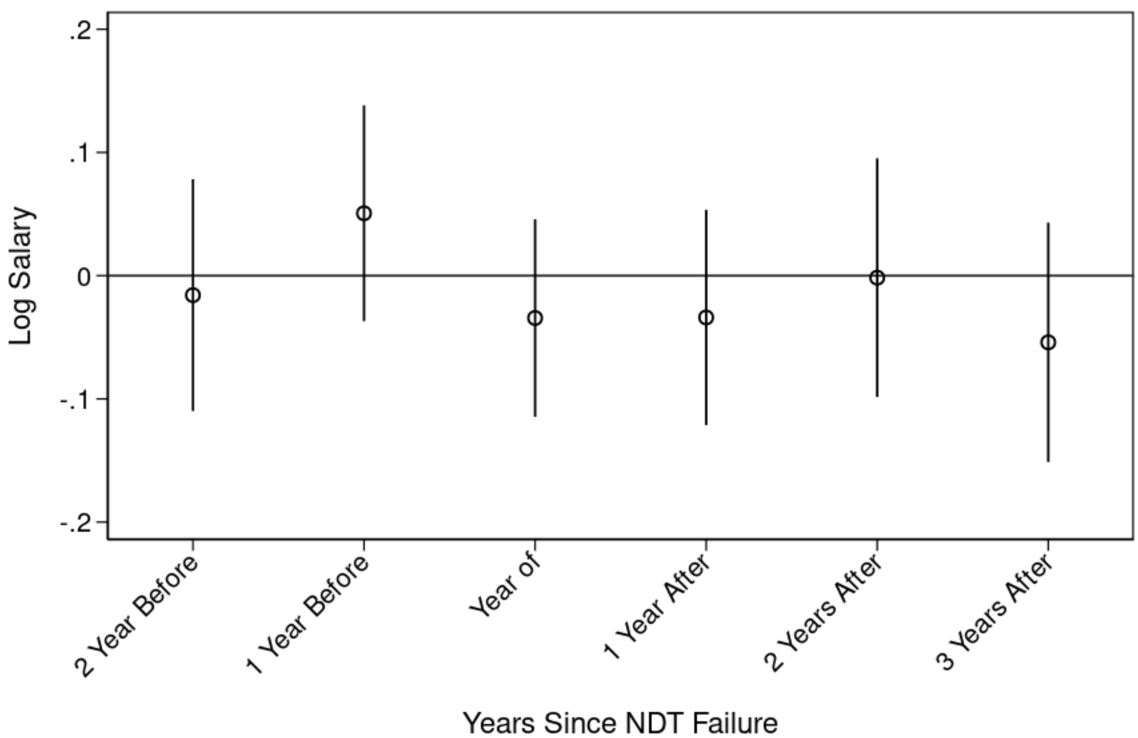

Notes: Control group is firms that do not fail NDT with the median year taken as year zero. Treated firms are those that fail in year zero. Regressions include industry by year fixed effects and controls for log number of employees and log dollars in assets in the retirement plan.Robust standard errors are clustered at the firm level. Confidence intervals are at the 95% significance level.

Figure 1-4: Percentage of Firms that are National Wage Setters

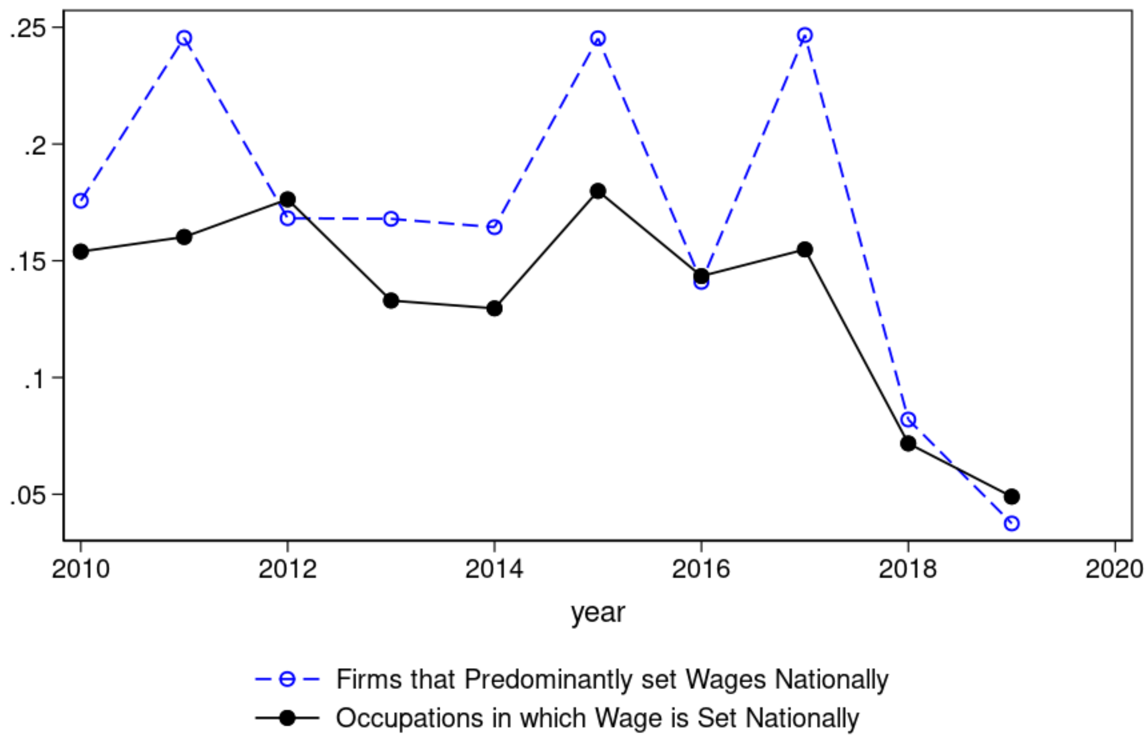

Notes: This figure shows the percentage of firms that set wages nationally in the matched sample of Lightcast wages, resumes and Form 5500. "% of Firms that Predominantly Set Wages Nationally" refers to firms that set at least 75% of their occupations across geographies at the same level. "% of Occupations in which Wages are set Nationally" refers to the any occupation for which a firm sets wages identically across geographies.

Figure 1-5: Comparison of IV Estimation Samples

(a) Geographical Distribution

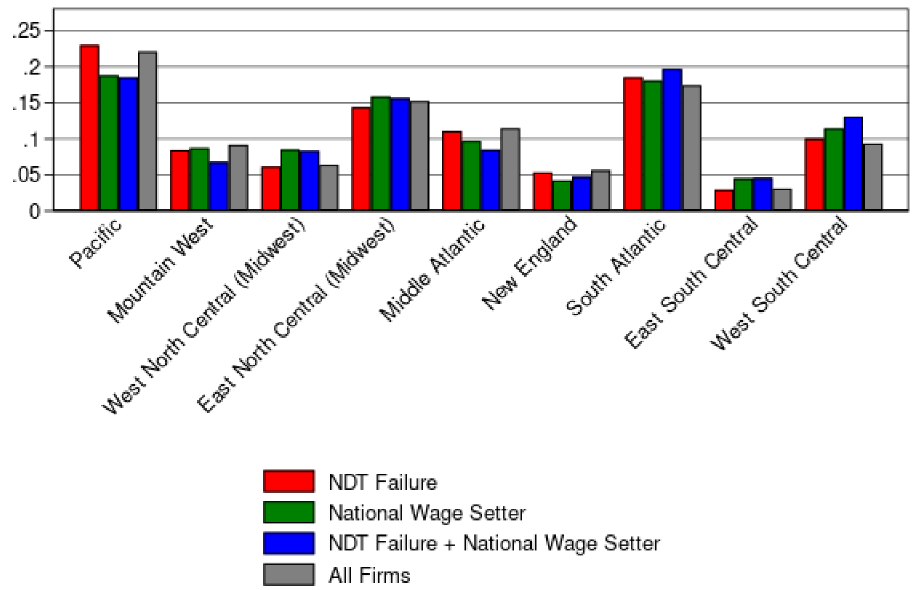

(b) Industry Distribution

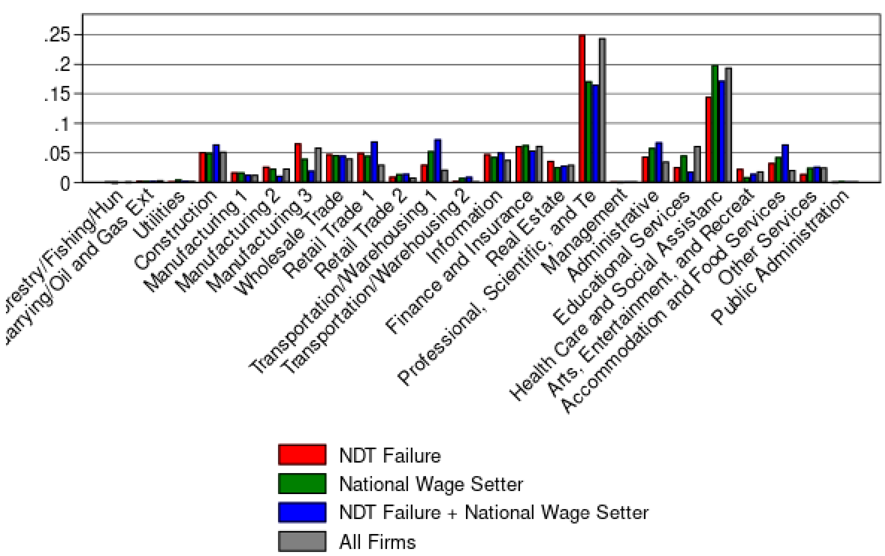

Figure 1-5: Comparison of IV Estimation Samples (continued)

(c) Occupation Distribution

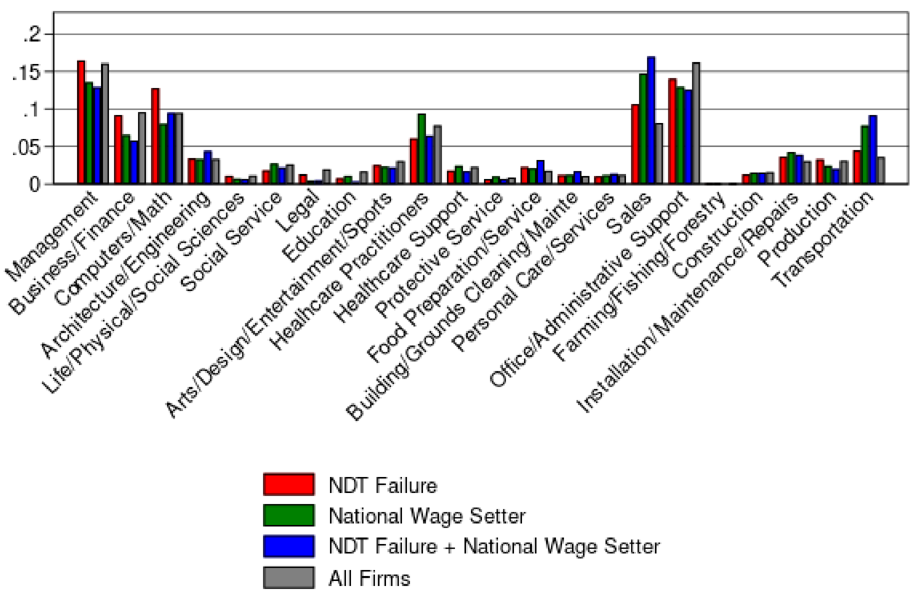

Notes: These figures compare the geographic, industry, and occupational distribution of firm in our estimation sample, split by firms affected by each of the instrumental variables.

Figure 1-6: Matched Sample Characteristics versus BLS

(a) Industry by Number of Firms

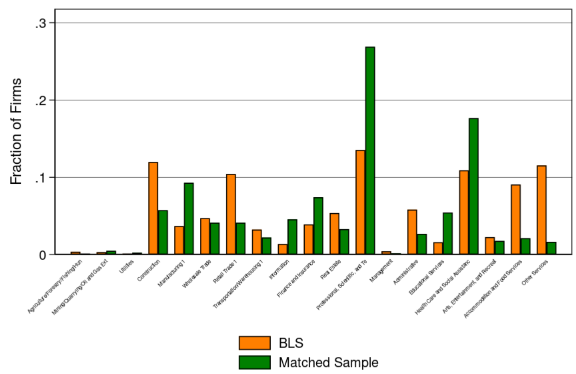

(b) Firm Size by Number of Firms

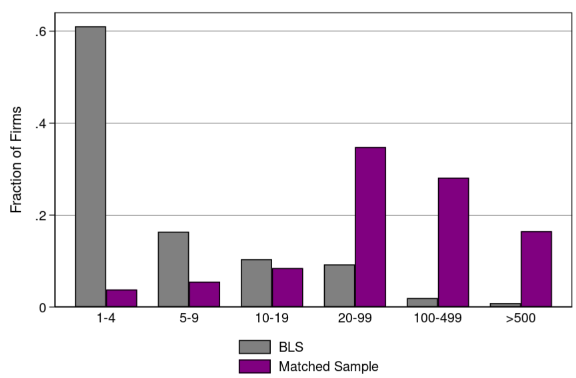

(c) Occupation by Number of Employees

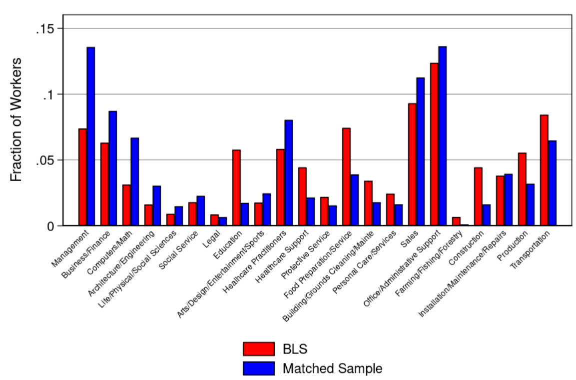

Notes: These figures shows the distribution of firms and employment by industry, firm size and occupation in the main sample versus the BLS. The main sample is all firms in the merged sample of Lightcast posted wage, Lightcast, resumes, and Form 5500. The BLS sample is from the Bureau of Labor Statistics database.

Figure 1-7: Example Survey Question

(a) 401(k) with a 3% match versus 401(k) with no match

Scenario 5

	Job 1	Job 2
Annual Earnings when working full time:	$51,500	$49,000
Retirement benefits:	Company sponsored 401(k) (no employer match)	Company sponsored 401(k) with matching of 100% up to 3% (for a total possible match of $1,470 per year)
Healthcare benefits:	No	No
Vacation:	10 days of paid vacation	10 days of paid vacation
Work flexibility:	2 days of remote work per week	2 days of remote work per week

Note that with Job 2, your pre-tax take home pay will be $2,500 lower annually, or approximately $208 lower per month.

Select which you would be more likely to accept.

Job 1

Job 2

(b) 401(k) with no match versus no 401(k)

Scenario 3

	Job 1	Job 2
Annual Earnings when working full time:	$49,500	$51,500
Retirement benefits:	Company sponsored 401(k) (no employer match)	None
Healthcare benefits:	Yes	Yes
Vacation:	20 days of paid vacation	20 days of paid vacation
Work flexibility:	2 days of remote work per week	2 days of remote work per week

Note that with Job 1, your pre-tax take home pay will be $2,000 lower annually, or approximately $167 lower per month.

Select which you would be more likely to accept.

Job 1

Job 2

Notes: These figures show example question from the survey.

Figure 1-8: Demographic Characteristics of Survey Participants

(a) Age

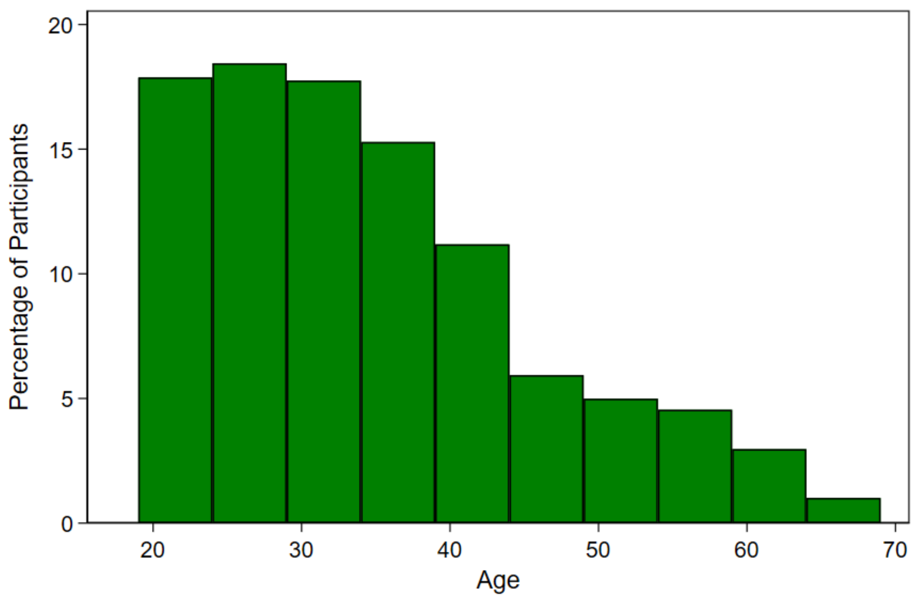

(b) Gender

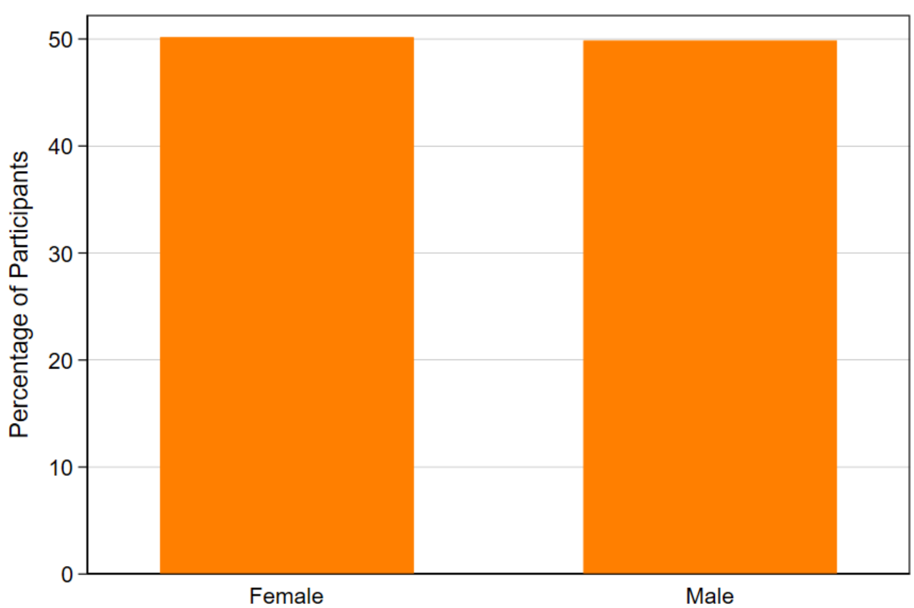

Figure 1-8: Demographic Characteristics of Survey Participants (continued)

(c) Education

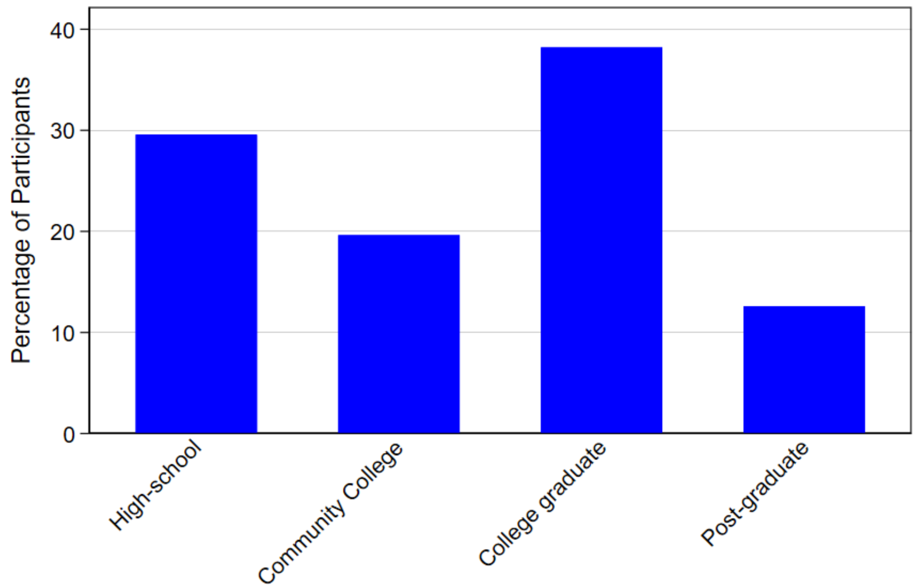

(d) Ethnicity

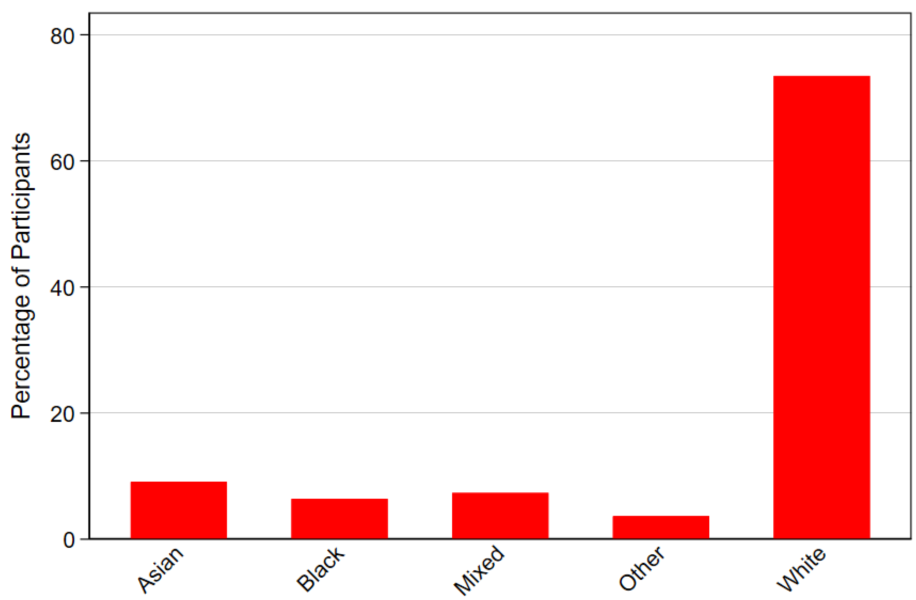

Notes: These figures show characteristics of the 1,600 survey participants recruited on Prolific. All participants are living in the U.S., speak English as a first language and are either working full-time or seeking work.

Figure 1-9: Willingness to Pay for Retirement Benefits

(a) Intensive Margin: Has 401(k) with 3% Match versus 401(k) with no match

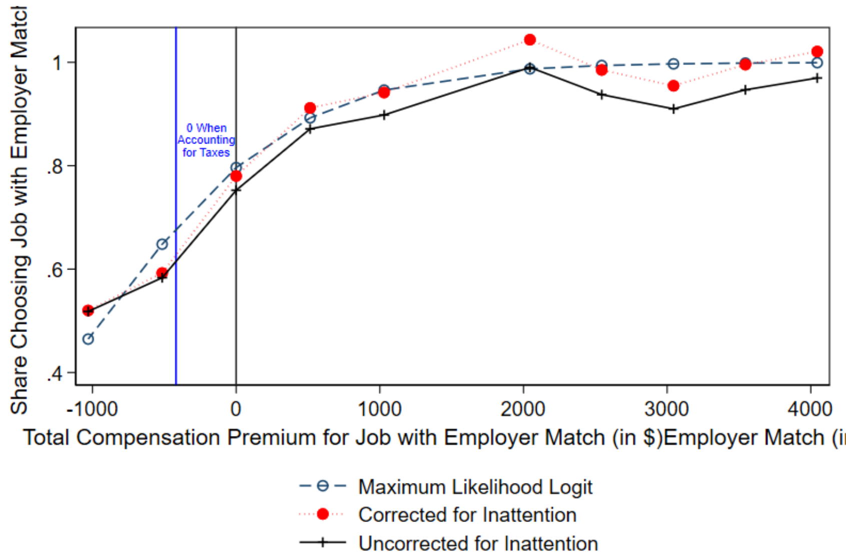

(b) Extensive Margin: Has 401(k) with no match versus no 401(k)

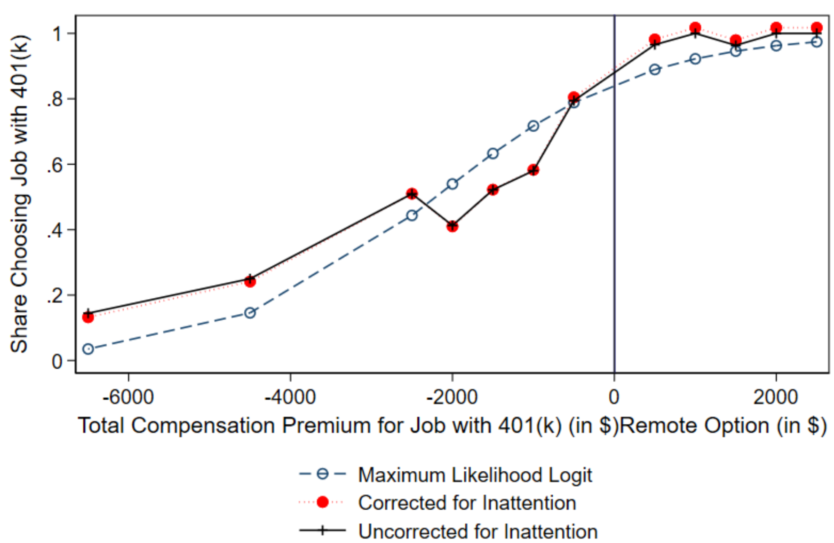

Notes: These plots show the fraction of participants who chose the job with the better retirement benefit plotted against the difference in total compensation. The total compensation gap is the total compensation for the job with the 401(k) in panel (b) and the job with the match in panel (a) minus the the total compensation for the job with no retirement in panel (b) or the job with no match in panel (a). Based on a a survey with 1,629 participants. The black line shows the raw data. The red dotted line shows the distribution corrected for inattention. The blue dashed line shows the distribution estimated by maximum likelihood.

Figure 1-10: Retirement Weights versus Salaries

(a) Intensive Margin of DC Plan

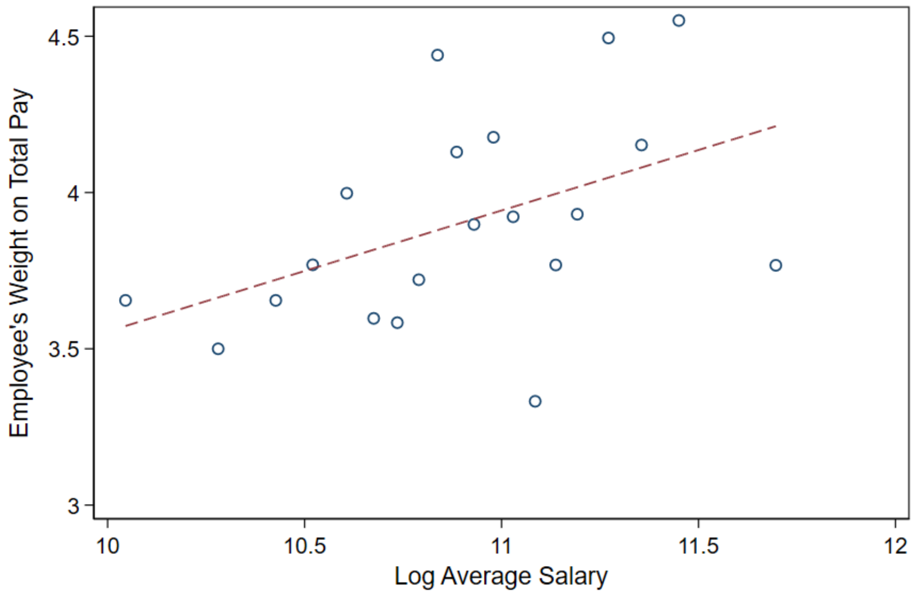

(b) Extensive Margin of DC Plan

Notes: These figures show binscatters of the average log salary in an industry by occupation group against a) the estimated weight on total pay $(1 - \gamma_l)$ and b) the weight on the extensive margin of having a retirement plan (β_l). The binscatters control for year fixed effects and are unweighted.

Figure 1-11: Firm-level Amenities around NDT Failure

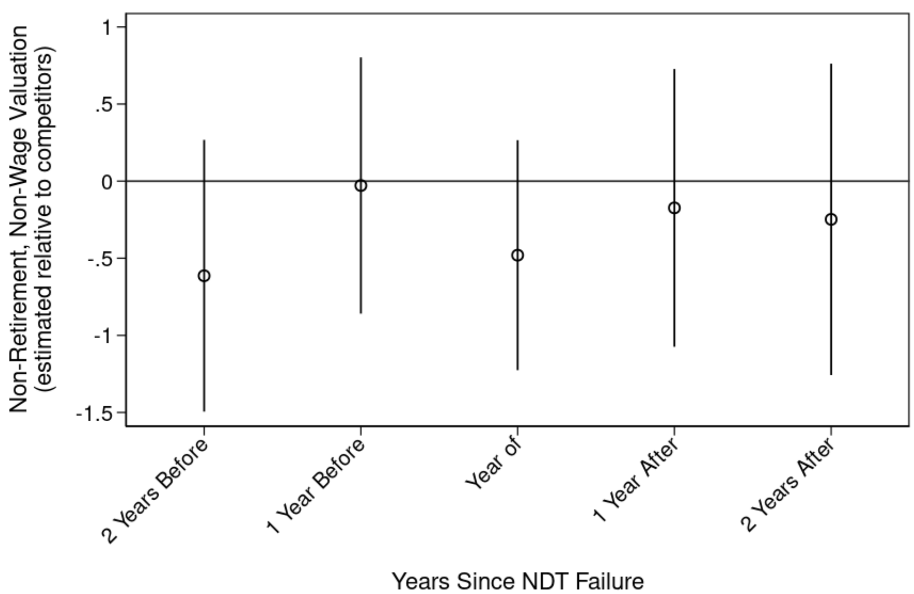

Notes: This figure shows difference-in-difference results for firm-level amenity valuations around NDT failure. The control group is firms that do not fail NDT with the median year taken as year zero. Treated firms are those that fail in year zero. Regressions include industry by year fixed effects and controls for log number of employees and log dollars in assets in the retirement plan. Robust standard errors are clustered at the firm level. Confidence intervals are at the 95% significance level.

Tables

Table I: Difference in Difference Results: Effect of NDT Failure on Plan Features

	(1) Employer Ratio of All Contributions	(2) Employer Contribution Rate	(3) Has Autoenrollment
Time	-0.00987***	-0.00295**	0.0145***
	(0.00244)	(0.000949)	(0.00401)
Treated	-0.0735***	-0.0146***	0.106***
	(0.00343)	(0.00120)	(0.00965)
Time x Treated	0.0240***	0.00407*	-0.0235
	(0.00604)	(0.00204)	(0.0185)
Observations	28043	28350	34952
R^2	0.148	0.161	.152

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

Notes: This regression shows the difference in difference results for firms that failed NDT tests compared to firms that did not. The sample period is 2 years before, the year of, and 3 years after the NDT Failure. The year of is set to the median sample year for non-failing firms. Regressions control for year by industry fixed effects, log number of employees and log dollar assets in the retirement plan. Only firms with DC plans are included. Standard errors, clustered at the firm level, are in parentheses

Table II: Difference in Difference Results: Effect of NDT Failure on Non-Retirement Firm Characteristics

	(1) Log Salary	(2) Log # of Postings	(3) Log # of New Hires	(4) Log Years of Experience Required	(5) Log Dollars Paid in Healthcare Benefits per Person
Time	0.00946	0.480***	0.256***	0.0219	0.314**
	(0.00616)	(0.0200)	(0.0153)	(0.0340)	(0.106)
Treated	0.0362***	0.145***	0.0521*	0.0935*	0.166
	(0.00854)	(0.0259)	(0.0217)	(0.0470)	(0.160)
Time x Treated	-0.00921	-0.0133	0.101*	-0.0335	0.398
	(0.0179)	(0.0550)	(0.0455)	(0.0910)	(0.334)
Observations	31678	31678	31678	21954	28239
R^2	0.158	0.125	0.272	0.100	.039

Standard errors, clustered at the firm level, are in parentheses

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

Notes: This regression shows the difference in difference results for firms that failed NDT tests compared to firms that did not. The sample period is 2 years before, the year of, and 3 years after the NDT Failure. The year of is set to the median sample year for non-failing firms. Regressions control for year by industry fixed effects, log number of employees and log dollar assets in the retirement plan. Only firms with DC plans are included.

Table III: Wages of National Wage Setters

	Log Average Salary		
	Firm level	Market Level	Market Level, Multi-establishment firms
National Wage Setter	0.0228**	0.0224***	0.0171***
	(0.00779)	(0.00170)	(0.00175)
Observations	47,348	468,011	398,544
R^2	0.151	0.452	0.472

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

Notes: The dependent variable is log average salary at the firm level (column 1) and the occupation by CBSA level (columns 2 and 3). The independent variable is an indicator equal to one if the firm sets at least 75% of its wages nationally and 0 otherwise. Additional controls include log employment, industry by year fixed effects (column 1) and industry and occupation by CBSA by year fixed effects (column 2). Robust standard errors are clustered at the firm level.

Table IV: Characteristics of National Wage Setters

	(1) Log Employment	(2) Employer Contribution Rate	(3) Log # of New Hires	(4) Turnover	(5) Log $ per Person Spent on Healthcare
National Wage Setter	-0.254***	-0.001	-0.238***	0.001	-0.204
	(0.041)	(0.001)	(0.025)	(0.002)	(.167)
Observations	21611	19511	22262	21901	19981
R^2	0.093	0.050	0.137	0.407	0.0240

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

Notes: Regressions are at the firm level. The independent variable is an indicator equal to one if the firm sets at least 75% of its wages nationally and 0 otherwise. Additional controls include industry by year fixed effects. Only firms with at least 2 establishments are included. Robust standard errors, clustered at the firm level, are in parentheses.

Table V: Comparison of IV Samples

	NDT Failure	National Wage Setter	NDT Failure + National Wage Setter	All Firms
Employment and Hiring Variables				
# of Job Postings per Year	31.01 (6.00)	166.50 (30.00)	122.17 (33.00)	33.91 (6.00)
# of New Hires per Year	5.38 (2.00)	12.33 (2.00)	11.58 (3.00)	4.72 (1.00)
Hire Success	0.09 (0.00)	0.08 (0.00)	0.07 (0.00)	0.11 (0.00)
Annual Salary	57927.65 (50000.00)	55208.22 (47465.64)	52886.59 (44905.46)	56487.93 (49000.00)
Hourly Wage	27.88 (24.04)	26.61 (22.82)	25.44 (21.59)	27.21 (23.56)
Turnover - Form 5500 Employment	0.24 (0.17)	0.23 (0.14)	0.22 (0.12)	0.24 (0.16)
Tenure (months)	29.64 (26.21)	29.77 (27.15)	29.39 (27.40)	29.59 (25.88)
Total Employees	376.51 (90.00)	1155.29 (103.00)	775.88 (138.50)	397.88 (66.00)
Retirement Plan Variables				
Total Plan Assets (Millions of $s)	30.07 (3.61)	113.27 (3.58)	70.22 (4.81)	37.58 (2.63)
Employer Contribution Rate (%)	3.30 (2.16)	4.34 (2.87)	2.73 (1.88)	5.12 (3.38)
Ratio of Employer Contribution to Total	0.22 (0.22)	0.28 (0.29)	0.21 (0.21)	0.30 (0.30)
Employer Contributions per Participant ($)	1848.11 (1017.00)	985.29 (1427.80)	2110.26 (845.68)	2943.37 (1695.64)
Total # of Plan Participants	777.23 (165.00)	1876.81 (182.00)	1854.79 (332.50)	608.04 (92.00)

Table V: Comparison of IV Samples (continued)

	NDT Failure	National Wage Setter	NDT Failure + National Wage Setter	All Firms
Other Variables				
% of Employees Captured in Resume Data	44.65	41.36	40.97	41.97
	(39.72)	(35.71)	(35.41)	(35.84)
% of Firms that set Wages Nationally	0.35	0.64	0.63	0.34
	(0.10)	(0.81)	(0.80)	(0.07)
% of Firms with $\geq$ 2 Establishments	0.48	1.00	1.00	0.41
	(0.00)	(1.00)	(1.00)	(0.00)
% of Firms with $\geq$ 4 Establishments	0.21	0.66	0.72	0.15
	(0.00)	(1.00)	(1.00)	(0.00)
# of Unique CBSAs per Firm	7.74	18.30	24.04	4.71
	(2.00)	(6.00)	(8.00)	(1.00)
# of Unique Occupations per Firm	9.52	19.97	19.98	7.44
	(4.00)	(6.00)	(9.00)	(2.00)
# of Unique Years per Firm	2.50	2.73	3.40	1.93
	(2.00)	(2.00)	(2.00)	(1.00)
# of Transitions	144499	264478	66938	555991
# of Unique Jobs	124438	227237	55786	486081
# of Firms	4328	2517	588	24554

Notes: This table shows summary statistics for them matched sample of firms that have job postings data in Lightcast, resume data in Lightcast, and Form 5500 data on health and retirement plans. Summary statistics are calculated at the firm level and averaged over all years the firm appears in the sample. Retirement summary statistics are conditional on having a retirement plan. Sample are split into groups affected by each of the instrumental variables. Medians are in parentheses.

Table VI: Distribution of Wages in the Matched Sample versus BLS

	Mean	P10	P25	P50	P75	P90
Matched Sample	49,554	21,840	27,560	40,951	60,000	90,000
BLS	50,629	24,000	30.970	42,880	61,690	85,570

Notes: This table shows the mean and percentiles for the annual salary in the matched sample of Lightcast job postings, Lightcast resumes, and Form 5500 versus from the Occupation Employment Statistics database for 2015.

Table VII: Summary Statistics for Matched Sample of Wages, Resumes, and Retirement Plans

	Mean	Median
Employment and Hiring Variables		
# of Job Postings per Year	33.91	6.00
# of New Hires per Year	4.72	1.00
Hire Success	0.11	0.00
Annual Salary	56488	49,900
Turnover - Form 5500 Employment	0.24	0.16
Tenure (months)	29.59	25.88
Total Employees	397.88	66.00
Retirement Plan Variables		
Total Plan Assets (Millions of $s)	37.58	2.63
Employer Contribution Rate (%)	5.12	3.38
Ratio of Employer Contribution to Total	0.30	0.30
Employer Contributions per Participant ($)	2943.37	1695.64
Total # of Plan Participants	608.04	92.00
Other Variables		
% of Employees Captured in Resume Data	41.97	35.85
# of Unique CBSAs per Firm	4.71	1.00
# of Unique Occupations per Firm	7.44	2.00
# of Unique Years per Firm	1.93	1.00
# of Transitions	555991	
# of Unique Jobs	486081	
# of Firms	24554	

Notes: This table shows summary statistics (means and medians) for the matched sample of firms that have both job postings data in Lightcast, resume data in Lightcast, and Form 5500 data on health and retirement plans. Summary statistics are calculated at the firm level and averaged over all years the firm appears in the sample. Retirement summary statistics are conditional on having a match in the Form 5500 data. Medians are in parentheses.

Table VIII: IV Results: Effect of Wages and Retirement on Recruiting Success

	(1)	(2)	(3)	(4)
	First Stage		OLS	IV
	Log Salary	Employer Contribution Rate	Successfully Hire	
National Wage Setting Instrument	0.477***	-0.161***		
	(0.0116)	(0.0228)		
After NDT Failure	-0.00967	0.330***		
	(0.00800)	(0.0274)		
Log Salary			0.0141***	0.0142*
			(0.00419)	(0.0079)
Employer Contribution Rate			-0.00198	0.0271**
			(0.00131)	(0.0124)
Observations	124533	124533	124533	124533
Cragg-Donald Wald F-Statistic				136.485
Kleinbergen-Paap rk LM Statistic				190.852
Additional Controls:				
Year by Log Employment	Y	Y	Y	Y
Year by Log Assets	Y	Y	Y	Y
Health Plan Indicator	Y	Y	Y	Y
Fixed Effects:				
Firm by CBSA by Occupation	Y	Y	Y	Y
Year	Y	Y	Y	Y

Standard errors, clustered at the firm-level, are in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is a dummy equal to one if the firm successfully filled a position in a given occupation and CBSA in the year it was posted. The baseline average hire success rate is 11%. All regressions include firm by CBSA by occupation (5-digit SOC code) fixed effects and year fixed effects. Regressions are weighted by occupation by CBSA employment size. Only firms with at least two establishments and a defined contribution retirement plan are included.

Table IX: IV Results by Income Group: Effect of Wages and Retirement on Recruiting Success

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Low-income Occupations				High-income Occupations			
	First Stage		OLS	IV	First Stage		OLS	IV
	Log Salary	Employer Contribution Rate	Successfully Hire		Log Salary	Employer Contribution Rate	Successfully Hire	
National Wage Setting Instrument	0.429***	-0.149***			0.509***	-0.173***		
	(0.0156)	(0.0291)			(0.0164)	(0.0330)		
After NDT Failure	-0.0035	0.388***			0.000732	0.189***		
	(0.00924)	(0.0461)			(0.0132)	(0.0312)		
Log Salary			0.00906	0.0187**			0.0172***	0.0203*
			(0.00577)	(0.0086)			(0.00583)	(0.0113)
Employer Contribution Rate			-0.00383**	0.0160*			-0.000173	0.0603**
			(0.00156)	(0.0097)			(0.00220)	(0.0279)
Observations	65022	65022	65022	65022	59511	59511	59511	59511
Cragg-Donald Wald F-Statistic				140.673				22.997
Kleinbergen-Paap rk LM Statistic				141.212				52.433
Fixed Effects:								
Firm by CBSA by Occupation	Y	Y	Y	Y	Y	Y	Y	Y
Year	Y	Y	Y	Y	Y	Y	Y	Y

Standard errors, clustered at the firm-level, are in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is a dummy equal to one if the firm successfully filled a position in a given occupation and CBSA in the year it was posted. The baseline average hire success rate is 8% for the low-income group and 14% for the high-income group. All regressions include firm by CBSA by occupation (5-digit SOC code), and year fixed effects, and controls for log employment by year, log assets by year, and a health plan indicator. Regressions are weighted by occupation by CBSA employment size. Only firms with at least two establishments and a defined contribution retirement plan are included. Low (high) income are occupations below (above) the median income in that year.

Table X: IV Results by Age: Effect of Wages and Retirement on Recruiting Success

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	High-age Occupations				Low-age Occupations			
	First Stage		OLS	IV	First Stage		OLS	IV
	Log Salary	Employer Contribution Rate	Successfully Hire		Log Salary	Employer Contribution Rate	Successfully Hire	
National Wage Setting Instrument	0.469***	-0.125***			0.458***	-0.158***		
	(0.0193)	(0.0446)			(0.0206)	(0.0395)		
After NDT Failure	0.00864	0.296***			-0.0042***	0.326***		
	(0.0112)	(0.0472)			(0.0158)	(0.0594)		
Log Salary			0.00810	0.0236*			0.00989	0.0205*
			(0.00791)	(0.0122)			(0.00768)	(0.0111)
Employer Contribution Rate			-0.00267	0.0720*			-0.00277	0.00638
			(0.00220)	(0.0416)			(0.00269)	(0.0074)
Observations	46159	46159	46159	46159	33009	33009	33009	33009
Cragg-Donald Wald F-Statistic				43.111				30.937
Kleinbergen-Paap rk LM Statistic				55.928				38.738
Fixed Effects:								
Firm by CBSA by Occupation	Y	Y	Y	Y	Y	Y	Y	Y
Year	Y	Y	Y	Y	Y	Y	Y	Y

Standard errors, clustered at the firm-level, are in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is a dummy equal to one if the firm successfully filled a position in a given occupation and CBSA in the year it was posted. The baseline average hire success rate is 10% for the low age occupations and 11% for the high age occupations. All regressions include firm by CBSA by occupation (5-digit SOC code), year fixed effects, and controls for log employment by year, log assets by year, and a health plan indicator. Regressions are weighted by occupation by CBSA employment size. Only firms with at least two establishments and a defined contribution retirement plan are included. Low (high) age occupations are those for which the median age, as measured by the BLS's Occupational Employment Statistics, is below (above) 42 years.

Table XI: IV Results by Gender: Effect of Wages and Retirement on Recruiting Success

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Male-dominated Occupations				Female-dominated Occupations			
	Log Salary	Employer Contribution Rate	Successfully Hire		Log Salary	Employer Contribution Rate	Successfully Hire	
National Wage Setting Instrument	0.502***	-0.181***			0.423***	-0.138***		
	(0.0168)	(0.0304)			(0.0190)	(0.0410)		
After NDT Failure	-0.0248*	0.230***			-0.00734	0.340***		
	(0.0136)	(0.0327)			(0.0115)	(0.0531)		
Log Salary			0.00351	0.0140*			0.00443	0.0198
			(0.00633)	(0.0083)			(0.00825)	(0.0122)
Employer Contribution Rate			-0.00452**	0.0370*			-0.00355	0.0315*
			(0.00180)	(0.0210)			(0.00223)	(0.0189)
Observations	44987	44987	44987	44987	48223	48223	48223	48223
Cragg-Donald Wald F-Statistic				27.462				49.845
Kleinbergen-Paap rk LM Statistic				63.435				54.683
Fixed Effects:								
Firm by CBSA by Occupation	Y	Y	Y	Y	Y	Y	Y	Y
Year	Y	Y	Y	Y	Y	Y	Y	Y

Standard errors, clustered at the firm-level, are in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is a dummy equal to one if the firm successfully filled a position in a given occupation and CBSA in the year it was posted. The baseline average hire success rate is 9% for the male-dominated occupations and 12% for the female-dominated occupations. All regressions include firm by CBSA by occupation (5-digit SOC code), year fixed effects and controls for log employment by year, log assets by year, and a health plan indicator. Regressions are weighted by occupation by CBSA employment size. Only firms with at least two establishments and a defined contribution retirement plan are included. Male (female) dominate occupations are those for which the estimated share of male (female) workers is greater than 50%, as measured by the BLS's Occupational Employment Statistics.

Table XII: Willingness to Pay for Retirement Benefits: Survey Evidence

	Intensive Margin	**Extensive Margin**
	401(k) with no match versus 401(k) with a 3% match	No 401(k) versus 401(k) with no match
Fraction that chose job with better retirement	0.846 (1004)	0.669 (1628)
Conditional on:		
Better retirement job has lower wage	0.711 (478)	0.491 (1036)
Better retirement job has higher wage	0.968 (526)	0.980 (592)
Better retirement job has lower total comp	0.555 (191)	0.491 (1036)
Better retirement job has higher total comp	0.942 (724)	0.980 (592)
Better retirement job has lower total comp, net of taxes	0.555 (1036)	0.491 (191)
Conditional on choosing job with better retirement and lower total comp:		
WTP	2226	1775
	(106)	(509)
WTP as a percent of wages	4.3 (106)	3.4 (509)
WTP in total comp	748 (106)	1775 (509)
WTP as a percent of total comp	1.5 (106)	3.4 (509)
WTP in total comp, net of taxes	258 (106)	1385 (509)
Cost of retirement plan to employer fixed/set-up costs)	1478 (106)	0 (509)

Notes: This table shows summary statistics for the survey conditions that test willingness to pay for the intensive and extensive margin of retirement benefits. The extensive margin condition asks participants to choose between similar jobs, one of which offers a 401(k) with no match and the other offers no 401(k). The intensive margin condition has participants to choose between similar jobs, one or which offers a 401(k) with a 3% match an the other offers a 401(k) with no match. Numbers in parentheses show the number of participants who answered for the relevant condition.

Table XIII: Maximum Likelihood Estimates of Willingness to Pay in Survey

Treatment	Mean	SD	P25	P50	P75
Willingness to Pay for the Intensive Margin of Employer Contributions					
401(k) with 3% match versus 401(k) with no match	909.17	1400.26	60.61	909.17	1757.72
	(122.64)	(154.02)	(94.29)	(122.64)	(196.51)
Willingness to Pay for the Extensive Margin of Retirement Plans					
401(k) with no match versus no 401(k)	2267.56	2383.31	823.29	2267.56	3711.84
	(129.99)	(173.83)	(82.303)	(129.99)	(221.84)

Notes: This table shows the distribution of the willingness to pay estimates from the survey. Estimates are from an inattention-corrected maximum likelihood logit model using data from the experiment. Bootstrapped standard errors based on 1000 samples are in parentheses.

Table XIV: Model Parameters

Parameter	Definition	Source
$w_{i,j,l}$	Wage (in dollars) for individual i at firm j in occupation and CBSA l	Data: Lightcast Job Postings
$\mathbb{1}_{h_j}$	An indicator equal to one if firm j has a healthcare plan	Data: Form 5500
α_l	The weight on the dollar value of healthcare benefits	Calibrated (Baicker and Chandra (2006), Miller (2004))
$f_{j,l}^{NE}$	The recruiting intensity of firm j in CBSA and occupation l to unemployed workers	Data: Lightcast Resumes
$g_{j,l}$	The number of workers at firm j in occupation and CBSA l	Data: Lightcast Resumes
$r_{j,l}$	The retirement benefit (in dollars per person) given to workers at firm j in occupation and CBSA l	Data: Form 5500
$\mathbb{1}_{r_j}$	An indicator equal to one if firm j has a retirement plan	Data: Form 5500
$\hat{\gamma}_l$	The weight employees place on total compensation, relative to a weight of one on wages	**Estimated**
$\hat{\beta}_l$	The weight employees place the presence of a retirement plan	**Estimated**
$\widehat{\Delta ln(a_l)}$	The residual portion of job transitions that is not explained by wages, healthcare, retirement or industry, occupation or CBSA differences - referred to as "amenities"	**Estimated**
σ_l	The standard deviation of the idiosyncratic match value for workers in market l	**Estimated**

Notes: This table defines all of the model parameters and how they are measured or estimated.

Table XV: Summary Statistics of Model Estimation Sample

	mean	p50	p10	p90
Number of Employees	3795.95	856.00	61.00	8278.00
Average Salary	55929.12	50000.00	28500.00	90496.82
# of Job Postings per Year	209.29	14.00	2.00	322.00
Has a Retirement Plan	0.58	1.00	0.00	1.00
Employer Contribution $ per Person	2967.00	2063.91	0.00	6604.01
Participation Rate in Retirement Plan	0.75	0.86	0.43	0.93
Employer Contribution Rate	0.06	0.04	0.00	0.14
Has a Healthcare Plan	0.87	1.00	0.00	1.00
Number of Transitions	34,656			
Number of Industry by Occupation Groups	308			
Number of Firms	9,511			
Number of Unique Jobs	18,751			

Notes: This table shows summary statistics for the estimation sample of the on-the-job search model. The sample is limited from the full matched sample of Lightcast wages, Lightcast resumes and Form 5500 to individuals who transitioned between firms within the same industry, occupation, and CBSA.

Table XVI: Inputs to Model Estimation

	Source	Mean	p50	p10	p90
$\Delta(ln(w_{j,l}(1+r_j)))$	Lightcast & Form 5500	0.01	0.01	-0.70	0.73
$\Delta w_{j,l}$ in dollars	Lightcast Job Postings	385.23	149.41	-40000.00	40821.33
$\Delta r_{j,l}$ in dollars	Form 5500	-51.60	48.02	-3774.77	4178.51
$\Delta \mathbb{1}_{h_j}$	Form 5500	0.02	0.00	-1.00	1.00
$\Delta \mathbb{1}_{r_j}$	Form 5500	0.01	0.00	-1.00	1.00
$f_{j,l}^{NE}$	Lightcast Resumes	0.07	0.01	0.00	0.17
$g_{j,l}$	Lightcast Resumes	806.16	82.00	6.00	1384.00
# of Firms		9,511			
# of Industry by Occupation Groups		308			
# of Transitions		34,656			
# of Jobs		28,036			

Notes: This table shows summary statistics for the inputs to the on-the-job random search model.

Table XVII: Search Model Parameter Estimates

	mean	p50	p10	p90
γ_l	-2.83	-2.45	-7.37	0.83
$(1-\gamma_l)$	3.83	3.45	8.37	0.17
Elasticity of Total Pay to Wages	-0.74	-0.71	-0.88	4.76
β_l	0.02	0.02	0.00	0.04
Δa_l	0.93	0.76	-0.11	1.67
# of Transitions		34,656		
# of Industry by Occupation Groups		308		
# of Firms		9,511		

Notes: This table shows the estimated parameter values from the on-the-job random search model.

Table XVIII: Required Compensating Differentials to Make up for 1% Decrease in Wage - by Salary Group

	Estimation Sample Salary Percentile		
	10th-25th%	Middle 10%	75th-90th %
Average Salary	41618.00	57024.73	74649.37
γ_l	0.80	-2.90	-3.35
Dollar Differences:			
Wage ($)	-414.11	-567.41	-742.77
Retirement ($)	2,239,50	107.56	129.69
Total Compensation ($)	1,825.40	-459.85	-613.09
Percentage Differences:			
Total Compensation (%)	4.11	-0.74	-0.77
Employer Contribution Rate (pp)	5.48	0.28	0.25

Notes: Table shows averages. $ values are based on the ex-ante average salary, retirement contributions, and total compensation in the industry by occupation group, weighted by the number of new employees. Each group represents 25-30 industry by occupation groups.

Table XIX: Effect of Changing Wages or Retirement on Recruiting Success

	p50	p25	p75
1% Increase in Wages			
% Change in number of new hires	0.16	0.07	0.32
Net cost per one new hire	543.86	387.17	781.67
Net cost per employee	0.06	0.00	1.00
1 pp Increase in Contribution Rates			
% Change in number of new hires	0.41	0.10	1.05
Cost for New Employees Only:			
Net cost per one new hire	502.69	357.86	722.50
Net cost per employee	0.10	0.00	1.68
Cost for All Employees:			
Net cost per one new hire	8282.43	2400.00	36824.67
Net cost per employee	34.79	12.00	91.11

Notes: This table shows the effect of changing wages or retirement contributions by 1% of one percentage point, respectively, on recruiting success. The top panel shows the effect of 1% increase in wage. The bottom panel shows the effect of a one percentage point increase in the employer contribution rate, with costs for only the new hires and costs for all employees. Statistics are at the firm-level.

Table XX: Effect of Increasing Wages or Retirement Contributions on Worker Valuations

	Percentile of Retirement Valuation:		
	Top 10th	Middle 10th	Bottom 10th
Average Salary	70027.95	67144.90	56933.18
γ_l	-7.84	-2.48	0.91
% Increase in Valuation due to:			
1 pp Increase in Contribution Rate	42.94	10.92	0.69
1% Increase in Wages	4.18	2.98	7.93
# of Industry by Occupation Groups	31	32	31
# of Jobs	3,374	5,561	5,520

Notes: This table shows the effect on worker valuations of increasing wages by 1% or increasing the employer contribution rate by 1 percentage point. The effect is averaged amongst all workers in a decile of retirement valuations.

Chapter 2

Household Portfolios and Retirement Saving over the Life Cycle

With Jonathan Parker, Antoinette Schoar, and Duncan Simester

2.1 Introduction

The financial environment facing the typical American household has changed dramatically over the past four decades, as defined-contribution pension plans have spread, the costs of investing have declined, and regulations have evolved. In this paper, we use individual investors account-level data from 2006 and 2018 to characterize the life-cycle portfolio and saving behavior of American middle-class and upper-middle class investors with some retirement saving. We have three main sets of results.

First, we document that the share of wealth held in the stock market by typical American investors is hump-shaped over the working life, increasing until around age 50 and declining thereafter as investors approach retirement. Investors also hold a relatively high share of their wealth in equity on average, roughly 70%. These two findings represent a significant change in behavior. In the 1990's, equity shares documented in similar data were lower and rebalancing was largely unrelated to age (Ameriks and Zeldes, 2004). These two findings are not as visible in survey data (the Survey of Consumer Finances, SCF) partly due to mis-estimation of equity shares in hybrid funds.

Second, we find that these changes in retail investor behavior were accelerated by financial innovation and regulation, specifically by Target Date Funds (TDFs) and the Pension Protection Act of 2006 which allowed the use of TDFs as default investment options in employer-sponsored retirement plans. We use the quasi-exogenous variation between investors who enroll in a given retirement plan just before and just after it switches its default investment funds to TDFs. Consistent with the age-dependent asset allocations of TDFs, the change to a TDF default leads lower-income, younger investors to invest a greater share in the stock market and older workers to invest a lower share. Both effects decrease over time, as more investors adopt TDFs and TDF-like strategies.

Third, in contrast to portfolio allocations, retirement saving rates have a monotone increasing lifecycle pattern that is relatively stable over time and across cohorts. In our sample of investors with some retirement wealth, saving rates increase steadily with age, almost doubling between age 25 and 65. Their stability over time suggests at most minor effects of retirement saving products and changes in regulations designed to increase retirement saving over the past two decades.

As nicely elucidated in Campbell (2016), the institutions and laws surrounding household finance can be structured to improve the financial decision making of market participants. Our findings suggest that the fund design and regulation of TDFs has moved the average portfolios of retail investors towards the age pattern recommended by most prescriptive models of portfolio choice. Over the same time period, changes in retirement plan regulation and structure, most notably default contribution rates and automatic re-enrollment, have seemingly had little effect on retirement saving rates, which are remain lower that those recommended by most prescriptive models of saving and wealth accumulation (e.g. Poterba, 2014; Gomes et al., 2018; Duarte et al., 2021).

Our findings are based on an analysis of anonymized, account-level data from a large financial services company. The data contain the portfolios, individual trades, and detailed characteristics of millions of investors covering more than a trillion dollars in investable wealth. The size of our data set allows us to identify a sub-sample that is reasonably representative of the "typical" American retail investor with some retirement savings: investors with retirement savings accounts in the middle 80% of the age-adjusted distribution of retirement wealth, who we call *retirement investors* (RIs). Our analysis focuses on *investable wealth*, defined as stocks, bonds, and investment funds in retirement accounts and non-retirement brokerage accounts, and excluding bank accounts, durable goods, and housing, and on *saving rates* defined primarily as (realized) retirement contribution rates (the vast majority of inflows into financial wealth among our retirement investors).

We first document that during the last two decades this sample of US middle-class and upper-middle class investors on average held 70% of their investable wealth in the stock market, a share that is significantly higher than in prior decades. Second, across age, controlling for birth cohorts, the lifecycle pattern in equity allocations is hump-shaped. A cohort's average equity share increases by 7% as people age from 25 to 50 (and moves through calendar time), and then falls by the same amount from age from 50 to 65, as people reallocate financial wealth into safer assets, such as fixed income or cash-like securities.

These behaviors are markedly different from earlier time periods. Using similar administrative data on retirement savings prior to 2000, Ameriks and Zeldes (2004) reports an average equity share of only 58% and no re-balancing out of equities as people aged. These results are also different from leading survey evidence. The 2016 SCF shows an average share of only 54% of investable wealth allocated to equities in a comparable sample. Similarly, the SCF shows a reduction in equity shares with age that is only a third as large. We present several pieces of evidence that this difference is partly driven by a combination of mis-reporting and SCF data processing assumptions. In particular, both appear to underestimate the average equity share and its age-dependence in hybrid funds such as TDFs. While these findings highlight the benefits of administrative data, particular for products designed for less sophisticated investors, some residual differences remain and may reflect the policies of our data provider.

We also observe significant changes in behavior across ten-year birth cohorts. Each younger cohort has a higher equity shares than the prior cohort did at the same age. For example, cohorts born after 1970 have higher equity shares *at every (overlapping) age* than the previous, older cohorts. We also find that younger cohorts rebalance more as they age than older cohorts. This pattern is similar across terciles of ex-ante income, and while log income differences explain about half of the *level* differences in equity shares across people, the lifecycle pattern of income does not change the life-cycle pattern in equity shares.

Why has investor behavior changed from the 1990s? We show that the increased allocation to equity and the rebalancing over the life-cycle into safe assets was significantly accelerated by the growth of TDFs facilitated by changes in pension law. TDFs are mutual funds that automatically rebalance portfolio shares across different asset classes as people age. For investors more than 25 years before their target retirement date, a typical TDF maintains 80% to 90% of its assets in diversified stock funds and the remainder in bond funds. About 20 years before the retirement date, a TDF typically starts re-balancing the portfolio towards safe assets so as to reach a roughly even allocation between stocks and bonds by retirement.

TDFs grew rapidly following the Pension Protection Act (PPA) of 2006 which sanctioned the use of TDFs as "Qualified Default Investment Alternatives" (QDIA) in employer-sponsored retirement plans.[1] Prior to the PPA, most QDIAs were money market funds (which do not invest in stocks). Following the PPA and the employer adopting a TDF as the default fund, new employees who did not choose allocations were defaulted into TDFs, which would hold 80-90% of its wealth in stocks for a younger worker. Additionally, new hires who did make active decisions may have been influenced by the default option. Initially, existing plan participants (existing employees) were unaffected by the change in the default option.[2]

We estimate the effect of having a TDF as the QDIA by comparing the portfolios of people who enroll in a retirement plan in the two years before the PPA to the portfolios of those who first enroll in the same plan in the two years after, holding constant the employer. Because the PPA permitted but did not require employers to change the default allocation, we focus on firms that adopted TDFs as defaults in 2007 or 2008.[3] The main

[1] As discussed for example in Parker et al. (2020), investments in TDFs increased dramatically after the PPA of 2006, reaching $2.3 trillion in 2019 (Investment Company Institute, 2020)

[2] Over time, many plans have adopted regular automatic re-enrollment, in which existing employees are re-enrolled and must either make active choices or be defaulted into their plan's defaults.

[3] We also repeat the analysis for all employers independent of their default option, though these estimates are much noisier since many employers took many years to change their default options.

identifying assumption is that the change in default options does not cause employers to choose workers with different portfolio preferences or workers to choose different employers. Both of these possibilities are unlikely given what we know about labor market participation.

The adoption of a TDF as the QDIA leads younger new enrollees to invest more of their financial wealth in the stock market compared to those starting prior to the switch at the same employer. The youngest cohort (people 25-35 years old at job start) increase the share of their portfolios invested in stocks by more than 5%. Consistent with the glide path of TDFs, the adoption of a TDF as the default fund reduces the share of older workers' portfolios invested in stocks.

These effects of default investments are persistent but decline over the five years following enrollment. For lower-income, younger workers, the increase in equity share halves over the five years following enrollment. For older workers, the relative reduction disappears entirely. This convergence is driven by the control group, those investors that were not defaulted into TDFs. Thus, investors who enrolled before the default changed still over time adopted TDFs or TDF-like strategies, perhaps in response to advertising, financial advice, or peer effects from their co-workers. We conjecture that the sanctioning of TDFs and their implied lifecycle glide paths by the PPA, as well as their subsequent adoption by retirement plan sponsors and administrators, led to the more widespread adoption of TDFs and TDF-type strategies through many channels besides default investment funds.

In the third and last part of our paper, we show that, unlike in the case of portfolio allocations, there has been little change in retirement savings behavior across cohorts over time. Measuring a person's retirement savings rate as their annual contribution to their retirement saving plans as a share of their income, average retirement saving rates across all birth cohorts are 4.5% at age 25 and 8.5% at 65 years of age.[4] Tracking birth cohorts or even individuals over time, we find that between the age of 25 and 65 the average person increases their contribution rate by about 4.5%. This pattern does not change much across cohorts. The oldest cohort, those born between 1943 and 1953, have contribution rates that are about 0.5% higher in levels than the other cohorts. However, younger cohorts see slightly higher within-person changes as they age. Thus average changes in savings rates across cohorts are negligible.

In contrast, there is significant heterogeneity in the level of contribution rates across income terciles, Across all ages, the bottom tercile of the income distribution has an almost 2% lower contribution rate than the top tercile: 3.9% compared to 5.7% at age 25, and 7.3% compared to 9.2% at age 65. This pattern again stayed constant across birth cohorts.

Finally, comparing people enrolling before and after the PPA — which included a number of provisions intended to increase savings (Choi et al., 2004) — we find that people enrolling after the Act had similar contribution rates to those enrolling before the Act. In sum, despite large changes over time in plan design and regulation as well as in portfolio holdings, contribution rates to retirement saving plans among our sample of retirement investors have remained remarkably stable.

Related Literature Our paper is most closely related to papers that use administrative data to measure household portfolio allocations over the lifecycle, in particular Ameriks and Zeldes (2004). Poterba and Samwick (2001) also finds significant cohort effects in portfolio allocations over the life cycle. Administrative data from Norway shows that

[4]This measure includes automatic payroll deductions or auto escalation programs, but excludes any rebalancing flows or portfolio appreciation. We also check that our main conclusions are not related to people hitting the legal limit on tax advantaged contributions in a year, which occurs for 6-9% of our sample.

Norwegian investors have a hump-shaped equity allocation (Fagereng et al., 2017). There are substantial differences in portfolios across countries, see (Guiso et al., 2003b,a), and for example Christelis et al. (2013) shows that U.S. households have higher levels of stock ownership and stock market participation than most European households (49.7% versus 26%). Gomes and Smirnova (2021) estimates a lifecycle model for U.S. households and also finds a hump-shaped pattern in age.

We also contribute to a growing literature on the institutional causes of portfolio behavior. McDonald et al. (2019) studies changes in fund selection by new participants following changes in default investment funds in retirement plans in 2012. Mitchell and Utkus (2020), using Vanguard data, looks at the effect of TDFs on existing employees and new entrants under both voluntary choice and automatic enrollment plans. That paper shows that in voluntary enrollment plans, 28.4% of new entrants adopted a TDF in their 401(k) portfolios, compared to only 10.2% of existing employees. But in plans with automatic enrollment, 79% of new entrants chose a TDF. Similar to our findings, TDF investors held substantially more in equity: 81% for TDF investors compared to 63% for those without TDFs. Gomes et al. (2020) shows that TDFs improve investment performances due to a reduction in risk-taking in anticipation of lower expected returns.

Our paper also informs models of optimal portfolio choice (see the surveys Curcuru et al., 2010; Wachter, 2010). Merton (1969) and Samuelson (1969) provide canonical models in which portfolio allocations are constant over the life cycle and scale-invariant. A large body of research derives optimal portfolio choice in more complex models, the most pertinent example of which is the case where investors receive realistic stochastic, non-tradable "endowment" income over their working lives, which generally implies that investors should reduce holdings of risky assets over their life cycle, see Viceira (2001), Heaton and Lucas (2000), Campbell and Viceira (2002), Benzoni et al. (2007), Gomes et al.

(2020), and Storesletten et al. (2007). Other examples include non-standard utility functions, differences in risk aversion, and differences in beliefs.[5]

For the lifecycle pattern of savings rate, we relate to a large prescriptive literature concerned with what amount of saving households should be doing (e.g. Lusardi and Mitchell, 2007; Scholz et al., 2006), and a large positive literature estimating models from saving profiles assuming optimal behavior (e.g. Gourinchas and Parker, 2002). When looking at contribution rates, Gomes et al. (2018) suggests that more than 75% of US retirement savers display a significant shortfall in their contributions relative to an optimal consumption model. Poterba et al. (2011) similarly shows that households have inadequate financial wealth to support retirement, and for more than 70% of households, social security is their major asset.

Finally, our findings contribute to a growing literature evaluating the overall impact of defaulting people into savings allocations. Madrian and Shea (2001) and Choi et al. (2004) find very large effects of default enrollment for participation and savings rates in DC plans early in the introduction of these automatic enrollment options. More recently a couple of papers have found that these allocations might be partially undone by other choices people make.

2.2 Data

This section describes the account-level data set and then how we create a subsample that is representative of typical American retirement investors over their working lives.[6]

[5]For utility functions see Carroll (2000), Wachter and Yogo (2010), and Meeuwis (2019); for risk aversion see Ameriks et al. (2015) and Ameriks et al. (2019), and for beliefs see Meeuwis et al. (2018) and Giglio et al. (2019).

[6]In a method closely related to Meeuwis et al. (2018)

2.2.1 Account-level data

Our main data set contains anonymized, account-level data on financial holdings from a large US financial institution. For each investor, the data contain information on all their accounts held at the firm. For these accounts, we observe end-of-month account balances and holdings, and all inflows, outflows, and transfers at a daily frequency. We observe assets at the CUSIP level for 87% of wealth. For the remaining 13% we observe the characteristics of the fund the wealth is invested in. We aggregate accounts at the (de-identified) individual level and track each individual's portfolio. The data cover millions of investors and trillions in financial wealth. Our sample uses information from December 31, 2006 to December 31, 2018. We use the data at an annual frequency. We measure balances and holdings at the end of each calendar year by aggregating all observed flows within each calendar year. When we observe joint accounts for married couples, we allocate the funds to the spouse who has more total individual assets.

We focus on *investable wealth* defined as money market funds, non-money market funds, individual stocks and bonds, certificates of deposit, quasi-liquid retirement wealth, and other managed accounts.[7] We classify fund and security holdings into equity, long-term bonds, short-term bonds, and alternative assets (e.g. real estate and precious metals). Multi-class funds, also known as Target Date Funds (TDFs) or hybrid funds, are split between equity and fixed income in proportion to the observed equity share of the fund. Table I provides detailed variable definitions.

In addition to account-level portfolio information, we observe each investor's age, gender, zip code, and marital status (and an (imperfect) link to the partner if they also have accounts at the firm). For a subsample of the data, we also observe an anonymized

[7]Excluded categories of financial wealth are checking and savings accounts, saving bonds, cash value of life insurance, and other financial assets.

employer indicator, 3-digit NAICS code of the employer's industry, employment tenure, and, for a further subsample, gross annual wage income. We annualize all income observations by scaling up part-year incomes to a full-year equivalent.

2.2.2 Retirement investor subsample

While these data provide a detailed view of portfolio allocations for a large number of US investors, there are two potential limitations of our data. First, while we observe a significant share of US investors, this is obviously not a randomly selected sample. In particular, most of the wealth we observe is held in retirement savings accounts and few investors have a very high net worth (as we document subsequently). We would like to understand the relationship between our sample and a similar subsample of the US population. The second potential limitation is that we do not necessarily observe all of people's investable wealth because we do not observe wealth at other institutions.[8]

In order to both minimize and evaluate the importance of these two concerns, we define a sub-sample of people that are well-represented in our data and we can confirm are broadly similar to the same sub-sample in the US population. Our firm's data mainly includes typical working Americans with retirement saving during their working lives. This allows us to define a sample of *retirement investors* (RIs) that we can compare to a similarly-defined sample in the Survey of Consumer Finances (SCF).

First, we restrict our retirement investors sample to investors that are between 25 and 65 years of age. We exclude the youngest members of the sample because they typically have very low levels of investable wealth. By selecting 65 as the upper-bound, we

[8]The only concern is missing *investable* wealth. In both our data and the SCF, we exclude wealth in savings and checking accounts, as well as net housing wealth, defined benefit pension plans, etc.

avoid the issue that there is significant attrition among older investors in our data. Thus, our analysis focuses on working-age investors and so mostly on investors with labor income. Second, we drop investors with extremely high or low levels of retirement wealth, where *retirement wealth* consists of all investable wealth in retirement saving accounts of all types (excluding defined benefit plans and Social Security). We drop low wealth investors because they may simply have wealth at other institutions. We drop high wealth individuals because they are under-represented in our data. We choose our sample based on retirement wealth rather than total investable wealth because our data has incomplete coverage of non-retirement wealth (as we discuss below).

To construct our sample of RIs, we use data from the 2016 SCF to define a consistent sample of households based on retirement wealth holding. To correspond to our data, we treat couples in the SCF as two individuals. The SCF data allows us to measure retirement wealth, wage income, and age at the individual level, but non-retirement wealth is only measured at the household level, an issue we address in the following subsection. Using individuals aged 25-85 with some retirement wealth, we run quantile regressions of the log of individuals' retirement wealth (comparable to the measures in our institution's data) on a third order polynomial in age. We then drop individuals with retirement wealth below the estimated 10th percentile or above the 90th by age.[9]

Individual retirement investors make up 28% of the population of US households and 38% of the population of households aged 25-65 according to the (representative) SCF. They hold 33% (39%) of all household investable (retirement) wealth and 52% (54%) of investable (retirement) wealth among households age 25-65. Approximately 33% of both retirement wealth and investable wealth is held by the top 10%. Approximately 30%

[9]For age 30, in the SCF data, the lower bound is $1,328 and the upper bound is $66,370. For age 60, the lower bound is $6,774 and the upper bound is $744,000. See Appendix Figure B.1

of retirement wealth and investable wealth is held by those aged 66-85.

In our data, retirement investors – individuals between 25 and 65 and in the middle 80% of the distribution of retirement wealth at each age – make up 73% of accounts that we observe and hold 75% of all retirement wealth. Our sample of retirement investors contains millions of individual investors and well more than a trillion in investable wealth.

2.2.3 Descriptive statistics and comparison to SCF

The top panel of Table II shows summary statistics for our sample of retirement investors in 2016.[10] In our RI sample, the average age is 45 years old, and the average (median) wage income is $101,384 ($74,230). About 55% of the sample are male and 70% are married. The average portfolio beta is 0.75, and nearly half of investable wealth, on average, is allocated to target date funds. The average retirement wealth is $96,000. The bottom panel of Table II shows analogous statistics for the population of US retirement investors as estimated from the 2016 SCF. The average age is 47, the average (median) wage income is a lower $66,459 ($50,000), about half are male, and a slightly higher 78% of investors are married.[11] In terms of wealth, the average investor in the SCF RI sample lives in a household with approximately $273,000 of investable wealth and has $98,000 in retirement wealth themselves (bottom panel of Table II), comparable to the average in our sample.[12]

[10] Appendix Table B.1 shows the same statistics for our entire sample period, 2006-2018.

[11] Because of the way heads of households are assigned in the SCF, about 78% of respondents are male in the SCF. When including partners, the sample is evenly split between males and females.

[12] The statistics in Table II are representative of our retirement sample of middle class Americans with retirement wealth. They are not representative of the assets under management for a typical firm, since we are explicitly dropping the highest wealth households whose wealth mostly lies outside of retirement accounts.

Figure 2-1 shows that the retirement wealth distribution of our RI sample lines up well with individual respondents measured by the SCF. The SCF captures a somewhat higher mass of high wealth individuals, but overall the distributions are similar, suggesting that we are missing little retirement wealth at other financial institutions. Because RIs in our sample typically have most or all of their investable wealth in retirement accounts, we conclude that our sample of RIs provides a good overview of how the investable wealth of typical U.S. retail investors is allocated.

However, Table II also shows that RIs in our data have significantly less non-retirement wealth than RIs in the SCF. Is this because we miss wealth held at other institutions or because our data measure *individual* wealth and so we miss some of the assets that are common to both spouses for partnered investors? Figure 2-2a shows the total investable wealth distribution for individuals in our sample compared to households in the SCF (as in Table II) and confirms that we miss non-retirement wealth for wealthier individuals/-households relative to the SCF.

However, the difference in wealth observed in the SCF and in our data is mainly due to the fact that the SCF measures household wealth rather than individual wealth. That is, little of the difference is due to our missing wealth held with other institutions. In our data, the sample of married households for which we observe both spouses has on average fifty percent more investable wealth. Figure 2-2b shows that the distributions of total household investable wealth are much more similar for this married subset of our sample and the sample of married investors in the SCF. For couples, our data matches the SCF more closely, though our sample has a slightly higher median and mean wealth than the SCF. Appendix Tables B.2 to B.3 confirm this rough similarity both for married couples and separately for single individuals. We conclude that the difference in the distribution of wealth between our RI sample and that of the SCF is primarily driven by the unit of observation – individual investor as opposed to households.

Table II also summarizes the retirement saving behavior of our sample. The average RI designates a contribution rate of 8.1% of their income. However, because many people choose high rates that exceed the legal maximum contribution limit, the average ex-post rate is 6.4% of income. The SCF does not measure or allow us to infer portfolio betas, employment tenure, or retirement plan contribution rates.

2.3 The equity shares of portfolios

2.3.1 The average equity share

Our first main result is that in our RI sample, middle-class American investors hold a large share of their portfolio in equity. The average equity share of investable wealth is 71.0% in 2016 (Table III) and the median is 77.3%. For retirement wealth, the average is 71.1% and the median is 77.7%. Figure 2-3a plots the average equity share by year and while it is higher when the stock market has done well and lower when the stock market has done poorly, the average equity share is reasonably stable over time.[13]

Table III shows that the equity shares calculated for RI's in the 2016 SCF are substantially lower, 54.5% of household investable wealth and 52% of individual retirement wealth.[14]

We hypothesize that the difference between the equity shares arises largely because

[13]Table B.4 shows the comparison of the 2016 SCF with our full sample from 2006-2018. The magnitudes change slightly, but the arguments that follow still hold.

[14]These figures are somewhat lower than commonly reported in the SCF because we are calculating the average equity share rather than taking the ratio of averages. Because equity shares increase with wealth, equity shares calculated as aggregate equity over aggregate wealth are larger (e.g. Bricker et al. (2016)).

our data allow us to measure equity share exactly. Specifically, the SCF data are based on a survey in which people under-report the share of their wealth invested in equity because they are unaware that TDFs are allocated nearly entirely to equity for at least the first half of peoples' working lives. The main alternative hypotheses are that our sample under-represents investors with low equity shares and that the investments at other financial institutions that we do not observe have lower equity shares. Five pieces of evidence support our hypothesis.

First, consistent with people not being aware of the high equity content of TDFs, the respondents in the SCF who have some of their retirement assets in "mixed" funds report having lower than average equity shares, as shown by comparing columns (2) and (4) of Table III.[15] In the SCF, the subset of retirement investors who report having some assets in a mixed fund report an equity share of only 47%, versus 54% for all retirement investors in the SCF. We observe the exact opposite in our data: the subset of investors with TDFs has a somewhat higher equity share (77% versus 71%) consistent with the under-reporting of the equity share by SCF respondents.

Second, SCF respondents also appear to under-report the decline in equity shares with age, and in a way that is correlated with TDF ownership. Households in the SCF report little rebalancing out of equity as they approach retirement relative to households in our data. What decline there is is primarily among those households not holding TDFs – those holding TDFs report quite flat equity shares despite significant automatic rebalancing. As a result, the difference between equity shares in our data and the SCF is the highest for young households who hold some TDFs. These patterns are what one would

[15]The SCF phrases this question as "How is it invested? Is it all in stocks, all in interest-earning assets, is it split between these, or something else?" and then offers a variety of choices. We infer that the participant has something in a target date fund if they report having a mixed allocation of assets or if they have assets in a mutual fund or ETF. Thus, the same way survey responses may misreport equity share, they may also misreport (or we may mischaracterize) investments as hybrid funds.

anticipate if misreporting were due to a lack of understanding of how much TDFs allocate to equity for younger investors.

Third, the difference between equity shares of investors holding TDFs and investors who do not occurs primarily in retirement wealth, where the vast majority of TDFs are held, and not in non-retirement accounts. Panel B of Table III shows that RIs in the SCF reported equity shares that are about 9% lower when they hold assets in a mixed fund, while RIs in our sample report an equity share that is approximately 5% higher when they hold assets in a TDF.

Fourth, it is notable that outside of retirement wealth, our investors hold significantly lower shares of their wealth in equity, both than they do in retirement accounts and than the SCF households report in non-retirement accounts, both of which further support our main hypothesis. Specifically, consider the argument that our sample overstates equity share because we omit non-retirement wealth that the SCF measures. This is possible because, as noted, the distribution of wealth in our sample does not perfectly match that in the SCF, mostly due to the SCF reporting a somewhat larger amount of wealth held in non-retirement accounts, 13% versus 4% (Figures 2-1 and 2-2a and Table II). But non-retirement assets in the SCF (Panel C of Table III), have an equity share of 73%, which is higher than in our sample (and than in SCF retirement accounts).[16] Because this wealth has a high equity share, if we were able to observe and add such wealth to our data, it would raise, not lower, the average equity share in our data.

Finally, we compare the time series of the SCF, starting in 2007, with our sample in the same years and find that the discrepancy worsens over time. This is consistent with the rise of TDFs, both in our sample and in the United States in general (see Figure 2-5

[16]This includes equity held in trusts and mutual funds or stocks held outside of retirement accounts as a fraction of all trusts, mutual funds, stocks, bonds, and CDs held outside of retirement.

and Parker et al. (2020)). Appendix Figures B.3 and B.4 show the time series of the equity share of retirement wealth and investable wealth, respectively in 2007, 2010, 2013, and 2016 (the years of the SCF that overlap with our data). While there is already a large difference between the average equity share reported in the SCF and our sample in 2007, the gap grows over time, which supports the notion that TDFs contribute to the difference. Moreover, the difference between those with and without TDFs in their portfolios grows over time in the SCF, further suggesting that the mismeasurement is worsened by the presence of TDFs in the portfolio.

Despite these five arguments, some of the difference in measured equity shares may come either from idiosyncrasies of our sample or the wealth held at the institution we observe. For example, we could in part be measuring a firm fixed effect that raises equity share. Another possibility is that our average equity share is affected by the gender composition of our sample or by the fact that our unit of analysis is individuals and the SCF measures households.

However, neither the unit of observation nor household composition appears to be responsible for our finding of high average equity shares. The above pieces of evidence that TDF equity shares are misreported in the SCF all appear in the subsample of only married investors (Appendix Table B.5). For these sub-samples — married investors in the SCF and investors in our data for which we observe both spouses — we are comparing groups of individuals with very similar gender compositions. We also find similar patterns in the sample of all single investors in the SCF and in our sample (Appendix Table B.6), although with the one exception. Single investors who hold TDFs in their retirement accounts have higher non-retirement equity shares than all single investors.[17] This fact suggests that the non-retirement wealth of younger, single investors, that we do

[17] See columns (2) and (4) of Panel C of Table B.6.

not observe in our data likely has a higher equity share than the retirement wealth we do observe. Lastly, the residuals from a regression of portfolio equity share on gender, wealth, and birth-year fixed effects shows the same patterns again in the residuals, further indicating that the results are not driven by differences in sample composition (Table B.7).

We conclude that, while the two samples do not match perfectly, the SCF likely significantly understates the average equity shares of our retirement investor sample. We now turn to the analysis of the lifecycle dynamics of equity shares and how this has changed relative to similar administrative data from the 1990's.

2.3.2 The cross-section of equity shares by age

In the cross-section, averaging across people and years in our sample, the age profile of equity shares is declining in age. Figure 2-3b shows that the average equity share is roughly constant across ages prior to age 50 and then declines rapidly with age after age 50.

To control for the effect of differences in labor earnings, we regress equity share on indicator variables for three-year age groups using the following specification:

$$y_{it} = \boldsymbol{\beta}_1' \boldsymbol{Age}_{it} + \beta_2 Inc_{it} + \epsilon_{it} \tag{2.1}$$

where y_{it} is portfolio equity share, Age_{it} is a vector of three-year age group indicators and Inc_{it} is the log deviation of the individual's income in each year from the sample mean income, which is only included in some specifications.

The first column of Table IV shows that equity shares actually decline monotonically with age, beginning at approximately 74% for 25-27 year-olds and decreasing to approximately 55% for 64-65 year-olds. But this decline is uneven. From age 26 to 48 the average equity share decreases by only 3% over 22 years, while after age 50 the equity share decreases by 2-3% per year.

The second column of Table IV shows that, comparing people with the same income across ages, the portfolio share of equity still declines monotonically with age but now more steadily, with a more significant decline before age 50. People with higher income tend to have higher equity shares. A two standard deviation change in income is associated with a nearly 8% higher equity share. This effect is identified from the primary source of variation which is the difference in income across individuals not age groups. Finally, comparing the last rows of Columns (1) and (2) shows that differences in income explain roughly as much variation in portfolios as age groups explain.

This cross-sectional age pattern holds widely across sub-groups of retirement investors. Columns (3)-(5) of Table IV show that the lifecycle patterns found in columns (1) and (2) are not driven by investor income: equity share decreases by approximately 25% over the lifecycle, regardless of one's initial income tercile.[18] Comparing columns (3)-(5), we see that the equity share is lower, by about 5%, in the lowest income group, but the decrease with age is similar in magnitude.

Finally, the lifecycle pattern does not appear to be driven by the passive appreciation of equity holdings. Appendix Table B.8 shows that the cross-sectional results are similar using price-constant equity shares. The price-constant equity share measures inflows and

[18]Initial income is based up on the first (or second, if first is not available) year in which the individual enters our sample. The first tercile of initial income covers those with income below $46,000 per year, approximately. The second covers those with income $46,000-75,000. The third tercile is those with an initial income greater than roughly $75,000 per year.

outflows to each asset, ignoring any change in price, making equity shares insensitive to asset returns.

2.3.3 The portfolio share of equity over the lifecycle

In contrast to the cross-sectional pattern, tracking the same individuals over time, equity shares are hump-shaped across the working life.

Table V shows analogous regressions to the specification in equation (2.1) but includes person fixed effects. Column (1) shows that young people tend to increase the equity share of their portfolios, but as they approach retirement they reduce their equity exposure. People increase their equity share by approximately 7% from age 25 to 50, then they decrease it by about the same amount from age 50 to 65. Changes in income do not drive this result. The same pattern holds in column (2) when controlling for income, with a slightly higher magnitude (10% instead of 7%).[19]

This lifecycle hump-shaped pattern holds across income groups. The last three columns of Table V show regression results for different levels of initial income. Each group increases its equity share by 5-7% from age 25-50, and then decreases it as they age. We observe more aggressive rebalancing away from equity in the higher income groups, with the richest decreasing their shares by about 7% relative to their position at age 25-27. Those in the lowest income group decrease their equity share by only about 2%. Of course, those with higher income also start out with higher equity shares, and thus have more room to decrease them.

[19] The coefficient on income measuring the effect of changes in income is also smaller than the coefficient on the level of income in Table IV measured in the cross-section, an effect examined in detail in Meeuwis (2019).

2.3.4 Changes across cohorts

Third we document that equity shares are increasing across cohorts and the hump-shaped pattern of equity over the lifecycle is a relatively new phenomenon. We focus on cohorts of people born in 10-year periods. We have five such cohorts, from those born between 1943 and 1952, to those born between 1983 and 1992. Looking across ages and across cohorts reveals two patterns.

First, looking across years, Figure 2-4a shows that the three cohorts born more recently (those born around 1965, 1975, and 1985) have slightly higher equity shares to those born around 1945 and 1955 in the 2000's, but fifteen to twenty percent higher equity shares in 2018.[20] The oldest cohort start with roughly 10% less of their portfolio allocated to equity and ends the sample with 15% to 25% less than the youngest three cohorts.

Second, and more importantly, younger cohorts have higher equity shares than older cohorts at the same age. This fact can be seen at the ages for which cohorts overlap in Figure 2-4b, which shows equity share by age for different cohorts.[21] More precisely (based on Appendix Table B.9), at any age there is a monotone increase in average equity share as one looks at younger cohorts (excluding 'endpoints' for each cohort where age composition potentially plays a big role) at the same age.

These patterns imply that cohort-specific lifecycle profiles of equity shares have both risen and become more hump-shaped for younger cohorts. Table VI makes this point by tracking the same investor over time, broken out by birth year cohort.

First, the older cohorts allocate away from equity sooner and more quickly than

[20]Figure B.6a shows a similar pattern for portfolios betas. Moreover, portfolios betas within cohort are relatively stable over time.

[21]Figure B.6b shows that this pattern is also true for portfolio betas. In other words, more recent cohorts have higher-beta portfolios than older cohorts at the same age.

younger cohorts. Comparing the results in column (1) to column (2), from ages 55-65, the 1943 cohort decreases its equity share (relative to their own shares at age 52-54) by 19%. Meanwhile, those in the 1953 cohort decreased their equity share (again, relative to their own shares at age 52-54) by only about 5%. Similarly large difference appears when comparing the 1953 cohorts (column 2) to the 1963 cohorts (column 3). Those born 1953-1962 decrease their equity shares, by about 8% from age 43-52, while those born 1963-1972 actually increase their equity share by nearly 2% over the same age range.[22]

Second, among the youngest cohorts, shown in Columns (4) and (5), those born more recently increase their equity shares more quickly in their earliest years of investing, by approximately 1% more from ages 25-36 than those born 10 years earlier.

What has driven these changes in portfolio behavior? The greater investment in the stock market by the young and the decline over the latter half of working life are both consistent with an increased adoption of TDFs. As shown in Figure 2-5, younger cohorts are much more heavily invested in TDFs. Columns (6)-(8) of Table VI shows the lifecycle pattern of stock market investment by initial allocation to TDFs. Investors who begin with a high allocation of their portfolio to TDFs (75-100%) start life with higher initial equity shares and exhibit much stronger rebalancing behavior than those who start life with less invested in TDFs. In contrast, those with initially low allocations to TDFs (0-25%) follow more of a hump-shaped pattern in their equity shares, starting with approximately 62% equity, increasing it to 70% by mid-life, and lowering it only modestly to 66% at age 65. We discuss the role of TDFs in more detail in the next section.

[22]Put differently, comparing the trend in column (2) from ages 43-57 we see that those born from 1953-1962 decrease equity shares at about 2-4% per year. On the other hand, column (3) shows that those born from 1963-1972, at the same age, hold their equity share almost constant until they reach age 52, when they start to decrease it by only 1-2% per year.

2.3.5 Relation to portfolios during the 1990s

In this subsection, we present a final piece of direct evidence that portfolio behavior has changed over time by comparing the lifecycle portfolio holdings in our data covering 2006-2018 to the portfolio holdings in very similar administrative from the 1990s. Specifically, we replicate the central analysis from Ameriks and Zeldes (2004) which is based on administrative data from a large financial institution, where both the type of data and institution are quite similar to ours. Figure 12 in Ameriks and Zeldes (2004) shows that investors held less of their wealth in stocks and did not reduce their equity shares with age.

Figure 2-6 replicates Figure 12 from Ameriks and Zeldes (2004) and visually summarizes three main points. First, the top panel of Figure 2-6 shows that equity shares in more recent years are high and decline with age across households, with a steeper slope after age 55. In contrast, in earlier data, equity shares decrease with age from age 25 to 35 and are roughly the same for all ages after 35 (Ameriks and Zeldes, 2004, Figure 12, top panel).

Second, the middle panel of Figure 2-6 shows that tracking cohorts as they age, equity shares are roughly independent of age during the first half of working life and then decrease with age during the second half of working life. In contrast, the middle panel of Figure 12 in Ameriks and Zeldes (2004) shows that each cohort's equity share was *upward sloping* in age in the 1990's.

Third, in the last panel of Figure 2-6, the solid, red line shows that equity shares are hump-shaped in age controlling for differences across cohorts. The analogous figure in Ameriks and Zeldes (2004) shows a linear upward sloping line. Controlling for differences across years, the dashed blue line shows that equity shares are decreasing, more rapidly later in life, as in a TDF glide path. The analogous figure in Ameriks and Zeldes (2004) shows a flat line.

Finally, to confirm that we are documenting a real change in behavior rather than something specific to administrative datasets or these two firms, we replicate the Ameriks and Zeldes (2004) analysis of portfolios in the SCF. We compare Figure 9 in Ameriks and Zeldes (2004) which is based on SCF data from 1989-1999 with our own version of Figure 9 based on SCF data from 2007-2016 (Figure B.5). We find the same changes as we show between the two administrative data sets, but of a smaller magnitude.[23] Specifically, equity shares have risen in the SCF and there is now a lifecycle pattern of rebalancing out of equity in the second-half of working life.

2.4 Pension regulation, TDFs, and portfolio allocations

This section provides evidence that the the rapid rise of TDFs following the Pension Protection Act (PPA) of 2006 contributed to both of the main new facts we document: that equity shares that are high earlier in life and decline linearly over the second half of investor's working lives.

The Pension Protection Act of 2006 permitted Target Date Funds to be "Qualified Default Investment Alternatives" (QDIA) in employer-sponsored, defined-contribution retirement plans. The act provided a "safe harbor provision" that clarified that the use of a TDFs as the default investment vehicles in a plan was consistent with the fiduciary responsibilities of the plan sponsor (the employer) and the plan administrator.[24] Prior

[23] This smaller magnitude is consistent with the SCF not correctly measuring equity shares in TDFs, as we discussed in relation to Table III's evidence that equity shares decline with age more in our data than in the SCF.

[24] https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/resource-center/fact-

to this provision, both employers and administrators faced potential legal liability for replacing existing default options — primarily safe money market funds — with TDFs. Following the PPA, plans increasingly adopted TDFs as defaults, which moved employees who passively accepted or chose the default investment out of very safe, low-return funds and into largely equity funds. Following the PPA, the availability, adoption, and use of TDFs accelerated rapidly in the U.S. (Figure 2-5 shows the rise in our data and Parker et al. (2020) documents the aggregate growth in assets under management).

2.4.1 Identification Strategy

To identify the effect of TDFs on investors' portfolio allocations, we compare the life-cycle investment behavior of workers hired by a given firm just before and after 2006 at firms that switched their default investment at this time to a TDF. This analysis identifies the exogenous effect of the PPA on investors' portfolios assuming that people (or employees) did not endogenously change the jobs they take due to the introduction of the PPA or their (potential) employers response to it. This is a reasonable assumption since employees typically are not aware of these regulatory changes and base their employment decision on many other factors. Employees who joined their employers before 2006 almost exclusively entered into plans that did not have TDFs as a default option, since without the safe harbor provision of the PPA, employers found it very risky to use this option. After 2006, many employers adopted TDFs as defaults and the employees joining the firm after this change then saw a different default investment vehicle. We focus specifically on individuals who enrolled in plans in 2007 or 2008 that switched to having a TDF as a default following the PPA.

---

sheets/default-investment-alternatives-under-participant-directed-individual-account-plans

We first analyze the short term – two year – effect of the adoption of a TDF as the default investment fund. We take the sample of employees who start a new job between 2005 and 2008 at a firm that switched to having a TDF as the default option after 2006. The specification is:

$$y_{it} = \beta_1 \times D_{treated} + \beta_2 \times D_{treated} \times AgeEnrolled_i + \beta_3 \times AgeEnrolled_i + \lambda_f + \epsilon_{it} \quad (2.2)$$

where y_{it} is the portfolio equity share. $D_{treated}$ is an indicator variable equal to one if an investor is enrolled in a retirement plan that switched to having a TDF default immediately after the PPA, in 2007 or 2008. λ_f is an employer fixed effect, so we compare individuals enrolling before the Act (in 2005 or 2006) to those enrolling after (in 2007 or 2008) at the same employer. Since TDFs by definition change their target allocation for people of different ages, we also include categorical variables for 10-year age groups at enrollment and interactions of these age groups with the treatment variable. These regressions include only firm-level fixed effects and not individual-level fixed effects since we estimate the effect of the PPA by comparing across investors who enrolled just before and after 2006. Table VII shows the results of estimation of equation (2.2) restricting the sample to the two years after enrollment for each new employee.

Column (1) of Table VII shows that someone aged 25 - 35 who is enrolled into a plan with a TDF default in the two years after the act (compared to those enrolled in the two years prior to the act in the same firm) has a 5.5% higher equity share during the first two years of their employment (the coefficient on the *treated* indicator variable). The effect is statistically significant and economically large relative to the average increase and decrease of average equity over the lifecycle. The coefficients on the interaction terms show that the effect of a TDF default on older individuals is to decrease their equity shares. For example, while those aged 55-65 at enrollment have 13% lower equity shares than those aged 25-34 (row 4), those treated by the change in default have equity shares that are lower again by nearly 15% more (row 7). This age pattern in the treatment effect – from positive when young to negative when old – is consistent with the change in behavior that we document. Relative to the low and (roughly) age-invariant equity profiles of the 1990's, TDFs raise the allocation to stock for younger workers and lower it for older workers. Controlling for income (column (2)) the effect on young households declines only slightly and that the lifecycle pattern remains quite similar.

Across income groups, the effect of a TDF default is generally larger for young investors with lower incomes, and generally smaller for all investors with high income. In Columns (3)-(4) of Table VII, we repeat the analysis from Column (1) for the subsamples of people with the lowest and highest initial income. Column (3) shows that the initial impact on equity share at young ages is almost 6%, compare to 5.5% in the full sample. The effect on those in the bottom income tercile is similar for the older age groups as in the full sample. For the highest income tercile (column (4)), the treatment effects are significantly muted. For the youngest group (people 25-35 in age) the magnitude of treatment effect on equity shares is less than 2%. There are two possible reasons for this small effect. First, even in the control group, young higher-income households have high equity shares. Second, higher-income households make more active decisions make less use of TDFs. Consistent with the latter, the rebalancing effect of enrolling with a TDF default is much less pronounced for the highest income tercile.

These results are not driven by differences in investors' pre-existing portfolio allocations or their experience with assets or asset managers prior to enrolling with their new employer. Columns (5)-(6) display the results of the same analysis as displayed in columns (1)-(2), but conducted only on those individuals who have no other retirement assets or rollover funds prior to enrollment at our institution. It turns out that this sample restriction drops very few households. As a result, the results are virtually unchanged from those in the first two columns, even for older new employees.

2.4.2 Medium-run impact and convergence

In order to analyze the impact of the PPA on the investment dynamics and persistence of portfolio allocation over the medium term, we now repeat the analysis above but track individuals for five years after enrollment in a retirement plan. We expand our specification to:

$$\begin{aligned} y_{it} = &\beta_1 \times D_{treated} + \beta_2 \times D_{treated} \times AgeEnrolled_i + \beta_3 \times AgeEnrolled_i \\ &+ \beta_4 \times D_{treated} \times \lambda_t + \beta_5 \times \lambda_t + \lambda_f + \epsilon_{it} \end{aligned} \tag{2.3}$$

where the notation is the same as in equation (2.2) and we have added year fixed effects λ_t and interactions of these year fixed effects with the treatment indicator to trace out the investment trend for both the control and treatment group over time, respectively. In this specification, to study differences by age, we run separate regressions using different age groups, rather than including the full set of age fixed effects and interactions.

Column (1) of Table VIII confirms that the effect of the PPA on equity shares is positive even averaged over the first five years after enrollment and based on the full sample enrolled in 2005-2008.[25] More interestingly, the dynamic analysis however reveals that the difference in portfolios declines over time. As before, the positive effect of the PPA is much larger for the low income group than the high income group (columns 2 and 3), however, this difference is not persistent, shrinking to nearly zero five years after treat-

[25] Appendix Table B.10 shows the results only using those who enrolled in their plan in 2007 as the control group. This minimizes possible spurious correlation due to the financial crisis.

ment.

More importantly still, TDFs tend to raise equity shares for the young and lower equity shares for those near retirement. Column (4) shows that for those aged 25-34 at enrollment the change to a TDF default increases equity share on average by 1.5% the year of the change and rises to almost 4% in the following two years. In the last two years we see some convergence between the treatment and control groups. But at the end of the five-year period, the treated individuals still have equity shares that are nearly 3% higher than those of the control group.

Table VIII further shows that for the older age groups, the effect is the opposite: the PPA decreases equity shares immediately following treatment, which is in line with the prescribed glide path of TDFs. As with the youngest group, this difference tends to decrease over time as the two groups converge. For those aged 55-65, the difference is persistent, with the treated group's equity shares still being 2% lower than the control group's five years after treatment.

The PPA also played a role in the convergence of portfolios allocations between income groups, particularly for those that were enrolled at a young age. Figure 2-7 plots the predicted equity shares of treatment and control groups broken out by age and income tercile.[26] Looking first at the youngest age group (25-34) in Figure 2-7a, the adoption of a TDF default significantly increased equity shares for the low income group. And again we see that the control group converges somewhat to the treatment group over time. In contrast, the effect on equity shares for the high income group is positive but much smaller.

Figure 2-7b shows the results for those enrolled from age 55-65 for high-income in-

[26]These are estimated in unreported regressions that repeat columns (4) and (7) from Table VIII on the income subsamples.

vestors and for low-income investors. The PPA significantly decreased equity shares for both groups. Over time, the two treated groups become more similar when compared to the two untreated groups, implying a similar convergence effect of the changes in defaults facilitated by the PPA.

In conclusion, middle-class and upper-middle class working-age American investors with retirement wealth now hold a large share of their financial wealth in equity and reduce the share as they age, following a concave rather than a linear lifecycle pattern. This is relatively new behavior, not visible prior to 2000. This large change appears to be due to the combination of industry development and regulatory approval of target date funds as defaults in retirement saving plans. The new portfolio behavior follows the prescription embedded in TDFs investment strategies, to invest mostly in stocks when young and to decrease this share significantly in the second half of working life as retirement approaches.

2.5 Contribution Rates

This section presents an analysis of the average contribution rates that investors make to their retirement plans over their working lives. The analysis, which mirrors our analysis of portfolio composition, shows three main results. First, contribution rates increase linearly with age, increasing by 4-6% over the working life. Second, and more importantly, unlike portfolio behavior, this behavior has been relatively stable over time. Third, average contribution rates responded only minimally to the Pension Protection Act of 2006, and actually decreased slightly following the Act. Thus, contribution rates are less sensitive than equity shares to the changing regulatory environment or investment trends over time.

2.5.1 Realized Contribution Rates

We measure realized contribution rate as the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year.

Table IX presents the coefficients from estimation of equation (2.1) with realized contribution rates as the dependent variable. Column (1) shows that, in the cross-section, contribution rates increase monotonically with age, from about 4.6% at age 25 to 8.5% at age 65. Columns (2) - (5) show that contribution rates increase by a similar 4% over the working life when controlling for income in two different ways. In column (2), the coefficient on current log income deviation from the average implies that each 1% deviation in income from the average is associated with a nearly 2 basis point increase in reported contribution rate. Instead, looking across initial income groups, those in the highest income group (column (5)) save nearly 2% more, on average, than those in the lowest income group (column (3)) at every age. The increases over the lifecycle, however, are parallel: each group increases its total saving rate by about 3.5% from age 25-65.

This relatively stable, increasing age pattern of saving is not only a feature of the cross-section, but characterizes the average behavior of investors as they age.[27] Table X shows the regression of realized contribution rate on age group indicators, including a person fixed effect which effectively include a cohort effects. In the baseline results in column (1), contribution rates increase by just over 5% over the lifecycle, the same increase as in the cross-sectional age pattern. In column (2), we also control for log income which does little to change the age effects nor the R-squared of the regression. This confirms that, when controlling for the person fixed effect, income is less important for determining

[27] We find analogous results to those in this subsection controlling for cohort effects rather than individual effects (see Appendix Table B.11).

contribution rates and that increasing savings rates with age are unlikely to be due to income profiles.

Splitting households by initial income, all income groups also show a similar life-cycle pattern, though those with higher incomes have higher contribution rates overall. Columns (3) and (5) show parallel increases of 4.6% and 5.1% over working life in average contribution rates for the bottom and top income groups. The average contribution rate of the middle initial income group increases by a slightly larger 6.2%. These differences in contribution rates across cohorts in part relate to contribution limits set by the IRS, as we analyze in the next subsection.

The within-person results broken out by cohort again show that each cohort increases its savings rate with age. Columns (1)-(5) of Table XI show that the younger cohorts increase their rate of contribution at a slightly faster pace. For example, comparing column (4) to column (5), we see that 28-30 year-olds born from 1983-1992 increase their contribution rate by 0.81%, relative to at age 25-27, while 28-30 year-olds born from 1973-1982 increase it by 0.52%. A similar pattern holds when differencing across rows for the other age groups that are common to multiple cohorts. In summary, although older cohorts start at a higher savings rate, the younger cohorts increase their rate slightly faster as they age, even when controlling for income.

Finally, the stable pattern of saving behavior holds regardless of the share of TDFs initially held by the investor, although investors with large initial investments in TDFs increases their savings rates by more than other investors. Columns (6)-(8) of Table XI break out the within-person results by initial TDF share. Investors with intermediate investments in TDFs (column (7)) have the highest level contribution rates. Investors with TDF shares below of 25% increase their contribution rates over their working lives by a large amount, 6.6%, or about 1.5 to 2% more than the 4.9% to 4.6% increase of the two other groups that start with a larger allocation to TDFs.

2.5.2 Realized versus Reported Contribution Rates

Thus far, we have limited our discussion of contribution rates to *realized* contribution rates: the percentage of income that is actually saved for retirement, ex-post (on a year-by-year basis). However, there is a distinction between the realized rate of savings and the designated or *reported* rate of savings that investors decide upon ex-ante. The difference between reported and realized contribution rates is the result of retirement contribution limits, set by the IRS. [28] Depending on their income and reported contribution rate, some people will hit their maximum contribution before the end of the year, and thus their actual *realized* contribution will be less than what they designated at the beginning of the year. This may occur if someone has a very high income, or if someone sets a very high contribution rate. We address this discrepancy in two ways, both of which confirm the results of the previous subsection.

First, we condition our analysis on an indicator variable equal to one if an individual hits their contribution limit in the given year. We set this indicator to one if the investor's *reported* contribution rate times their income is larger than the allowed amount by the IRS in that year.[29] We find that 6-9% of our sample with available income data max out on their contribution in a given year. We then conduct our analysis again using a

[28]https://www.irs.gov/newsroom/401k-contribution-limit-increases-to-19500-for-2020-catch-up-limit-rises-to-6500

[29]Base contribution limits increased from $15,000 in 2006 to $18,000 in 2017. In addition, contribution limits for individuals older than 50 are higher by a "catch-up" contribution amount that rose from $5,000 in 2006 to $6,000 in 2017. We use the age- specific limit in the corresponding year to calculate the limit for each investor.

specification that controls for hitting the contribution limit:

$$y_{it} = \beta_1 \times D_{maxout} + \beta_2 \times D_{maxout} \times Age_{it} + \beta_3 \times Age_{it} + \epsilon_{it} \tag{2.4}$$

where y_{it} is the realized contribution rate. D_{maxout} is an indicator equal to one if the individual investor, i, maxes out on their retirement contribution, as described above, in a given year. Age_{it} are indicators for ten-year age groups. In some specifications, we also include a control for the deviation of the investor's current income from the average.

In the cross-section, retirement contribution rates increase by under 3% over the working life when controlling for hitting the contribution limit and whether controlling for income or not (Table XII). Note that the coefficient on maximizing retirement contributions is positive implying that on average investors that contribute a larger share of their incomes are more likely to be hitting the legal limits. The coefficients on the interaction terms between hitting the cap and age show that the cap on contributions lowers realized contributions most strongly for those at prime earning age (age 35-54, Columns (1) and (2)). Columns (3)-(4) show results replacing the indicator variable in equation with *max out ever*, an indicator equal to one if the individual maxes out on their contribution during any year that we observe them in the sample (before or after the current year). The cross-sectional patterns lie between those without any control for maximizing retirement contributions and those withe contemporaneous control in Columns (1) and (2).

Our second method of addressing the discrepancy between realized and reported contribution rates is by repeating our analysis on the *reported* contribution rate rather than the realized rate. The reported contribution rate is the percentage of income that the individual designates to their retirement account at the beginning of each year.

Using designated rather than realized contribution rate largely confirm our results using realized rates.[30] In the cross-section, designated savings increases monotonically with age from about 6% to 10% over the lifecycle (Table B.12). Hence, reported contribution rates are about 1% higher than the realized rates observed in Table IX, confirming that some individuals set a rate that is too high and hence save at a rate lower than anticipated. Column (2) of Table B.12 shows the same age pattern controlling for log income. As before, the coefficient on income implies that each 1% deviation of income from the average is associated with a nearly 5% increase in reported contribution rate. Note that income has significantly more explanatory power for designated contribution rates (adding income doubles the R-squared) than for realized rates, consistent with the contribution cap distorting an otherwise relatively stable desired contribution rate over the income distribution. We also find cross-sectional patterns in designated contribution rates across different cohorts and different TDF allocations that are similar to those for realized contribution rates (Appendix Table B.13).

Finally, and most importantly, the baseline results including a person fixed effect, shown in Table B.14 confirm that individuals increase their contribution with age at a magnitude that explains nearly all of the aggregate variation. Similar to realized contribution rates, higher reported contribution rates are not driven by people earning higher incomes as they age (column (2)). Each cohort behaves similarly, but younger cohorts increase their contributions at a slightly quicker pace (Table B.15, columns (1)-(5)). Additionally, those with the lowest allocation to TDFs (column (8)) increase their contribution rate by more than those with higher allocations to TDFs at every age. All these results are consistent with conclusions in the previous subsection where we use realized contribution rates to conduct analogous analysis.

[30] Tables in the Appendix.

2.5.3 The Effect of the Pension Protection Act of 2006

Our results so far suggest that the Pension Protection Act of 2006, which included several provisions designed to encourage savings in retirement funds (Beshears et al., 2010), had little to no impact on actual retirement saving rates by age or across cohorts of savers. In this subsection, we present evidence that the immediate effects of the PPA on retirement saving rates were, if anything, negative.

We replicate the difference-in-difference analysis of subsection 2.4 but comparing the retirement saving rates of new enrollees at the same employer in the two years before and after the PPA of 2006. Unlike in our analysis of portfolios, we designate anyone enrolled during 2007-2008 as a treated investor, regardless of whether or not their plan's default investment allocation changed and estimate equations (2.2) and (2.3) with reported contribution rate as the dependent variable.[31]

First, as shown in column (1) of Table XIII, those enrolling at an employer after PPA 2006 have lower contribution rates in the two years following enrolment. The effect starts at a large -0.43% of income for those age 25-35, and becomes increasingly negative with age, reaching -1.2% for those age 55-65. The negative sign, magnitude, and pattern are similar when controlling for income (column (2)) and across income groups (Columns (3)-(4)). Finally, this decrease in saving is similar for those with no other retirement assets at the institution prior to enrollment, as shown in columns (5)-(6). This result implies that the finding is not driven by those who have some wealth at the institution prior to

[31] In our analysis of portfolios equity shares, only those who were enrolled in a plan that changed its default investment to a TDF after the PPA are considered treated. In that case, we measured the effect of TDF default allocation, induced by the PPA on portfolio allocation. In the case of contribution rates, we want to measure the impact of the PPA overall. The PPA had a significant number of provisions intended to increase savings rates, but we are not able to isolate those plan features in our regressions, due to data limitations. Hence, we simply designate anyone enrolled in a plan from 2007-2008 as a treated investor, regardless of which plan features changed following the PPA.

enrolling in a new plan.

Tracing out the effect over the five years following enrollment, the PPA had only a transitory negative effect on average contribution rates and it is largest for the oldest investors. We repeat our analysis tracking investors for five years after they enroll and including interactions of individual indicator variables for each year after treatment with an indicator for being treated by the PPA (enrolled in 2007 or 2008, versus 2005 or 2006). As shown in Table XIV, Column (1), the PPA has a negative initial effect on contribution rates, but the magnitude decreases over time and is essentially zero five years after treatment. Column (3) shows that the decrease in retirement contribution rates is slightly more persistent for households with higher (initial) income. Splitting the result by age group, columns (4)-(7) shows that the effect is negative for each age group, and largest in magnitude and most persistent for the older age groups. For example, those aged 55-65 when enrolling after PPA have contribution rates that are 1.3% lower than those enrolling before during the year they enroll, a difference that declines to 0.5% after five years. In contrast, those in the age group 25-34 enrolling after PPA contributed only -0.7% less of their income the year of enrollment and -0.1% five years after enrollment relative to those that enrolled just before PPA.

It is possible that the interaction terms are picking up some of the differences in year fixed effects and that saving rates are lower for the treatment group due to the timing of the financial crisis. However, Appendix Table B.16 shows that the results are similar if we include only those enrolled in 2007, rather than 2007-2008. This eliminates any possible spurious effects on saving rates due to the financial crisis. Moreover, the income control in column (1) indicates that the results hold even for those who did not experience significant changes in income due to the financial crisis.

Figure 2-10 shows the predicted contribution rates for those in the youngest and oldest age group, split out by those in the lowest and highest income tercile.[32] Looking first at those aged 25-34 in Figure 2-10a, the PPA significantly decreased contribution rates for both income groups, initially by about 0.7-0.9%.[33] However, the difference between the treated and control groups converges to zero over time.

For those enrolled when aged 55-65, investors with lower incomes are more affected by the PPA, as shown in Figure 2-10b. The treated group with high incomes decreases their contribution rate by about 0.7% following treatment. For the lower income group, the immediate effect is larger: 1.3%. For both income groups, the difference between treated and control after five years is nearly zero.

2.6 Conclusion

The results in this paper show that the portfolios of typical retirement investors in the U.S. have changed significantly over the last few decades. Investors hold more of their investable wealth in the stock market than they did in the 1990s, and they reduce the share of their portfolio invested in the stock market as they age when they used to maintain a relatively constant share as they aged. These two changes are consistent with the rise of new retirement savings products, such as TDFs, and the advent of new regulations, such as the PPA of 2006. We show that the adoption of TDFs as default investments in employer-sponsored retirement plans has a causal effect on portfolios in the direction

[32] These are estimated in unreported regressions that repeat columns (4) and (7) from Table XIV on the income subsamples.

[33] The reason that both the treatment and control group decrease their contribution rate over the five-year time period following enrollment is that the five-year period that we analyze happens to take place during the Great Recession and its aftermath. This pattern is consistent with the fact that contribution rates increase with age (Table B.12 (cross-section) and B.14 (within-person)) and at the same time contribution rate decreased uniformly across birth cohorts from 2007-2009 (Figure 2-9).

of the observed changes in portfolio holdings for younger workers, particularly those with lower incomes. But the causal effects are short-lived. We see that individuals in the control group who are not defaulted into TDF over the next five years start catching up to the treatment group.

Did the PPA and the rise of TDFs fundamentally change investor behavior, or, even without these changes, would investors have increased the share share of their portfolios invested in stocks when young stocks and decreased this share as they aged? On the one hand, greater dissemination of prescriptive, model-based portfolio advice and a recognition of the equity premium puzzle, might over time have led to this shift in investor behavior even absent the PPA. On the other hand, the PPA, by sanctioning TDFs as default investment options, may have been a critical catalyst in making higher equity allocations acceptable advice for retail investor portfolios. The timing is certainly consistent with this second interpretation, as is the observation that the PPA seems to have led the investment advice industry (many of whom, like retirement plan sponsors, have a fiduciary duty to their investors) to recommend TDF-like investment strategies more broadly. In either case, these changes in behavior may be part of the recent high valuation of stocks (high price to earnings ratios) as well as potentially leading to greater stability in asset class returns, as suggested in Parker et al. (2020).

Finally, our result that the lifecycle pattern of retirement contribution rates has been relatively stable suggests that the many changes both in the design of retirement saving plans and in the regulatory environment have had little effect on retirement wealth accumulation other than through portfolio choices and resulting returns.

Figures

Figure 2-1: Individual Retirement Wealth Distribution in Firm Data and the SCF in 2016

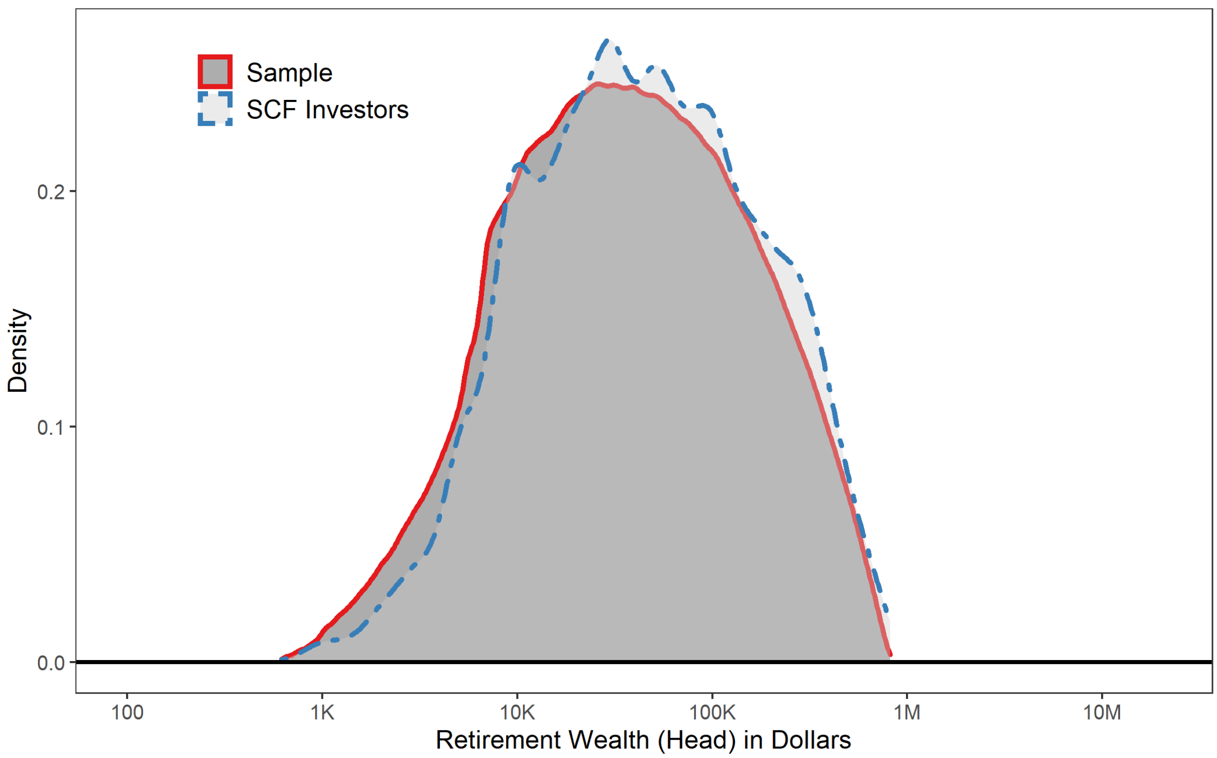

Notes: This figure plots the distribution of retirement wealth in the sample of retirement investors (RIs) versus the distribution of retirement wealth for RIs in the SCF in 2016. Retirement wealth is defined as any wealth in retirement saving accounts of all types (excluding defined benefit plans and Social Security).

Figure 2-2: Investable Wealth Distribution in Firm Data and the SCF in 2016

(a) Investable wealth, individual versus household

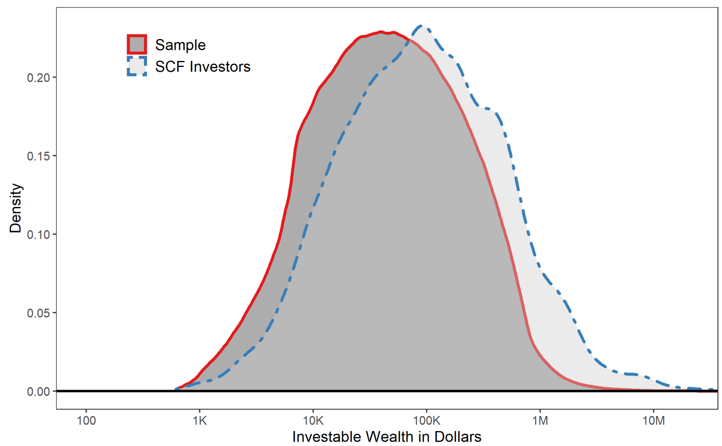

(b) Investable wealth, household versus household

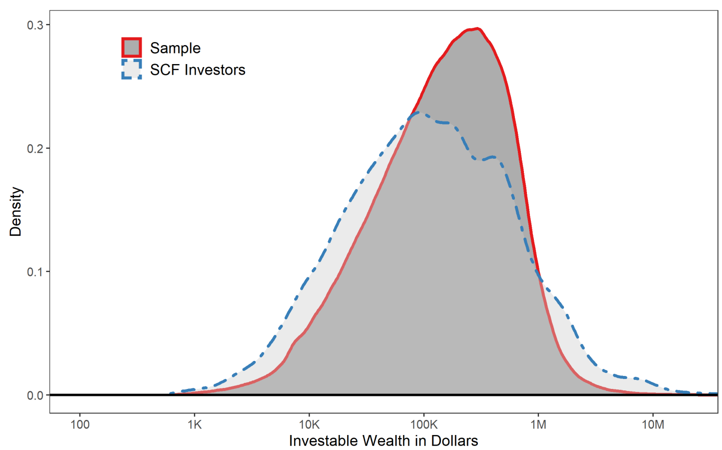

Notes: This figure plots the distribution of investable wealth of retirement investors (RIs) versus the distribution of investable wealth for RIs in the SCF in 2016. The top panel shows individual investable wealth in our sample versus household investable wealth in the SCF. The bottom panel shows household investable wealth in our sample for the subset of households in which we observe both spouses versus household wealth in the SCF for the subsample of investors who are married. Investable wealth is defined as money market funds, non-money market funds, individual stocks and bonds, certificates of deposit, quasi-liquid retirement wealth, and other managed accounts.

Figure 2-3: Portfolio Equity Share by Year and Age

(a) Equity Share by Year

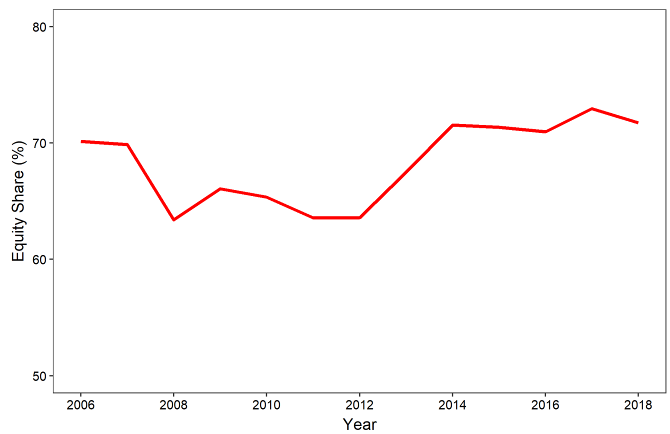

(b) Equity Share by Age

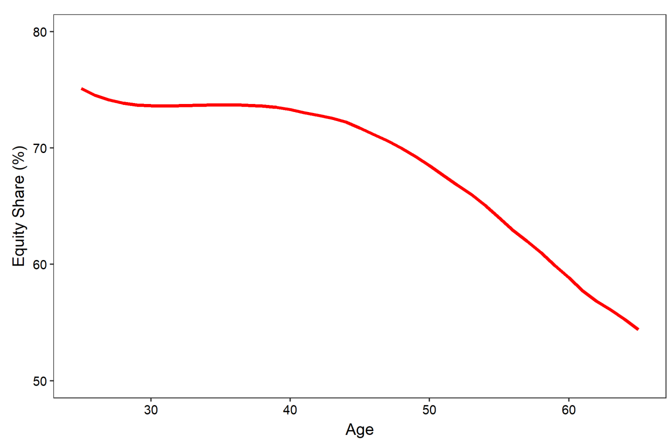

Notes: This figure shows the portfolio equity share in our sample. The top panel shows the equity share for the entire sample, averaged by year. The bottom panel shows the equity share averaged by age. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. The sample is our full set of retirement investors (RI).

Figure 2-4: Portfolio Equity Share by Birth Cohort

(a) Equity Share by Birth Cohort and Year

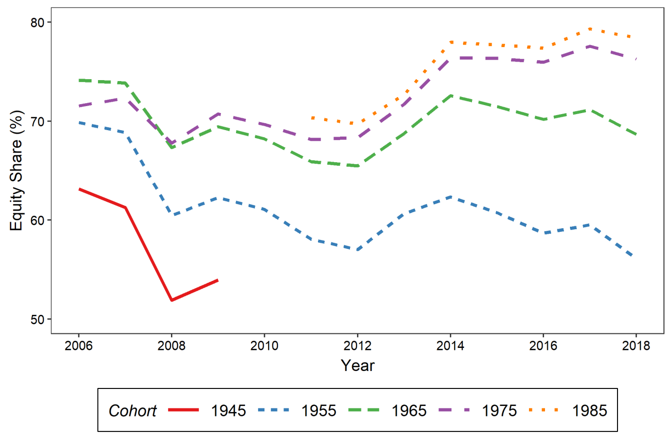

(b) Equity Share by Birth Cohort and Age

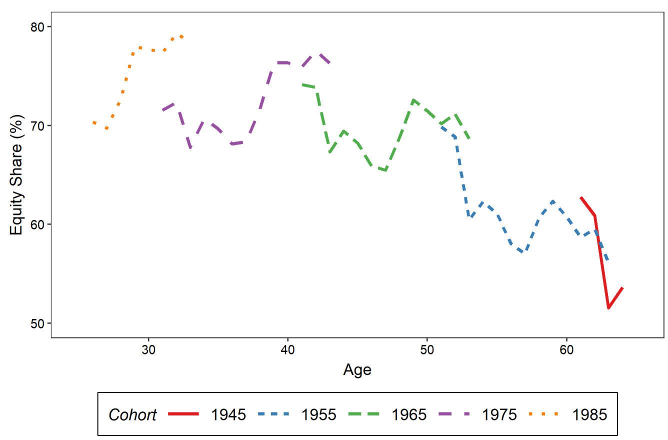

Notes: These figures show the portfolio equity share averaged by birth year cohorts. The top panel shows the averages by year over our sample period. We include only years during which each member of the cohort is aged 25-65, unless otherwise indicated. The bottom panel shows the averages by age, where age is the median age of the cohort. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. A cohort is defined as having been born in the three-year period centered around the year indicated. The sample is our full set of retirement investors (RI).

Figure 2-5: Target Date Fund Share by Birth Cohort

(a) TDF Share by Birth Cohort and Year

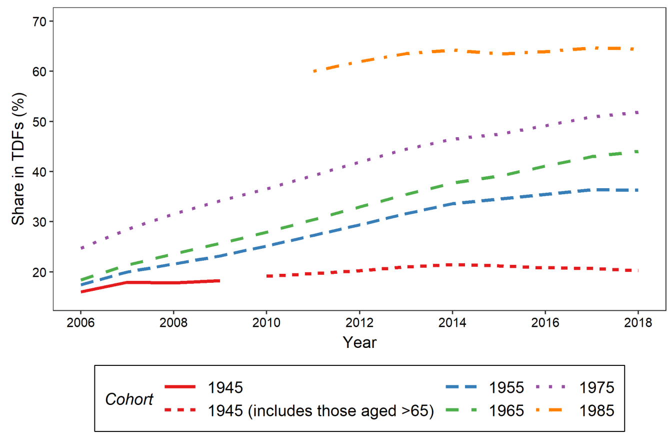

(b) TDF Share by Birth Cohort and Age

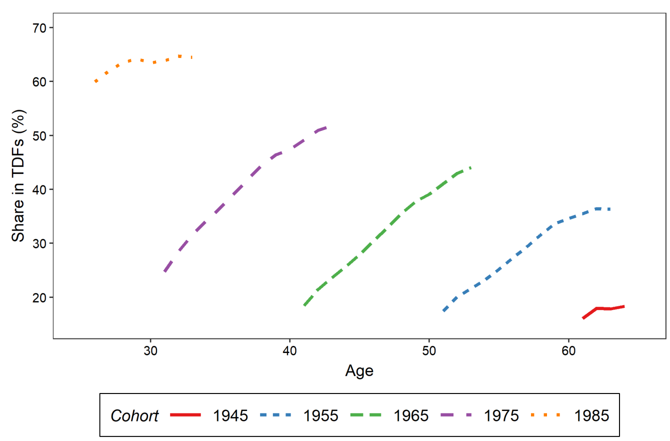

Notes: These figures show the share of the portfolios that is invested in Target Date Funds (TDF) averaged by birth year cohorts. The top panel shows the averages by year over our sample period. We include only years during which each member of the cohort is aged 25-65, unless otherwise indicated. The bottom panel shows the averages by age, where age is the median age of the cohort. TDFs are mutual funds that maintain a given portfolio share of assets invested in different asset classes, where the shares change with the number of years until ‘target date,’ the expected retirement date of the investor. A cohort is defined as having been born in the three-year period centered around the year indicated. The sample is our full set of retirement investors (RI).

Figure 2-6: Equity Share Among Equity Owners

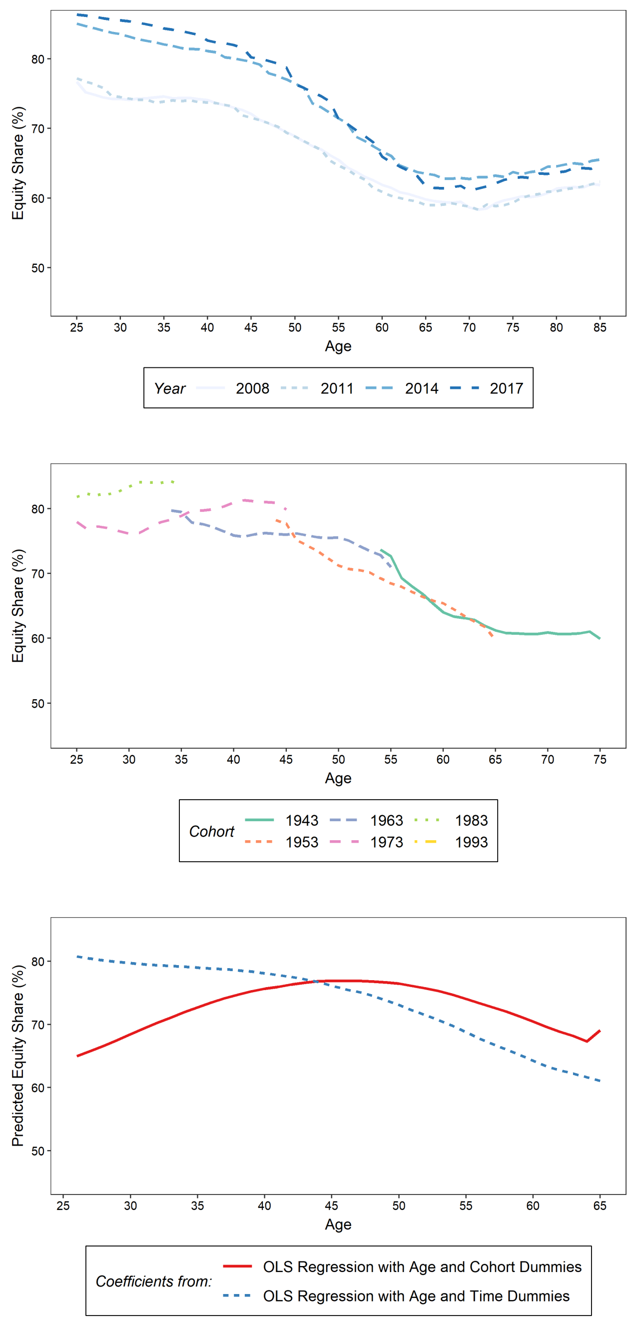

Notes: This figure replicates the results shown in Figure 12 of Ameriks and Zeldes (2004). The top figure shows the observed equity share by age in four different years of our sample. The middle figure shows the observed equity share by age in each cohort in our sample. A cohort is defined as having been born in the ten-year period beginning with the year indicated. The bottom figure shows the predicted values from a regression of equity share on indicator variables for age and either cohort or time. We obtain the predicted values by adding the median cohort or year coefficient, respectively, to each age coefficient. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. The sample is our set of retirement investors (RI) who own at least some equity.

Figure 2-7: Predicted Equity Share: Pension Protection Act

(a) Age Enrolled 25-34

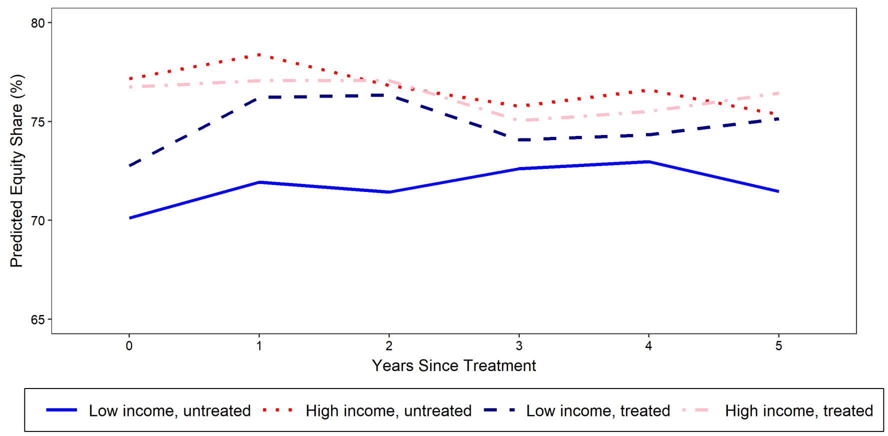

(b) Age Enrolled 55-65

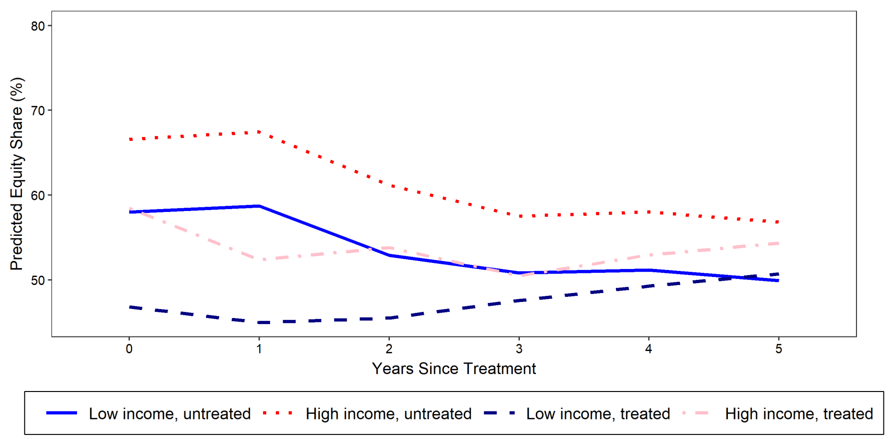

Notes: This figure shows the predicted equity share for those treated by the Pension Protection Act of 2006 and those not treated by the act, split out by age and income groups. The top panel shows the results for those aged 25-34 when enrolled. The bottom panel shows the results for those aged 55-65 when enrolled. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. The sample is our set of retirement investors (RI) who were enrolled between 2005-2008.

Figure 2-8: Realized Contribution Rate by Birth Cohort

(a) Realized Contribution Rate by Birth Cohort and Year

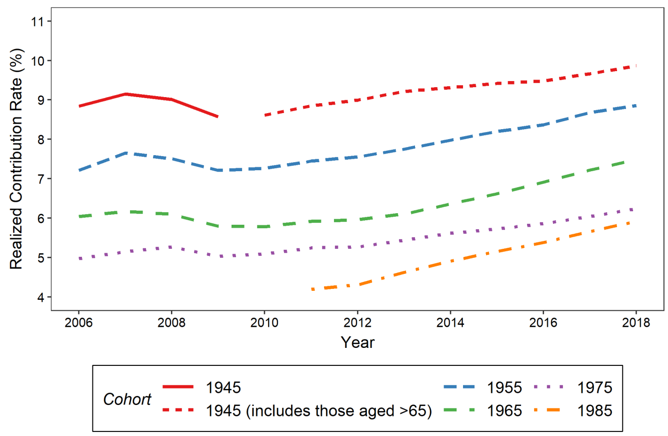

(b) Realized Contribution Rate by Birth Cohort and Age

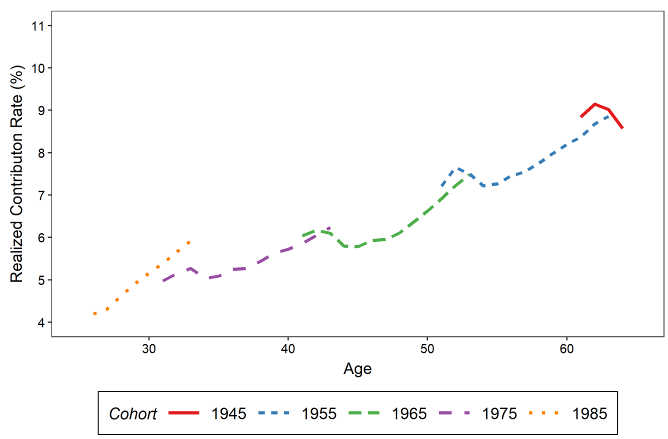

Notes: These figures show the realized contribution rate averaged by birth year cohorts. The top panel shows the averages by year over our sample period. We include only years during which each member of the cohort is aged 25-65, unless otherwise indicated. The bottom panel shows the averages by age, where age is the median age of the cohort. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. A cohort is defined as having been born in the three-year period centered around the year indicated. The sample is our full set of retirement investors (RI).

Figure 2-9: Reported Contribution Rate by Birth Cohort

(a) Reported Contribution Rate by Birth Cohort and Year

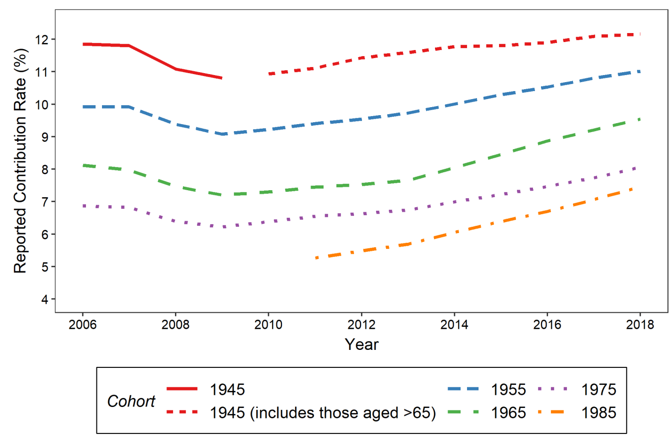

(b) Reported Contribution Rate by Birth Cohort and Age

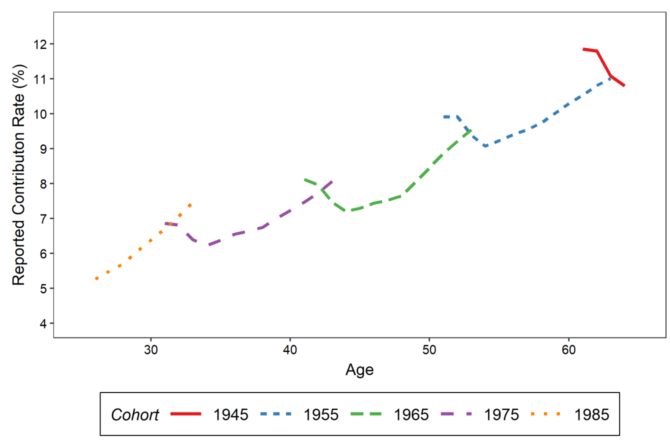

Notes: These figures show the reported contribution rate averaged by birth year cohorts. The top panel shows the averages by year over our sample period. We include only years during which each member of the cohort is aged 25-65, unless otherwise indicated. The bottom panel shows the averages by age, where age is the median age of the cohort. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. A cohort is defined as having been born in the three-year period centered around the year indicated. The sample is our full set of retirement investors (RI).

Figure 2-10: Predicted Contribution Rate: Pension Protection Act

(a) Age Enrolled 25-34

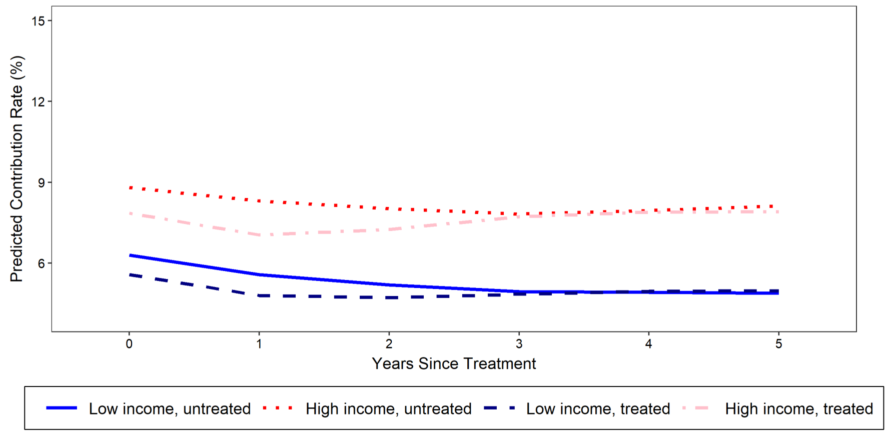

(b) Age Enrolled 55-65

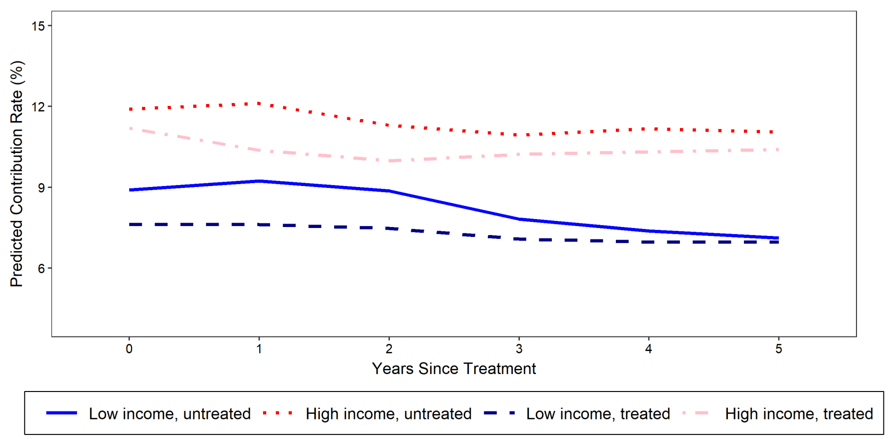

Notes: This figure shows the predicted contribution rate for those treated by the Pension Protection Act of 2006 and those not treated by the act, split out by age and income groups. The top panel shows the results for those aged 25-34 when enrolled. The bottom panel shows the results for those aged 55-65 when enrolled. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The sample is our set of retirement investors (RI) who were enrolled between 2005-2008.

Tables

Table I: Definitions of Key Variables

Variable	**Definition**
Investable wealth	The dollar value of the following assets, measured at the end of each calendar year and summed across retirement funds, individual brokerage accounts, and accounts managed by a financial advisor: money market mutual funds, non-money market funds (including mutual funds and ETFs), individual stocks and bonds, certificate of deposits, and trusts. The measure excludes bank accounts (checking and saving), savings bonds, cash value of life insurance, durable goods, and housing.
Retirement wealth	The dollar value of all wealth in retirement saving accounts of all types, measured at the end of each calendar year. This includes 401K and 403B plans, IRAs, and other Thrift plans. It excludes defined benefit plans and social security.
Non-retirement wealth	The dollar value of all investable wealth that is not retirement wealth, measured at the end of each calendar year. It includes individual stocks, bonds, money market mutual funds, and non-money market funds (including mutual funds and ETFs), certificates of deposit and trusts that are not held in retirement accounts.
Labor income	The dollar value of gross labor/wage income (pre-tax) earned by the head of household, annualized by scaling up part-year incomes to a full-year equivalent. In the SCF, the sum of wages from the head of household's first and second job and self-employment income. Both measures exclude rental income, dividends, royalties, and any income that is not payment for labor. When included in regressions, we normalize income by taking the log deviation of labor income from the RI sample average in the same year.
Retirement Share of Wealth	Total retirement wealth divided by total investable wealth at the end of each calendar year.
Target Date Funds (TDF)	Mutual funds that maintain a given portfolio share of assets invested in different asset classes, where the shares change with the number of years until 'target date,' the expected retirement date of the investor, sometimes referred to as "hybrid", "combination", "auto-rebalancing", or "mixed" funds.
TDF Share of Investable Wealth	Total dollar value of TDFs in the portfolio divided by total investable wealth at the end of each calendar year.
Employment Tenure	The number of years that an employee has been working for their current employer, available for a subset of our sample for which labor income is available.

Definitions of key variables, continued

Variable	Definition
Equity share of retirement wealth	The percentage share of the retirement wealth at the end of each calendar year that is invested in equities and equity-like securities such as individual stocks, equity mutual funds, and the equity component of blended funds (TDFs and auto-rebalancing funds).
Equity share of investable wealth	The percentage share of investable wealth at the end of the calendar year that is invested in equities and equity-like securities such as individual stocks, equity mutual funds, and the equity component of blended funds (TDFs and auto-rebalancing funds).
Equity share of non-retirement wealth	The percentage share of non-retirement wealth at the end of the calendar year that is invested in equities and equity-like securities such as individual stocks, equity mutual funds, and the equity component of blended funds (TDFs and auto-rebalancing funds).
Long-term bonds (fixed income)	Bond funds, long-term government and corporate bonds, and the portion of funds that invest across asset classes (TDFs and auto-rebalancing funds) that is not allocated to equity.
Short-term bonds (cash-like securities)	Money market mutual funds, short-term treasury bonds, and CDs.
Market betas	Using all available return data from 2006 to 2018, we estimate betas from monthly regressions of excess asset returns on excess market returns. We require at least 24 monthly return observations. We set the market beta of short-term bonds to zero. We use the estimated beta on a corresponding ETF as a proxy for individual betas on agency bonds (ticker: AGZ), municipal bonds (MUB), TIPS (TIP), gold (IAU), silver (SLV), and platinum (PPLT). For mixed-asset funds, we account for time variation in betas due to a changing equity share of the portfolio (especially for lifecycle funds) by assuming that the fund market beta is affine in the fund equity share with a fund-specific intercept and a common slope. We estimate the common slope in a pooled regression that includes all mixed-asset funds in an investor's portfolio.
Reported contribution rate	The elected retirement saving rate as a fraction of labor income in employment-based accounts, reported at a monthly frequency. We use the value reported in January for our annual data. This is available only for the subset of the sample for which labor income is observed.
Realized contribution rate	The sum of all flows into retirement accounts in a given year, as a fraction of annual realized labor income. This is calculated only for the subset of the sample for which labor income is observed.

Table II: Characteristics of Sample of Retirement Investors in 2016

Retirement Investors				
	Summary Statistics			
	Mean	Median	SD	Percentage of RI Sample with Observed Data
Age (Years)	45.38	46	11.28	100%
Female	0.46	0	0.50	94.0%
Married	0.72	1	0.45	89.5%
Labor Income ($)	101,384	74,230	195,060	41.0%
Investable Wealth ($)	116,938	38,394	367,156	100%
Retirement Wealth ($)	95,654	35,451	155,237	100%
Retirement Share of Wealth (%)	96.3	100	13.9	100%
Portfolio Beta	0.75	0.84	0.32	86.9%
TDF Share of Invest. Wealth (%)	47.9	37.3	44.7	99.6%
Employment Tenure (Years)	10.50	7.94	9.17	60.0%
Reported Contribution Rate (%)	8.1	6.0	7.3	53.2%
Realized Contribution Rate (%)	6.4	5.5	5.3	47.1%
Retirement Investors - Survey of Consumer Finance				
	Summary Statistics			
	Mean	Median	SD	Number of Observations
Age	46.78	47	10.63	3130
Female	0.50	0	0.50	3130
Married	0.78	1	0.39	3130
Labor Income (Individual, $)	66,459	50,000	1,129,486	3130
Labor Income (Household, $)	101,349	77,000	1,445,913	1889
Investable Wealth (Household, $)	273,282	72,000	17,019,097	1889
Retirement Wealth (Household, $)	193,568	76,830	659,727	1889
Retirement Wealth (Individual, $)	97,658	41,500	155,503	3130
Retirement Share of Investable Wealth (Individual, %)	65.32	76.19	38.43	3130
Retirement Share of Investable Wealth (Household, %)	87.81	100.00	31.33	1889

Notes: This table presents summary statistics on demographics, wealth, and portfolio allocations for our Retirement Investor (RI) sample in 2016 and a comparable sample of the 2016 Survey of Consumer Finance (SCF). Detailed definitions for retirement wealth and investable wealth are provided in Table I. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Market betas are obtained by regressing monthly fund or security excess returns on the value-weighted CRSP market excess return over the period 2007–2017 with at least 24 observations. Income is the labor income of the respondent in 2015. The sample is not representative of the assets under management of our financial service firm, since by design we drop the highest and lowest income groups.

Table III: Average Share of Equity in Portfolios Among Retirement Investors

	All Retirement Investors		Retirement Investors with Hybrid Fund (e.g. TDF) in Retirement Account	
Panel A: All Investable Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	71.0	54.5	76.6	46.9
Age 25-34	77.6	59.1	84.8	49.6
Age 35-44	76.0	55.9	82.2	47.9
Age 45-54	71.2	53.8	74.7	45.5
Age 55-65	60.5	51.2	61.2	45.4
Respondents		54.3		47.0
Partners		54.8		46.9
Panel B: Retirement Wealth	Main Sample (Individuals)	SCF (Individuals)	Main Sample (Individuals)	SCF (Individuals)
All RIs	71.1	51.7	76.7	42.1
Age 25-34	77.7	56.2	85.0	44.2
Age 35-44	76.2	54.1	82.4	43.5
Age 45-54	71.4	50.5	74.8	40.2
Age 55-65	60.6	48.0	61.2	41.2
Respondents		52.1		43.3
Partners		50.8		39.8
Panel C: Non-Retirement Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	51.1	73.4	53.2	73.2
Age 25-34	52.0	87.5	53.0	86.9
Age 35-44	53.5	68.9	55.5	68.3
Age 45-54	51.1	74.5	52.9	73.6
Age 55-65	48.8	69.6	50.8	69.6
Respondents		73.9		74.4
Partners		72.7		71.2

Notes: This table presents the share of equity in the portfolio allocations for various samples of our Retirement Investors (RI) sample in 2016 and the comparable RI sample of the 2016 Survey of Consumer Finance (SCF). Panel A shows equity shares of total investable wealth at the individual level in our sample and the household level in the SCF. Panel B shows equity shares of retirement wealth, at the individual level in both datasets. Panel C shows equity shares of non-retirement wealth at the individual level in our sample and the household level in the SCF. The figures in Panel C are conditional on owning some non-retirement wealth, which is approximately 40% of the SCF RI sample and 16% of our RI sample. The first two columns show the means for the full sample of RIs in each dataset. The last two columns show the means for the subsample of the RI sample that has some of their retirement assets in a target date fund (TDF). Investable wealth is defined as money market funds, non-money market funds, individual stocks and bonds, Retirement wealth is defined as any wealth in retirement saving accounts of all types (excluding defined benefit plans and Social Security). certificate of deposits, quasi-liquid retirement wealth, and other managed accounts. The equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets.

Table IV: Cross-Sectional Regressions of Equity Share, Full Sample and by Income Terciles

	Portfolio equity share				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.7366	0.8031	0.7489	0.7915	0.7943
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0005)
Age 28-30	0.7326	0.7964	0.7321	0.7797	0.7865
	(0.0001)	(0.0002)	(0.0003)	(0.0002)	(0.0004)
Age 31-33	0.7331	0.7888	0.7272	0.7724	0.7790
	(0.0001)	(0.0001)	(0.0003)	(0.0002)	(0.0003)
Age 34-36	0.7348	0.7816	0.7253	0.7674	0.7730
	(0.0001)	(0.0001)	(0.0003)	(0.0002)	(0.0003)
Age 37-39	0.7344	0.7731	0.7208	0.7614	0.7681
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Age 40-42	0.7296	0.7615	0.7118	0.7515	0.7607
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0002)
Age 43-45	0.7209	0.7479	0.6990	0.7383	0.7509
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0002)
Age 46-48	0.7053	0.7280	0.6787	0.7172	0.7341
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0002)
Age 49-51	0.6844	0.7022	0.6542	0.6903	0.7102
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0002)
Age 52-54	0.6598	0.6738	0.6263	0.6602	0.6818
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0002)
Age 55-57	0.6304	0.6402	0.5923	0.6244	0.6482
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Age 58-60	0.6002	0.6063	0.5593	0.5869	0.6121
	(0.0001)	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Age 61-63	0.5702	0.5730	0.5250	0.5486	0.5765
	(0.0002)	(0.0002)	(0.0004)	(0.0004)	(0.0004)
Age 64-65	0.5496	0.5482	0.4969	0.5173	0.5485
	(0.0002)	(0.0003)	(0.0005)	(0.0005)	(0.0005)
Log income		0.0761			
		(0.0003)			
Person fixed effect?	N	N	N	N	N
% of RI Sample	93.4	40.9	15.8	16.7	16.2
R-squared	0.0379	0.0751	0.0553	0.0744	0.0609

Notes: This table presents regression coefficients of annual individual portfolio equity shares on a set of demographic controls. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. The baseline specification in column (1) shows the coefficients for the regression of equity share on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table V: Within-Person Regressions of Equity Share, Full Sample and by Income Terciles

	Portfolio equity share				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.6758	0.6624	0.6336	0.6777	0.6784
	(0.0004)	(0.0006)	(0.0009)	(0.0009)	(0.0010)
Age 28-30	0.6854	0.6775	0.6273	0.6862	0.6945
	(0.0004)	(0.0005)	(0.0009)	(0.0008)	(0.0009)
Age 31-33	0.7042	0.7003	0.6395	0.7006	0.7099
	(0.0003)	(0.0005)	(0.0009)	(0.0008)	(0.0008)
Age 34-36	0.7227	0.7219	0.6572	0.7149	0.7213
	(0.0003)	(0.0005)	(0.0009)	(0.0008)	(0.0007)
Age 37-39	0.7370	0.7394	0.6735	0.7267	0.7302
	(0.0003)	(0.0005)	(0.0008)	(0.0008)	(0.0007)
Age 40-42	0.7460	0.7519	0.6864	0.7344	0.7348
	(0.0003)	(0.0005)	(0.0008)	(0.0007)	(0.0007)
Age 43-45	0.7517	0.7613	0.6968	0.7397	0.7364
	(0.0003)	(0.0004)	(0.0008)	(0.0007)	(0.0006)
Age 46-48	0.7519	0.7647	0.7011	0.7391	0.7325
	(0.0003)	(0.0004)	(0.0007)	(0.0007)	(0.0006)
Age 49-51	0.7486	0.7637	0.7026	0.7344	0.7242
	(0.0002)	(0.0004)	(0.0007)	(0.0006)	(0.0006)
Age 52-54	0.7397	0.7559	0.6964	0.7227	0.7102
	(0.0002)	(0.0004)	(0.0006)	(0.0006)	(0.0005)
Age 55-57	0.7253	0.7412	0.6833	0.7029	0.6907
	(0.0002)	(0.0004)	(0.0006)	(0.0006)	(0.0005)
Age 58-60	0.7071	0.7220	0.6661	0.6771	0.6664
	(0.0002)	(0.0003)	(0.0005)	(0.0005)	(0.0005)
Age 61-63	0.6845	0.6987	0.6420	0.6465	0.6389
	(0.0001)	(0.0002)	(0.0004)	(0.0004)	(0.0004)
Age 64-65	0.6635	0.6752	0.6159	0.6168	0.6132
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log income		0.0365			
		(0.0003)			
Person fixed effect?	Y	Y	Y	Y	Y
% of RI Sample	93.4	40.9	15.8	16.7	0.162
R-squared	0.7561	0.7742	0.7742	0.7372	0.6876

Notes: This table presents regression coefficients of annual individual portfolio equity shares on a set of demographic controls. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. The baseline specification in column (1) shows the coefficients for the regression of equity share on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. All regressions include a person fixed effect. The age group coefficients are normalized by adding the average fixed effect back to the estimated coefficients. The excluded age group is those aged 64-65. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table VI: Within-Person Regressions of Equity Share on Age Groups by Cohort and TDF Share

	Portfolio equity share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1943 Cohort	1953 Cohort	1963 Cohort	1973 Cohort	1983 Cohort	Initial TDF Share 75-100 %	Initial TDF Share 25-75 %	Initial TDF Share 0-25 %
Age 25-27				0.7153	0.7953	0.6582	0.5457	0.6264
				(0.0006)	(0.0004)	(0.0007)	(0.0011)	(0.0009)
Age 28-30				0.7208	0.8122	0.6765	0.5607	0.6276
				(0.0004)	(0.0004)	(0.0006)	(0.0010)	(0.0008)
Age 31-33				0.7379	0.8420	0.6930	0.5769	0.6378
				(0.0004)	(0.0003)	(0.0005)	(0.0009)	(0.0008)
Age 34-36			0.7268	0.7594	0.8629	0.7085	0.5956	0.6553
			(0.0005)	(0.0004)	(0.0001)	(0.0005)	(0.0008)	(0.0008)
Age 37-39			0.7142	0.7802		0.7169	0.6106	0.6717
			(0.0004)	(0.0003)		(0.0004)	(0.0008)	(0.0007)
Age 40-42			0.7101	0.8094		0.7193	0.6222	0.6851
			(0.0003)	(0.0002)		(0.0004)	(0.0007)	(0.0007)
Age 43-45		0.7859	0.7215	0.8176		0.7185	0.6316	0.6963
		(0.0008)	(0.0003)	(0.0001)		(0.0004)	(0.0007)	(0.0007)
Age 46-48		0.7449	0.7269			0.7128	0.6356	0.7011
		(0.0006)	(0.0003)			(0.0003)	(0.0007)	(0.0006)
Age 49-51		0.7182	0.7415			0.7030	0.6358	0.7041
		(0.0006)	(0.0002)			(0.0003)	(0.0006)	(0.0006)
Age 52-54	0.7366	0.7116	0.7377			0.6886	0.6304	0.7015
	(0.0010)	(0.0006)	(0.0001)			(0.0003)	(0.0006)	(0.0006)
Age 55-57	0.6802	0.6974				0.6690	0.6184	0.6927
	(0.0006)	(0.0006)				(0.0003)	(0.0006)	(0.0005)
Age 58-60	0.6192	0.6876				0.6448	0.6015	0.6831
	(0.0004)	(0.0005)				(0.0002)	(0.0005)	(0.0005)
Age 61-63	0.5880	0.6749				0.6153	0.5770	0.6701
	(0.0003)	(0.0005)				(0.0002)	(0.0004)	(0.0004)
Age 64-65	0.5505	0.6615				0.5887	0.5518	0.6557
	(0.0000)	(0.0000)				(0.0001)	(0.0001)	(0.0001)
Log income	0.0274	0.0226	0.0256	0.0407	0.0662			
	(0.0012)	(0.0006)	(0.0006)	(0.0006)	(0.0008)			
Person fixed effect?	Y	Y	Y	Y	Y	Y	Y	Y
% of RI Sample	3.1	10.9	11.5	10.3	5.0	39.5	7.9	10.1
R-squared	0.7948	0.7627	0.7537	0.7420	0.7343	0.7457	0.6769	0.6892

Notes: This table presents regression coefficients of annual individual portfolio equity shares on a set of demographic controls. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. Columns (1)-(5) show the results including age-group controls and a control for log income, broken out by birth cohort groups. Log income is measured as the log deviation of the individual's income from the average income of the RI sample. A cohort is defined as having been born in the ten year period beginning with the year indicated. Columns (6)-(8) show the results for different groups based on the initial share of their portfolio that is invested in target date funds (TDFs). All regressions include a person fixed effect. The age group coefficients are normalized by adding the average fixed effect back to the estimated coefficients. The excluded age group is those aged 64-65. The sample is our full RI sample from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table VII: Regressions of Equity Share on Automated Investment Allocation: Average Effect Two Years After Entering Sample

	Portfolio equity share					
	(1) All	(2) All	(3) Bottom Income Tercile	(4) Top Income Tercile	(5) No prior non-retirement wealth + no rollover assets	(6) No prior non-retirement wealth + no rollover assets
Treated	0.0552	0.0533	0.0599	0.0186	0.0578	0.0555
	(0.0007)	(0.0008)	(0.0011)	(0.0024)	(0.0007)	(0.0008)
Age 35-44	-0.0134	-0.0271	-0.0339	-0.0112	-0.0140	-0.0272
	(0.0006)	(0.0007)	(0.0013)	(0.0012)	(0.0006)	(0.0007)
Age 45-54	-0.0700	-0.0875	-0.1011	-0.0627	-0.0720	-0.0887
	(0.0007)	(0.0008)	(0.0015)	(0.0014)	(0.0007)	(0.0009)
Age 55-65	-0.1325	-0.1502	-0.1658	-0.1254	-0.1352	-0.1520
	(0.0012)	(0.0014)	(0.0026)	(0.0023)	(0.0013)	(0.0015)
Age 35-44 x Treatment	-0.0581	-0.0542	-0.0508	-0.0366	-0.0600	-0.0549
	(0.0014)	(0.0016)	(0.0024)	(0.0041)	(0.0015)	(0.0017)
Age 45-54 x Treatment	-0.1029	-0.0885	-0.0717	-0.0809	-0.1042	-0.0895
	(0.0018)	(0.0021)	(0.0030)	(0.0050)	(0.0018)	(0.0021)
Age 55-65 x Treatment	-0.1479	-0.1314	-0.1235	-0.1173	-0.1495	-0.1322
	(0.0032)	(0.0038)	(0.0055)	(0.0091)	(0.0033)	(0.0038)
Log income		0.1031				0.1072
		(0.0012)				(0.0013)
Constant	0.7352	0.7476	0.7180	0.7353	0.7335	0.7468
	(0.0003)	(0.0004)	(0.0006)	(0.0003)	(0.0003)	(0.0004)
Firm Fixed Effect?	Y	Y	Y	Y	Y	Y
% of Total Sample	1.3	0.9	0.3	0.3	1.2	0.9
% of Sample Enrolled 2005-2008	18.1	12.8	5.0	3.9	17.0	12.2
R-squared	0.1543	0.1502	0.2266	0.1044	0.1620	0.1565

Notes: This table presents regression coefficients of annual household portfolio equity shares on a treatment dummy for being enrolled into a plan with a target date fund (TDF) as the default after the Pension Protection Act of 2006. We set this treatment dummy equal to one for those enrolled in their firm's retirement plan in 2007 or 2008 when that plan had a TDF as a default and zero for those enrolled in 2005 or 2006. Columns (1)-(2) show the results for the first two years of data after the individual enters our sample. Columns (3)-(4) repeat column (1) for those in the lowest and highest tercile of initial income, respectively. Columns (5)-(6) repeat columns (1)-(2) including only individuals who had no prior retirement wealth before enrollment and no rollover assets of any kind. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. Log income, when included, is the log deviation of the individual's current income from the average income of the RI sample. The sample is our set of retirement investors (RI) who enrolled in their plan from 2005-2008. Standard errors, in parentheses, are clustered at the household level.

Table VIII: Regressions of Equity Share on Automated Investment Allocation: Long-run Effect

	Portfolio equity share						
	(1) Full Sample	(2) Bottom Income Tercile	(3) Top Income Tercile	(4) Age Enrolled 25-34	(5) Age Enrolled 35-44	(6) Age Enrolled 45-54	(7) Age Enrolled 55-65
Year of x Treatment	0.0198	0.0463	0.0071	0.0161	-0.0263	-0.0584	-0.0618
	(0.0023)	(0.0036)	(0.0051)	(0.0022)	(0.0066)	(0.0092)	(0.0171)
1 Year After x Treatment	0.0430	0.0655	-0.0065	0.0295	-0.0154	-0.0750	-0.1321
	(0.0011)	(0.0015)	(0.0029)	(0.0011)	(0.0036)	(0.0068)	(0.0097)
2 Years After x Treatment	0.0683	0.0861	0.0287	0.0363	0.0186	-0.0181	-0.0515
	(0.0009)	(0.0013)	(0.0025)	(0.0011)	(0.0028)	(0.0050)	(0.0116)
3 Years After x Treatment	0.0032	0.0254	-0.0269	0.0005	-0.0261	-0.0268	-0.0524
	(0.0010)	(0.0015)	(0.0021)	(0.0013)	(0.0018)	(0.0025)	(0.0046)
4 Years After x Treatment	-0.0244	-0.0185	-0.0292	0.0003	-0.0173	-0.0274	-0.0410
	(0.0010)	(0.0015)	(0.0019)	(0.0013)	(0.0014)	(0.0020)	(0.0039)
5 Years After x Treatment	0.0036	0.0104	-0.0055	0.0257	0.0099	-0.0060	-0.0201
	(0.0010)	(0.0015)	(0.0023)	(0.0014)	(0.0016)	(0.0022)	(0.0043)
1 Year After	0.0087	0.0209	0.0126	0.0136	0.0133	0.0096	0.0088
	(0.0009)	(0.0018)	(0.0013)	(0.0012)	(0.0014)	(0.0016)	(0.0026)
2 Years After	-0.0194	0.0022	-0.0201	0.0067	-0.0207	-0.0445	-0.0564
	(0.0010)	(0.0019)	(0.0014)	(0.0012)	(0.0015)	(0.0017)	(0.0029)
3 Years After	-0.0272	0.0046	-0.0383	0.0142	-0.0282	-0.0675	-0.0876
	(0.0010)	(0.0019)	(0.0015)	(0.0013)	(0.0016)	(0.0018)	(0.0030)
4 Years After	-0.0221	0.0074	-0.0334	0.0243	-0.0209	-0.0629	-0.0810
	(0.0010)	(0.0019)	(0.0015)	(0.0013)	(0.0016)	(0.0018)	(0.0031)
5 Years After	-0.0348	-0.0091	-0.0459	0.0134	-0.0373	-0.0824	-0.0955
	(0.0011)	(0.0020)	(0.0016)	(0.0014)	(0.0017)	(0.0019)	(0.0034)
Log income	0.0487						
	(0.0012)						
Constant	0.7279	0.6751	0.7473	0.7255	0.7432	0.7059	0.6374
	(0.0010)	(0.0020)	(0.0014)	(0.0013)	(0.0015)	(0.0017)	(0.0028)
Firm Fixed Effect?	Y	Y	Y	Y	Y	Y	Y
% of RI Sample	1.6	0.5	0.5	0.8	0.6	0.4	0.1
% of Sample Enrolled 2005-2008	22.4	7.8	7.1	11.9	8.2	5.8	1.9
R-squared	0.0969	0.1727	0.0716	0.1537	0.1013	0.0942	0.1181

Notes: This table presents regression coefficients of annual household portfolio equity shares on being treated with the Pension Protection Act (PPA) of 2006. "Year of" means the year the individual enrolled in their retirement plan and "x years after" is x years after they enrolled in the plan. Each column includes year dummies for each year after enrollment, and interactions of these dummies with the treatment dummy. The treatment dummy is equal to one if the individual enrolled in 2007 or 2008 to a plan that switched to having a target date fund as the default following the PPA and zero if they enrolled in 2005 or 2006. The full sample is those enrolled from 2005-2008 who otherwise meet the RI sample criteria. The bottom (top) income tercile includes those whose initial income is in the lowest (highest) tercile. Columns (4)-(7) break out the result for all individuals enrolled from 2005-2008 by age at enrollment. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. Log income, when included, is the log deviation of the individual's current income from the average income of the RI sample. Standard errors, in parentheses, are clustered at the household level.

Table IX: Cross-Sectional Regressions of Realized Contribution Rate, Full Sample and by Income Terciles

	Realized contribution rate				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.0456	0.0512	0.0393	0.0520	0.0569
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Age 28-30	0.0497	0.0540	0.0425	0.0545	0.0613
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Age 31-33	0.0526	0.0558	0.0445	0.0555	0.0629
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Age 34-36	0.0545	0.0568	0.0461	0.0558	0.0632
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Age 37-39	0.0560	0.0578	0.0474	0.0564	0.0634
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Age 40-42	0.0576	0.0590	0.0490	0.0576	0.0639
	(0.0000)	(0.0000)	(0.0001)	(0.0000)	(0.0000)
Age 43-45	0.0596	0.0608	0.0514	0.0596	0.0650
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0000)
Age 46-48	0.0617	0.0629	0.0538	0.0622	0.0664
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0000)
Age 49-51	0.0662	0.0674	0.0569	0.0662	0.0719
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0000)
Age 52-54	0.0713	0.0727	0.0604	0.0711	0.0782
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 55-57	0.0752	0.0768	0.0637	0.0756	0.0822
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 58-60	0.0792	0.0811	0.0671	0.0805	0.0863
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 61-63	0.0833	0.0855	0.0712	0.0857	0.0902
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 64-65	0.0848	0.0873	0.0734	0.0877	0.0915
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Log income		0.0185			
		(0.0000)			
Person fixed effect?	N	N	N	N	N
% of RI Sample	41.4	41.1	12.0	13.3	12.6
R-squared	0.0472	0.0578	0.0446	0.0390	0.0385

Notes: This table presents regression coefficients of realized contribution rate on a set of demographic controls. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. The baseline specification in column (1) shows the coefficients for the regression of realized contribution rate on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table X: Within-Person Regressions of Realized Contribution Rate, Full Sample and by Income Terciles

	Realized contribution rate				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.0271	0.0142	0.0270	0.0282	0.0436
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 28-30	0.0344	0.0221	0.0331	0.0365	0.0534
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 31-33	0.0400	0.0283	0.0379	0.0427	0.0595
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 34-36	0.0439	0.0325	0.0416	0.0469	0.0626
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 37-39	0.0468	0.0357	0.0444	0.0504	0.0643
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0001)
Age 40-42	0.0493	0.0384	0.0470	0.0538	0.0653
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0001)
Age 43-45	0.0516	0.0409	0.0495	0.0570	0.0663
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0001)
Age 46-48	0.0541	0.0434	0.0520	0.0604	0.0675
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0001)
Age 49-51	0.0586	0.0480	0.0551	0.0650	0.0724
	(0.0001)	(0.0001)	(0.0002)	(0.0001)	(0.0001)
Age 52-54	0.0640	0.0534	0.0588	0.0704	0.0782
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 55-57	0.0686	0.0581	0.0625	0.0757	0.0824
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 58-60	0.0736	0.0630	0.0664	0.0813	0.0870
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 61-63	0.0786	0.0681	0.0707	0.0870	0.0915
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 64-65	0.0817	0.0711	0.0731	0.0900	0.0944
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log income		-0.0108			
		(0.0001)			
Person fixed effect?	Y	Y	Y	Y	Y
% of RI Sample	41.4	41.1	12.0	13.3	12.6
R-squared	0.7684	0.7709	0.7755	0.7635	0.7355

Notes: This table presents regression coefficients of realized contribution rate on a set of demographic controls. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. The baseline specification in column (1) shows the coefficients for the regression of realized contribution rate on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. All regressions include a person fixed effect. The age group coefficients are normalized by adding the average fixed effect back to the estimated coefficients. The excluded age group is those aged 64-65. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table XI: Within-Person Regressions of Realized Contribution Rate on Age Groups by Cohort and TDF Share

	Realized contribution rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1943 Cohort	1953 Cohort	1963 Cohort	1973 Cohort	1983 Cohort	Initial TDF Share 75-100 %	Initial TDF Share 25-75 %	Initial TDF Share 0-25 %
Age 25-27				0.0376	0.0570	0.0448	0.0926	0.0583
				(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
Age 28-30				0.0428	0.0651	0.0510	0.1013	0.0682
				(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
Age 31-33				0.0474	0.0725	0.0550	0.1067	0.0756
				(0.0001)	(0.0001)	(0.0001)	(0.0003)	(0.0003)
Age 34-36			0.0468	0.0513	0.0777	0.0577	0.1101	0.0809
			(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0003)	(0.0003)
Age 37-39			0.0474	0.0549		0.0598	0.1124	0.0851
			(0.0001)	(0.0001)		(0.0001)	(0.0003)	(0.0003)
Age 40-42			0.0488	0.0588		0.0616	0.1145	0.0886
			(0.0001)	(0.0000)		(0.0001)	(0.0003)	(0.0003)
Age 43-45		0.0688	0.0510	0.0607		0.0633	0.1164	0.0919
		(0.0002)	(0.0001)	(0.0000)		(0.0001)	(0.0003)	(0.0003)
Age 46-48		0.0689	0.0538			0.0653	0.1183	0.0950
		(0.0001)	(0.0001)			(0.0001)	(0.0003)	(0.0003)
Age 49-51		0.0716	0.0594			0.0698	0.1223	0.0996
		(0.0001)	(0.0000)			(0.0001)	(0.0002)	(0.0003)
Age 52-54	0.0946	0.0764	0.0690			0.0751	0.1272	0.1051
	(0.0002)	(0.0001)	(0.0001)			(0.0001)	(0.0002)	(0.0003)
Age 55-57	0.0964	0.0811				0.0795	0.1314	0.1098
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 58-60	0.0974	0.0868				0.0841	0.1356	0.1150
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 61-63	0.1006	0.0936				0.0884	0.1397	0.1205
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 64-65	0.1039	0.0941				0.0906	0.1414	0.1238
	(0.0001)	(0.0000)				(0.0001)	(0.0001)	(0.0000)
Log income	-0.0131	-0.0159	-0.0155	-0.0079	0.0039			
	(0.0003)	(0.0002)	(0.0001)	(0.0001)	(0.0002)			
Person fixed effect?	Y	Y	Y	Y	Y	Y	Y	Y
% of RI Sample	3.2	11.0	11.5	10.4	5.0	15.4	3.5	5.2
R-squared	0.8156	0.7798	0.7505	0.7139	0.7412	0.7511	0.7482	0.7396

Notes: This table presents regression coefficients of annual individual realized contribution rates on a set of demographic controls. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Columns (1)-(5) show the results including age-group controls and a control for log income, broken out by birth cohort groups. Log income is measured as the log deviation of the individual's income from the average income of the RI sample. A cohort is defined as having been born in the ten year period beginning with the year indicated. Columns (6)-(8) show the results for different groups based on the initial share of their portfolio that is invested in target date funds (TDFs). All regressions include a person fixed effect. The age group coefficients are normalized by adding the average fixed effect back to the estimated coefficients. The excluded age group is those aged 64-65. The sample is our full RI sample from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table XII: Regressions of Realized Contribution Rate on Maxing Out on Contribution Limit

	Realized contribution rate			
	(1) All Observations	(2) All Observations	(3) All Observations	(4) All Observations
Max Out	0.0584	0.0590		
	(0.0001)	(0.0001)		
Max Out Ever			0.0403	0.0475
			(0.0001)	(0.0001)
Age 25-34	0.0477	0.0475	0.0458	0.0430
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Age 35-44	0.0528	0.0528	0.0495	0.0479
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Age 45-54	0.0613	0.0613	0.0574	0.0561
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Age 55-65	0.0735	0.0735	0.0696	0.0681
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log income		-0.0014		-0.01235
		(0.0001)		(0.0001)
Age 35-44 x Max Out	-0.0173	-0.0174		
	(0.0001)	(0.0001)		
Age 45-54 x Max Out	-0.0184	-0.0185		
	(0.0001)	(0.0001)		
Age 55-65 x Max Out	-0.0056	-0.0057		
	(0.0001)	(0.0002)		
Age 35-44 x Max Out Ever			-0.0031	-0.0027
			(0.0001)	(0.0001)
Age 45-54 x Max Out Ever			-0.0005	0.0003
			(0.0001)	(0.0001)
Age 55-65 x Max Out Ever			0.0104	0.0111
			(0.0001)	(0.0001)
Person fixed effect?	N	N	N	N
Percentage of Total Sample	44.9	41.3	44.9	41.3
R-squared	0.1118	0.1123	0.1473	0.1518

Notes: This table presents regression coefficients of annual realized contribution rates on measures of maxing out on retirement contributions. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Maxing out is defined as when an individual exceeds the dollar amount that is allowed for 401(k) contributions in a year, set by the IRS. Columns (1)-(2) contain a dummy for maxing out that it set to one if the individual maxes out their contribution in the current year. Columns (3)-(4) contain a dummy for maxing out that is set to one if the individual has *ever* maxed out their contribution while we observe them in our sample. Each specification also contains interactions of the corresponding max out measure with age group dummies. Log income is measured in the first (or second, if first is not available) year that we observe the individual. We then take the log deviation of the first year's income from the RI sample's average. The sample is our full RI sample from 2006-2017. Standard errors, in parentheses, are clustered at the household level.

Table XIII: Regressions of Reported Contribution Rate on the Pension Protection Act: Average Effect Two Years After Entering Sample

	Reported contribution rate					
	(1) All	(2) All	(3) Bottom Income Tercile	(4) Top Income Tercile	(5) No prior non-retirement wealth + no rollover assets	(6) No prior non-retirement wealth + no rollover assets
Treated	-0.0043	-0.0034	-0.0028	-0.0073	-0.0042	-0.0034
	(0.0001)	(0.0001)	(0.0001)	(0.0003)	(0.0001)	(0.0001)
Age 35-44	0.0117	0.0084	0.0103	0.0077	0.0112	0.0082
	(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0001)	(0.0001)
Age 45-54	0.0239	0.0203	0.0204	0.0211	0.0229	0.0196
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0002)
Age 55-65	0.0406	0.0367	0.0339	0.0406	0.0389	0.0354
	(0.0003)	(0.0003)	(0.0005)	(0.0005)	(0.0003)	(0.0003)
Age 35-44 x Treatment	-0.0023	-0.0026	-0.0015	-0.0002	-0.0021	-0.0024
	(0.0002)	(0.0002)	(0.0003)	(0.0004)	(0.0002)	(0.0002)
Age 45-54 x Treatment	-0.0045	-0.0047	-0.0028	-0.0029	-0.0038	-0.0042
	(0.0002)	(0.0002)	(0.0004)	(0.0004)	(0.0002)	(0.0002)
Age 55-65 x Treatment	-0.0077	-0.0083	-0.0058	-0.0078	-0.0067	-0.0075
	(0.0004)	(0.0004)	(0.0007)	(0.0008)	(0.0004)	(0.0004)
Log income		0.0314				0.0307
		(0.0002)				(0.0002)
Constant	0.0619	0.0623	0.0508	0.0705	0.0613	0.0621
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Firm Fixed Effect?	Y	Y	Y	Y	Y	Y
% of Total Sample	1.9	1.4	0.6	0.4	1.8	1.3
% of Sample Enrolled 2005-2008	27.5	20.3	7.9	6.1	25.9	19.3
R-squared	0.1737	0.1915	0.1718	0.1242	0.1726	0.1888

Notes: This table presents regression coefficients of reported contribution rate on a treatment dummy for being enrolled into a plan following the Pension Protection Act (PPA) of 2006. We set this treatment dummy equal to one for those enrolled in their firm's retirement plan in 2007 or 2008 and zero for those enrolled in 2005 or 2006. Columns (1)-(2) show the results for the first two years that we observe the individual in our sample. Columns (3)-(4) repeat column (1) for those in the lowest and highest tercile of initial income, respectively. Columns (5)-(6) repeat columns (1)-(2) including only individuals who had no prior retirement wealth before enrollment and no rollover assets of any kind. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. Log income, when included, is the log deviation of the individual's current income from the average income of the RI sample. The sample is our set of retirement investors (RI) who enrolled in their plan from 2005-2008. Standard errors, in parentheses, are clustered at the household level.

Table XIV: Regressions of Reported Contribution Rate on the Pension Protection Act: Long-run Effect

	Reported contribution rate						
	(1) Full Sample	(2) Bottom Income Tercile	(3) Top Income Tercile	(4) Age Enrolled 25-34	(5) Age Enrolled 35-44	(6) Age Enrolled 45-54	(7) Age Enrolled 55-65
Year of x Treatment	-0.0085	-0.0092	-0.0092	-0.0069	-0.0088	-0.0093	-0.0127
	(0.0002)	(0.0004)	(0.0004)	(0.0003)	(0.0004)	(0.0005)	(0.0009)
1 Year After x Treatment	-0.0116	-0.0098	-0.0140	-0.0087	-0.0119	-0.0143	-0.0167
	(0.0001)	(0.0002)	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0006)
2 Years After x Treatment	-0.0072	-0.0074	-0.0091	-0.0055	-0.0071	-0.0101	-0.0124
	(0.0001)	(0.0002)	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0005)
3 Years After x Treatment	-0.0026	-0.0034	-0.0033	-0.0014	-0.0029	-0.0056	-0.0071
	(0.0001)	(0.0001)	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0005)
4 Years After x Treatment	-0.0012	-0.0016	-0.0026	-0.0007	-0.0025	-0.0054	-0.0072
	(0.0001)	(0.0002)	(0.0003)	(0.0001)	(0.0002)	(0.0003)	(0.0006)
5 Years After x Treatment	-0.0003	0.0001	-0.0028	-0.0009	-0.0023	-0.0042	-0.0051
	(0.0001)	(0.0002)	(0.0003)	(0.0002)	(0.0002)	(0.0003)	(0.0007)
1 Year After	-0.0041	-0.0064	-0.0036	-0.0062	-0.0034	0.0029	0.0017
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0006)
2 Years After	-0.0092	-0.0093	-0.0093	-0.0094	-0.0114	-0.0050	-0.0075
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
3 Years After	-0.0126	-0.0125	-0.0123	-0.0116	-0.0156	-0.0104	-0.0142
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
4 Years After	-0.0130	-0.0135	-0.0117	-0.0112	-0.0163	-0.0113	-0.0149
	(0.0002)	(0.0003)	(0.0004)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
5 Years After	-0.0133	-0.0145	-0.0111	-0.0113	-0.0170	-0.0127	-0.0171
	(0.0002)	(0.0003)	(0.0004)	(0.0002)	(0.0003)	(0.0004)	(0.0008)
Log income	0.0424						
	(0.0002)						
Constant	0.0806	0.0698	0.0982	0.0706	0.0826	0.0891	0.1047
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0003)
Firm Fixed Effect?	Y	Y	Y	Y	Y	Y	Y
% of RI Sample	2.5	0.9	0.8	1.3	0.9	0.6	0.2
% of Sample Enrolled 2005-2008	35.7	12.7	11.3	18.7	12.6	8.7	2.9
R-squared	0.1509	0.1169	0.0846	0.1367	0.1201	0.1096	0.1423

Notes: This table presents regression coefficients of reported contribution rate on being treated with the Pension Protection Act (PPA) of 2006. "Year of" means the year the individual enrolled in their retirement plan and "x years after" is x years after they enrolled in the plan. Each column includes year dummies for each year after enrollment, and interactions of these dummies with the treatment dummy. The treatment dummy is equal to one if the individual enrolled in 2007 or 2008, after the PPA, and zero if they enrolled in 2005 or 2006. The full sample is those enrolled from 2005-2008 who otherwise meet the RI sample criteria. The bottom (top) income tercile includes those whose initial income is in the lowest (highest) tercile. Columns (4)-(7) break out the result for all individuals enrolled from 2005-2008 by age at enrollment. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. Log income, when included, is the log deviation of the individual's current income from the average income of the RI sample. Standard errors, in parentheses, are clustered at the household level.

Chapter 3

The Impact of the Paycheck Protection Program on (Really) Small Businesses

3.1 Introduction

The global economic crisis associated with the COVID-19 pandemic has put unprecedented strain on small businesses. Following the arrival of the virus in the United States and the declaration of a state of emergency on March, 13 2020, many small businesses experienced large declines in demand due to social distancing and the government mandated shut-downs that followed. In response, the federal government swiftly enacted the Coronavirus Aid, Relief and Economic Security (CARES) Act on March 27, 2020, a main component of which was the Paycheck Protection Program (PPP) designed to provide relief to small businesses. The PPP offered loans to small businesses, typically those with less than 500 employees, to cover up to eight weeks of payroll costs. The program was specifically designed to encourage small businesses to keep employees on their payroll in the face of the pandemic (SBA (2021a)). The loans are eligible to be forgiven if firms maintain most of their previous employment and wage payments. As of the end of the program on May 21, 2021, $793 billion in loans had been dispersed to nearly twelve million firms.[1] Did this spending achieve its goals of maintaining employment and wages of small businesses? Does it provide a reasonable model for future policy responses to the ongoing pandemic?

I use monthly administrative data from a private payroll processor in the American southwest to document several findings on both PPP take-up and efficacy for small and very small firms.[2] First, I show that take-up amongst small businesses is consistent with previous findings in the literature, which looked primarily at larger small businesses. Take-up is strongly associated with four factors: size, industry, average wages, and banking relationships. Next, I show that the PPP had a large and positive effect on employment amongst these small firms, and that the effect was much larger for firms that are in industries with a lower number of hourly workers, a high number of workers with remote working capabilities, and essential businesses. My estimates imply a much larger effect on employment for small businesses than found by most previous studies on the PPP. My results are identified by using administrative data on PPP application status to compare firms that did apply for PPP versus those that did not, but are otherwise similar on observables. The identifying assumption is that PPP take-up is primarily driven by as-good-as random factors related to banking frictions and the complex roll-out of the program (Granja et al. (Forthcoming), Li and Strahan (2021), Joaquim and Netto (2021))

[1] In comparison, the American Recovery and Reinvestment Act of 2009 and the Troubled Asset Relief Program (TARP) disbursed roughly $800 billion and $700 billion, respectively, across different areas of the economy.

[2] Previous work on the PPP has primarily compared outcomes for firms on either side of the 500 employee cutoff for program eligibility. In contrast, the median firm in my data set has five employees. Thus, rather than using variation in eligibility, I worked with the payroll processor to collect data on whether or not each firm applied for PPP.

Before estimating the effects of the PPP, I validate previous findings on which firms are most likely to apply for relief through the PPP. I find four characteristics to be the strongest predictors of take-up. First, larger firms are more likely to take-up the PPP loans. My estimates imply that a firm that is one log-point larger has 1.8 times larger odds of applying, holding the other factors constant. In the context of my sample, this means that, all things equal, a firm at the 75th percentile of employment (12 employees) is 13% more likely to apply than a firm at the 25th percentile (3 employees).[3] Second, industry is a significant factor in determining take-up. Notably, the most affected industries were not necessarily the most likely to apply. For example, in my sample, Trade, Information, Professional Services and Healthcare were amongst the most likely applicants.[4] Third, firms that have higher average wages per worker are also more likely to apply. A firm at the 75th percentile of the average wage distribution is about 5.5% more likely to apply than a firm at the 25th percentile. Finally, firms that use a bank that does more PPP lending in their state are more likely to apply.[5]

In the next part of the paper, I measure the effect of the PPP on firm outcomes using a dynamic difference-in-differences event study (similar to Autor et al. (2022), Chetty et al. (2020), and Hubbard and Strain (2020)). My identification strategy relies on making comparisons among firms that are similar on observable characteristics and assuming that the remaining variation in PPP application is driven by as-good-as-random differences that are uncorrelated with the firms' outcomes. The observable factors I control for include indicator variables for firm size, NAICS sector by time fixed effects, and city by

[3]Previous evidence (Barrios et al. (2020), Bartik et al. (2020a), Balyuk et al. (2020)) has shown that large firms found it easier to access the loans, consistent with my findings.

[4]This is consistent with Papanikolaou and Schmidt (2021) and Granja et al. (Forthcoming), which found the most affected industries and areas received a disproportionately small number of loans.

[5]Granja et al. (Forthcoming) also shows that the bank at which the loan was processed is important for both the take-up decision and how quickly firms received funds. Joaquim and Netto (2021) theoretically show that banks are incentivized to allocate loans to firms that are less impacted by the pandemic. See also Li and Strahan (2021), Balyuk et al. (2020), Erel and Liebersohn (Forthcoming), Ben-David et al. (2021)

time fixed effects. In other words, I compare firms within the same city, 2-digit NAICS industry, and employment size group, tracing the effect of the PPP over time at a monthly frequency.

In this event study, I find that PPP increased average employment by 7.5% at treated firms in the five months following application to the PPP.[6,7] The estimate is significantly different from zero at the 1% level. I also find an insignificant effect on wages. Moreover, I find that that firms that received PPP did not hire more workers in general or rehire more former employees than non-PPP firms. Moreover, they were no more likely to rehire or hire employees on the intensive margin. My estimates also indicate that PPP applicants were slightly less likely to reduce employment in the entire post-PPP period, though the estimate is not significantly different from zero. Together, these results suggest that the program largely worked by preventing layoffs, as was its primary intent.

To the best of my knowledge, my findings are the first to show that the PPP worked as intended, largely by preventing layoffs, rather than inducing more hiring or rehiring of former employees, To confirm this hypothesis, I run cross-sectional regressions on employee turnover on firms characteristics. I separate the analysis by total hires, new hires, or rehires or former employees. A rehire of a previous employee implies that the the worker is known to be a good match to the firm and was likely laid off due to financial reasons, not for performance issues. Although the intention of the PPP was to preserve employee-employer relationships, I find no evidence that the PPP caused firms to rehire former employees at higher rates. However, I do find suggestive evidence that PPP firms were less likely to reduce employment, implying that more matches were preserved at

[6]In the first month following PPP, my estimate is 13.7% which is quadruple the magnitude of that found in Autor et al. (2022) for the largest small businesses. The preferred estimate in that paper is a 3.25% increase in employment.

[7]On average, firms that applied for PPP and firms that did not both experienced employment declines over this period. The positive estimate implies that treated firms experienced a smaller decline.

PPP-applicants versus non-applicants.

Within these average effects there is significant heterogeneity by industry. The positive effect of the PPP on employment was larger in industries with fewer hourly workers, more employees who can work remotely, and essential businesses. These results suggest that the effectiveness of the PPP depends on interactions with local economic conditions and restrictions on business activities. Firms were only able to reap the benefits of the program when they were able to continue operations and had sufficient consumer demand. There were no significant differences in hiring or rehiring trends between the different types of firms I examined.

There are several sources of variation that may drive differences in take-up, conditional on observables, but may or may not be related to firm outcomes outside of the PPP take-up channel. The primary source of such variation that I use for identification are the banking frictions associated with access to PPP. These frictions have been well documented in other papers (Granja et al. (Forthcoming), Erel and Liebersohn (Forthcoming), Barrios et al. (2020), Li and Strahan (2021), Ben-David et al. (2021)). Firms that had pre-existing relationships with more PPP-savvy banks were more likely to have quick access to a loan.[8] Thus the main identifying assumption is that PPP-savvy banks do not influence firm outcomes through other channels. This is consistent with the findings in Granja et al. (Forthcoming) that banks that did more PPP lending were actually smaller banks that typically did less SBA lending in general, thus suggesting that they are not necessarily more adept at small business financing. Moreover, the most commonly used banks (for payroll processing) in my sample were those that did proportionally small amounts of PPP lending, according to Granja et al. (Forthcoming). This suggests that

[8]While I observe which bank a firm uses to process their payroll, I do not observe which bank they used to process their PPP loan. This fact, combined with the small sample size, means that an instrumental variables approach is not possible in this setting.

banking frictions were likely to apply in this sample of firms and are thus a plausible source of exogenous variation.[9]

Of course there are other possible drivers of take-up. The complexity associated with the rapid roll-out of the program has also been well documented as a factor in the take-up decision. Three pieces of evidence support this. First, the initial wave of applications for PPP was rationed because the PPP ran out of money and stopped accepting applications for several weeks (Kimball (2020), Doniger and Kay (2021)). Being rationed in the first round may have discouraged some firms from applying later. Second, the requirements about how and when a loan would be forgiven were not only difficult to interpret, but also changed several times while the program was running, which presumably led to temporal variation in the attractiveness of the program across firms over time (Hayashi (2020), Brewer et al. (2020)). A number of businesses even took and subsequently returned loans. Third, given the confusion and complexity associated with the program, it is feasible that some firms chose not to apply simply to avoid navigating the process. Bartik et al. (2020b) find that over 10% of firms did not plan on applying for PPP because "it's a hassle". Another 10% did not think that they would get the money in time. Nearly one-third did not believe they were eligible, despite the survey only being administered to firms with less than 500 employees, most of which were presumably eligible. Balyuk et al. (2020) find hesitancy to take up PPP amongst small firms, due to concerns over government scrutiny. Together, these pieces of evidence support the idea that, controlling for observables, the take-up decision was driven by as-good-as-random differences, in which case my results are unbiased.

Two additional robustness checks address concerns about variation driven by private

[9] In section 3.5, I'll discuss in more detail the validity of the as-good-as random assumption when it comes to cross-bank variation.

information or unobservable firm health prior to the decision to take on a PPP loan. I repeat all analyses using the balanced panel and again controlling for yearly growth in the first two months of 2020. The results are similar to the main sample. There is a positive effect on employment that declines over time and no statistically significant effect on wages for PPP-applicants. There is also no effect on hiring or rehiring in these subsamples. These results support the underlying identifying assumption that application was not selected on firm health.

3.1.1 Relation to Literature

This paper most directly relates to recent studies using differences-and-differences designs to evaluate the effectiveness of the PPP. There are two broad categories of paper that do so. First are those using firm-level data and second are those that use the SBA's loan-level PPP data combined with some measure of aggregate geographic outcomes. The first set of papers primarily uses the program's design for identification by comparing firms just above and below the 500-employee eligibility cutoff. Autor et al. (2022) finds a positive effect on both wages and employment, on the order of 2-4%. Chetty et al. (2020) finds a null effect on employment. Hubbard and Strain (2020) find a positive effect on employment of less than 1%. My work is complementary to these papers, but meaningfully expands upon their findings in three ways.

First, I study much smaller firms. Rather than focusing on firms around the 500-employee cutoff my data has a median size of five firms. As shown in Faulkender et al. (2020), 95% of PPP loans went to firms with 38 or fewer employees. Given the different operational structures and financing constraints (see Bartik et al. (2020a)) of firms of these different sizes, it is natural to think that their application decision as well as use of funds might differ. Indeed, given that I find much larger effects than the papers that use the cut-off strategy, financing constraints may be a significant factor in the effectiveness of the program at the firm level. Thus, it is essential to look at these very small firms when assessing the success of the program.

Second, this paper uses administrative data on PPP application status, rather than imputing application status based on size cutoffs. I know with certainty which firms in my sample applied for the program, thus I can measure treatment effects more precisely. In the papers that assume any firm with under 500 employees applied, there are inevitably some non-PPP firms in the treatment group. If the PPP truly had a positive impact on employment outcomes, the results using this that identification strategy could be downward biased. This is consistent with my finding of much larger employment effects in my sample.

Third, my data provides details not available in larger administrative payroll data sets. Namely, I have information on hiring and rehiring in the months following PPP. Thus, I observe not just the level of employment, but also whether the level changes were driven by hiring, re-hiring, or layoffs. As a result, I can speak directly to whether or not the program preserved jobs. Other papers that show only level changes cannot say if the PPP preserved pre-existing firm-employee matches.

Another paper that falls in this category is Bartik et al. (2020c). The data used in this paper is perhaps most similar to my data given the average number of employees: 7 as of January 2020. The main finding is that PPP increased firms' survival percentage significantly. However, this paper uses survey data and thus may be more subject to measurement error than my administrative data. It also lacks detailed information on wages and hiring. Another is Denes et al. (2021), which looks as firm financial fragility as an outcome, rather than employment outcomes.

The second set of papers, which do not have firm-level data, use identification strategies primarily related to the banking frictions and timing discontinuities associated with the PPP's roll out. Granja et al. (Forthcoming) uses geographical variation in banks, which had differing propensities to engage in PPP lending, and finds no significant effect of the PPP on employment or shutdowns at the local level.[10] Doniger and Kay (2021) uses the variation in timing due to funds running out for several weeks in April of 2020 to measure the effect of PPP, finding that it saved millions of jobs. Faulkender et al. (2020) also exploits variation in geographical banking structure and the timing of the loan roll-outs and find that a 10 percentage point increase in eligible payroll covered by PPP resulted in a 1 to 2 percentage point smaller jump in weekly initial unemployment insurance (UI) claims. Li and Strahan (2021) show that areas with better supply of PPP loans (due to presence of certain banks) have better employment outcomes. Bartik et al. (2020a) finds that states which received more PPP funding saw a faster employment recovery.

Compared to these papers, the focus of this paper is to assess detailed, firm-level outcomes, particularly for very small firms. Again, knowing application status provides the advantage of precise estimates that don't rely on being able to predict PPP-status. Firm-level data also allows me to look at the heterogeneous impact of the program across different types of firms. My data cannot speak as much to aggregate outcomes or random drivers of take-up. Both strategies provide unique information that is relevant for assessing the success of the policy.

My work is also related to other studies that assess the effectiveness of the PPP on

[10]Granja et al. (Forthcoming) has some firm-level data from Homebase. Homebase disproportionately covers small firms in food and beverage service and retail; therefore, it is not representative of aggregate employment.

various dimensions. Papanikolaou and Schmidt (2021) shows that PPP funds flowed to industries that potentially needed them less, as they were not as hard hit by the pandemic. Barrios et al. (2020) develops a framework for predicting demand for PPP loans and finds that, to date, disbursements have been quite similar to their model's predictions. Balyuk et al. (2020) finds that smaller firms were more hesitant to take up PPP and that banking relationships with small banks contributed to this. This paper substantiates these findings within a sample of very small firms.

This paper also contributes to the study of small business financing by providing another cost estimate for small business loan programs, particularly relief programs that occur during times of emergency. Brown and Earle (2017) studies the impact of SBA loans during normal times and finds that every million dollars in loans results in 3-3.5 new jobs created. Feyrer and Sacerdote (2011) examines the overall effectiveness of government stimulus during the 2008 financial crisis. The main finding is that the American Recovery and Reinvestment Act of 2009 created additional jobs at a cost of $170,000 of stimulus per job. In contrast, I find a much higher cost-per-job estimate for the smallest firms, around $270,000 per job per year.[11]

The paper proceeds as follows. Section 3.2 describes the details of the PPP. Section 3.3 describes the data. Section 3.4 substantiates the stylized facts about PPP take-up in my sample. Section 3.5 details the research design for evaluating the effect of the PPP on firms outcomes. Section 3.6 describes my results on firm outcomes, using the differences-in-differences design. Section 3.7 briefly analyzes the aggregate effect and cost of the PPP

[11]More generally, other papers have studied the importance of loan programs on business outcomes. Petersen and Rajan (1994) uses a survey administered by the SBA to study lending relationships. The main finding is that preexisting lending relationships can lead to easier access to funds; this is especially important in smaller firms where there are likely large information asymmetries between firms and lenders. Lelarge et al. (2010) finds that firms that received a loan guarantee from the French government had higher growth rates.

that is implied by my results. Section 3.8 concludes.

3.2 Policy Details

The Paycheck Protection Program (PPP) was a specific provision of the Coronavirus, Aid, Security, and Economic Relief (CARES) act intended to provide small business relief. Signed into law of March 27, 2020, the PPP program, implemented by the Small Business Administration (SBA) provides loans to small businesses, typically those with under 500 employees.[12]

Originally, $349 billion was allocated to the program, to be dispensed through the end of the year. However, the funding ran out less than two weeks after loans began to be dispensed. On April 24, 2020, an additional $320 billion was added to the PPP and the SBA began accepting applications again on April 27th, 2020. The first iteration of the program officially ended on August 8, 2020.[13] According to the SBA, the second round of loans dispensed smaller loans on average and the vetting process for approval was more stringent (Hare (2020)). For example, public companies were essentially excluded from the second round of loans.

The loans are low interest, with a rate of 1% and with a maturity of either two or five years.[14] Firms can apply to receive up to 10 weeks of payroll costs, based on their 2019 average payroll. Moreover, the loans are eligible to be forgiven in full if firms follow certain guidelines on how the funds are spent. The rules of the PPP have changed substantially

[12] There were some additional industry and tax based rules on eligibility. See SBA (2021b) for more details.

[13] The program was later extended, with extensive modifications, following the passage of the Consolidated Appropriations Act of 2021. That version of the program ran from January-May 2021.

[14] Originally, loans had a two year maturity. This was extended to five years for loans issued after June 5, 2020 with the Paycheck Protection Program Flexibility Act (see below).

since the program was first announced. In particular, the Paycheck Protection Program Flexibility Act (PPPFA), signed into law on June 5, 2020, changed several rules. Most of the changes related to forgiveness criteria applied retroactively to businesses that already had loans. I'll describe the rules in their form as of September 2020, which is the end of my sample period, but note that many businesses made the decision to apply with the original, more restrictive, rules in place. Appendix Table C.2 describes the changes that occurred after the PPPFA.

In order to have the loan forgiven, at least 60% must be used to cover payroll costs. Additionally, the firm must maintain 75% of their full time equivalent (FTE) employment in order for the loan to be forgiven. The safe harbor rule gives firms until December 31, 2020 to satisfy this requirement.[15] The amount forgiven is "all or nothing": if a firm uses 60% of the funds for payroll, it is all forgiven while if they use less than 60% for payroll, none of it is forgiven.[16] This must occur over a 24-week covered period, meaning that firms must spend 60% of the loan, which is typically equivalent to 10-weeks of payroll costs, over a 24-week period.[17] Lastly, firms cannot reduce wages by more than 25% in order to qualify for loan forgiveness.

If the loan is not forgiven in full, payments can be deferred for six months plus the amount of time it takes for the SBA to give the lender the appropriate forgiveness amount. No personal collateral was required. However, despite the policy design, intended to make the loans more like a grant, business owners remained skeptical. A survey by Bartik et al. (2020b) found that over 30% of small business respondents said that they did not

[15]If firms can document that they are unable to rehire individuals employed as of February 15, 2020 (or similarly qualified replacements) by December 31, 2020, this will not count against their forgiveness amount. This can either be due to individual choice on the part of the employees or because their business activity is limited by local governments.

[16]Prior to the PPPFA, the amount to be forgiven would be proportional to the amount that was spent on payroll. See Appendix Table C.2.

[17]This is a substantial relaxation of the original rules, in which the covered period was only eight weeks.

trust the government and/or the bank to forgive the loan. Moreover, business owners may have faced difficulties in spending even 60% of the funds on payroll and maintain 75% of their pre-pandemic FTE when many states had required businesses to remain closed and were offering generous unemployment insurance (UI) benefits during this time.[18,19]

The loans are financially attractive on several dimensions. Ignoring the forgiveness provisions, they are extremely low interest relative to typical SBA loans, which have an interest rate of 7-8%. Moreover, the deferral period, zero-collateral requirement, and the long maturity, especially for post-PPPFA loans, made the loans very low cost. The provisions enacted with the PPPFA also made forgiveness much more attainable. Since the covered period is 24-weeks, firms have nearly 6 months to spend 60% of 10 weeks of payroll on wages. This means that firms only need to maintain a quarter of their pre-pandemic level of employment to qualify for forgiveness. The rules around wage reduction and FTE maintenance are somewhat alleviated by allowing provisions for businesses that face government shut-downs or employees that refused good faith offers of employment.

Consistent with these favorable terms, take-up of PPP loans was quite high. As of August 8, 2020, when the second round officially ended, $525 billion in loans had been disbursed to over five millions firms. Barrios et al. (2020) estimate that demand would be $750 billion if all eligible firms took up the loan. Appendix Figure C.1 shows the distribution by firm size of firms that received PPP versus all firms in the U.S. The calculated take-up rate is 75%. The SBA estimated that PPP loans had covered 72-96% of

[18] Another source of confusion has been what will occur if a company goes bankrupt after receiving the loans, or is currently in bankruptcy proceedings (Moore (2020), Iacurci (2020))

[19] It is worth noting that Bartik et al. (2020a) and Altonji et al. (2020) find no effect of increased UI benefits on aggregate employment.

all payroll across states as of June 2020 (SBA (2020)). In this paper, I focus primarily on the first and second rounds of PPP, which stopped accepting applications on August 8, 2020. Following the passage of the Consolidated Appropriations Act of 2021, the PPP restarted in January 2021 with significant modifications, and continued through May 31, 2021.

3.3 Administrative Payroll Data

My data is from a private payroll processor that is headquartered in the American southwest. The company services nearly 400 firms, representing approximately 6,000 employees prior to the pandemic's start. The firm has clients located in nine U.S. states, but the majority of the clients (and employees) are located in one U.S. state in the Southwest. The company has provided me access to firm level data at the monthly frequency from January 2019-September 2020. The data contains information about industry, number of employees, hiring, total wage bill and whether or not the firm applied for a PPP loan during the pandemic.[20,21]

The firms represented in the data set are quite representative of all U.S. small businesses, both in terms of industry composition and size. The top panel of Figure 3-1 shows the breakdown of firms by industries, measured at the 2-digit NAICS code level. The bottom panel shows the distribution of employment at these firms prior to the pandemic (February 2020). Despite the small sample size, the firm has broad coverage of many different industries. Moreover, it matches the distribution of firms in the U.S. and the states in which it has clients reasonably well. It is over-represented in the healthcare

[20] I assume that if a firm applied, they received a loan. This is likely to be the case for most firms, as the acceptance rate was high (Horan (2020)).

[21] Note that I observe if a firm applied, regardless of which bank they applied through. Thus, I know application status even if they applied via a non-traditional lender, like FinTech. This may be particularly important for these very small firms (Erel and Liebersohn (Forthcoming)).

sector and somewhat underrepresented in retail and wholesale trade, relative to the U.S. total.

Table I shows summary statistics for the sample versus all businesses in the U.S. with under 500 employees. The average firm had a mean of 14 employees from February 2019-February 2020, versus 9 in the 2016 SUSB data. Average monthly wages are slightly higher than in the US, but the mean wage is slightly lower, perhaps due to the differing sector compositions.[22] In sum, the sample is quite comparable to the average small firm in the U.S. and has a good representation across sectors.

The sample provides a snapshot of how effective the PPP was for some small and very small businesses, particularly those in states that did not have a severe case load during the first wave (March and April 2020) of COVID-19 the United States. During the first wave, which coincided with the start of the PPP, there was a much lower case load in the states where my firms operate compared to the U.S.[23] Around mid-June, the states represented in my data see an increase in case-load. Regarding the timeline of events affecting business activity, there was a shutdown that took place from early April to mid-May. However, the requirements of the shut-down were relatively relaxed compared to states that had more severe outbreaks. Moreover, many businesses were allowed to start reopening before the shutdown technically ended.

The sample is clearly selected in several ways. First, the firms are geographically concentrated in the American southwest, where the first wave of the virus was less severe. Second, these firms are primarily small enterprises, with the median firm having only

[22]The sample over-represents the health service sectors, in which the average wage is below the national average (BLS (2020b)).

[23]Appendix Figure C.2 shows the number of cases per person in the states in which the firm has clients versus the U.S. total.

5 employees (prior to the pandemic). The largest client has just over 300 employees.[24] Thus, the sample does not cover large firms, but rather captures a sample of small to medium sized enterprises. Third, the firm's clients are concentrated in specific industries (see Figure 3-1). In particular, the healthcare industry is over-represented relative to the U.S. population. Lastly, the firms are those which would hire a payroll processor, rather than do their payroll in house. Despite these caveats, the summary statistics compared to the U.S. total are reassuring in the sense that these firms are, at least based on size, fairly representative of the average small business in the U.S.

The data have been provided to me on a monthly basis. I observe time varying data on wage bills, number of employees, and the composition of new hires (i.e. whether a new hire is a rehire of a previous employee, or a completely new person).[25] I also observe several fixed characteristics of the firm: zip code, 6-digit NAICS code, and the bank where the firm processes payroll. In most of the analysis, I keep firms that were present in the data beginning in February 2020 and remained until September 2020.[26,27]

3.4 PPP Loan Take-up

In this section, I show that take-up in my sample is consistent with the stylized facts on PPP take-up that have been documented elsewhere in the literature. The advantage

[24] The firm has one client with just over 500 employees, but it is dropped from all analyses as it is ineligible for the PPP.

[25] Wage bills are normalized to the monthly level based on the firm's pay frequency and how many paydays they have in a given month. Most firms in the data pay at a biweekly frequency, so the wage bills must be adjusted to account for the fact that some months have five Fridays while some have four.

[26] Unreported robustness checks impute zero for firms that drop out, rather than dropping them from the analysis. This occurs only for a small number of firms (< 10). Results are similar.

[27] Note that the Treasury has released detailed information on all firms, including name and the dollar value of the loan received. However, due to a privacy agreement with my data provider, I am not able to match my data to the Treasury data.

of my data is that I have more details than are provided in the Treasury data release and most other administrative datasets. In particular, I know exact employment and wage bill and the firms' histories over the year and half prior to the pandemic. Within my sample, I can explore the factors which contribute to a firm not applying for the PPP. As discussed in Section 3.2, these loans have extremely favorable terms, have the chance of turning into a grant, and are being offered in the middle of a large economic crisis that has impacted firms across the board. Hence, not applying is a puzzle.[28]

To estimate the effects of variables found to be associated with take-up in previous analyses, I run a logistic regression of the form[29]:

$$PPP_i = \alpha + \gamma_j + \gamma_s + \gamma_b + X_i + \epsilon_i \tag{3.1}$$

where γ_j is an industry fixed effect, γ_s is a state fixed effect, γ_b is a bank-PPP lending tercile fixed effect and X_i is a vector of variables of interest: baseline size, average salary, and a measure of the size of the firm's shock in April of 2020, either at the individual firm level or the industry-wide level.

The results are shown in Table II. First, I'll discuss the results on size. Size has been well documented as a factor in take-up in other studies, with larger firms being more likely to apply for and receive PPP loans (Barrios et al. (2020), Bartik et al. (2020a), Balyuk et al. (2020), Erel and Liebersohn (Forthcoming)). This is true in my sample as well. Even with the full set of controls in column (5), the odds ratio is nearly 1.8, implying that a firm that is one log-point larger, in terms of employment, has 1.8 times the odds of applying.

[28]An important agenda for future research is to analyze why 25% of firms in total did not apply. Given the sample selection, my results here may or may not generalize to the whole population.

[29]The variables of interest are also motivated by summary statistics in my own sample. Appendix Figure C.3 shows several factors that seem to be important for take-up based on the previous research, split by PPP-applying and non-applying firms. The figures confirm that larger firms with more highly paid employees, and firms that are in certain industries are more likely to apply.

In the context of this sample, the interpretation is that a firm at the 75th percentile of the employment distribution (12 employees) is 13% more likely to apply than a firm at the 25th percentile (3 employees), all other factors constant.[30]

Second, industry is a significant factor in take-up. However, this relationship is not a one-to-one positive correlation with pandemic exposure. That is, not all firms that were most exposed to the effects of the pandemic were most likely to apply. Figure 3-2 shows the odds ratios on the industry fixed effects from the regression in the fifth column of Table II. The reference category is the Food and Accommodation Services sector. Almost all industries were less likely to apply than those in Food and Accommodation Services, with the exception of Trade, Information, Professional Services, Educational Services, and Healthcare and Social Assistance. Somewhat surprisingly, these are actually not the industries that had the largest employment declines during the early pandemic, as measured by Cajner et al. (2020). They are however, some of the most exposed to the pandemic, as measured both by the percentage of workers that are able to work from home, as in Papanikolaou and Schmidt (2021) and the amount of face-to-face interaction required with customers, as in Leibovici et al. (2020).

The coefficient on average wage is also significant and greater than one across each specification. This implies that firms within the same industry, the same state, that use a similar bank and have similar employment levels and similar employment growth in April of 2020 are more likely to apply if they have a higher average wage per worker. A firm at the 75th percentile of the average wage distribution is about 5.5% more likely to apply than a firm at the 25th percentile. This, combined with the industry results, is consistent with the evidence documented in Granja et al. (Forthcoming), Papanikolaou

[30] I also find significant differences between those firms that applied early, during the first round of PPP that took place from April 3, 2020-April, 17, 2020, versus those that applied during the second round, in May or June. Appendix Tables C.3-C.5 show the results.

and Schmidt (2021) and Joaquim and Netto (2021) that funds flowed disproportionately toward firms that were less in need.

Third, column (5) shows that firms that use a primary bank that did more PPP lending in their state overall are more likely to apply. This is consistent with the findings in Granja et al. (Forthcoming) and Faulkender et al. (2020). Specifically, a firm with a bank in the top tercile of PPP lending in the state has 1.5 times the odds of applying compared to a firm in the bottom tercile.[31]

Finally, there is no significant effect on the monthly employment growth variables in any specification, indicating that there is not adverse selection into PPP application after controlling for industry, size, and bank.[32,33]

The results are robust to using an industry fixed effect (columns (1), (2), and (5)), or a control for the industry's employment decline in the early pandemic (from February 15, 2020-April 11, 2020, as measured in Cajner et al. (2020) (columns (3) and (4)). The odds ratio of less than one on this variable implies that firms in industries with higher employment growth during the early pandemic were less likely to apply, as expected.

Consistent with studies to date that have analyzed PPP loan take-up, I find significant effects of size, industry, bank, and average wage per worker. The results confirm that firms in industries more exposed to the pandemic, but not necessarily firms that were individually more adversely affected, were more likely (and quicker) to apply. Moreover,

---

[31] See also Li and Strahan (2021), Balyuk et al. (2020), Erel and Liebersohn (Forthcoming).

[32] Hence, the difference observed in Appendix Figure C.3b is likely a function of size and/or industry composition amongst appliers versus non-appliers. In other words, industries that had larger drops in employment in April were also more likely to apply.

[33] The results also hold if I control for firms' long-term (yearly) growth prior to the pandemic (Appendix Table C.6) The results are also qualitatively similar looking only at the balanced panel of firms that have been in the data set since January of 2019 (Appendix Table C.7). These robustness checks serve as additional evidence that firms likely didn't change their application decision based on private information.

larger firms and firms with higher wages per worker were more likely to apply. Lastly, firms with banking relationships with banks that did more PPP lending overall were more likely to apply.

3.5 Identification of PPP Treatment Effects

My methodology to assess the effectiveness of the PPP is a dynamic difference-in-differences event study, similar to Autor et al. (2022). In order to assess the PPP's effectiveness, we need a valid counterfactual. In other words, what would have happened to firms if they hadn't applied? While no such counterfactual exists, I can estimate it by looking at firms that are otherwise similar on observables but did not apply for the PPP. Because I observe whether or not a firm applied for the PPP, I can directly compare the outcomes of firms that applied to those that did not. The specification is:

$$y_{it} = \alpha + \theta PPP_i + \sum_{t \in T} \beta_t(\theta_t \times PPP_i) + \gamma_e + \gamma_{jt} + \gamma_{st} + \epsilon_{it} \tag{3.2}$$

where PPP_i is an indicator equal to one if the firm applied for the PPP, θ_t is a time dummy that corresponds to months relative to when the firm applied for PPP[34], γ_e is a dummy for the firm's employment size[35],γ_{jt} is a 2-digit NAICS code by month fixed effect, and γ_{st} is a month by city fixed effect. y_{it} is an outcome variable, employment, total wage bill, or a measure of hiring divided by the value of the variable in February 2020. The large set of controls ensures that I compare outcomes of treated firms to similar untreated firm,

[34]Firms in my sample applied in either April, May or June. θ_t indexes months since application, with $t = 0$ corresponding to the month of application, thus this varies across firms. I set $t = 0$ to April for all control firms. The industry by time and city by time fixed effects refer to calendar time, thus these controls account for differences in economic conditions from month-to-month.

[35]I group firms into buckets with between one and five employees, between five and 25 employees and greater than 25 employees. This roughly corresponds to terciles.

hence I can estimate the average effect of treatment on the treated (ATT).

The assumption underlying the estimation is that, in the absence of the PPP, control and treated firms would have had similar outcomes. In other words, differences in take-up are driven by as-good-as random variation, primarily related to the complexity, speed, and subsequent banking frictions with which the program was introduced. The regression itself takes several steps to ensure that this is the case. First, the employment size fixed effects ensure that I compare firms of similar size. Second, the time by industry fixed effects control for time varying shocks within a given industry that are common to all firms in that industry. Given the varied exposure to the pandemic across industries (Papanikolaou and Schmidt (2021)) this is quite important. Lastly, the city by time fixed effects control for time varying shocks within a city that are common to all firms in that city. This is especially important, given that the pandemic affected all areas differently and local governments had heterogeneous responses to contain/prevent outbreaks.

The banking frictions, from which my results are partially identified, warrant additional discussion. Granja et al. (Forthcoming) and others have documented that there was significant heterogeneity in the administration of PPP loans across banks. Consequently, firms in locations with access to those banks were more likely to receive loans. Some firms that wanted PPP may not have applied simply because they didn't have a banking relationship with the right bank. However, one may worry that a banking relationship is correlated with firm outcomes, independent of its impact on PPP status (Petersen and Rajan (1994) and Brown and Earle (2017)). Granja et al. (Forthcoming) find that the banks that did the most PPP lending were actually banks that typically did less SBA lending overall. This suggests that PPP-heavy banks are not necessarily more adept at small business financing in general. Indeed, the three largest banks by deposit share in my sample are JP Morgan Chase, Bank of America, and Wells Fargo, which did a disproportionately small amount of PPP lending (relative to their usual small business lending). This alleviates some concern about how the bank may influence the firm's outcome, outside of PPP lending.

Which additional sources of variation might drive differences in take-up, conditional on observables, but are not related to firm outcomes (other than through PPP take-up)? There are two such sources that may call into question the main assumption underlying my estimation. In what follows, I'll discuss each and argue that they are not significant sources of bias in my sample. First, is a CEO or manager's beliefs about the severity and duration of the pandemic. Bartik et al. (2020b) found in an April 2020 survey that about half of firms believed that the crisis would be over by July of 2020; given this belief, some firms may not think it worth it to apply. If this is the case, these businesses would likely appear better off in the short-run, but then would be more negatively impacted as the pandemic continued. This is the opposite of what I observe: I see untreated firms becoming more similar to treated firms over time. This indicates that this is not a primary driver of take-up in my sample.[36]

Second, the manager may have some private information about the firm's prospects.[37] That is, a CEO/manager may have known that his firm was likely to fail with or without the PPP. This is obviously correlated with the ex-post outcomes. While I can't observe

[36] A likely alternative story is that if a manager believes that the pandemic is not serious and hence wouldn't expect it to impact their business much, this makes the PPP essentially "free money". Especially given that they have a payroll processor dedicated to helping them apply, these firms may even be more likely to apply given their optimism (see Landier and Thesmar (2009)). Hence, if many firms that took the loans did not direly need the assistance, my estimates might overstate the effect of the PPP. This is consistent with Joaquim and Netto (2021), which shows that firm-level regressions are likely to overestimate the effect of the PPP, due the the banking system's distorted incentives.

[37] Another possible potential source of variation is the firm's financial sophistication. Firms with a dedicated CFO or accountant, or better financial records may be more likely to apply. Then, one might believe that such firms may be more or less successful in navigating the pandemic. I rule this out as primary explanation for the take-up decision, because I am dealing with a sample of firms that already use a payroll processor. By definition, these firms have someone dedicated to maintaining their wage and employment records.

this directly, I can control for it with several measures available in my data. First of all, I can observe how long the firm has been with the payroll processor. Given that younger firms are more likely to fail, limiting the analysis to firms who have been with the payroll processor since the beginning of my data set (January 2019) somewhat alleviates this concern. In robustness checks, I repeat all analyses using the balanced panel and again controlling for yearly growth in the first two months of 2020. The results show a positive, significant, and decreasing of the PPP on employment in the five months following PPP application, similar in magnitude to the effect observed in the full sample. Hence excluding firms that are statistically most likely to fail, based on age, maintains the results. However, the general direction of the bias due to this source of variation is ambiguous. It could be that firms with this private information do not apply, thus overestimating the results, or it could be that they do apply, thus attenuating the results. The fact that I don't see large differences in employment or wage levels between the control and treatment group five months after application suggests that this is not a significant source of bias in my sample. Moreover, as of September 2020, very few firms had left the sample, suggesting there had been minimal business failures.[38]

A second version of the analysis collapses all observations by either pre or post-PPP:

$$y_{it} = \alpha + \theta PPP_i + \beta_t(\theta_{t=POST} \times PPP_i) + \gamma_e + \gamma_{jt\in\{Pre,Post\}} + \gamma_{st\in\{Pre,Post\}} + \gamma_t + \epsilon_{it} \tag{3.3}$$

This measures the average effect of the PPP in all months following application for treated firms, relative to controls after April. The same control variables are used in order to compare firms in the same industry and in the same city. An additional control for

[38]Note that business failures are difficult to measure. First, some firms did undertake furloughs at various points in the sample, only to reenter at more normal wage and employment levels later. Thus a month of zero wages does not necessarily mean failure. Second, just because a firms leaves my sample, doesn't mean they have failed; they may simply have ended their relationship with the payroll processor.

calendar time fixed effects, γ_t is added in order to account for the differences in economic conditions each month.

The coefficient of interest in both cases is the β_t vector. This traces the effect of the PPP over time (or pre versus post-PPP in the case of equation (3.3)). I begin the regressions in February 2020, using all firms that were present as of January 2020.[39] If a firm drops out at any point after February 2020, I drop them from the sample entirely.[40] Examining the coefficient prior to the start of the pandemic serves as an additional check on the comparability of firms. If PPP application is not correlated with unobservables, there should be no trend prior to the start of the pandemic.

3.6 Results

Using the difference-in-differences analysis, I first show that the PPP had a positive but transient effect on employment for treated firms in the full sample, on the order of nearly 14% in the first month following application to the program. The effect decreases over time and is statistically equivalent to zero five months after application. Then, I show that there was no significant effect on total wage bill. Next, I show that the positive employment effects occur primarily in industries that are less exposed to the pandemic: industries with fewer hourly workers, more remote workers, and essential businesses. Hence, the program was only effective for firms least affected by government shut-downs and social distancing. Lastly, I show that there is no evidence that firms receiving PPP

[39]I exclude January due to seasonality: firms often make adjustments to employment and compensation during this month. I exclude the data from 2019 to preserve sample size.

[40]Unreported robustness checks impute zero for firms that drop out, rather than dropping them from the analysis, as in Autor et al. (2022). This paper says that to do so is “conservative", because firms that drop out of their sample may not have failed, but just left the payroll processor. However, whether or not this is “conservative" depends on if the firm is in the treatment or control group, thus I elect to just drop firms that leave the sample entirely. This occurs only for a small number of firms (< 10). Results are similar.

funds were more likely to rehire former employees or hire more employees in general; the increase in employment at PPP firms is driven primarily by fewer layoffs.

3.6.1 Employment and Wage Effects

The PPP had an initial positive effect on employment at treated firms that is significantly larger than most other studies have found. The effect declines to zero over time, but is positive and significantly different from zero in the entire post-PPP period that I study. The first panel of Figure 3-3, shows the results from equation (3.2) with employment as the dependent variable.[41] The point estimate for one and two months prior to the onset of the pandemic and the month of the onset are all statistically zero. This indicates that firms are ex-ante similar, holding the covariates constant, supporting a causal interpretation of the PPP on employment. One month following PPP application, average employment increased by 13.7% at treated firms. The effect is significant at the 1% level. Two to four months after the PPP, the estimate is again positive, significant and around 12-14%. The effect declines to zero five months after application. The estimate for one month after substantially larger in magnitude than that of Autor et al. (2022), which estimates an effect between 2-4.5% as of the end of May on much larger firms (around 500 employees). Figure 3-3c shows the results from equation (3.3), which aggregates each month after application to one “post-PPP" period. The aggregate effect on employment is 7.5% and is statistically different from zero at the 1% level.

I also find that the PPP had no corresponding impact on the firm’s total wage bill

[41] Appendix Figure C.6 shows bootstrapped results of my main specification. Due to both my small sample and the lack of research on the bias in dynamic difference in difference models, I bootstrap the estimator for robustness. The procedure is to reconstruct the dependent variable by randomly sampling from residuals of the main estimation and re-estimate the main specification. I repeat this 1,000 times and plot the 10th percentile, median, and 90th percentile of the resulting distribution of estimated treatment effects. The bootstrapped results show a similar margin of error to the main specification.

(Figure 3-3b).[42] Again, the point estimate one and two months prior to PPP application is close to zero, implying that firms are similar on the observables, ex-ante. Total wage bill decreased by 12.8% the month of application, though this estimate is not statistically different from zero. In the following months, total wage bill increased relative to controls in each month after application. However, none of these estimates are significantly different from zero. The aggregate result in 3-3c suggests a near zero effect on wages, on average.

These results are not driven by positive selection into PPP application. I have already presented several pieces of evidence to show that firms that were not as affected by the pandemic (within an industry and city) were not more likely to apply. First, the results on take-up suggested that firms that applied were in fact in industries more exposed to the pandemic. Thus, it seems unlikely that firms more affected on some unobservable dimension would be more likely to apply. Second, the point estimates of the difference-and-difference specification in the month of application are statistically zero for both employment and wage bills.

I address remaining concerns about selection in three ways.[43] First, I repeat the analysis only on the balanced panel of firms that have been in the data set from January 2019-June 2020. This provides some reassurance that more established firms are not being compared to fledgling firms in the analysis. The results, shown in Appendix Figure C.7 suggests a similar effect in this sample of firms. Second, I repeat the analysis adding a control for average yearly growth in January and February of 2020.[44] The results, shown in Appendix Figure C.8, also show a similar effect for both employment and wages. These

[42] In the following subsection, I will show that this ambiguous result is partially driven by heterogeneity across types of firms.

[43] As discussed in Section 3.5, firms may have private information about their firm's prospects and apply for the PPP even if that private information is negative. This would lead to attenuation of my results, as such firms would pollute the treatment groups with poorer outcomes. Moreover, if only firms that were very negatively affected by the pandemic early on apply, the results are also attenuated.

[44] Note that this includes only firms in the balanced panel starting in February 2020, by construction.

results suggest that the PPP is similarly effective for firms that are more established, suggesting that adverse selection is not present.

Third, I present suggestive evidence that the results are not driven by unobservable differences between firms. Appendix Figures C.9-C.10 show the results for each outcome variable, respectively, but with the controls added progressively. Appendix Figure C.9 shows that, even with no controls, the difference between treated and controls is nearly zero before the pandemic and the PPP. Even adding the size control barely changes the results, which is intuitive given that the dependent variables are measured relative to February for all firms. However, the effect on employment is zero before adding industry controls (Appendix Figure C.9c). It is not until we compare firms within the same industry that we see the PPP have a positive effect on employment ex-post. Then adding city by time fixed effects (Appendix Figure C.9d) results in firms that are more comparable ex-ante and a slightly larger effect. The fact that all of the point estimates are close to zero prior to PPP application, is reassuring that there are likely not unobservables that would significantly alter the firms' comparability. The results on wage bill in Appendix Figure C.10 shows a similar pattern.

3.6.2 Heterogeneity by Firm Type

In this subsection, I show that the PPP is much more effective for firms in industries with a smaller number of hourly workers, more workers that are able to work from home (remotely), and for essential businesses. The results suggest that the PPP in isolation is insufficient to buffer firms from the effects of the pandemic. However, if firms are able to remain open, or have employees who can work from home, then the PPP has positive effects on both employment and wages.

First, I establish that firms in industries with a low amount of hourly workers are positively impacted by the PPP, while firms in industries with a high number of hourly workers are not. I do this by first classifying industries based on the prevalence of hourly workers. Using data from the BLS, I calculate the percentage of workers that are hourly in each two-digit NAICS code (BLS (2020a), BLS (2016)). I then repeat the analysis on the firms that are in industries in the top tercile of the distribution (have the most hourly workers) versus those that are in the bottom tercile, (have the fewest hourly workers).[45]

As shown in Figure 3-4a, the employment gains observed in the full sample take place entirely in industries with a low percentage of hourly workers. Average employment in these industries increased by 8.5% in the five months following application to the PPP. As in the full sample the effect is initially higher and decreases over time (see Appendix Figure C.11). On the other hand, firms in industries with many hourly workers experienced no effect on employment.

There is no effect of PPP on wages for firms with many or few hourly workers (Figures 3-4a, C.11c, C.11d). The positive effect on employment and zero effect on wages for firms with few hourly workers suggests that these firms may have maintained employment but reduced wages.

In sum, industries with many hourly workers benefited less from the PPP than firms with more salaried workers. It is also true that industries in which most workers are not hourly are those that are less affected by stay-at-home orders. For example, many professional services firms, such as finance, insurance and tech, continued operating during the pandemic with their employees working remotely. On the other hand, restaurants, hotels, and retail stores, where many workers are hourly, were not able to do so (Papanikolaou

[45] Appendix Table C.8 shows the percentages for each industry.

and Schmidt (2021)). Hence, even firms in the latter industries that received PPP may not have been able to put the money to full use with shut down orders in place.

To test this hypothesis, I break out firms based on the percentage of workers in an industry that are typically able to work from home, as measured in Papanikolaou and Schmidt (2021).[46] I repeat the difference-in-difference analysis for firms in the top versus bottom half of the distribution for workers that can work from home.

The positive effects of the PPP on employment are much larger in industries in which many workers can work from home (Figure 3-4b). Average employment at the most remote capable firms increased by 19% in the five months following PPP. Average employment at the least remote capable firms, also increase, but by less than half as much, 9.9%. Both effects are statistically significantly different from zero. Appendix Figures C.12a and C.12b show that the employment effects declined to zero over time for the highly remote firm, but remained positive throughout the sample period for the low-remote firms.

The findings on wages in this subsample also reveal significant heterogeneity, which partially drives the ambiguous result in the full sample. Figure 3-4b shows that the PPP had a positive effect on wages for the highly remote workers of 19.1%, but no effect on wages for the low-remote workers. C.12c and C.12d show the results on a monthly basis. As with employment, the effect of PPP is initially positive but declines over time for the highly-remote firms.

The fact that PPP raised employment primarily in firms where workers are not paid hourly and in industries where they can work from home suggests that PPP's effectiveness was limited by stay at home orders and decreased demand/foot traffic due to social

[46]Note that the remote measure and the hourly measure are highly correlated. See Appendix Table C.9.

distancing. To test this hypothesis, I split out the results by essential versus non-essential businesses.[47]

Indeed, the positive effects of the PPP on employment and wages occurred mostly within essential businesses (Figure 3-4c). Essential businesses increased employment by 11.6% in the five months following PPP. Non-essential businesses increased average employment by only 4.6% over the same period. The month by month results (Appendix Figures C.13b and C.13a) show that the effect declined over time for both types of firm. The effects on wage bills and indicate that PPP had no effect for both essential and non-essential businesses

PPP-treated firms that were not allowed to remain physically open (non-essential businesses) did not benefit from the program, while essential businesses did. At the same time, only firms with with many remote-capable workers and fewer hourly workers benefited from the program. This suggests that the PPP was only effective for these small businesses if they were able to remain operational in some capacity, either due to being considered essential or from being able to conduct business remotely.[48]

3.6.3 Employee Turnover

In this subsection, I show new evidence that differences in employment at PPP versus non-PPP firms were driven by fewer layoffs, not increased hiring or rehiring of former employees. This holds in the full sample as well as the subsamples comparing more

[47] I categorize firms using the list provided in Papanikolaou and Schmidt (2021). This list is based on a conservative definition intended be applied for the whole U.S. The definition of essential in the state where most firms in my sample are located is certainly more generous. Hence, I am likely classifying some businesses that were considered "essential" as non-essential.

[48] I also test for differences between tradable and non-tradable industries, but do not find significant differences. Results available upon request.

hourly versus less hourly, remote versus non-remote, and essential versus non-essential industries. To show this, I utilize the data on employee turnover. For every new employee that a firm hires, I am given data on whether the employee is a new hire or a rehire of a previous employee.[49]

I first repeat the difference-in-differences analysis to show that hiring behavior was not significantly different between PPP and non-PPP firms in the months following application. The results using the number of total hires, new hires, and rehires as a percentage of employment in February 2020 as dependent variables are shown in Figure 3-5. For total hires (Figure C.23a), there is a positive effect of the PPP in the month of application, though it is not significantly different from zero. The effect decreases one and two months following application and then increases again, but its never significantly different form zero. The pattern is similar for number of new hires (Figure C.23b). Figure C.23c shows that PPP-treated firms did not re-hire more former employees in the months following treatment either.

I supplement the analysis by running cross-sectional regressions, looking specifically at the month of, one, and two months following application as well as the entire-post PPP period to compare treated versus control firms. Specifically, I regress the number of total, new, and rehires as a percentage of February 2020 employment on an indicator variable for PPP application, a control for lagged monthly employment growth (in order to compare firms that would need to hire a similar number of employees to reach their previous level of employment), two-digit industry fixed effects, size fixed effects, and state fixed effects. In the monthly regressions, I also include a month fixed effect to account for the fact that different firms applied in different months.

[49] Appendix Figure C.14 plots the number of total hires, new hires, and rehires as a percentage of the firm's employment in February 2020, split by PPP and non-PPP firms.

I find that there is no effect of the PPP on total or new hires in any month following application, or in the entire post period. These results are shown in the top and middle panels of Table III. In fact, the point estimates on total hires and new hires are almost all negative. Moreover, the only significant effects on rehiring are a negative effect two months after application. These point estimates suggest that PPP firms re-hired .59% fewer employees (as a percentage of their February employment) in the entire post-PPP period and did no more hiring overall.

I also find that there is no effect of PPP on the intensive margin of rehiring, hiring, or reducing employment. That is, PPP firms are not more likely to rehire former employees, hire anyone, or reduce employment in general. To show this, I run a logistic regression of the extensive margin of rehiring, hiring and reducing employment on the same set of controls. I set an indicator equal to one if the firm did the respective adjustment in a given month and zero otherwise. The results are shown in Table IV, with the coefficients reported as odds ratios. Beginning with rehiring in the top panel, the month of PPP application, treated firms are not significantly more likely to rehire former employees, though the odds ratio is greater than one, implying a positive, but statistically insignificant effect of the PPP. One month after PPP, there is also no significant effect and the odds ratio is nearly one. Two months after PPP application, treated firms are less likely to rehire former employees; the odds ratio is meaningful in magnitude and highly significant.[50] In the entire post-PPP period, treated firms are less likely to rehire former employees, but not significantly so.

Turning to hiring, reported in in the second panel of Table IV, PPP firms are not significantly more likely to hire in the months following PPP. Though the odds ratios are

[50]This may be because they have already done so, given the odds-ratio greater than one in the month of application.

larger than one, they are not significantly different from zero.

Lastly, PPP firms were not more (or less) likely to reduce employment in the months following PPP. The odds ratio for the entire post-PPP period is less than one, indicating that they were slightly less likely, on average, to reduce employment than non-PPP firms, However the estimate is not significantly different from zero.

The results imply that the positive effect on employment observed in the difference-in-differences results (Figure 3-3) is largely driven by PPP firms laying off fewer workers than controls, rather than hiring more employees. This indicates that the program worked as intended, by preserving employer-employee matches. Moreover, I find no evidence that PPP firms were more likely to bring back previously laid off employees than the controls.

Splitting up the firms by business-type, as in the previous section, I find no significant effects on hiring or rehiring for any type of business on a monthly basis (Appendix Figures C.20-C.22). The average estimates for the entire post-PPP period suggest that some non-essential businesses and low-remote businesses hired more workers in the months following PPP (Appendix Figure C.23). This is consistent with the narrative that these firms did not have positive employment effects from the PPP. Thus they likely laid off more workers early on and needed to do more hiring than the businesses for which the PPP helped to prevent layoffs. Overall, these results confirm that fewer layoffs drive the positive employment results across the board.

3.7 Aggregate Effects

How effective was the PPP, on a per dollar basis? In this section, I calculate the aggregate effect of the PPP through September of 2020 based on my estimates in the previous section. I estimate that the program maintained employment at an expense of between $88,000-157,000 per job as of September 2020. The effect is estimated using:

$$\text{Total payroll effect}_t = \beta_t \times \gamma \times N \tag{3.4}$$

where β_t is my average effect of treatment on the treated (ATT) effect from equation (3.3), γ is the percentage of the eligible population that applied, and N is the number of employees at eligible firms. For this section, I use unreported results in which the regressions are weighted by the firm's size employment share in the BLS. In the primary analysis, regressions are unweighted as I am intending to estimate causal effects. However, when estimating population descriptive statistics, weighting to correct for representativeness is needed (see Solon et al. (2015)).

First, I limit the analysis to the local effect in my sample. In this case, $N = 5,525$ (as of January 2020) and $\gamma = .57$. Using my primary estimate of β_t=8.0% for the entire post-PPP period from April-September, this implies 253 jobs preserved in the sample through September 2020. Noting that the PPP disbursed an estimated $39.6 million to firms in my sample, each job costs $157,000.[51] The results using the estimates from my main specification are summarized in Table V, in the "In-sample - main results" column.

Next, I assume that the estimate applies to all small firms (<300 employees) to be

[51] The estimate of funds distributed to firms in my sample is based on their total wage bill. The PPP allows firms to apply for 2.5 times their average monthly wage bill, based on the average in the year prior to the start of the pandemic.

consistent with the firms represented in my sample.[52] This results in an estimate of N of 57 million.[53] I use an estimated γ of 75%, consistent with my calculation in Section 3.4. Using again my primary (weighted) estimate of β_t=8.0% for the entire post-PPP period from April-September, this implies 3.43 million jobs preserved through September 2020. In aggregate the PPP disbursed \$465 billion to firm with fewer than 300 employees.[54] This implies a cost-per-job of \$135,000 as of September 2020. The results for this sub-sample are shown in the second column of Table V.

Thus, my in-sample estimates imply a slightly higher estimate than in the aggregate calculation. All in all, these estimates correspond to a cost of \$271,000-313,000 per job-year, assuming the cost is a monthly flow. Further, the estimates suggests that only 10-20% of the total funds disbursed went to wages at preserved jobs.[55]

Recall however that there is significant heterogeneity in treatment effects; only firms that were less affected by the pandemic, namely firms with remote capable workers, truly benefited from the programs. Aggregate estimates for this subsample on which the program was most effective provide an upper bound for the true effect of the program.

Consider the aggregate estimate for only the firms that benefited the most in my sample, those in industries with many employees who can work remotely, where $\beta =$ 13.1 in the main specification (when weighted by firm size). In my sample, these firms

[52]The largest firm in my sample, besides the firm with 500+ employees that was excluded from the analysis, has just over 300 employees.

[53]The Statistics of U.S. Businesses (SUSB) estimates that there are approximately 55 million employees at firms with less than 300 employees as of 2017. Following Autor et al. (2022), I scale up this number by an additional 3 percent, corresponding to the growth in private payrolls between December 2017 (the last year of the SUSB data) and December 2019 in the BLS's Current Employment Statistics data.

[54]To calculate this, I use the microdata provided by the Treasury. I calculate the total value of loans disbursed for each firm that indicated that the number of jobs retained was less than 300. Firms with jobs retained of less than 300 represent 99.8% of all firms in the data.

[55]To calculate this, I multiply the number of jobs preserved by the average wage that would have been earned from April-September 2020 by employees in each sample, then divide by total funds disbursed.

represent $N = 3,533$ and had an estimated total PPP loan value of $23.7 million. This implies a job-year cost of $88,000 as of September 2020 for the firms that benefited most from the program. In this case, approximately 19% of the total funds disbursed went directly to wages for preserved jobs. The results for this subsample are shown in the final column of Table V,

These estimates of course do not consider the cost of the alternative, had these employees been laid off. As an extreme example, suppose that each employee had been unable to find another job and instead received unemployment insurance (UI) for the six months from April to September.[56] Such an individual in my sample would have received approximately $4,284 over this period in state UI and approximately $9,600 in federal UI, for a total of $13,884.[57] Hence, even considering the cost of UI, the per-person estimate is reduced only by approximately $14,000. This is likely an upper-bound given most people did not remain continuously unemployed for six months.

My estimates imply a smaller per job-year cost than found in Autor et al. (2022)'s, where the preferred estimate is $224,000 per job as of the end of May 2020.[58] In comparison, Brown and Earle (2017) estimate that typical SBA loans cost only $21,580 to $25,450 per job per year. Feyrer and Sacerdote (2011) estimate a cost of $170,000 per job-year due to the American Recovery and Reinvestment Act of 2009. Thus the PPP appears to be more expensive than typical SBA loans, but similar to measures of other relief programs, though the PPP intends to *preserve* rather than create jobs. Moreover,

[56]This is an extreme, but still reasonable assumption given the high rate of unemployment over this time period. The BLS estimates that the majority of workers who lost their jobs during the pandemic had returned to work as of September 2020, though the labor force had also shrunk.

[57]The state UI estimate is based on the average amount per week given in the state in which most of the employees in my sample are located, and assuming they would be eligible for benefits for all six months of unemployment. The federal estimate is based on the additional $600 per week that was specified in CARES act from 3/29/2020-7/25/2020.

[58]This implies a per job-year cost of $1,2942,222.

the jobs are quite expensive even when focusing on the subsamples that the program appears to have targeted most effectively. While my estimates imply quite a high cost per job of the PPP, even for the most affected firms, it would be preliminary to offer a complete welfare analysis of the PPP from this small sample. The overall effects of the program must be considered along with other aid provided as part of the CARES act, such as the individual stimulus checks and increased unemployment benefits as well as the loss of healthcare benefits associated with losing a job. It would also be remiss to ignore the potential longer-run effects that the program might have on job preservation and business survival. Future work should seek to address this.

3.8 Discussion and Conclusion

Using administrative data from a payroll processor, I find that the PPP was largely successful in increasing employment at treated firms. Average employment at treated firms increased by 14% the month after application and 7.5% cumulatively over the five month period following PPP. On the other hand, there is no significant effect on wage bills. My results suggest a high per-job cost of the PPP, as much as $157,000 as of September 2020. At best, less than 20% of all funds disbursed went to wages at preserved jobs.

I also find that there is substantial heterogeneity in the effect of the PPP across types of firms. The employment preservation occurs almost entirely in industries that have a low percentage of hourly workers and more employees who can work remotely. The effect is also larger in essential businesses. This implies that the PPP does not act alone, but interacts with local economic conditions and restrictions in business activities.

This paper also provides insights into firm take-up of PPP loans. Take-up is estimated at 75% nationally, thus a non-significant number of eligible firms opted not to take these very financially attractive loans. My results suggest that smaller firms with lower average salaries are less likely to take-up the loans. This suggests that the least organizationally complex firms elected not to take the loans.

The results come with several caveats. First, the sample is not geographically diverse, with most firms being in one state. Moreover, this state was not particularly hard hit by the pandemic in the first wave in April and May. Second, the sample is small relative to all firms in the U.S. However, the sample is quite comparable to all small businesses in the U.S. in terms of its distribution by industries and average size.

There are several avenues for future work. First, after the end of my sample period, the PPP was adjusted in several ways to target smaller, more affected businesses. My results suggest that exactly these firms were less likely to take the loans. It remains an open question if those adjustments, like increasing assurances of the loans being forgivable and increasing the range of expenditures that the loans could be used for, alleviated this. Second, we should seek to understand why some firms did not apply. Though my results suggest that smaller, less complex firms were less likely to apply, more detailed data, including surveys of small business owners would help to understand better why this is the case. Lastly, examining the long-run outcomes of the firms that received the PPP versus those that did not will speak to the effectiveness and overall welfare effect of the program.

Figures

Figure 3-1: Distribution by NAICS Industries

(a) Number of Firms

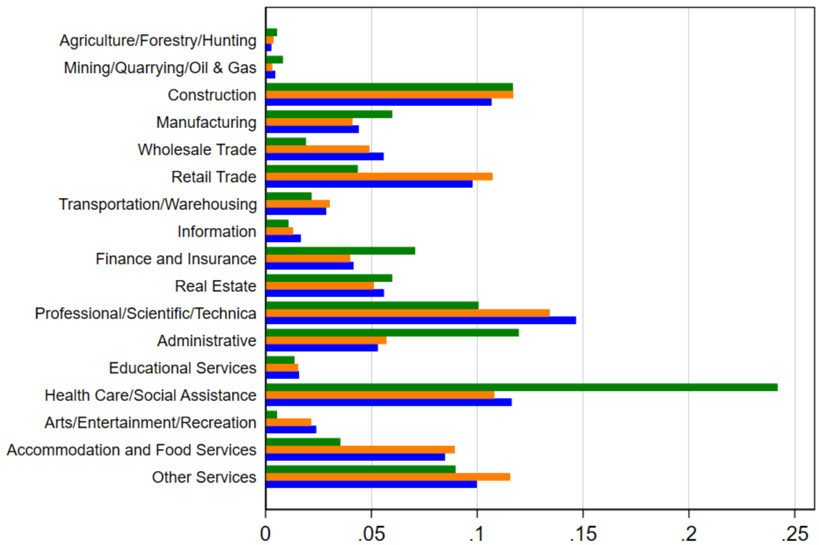

(b) Number of Employees

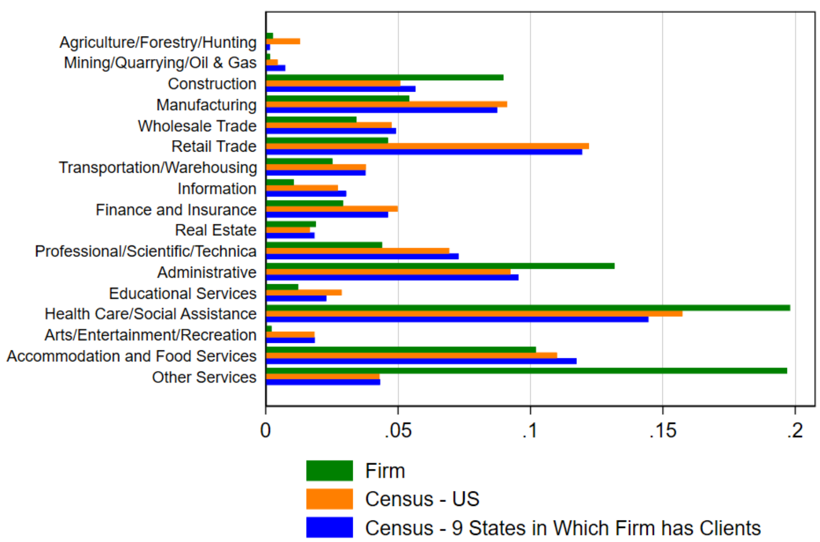

Notes: These figures show the distribution of the number (top) and employment (bottom) of firms in my sample by 2-digit NAICS code compared to the total U.S. and the nine states in which the firm has clients. Appendix Table C.1 shows the top panel in table format, with sample firms split out by PPP application status.

Figure 3-2: Odds Ratios of NAICS Fixed Effects

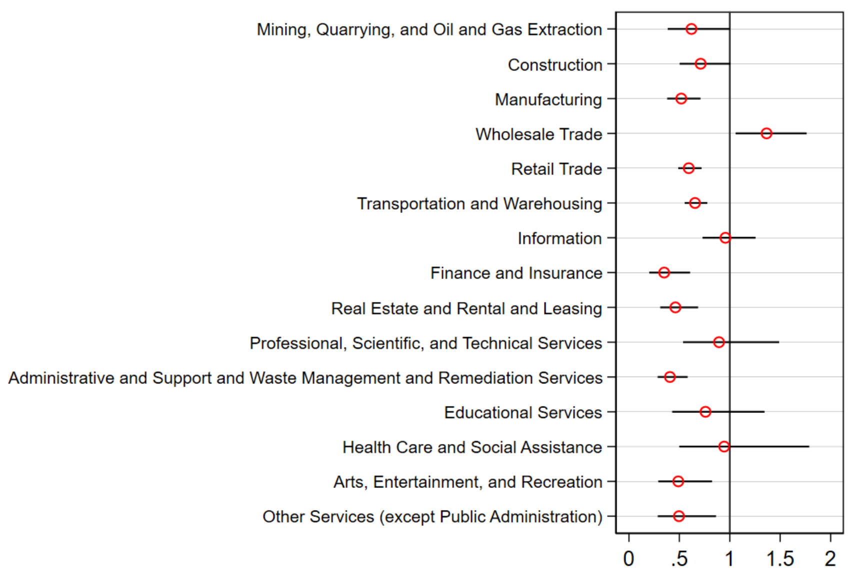

Notes: This figure shows the odds ratios for the NAICS industry fixed effects from the logistic regression of take-up on firms characteristics (equation (3.1)). The reference category is NAICS 72, Food and Accommodation Services. Standard errors are clustered at the 2-digit NAICS level. The coefficients correspond to the final column in Table II.

Figure 3-3: Effect of PPP on Employment, Wage Bill, and Wage per Worker

(a) Employment

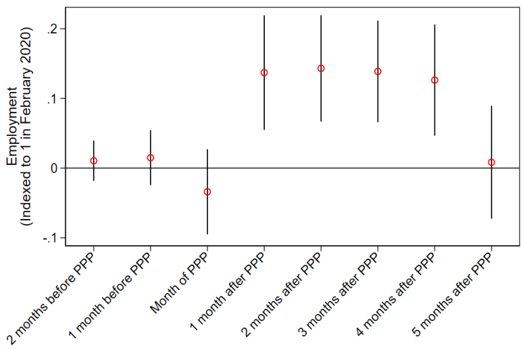

(b) Wage Bill

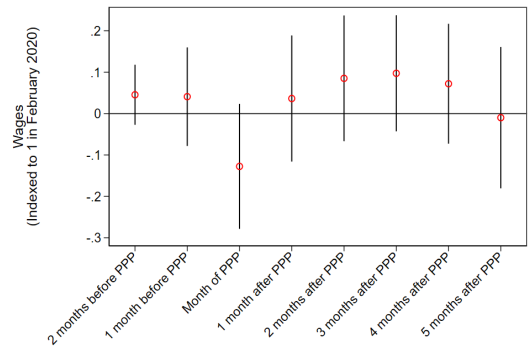

(c) Employment & Wage Bill, Average Effect

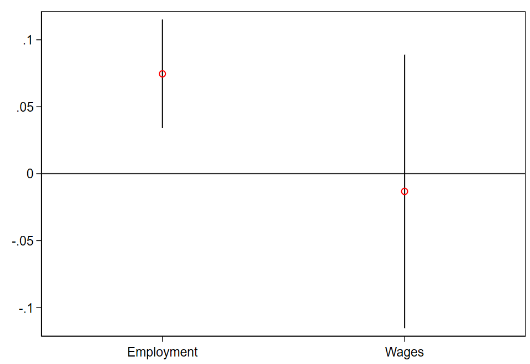

Notes: These figures show the results from the difference-in-differences specifications. Panels (a)-(b) correspond to equation (3.2) and panel (c) corresponds to equation (3.3). I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals.

Figure 3-4: Average Effect of PPP, by Firm Type

(a) By Prevalence of Hourly Workers

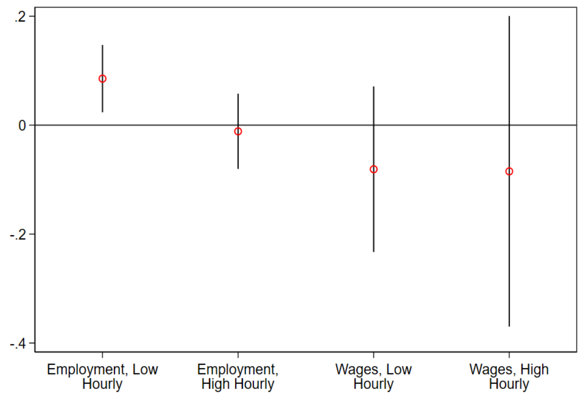

(b) By Remote Capability

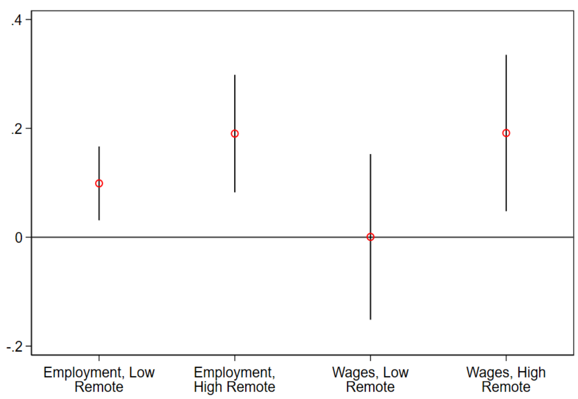

(c) By Essential Businesses

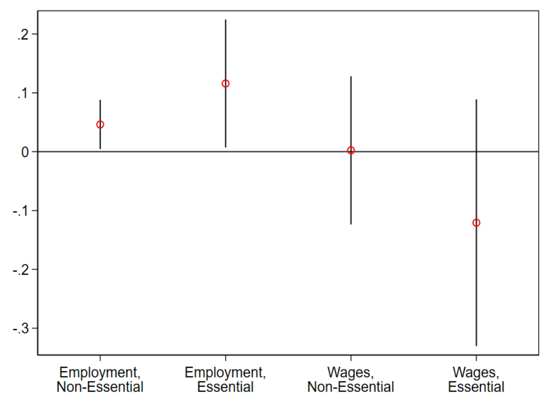

Notes: These figures show the results from the difference-in-differences specification in equation (3.3), split across different types of firms. I plot the coefficient on β_t, which estimates the difference between firms that applied for PPP and firms that did not after application to the PPP. The dependent variables are normalized to one in February 2020.

Figure 3-5: Effect of PPP on Employee Turnover

(a) Total Number of Hires

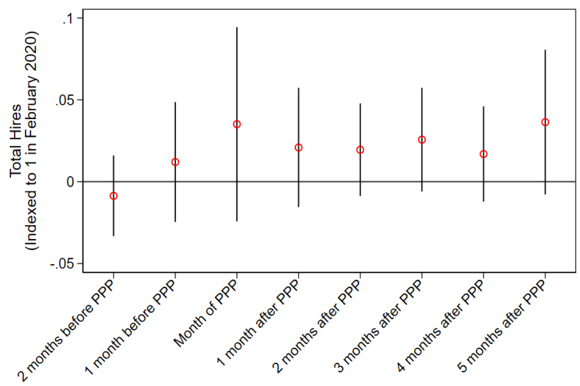

(b) Number of New Hires

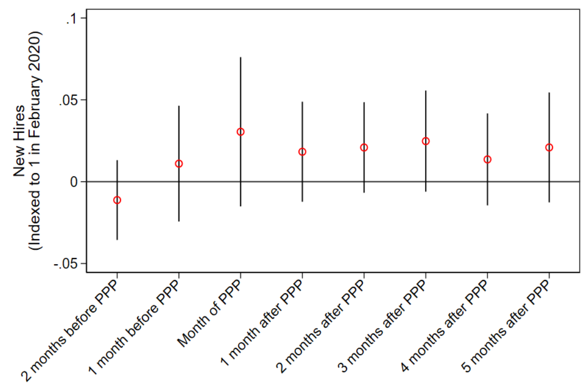

(c) Number of Rehires

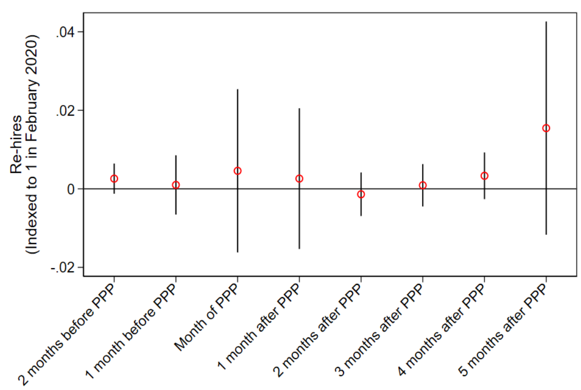

Notes: These figures show the results from the difference-in-differences specification in equation (3.2). I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variables are measured as fraction of the firm's base employment (average February 2019-February 2020). The black lines represent 90% confidence intervals.

Tables

Table I: Summary Statistics of Sample Firms

	Applied for PPP?			
	No	Yes	All	All Small Businesses < 500 Employees
Employment	8.90	17.82	13.95	9.32
	(2.67)	(8.15)	(5.31)	
Monthly Wage Bill ($)	33,276.77	49,518.42	42,477.50	36,737.91
	(6,868.77)	(26,817.52)	(16,229.22)	
Average monthly wages per worker	3,200.01	3,364.12	3,292.98	3,639.12
	(2,244.35)	(2,983.48)	(2,641.47)	
Observations	163	214	377	

Notes: This table shows summary statistics (means) for the firms in my sample. Medians are in parentheses. Employment and wage figures are based on the firms' averages from February 2019-February 2020 (for one year prior to the pandemic in the U.S.). The first and second columns show the statistics for firms that did not apply and applied for PPP, respectively. The third column shows the full sample. One firm in the sample has over 500 employees and therefore is not eligible for the PPP, so it is dropped from all analyses. The final column shows a comparison for all firms with under 500 employees in U.S., which are calculated from the 2016 Statistics for U.S. Businesses (Census Bureau).

Table II: Take-up Regressions: Full Sample

	(1)	(2)	(3)	(4)	(5)
			Odds Ratios		
Log Base Employment	1.7912	1.7654	1.8254	1.7903	1.7847
	(0.2620)	(0.2964)	(0.2391)	(0.2677)	(0.4841)
Log Average Salary (Base Period)	1.4498	1.4290	1.4654	1.4514	1.3958
	(0.2349)	(0.2285)	(0.1904)	(0.1755)	(0.1345)
Monthly Employment Growth (April 2020)		0.9643		0.9204	0.8478
		(0.3443)		(0.2785)	(0.3526)
Average industry employment change during pandemic			0.1994	0.2027	
			(0.1325)	(0.1438)	
Bank in 2nd Tercile of Overall PPP Lending					1.2385
					(0.5479)
Bank in 3rd Tercile of Overall PPP Lending					1.5148
					(0.2865)
Industry FE?	Y	Y	N	N	Y
State FE?	Y	Y	Y	Y	Y
N	345	342	342	339	291
pseudo R^2	0.126	0.120	0.112	0.107	0.121

Notes: This table shows the results of a logistic regression of an indicator equal to one if the firm applied for PPP on a set of controls for firm characteristics (equation (3.1)). The odds ratios are reported. Log Base Employment is the log of the firm's average employment from February 2019-February 2020 (or the months in which the firm appears in that data set over that time period). Log Base Average Salary is the firm's average monthly wage per worker over the same time period. Monthly Employment Growth (April 2020) is the firm's log employment growth from March 2020 to April 2020. Average industry employment change during pandemic is the percentage decline in paid employment for the firm's two-digit NAICS code from February 15, 2020- April 25, 2020 (taken from Cajner et al. (2020)). 2nd and 3rd tercile of overall PPP lending are dummies equal to one if the firm's primary bank is in the second or third (highest) tercile, respectively, of share of PPP lending less than $150,000 in the firm's state. Standard errors, clustered at the 2-digits NAICS code level, in parentheses.

Table III: Cross-Sectional Regressions on Employee Turnover

	Month of PPP Application	One Month After PPP Application	Two Months After PPP Application	Entire post-PPP period
		# of Total Hires		
Applied for PPP Loan	-0.6209	0.1600	-1.2754	0.0281
	(4.6565)	(4.4492)	(2.2023)	(0.0894)
N	355	348	344	334
R^2	0.024	0.047	0.120	0.053
		# of New Hires		
Applied for PPP Loan	-1.3991	-0.0057	-0.4816	-0.4580
	(3.5959)	(4.2099)	(2.1833)	(1.6480)
N	355	348	344	338
R^2	0.031	0.033	0.115	0.265
		# of Re-hires		
Applied for PPP Loan	0.7782	0.1658	-0.7938***	-0.5918
	(1.3467)	(1.4251)	(0.2819)	(0.5529)
N	355	348	344	338
R^2	0.055	0.125	0.172	0.351
Industry FE?	Y	Y	Y	Y
State FE?	Y	Y	Y	Y
Employment Size Controls?	Y	Y	Y	Y
Date Controls?	Y	Y	Y	N
Control for lagged employment growth?	Y	Y	Y	Y

Notes: This table shows the results of cross-sectional regressions of percentage of total, new and rehires on a PPP indicator. Coefficients are shown as percentages of the firm's employment in February 2020. The regressions also control for one-month lagged monthly employment growth, industry fixed effects, state fixed effects, employment size fixed effects and month fixed effects. The first column shows results for the month of PPP application (April for controls firms, April-June for treated firms). The second column shows results for one month following PPP application (May for control firms, May-July for treated firms). The third column shows results for two months following PPP application (June for control firms, June-August for treated firms). The fourth column is for the entire post-PPP period from April-September. Standard errors in parentheses. $p < 0.10$, $p < 0.05$, $p < .01$.

Table IV: Cross-Sectional Regressions on the Effect of PPP on the Likelihood of Rehiring, Hiring, and Reducing Employment

	Month of PPP Application	One Month After PPP Application	Two Months After PPP Application	Entire post-PPP period
		Rehired a Former Employee		
Applied for PPP Loan	1.6672	0.9246	0.0658***	0.5565
	(1.1116)	(0.4916)	(0.0627)	(0.2267)
N	309	321	273	326
pseudo R^2	0.388	0.371	0.584	0.403
		Hired Any Employee		
Applied for PPP Loan	1.0153	1.5492	1.0583	1.1839
	(0.4130)	(0.6224)	(0.4703)	(0.3772)
N	337	329	330	328
pseudo R^2	0.237	0.307	0.400	0.359
		Reduced Employment		
Applied for PPP Loan	1.0653	0.6291	1.5359	0.9230
	(0.4091)	(0.2612)	(0.7653)	(0.2705)
N	342	336	300	325
pseudo R^2	0.252	0.178	0.179	0.204
Industry FE?	Y	Y	Y	Y
State FE?	Y	Y	Y	Y
Employment Size Controls?	Y	Y	Y	Y
Date Controls?	Y	Y	Y	N
Control for lagged employment growth?	Y	Y	Y	Y

Notes: This table shows the results of a logistic regression on the extensive margin of employment outcomes. The dependent variable in the to panel is an indicator equal to one if the firm rehired former employees in a given month and zero otherwise. The dependent variable in the second panel is an indicator equal to one if the firm hired any employees in a given month and zero otherwise. The dependent variable in the bottom panel is an indicator equal to one if the firm reduced employment from February 2020 to September 2020 and zero otherwise. The regressions also control for one-month lagged employment growth, industry fixed effects, state fixed effects, employment size, and month fixed effects. Odds-ratios are reported. Standard errors in parentheses.

Table V: Estimate of Aggregate Costs

	In-sample - main results	Aggregate - Firms with < 300 Employees	In-sample - most remote-capable industries
β_t	8.0%	8.0%	13.1
γ	57%	75%	58%
N	5,525	58 million	3,533
Total jobs preserved	253	3.43 million	268
Total funds disbursed	$39.6 million	$465 billion	$23.7 million
Cost per job			
April-September 2020	$157,000	$135,000	$88,000
With UI Benefits	$143,000	$121,000	$74,000
Per year	$313,000	$271,000	$177,000
% of total funds that goes to wages of preserved jobs	12.6%	19.3%	19.4%

Notes: This table shows the estimates for the aggregate cost and cost-per-job of the PPP, quantified in equation (3.4). Cost per job April-September is the total funds disbursed divided by the total jobs preserved. Cost per job April-September 2020 (accounting for UI benefits) is the total cost per-job less the estimated amount than an individual would have received in UI benefits from both their state and the federal government, had they remained unemployed from April-September. This estimate is based on an average of state UI benefits corresponding to the sample in each column and the $600 per week provided by the Federal Government via the CARES Act through July 25, 2020. Cost-per job year is a scaled up estimate of Cost per-job April-September 2020, assuming the cost is a monthly flow. Percentage of total funds that goes to wages of preserved jobs is the number of jobs preserved multiplied by the average dollar value of wages a worker in each sample would have earned from April-September 2020, divided by total funds disbursed.

Bibliography

Altonji, Joseph, Zara Contractor, Lucas Finamor, Ryan Haygood, Ilse Lindenlaub, Costas Meghir, Cormac O'Dea, Dana Scott, Liana Wang, and Ebonya Washington, "Employment Effects of Unemployment Insurance Generosity During the Pandemic," Working Paper, Tobin Center for Economic Policy at Yale University 2020.

Ameriks, John and Stephen Zeldes, "How Do Household Portfolio Shares Vary with Age," TIAA-CREF Working Paper, Columbia University, New York, 2004.

_ , Andrew Caplin, Minjoon Lee, Matthew D Shapiro, and Christopher Tonetti, "The Wealth of Wealthholders," Working Paper 20972, National Bureau of Economic Research February 2015.

_ , Gábor Kézdi, Minjoon Lee, and Matthew D. Shapiro, "Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares: The Attenuation Puzzle," Journal of Business and Economic Statistics, 2019, pp. 1–27.

Arnoud, Antoine, Taha Choukhmane, Jorge Colmenares, Cormac O'Dea, and Aneesha Parvathaneni, "The Evolution of U.S. Firms' Retirement Plan Offerings: Evidence from a New Panel Data Set," NBER Working Paper, 2021, NB20-14.

Autor, David, David Cho, Leland D. Crane, Mita Goldar, Byron Lutz, Joshua Montes, William B. Peterman, David Ratner, Daniel Villar, and Ahu Yildirmaz, "An evaluation of the Paycheck Protection Program using administrative payroll microdata," Journal of Public Economics, 2022, 211, 104664.

_ , David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms*," The Quarterly Journal of Economics, 02 2020, 135 (2), 645–709.

Azar, José A, Steven T Berry, and Ioana Marinescu, "Estimating Labor Market Power," Working Paper 30365, National Bureau of Economic Research August 2022.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska, "Concentration in US labor markets: Evidence from online vacancy data," Labour Economics, 2020, 66, 101886.

Baicker, Katherine and Amitabh Chandra, "The Labor Market Effects of Rising Health Insurance Premiums," Journal of Labor Economics, 2006, 24 (3), 609–634.

Balyuk, Tetyana, Nagpurnanand R Prabhala, and Manju Puri, "Indirect Costs of Government Aid and Intermediary Supply Effects: Lessons From the Paycheck Protection Program," Working Paper 28114, National Bureau of Economic Research November 2020.

Barrios, John Manuel, Michael Minnis, William C. Minnis, and Joost Sijthoff, "Assessing the Payroll Protection Program: A Framework and Preliminary Results (May 13, 2020)," Working Paper 2020.

Bartik, Alexander W., Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matthew Unrath, "Measuring the Labor Market at the Onset of the COVID-19 Crisis," Brookings Papers on Economic Activity, BPEA Conference Drafts, 2020.

Bartik, Alexander W, Marianne Bertrand, Zoë B Cullen, Edward L Glaeser, Michael Luca, and Christopher T Stanton, "How Are Small Businesses Adjusting to COVID-19? Early Evidence from a Survey," Working Paper 26989, National Bureau of Economic Research April 2020.

_ , Zoe B Cullen, Edward L Glaeser, Michael Luca, Christopher T Stanton, and Adi Sunderam, "The Targeting and Impact of Paycheck Protection Program Loans to Small Businesses," Working Paper 27623, National Bureau of Economic Research July 2020.

Becker, Dan, "Non-Wage Job Characteristics and the Case of the Missing Margin," SSRN Electronic Journal, 04 2011.

Ben-David, Itzhak, Mark J Johnson, and René M Stulz, "Why Did Small Business Fintech Lending Dry Up During March 2020?," Working Paper 29205, National Bureau of Economic Research September 2021.

Bennett, Nathan, Terry C. Blum, Rebecca G. Long, and Paul M. Roman, "A Firm-Level Analysis of Employee Attrition," Group & Organization Management, 1993, 18 (4), 482–499.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, "Portfolio choice over the life-cycle when the stock and labor markets are cointegrated," 2007, (WP-07-11).

Bernheim, B. Douglas, Andrey Fradkin, and Igor Popov, "The Welfare Economics of Default Options in 401(k) Plans," American Economic Review, September 2015, 105 (9), 2798–2837.

Beshears, John, James Choi, David Laibson, Brigitte C. Madrian, and Brian Weller, Public Policy and Saving for Retirement: The Autosave Features of the Pension Protection Act of 2006, Harvard University Press, BLS, https://www.bls.gov/opub/reports/minimum-wage/2015/home.htm 2016. Accessed: 2021-08-31.

_, "Employment by major industry sector," https://www.bls.gov/emp/tables/employment-by-major-industry-sector.htm 2020. Accessed: 2021-08-31.

_, "Occupational Employment and Wages Summary," https://www.bls.gov/news.release/ocwage.nr0.htm#modal-box 2020. Accessed: 2021-08-31.

Bonhomme, Stéphane and Grégory Jolivet, "The pervasive absence of compensating differentials," Journal of Applied Econometrics, 2009, 24 (5), 763–795.

Brewer, William L., Douglas R. Brown, and S. Graham Simmons III, "Government Provides Relief in the Paycheck Protection Program Flexibility Act," National Law Review, 2020, X (162).

Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus, "Measuring income and wealth at the top using administrative and survey data," Brookings Papers on Economic Activity, 2016, 1, 261–331.

Brown, J. David and John S. Earle, "Finance and Growth at the Firm Level: Evidence from SBA Loans," Journal of Finance, 2017, 72, 1039–1080.

Bubb, Ryan and Patrick L. Warren, "An Equilibrium Theory of Retirement Plan Design," American Economic Journal: Economic Policy, May 2020, 12 (2), 22–45.

_, Patrick Corrigan, and Patrick L. Warren, "A Behavioral Contract Theory Perspective on Retirement Savings," Connecticut Law Review, 2015, 47.

Burdett, Kenneth and Dale Mortensen, "Wage Differentials, Employer Size, and Unemployment," International Economic Review, 1998, 39 (2), 257–73.

Cajner, Tomaz, Leland D Crane, Ryan A Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz, "The U.S. Labor Market during the Beginning of the Pandemic Recession," Working Paper 27159, National Bureau of Economic Research May 2020.

Campbell, John Y., "Restoring Rational Choice: The Challenge of Consumer Financial Regulation," American Economic Review: Papers & Proceedings, 2016, 106 (5), 1–30.

_ and Luis M. Viceira, Strategic Asset Allocation, Oxford University Press, 2002.

Carroll, Christopher, "Portfolios of the Rich," NBER Working Papers 7826, National Bureau of Economic Research, Inc 2000.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team, "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data," Working Paper 27431, National Bureau of Economic Research June 2020.

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" NBER Chapters, National Bureau of Economic Research, Inc, September 2004, pp. 81–126.

Choukhmane, Taha, "Default Options and Retirement Saving Dynamics," Working Paper, 2019.

_ , Lucas Goodman, and Cormac O'Dea, "Efficiency in Household Decision Making: Evidence from the Retirement Savings of US Couples," Working Paper, 2022.

Christelis, Dimitris, Dimitris Georgarakos, and Michael Haliassos, "Differences in Portfolios across Countries: Economic Environment versus Household Characteristics," The Review of Economics and Statistics, March 2013, 95 (1), 220–236.

Curcuru, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, "Heterogeneity and Portfolio Choice: Theory and Evidence," in Yacine Ait-Sahalia and Lars Peter Hansen, eds., Handbook of Financial Econometrics: Tools and Techniques, London: North-Holland, 2010, chapter 6.

Denes, Matthew, Spyridon Lagaras, and Margarita Tsoutsoura, "First Come, First Served: The Timing of Government Support and Its Impact on Firms," Working Paper 2021.

Doniger, Cynthia L. and Benjamin S. Kay, "Ten Days Late and Billions of Dollars Short: The Employment Effects of Delays in Paycheck Protection Program Financing," Finance and Economics Discussion Series 2021-003, Board of Governors of the Federal Reserve System (U.S.) January 2021.

Duarte, Victor, Julia Fonseca, Aaron Goodman, and Jonathan A. Parker, "Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle," NBER Working Paper w29559, December 2021.

Engelhardt, Gary V. and Anil Kumar, "Employer matching and 401(k) saving: Evidence from the health and retirement study," Journal of Public Economics, 2007, 91 (10), 1920–1943. Special Issue published in cooperation with the National Bureau of Economic Research and Uppsala University: Proceedings of the Trans-Atlantic Public Economics Seminar on Public Policy and Retirement Behavior June 12-14 2006.

Erel, Isil and Jack Liebersohn, "Does FinTech Substitute for Banks? Evidence from the Paycheck Protection Program," Journal of Financial Economics, Forthcoming.

Fadlon, Itzik, Jessica A Laird, and Torben Heien Nielsen, “Do Employer Pension Contributions Reflect Employee Preferences? Evidence from a Retirement Savings Reform in Denmark,” American Economic Journal: Applied Economics, 2016, 8 (3), 196–216.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso, “Asset Market Participation and Portfolio Choice over the Life-Cycle,” The Journal of Finance, 04 2017, 72.

Faulkender, Michael, Robert Jackman, and Stephen Miran, “The Job Preservation Effects of Paycheck Protection Program Loans,” Working Paper January 2020.

Feyrer, James and Bruce Sacerdote, “Did the Stimulus Stimulate? Real Time Estimates of the Effects of the American Recovery and Reinvestment Act,” Working Paper 16759, National Bureau of Economic Research February 2011.

Galuscak, Kamil, Mary Keeney, Daphne Nicolitsas, Frank Smets, Pawel Strzelecki, and Matija Vodopivec, “The determination of wages of newly hired employees: Survey evidence on internal versus external factors,” Labour Economics, 2012, 19 (5), 802–812. Special Section on: Price, Wage and Employment Adjustments in 2007-2008 and Some Inferences for the Current European Crisis.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, “Five Facts about Beliefs and Portfolios,” NBER Working Paper No. 25744 2019.

Goda, Gopi Shah, Damon Jones, and Colleen Flaherty Manchester, “Retirement Plan Type and Employee Mobility: The Role of Selection,” Journal of Human Resources, 2017, 52 (3), 654–679.

Goldhaber, Dan, Cyrus Grout, and Kristian L. Holden, “Pension Structure and Employee Turnover: Evidence from a Large Public Pension System,” ILR Review, 2017, 70 (4), 976–1007.

Gomes, Francisco, Alexander Michaelides, and Yuxin Zhang, “Tactical target date funds,” Management Science, December 2020.

_ and Oksana Smirnova, “Stock Market Participation and Portfolio Shares Over the Life-Cycle,” Working Paper, 01 2021.

_ , Kenton Hoyem, Wei-Yin Hu, and Enrichetta Ravina, “Retirement Savings Adequacy in U.S. Defined Contribution Plans,” Working Paper, 01 2018.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, “Consumption Over the Lifecycle,” Econometrica, January 2002, 70 (1), 47–89.

Granja, João, Christos Makridis, Constantine Yannelis, and Eric Zwick, “Did the Paycheck Protection Program Hit the Target?,” Journal of Financial Economics, Forthcoming.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz, "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," American Economic Review, February 2021, 111 (2), 428–71.

Guiso, Luigi, Michael Haliassos, and Tullio Jappelli, "Household stockholding in Europe: where do we stand and where do we go?," Economic Policy, April 2003, 18 (36), 123–170.

_ , _ , and _ , "Stockholding: A Europen Comparison," in Luigi Guiso, Michael Haliassos, and Tullio Jappelli, eds., Stockholding in Europe, London: Palgrave Macmillans, 2003, chapter 1.

Gustman, Alan L. and Thomas L. Steinmeier, "The Influence of Pensions on Behavior: How Much Do We Really Know?," Working Paper RD71, TIAA-CREF Institute 2002.

_ , Olivia S. Mitchell, and Thomas L. Steinmeier, "The Role of Pensions in the Labor Market: A Survey of the Literature," Industrial and Labor Relations Review, 1994, 47 (3), 417–438.

Hall, Robert E. and Andreas I. Mueller, "Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values," Journal of Political Economy, 2018, 126 (4), 1594–1637.

Hare, Neil, "Loan Forgiveness Under The PPP And SBA EIDL Programs: 10 Things Small Businesses Need To Know," Forbes, 2020.

Hayashi, Yuka, "Demand for Small-Business Loans Cools," The Wall Street Journal, 2020.

Hazell, Joe, Christina Patterson, Heather Sarsons, and Bledi Taska, "National Wage Setting," Working Paper, 2022.

Hazell, Jonathon and Bledi Taska, "Downward Rigidity in the Wage for New Hires," ERN: Wages; Intergenerational Income Distribution (Topic), 2020.

Heaton, John and Deborah Lucas, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," The Journal of Finance, 2000, 55 (3), 1163–1198.

Hershbein, Brad and Lisa B. Kahn, "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," American Economic Review, July 2018, 108 (7), 1737–72.

Horan, Stephanie, "PPP Loan Acceptance Rates in the First and Second Rounds of Funding – 2020 Study," yahoo!, 2020.

Hubbard, Glenn and Michael R. Strain, "Has the Paycheck Protection Program Succeeded?," Conference Draft, Brookings Paper on Economic Activity Fall 2020.

Iacurci, Greg, "Small businesses expect to fail. Here's what will happen to their PPP loans.," CNBC, 2020.

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," Econometrica, 1994, 62 (2), 467–475.

Investment Company Institute, Investment Company Fact Book, Investment Company Institute, Washington, DC, 2020.

Joaquim, Gustavo and Felipe Netto, "Bank incentives and the effect of the paycheck protection program," Working Papers 21-15, Boston, MA 2021.

Johnson, Brian R., "The Golden Goose in the Crosshairs: The Transition to Defined Contribution Pension Plans in the Public Sector: Unintended Consequences," Journal of Health and Human Services Administration, 2013, 35 (4), 414–468.

Kimball, Spencer, "Congress approved $370 billion more in small business loans, but the money could run out in days — again," CNBC, 2020.

Kristal, Tali, Yinon Cohen, and Edo Navot, "Workplace Compensation Practices and the Rise in Benefit Inequality," American Sociological Review, 2020, 85 (2), 271–297.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," American Economic Review, January 2022, 112 (1), 169–212.

Landier, Augustin and David Thesmar, "Financial Contracting with Optimistic Entrepreneurs," The Review of Financial Studies, 2009, 22 (1), 117–150.

Lee, Chun-Hsien, Mu-Lan Hsu, and Nai-Hwa Lien, "The impacts of benefit plans on employee turnover: a firm-level analysis approach on Taiwanese manufacturing industry," The International Journal of Human Resource Management, 2006, 17 (11), 1951–1975.

Lehmann, Tobias, "Non-Wage Job Values and Implications for Inequality," Working Paper 2022.

Leibovici, Fernando, Ana Maria Santacreu, and Matthew Famiglietti, "How the Impact of Social Distancing Ripples through the Economy," On The Economy Blog, Federal Reserve Bank of St. Louis, April 2020.

Lelarge, Claire, David Sraer, and David Thesmar, Entrepreneurship and Credit Constraints: Evidence from a French Loan Guarantee Program, University of Chicago Press, May

Li, Lei and Philip E. Strahan, "Who Supplies PPP Loans (and Does It Matter)? Banks, Relationships, and the COVID Crisis," Journal of Financial and Quantitative Analysis, 2021, 56 (7), 2411–2438.

Lusardi, Annamaria and Olivia S. Mitchell, "Baby Boomer retirement security: The roles of planning, financial literacy, and housing wealth," Journal of Monetary Economics, January 2007, 54 (1), 205–224.

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

Maestas, Nicole, Kathleen J Mullen, David Powell, Till von Wachter, and Jeffrey B Wenger, "The Value of Working Conditions in the United States and Implications for the Structure of Wages," Working Paper 25204, National Bureau of Economic Research October 2018.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," American Economic Review, December 2017, 107 (12), 3722–59.

McDonald, Robert L., David P. Richardson, and Thomas A. Rietz, "The effect of default target date funds on retirement savings allocations," TIAA Institute, Research Dialog, June 2019, No. 150.

Meeuwis, Maarten, "Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios," 2019. MIT Sloan Working Paper.

_ , Jonathan A. Parker, Antoinette Schoar, and Duncan Simester, "Belief Disagreement and Portfolio Choice," NBER Working Paper No. 25108 2018.

Merton, Robert C., "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," Review of Economics and Statistics, 1969, 51 (3), 247–57.

Miller-Merrell, Jessica, "Companies Boost 401(k) Benefits to Retain Workers - The Wall Street Journal Google Your News Update - WSJ Podcasts," `https://www.wsj.com/podcasts/google-news-update/companies-boost-401k-benefits-to-retain-workers/79ab3386-021d-4654-9fa5-6db60a672c8b?mod=error_page` April 2013. (Accessed on 10/18/2022).

Miller, Richard D., "Estimating the Compensating Differential for Employer-Provided Health Insurance," International Journal of Health Care Finance and Economics, 2004, 4 (1), 27–41.

Mitchell, Olivia S and Stephen Utkus, "Target Date Funds and Portfolio Choice in 401(k) Plans," January 2020, (26684).

Mitchell, Olivia S., Stephen P. Utkus, and Tongxuan (Stella) Yang, "Turning Workers into Savers? Incentives, Liquidity, and Choice in 401 (k) Plan Design," National Tax Journal, 2007, 60 (3), 469–489.

Moore, Brett, “PPP Loans Sparking Disputes and Differing Results in Bankruptcy Proceedings,” New York Law Journal, 2020.

Ouimet, Paige and Geoffrey A. Tate, “Firms with Benefits? Nonwage Compensation and Implications for Firms and Labor Markets,” 2022.

Papanikolaou, Dimitris and Lawrence D W Schmidt, “Working Remotely and the Supply-Side Impact of COVID-19,” The Review of Asset Pricing Studies, 11 2021, 12 (1), 53–111.

Parker, Jonathan A, Antoinette Schoar, Allison T Cole, and Duncan Simester, “Household Portfolios and Retirement Saving over the Life Cycle,” Working Paper 29881, National Bureau of Economic Research March 2022.

_ , _ , and Yang Sun, “Retail Financial Innovation and Stock Market Dynamics: The Case of Target Date Funds,” October 2020, (28028).

Péer, Eyal, David M. Rothschild, Andrew Gordon, Zak Evernden, and Ekaterina Damer, “Data quality of platforms and panels for online behavioral research,” Behavior Research Methods, 2021, 54, 1643 – 1662.

Peer, Eyal, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti, “Beyond the Turk: Alternative platforms for crowdsourcing behavioral research,” Journal of Experimental Social Psychology, 2017, 70, 153–163.

Petersen, Mitchell A. and Raghuram G. Rajan, “The Benefits of Lending Relationships: Evidence from Small Business Data,” The Journal of Finance, 1994, 49 (1), 3–37.

Poterba, James M., “Retirement Security in an Aging Population,” American Economic Review, May 2014, 104 (5), 1–30.

_ and Andrew Samwick, “Household Portfolio Allocation over the Life Cycle,” in “Aging Issues in the United States and Japan” NBER Chapters, National Bureau of Economic Research, Inc, September 2001, pp. 65–104.

Poterba, James, Steven Venti, and David Wise, “The Composition and Drawdown of Wealth in Retirement,” Journal of Economic Perspectives, Fall 2011, 25 (4), 95–118.

Rosen, Sherwin, “Chapter 12 The theory of equalizing differences,” in “in,” Vol. 1 of *Handbook of Labor Economics*, Elsevier, 1986, pp. 641–692.

Samuelson, Paul, “Lifetime Portfolio Selection by Dynamic Stochastic Programming,” The Review of Economics and Statistics, 1969, 51 (3), 239–46.

SBA, “Paycheck Protection Program Report through June 30, 2020,” `https://www.sba.gov/document/report-paycheck-protection-program-report-through-june-30-2020` 2020. Accessed: 2021-08-31.

—, "First Draw PPP loan," https://www.sba.gov/funding-programs/loans/covid-19-relief-options/paycheck-protection-program/first-draw-ppp-loan#section-header-2 2021. Accessed: 2021-08-31.

—, "Paycheck Protection Program," https://www.sba.gov/funding-programs/loans/covid-19-relief-options/paycheck-protection-program 2021. Accessed: 2021-08-31.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, "Are Americans Saving "Optimally" for Retirement?," Journal of Political Economy, 2006, 114 (4), 607–43.

Schubert, Gregor, Anna Stansbury, and Bledi Taska, "Employer Concentration and Outside Options," Working Paper RD71, HBS 2020.

Sheiner, Louise, "Health care costs, wages, and aging," Finance and Economics Discussion Series 1999-19, Board of Governors of the Federal Reserve System (U.S.) 1999.

Sialm, Clemens, Laura T. Starks, and Hanjiang Zhang, "Defined Contribution Pension Plans: Sticky or Discerning Money?," The Journal of Finance, 2015, 70 (2), 805–838.

Simon, Kosali Ilayperuma and Robert Kaestner, "Do Minimum Wages Affect Non-wage Job Attributes? Evidence on Fringe Benefits and Working Conditions," Working Paper 9688, National Bureau of Economic Research May 2003.

Sockin, Jason, "Show me the Amenity: Compensating Differentials and the Minimum Wage," Working Paper, 2021.

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge, "What Are We Weighting For?," The Journal of Human Resources, 2015, 50 (2), 301–316.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," The Quarterly Journal of Economics, 01 2018, 133 (3), 1331–1393.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron, "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," Review of Economic Dynamics, 2007, 10 (4), 519–548.

Sullivan, Paul and Ted To, "Search and Nonwage Job Characteristics," The Journal of Human Resources, 2014, 49 (2), 472–507.

Taber, Christopher and Rune Vejlin, "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market," Econometrica, 2020, 88 (3), 1031–1069.

Viceira, Luis M, "Optimal portfolio choice for long-horizon investors with nontradable labor income," The Journal of Finance, 2001, 56 (2), 433–470.

Wachter, Jessica, "Asset Allocation," Annual Reviews of Financial Economics, 2010, 2, 175–206.

Wachter, Jessica A. and Motohiro Yogo, "Why Do Household Portfolio Shares Rise in Wealth?," Review of Financial Studies, November 2010, 23 (11), 3929–3965.

Wasick, John F., "Think a 401(k) Is Not a Sexy Benefit? Competition May Change That - The New York Times," https://www.nytimes.com/2016/09/22/business/smallbusiness/think-a-401-k-is-not-a-sexy-benefit-competition-may-change-that.html September 2016. (Accessed on 10/18/2022).

Weber, Lauren, "In Battle for Workers, the Humble 401(k) Gets Richer in 2022 - WSJ," https://www.wsj.com/articles/in-battle-for-workers-the-humble-401-k-gets-richer-in-2022-11642501803 January 2022. (Accessed on 10/18/2022).

Whalen, J.R. and Anne Tergensen, "Companies Boost 401(k) Benefits to Retain Workers - The Wall Street Journal Google Your News Update - WSJ Podcasts," https://www.wsj.com/podcasts/google-news-update/companies-boost-401k-benefits-to-retain-workers/79ab3386-021d-4654-9fa5-6db60a672c8b?mod=error_page January 2022. (Accessed on 10/18/2022).

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," The Quarterly Journal of Economics, 2018, 133 (1), 457–507.

Appendix A

Appendix for Chapter 1

A.1 Appendix Figures

Figure A.1: NDT Failure over Time

(a) % of DC Plans with a Corrective Distribution or NDT Failure, by Size

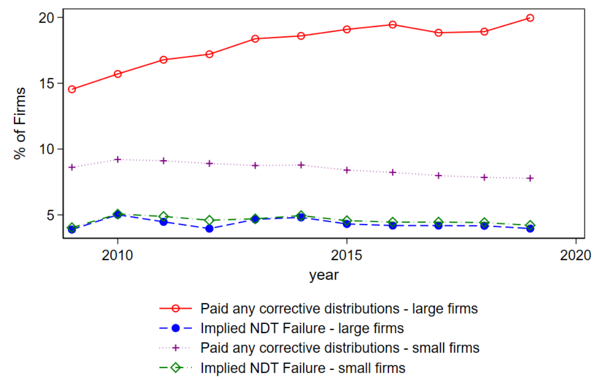

(b) $ in Corrective Distributions paid per Participant, Conditional on Paying Any

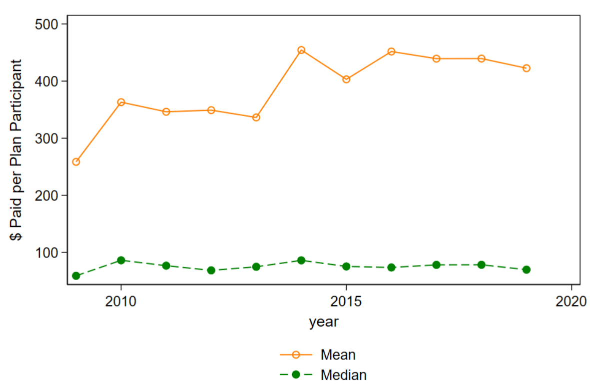

(c) $ in Corrective Distributions paid per Participant, Conditional on Paying Any, by Size

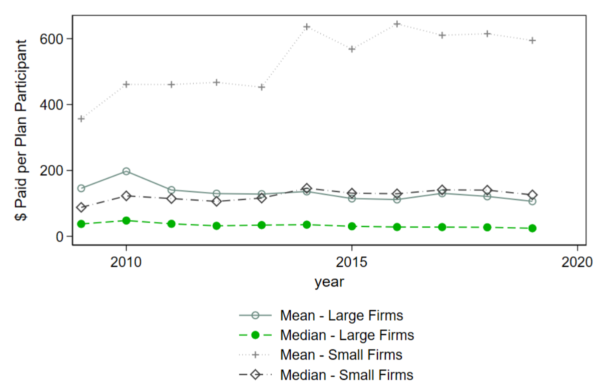

Notes: My calculations from the Form-5500 data, 2010-2019. Includes only DC plans. Dollar amounts in figures c and d are conditional on firms that paid some corrective distributions. Large plans are those with greater than 100 participants.

Figure A.2: National Wage Setting Instrument

(a) By Cost of Living

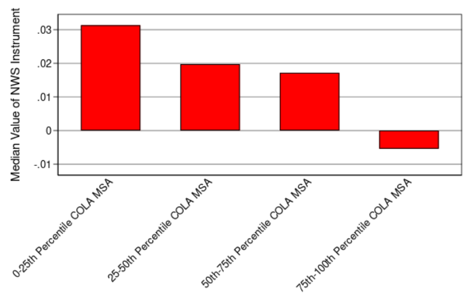

(b) By Population

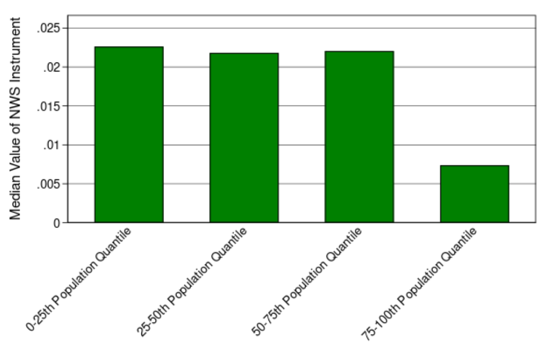

(c) By Firm Size

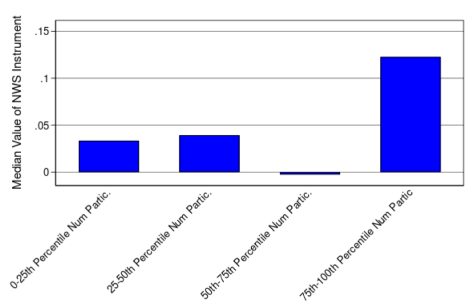

Notes: These figures show the median value of the national wage setting instrument, which measures the deviation of a given job's wage from the expected wage that is due to national wage setting.

Figure A.3: Percentage of Labor Force Captured in Lightcast Resume Data

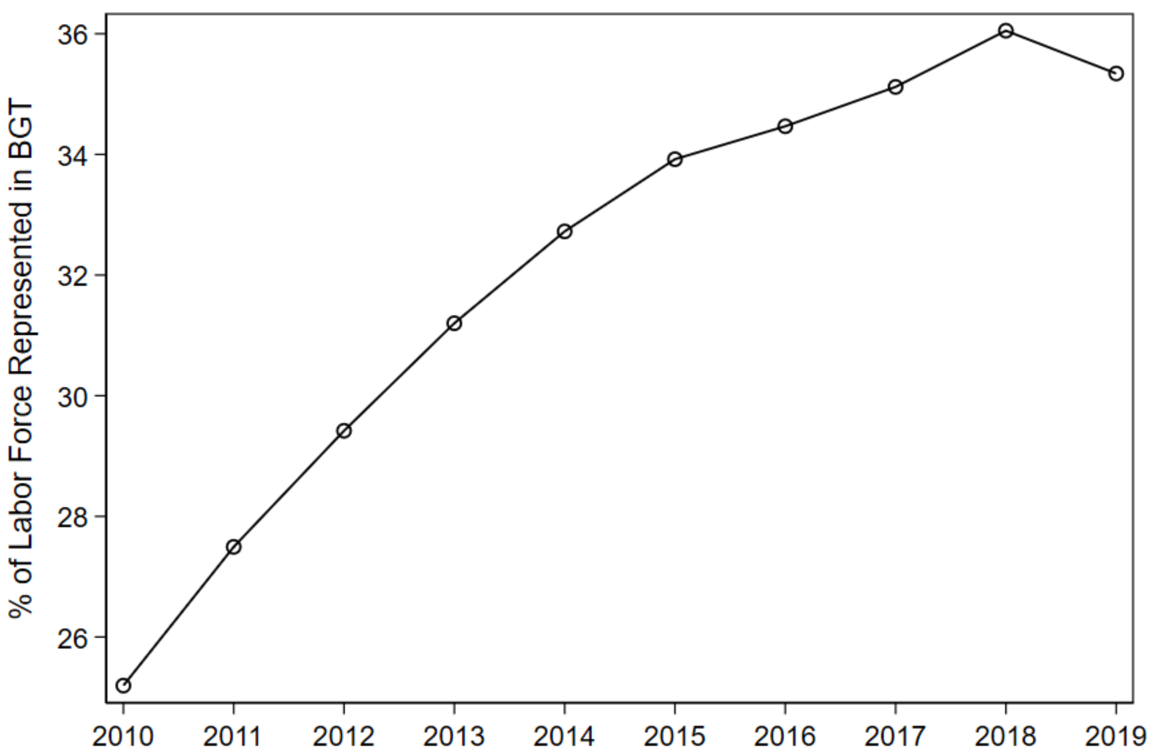

Notes: This figure shows the percentage of the workforce that is captured in the Lightcast resume data each year in our sample. The total workforce is from the Bureau of Labor Statistic's Occupation Employment Statistics. The Lightcast count includes all resumes.

Figure A.4: Lightcast Resume Representativeness

(a) Age in Lightcast versus BLS

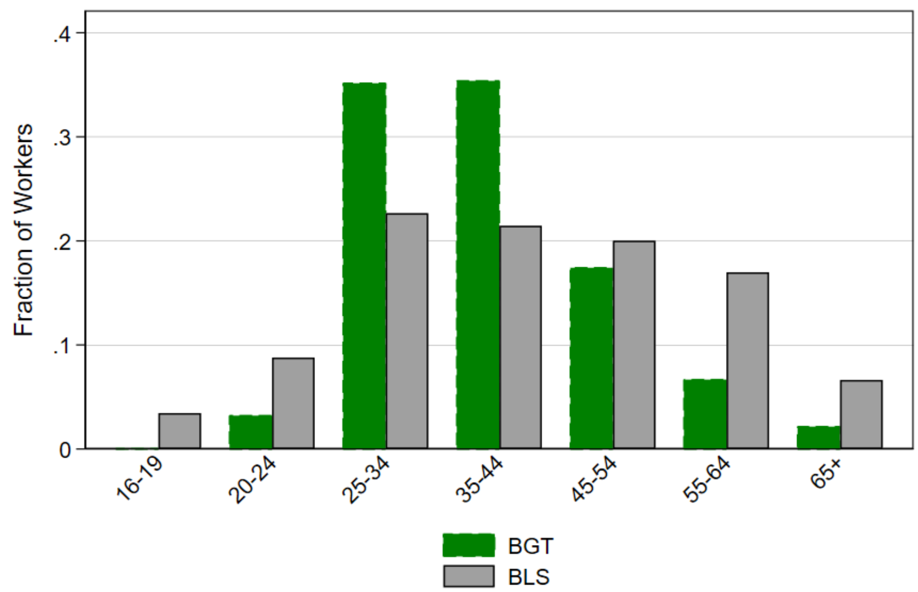

(b) Education in Lightcast versus BLS

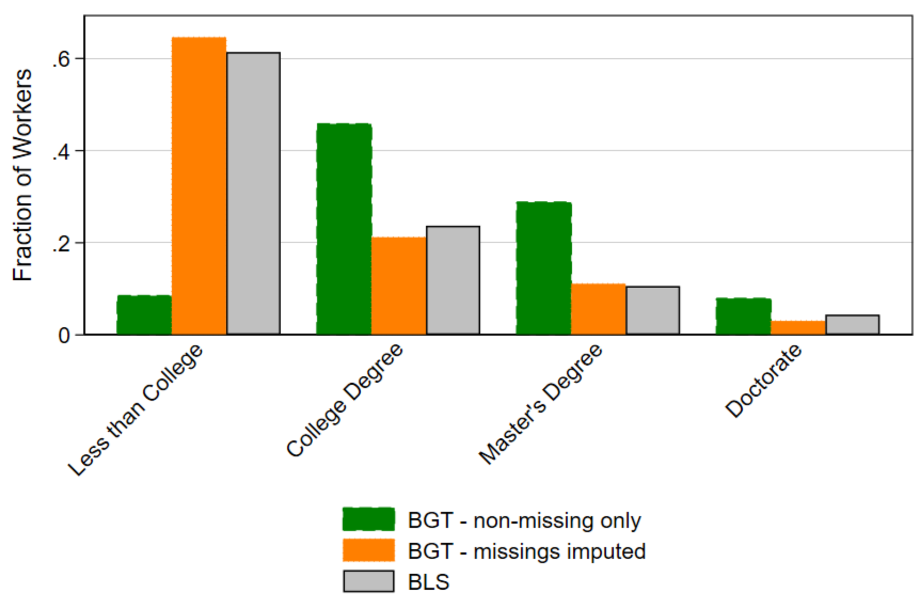

Figure A.4: Lightcast Resume Representativeness (continued)

(c) Occupation in Lightcast versus BLS

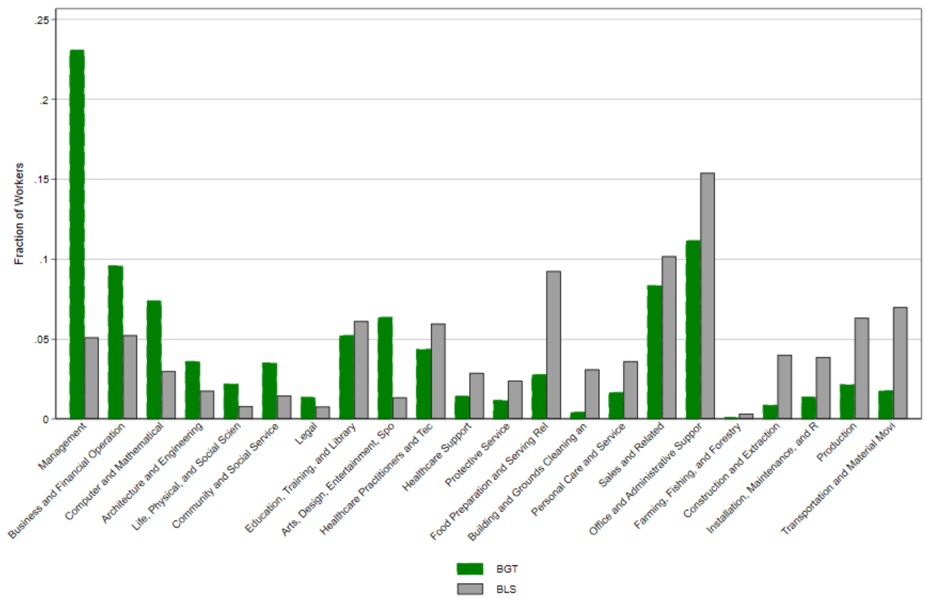

(d) State in Lightcast versus BLS

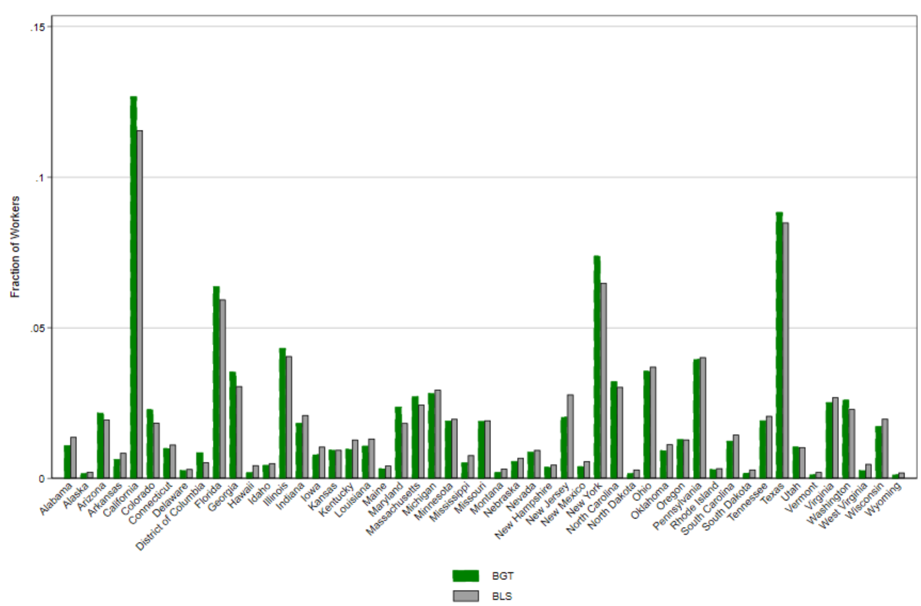

Notes: These figures shows the distribution of demographic characteristics in Lightcast (formerly Burning Glass Technologies, labeled BGT) versus the Bureau of Labor Statistic's Occupation Employment Statistics. Resumes from Lightcast include only those with non-missing current job-info.

Figure A.5: Form 5500 Firms versus BLS

(a) Industry by Number of Firms

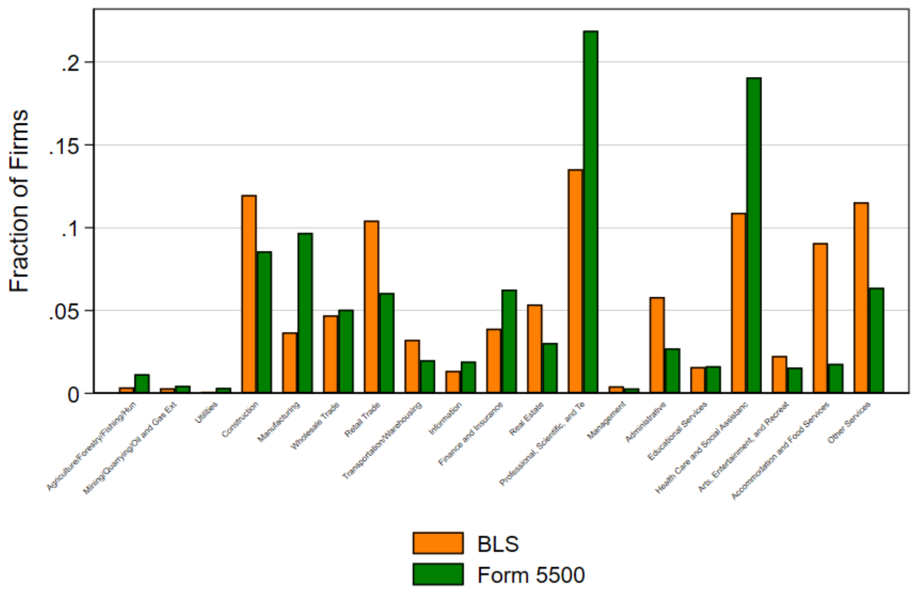

(b) Firm Size by Number of Firms

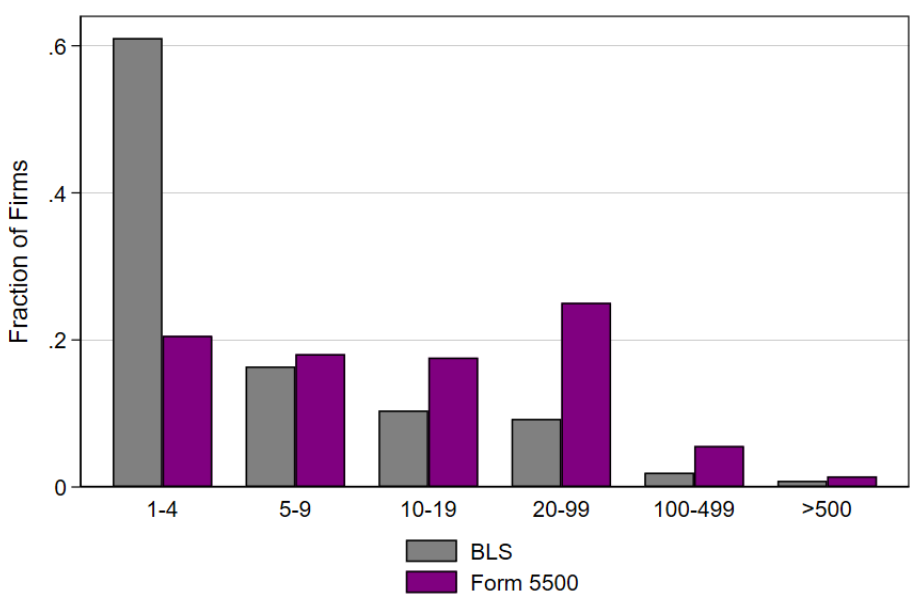

Notes: These figures shows the distribution of firms and employment by industry and by firm size. The Form 5500 sample is all firms that filed Form 5500 in 2019. The BLS sample is from the Bureau of Labor Statistics database.

Figure A.6: Example Survey Questions

(a) 401(k) with 3% match versus no 401(k)

Scenario 1

	Job 1	Job 2
Annual Earnings when working full time:	$51,500	$50,500
Retirement benefits:	None	Company sponsored 401(k) with matching of 100% up to 3% (for a total possible match of $1,515 per year)
Healthcare benefits:	No	No
Vacation:	20 days of paid vacation	20 days of paid vacation
Work flexibility:	No remote work option	No remote work option

Note that with Job 2, your pre-tax take home pay will be $1000 lower annually, or approximately $83 lower per month.

Select which you would be more likely to accept.

Job 1

Job 2

(b) 401(k) with 5% match versus no 401(k)

Scenario 4

	Job 1	Job 2
Annual Earnings when working full time:	$51,000	$52,500
Retirement benefits:	Company sponsored 401(k) with matching of 100% up to 5% (for a total possible match of $2,550 per year)	None
Healthcare benefits:	No	No
Vacation:	20 days of paid vacation	20 days of paid vacation
Work flexibility:	No remote work option	No remote work option

Note that with Job 1, your pre-tax take home pay will be $1,500 lower annually, or approximately $125 lower per month.

Select which you would be more likely to accept.

Job 1

Job 2

Figure A.6: Example Survey Questions (continued)

(c) 2 Days of Remote Work per Week versus No Remote Work

Scenario 2

	Job 1	Job 2
Annual Earnings when working full time:	$50,000	$51,500
Retirement benefits	Company sponsored 401(k)	Company sponsored 401(k)
Healthcare benefits:	Yes	Yes
Vacation:	20 days of paid vacation	20 days of paid vacation
Work flexibility:	2 days of remote work per week	No remote work option

Note that with Job 1, your pre-tax take home pay will be $1,500 lower annually, or approximately $125 lower per month.

Select which you would be more likely to accept.

Job 1

Job 2

Figure A.7: Willingness to Pay for Retirement Benefits and Remote Work Option

(a) 401(k) offers 3% match versus no 401(k)

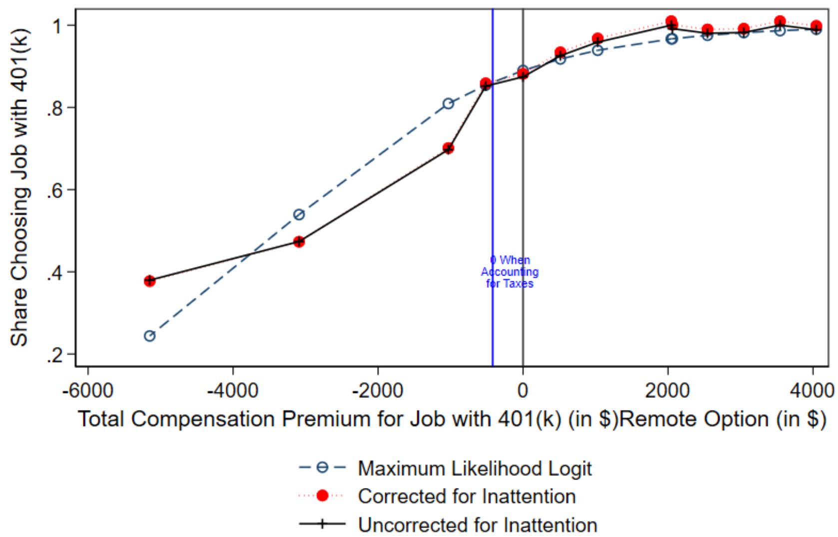

(b) 401(k) offers 5% match versus no 401(k)

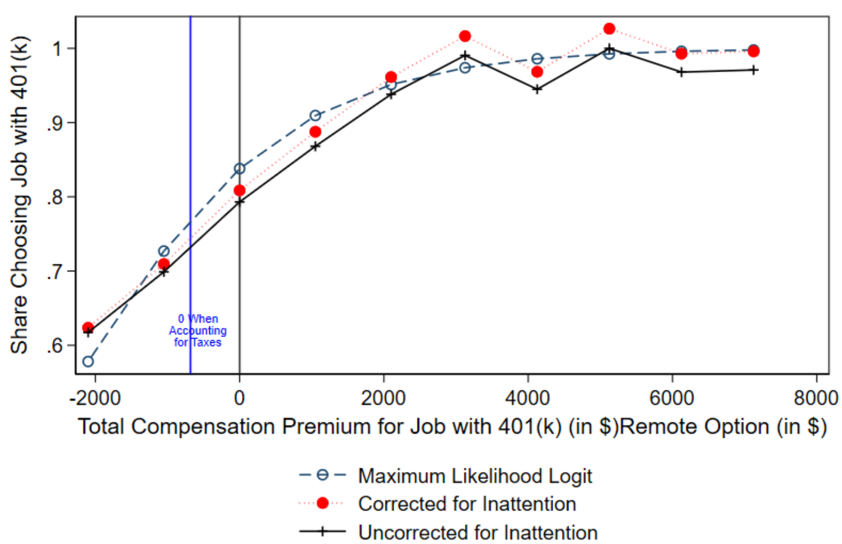

(c) 2 Days of Remote Work per Week versus no Remote Work

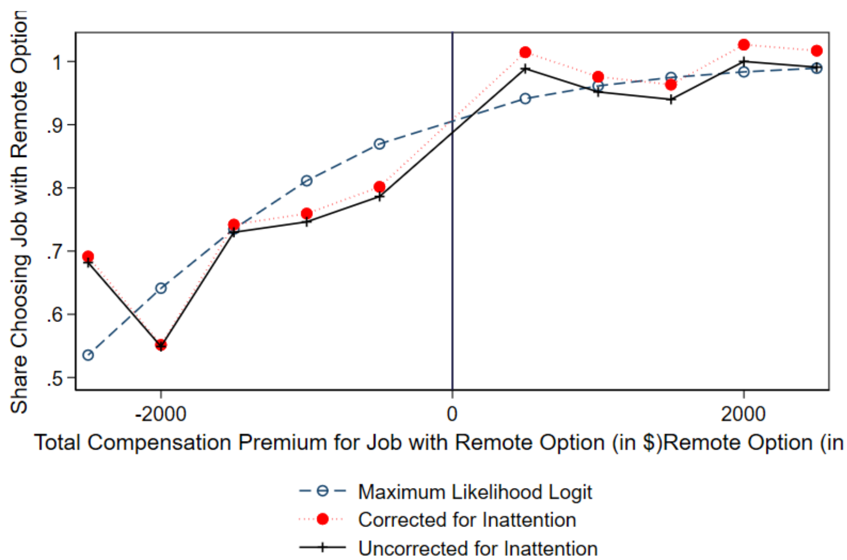

Notes: These plots show the fraction of participants who chose the job with the better benefits plotted against the difference in total compensation. The total compensation gap is the total compensation for the job with the better benefit minus the total compensation for the job with the worse benefit. Based on a survey of 1,600 participants.

Figure A.7: Reasons for Choosing Selected Job - Extensive Margin of Retirement Benefits

(a) Chose job with 401(k) versus job with no 401(k) when the 401(k) job pays less

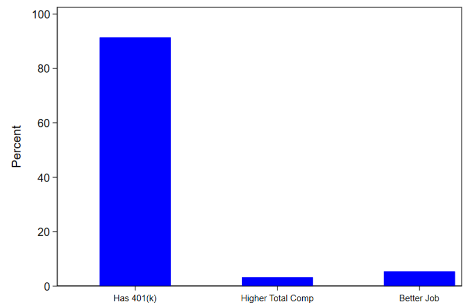

Notes: These plots show the distribution of reasons for choosing a given job, conditional on that job paying less than the alternative. Based on a subsample of 300 participants.

Figure A.8: Histogram of Wage Weights

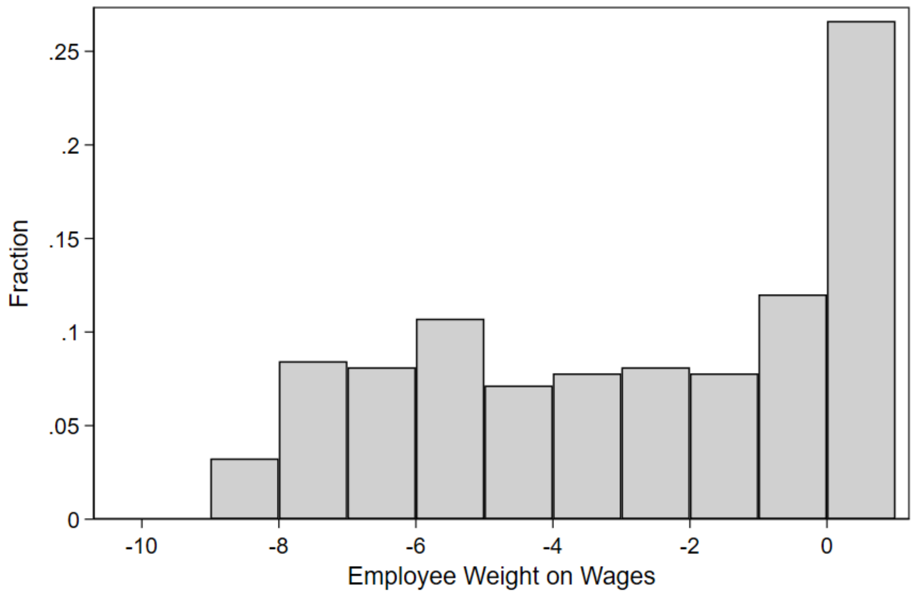

Notes: This figure shows a histogram of the values of γ_l, the weight on wages, estimated in the on-the-job random search model.

A.2 Appendix Tables

Table A.1: Example Non-discrimination Tests

Passing Firm				
Employee	Salary	Employee Deferral	Employer Contribution (100% of the first 3%)	Actual Contribution Percentage
Person 1	$150,000	$15,000	$4,500	3%
Person 2	$30,000	$0	$0	0%
Person 3	$30,000	$1,500	$900	3%
Person 4	$30,000	$1,200	900	3%
Mean ACP of HCEs				3%
Mean ACP of NHCEs				2%
Failing Firm				
Employee	Salary	Employee Deferral	Employer Contribution (100% of the first 3%)	Actual Contribution Percentage
Person 1	$150,000	$15,000	$4,500	3%
Person 2	$30,000	$0	$0	0%
Person 3	$30,000	$400	$400	1.3%
Person 4	$30,000	$200	200	.67%
Mean ACP of HCEs				3%
Mean ACP of NHCEs				.67%

Notes: This table shows example calculations for a non-discrimination test.

Table A.2: OES Average Wage by CBSA and 5-digit SOC

	Log OES Salary			Log OES Hourly Wage		
	(1) Median	(2) Mean	(3) Quantile within occupation	(4) Median	(5) Mean	(6) Quantile within occupation
Log Lightcast Salary	1.034***	1.110***	0.912***			
	(0.00600)	(0.00573)	(0.00435)			
Log Lightcast Hourly Wage				1.034***	1.110***	0.912***
				(0.00598)	(0.00572)	(0.00435)
Observations	85840	85949	85949	85696	85805	85805
Adjusted R^2	0.812	0.868	0.789	0.812	0.868	0.789

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Means and medians are calculated within CBSAs and 5-digit SOC codes. Quantiles are within 5-digit SOC codes across CBSAs. Regressions are weighted by employment share in the OES. Lightcast data is from all postings from 2010-2019 with wage and/or salary data available. When a range is given, we use the midpoint of the range.

Table A.3: Summary Statistics of Lightcast Postings Data with Wages

	(1)
Average Annual Salary	50247.39
	(29205.36)
Average Hourly Wage	24.17
	(14.16)
# of Positings per Firm	645.90
	(2985.64)
# of Firms	1,246,673
# of 5-digit SOC codes	437
# of CBSAs	929
# of Unique Jobs	8,014,427

Notes: This table shows summary statistics of the Lightcast postings data with available wage information from 2010-2019. Standard deviations, where applicable, are in parentheses. Shows only jobs with non-missing, wage, industry, occupation, and location information.

Table A.4: Resume Summary Statistics

	Mean	P10	P25	P50	P75	P90
Job Length (excluding current job, in months)	46	5	12	26	57	115
Job Length (including current job, in months)	36	4	11	23	47	82
# of Jobs per Resume	2.57	1	1	1	4	6
# of Years on Resume	15	4	8	14	21	30
# of Transitions per Resume	1.57	0	0	0	3	5
# of Resumes with Non-Missing Current Job	83,811,149					
# of Transitions:						
All	106,028,627					
Not within-firm	65,583,761					
Within Occupation, Industry and CBSA	860,230					

Notes: This table shows summary statistics from the Lightcast Resume Data. Only resumes with non-missing current job info that had been updated after 2019 are included.

Table A.5: Summary Statistics of DC Plans in Form 5500, 2010-2019

	Mean	Median	Standard Deviation
Total # of Active Plan Participants	120.39	6.00	2741.93
Total Current Employees	88.50	12.00	1882.92
Total Plan Assets (Millions of $s)	8.45	0.56	213.31
Ratio of Employer Contribution to Total Contributions	0.31	0.30	0.33
Employer Contributions per Participant ($)	3122.69	1399.86	7775.97
Implied NDT Failure (ever in sample)	0.06	0.00	0.23
Corrective Distributions per Person ($, Conditional on Paying Some)	111.34	28.67	1556.83
Has a Health Plan (%)	14.68	0.00	35.39
Health Spend ($) per Person (Conditional on Having a Health Plan, Large Plans Only)	4805.75	4286.99	22334.05
# of Unique Years per Firm	6.95	8.00	2.97
# of Unique Firms	1174495		
% of Firms that are Small Plans	80		

Notes: This table show summary statistics for all defined contribution plans in Form-5500 from 2010-2019. Small plans are those with less than 100 participants.

Table A.9: Transitions in Resume Waterfall

All	106,028,627
Non-missing industry, occupation, location info	88,162,464
Not within firm	65,583,761
Within Occupation	12,720,065
Within Industry	3.024,264
Within CBSA	860,230

Notes: This table shows the number of transitions in the Lightcast resume data that correspond to each criteria. These numbers are prior to matching with the other data sources.

Table A.6: Willingness to Pay for Retirement Benefits: Survey Evidence, Other Conditions

	401(k) with a 3% match versus no 401(k)	401(k) with a 5% match versus no 401(k)	2 days of remote work per week versus no remote work option)
Fraction that chose job with benefit	0.851 (1628)	0.885 (1004)	0.842 (1004)
Conditional on:			
Job with benefit having a lower wage	0.767 (1003)	0.796 (499)	0.740 (535)
Job with benefit having a higher wage	0.986 (625)	0.974 (505)	0.957 (469)
Job with benefit having lower total comp	0.635 (513)	0.659 (182)	0.7405 (535)
Job with benefit having higher total comp	0.965 (918)	0.960 (703)	0.957 (469)
Job with benefit having lower total comp, net of taxes	0.635 (513)	0.659 (182)	0.740 (535)
Conditional on choosing job with lower total compensation and with better benefit:			
WTP	3245 (326)	3950 (120)	1400 (396))
WTP as a percent of wages	6.3 (326)	7.5 (120)	2.7 (396))
WTP in total comp	1797 (326)	1522 (120)	1400 (396)
WTP as a percent of total comp	3.5 (326)	3.9 (120)	2.7 (396)
Tax saving	713 (326)	869 (120)	308 (396)
WTP in total comp, net of taxes	1083 (326)	653 (120)	1092 (396)
Cost of retirement plan to employer (excluding fixed/set-up costs)	1447 (326)	2427 (120)	0 (396)

Notes: This table shows summary statistics for the survey conditions that test willingness to pay simultaneously for the intensive and extensive margin of retirement benefits and for remote work capability. Numbers in parentheses show the number of participants who answered for the relevant condition. Based on a survey of 1,600 participants.

Table A.10: Effect of Changing Wages or Recruiting Success - Equivalent Effects to a 1pp Increase in Employer Contribution Rate

	p50	p25	p75
Equivalent Effect on Increase in Hiring			
% Increase in Wage Required	2.7		
% Change in number of new hires	0.41	0.19	0.81
Net cost per one new hire	1480.99	1054.32	2128.59
Net cost per employee	0.17	0.00	2.69
Equivalent Cost per New Hire			
% Increase in Wage Required	14		
% Change in number of new hires	1.98	0.88	3.76
Net cost per one new hire	8131.94	5789.15	11687.86
Net cost per employee	0.92	0.00	14.71

Notes: This table shows the effect of changing wages to get a) the same effect on recruiting success a 1 percentage point increase in employer contribution rate b) the same cost as a 1 percentage point increase in employer contribution rate.

Table A.7: Maximum Likelihood Estimates of Willingness to Pay in Survey, Other Conditions

Treatment	Mean	SD	P25	P50	P75
Willingness to Pay for the Intensive and Extensive Margin of Employer Contributions					
401(k) with 3% match versus no 401(k)	3589.41	3130.58	1692.29	3589.41	5486.52
	(246.83)	(192.70)	(169.17)	(246.83)	(347.14)
401(k) with 5% match versus no 401(k)	2492.70	2598.19	918.211	2492.70	4067.20
	(283.70)	(282.87)	(197.19)	(283.70)	(425.28)
Willingness to Pay for Remote Work Capability					
2 Days of Remote Work per Week versus no Remote Work Option	2935.24	2022.59	1709.56	2935.24	4160.92
	(208.16)	(199.08)	(137.10)	(208.16)	(311.41)

Notes: This table shows the distribution of the willingness to pay estimates from the survey. Estimates are from an inattention-corrected maximum likelihood logit model using data from the experiment. Bootstrapped standard errors based on 1000 samples are in parentheses.

Table A.8: Evidence for Heuristics in Experimental Choices

	Fraction
All experiments:	
Always chooses job with benefit	0.468
Always chooses job without benefit	0.071
Always chooses job with higher salary	0.203
Always chooses job with higher total comp	0.257
Retirement experiments only:	
Always chooses job with benefit	0.523
Always chooses job without benefit	0.075
Always chooses job with higher salary	0.241
Always chooses job with higher total comp	0.326
Observations	1629

Notes: This table shows the fraction of participants that followed the listed heuristic when completing the survey. The top panel includes the condition which tested for willingness to pay for remote work. The bottom panel shows only the conditions that tested for retirement benefits.

Appendix B

Appendix for Chapter 2

B.1 Appendix Figures

Figure B.1: Retirement Wealth Cutoffs

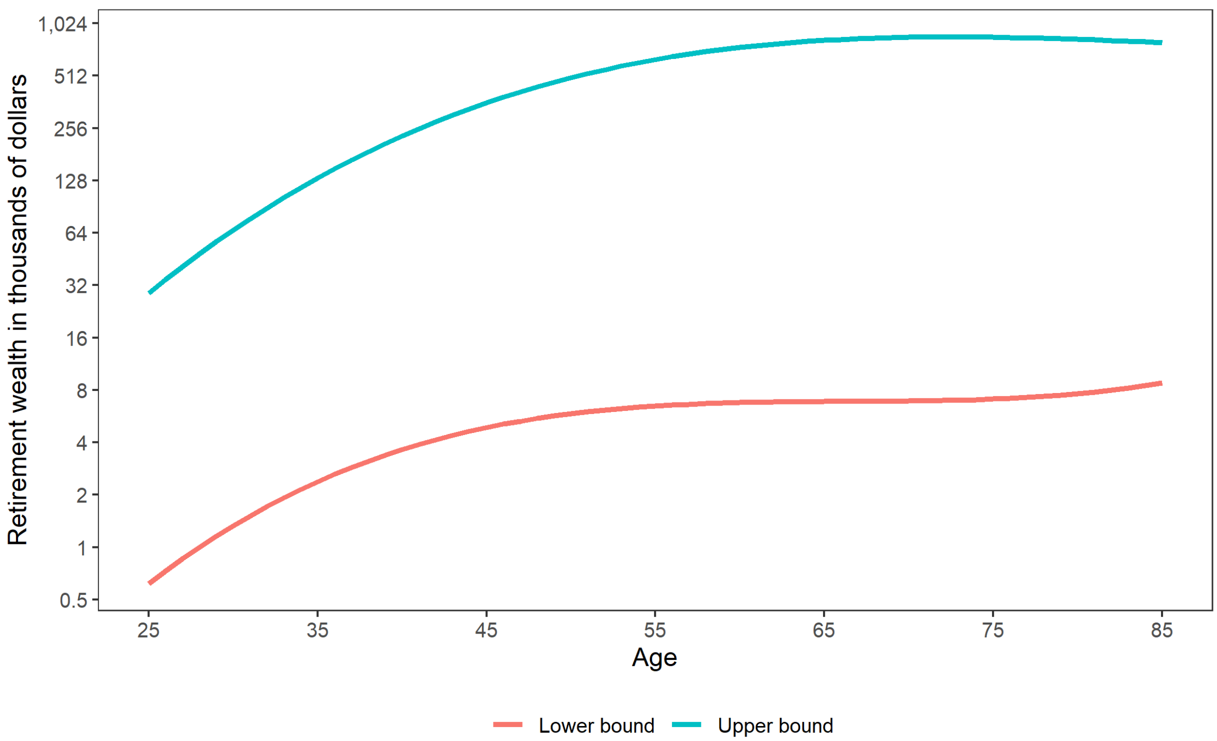

Notes: This figure shows the cutoffs on retirement wealth that are used to determine our retirement investor (RI) sample, described in Section 2.2.2. The cutoffs are determined by running quantile regressions of log of individual's retirement wealth on a third order polynomial in age in the 2016 Survey of Consumer Finance. We then drop individuals with retirement wealth below the estimated 10th percentile or above the 90th by age.

Figure B.2: Individual Labor Income Distribution in Firm Data and the SCF in 2016

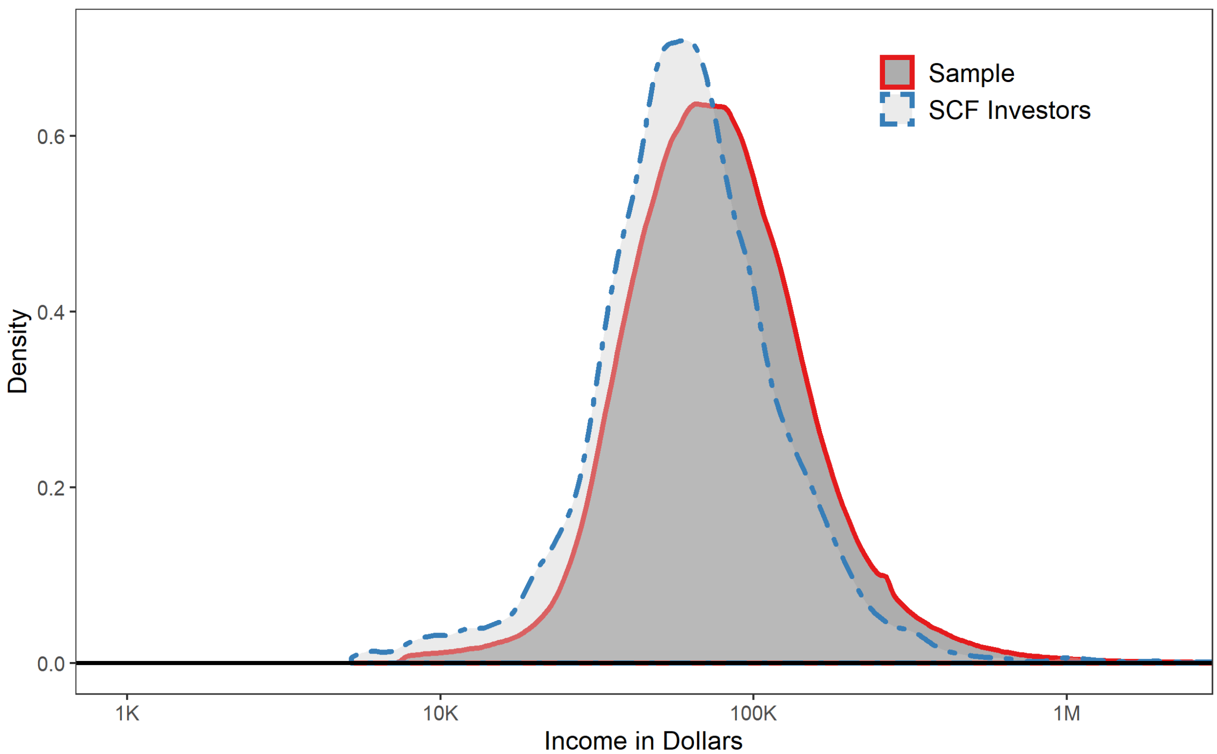

Notes: This figure plots the distribution of labor income in the sample of retirement investors (RIs) versus the distribution of labor income for RIs in the SCF in 2016.

Figure B.3: Equity Share of Retirement Wealth

(a) All ages

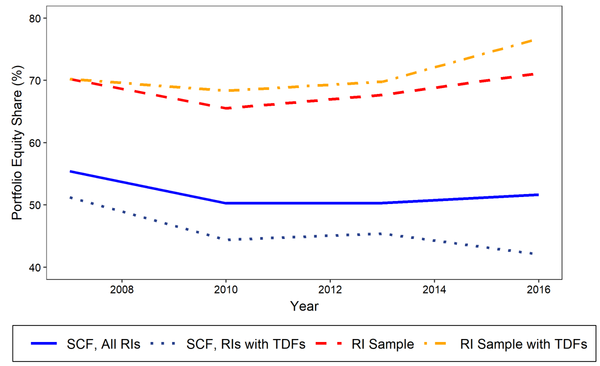

(b) Age 25-34

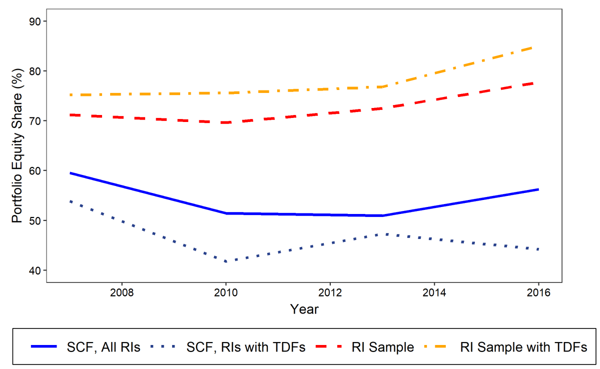

(c) Age 55-65

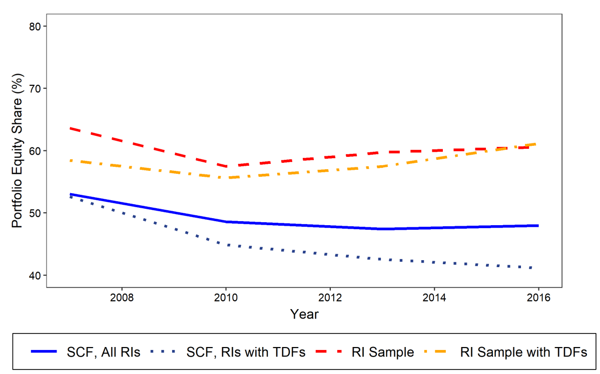

Notes: These figures show the portfolio equity share of retirement wealth over time. The SCF data is every three years, in 2007, 2010, 2013, and 2016. We show the same years in our sample. We also show the equity share for all RIs and for RIs who hold some assets in a TDF separately. Panel a shows all RIs, aged 25-65. Panel b shows RIs aged 25-34. Panel c shows RIs aged 55-65.

Figure B.4: Equity Share of Investable Wealth

(a) All ages

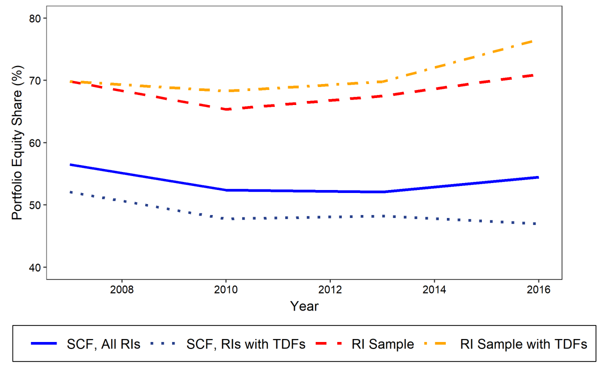

(b) Age 25-34

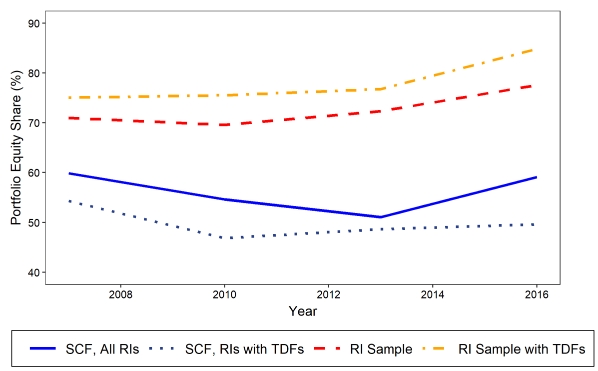

(c) Age 55-65

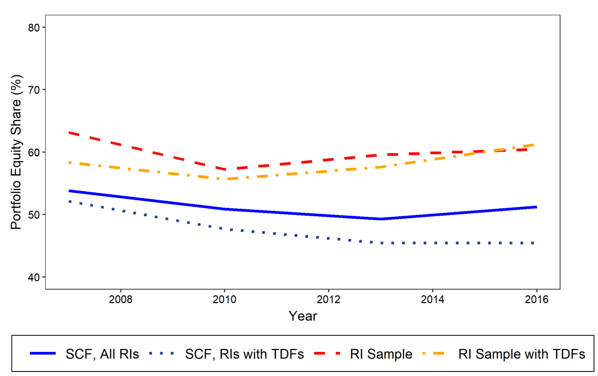

Notes: These figures show the portfolio equity share of investable wealth over time. The SCF data is every three years, in 2007, 2010, 2013, and 2016. We show the same years in our sample. We also show the equity share for all RIs and for RIs who hold some assets in a TDF separately. Panel a shows all RIs, aged 25-65. Panel b shows RIs aged 25-34. Panel c shows RIs aged 55-65.

Figure B.5: Equity Share Among Equity Owners (SCF)

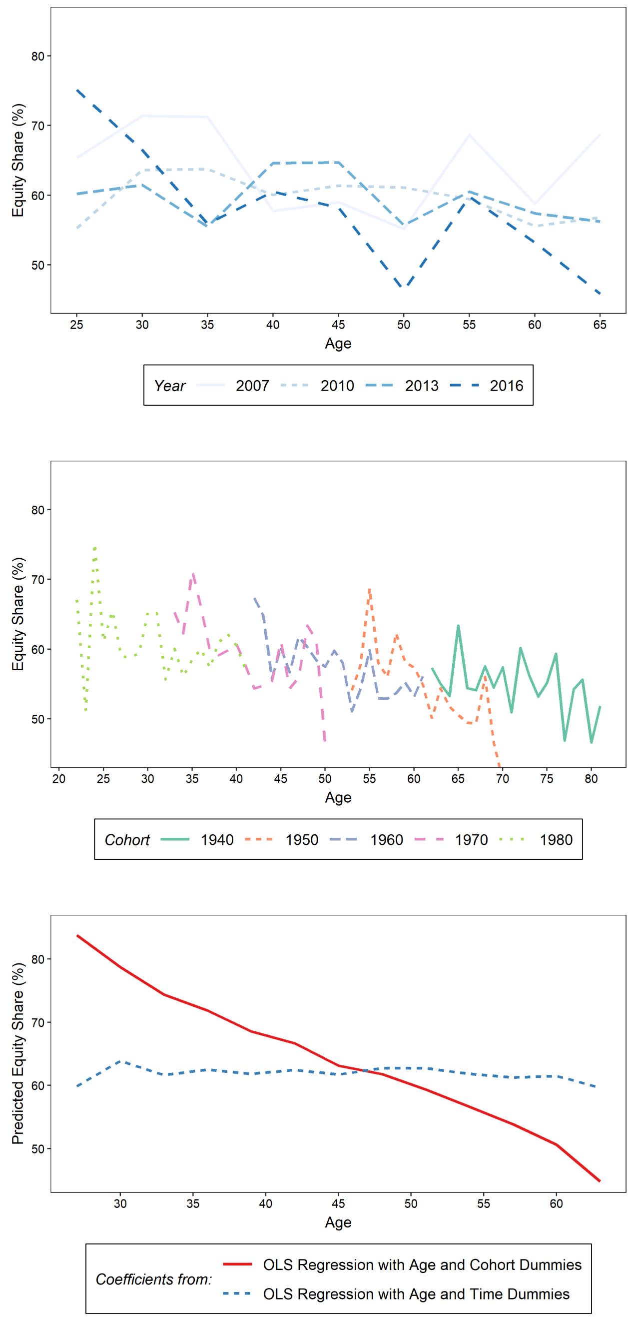

Notes: This figure replicates the results shown in Figure 9 of Ameriks and Zeldes (2004) using the SCF from 2007, 2010, 2013, and 2016. The top figure shows the observed equity share by age in each year. The middle figure shows the observed equity share by age in each cohort in our sample. A cohort is defined as having been born in the ten-year period beginning with the year indicated. The bottom figure shows the predicted values from a regression of equity share on indicator variables for age and either cohort or time. We obtain the predicted values by adding the median cohort or year coefficient, respectively, to each age coefficient. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. The sample is SCF retirement investors (RI) who own at least some equity.

Figure B.6: Portfolio Beta by Birth Cohort

(a) Portfolio Beta by Birth Cohort and Year

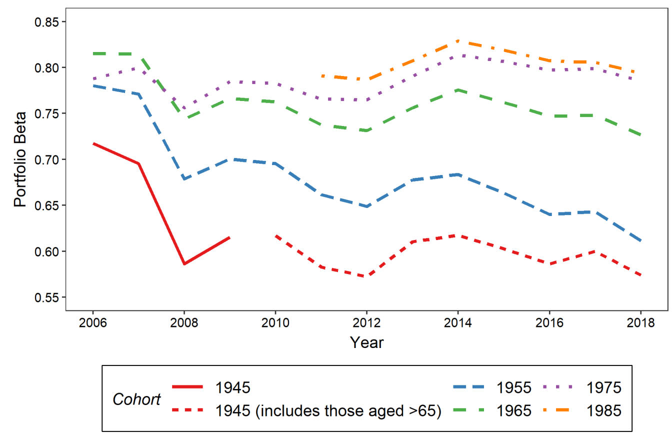

(b) Portfolio Beta by Birth Cohort and Age

Notes: These figures show the portfolios betas averaged by birth year cohorts. The top panel shows the averages by year over our sample period. We include only years during which each member of the cohort is aged 25-65, unless otherwise indicated. The bottom panel shows the averages by age, where age is the median age of the cohort. Portfolio betas are CAPM market betas calculated from all available return data from 2006-2018. A cohort is defined as having been born in the three-year period centered around the year indicated. The sample is our full set of retirement investors (RI).

Figure B.7: Incidence of Maxing Out on Contribution Limits, by Age Group

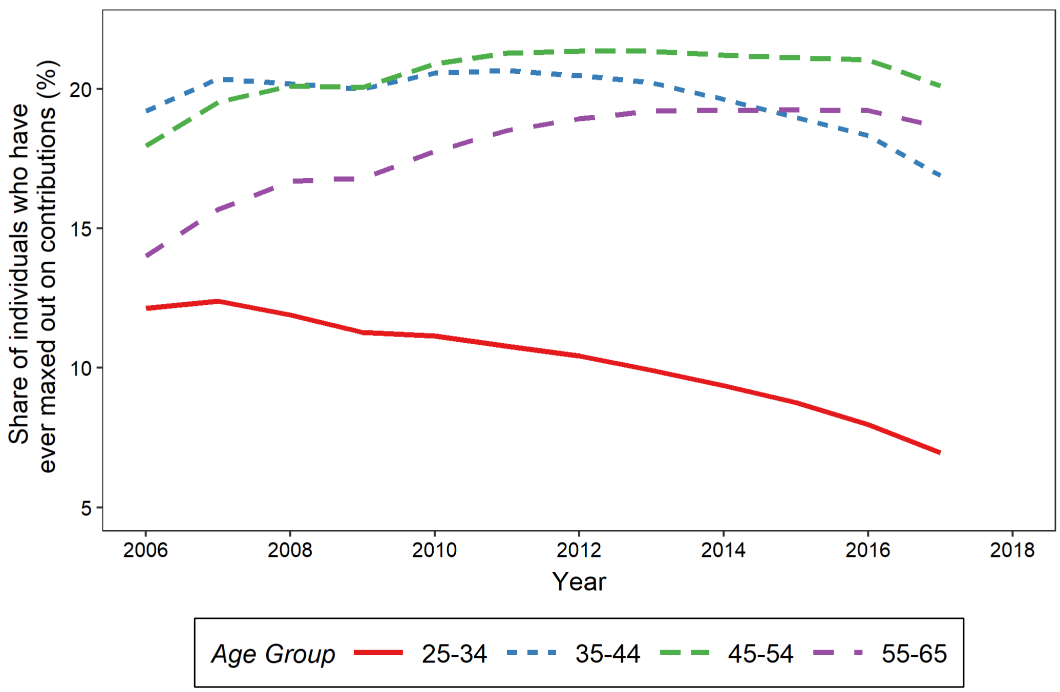

Notes: This figure shows the percentage of households that have ever hit their contribution limit in a given year, split by age groups. Maxing out is defined as when an individual exceeds the dollar amount that is allowed for 401(k) contributions in a year, set by the IRS. The sample is our full set of retirement investors (RI) which have income data available.

B.2 Appendix Tables

Table B.1: Characteristics of Sample of Retirement Investors

	Retirement Investors			
	Summary Statistics			
	Mean	Median	SD	Percentage of RI Sample with Observed Data
Age (Years)	45.38	46	11.01	100.0%
Share Female (%)	45.0	0	49.7	93.4%
Share Married (%)	73.8	100	44.0	88.6%
Labor Income ($)	94,044	69,506	214,798	44.5%
Employment Tenure (Years)	10.77	8.09	9.12	58.2%
Investable Wealth ($)	100,365	36,114	318,490	100%
Retirement Wealth ($)	81,349	32,922	131,540	100%
Retirement Share of Wealth (%)	96.1	100	14.5	100%
Portfolio Beta	0.75	0.84	0.34	85.7%
TDF Share of Invest. Wealth (%)	38.4	15.3	42.9	99.3%
Reported Contribution Rate (%)	8.0	6.0	7.2	49.1%
Realized Contribution Rate (%)	6.3	5.3	6.4	44.9%

Notes: This table presents summary statistics on demographics, wealth, and portfolio allocations for our Retirement Investor (RI) sample from 2006-2018. Detailed definitions for retirement wealth and investable wealth are provided in Table I. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Market betas are obtained by regressing monthly fund or security excess returns on the value-weighted CRSP market excess return over the period 2007–2017 with at least 24 observations. Income is the labor income of the head of household in 2015. The sample is not representative of the assets under management of our financial service firm, since by design we drop the highest and lowest income groups.

Table B.2: Characteristics of Sample of Retirement Investors - Married Subsample

Retirement Investors				
	Summary Statistics			
	Mean	Median	SD	Percentage of Married RI Sample with Observed Data
Age (Years)	49.48	51	10.30	100%
Share Female (%)	48.2	0	50	94.9%
Share Married (%)	100	100	0	100%
Labor Income ($)	113,105	83,213	202,430	34.5%
Employment Tenure (Years)	12.56	10.18	9.95	49.2%
Investable Wealth (Individual, $)	188,503	75,410	492,778	100%
Investable Wealth (Household, $)	285,085	126,205	669,449	100%
Retirement Wealth ($)	143,681	66,425	200,599	100%
Retirement Share of Wealth (%)	96.7	100	12.9	100%
Portfolio Beta	0.73	0.81	0.33	88.0%
TDF Share of Invest. Wealth (%)	36.4	11.9	42.1	99.7%
Reported Contribution Rate (%)	9.8	8	8.7	43.5%
Realized Contribution Rate (%)	7.7	6	6.4	37.7%
Retirement Investors - Survey of Consumer Finance				
	Summary Statistics			
	Mean	Median	SD	Number of Observations
Age	46.87	47.00	10.48	2556
Female (%)	46.20	0.00	49.99	2556
Married (%)	100.00	100.00	0.00	2556
Labor Income (Individual, $)	68,380	51,000	1,203,245	2556
Labor Income (Household, $)	101,349	77,000	1,445,913	2556
Investable Wealth (Household, $)	273,282	72,000	17,019,097	2556
Retirement Wealth (Household, $)	225,166	94,000	718,051	2556
Retirement Wealth (Individual, $)	100,805	45,000	154,875	2556
Retirement Share of Investable Wealth (Individual, %)	58.35	56.60	37.23	2556
Retirement Share of Investable Wealth (Household, %)	86.55	100.00	34.04	2556

Notes: This table presents summary statistics on demographics, wealth, and portfolio allocations for a subsample or our Retirement Investors (RI) sample who are married and for whom we observe both partners in our data set. We use 2016 data to compare with the 2016 Survey of Consumer Finance (SCF), in which we include only married investors here. Detailed definitions for retirement wealth and investable wealth are provided in Table I. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Market betas are obtained by regressing monthly fund or security excess returns on the value-weighted CRSP market excess return over the period 2007–2017 with at least 24 observations. Income is the labor income of the respondent in 2015. The sample is not representative of the assets under management of our financial service firm, since by design we drop the highest and lowest income groups.

Table B.3: Characteristics of Sample of Retirement Investors - Single Subsample

Retirement Investors				
	Summary Statistics			
	Mean	Median	SD	Percentage of Single RI Sample with Observed Data
Age (Years)	42.43	41	11.56	100%
Share Female (%)	49.6	0	50.0	98.1%
Share Married (%)	0	0	0	100%
Labor Income ($)	83,344	63,346	129,726	42.2%
Employment Tenure (Years)	9.58	6.74	8.56	61.9%
Investable Wealth ($)	83,535	25,156	284,098	100%
Retirement Wealth ($)	69,227	23,547	122,048	100%
Retirement Share of Wealth (%)	93.6	100	18.0	100%
Portfolio Beta	0.76	0.85	0.32	89.7%
TDF Share of Invest. Wealth (%)	53.2	60.3	45.2	99.7%
Reported Contribution Rate (%)	7.3	6.0	6.5	56.0%
Realized Contribution Rate (%)	5.9	5.0	4.8	50.0%
Retirement Investors - Survey of Consumer Finance				
	Summary Statistics			
	Mean	Median	SD	Number of Observations
Age (Respondent)	46.43	47.00	11.12	574
Share Female (%)	61.9	100	48.9	574
Share Married (%)	0	0	0	574
Labor Income (Respondent, $)	59,725	50,000	710,967	574
Investable Wealth (Household, $)	133,613	42,900	6,409,202	574
Retirement Wealth (Respondent, $)	82,806	35,000	124,628	574
Retirement Share of Wealth (%)	89.7	100	26.5	574

Notes: This table presents summary statistics on demographics, wealth, and portfolio allocations for a subsample or our Retirement Investors (RI) sample who are not married and for whom we observe only one member in the household. We use 2016 data to compare with the 2016 Survey of Consumer Finance (SCF), in which we include only unmarried investors here. Detailed definitions for retirement wealth and investable wealth are provided in Table I. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Market betas are obtained by regressing monthly fund or security excess returns on the value-weighted CRSP market excess return over the period 2007–2017 with at least 24 observations. Income is the labor income of the respondent in 2015. The sample is not representative of the assets under management of our financial service firm, since by design we drop the highest and lowest income groups.

Table B.4: Average Share of Equity in Portfolios Among Retirement Investors - Full Sample versus 2016 SCF

	All Retirement Investors		Retirement Investors with Hybrid Fund (e.g. TDF) in Retirement Account	
Panel A: All Investable Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	68.6	54.5	73.0	46.9
Age 25-34	73.9	59.1	80.7	49.6
Age 35-44	73.2	55.9	77.7	47.9
Age 45-54	68.6	53.8	70.5	45.5
Age 55-65	59.6	51.2	59.2	45.4
Panel B: Retirement Wealth	Main Sample (Individuals)	SCF (Individuals)	Main Sample (Individuals)	SCF (Individuals)
All RIs	68.9	51.7	73.1	42.1
Age 25-34	74.1	56.2	80.8	44.2
Age 35-44	73.4	54.1	77.9	43.5
Age 45-54	68.9	50.5	70.6	40.2
Age 55-65	59.8	48.0	59.1	41.2
Panel C: Non-Retirement Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	54.1	73.4	54.3	73.2
Age 25-34	53.6	87.5	53.9	86.9
Age 35-44	55.5	68.9	56.5	68.3
Age 45-54	53.9	74.5	54.5	73.6
Age 55-65	50.1	69.6	51.5	69.6

Notes: This table presents the share of equity in the portfolio allocations for various samples of our Retirement Investors (RI) sample and the comparable RI sample of the 2016 Survey of Consumer Finance (SCF). Panel A shows equity shares of total investable wealth at the individual level in our sample and the household level in the SCF. Panel B shows equity shares of retirement wealth, at the individual level in both datasets. Panel C shows equity shares of non-retirement wealth at the individual level in our sample and the household level in the SCF. The figures in Panel C are conditional on owning some non-retirement wealth, which is approximately 43% of the SCF RI sample and 16% of our RI sample. The first two columns show the means for the full sample of RIs in each dataset. The last two columns show the means for the subsample of the RI sample that has some of their retirement assets in a target date fund (TDF). Investable wealth is defined as money market funds, non-money market funds, individual stocks and bonds, Retirement wealth is defined as any wealth in retirement saving accounts of all types (excluding defined benefit plans and Social Security). certificates of deposit, quasi-liquid retirement wealth, and other managed accounts. The equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets.

Table B.5: Average Share of Equity in Portfolios Among Retirement Investors - Married Subsample

	All Retirement Investors		Retirement Investors with Hybrid Fund (e.g. TDF) in Retirement Account	
Panel A: All Investable Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All Ris	69.0	55.2	73.9	47.3
Age 25-34	77.0	59.3	84.4	50.0
Age 35-44	76.1	56.0	81.7	48.0
Age 45-54	71.0	54.6	74.3	46.5
Age 55-65	61.5	52.8	61.9	45.7
Respondents		55.5	0.0	47.7
Partners		54.8	0.0	46.9
Panel B: Retirement Wealth	Main Sample (Individuals)	SCF (Individuals)	Main Sample (Individuals)	SCF (Individuals)
All Ris	69.4	52.1	74.1	41.9
Age 25-34	77.2	56.4	84.6	44.4
Age 35-44	76.4	54.0	82.0	42.8
Age 45-54	71.5	50.8	74.6	40.3
Age 55-65	61.7	49.2	61.9	40.9
Respondents		53.2	0.0	43.7
Partners		50.8	0.0	39.8
Panel C: Non-Retirement Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	52.0	73.3	54.0	72.1
Age 25-34	53.1	89.9	53.5	86.2
Age 35-44	55.3	68.5	57.3	66.5
Age 45-54	52.7	75.0	54.3	74.2
Age 55-65	49.9	69.5	51.5	68.8
Respondents		74.0	0.0	72.7
Partners		72.7	0.0	71.2

Notes: This table presents the share of equity in the portfolio allocations for various samples of our Retirement Investors (RI) sample in 2016 and the comparable RI sample of the 2016 Survey of Consumer Finance (SCF). From our sample, this table shows summary statistics for the subset of investors who are married and for whom we observe both partners in our data set. In the SCF, this table shows only summary statistics of married investors.

Table B.6: Average Share of Equity in Portfolios Among Retirement Investors - Single Subsample

	All Retirement Investors		Retirement Investors with Hybrid Fund (e.g. TDF) in Retirement Account	
Panel A: All Investable Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	71.6	51.9	78.3	45.6
Age 25-34	77.4	58.4	85.0	48.5
Age 35-44	74.9	55.4	82.2	47.8
Age 45-54	70.4	51.3	74.5	43.0
Age 55-65	59.5	45.6	60.5	44.5
Panel B: Retirement Wealth	Main Sample (Individuals)	SCF (Individuals)	Main Sample (Individuals)	SCF (Individuals)
All RIs	71.8	50.2	78.4	42.5
Age 25-34	77.6	55.7	85.2	43.5
Age 35-44	75.1	54.5	82.4	45.8
Age 45-54	70.5	49.7	74.6	40.0
Age 55-65	59.6	43.7	60.5	42.1
Panel C: Non-Retirement Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Household)
All RIs	51.2	73.6	53.5	78.4
Age 25-34	52.3	81.2	54.0	88.3
Age 35-44	53.8	71.9	56.1	81.4
Age 45-54	50.6	71.6	52.3	71.0
Age 55-65	48.0	70.6	49.4	74.8

Notes: This table presents the share of equity in the portfolio allocations for various samples of our Retirement Investors (RI) sample in 2016 and the comparable RI sample of the 2016 Survey of Consumer Finance (SCF). From our sample, this table shows summary statistics for the subset of investors who are single and for whom we observe only one member of the household. In the SCF, this table shows only summary statistics of non-married investors.

Table B.7: Average Residual Share of Equity in Portfolios Among Retirement Investors

	All Retirement Investors		Retirement Investors with Hybrid Fund (e.g. TDF) in Retirement Account	
Panel A: All Investable Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	0.0011	-0.0082	0.0497	-0.0861
Age 25-34	0.0050	0.0032	0.0790	-0.0916
Age 35-44	0.0027	-0.0057	0.0653	-0.0886
Age 45-54	0.0024	-0.0047	0.0393	-0.0867
Age 55-65	-0.0051	-0.0199	0.0051	-0.0781
Panel B: Retirement Wealth	Main Sample (Individuals)	SCF (Individuals)	Main Sample (Individuals)	SCF (Individuals)
All RIs	0.0011	-0.0110	0.0494	-0.1102
Age 25-34	0.0052	-0.0039	0.0796	-0.1242
Age 35-44	0.0027	-0.0007	0.0656	-0.1101
Age 45-54	0.0025	-0.0123	0.0386	-0.1156
Age 55-65	-0.0051	-0.0216	0.0034	-0.0911
Panel C: Non-Retirement Wealth	Main Sample (Individuals)	SCF (Households)	Main Sample (Individuals)	SCF (Households)
All RIs	0.0004	-0.0049	0.0195	-0.0075
Age 25-34	0.0032	0.0269	0.0141	0.0249
Age 35-44	0.0009	-0.0390	0.0219	-0.0449
Age 45-54	-0.0002	0.0200	0.0186	0.0141
Age 55-65	-0.0007	-0.0169	0.0221	-0.0170

Notes: This table presents the residuals from a regression of the equity share on gender, investable wealth, and birth year cohort. We use the share of equity in the portfolio allocations for various samples of our Retirement Investors (RI) sample in 2016 and the comparable RI sample of the 2016 Survey of Consumer Finance (SCF). Panel A shows residuals of equity shares of total investable wealth at the individual level in our sample and the household level in the SCF. Panel B shows residuals of equity shares of retirement wealth, at the individual level in both datasets. Panel C shows residuals of equity shares of non-retirement wealth at the individual level in our sample and the household level in the SCF. The figures in Panel C are conditional on owning some non-retirement wealth, which is approximately 43% of the SCF RI sample and 16% of our RI sample. The first two columns show the means for the full sample of RIs in each dataset. The last two columns show the means for the subsample of the RI sample that has some of their retirement assets in a target date fund (TDF). Investable wealth is defined as money market funds, non-money market funds, individual stocks and bonds, Retirement wealth is defined as any wealth in retirement saving accounts of all types (excluding defined benefit plans and Social Security), certificate of deposits, quasi-liquid retirement wealth, and other managed accounts. The equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets.

Table B.8: Cross-Sectional Regressions of Price Constant Equity Share, Full Sample and by Income Terciles

	Price constant portfolio equity share				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.7357	0.8061	0.7433	0.7941	0.8012
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0006)
Age 28-30	0.7317	0.7984	0.7296	0.7811	0.7894
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0004)
Age 31-33	0.7313	0.7894	0.7244	0.7722	0.7794
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Age 34-36	0.7327	0.7813	0.7229	0.7666	0.7716
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Age 37-39	0.7327	0.7725	0.7193	0.7607	0.7660
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Age 40-42	0.7281	0.7606	0.7108	0.7509	0.7584
	(0.0001)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Age 43-45	0.7203	0.7473	0.6992	0.7384	0.7489
	(0.0001)	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Age 46-48	0.7054	0.7278	0.6795	0.7179	0.7331
	(0.0001)	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Age 49-51	0.6844	0.7016	0.6549	0.6906	0.7091
	(0.0001)	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Age 52-54	0.6595	0.6726	0.6264	0.6596	0.6802
	(0.0001)	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Age 55-57	0.6297	0.6383	0.5919	0.6228	0.6458
	(0.0001)	(0.0002)	(0.0004)	(0.0003)	(0.0003)
Age 58-60	0.5988	0.6038	0.5585	0.5845	0.6087
	(0.0001)	(0.0002)	(0.0004)	(0.0004)	(0.0003)
Age 61-63	0.5690	0.5705	0.5240	0.5459	0.5731
	(0.0002)	(0.0003)	(0.0004)	(0.0004)	(0.0004)
Age 64-65	0.5477	0.5454	0.4956	0.5144	0.5444
	(0.0002)	(0.0003)	(0.0006)	(0.0005)	(0.0005)
Log income		0.0743			
		(0.0003)			
Person fixed effect?	N	N	N	N	N
% of RI Sample	80.0	34.0	13.3	14.2	13.8
R-squared	0.0386	0.0781	0.0544	0.0769	0.0629

Notes: This table presents regression coefficients of annual individual price-constant portfolio equity shares on a set of demographic controls. The price-constant portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets, ignoring any changes in the price of these assets. These hypothetical portfolio shares track the inflows and outflows into these assets and are insensitive to passive appreciation. The baseline specification in column (1) shows the coefficients for the regression of equity share on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.9: Cross-Sectional Regressions of Equity Share on Age Groups by Cohort and TDF Share

	Portfolio equity share							
	(1) 1943 Cohort	(2) 1953 Cohort	(3) 1963 Cohort	(4) 1973 Cohort	(5) 1983 Cohort	(6) Initial TDF Share 75-100 %	(7) Initial TDF Share 25-75 %	(8) Initial TDF Share 0-25 %
Age 25-27				0.7376	0.8110	0.6868	0.7321	0.7963
				(0.0005)	(0.0002)	(0.0006)	(0.0007)	(0.0002)
Age 28-30				0.7404	0.8234	0.6958	0.7347	0.7857
				(0.0003)	(0.0002)	(0.0004)	(0.0005)	(0.0002)
Age 31-33				0.7533	0.8401	0.7094	0.7326	0.7758
				(0.0002)	(0.0002)	(0.0003)	(0.0004)	(0.0002)
Age 34-36			0.7545	0.7766	0.8406	0.7244	0.7321	0.7669
			(0.0005)	(0.0002)	(0.0003)	(0.0003)	(0.0004)	(0.0002)
Age 37-39			0.7379	0.7890		0.7336	0.7284	0.7532
			(0.0003)	(0.0002)		(0.0002)	(0.0003)	(0.0002)
Age 40-42			0.7283	0.8047		0.7357	0.7204	0.7330
			(0.0002)	(0.0002)		(0.0002)	(0.0003)	(0.0002)
Age 43-45		0.7470	0.7359	0.8038		0.7320	0.7103	0.7089
		(0.0006)	(0.0002)	(0.0003)		(0.0002)	(0.0003)	(0.0003)
Age 46-48		0.7020	0.7379			0.7205	0.6947	0.6763
		(0.0003)	(0.0002)			(0.0002)	(0.0003)	(0.0003)
Age 49-51		0.6694	0.7365			0.7025	0.6766	0.6411
		(0.0003)	(0.0002)			(0.0002)	(0.0003)	(0.0003)
Age 52-54	0.6948	0.6572	0.7226			0.6800	0.6566	0.6070
	(0.0010)	(0.0002)	(0.0003)			(0.0002)	(0.0003)	(0.0003)
Age 55-57	0.6363	0.6408				0.6531	0.6319	0.5687
	(0.0005)	(0.0002)				(0.0002)	(0.0003)	(0.0003)
Age 58-60	0.5809	0.6226				0.6230	0.6056	0.5346
	(0.0004)	(0.0002)				(0.0002)	(0.0004)	(0.0003)
Age 61-63	0.5558	0.6005				0.5928	0.5773	0.5035
	(0.0004)	(0.0003)				(0.0002)	(0.0004)	(0.0004)
Age 65-65	0.5430	0.5733				0.5720	0.5564	0.4796
	(0.0004)	(0.0006)				(0.0003)	(0.0005)	(0.0005)
Log income	0.1017	0.0959	0.0724	0.0572	0.0526			
	(0.0010)	(0.0005)	(0.0005)	(0.0005)	(0.0005)			
Person fixed effect?	N	N	N	N	N	N	N	N
% of RI Sample	3.1	10.9	11.5	10.3	5.0	39.5	7.9	10.1
R-squared	0.0220	0.0226	0.0291	0.0154	0.0095	0.0228	0.0523	0.2024

Notes: This table presents regression coefficients of annual individual portfolio equity shares on a set of demographic controls. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. Columns (1)-(5) show the results including age-group controls and a control for log income, broken out by birth cohort groups. Log income is measured as the log deviation of the individual's income from the average income of the RI sample. A cohort is defined as having been born in the ten-year period beginning with the year indicated. Columns (6)-(8) show the results for different groups based on the initial share of their portfolio that is invested in target date funds (TDFs). The sample is our full RI sample from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.10: Regressions of Equity Share on Automated Investment Allocation: Long-run Effect, Treated in 2007 Only

	Portfolio equity share						
	(1) Full Sample	(2) Bottom Income Tercile	(3) Top Income Tercile	(4) Age Enrolled 25-34	(5) Age Enrolled 35-44	(6) Age Enrolled 45-54	(7) Age Enrolled 55-65
Year of x Treatment	0.0406	0.0515	0.0467	0.0235	-0.0003	-0.0267	0.0003
	(0.0031)	(0.0052)	(0.0065)	(0.0029)	(0.0106)	(0.0120)	(0.0253)
1 Year After x Treatment	0.0245	0.0377	-0.0137	-0.0019	-0.0330	-0.0864	-0.1182
	(0.0014)	(0.0018)	(0.0041)	(0.0015)	(0.0059)	(0.0104)	(0.0152)
2 Years After x Treatment	0.0929	0.1067	0.0554	0.0478	0.0514	0.0115	-0.0207
	(0.0012)	(0.0017)	(0.0037)	(0.0014)	(0.0046)	(0.0077)	(0.0176)
3 Years After x Treatment	0.0614	0.0626	0.0427	0.0071	0.0106	0.0041	-0.0087
	(0.0012)	(0.0016)	(0.0033)	(0.0015)	(0.0037)	(0.0070)	(0.0134)
4 Years After x Treatment	-0.0210	-0.0154	-0.0212	-0.0068	-0.0121	-0.0123	-0.0205
	(0.0012)	(0.0020)	(0.0022)	(0.0017)	(0.0017)	(0.0025)	(0.0047)
5 Years After x Treatment	0.0107	0.0015	0.0064	0.0369	0.0146	-0.0168	-0.0175
	(0.0012)	(0.0019)	(0.0027)	(0.0017)	(0.0019)	(0.0026)	(0.0052)
1 Year After	0.0090	0.0210	0.0121	0.0135	0.0130	0.0092	0.0080
	(0.0009)	(0.0018)	(0.0013)	(0.0012)	(0.0014)	(0.0016)	(0.0026)
2 Years After	-0.0192	0.0024	-0.0207	0.0067	-0.0210	-0.0451	-0.0575
	(0.0010)	(0.0019)	(0.0014)	(0.0012)	(0.0015)	(0.0017)	(0.0029)
3 Years After	-0.0272	0.0048	-0.0388	0.0141	-0.0285	-0.0680	-0.0886
	(0.0010)	(0.0019)	(0.0015)	(0.0013)	(0.0016)	(0.0018)	(0.0030)
4 Years After	-0.0221	0.0077	-0.0340	0.0242	-0.0211	-0.0634	-0.0820
	(0.0010)	(0.0019)	(0.0015)	(0.0013)	(0.0016)	(0.0018)	(0.0031)
5 Years After	-0.0382	-0.0092	-0.0501	0.0132	-0.0375	-0.0855	-0.1016
	(0.0010)	(0.0019)	(0.0015)	(0.0013)	(0.0016)	(0.0018)	(0.0032)
Log income	0.0559						
	(0.0013)						
Constant	0.7279	0.6751	0.7473	0.7255	0.7432	0.7059	0.6374
	0.0010	0.0020	0.0014	0.0013	0.0015	0.0017	0.0028
% of RI Sample	1.5	0.5	0.5	0.8	0.6	0.4	0.1
% of Sample Enrolled 2005-2008	21.6	7.4	7.0	11.3	8.0	5.7	1.9
R-squared	0.0991	0.1753	0.0715	0.1583	0.1023	0.0925	0.1155

Notes: This table presents regression coefficients of annual household portfolio equity shares on being treated with the Pension Protection Act (PPA) of 2006. "Year of" means the year the individual enrolled in their retirement plan and "x years after" is x years after they enrolled in the plan. Each column includes year dummies for each year after enrollment, and interactions of these dummies with the treatment dummy. The treatment dummy is equal to one if the individual enrolled in 2007 to a plan that switched to having a target date fund as the default following the PPA and zero if they enrolled in 2005 or 2006. The full sample is those enrolled from 2005-2008 who otherwise meet the RI sample criteria. The bottom (top) income tercile includes those whose initial income is in the lowest (highest) tercile. Columns (4)-(7) break out the result for all individuals enrolled from 2005-2008 by age at enrollment. The portfolio equity share is defined as the sum of equity securities, pure equity funds, and the equity portion of hybrid funds, relative to total portfolios assets. Log income, when included, is the log deviation of the individual's current income from the average income of the RI sample. Standard errors, in parentheses, are clustered at the household level.

Table B.11: Cross-Sectional Regressions of Realized Contribution Rate on Age Groups by Cohort and TDF Share

	Realized contribution rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1943 Cohort	1953 Cohort	1963 Cohort	1973 Cohort	1983 Cohort	Initial TDF Share 75-100 %	Initial TDF Share 25-75 %	Initial TDF Share 0-25 %
Age 25-27				0.0496	0.0565	0.0492	0.0474	0.0418
				(0.0001)	(0.0000)	(0.0001)	(0.0001)	(0.0000)
Age 28-30				0.0517	0.0598	0.0533	0.0535	0.0480
				(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0000)
Age 31-33				0.0544	0.0625	0.0561	0.0566	0.0513
				(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 34-36			0.0569	0.0568	0.0651	0.0580	0.0583	0.0529
			(0.0001)	(0.0000)	(0.0001)	(0.0000)	(0.0001)	(0.0001)
Age 37-39			0.0568	0.0586		0.0595	0.0594	0.0539
			(0.0000)	(0.0000)		(0.0000)	(0.0001)	(0.0001)
Age 40-42			0.0578	0.0606		0.0611	0.0606	0.0549
			(0.0000)	(0.0000)		(0.0000)	(0.0001)	(0.0001)
Age 43-45		0.0646	0.0597	0.0632		0.0630	0.0622	0.0565
		(0.0001)	(0.0000)	(0.0001)		(0.0000)	(0.0001)	(0.0001)
Age 46-48		0.0645	0.0619			0.0650	0.0643	0.0582
		(0.0001)	(0.0000)			(0.0000)	(0.0001)	(0.0001)
Age 49-51		0.0673	0.0670			0.0697	0.0689	0.0621
		(0.0000)	(0.0000)			(0.0000)	(0.0001)	(0.0001)
Age 52-54	0.0772	0.0719	0.0738			0.0753	0.0747	0.0673
	(0.0002)	(0.0000)	(0.0001)			(0.0001)	(0.0001)	(0.0001)
Age 55-57	0.0790	0.0759				0.0793	0.0793	0.0716
	(0.0001)	(0.0000)				(0.0001)	(0.0001)	(0.0001)
Age 58-60	0.0810	0.0807				0.0834	0.0838	0.0762
	(0.0001)	(0.0001)				(0.0001)	(0.0001)	(0.0001)
Age 61-63	0.0844	0.0859				0.0876	0.0881	0.0809
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 65-65	0.0859	0.0894				0.0894	0.0894	0.0826
	(0.0001)	(0.0002)				(0.0001)	(0.0002)	(0.0002)
Log income	0.0145	0.0156	0.0132	0.0241	0.0336			
	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0001)			
Person fixed effect?	N	N	N	N	N	N	N	N
% of RI Sample	3.2	11.0	11.5	10.4	5.0	15.4	3.5	5.2
R-squared	0.0058	0.0182	0.0363	0.0358	0.0572	0.0390	0.0474	0.0486

Notes: This table presents regression coefficients of annual individual realized contribution rates on a set of demographic controls. The realized contribution rate is the percentage of an individual's annual income that has been invested into a retirement account over the previous year, calculated at the end of each calendar year. Columns (1)-(5) show the results including age-group controls and a control for log income, broken out by birth cohort groups. Log income is measured as the log deviation of the individual's income from the average income of the RI sample. A cohort is defined as having been born in the ten year period beginning with the year indicated. Columns (6)-(8) show the results for different groups based on the initial share of their portfolio that is invested in target date funds (TDFs). The sample is our full RI sample from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.12: Cross-Sectional Regressions of Reported Contribution Rate, Full Sample and by Income Terciles

	Reported contribution rate				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.0573	0.0715	0.0497	0.0632	0.0752
	(0.0000)	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Age 28-30	0.0611	0.0720	0.0519	0.0654	0.0797
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Age 31-33	0.0643	0.0725	0.0535	0.0664	0.0815
	(0.0000)	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Age 34-36	0.0668	0.0729	0.0549	0.0666	0.0825
	(0.0000)	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Age 37-39	0.0691	0.0736	0.0562	0.0670	0.0834
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 40-42	0.0716	0.0751	0.0580	0.0681	0.0847
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 43-45	0.0742	0.0773	0.0606	0.0702	0.0863
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 46-48	0.0770	0.0798	0.0634	0.0730	0.0882
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 49-51	0.0822	0.0853	0.0667	0.0775	0.0950
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 52-54	0.0876	0.0912	0.0703	0.0826	0.1019
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 55-57	0.0920	0.0960	0.0738	0.0875	0.1067
	(0.0000)	(0.0000)	(0.0001)	(0.0001)	(0.0001)
Age 58-60	0.0962	0.1010	0.0774	0.0928	0.1113
	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 61-63	0.1000	0.1055	0.0815	0.0975	0.1152
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 64-65	0.0997	0.0927	0.0880	0.1052	0.1276
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0001)
Log income		0.0492			
		(0.0001)			
Person fixed effect?	N	N	N	N	N
% of RI Sample	45.6	33.7	10.5	12.0	11.6
R-squared	0.0507	0.1040	0.0447	0.0376	0.0372

Notes: This table presents regression coefficients of reported contribution rate on a set of demographic controls. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The baseline specification in column (1) shows the coefficients for the regression of reported contribution rate on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.13: Cross-Sectional Regressions of Reported Contribution Rate on Age Groups by Cohort and TDF Share

	Reported contribution rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1943 Cohort	1953 Cohort	1963 Cohort	1973 Cohort	1983 Cohort	Initial TDF Share 75-100 %	Initial TDF Share 25-75 %	Initial TDF Share 0-25 %
Age 25-27				0.0699	0.0728	0.0637	0.0604	0.0536
				(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 28-30				0.0684	0.0749	0.0657	0.0644	0.0585
				(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0000)
Age 31-33				0.0700	0.0776	0.0687	0.0676	0.0622
				(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 34-36			0.0752	0.0719	0.0808	0.0714	0.0699	0.0645
			(0.0001)	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 37-39			0.0728	0.0740		0.0741	0.0718	0.0665
			(0.0001)	(0.0000)		(0.0001)	(0.0001)	(0.0001)
Age 40-42			0.0737	0.0768		0.0767	0.0739	0.0684
			(0.0000)	(0.0001)		(0.0001)	(0.0001)	(0.0001)
Age 43-45		0.0846	0.0757	0.0807		0.0795	0.0763	0.0705
		(0.0001)	(0.0000)	(0.0001)		(0.0001)	(0.0001)	(0.0001)
Age 46-48		0.0821	0.0788			0.0823	0.0788	0.0726
		(0.0001)	(0.0000)			(0.0000)	(0.0001)	(0.0001)
Age 49-51		0.0854	0.0850			0.0882	0.0843	0.0771
		(0.0001)	(0.0001)			(0.0001)	(0.0001)	(0.0001)
Age 52-54	0.0994	0.0904	0.0928			0.0941	0.0903	0.0823
	(0.0002)	(0.0001)	(0.0001)			(0.0001)	(0.0001)	(0.0001)
Age 55-57	0.1005	0.0952				0.0988	0.0954	0.0871
	(0.0001)	(0.0001)				(0.0001)	(0.0001)	(0.0001)
Age 58-60	0.1016	0.1009				0.1032	0.1000	0.0918
	(0.0001)	(0.0001)				(0.0001)	(0.0001)	(0.0002)
Age 61-63	0.1049	0.1064				0.1072	0.1037	0.0961
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 65-65	0.1065	0.1100				0.1088	0.1050	0.0976
	(0.0001)	(0.0002)				(0.0001)	(0.0002)	(0.0003)
Log income	0.0496	0.0510	0.0450	0.0507	0.0520			
	(0.0003)	(0.0001)	(0.0001)	(0.0001)	(0.0001)			
Person fixed effect?	N	N	N	N	N	N	N	N
% of RI Sample	2.4	9.0	9.6	8.6	4.2	16.9	4.0	5.9
R-squared	0.0410	0.0597	0.0358	0.0834	0.0873	0.0440	0.0482	0.0515

Notes: This table presents regression coefficients of annual individual reported contribution rates on a set of demographic controls. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. Columns (1)-(5) show the results including age-group controls and a control for log income, broken out by birth cohort groups. Log income is measured as the log deviation of the individual's income from the average income of the RI sample. A cohort is defined as having been born in the ten-year period beginning with the year indicated. Columns (6)-(8) show the results for different groups based on the initial share of their portfolio that is invested in target date funds (TDFs). The sample is our full RI sample from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.14: Within-Person Regressions of Reported Contribution Rate, Full Sample and by Income Terciles

	Reported contribution rate				
	(1) All Observations	(2) All Observations	(3) First Tercile of Initial Income	(4) Second Tercile of Initial Income	(5) Third Tercile of Initial Income
Age 25-27	0.0345	0.0306	0.0373	0.0380	0.0547
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0003)
Age 28-30	0.0407	0.0359	0.0422	0.0453	0.0645
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 31-33	0.0463	0.0406	0.0467	0.0515	0.0714
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 34-36	0.0507	0.0444	0.0505	0.0562	0.0763
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 37-39	0.0545	0.0476	0.0538	0.0602	0.0799
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 40-42	0.0582	0.0507	0.0568	0.0640	0.0832
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 43-45	0.0616	0.0538	0.0597	0.0677	0.0862
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 46-48	0.0653	0.0571	0.0628	0.0716	0.0893
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 49-51	0.0714	0.0631	0.0664	0.0768	0.0970
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
Age 52-54	0.0778	0.0696	0.0707	0.0825	0.1048
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0001)
Age 55-57	0.0837	0.0756	0.0750	0.0885	0.1111
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 58-60	0.0899	0.0819	0.0797	0.0949	0.1174
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 61-63	0.0959	0.0884	0.0848	0.1013	0.1236
	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Age 64-65	0.0997	0.0927	0.0880	0.1052	0.1276
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Log income		0.0212			
		(0.0001)			
Person fixed effect?	Y	Y	Y	Y	Y
% of RI Sample	45.6	33.7	10.5	12.0	11.6
R-squared	0.7764	0.7870	0.7809	0.7649	0.7438

Notes: This table presents regression coefficients of reported contribution rate on a set of demographic controls. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. The baseline specification in column (1) shows the coefficients for the regression of reported contribution rate on age group dummies. In the second column, we add a control for the log of income in the current year, measured as the individual's log deviation from the average income in the RI sample. Columns (3)-(5) show the results of the baseline specification for the first (lowest) through the third tercile of initial income, respectively. Initial income is based upon the income observed in the first (or second, if first is not available) year that we observe the individual. All regressions include a person fixed effect. The age group coefficients are normalized by adding the average fixed effect back to the estimated coefficients. The excluded age group is those aged 64-65. The sample is our full set of retirement investors (RI) from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.15: Within-Person Regressions of Reported Contribution Rate on Age Groups by Cohort and TDF Share

	Reported contribution rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1943 Cohort	1953 Cohort	1963 Cohort	1973 Cohort	1983 Cohort	Initial TDF Share 75-100 %	Initial TDF Share 25-75 %	Initial TDF Share 0-25 %
Age 25-27				0.0577	0.0791	0.0598	0.1062	0.0632
				(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
Age 28-30				0.0586	0.0856	0.0633	0.1121	0.0707
				(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
Age 31-33				0.0617	0.0927	0.0670	0.1172	0.0774
				(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
Age 34-36			0.0641	0.0650	0.0986	0.0703	0.1214	0.0827
			(0.0001)	(0.0001)	(0.0000)	(0.0001)	(0.0003)	(0.0003)
Age 37-39			0.0618	0.0693		0.0733	0.1248	0.0873
			(0.0001)	(0.0001)		(0.0001)	(0.0003)	(0.0003)
Age 40-42			0.0631	0.0743		0.0762	0.1283	0.0916
			(0.0001)	(0.0001)		(0.0001)	(0.0003)	(0.0003)
Age 43-45		0.0877	0.0656	0.0817		0.0790	0.1315	0.0956
		(0.0002)	(0.0001)	(0.0000)		(0.0001)	(0.0003)	(0.0003)
Age 46-48		0.0855	0.0695			0.0820	0.1347	0.0995
		(0.0002)	(0.0001)			(0.0001)	(0.0003)	(0.0003)
Age 49-51		0.0888	0.0768			0.0881	0.1407	0.1053
		(0.0001)	(0.0001)			(0.0001)	(0.0002)	(0.0003)
Age 52-54	0.1206	0.0944	0.0802			0.0943	0.1471	0.1116
	(0.0002)	(0.0001)	(0.0001)			(0.0001)	(0.0002)	(0.0003)
Age 55-57	0.1209	0.1005				0.0997	0.1531	0.1176
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 58-60	0.1215	0.1081				0.1053	0.1591	0.1240
	(0.0001)	(0.0001)				(0.0001)	(0.0002)	(0.0002)
Age 61-63	0.1260	0.1166				0.1104	0.1643	0.1303
	(0.0001)	(0.0001)				(0.0001)	(0.0001)	(0.0002)
Age 64-65	0.1316	0.1194				0.1133	0.1674	0.1342
	(0.0000)	(0.0000)				(0.0000)	(0.0000)	(0.0000)
Log income	0.0231	0.0237	0.0180	0.0182	0.0241			
	(0.0004)	(0.0002)	(0.0002)	(0.0002)	(0.0002)			
Person fixed effect?	Y	Y	Y	Y	Y	Y	Y	Y
% of RI Sample	2.4	9.0	9.6	8.6	4.2	16.9	4.0	5.9
R-squared	0.8488	0.7977	0.7666	0.7325	0.7472	0.7599	0.7601	0.7497

Notes: This table presents regression coefficients of annual individual reported contribution rates on a set of demographic controls. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. Columns (1)-(5) show the results including age-group controls and a control for log income, broken out by birth cohort groups. Log income is measured as the log deviation of the individual's income from the average income of the RI sample. A cohort is defined as having been born in the ten year period beginning with the year indicated. Columns (6)-(8) show the results for different groups based on the initial share of their portfolio that is invested in target date funds (TDFs). All regressions include a person fixed effect. The age group coefficients are normalized by adding the average fixed effect back to the estimated coefficients. The excluded age group is those aged 64-65. The sample is our full RI sample from 2006-2018. Standard errors, in parentheses, are clustered at the individual level.

Table B.16: Regressions of Reported Contribution Rate on the Pension Protection Act: Long-run Effect, Treated in 2007 Only

	Reported contribution rate						
	(1) Full Sample	(2) Bottom Income Tercile	(3) Top Income Tercile	(4) Age Enrolled 25-34	(5) Age Enrolled 35-44	(6) Age Enrolled 45-54	(7) Age Enrolled 55-65
Year of x Treatment	-0.0049	-0.0059	-0.0047	-0.0034	-0.0059	-0.0067	-0.0095
	(0.0003)	(0.0004)	(0.0005)	(0.0003)	(0.0005)	(0.0006)	(0.0010)
1 Year After x Treatment	-0.0104	-0.0084	-0.0120	-0.0078	-0.0105	-0.0121	-0.0140
	(0.0001)	(0.0002)	(0.0003)	(0.0001)	(0.0002)	(0.0003)	(0.0007)
2 Years After x Treatment	-0.0074	-0.0070	-0.0086	-0.0055	-0.0069	-0.0098	-0.0117
	(0.0001)	(0.0002)	(0.0003)	(0.0001)	(0.0002)	(0.0003)	(0.0006)
3 Years After x Treatment	-0.0026	-0.0030	-0.0032	-0.0015	-0.0025	-0.0046	-0.0058
	(0.0001)	(0.0002)	(0.0003)	(0.0002)	(0.0002)	(0.0003)	(0.0006)
4 Years After x Treatment	-0.0016	-0.0016	-0.0027	-0.0010	-0.0027	-0.0055	-0.0070
	(0.0001)	(0.0002)	(0.0003)	(0.0002)	(0.0002)	(0.0003)	(0.0007)
5 Years After x Treatment	-0.0008	0.0001	-0.0032	-0.0012	-0.0025	-0.0041	-0.0045
	(0.0001)	(0.0002)	(0.0004)	(0.0002)	(0.0002)	(0.0004)	(0.0008)
1 Year After	-0.0039	-0.0064	-0.0036	-0.0062	-0.0034	0.0029	0.0017
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0006)
2 Years After	-0.0091	-0.0093	-0.0093	-0.0094	-0.0115	-0.0051	-0.0075
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
3 Years After	-0.0126	-0.0125	-0.0123	-0.0117	-0.0157	-0.0105	-0.0142
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
4 Years After	-0.0129	-0.0135	-0.0117	-0.0113	-0.0163	-0.0114	-0.0148
	(0.0002)	(0.0003)	(0.0004)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
5 Years After	-0.0131	-0.0144	-0.0110	-0.0114	-0.0170	-0.0127	-0.0170
	(0.0002)	(0.0003)	(0.0004)	(0.0002)	(0.0003)	(0.0004)	(0.0008)
Log income	0.0429						
	(0.0002)						
Constant	0.0806	0.0698	0.0982	0.0706	0.0826	0.0891	0.1047
	(0.0002)	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0004)	(0.0003)
Firm Fixed Effect?	Y	Y	Y	Y	Y	Y	Y
% of RI Sample	1.9	0.7	0.6	1.0	0.7	0.5	0.2
% of Sample Enrolled 2005-2008	27.7	9.7	8.9	14.4	10.0	7.0	2.3
R-squared	0.1512	0.1213	0.0840	0.1355	0.1232	0.1120	0.1476

Notes: This table presents regression coefficients of reported contribution rate on being treated with the Pension Protection Act (PPA) of 2006. "Year of" means the year the individual enrolled in their retirement plan and "x years after" is x years after they enrolled in the plan. Each column includes year dummies for each year after enrollment, and interactions of these dummies with the treatment dummy. The treatment dummy is equal to one if the individual enrolled in 2007 immediately after the PPA, and zero if they enrolled in 2005 or 2006. The full sample is those enrolled from 2005-2007 who otherwise meet the RI sample criteria. The bottom (top) income tercile includes those whose initial income is in the lowest (highest) tercile. Columns (4)-(7) break out the result for all individuals enrolled from 2005-2007 by age at enrollment. The reported contribution rate is the percentage of their income that an individual designates to be allocated into their retirement accounts at the beginning of each calendar year. Log income, when included, is the log deviation of the individual's current income from the average income of the RI sample. Standard errors, in parentheses, are clustered at the household level.

Appendix C

Appendix for Chapter 3

C.1 Appendix Figures

Figure C.1: Distribution of PPP Loan Recipients and All Firms, by Firm Employment

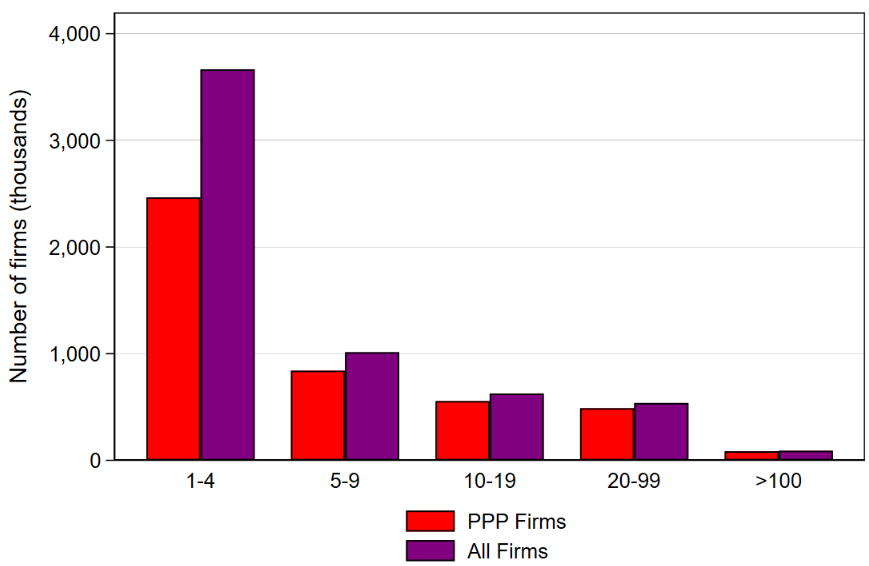

Notes: This figure shows the distribution of firms that received the PPP (red bars) and all firms in the U.S. (purple bars). Based on the author's calculations from the Treasury micro data on SBA loans and the 2016 Statistics of U.S. Businesses from the Census Bureau.

Figure C.2: Cases per Person, Sample versus U.S.

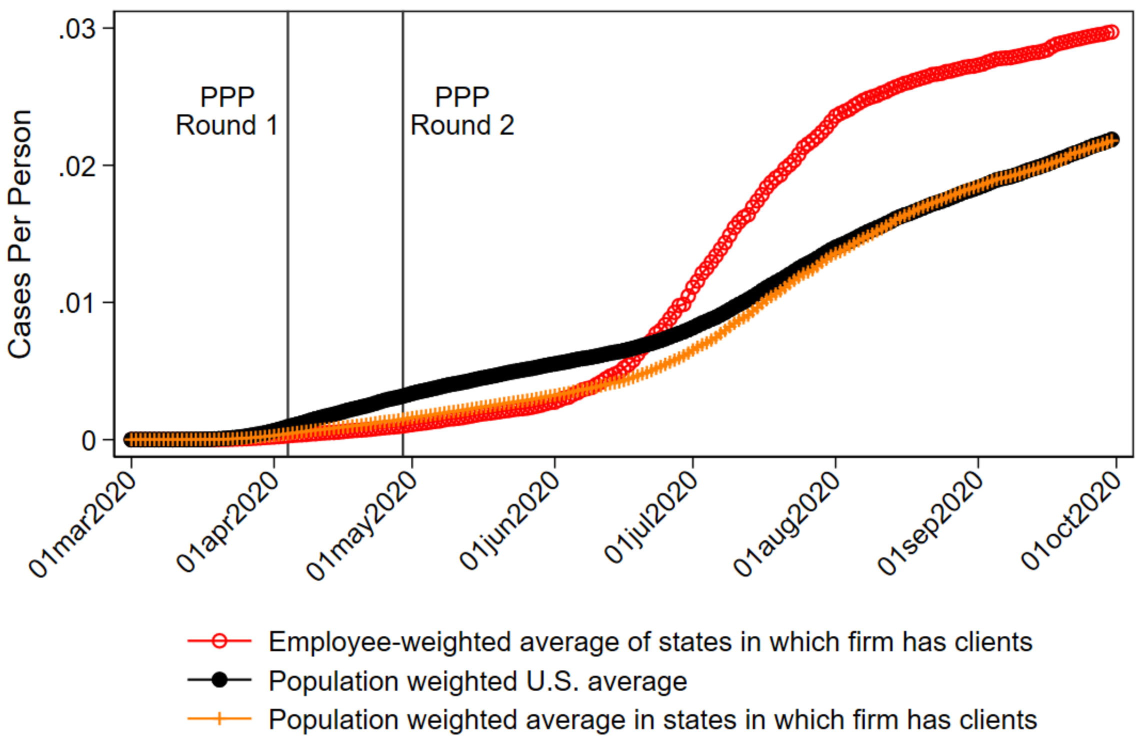

Notes: This figure shows the daily number of cases per person. First, I show the the average, weighted by number of employees in each state in the sample, for the states in which the firm has clients. Second, I show the population-weighted U.S. average in all 50 states. Lastly, I show the population-weighted average in the states in which the firm has clients. Case count data are from the Johns Hopkins Center for Systems Science and Engineering. Population data are from the Census.

Figure C.3: NAICS Distribution and Time Series of Selected Variables, PPP versus Non-PPP firms

(a) NAICS Distribution

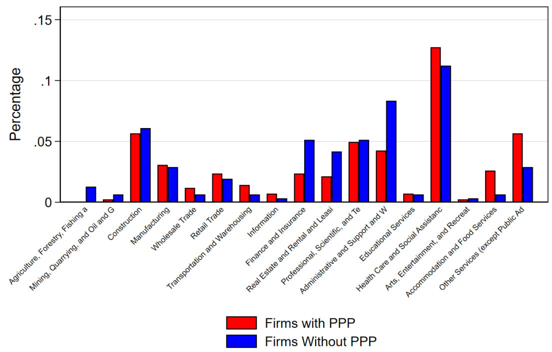

(b) Monthly employment growth

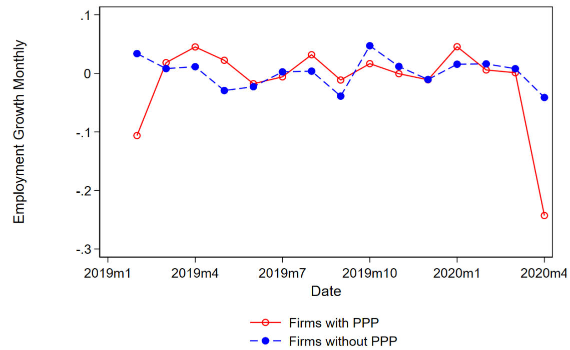

(c) Employment

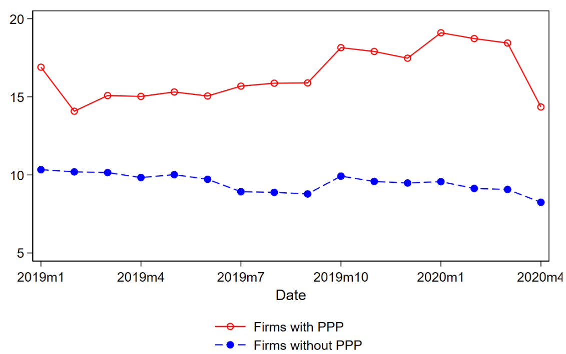

(d) Average wage per worker

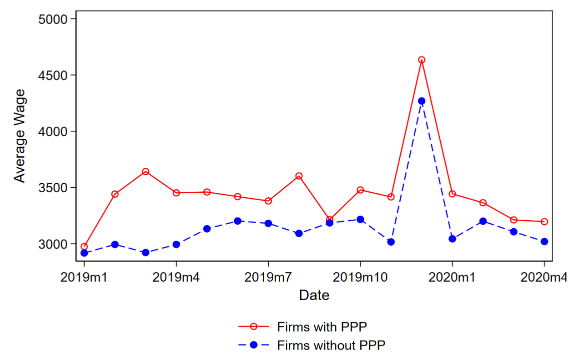

Notes: These figures show characteristics of each variable for the firms that received PPP (red) versus the firms that did not receive PPP (blue). The spike in average wage per worker at the end of 2019 is due to end-of-year/holiday bonus compensation.

Figure C.4: Average Wage per Worker, by Firm Size

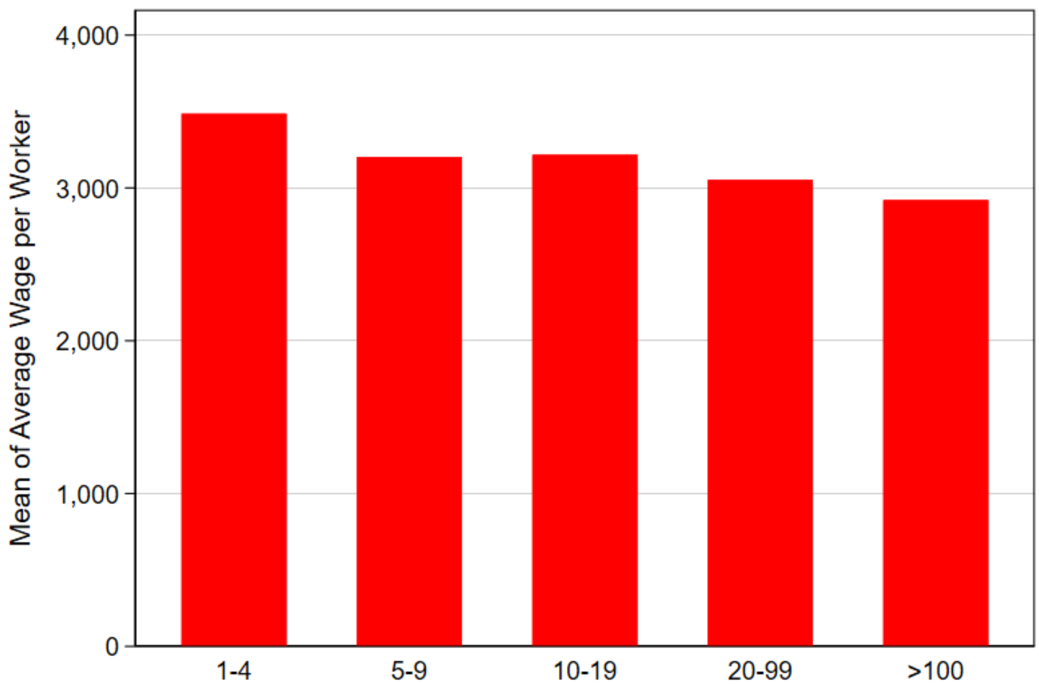

Notes: This figure shows the average monthly wages per worker at firms within each bucket of number of employees, as of February 2020.

Figure C.5: Odds Ratios of NAICS Fixed Effects - Early versus Late Applicants

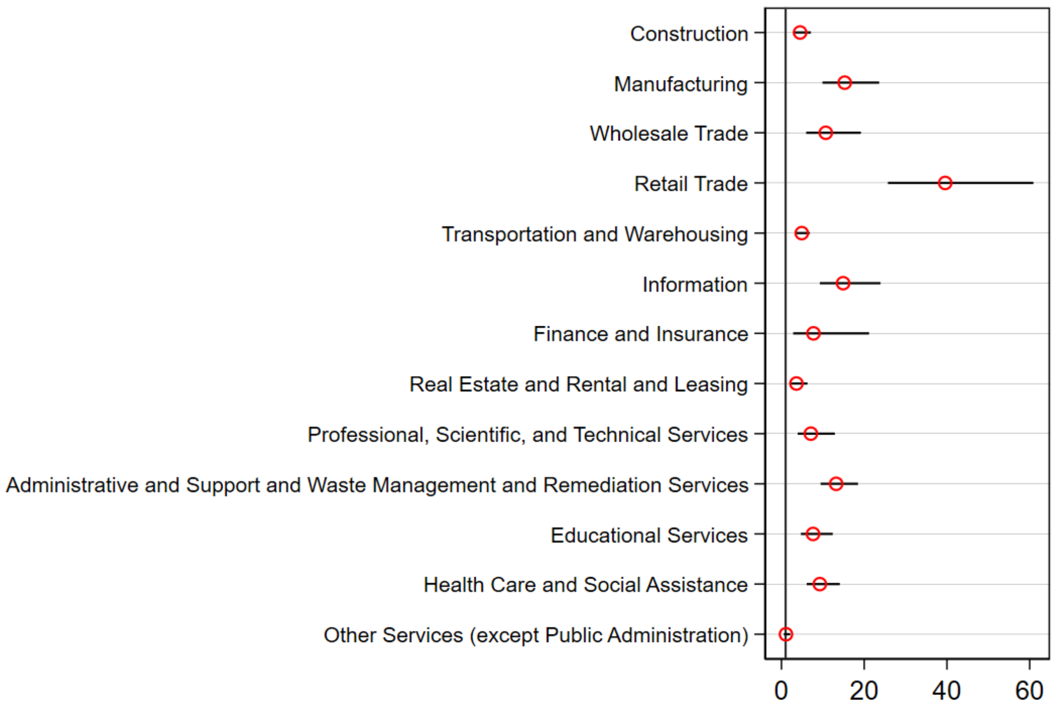

Notes: This figure shows the odds ratios for the NAICS industry fixed effects from the logistic regression of take-up on firm characteristics. This version compares late applicants versus early applicants. The reference category is NAICS 72, Food and Accommodation Services. Standard errors are clustered at the 2-digit NAICS level. The coefficients correspond to the final column in Table C.3.

Figure C.6: Effect of PPP on Employment and Wage Bill: Bootstrapped Results

(a) Employment

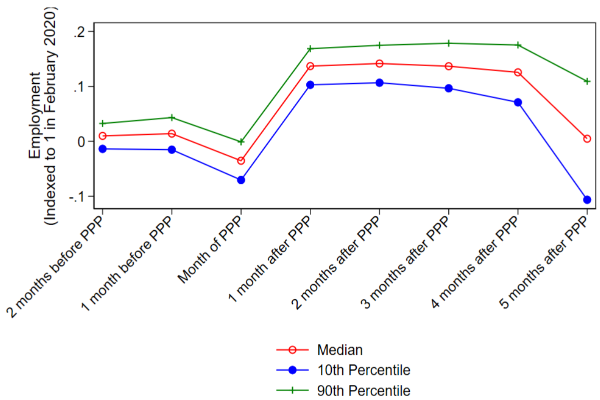

(b) Wage Bill

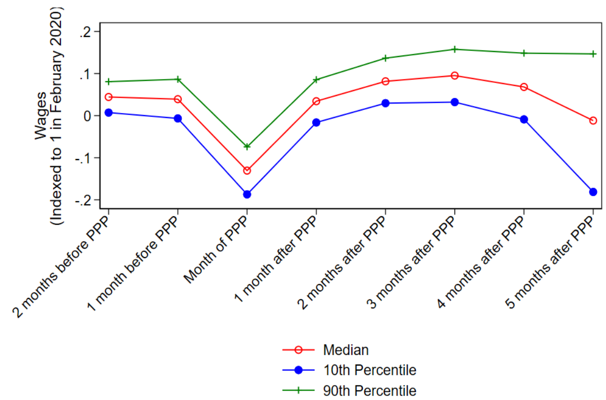

(c) # of Total Hires

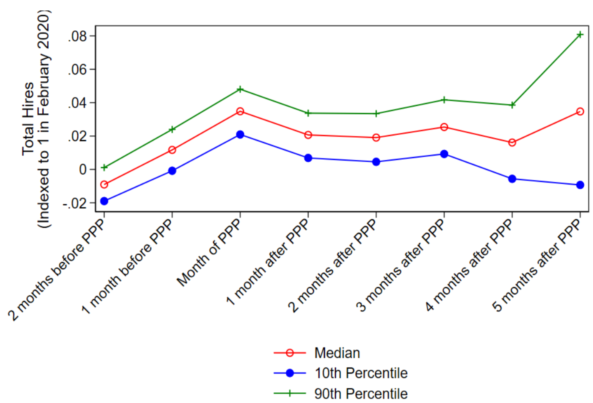

(d) # of New Hires

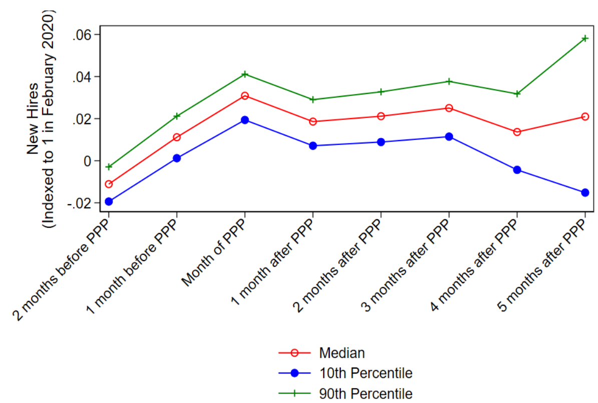

(e) # of Re-hires

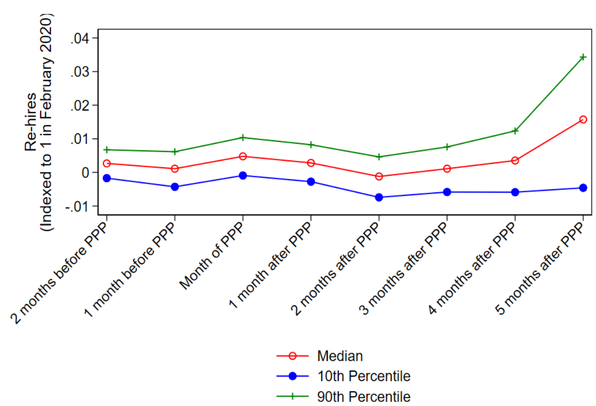

Notes: These figures show the results from the bootstrapped difference-in-differences specification in equation (3.2). The bootstrapping procedure reconstructs the dependent variable by randomly sampling from residuals of the main estimation, and then re-estimating the main specification. I repeat this 1,000 times and plot the 10th percentile, median, and 90th percentile. The dependent variable is normalized to one in February 2020.

Figure C.7: Effect of PPP on Employment and Wage Bill, Balanced Panel

(a) Employment

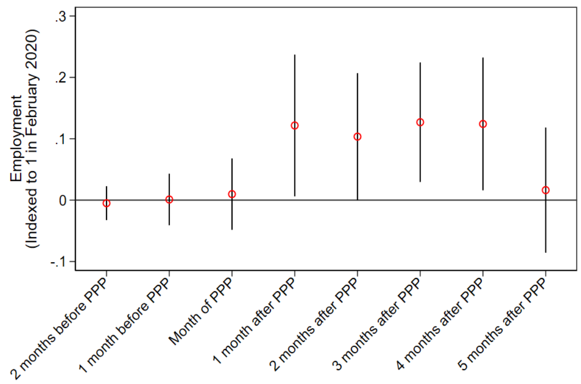

(b) Wage Bill

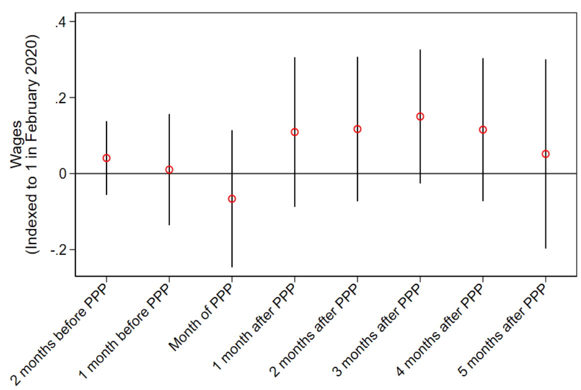

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) for the balanced panel of firms. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variable is normalized to one in February 2020.

Figure C.8: Effect of PPP on Employment and Wage Bill, Controlling for Yearly Growth

(a) Employment

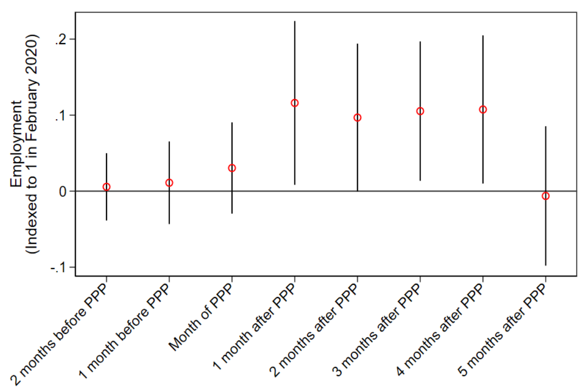

(b) Wage Bill

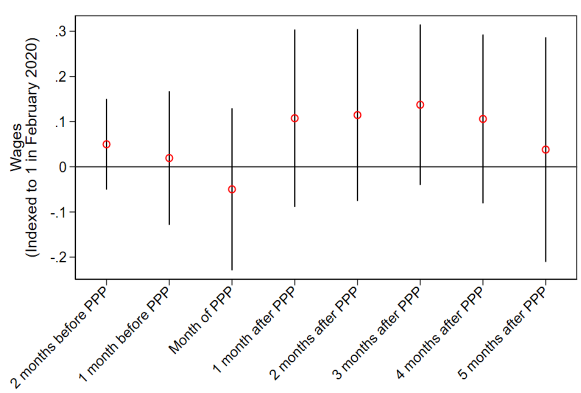

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) controlling for the firm's yearly growth in the first two months of 2020. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variable is normalized to one in February 2020.

Figure C.9: Effect of PPP on Employment, Difference-in-Differences Results, Progressively Adding Controls

(a) No Controls

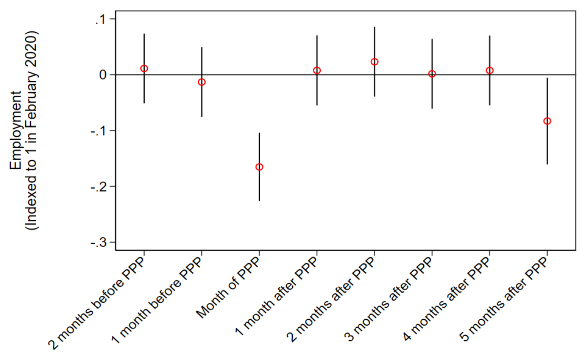

(b) Size Control Only

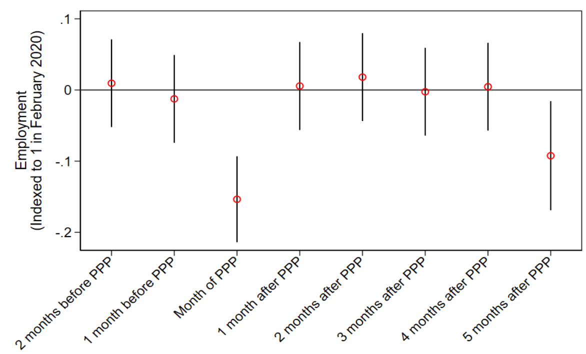

(c) Size and Industry by Time Controls

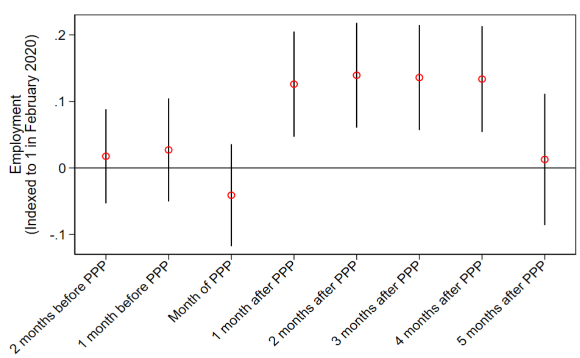

(d) Size, Industry by Time, and City by Time controls

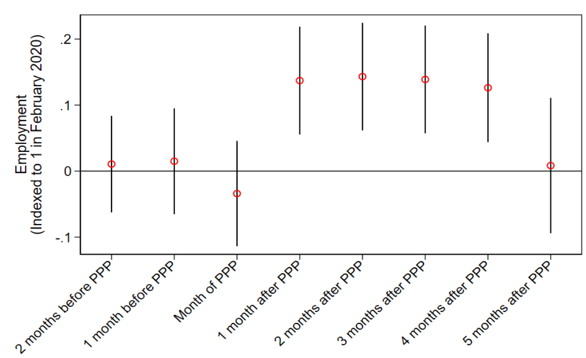

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) with the controls progressively added. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. Panel a shows the regression with no controls. Panel b shows the regression with only the employment size fixed effects. Panel c shows the results with the size control and the industry by time fixed effects. Panel d shows the results with the full set of results, as shown in Figure 3-3. The dependent variable is normalized to one in February 2020.

Figure C.10: Effect of PPP on Wage Bill, Difference-in-Differences Results, Progressively Adding Controls

(a) No Controls

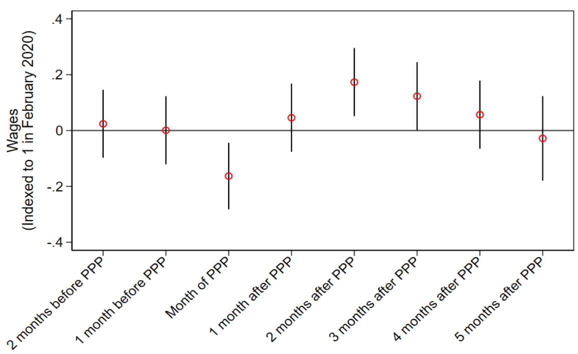

(b) Size Control Only

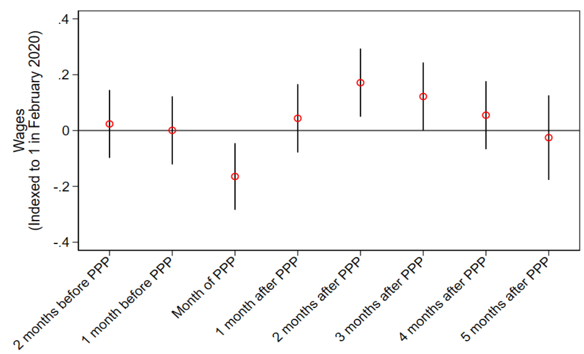

(c) Size and Industry by Time Controls

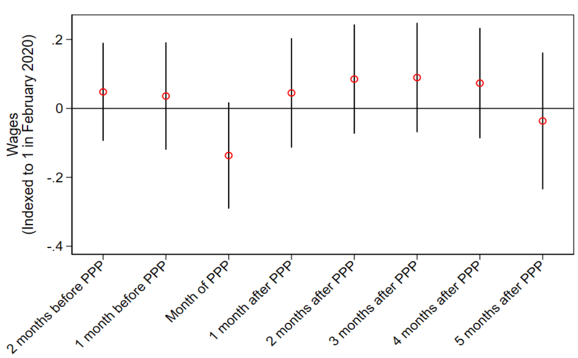

(d) Size, Industry by Time, and City by Time controls

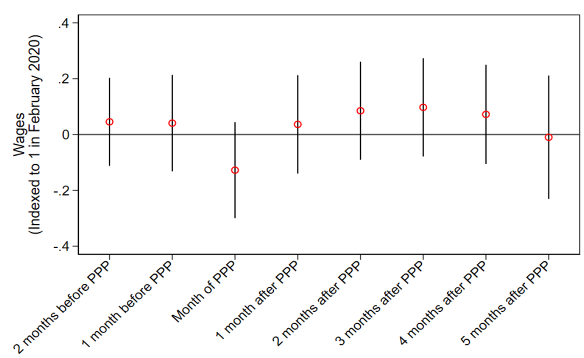

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) with the controls progressively added. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. Panel a shows the regression with no controls. Panel b shows the regression with only the employment size fixed effects. Panel c shows the results with the size control and the industry by time fixed effects. Panel d shows the results with the full set of results, as shown in Figure 3-3. The dependent variable is normalized to one in February 2020.

Figure C.11: Effect of PPP, by Prevalence of Hourly Workers in the Industry

(a) Employment: bottom tercile of hourly workers

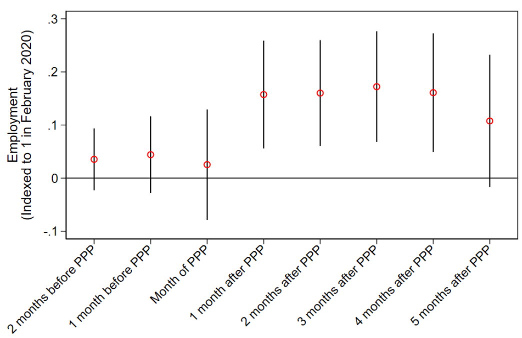

(b) Employment: top tercile of hourly workers

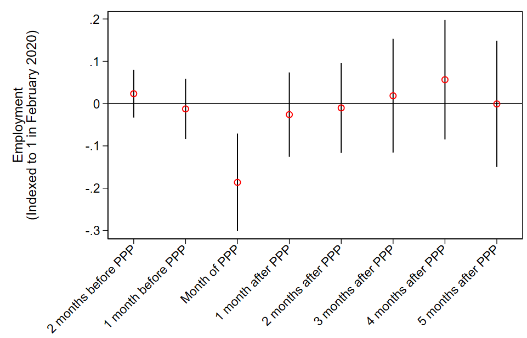

(c) Wage Bill: bottom tercile of hourly workers

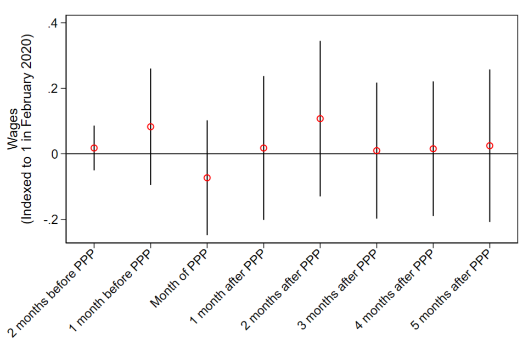

(d) Wage Bill: top tercile of hourly workers

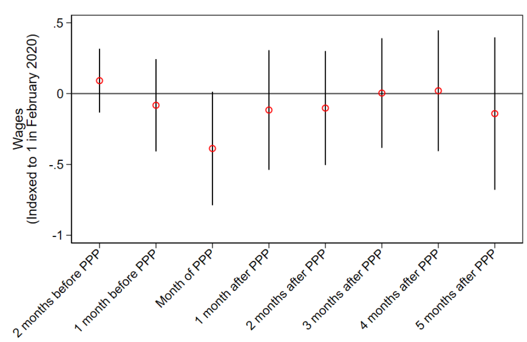

Notes: These figures show the results from the difference-in-differences specification in equation (3.2). I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The figures on the left show the results for firms that are in industries which are in the lowest tercile of percentage of hourly workers (i.e. have few hourly workers). The figures on the right show the results for firms that are industries that are in the top tercile of hourly workers (i.e. have many hourly workers). The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals.

Figure C.12: Effect of PPP, by Percentage of Workers that can Work Remotely

(a) Employment: top half of remote workers

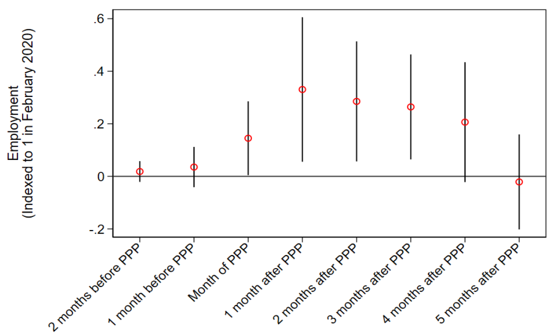

(b) Employment: bottom half of remote workers

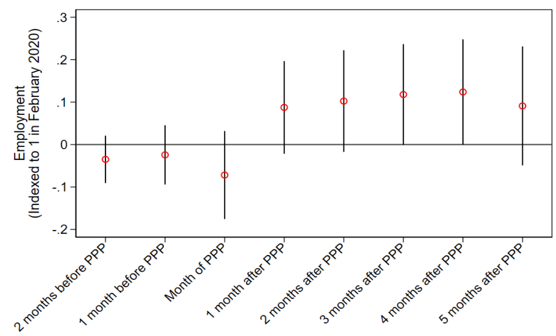

(c) Wage Bill: top half of remote workers

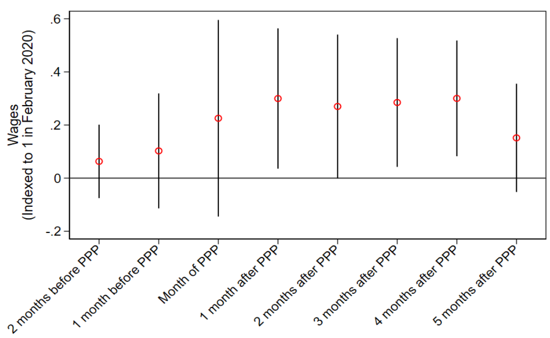

(d) Wage Bill: bottom half of remote workers

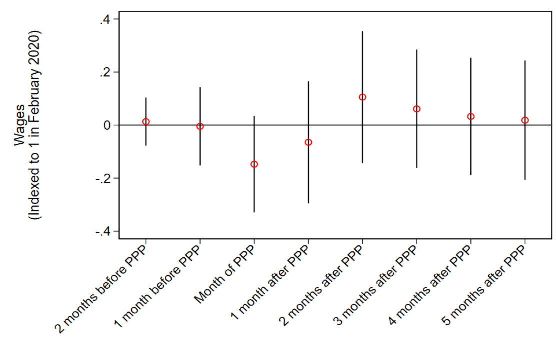

Notes: These figures show the results from the difference-in-differences specification in equation (3.2). I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The figures on the left show the results for firms that are in industries which are in the top half of percentage of workers that can work remotely (i.e. have more remote-capable workers). The figures on the right show the results for firms that are industries that are in the bottom tercile of remote workers (i.e. have fewer remote-capable workers). The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals.

Figure C.13: Effect of PPP, by Essential versus Non-Essential Businesses

(a) Employment: non-essential businesses

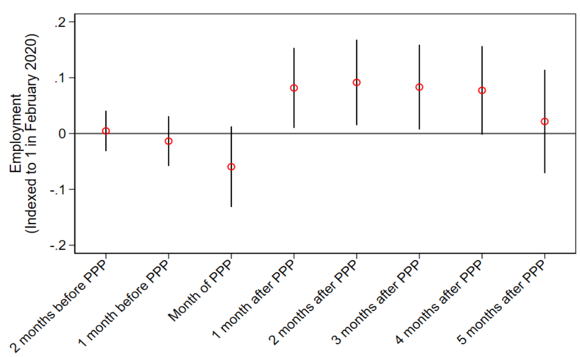

(b) Employment: essential businesses

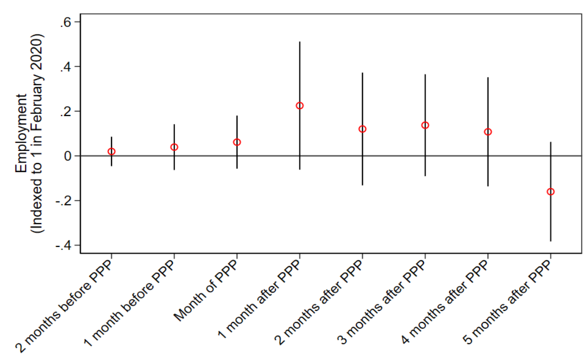

(c) Wage Bill: non-essential businesses

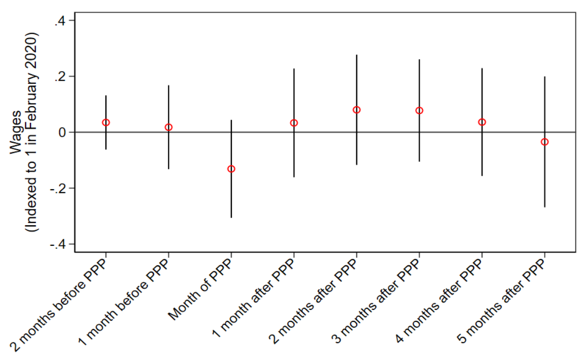

(d) Wage Bill: essential businesses

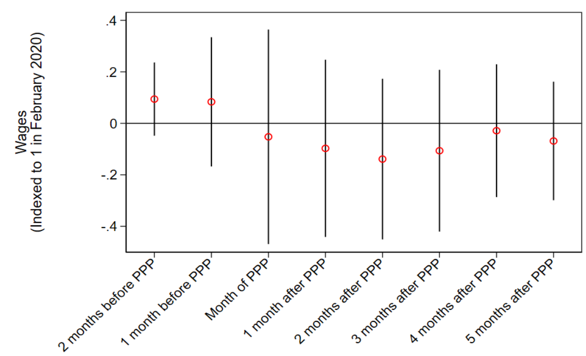

Notes: These figures show the results from the difference-in-differences specification in equation (3.2), I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The figures on the left show the results for firms that are considered essential businesses. The figures on the right show the results for non-essential businesses. The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals.

Figure C.14: Employee Turnover, PPP versus Non-PPP Firms

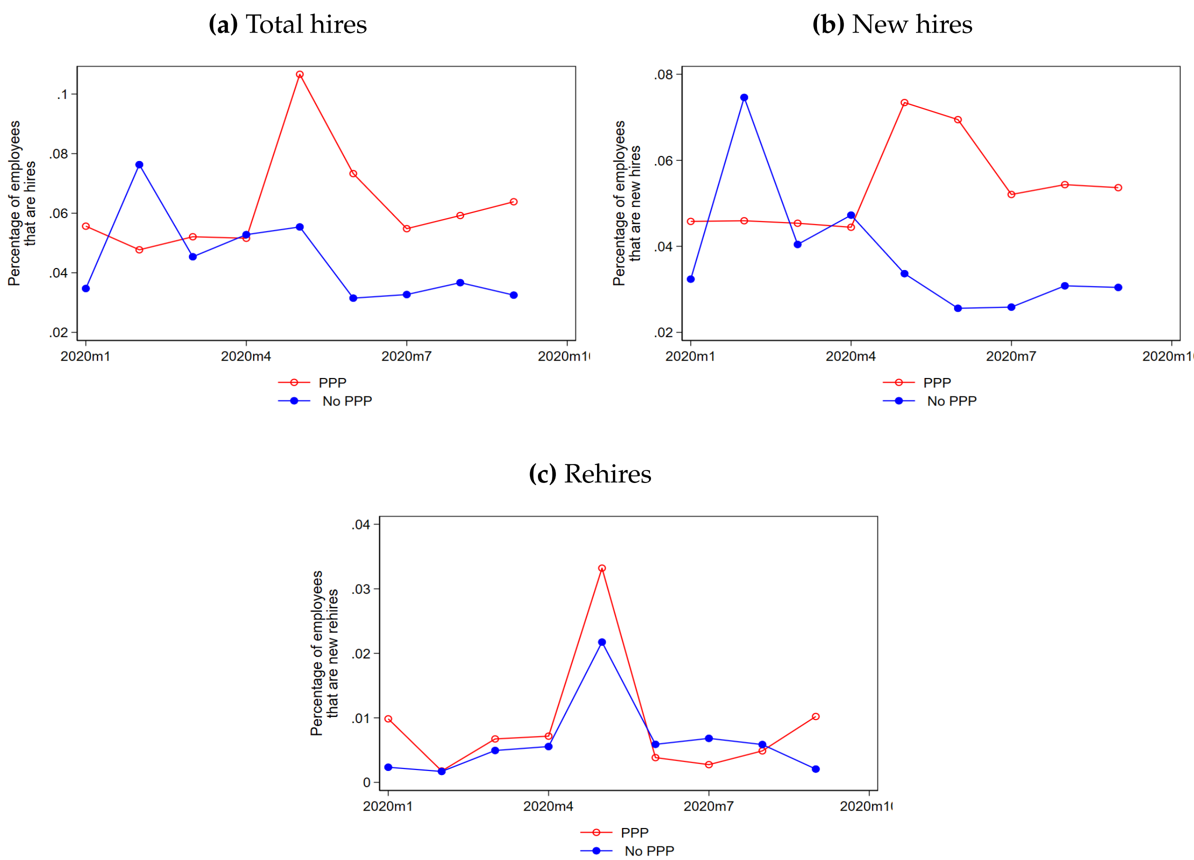

Notes: This figures shows the average number of total hires (a), new hires (b) and rehires (c), as a percentage of base employment in PPP versus non-PPP firms. Base employment is the average number of employees from February 2019-February 2020.

Figure C.15: Effect of PPP on Employee Turnover, Balanced Panel

(a) # of Total Hires

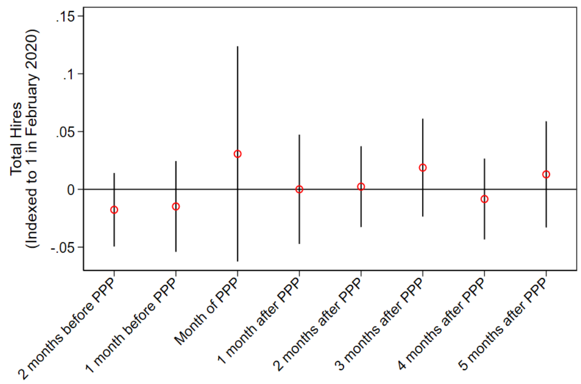

(b) # of New Hires

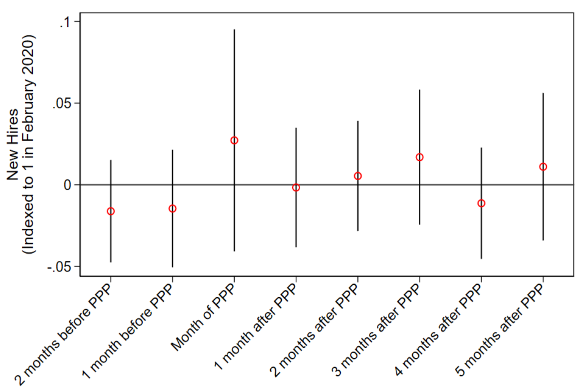

(c) # of Rehires

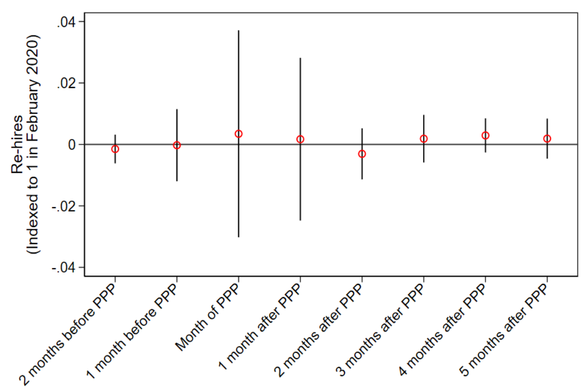

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) for the balanced panel of firms. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variable is normalized to one in February 2020.

Figure C.16: Effect of PPP on Employee Turnover, Controlling for Yearly Growth

(a) # of Total Hires

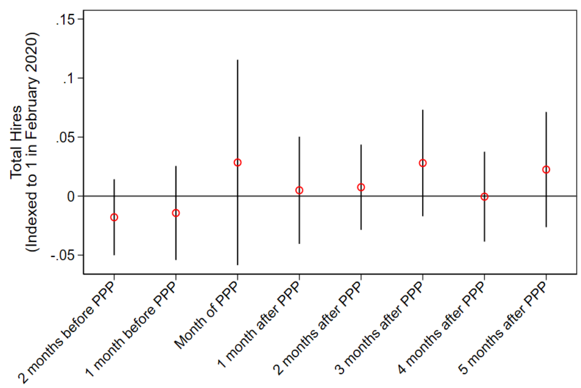

(b) # of New Hires

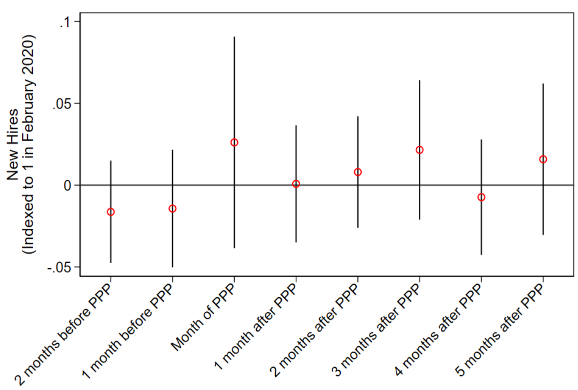

(c) # of Rehires

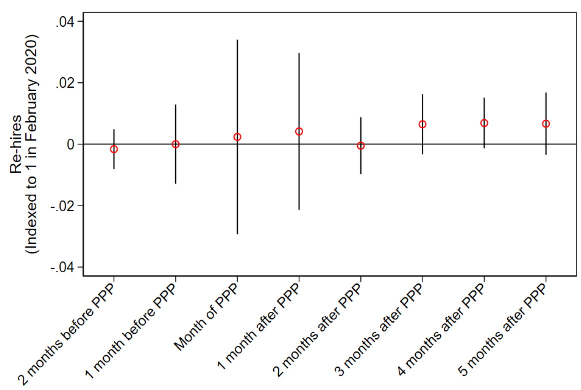

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) controlling for the firm's yearly growth in the first two months of 2020. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variable is normalized to one in February 2020.

Figure C.17: Effect of PPP on Total Hires, Difference-in-Differences Results, Progressively Adding Controls

(a) No Controls

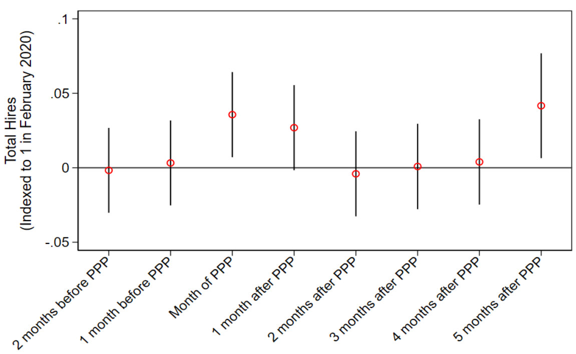

(b) Size Control Only

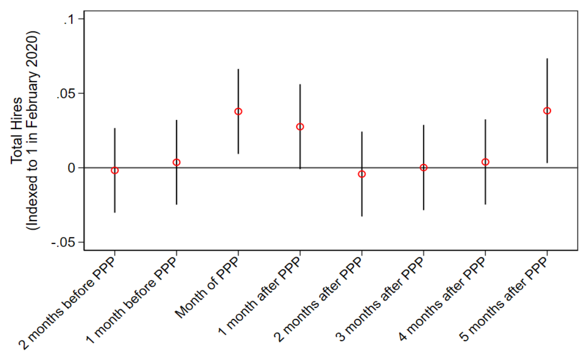

(c) Size and Industry by Time Controls

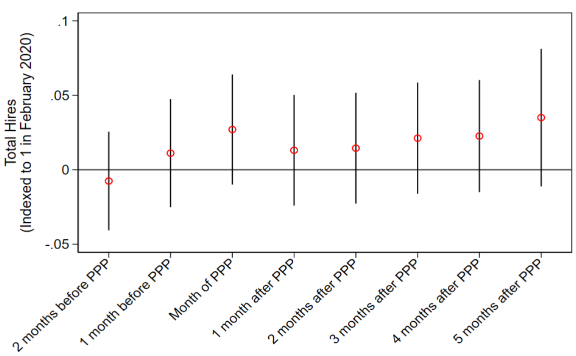

(d) Size, Industry by Time, and City by Time controls

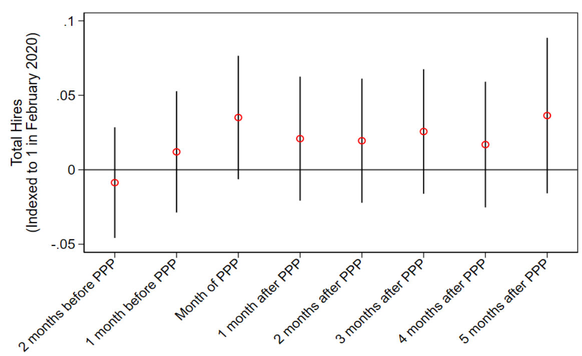

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) with the controls progressively added. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. Panel a shows the regression with no controls. Panel b shows the regression with only the employment size fixed effects. Panel c shows the results with the size control and the industry by time fixed effects. Panel d shows the results with the full set of results, as shown in Figure 3-5. The dependent variable is normalized to one in February 2020.

Figure C.18: Effect of PPP on New Hires, Difference-in-Differences Results, Progressively Adding Controls

(a) No Controls

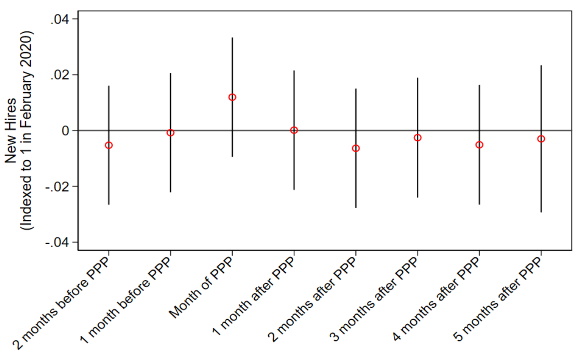

(b) Size Control Only

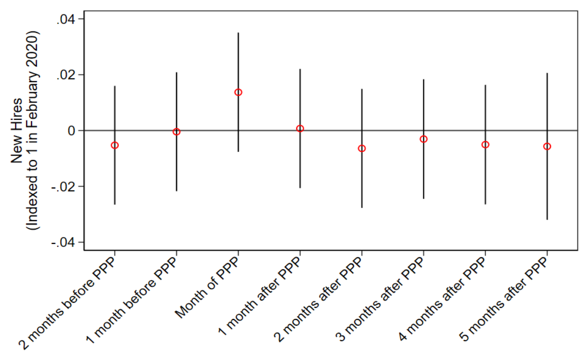

(c) Size and Industry by Time Controls

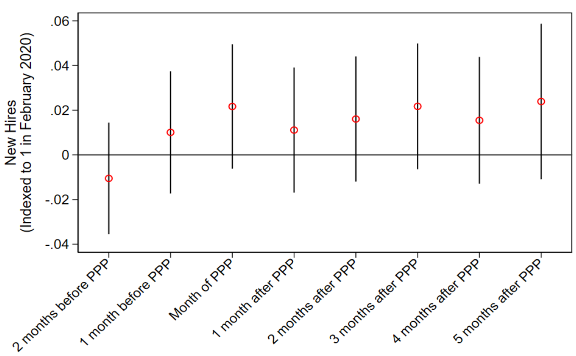

(d) Size, Industry by Time, and City by Time controls

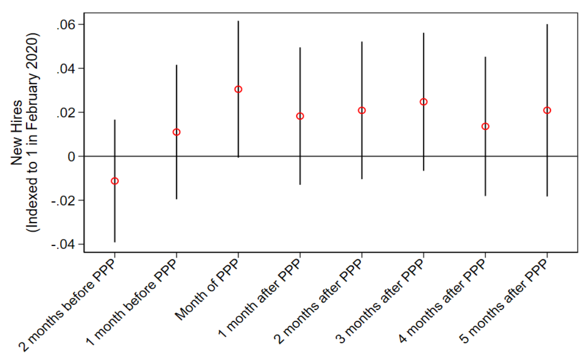

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) with the controls progressively added. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. Panel a shows the regression with no controls. Panel b shows the regression with only the employment size fixed effects. Panel c shows the results with the size control and the industry by time fixed effects. Panel d shows the results with the full set of results, as shown in Figure 3-5. The dependent variable is normalized to one in February 2020.

Figure C.19: Effect of PPP on Re-hires, Difference-in-Differences Results, Progressively Adding Controls

(a) No Controls

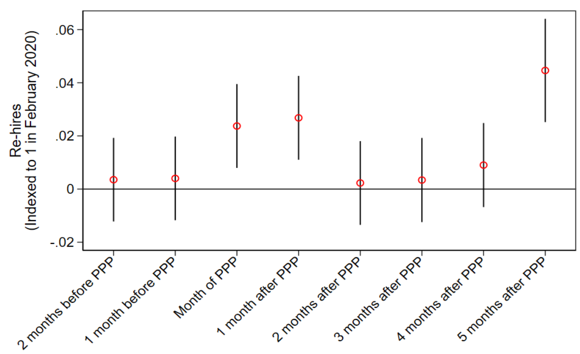

(b) Size Control Only

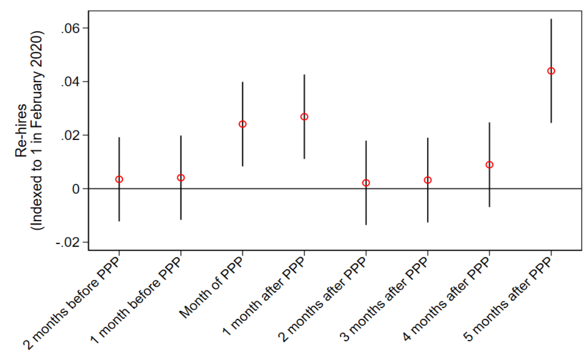

(c) Size and Industry by Time Controls

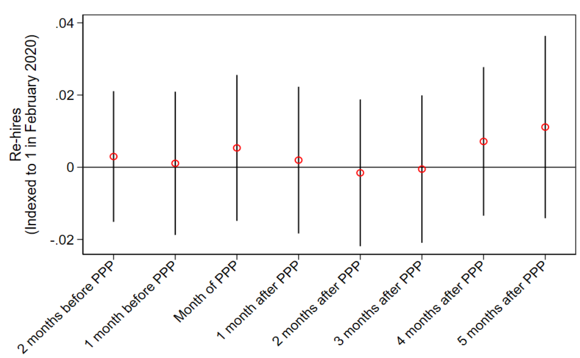

(d) Size, Industry by Time, and City by Time controls

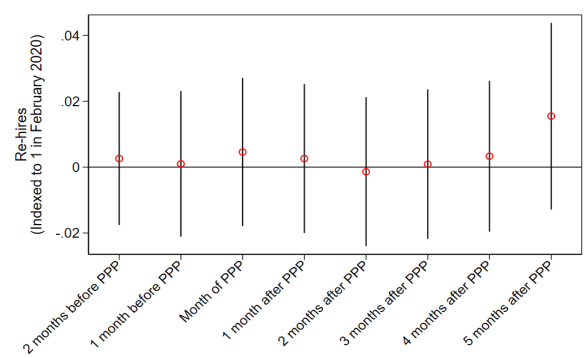

Notes: These figures show the results from the difference-in-differences specification in equation (3.2) with the controls progressively added. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. Panel a shows the regression with no controls. Panel b shows the regression with only the employment size fixed effects. Panel c shows the results with the size control and the industry by time fixed effects. Panel d shows the results with the full set of results, as shown in Figure 3-5. The dependent variable is normalized to one in February 2020.

Figure C.20: Effect of PPP on Employee Turnover, by Prevalence of Hourly Workers in the Industry

(a) Total Hires: bottom tercile of hourly workers

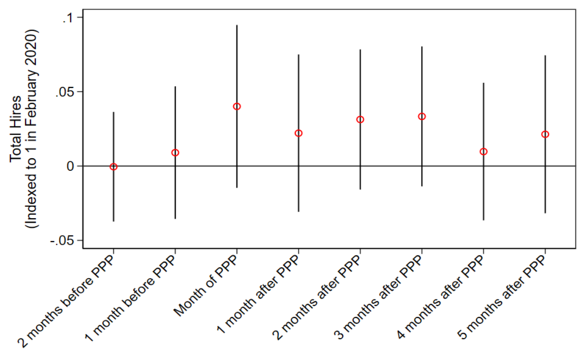

(b) Total Hires: top tercile of hourly workers

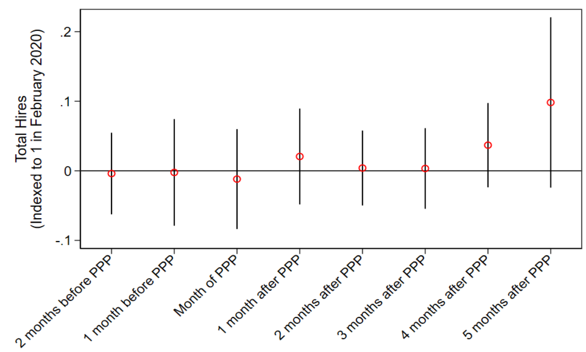

(c) New Hires: bottom tercile of hourly workers

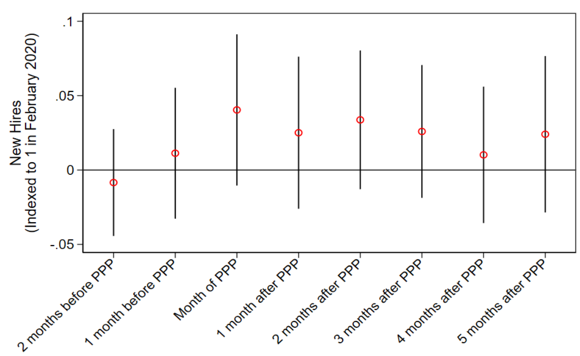

(d) New Hires: top tercile of hourly workers

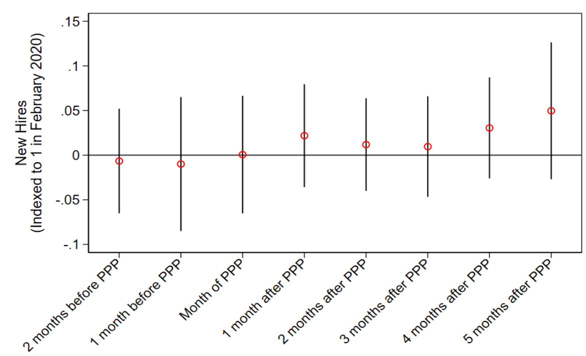

(e) Rehires: bottom tercile of hourly workers

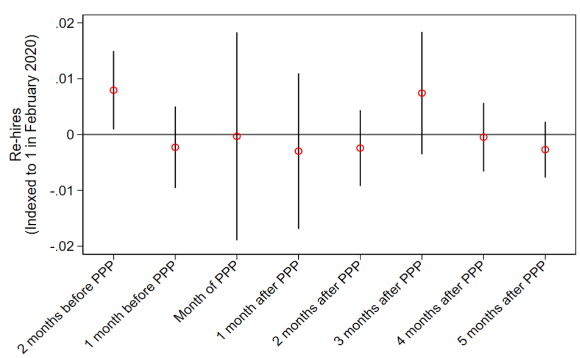

(f) Rehires: top tercile of hourly workers

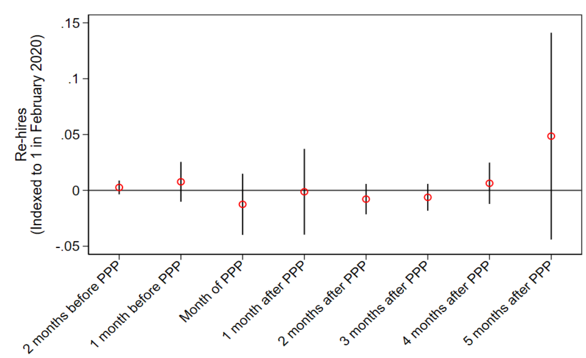

Notes: These figures show the results from the difference-in-differences specification in equation (3.2), I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The figures on the left show the results for firms in industries with the fewest hourly workers. The figures on the right show the results for firms in industries that have the most hourly workers. The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals

Figure C.21: Effect of PPP on Employee Turnover, by Remote Capability of Employees

(a) Total Hires: high remote capability

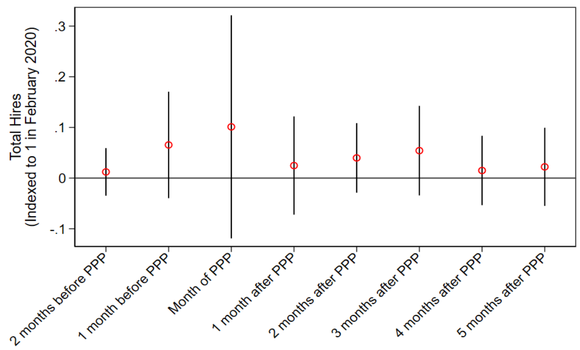

(b) Total Hires: low remote capability

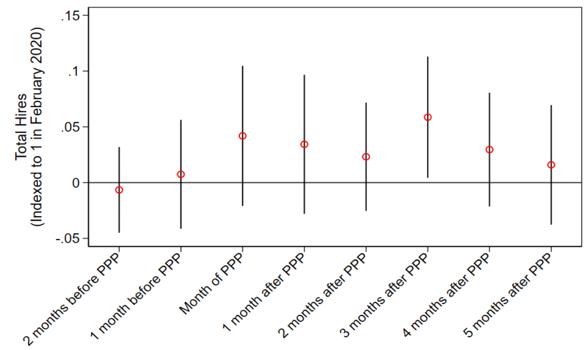

(c) New Hires: high remote capability

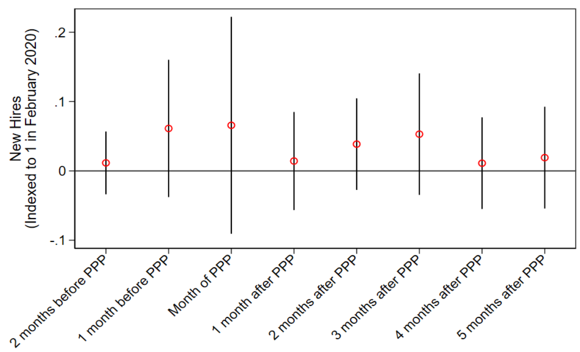

(d) New Hires: low remote capability

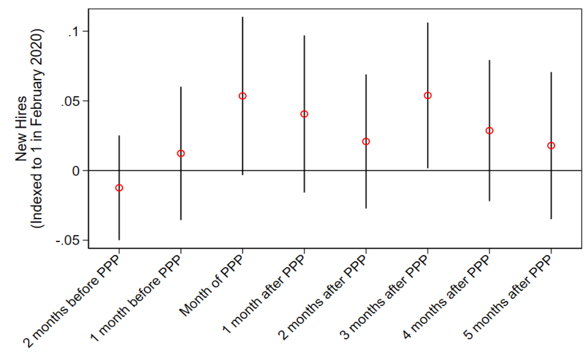

(e) Rehires: high remote capability

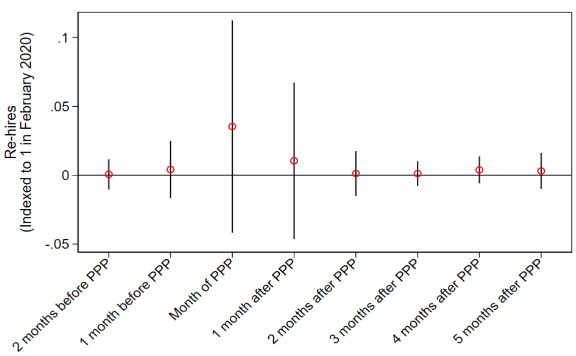

(f) Rehires: low remote capability

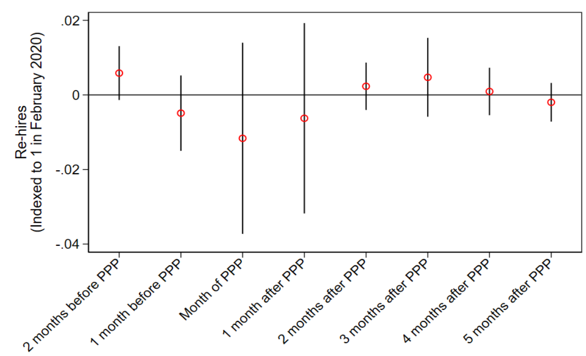

Notes: These figures show the results from the difference-in-differences specification in equation (3.2), I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The figures on the left show the results for firms which are in industries in which employees are less likely to be able to work from home. The figures on the right show the results for firms which are in industries in which employees are more likely to be able to work from home.The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals

Figure C.22: Effect of PPP on Employee Turnover, by Essential versus Non-Essential Businesses

(a) Total Hires: non-essential businesses

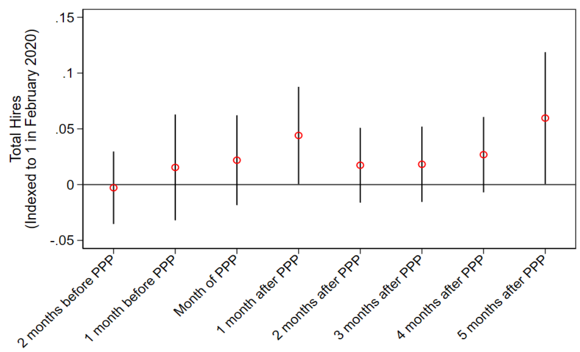

(b) Total Hires: essential businesses

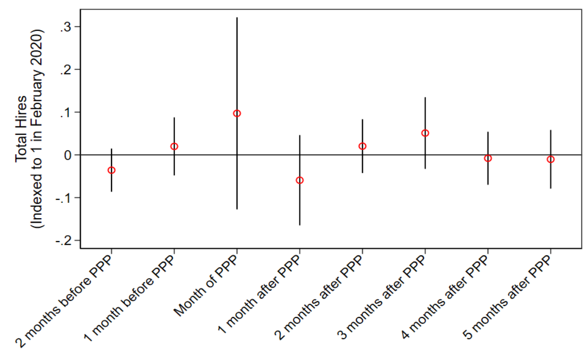

(c) New Hires: non-essential businesses

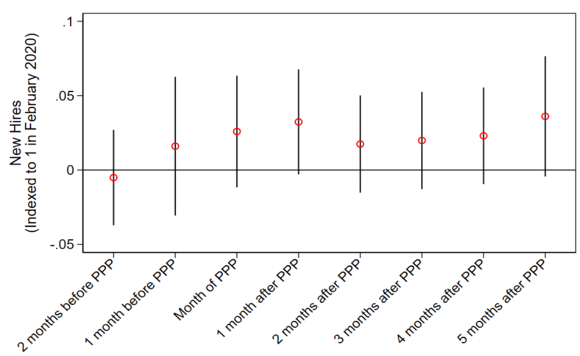

(d) New Hires: essential businesses

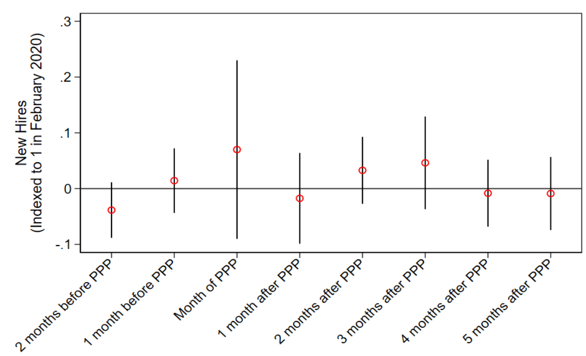

(e) Rehires: non-essential businesses

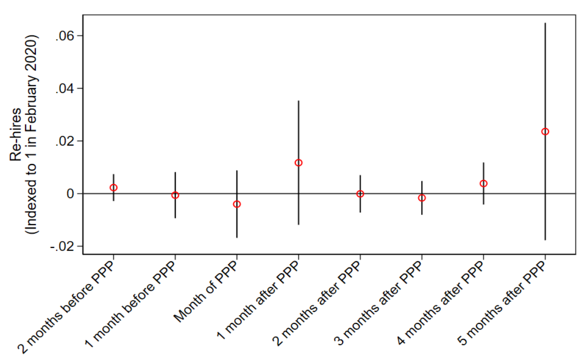

(f) Rehires: essential businesses

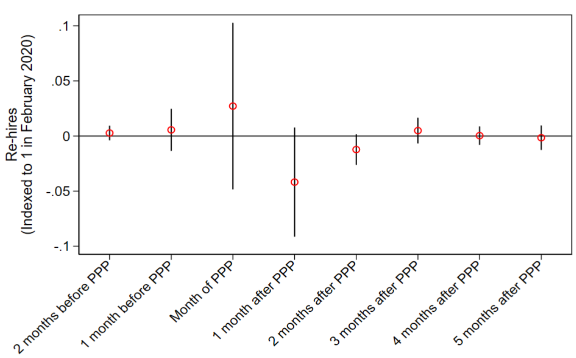

Notes: These figures show the results from the difference-in-differences specification in equation (3.2), I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The figures on the left show the results for firms which are typically classified as non-essential. The figures on the right show the results for firms which are typically classified as essential. The dependent variables are normalized to one in February 2020. Black lines represent 90% confidence intervals

Figure C.23: Effect of PPP on Employee Turnover, Average Effects

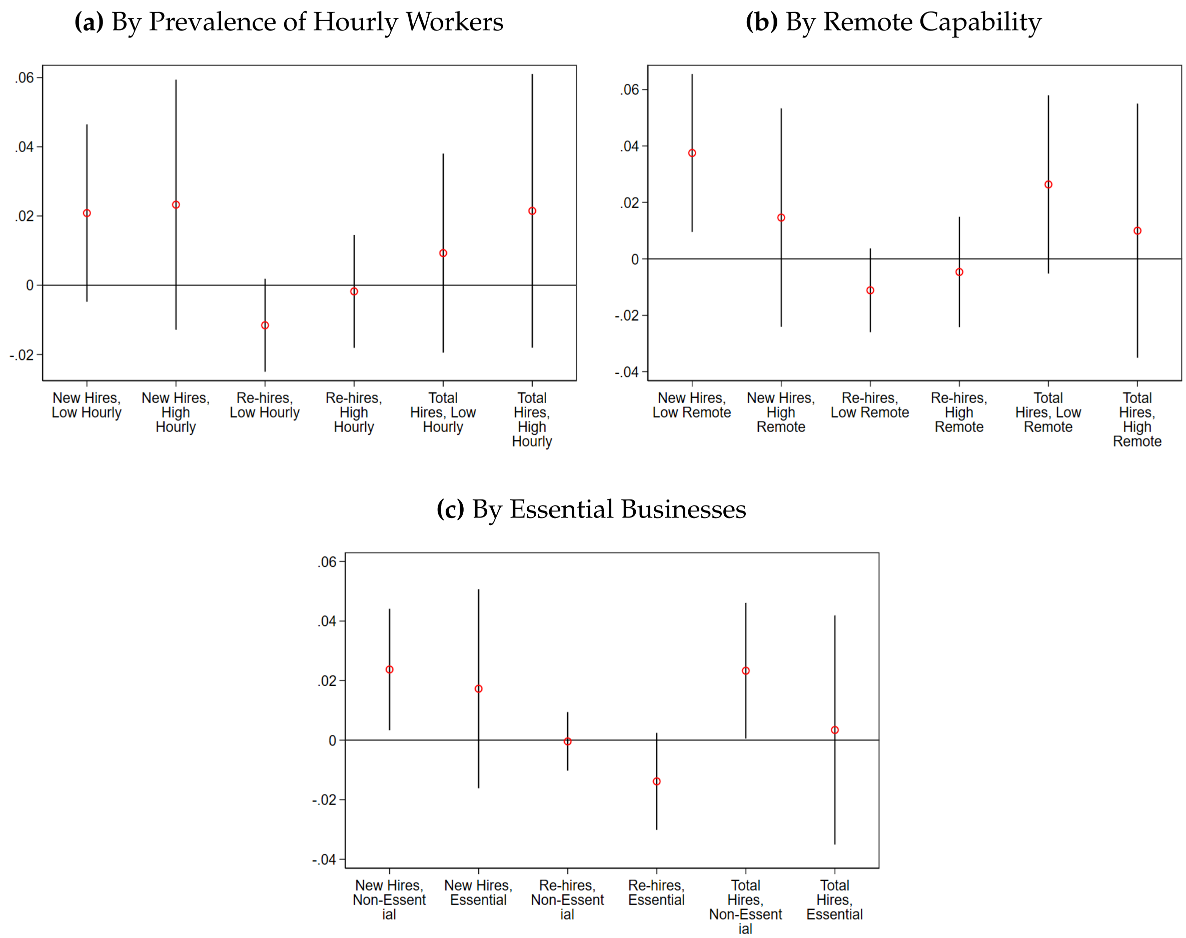

Notes: These figures show the results from the difference-in-differences specification in equation (3.2, split by different types of firms. I plot the coefficient on β_t, which traces the differences between firms that applied for PPP and firms that did not over time. The dependent variables are measured as fraction of the firm's base employment (average February 2019-February 2020). The black lines represent 90% confidence intervals.

C.2 Appendix Tables

Table C.1: NAICS Distributions of Sample Firms

	Applied for PPP?			
	No	Yes	All	All Small Businesses < 500 Employees
NAICS sectors (%)				
Construction	11.7	11.3	11.4	11.7
Manufacturing	5.5	6.1	5.9	4.1
Wholesale Trade	1.2	2.3	1.9	4.9
Retail Trade	3.7	4.7	4.3	10.7
Transportation/Warehousing	1.2	2.8	2.1	3.1
Finance/Insurance	9.8	4.7	6.9	4.0
Real Estate	8.0	4.2	5.9	5.2
Professional/Scientific/Tech	9.8	9.9	9.9	13.5
Administrative/Support/Waste	16.0	8.5	11.7	5.8
Health Care	21.5	25.4	23.7	10.9
Accommodation/Food	1.2	5.2	3.5	9.0
Other Services	5.5	11.3	8.8	11.6
Observations	163	214	377	

Notes: This table shows the 2-digit NAICS industry distribution for my sample firms. The first and second columns show the statistics for firms that did not apply and applied for PPP, respectively. The third column shows the full sample. One firm in the sample has over 500 employees and therefore is not eligible for the PPP, so it is dropped from all analyses. The final column shows a comparison for all firms with under 500 employees in U.S., which are calculated from the 2016 Statistics for U.S. Businesses (Census Bureau).

Table C.2: PPP Loan Terms and Forgiveness Rules, Pre-PPPFA and Post-PPPFA

	Loans issued:	
	March 27, 2020-June 4, 2020	June 5, 2020-August 8, 2020
Loan term	2 years	5 years
Interest Rate	1%	1%
% that must be used for payroll	75%	60%
Prorated forgiveness?	Yes	No
Covered period	8 weeks	24 weeks
FTE restoration deadline	June 30, 2020	December 31, 2020
Payment deferral period	6 months	6 months + administrative lag time

Notes: This table describes the rules on loan terms and forgiveness for the PPP prior to the Paycheck Protection Flexibility Act (June 5, 2020) and afterward. Note that the new rules on forgiveness apply retroactively to firms that applied before the PPPFA was enacted.

Table C.3: Take-up Regressions: Late versus Early appliers

	(1)	(2)	(3)	(4)	(5)
			Odds Ratios		
Log Base Employment	0.6387***	0.6471***	0.6588***	0.6916***	0.7390***
	(0.0664)	(0.0734)	(0.0614)	(0.0651)	(0.0784)
Log Average Salary (Base Period)	0.9621	0.9613	0.9147	0.8935	0.9928
	(0.4412)	(0.4370)	(0.3436)	(0.3402)	(0.4805)
Monthly Employment Growth (April 2020)		0.9670		1.4155*	0.9320
		(0.1846)		(0.2913)	(0.2033)
Average industry employment change during pandemic			11.1326	8.5239	
			(30.8182)	(23.6085)	
Bank in 2nd Tercile of Overall PPP Lending					1.3178
					(0.7744)
Bank in 3rd Tercile of Overall PPP Lending					1.5910
					(1.0898)
Industry FE?	Y	Y	N	N	Y
State FE?	Y	Y	Y	Y	Y
N	194	193	195	194	168
pseudo R^2	0.135	0.134	0.053	0.058	0.119

Notes: This table shows the results of a logistic regression of an indicator equal to one if the firm applied for PPP in the second round on a set of controls for firm characteristics (equation (3.1)). I drop all firms that did not apply, so this regression compares those that applied late to those that applied early. The odds ratios are reported. Log Base Employment is the log of the firm's average employment from February 2019-February 2020 (or the months in which the firm appears in that data set over that time period). Log Base Average Salary is the firm's average monthly wage per worker over the same time period. Monthly Employment Growth (April 2020) is the firm's log employment growth from March 2020 to April 2020. Average industry employment change during pandemic is the percentage decline in paid employment for the firm's two-digit NAICS code from February 15, 2020- April 25, 2020 (taken from Cajner et al. (2020)). 2nd and 3rd tercile of overall PPP lending are dummies equal to one if the firm's primary bank is in the second or third (highest) tercile, respectively, of share of PPP lending less than $150,000 in the firm's state. Standard errors, clustered at the 2-digits NAICS code level, in parentheses. *$p < 0.10$, **$p < 0.05$, *** $p < .01$

Table C.4: Take-up Regressions: Early Applicants

	(1)	(2)	(3)	(4)	(5)
			Odds Ratios		
Log Base Employment	2.1305***	2.0410***	2.1098***	1.9962***	1.8450***
	(0.1571)	(0.1344)	(0.1430)	(0.1234)	(0.1684)
Log Average Salary (Base Period)	1.1429	1.1918	1.1738	1.2445	1.1754
	(0.4427)	(0.4754)	(0.3685)	(0.3979)	(0.4392)
Monthly Employment Growth (April 2020)		0.7634		0.6276*	0.6942
		(0.2258)		(0.1644)	(0.2246)
Average industry employment change during pandemic			0.0897	0.1170	
			(0.2109)	(0.2798)	
Bank in 2nd Tercile of Overall PPP Lending					0.8779
					(0.6142)
Bank in 3rd Tercile of Overall PPP Lending					0.8552
					(0.5358)
Industry FE?	Y	Y	N	N	Y
State FE?	Y	Y	Y	Y	Y
N	327	324	342	339	276
pseudo R^2	0.196	0.195	0.149	0.153	0.160

Notes: This table shows the results of a logistic regression of an indicator equal to one if the firm applied during the first round for the PPP on a set of controls for firm characteristics (equation (3.1)). The odds ratios are reported. Log Base Employment is the log of the firm's average employment from February 2019-February 2020 (or the months in which the firm appears in that data set over that time period). Log Base Average Salary is the firm's average monthly wage per worker over the same time period. Monthly Employment Growth (April 2020) is the firm's log employment growth from March 2020 to April 2020. Average industry employment change during pandemic is the percentage decline in paid employment for the firm's two-digit NAICS code from February 15, 2020- April 25, 2020 (taken from Cajner et al. (2020)). 2nd and 3rd tercile of overall PPP lending are dummies equal to one if the firm's primary bank is in the second or third (highest) tercile, respectively, of share of PPP lending less than $150,000 in the firm's state. Standard errors, clustered at the 2-digits NAICS code level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < .01$

Table C.5: Take-up Regressions: Late Applicants

	(1)	(2)	(3)	(4)	(5)
			Odds Ratios		
Log Base Employment	1.4949**	1.5108**	1.5551***	1.5578***	1.6742**
	(0.2432)	(0.2511)	(0.2034)	(0.2040)	(0.4349)
Log Average Salary (Base Period)	1.5826**	1.5386**	1.4354*	1.3990	1.4662
	(0.2976)	(0.3004)	(0.2918)	(0.2890)	(0.3576)
Monthly Employment Growth (April 2020)		1.3190		1.1279	1.1754
		(0.5279)		(0.3403)	(0.4801)
Average industry employment change during pandemic			0.5441	0.5562	
			(0.8138)	(0.8322)	
Bank in 2nd Tercile of Overall PPP Lending					1.0090
					(0.3385)
Bank in 3rd Tercile of Overall PPP Lending					1.7539**
					(0.4559)
Industry FE?	Y	Y	N	N	Y
State FE?	Y	Y	Y	Y	Y
N	262	260	259	257	223
pseudo R^2	0.096	0.094	0.064	0.061	0.116

Notes: This table shows the results of a logistic regression of an indicator equal to one if the firm applied during the second round of the PPP (excluding firms that applied during the first round) for the PPP on a set of controls for firm characteristics (equation (3.1)). The odds ratios are reported. Log Base Employment is the log of the firm's average employment from February 2019-February 2020 (or the months in which the firm appears in that data set over that time period). Log Base Average Salary is the firm's average monthly wage per worker over the same time period. Monthly Employment Growth (April 2020) is the firm's log employment growth from March 2020 to April 2020. Average industry employment change during pandemic is the percentage decline in paid employment for the firm's two-digit NAICS code from February 15, 2020- April 25, 2020 (taken from Cajner et al. (2020)). 2nd and 3rd tercile of overall PPP lending are dummies equal to one if the firm's primary bank is in the second or third (highest) tercile, respectively, of share of PPP lending less than \$150,000 in the firm's state. Standard errors, clustered at the 2-digits NAICS code level, in parentheses. *$p < 0.10$, **$p < 0.05$, *** $p < .01$

Table C.6: Take-up Regressions: Controlling for Average Yearly Growth

	(1)	(2)
		Odds Ratios
	All Applicants	Late versus early applicants
Log Base Employment	1.9531**	0.8089
	(0.5515)	(0.1176)
Log Average Salary (Base Period)	1.1728	0.9748
	(0.2301)	(0.5211)
Monthly Employment Growth (April 2020)	0.9001	2.8864**
	(0.4133)	(1.4994)
Average Yearly Employment Growth (2020, pre-pandemic)	1.1299	0.6612
	(0.2314)	(0.4832)
Bank in 2nd Tercile of Overall PPP Lending	0.9442	1.9074
	(0.6246)	(1.3279)
Bank in 3rd Tercile of Overall PPP Lending	1.3629	1.6017
	(0.4872)	(1.5099)
Industry FE?	Y	Y
State FE?	Y	Y
N	194	115
pseudo R^2	0.118	0.142

Notes: This table shows the results of a logistic regression of an indicator equal to one if the firm applied for the PPP on a set of controls for firm characteristics (equation (3.1)). The first column shows the full sample. The second column shows the late applicants, relative to the early applicants only. The odds ratios are reported. Log Base Employment is the log of the firm's average employment from February 2019-February 2020 (or the months in which the firm appears in that data set over that time period). Log Base Average Salary is the firm's average monthly wage per worker over the same time period. Monthly Employment Growth (April 2020) is the firm's log employment growth from March 2020 to April 2020. Average industry employment change during pandemic is the percentage decline in paid employment for the firm's two-digit NAICS code from February 15, 2020- April 25, 2020 (taken from Cajner et al. (2020)). 2nd and 3rd tercile of overall PPP lending are dummies equal to one if the firm's primary bank is in the second or third (highest) tercile, respectively, of share of PPP lending less than $150,000 in the firm's state. Standard errors, clustered at the 2-digits NAICS code level, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, *** $p < .01$

Table C.7: Take-up Regressions: Balanced Sample

	(1)	(2)	(3)	(4)	(5)
			Odds Ratios		
Log Base Employment	1.7639***	1.6562**	1.8510***	1.7571***	1.9297*
	(0.3246)	(0.3456)	(0.2861)	(0.3014)	(0.7037)
Log Average Salary (Base Period)	1.2850	1.3785	1.3863	1.4535	1.1279
	(0.4757)	(0.4166)	(0.4817)	(0.4360)	(0.3284)
Monthly Employment Growth (April 2020)		0.5862*		0.5865	0.4206
		(0.1768)		(0.1909)	(0.2383)
Average industry employment change during pandemic			0.6416	1.0688	
			(0.9174)	(1.6979)	
Bank in 2nd Tercile of Overall PPP Lending					0.8017
					(0.5201)
Bank in 3rd Tercile of Overall PPP Lending					1.2698
					(0.4382)
Industry FE?	Y	Y	N	N	Y
State FE?	Y	Y	Y	Y	Y
N	206	206	211	211	175
pseudo R^2	0.110	0.116	0.081	0.086	0.135

Notes: This table shows the results of a logistic regression of an indicator equal to one if the firm applied for PPP on a set of controls for firm characteristics (equation (3.1)). This robustness check includes only the balanced panel of firms that have been in the data since January 2019. The odds ratios are reported. Log Base Employment is the log of the firm's average employment from February 2019-February 2020 (or the months in which the firm appears in that data set over that time period). Log Base Average Salary is the firm's average monthly wage per worker over the same time period. Monthly Employment Growth (April 2020) is the firm's log employment growth from March 2020 to April 2020. Average industry employment change during pandemic is the percentage decline in paid employment for the firm's two-digit NAICS code from February 15, 2020- April 25, 2020 (taken from Cajner et al. (2020)). 2nd and 3rd tercile of overall PPP lending are dummies equal to one if the firm's primary bank is in the second or third (highest) tercile, respectively, of share of PPP lending less than $150,000 in the firm's state. Standard errors, clustered at the 2-digits NAICS code level, in parentheses. *$p < 0.10$, **$p < 0.05$, *** $p < .01$

Table C.8: Hourly Workers by Industry

	Number of Hourly Employees	Number of Employees	Percentage of hourly employees
Manufacturing	9,094	12,688.7	71.7
Retail Trade	11,346	15,833.1	71.7
Mining/Quarrying/Oil and Gas Extraction	475	683.3	69.5
Construction	4,807	7,289.3	65.9
Transportation/Warehousing	3,395	5,419.1	62.6
Accommodation/Food	9,828	6,348.5	60.1
Educational and Health Services	13,283	23,666.5	56.2
Other Services	3,254	6,622.4	49.1
Finance/Insurance	3,499	8,568.8	40.8
Agriculture/Forestry/Fishing /Hunting	825	2,310	35.7
Professional Services	6,336	20,999.5	30.2
Wholesale Trade	1,675	5,852.5	28.6

Notes: This table show the number and percentage of employees in a given industry that are paid on an hourly basis. Employment numbers are in thousands. Data are from the Bureau of Labor Statistics.
https://www.bls.gov/emp/tables/employment-by-major-industry-sector.htm
https://www.bls.gov/opub/reports/minimum-wage/2015/home.htm

Table C.9: Correlations between Remote, Hourly, and Essential Workers

	Highly remote	Highly hourly	Essential
Highly remote	1.00	-0.77	0.08
Highly hourly	-0.77	1.00	-0.16
Essential	0.08	-0.16	1.00