Breaking the Implicit Contract: Using Pension Freezes to Estimate the Labor Supply Elasticity *

Dhiren Patki †

This version: July 2019

PRELIMINARY AND INCOMPLETE

Abstract

Although the labor market is characterized by long-term associations between employers and employees, the labor supply literature frequently models employer-employee interactions in the context of competitive spot markets. This study exploits firm's decisions to renege on implicit pension promises to study the labor supply behavior of older workers in the context of long-term employer-employee relationships. Combining rich administrative data and survey data, it finds an extensive margin labor supply elasticity of 0.47, which is two times larger than existing estimates. The findings in this paper illustrate the importance of long-term contracts in explaining labor supply behavior and provide evidence about the effectiveness of policies aimed at lengthening the working lives of older Americans.

JEL codes: J22, J26, J32, J41, M52

*I am indebted to Matthew Shapiro, Charles Brown, and Melvin Stephens for advice and support. I thank John Bound, R. Andrew Butters, Jamie Fogel, Parag Mahajan, and seminar participants at the University of Michigan for valuable comments. I am grateful to Kristin McCue for her help in understanding the data and assisting with the Census disclosure review process. This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. I acknowledge support by the CenHRS project grant to the University of Michigan from the Alfred P. Sloan Foundation, the Alfred P. Sloan Foundation Pre-doctoral Fellowship on the Economics of an Aging Workforce, awarded through the NBER, and the F. Thomas Juster Economic Behavior Research Award. All errors are my own.

†Department of Economics, University of Michigan, 611 Tappan Ave., Ann Arbor, MI 48109. Email: dpatki@umich.edu.

1 Introduction

Long-term relationships between employers and employees are ubiquitous in the labor market. These relationships can aid in the development of skills through on-the-job training, facilitate risk sharing, and be used to incentivize effort over long horizons (see, e.g., Becker (1962), Rosen (1985), and Lazear (1979)). In the literature, long-term associations between employers and employees have been regarded as implicit contracts because they are self-enforced through reputation rather than through explicit legal means.

A key distinction between implicit contracts and spot markets relates to how employers compensate employees. Unlike spot markets where compensation reflects worker productivity in each period, compensation with implicit contracts reflects cumulative worker productivity over the duration of the contract. As shown in Becker (1962), with long-term relationships, it is the present value of compensation over the duration of the relationship that determines workers employment choices. Indeed, holding the present value of the contract fixed, the time path of observed compensation can take any form without affecting employment allocations.

The difference between spot markets and implicit contracts is materially important for inferring how labor supply behavior responds to changes in compensation. To appreciate this distinction, consider that the spot market view attributes fluctuations in labor supply to observed fluctuations in compensation earned in each period. In contrast, in the implicit contract view, changes in the time path of observed compensation may end up having no real effect employment if the present value of the contract remains fixed. As a result, when employer-employee associations are characterized by implicit contracts, behavioral elasticities are identified by variation in the contract price of labor rather than the spot market price of labor.

Despite the fact that much of the labor market is governed by implicit long-term contracts, the literature estimating labor supply elasticities is based almost entirely on the spot market perspective.[1] One reason for this shortcoming is that typical survey and administrative data sets do not easily admit classification of jobs into implicit contract types or spot market types. Second, it is difficult to ascertain the extent to which implicit contracts are about training, risk sharing, or encouraging higher effort over long careers. In part this is because important aspects of implicit contracts such as work hours expectations, incentive payments, expected contract duration, and other contingencies are not directly observable in conventional data sources.

[1] Beaudry and DiNardo (1995), who examine risk sharing contracts, is a notable counterexample. Also see Abowd and Card (1987) and Abowd and Card (1989) who cast doubt on the spot market framework by showing that much of the variation in earnings and hours occurs at fixed wage rates.

Finally, identification of labor supply elasticities in the implicit contract setting requires plausibly exogenous variation in the present value of future compensation for a given employer-employee match that is difficult to isolate.

Defined benefit (DB) pensions have long been recognized as a type of implicit contract in which workers are paid less than their marginal product when young and promised more than their marginal product when old (see, e.g., Lazear (1984), Ippolito (1985), and Lazear and Moore (1988)). This backloaded compensation structure generates incentives for workers to invest in employer-specific human capital and can elicit more effort over the duration of a worker's career with an employer than a spot market arrangement might. Amidst the overall decline in DB pension provision in the United States, recent years have seen employers reneging on implicit contracts by eliminating the promise of backloaded DB pension accruals and replacing them with flatter defined contribution (DC) accruals. These actions, known as hard freezes, have the effect of unexpectedly changing the contract price of labor. For workers late in their careers, hard freezes reduce the amount of compensation that they had counted on earning at retirement age. As such, pension freezes provide a setting to examine how the labor supply behavior of older workers responds to employer-driven ruptures of implicit contracts.

To study why employers have increasingly reneged on implicit contracts and to examine how employees respond to these changes, I create a new dataset that combines information from Internal Revenue Service (IRS) administrative records on the universe of private sector pension plans with longitudinal employer-employee linked data from the Census Bureau. I exploit a quasi-experimental research design pooling together thousands of firm-level experiments and comparing the labor supply of workers whose employers freeze their DB plans with workers whose employers keep their DB plans intact.

On the firm side, the data reveal that freezes are driven primarily by increases in the level and volatility of DB pension costs. On the worker side, the data reveal that labor supply responses to these shocks vary in two important dimensions. First, workers initially affected at different ages exhibit differing labor supply responses because they experience compensation changes of varying magnitude. Second, holding age and the value of lost compensation fixed, some workers respond to freezes by retiring early while others respond by retiring later. Early retirements reflect the dominance of substitution effects while delayed retirements reflect the dominance of wealth effects. Workers' dynamic response to pension freezes therefore reveals heterogeneity in their preferences for longer careers.

The administrative microdata that I use describe behavioral responses to pension freezes (the numerator of the labor supply elasticity). These data do not include information on the dollar value of compensation changes induced by pension freezes (the denominator of the labor supply elasticity). To estimate elasticities, I turn to rich data from the Health and Retirement Study (HRS) that include precise measures of pension wealth as well as subjective information about retirement expectations for a representative sample of older, DB eligible, workers.[2] Using these data, I simulate the effect of pension freezes and estimate how they change the contract price of labor. I then divide observed behavioral changes by simulated price changes to recover the labor supply elasticity. I find that the two-sample implicit contract-based extensive margin labor supply elasticity is 0.47 for 56 to 64 year-old workers. For men in the same age group, the estimate is 0.34; for women it is 0.63. These estimates do not account for wealth effects and therefore represent uncompensated elasticities. Nevertheless, the pension freeze-based elasticities are two times larger than prior estimates thereby underscoring the importance of long-term contracts in better explaining labor supply behavior.[3]

This paper makes three contributions. First, it adds to the literature that has sought to estimate extensive margin labor supply elasticities using microdata and quasi-experimental methods.[4] Its chief contribution to this body of work is to re-frame the association between employees and employers as being governed by implicit contracts rather than spot markets and to estimate the labor supply elasticity within this richer framework. Second, because the literature on labor supply has studied prime-age workers almost exclusively, there is a paucity of estimates on the labor supply elasticity of older workers. By focusing on the labor supply choices of older Americans, this paper adds to a small set of available estimates for an important subpopulation (see, e.g., Gelber et al. (2017)). Lastly, by highlighting the impact of shifts from DB to DC pensions that occur within rather than between firms, this paper sheds new light on the changing pension landscape in the United States.

The findings in this paper have implications for the macroeconomic, public finance, and labor market implications of an aging workforce including the effectiveness of tax policy in influencing the continued labor force participation of older workers, an improved understanding of age-based variation in Social Security claiming, the impact of the shift from DB to DC pensions on labor supply behavior, and the role of

[2]The incidence of pension freezes in the HRS is too low to study systematically.

[3]The average of compensated (Hicksian) extensive margin elasticity estimates in the meta-analysis of Chetty et al. (2012) is 0.25.

[4]See, e.g., Eissa and Liebman (1996), Meyer and Rosenbaum (2001), Devereux (2004), Eissa and Hoynes (2004), Blundell et al. (2011) for Hicksian elasticities and Card and Hyslop (1995), Carrington (1996), Bianchi et al. (2001), Manoli and Weber (2011), and Brown (2013) for Frisch elasticities.

long-term employers in supporting longer working lives.

The remainder of this paper is structured as follows. Section 2 provides institutional details on the DB pension landscape in the United States and explains important recent changes influencing firm's decisions to continue sponsoring DB plans. Section 3 outlines a model of the worker's response to a freeze, and makes several testable predictions. Section 4 discusses three different sources of administrative data and provides details about HRS survey data. Section 5 outlines the empirical framework and presents summary statistics. Section 6 interprets individual level outcomes estimated using administrative data. Section 7 combines labor supply responses to freezes obtained from the administrative data with simulated estimates of compensation changes from the HRS to estimate behavioral elasticities. Section 8 concludes.

2 Institutional Background

In 1980 DB plans covered 61 percent of pension eligible private sector workers. By 2015 the same statistic had fallen to 16 percent. This overall decline occurred not only because of a surge in new 401(k) DC plans but also because of stagnation in the creation of new DB plans. For most of this period, DB plans continued to operate normally with only rare instances of distressed plan terminations triggered by firm bankruptcy.[5] Starting in the late 1990's, however, several firms began converting their DB plans to cash balance (CB) plans. CB conversions switched DB accruals away from formulas based on years of service and earnings, to account based plans that provided participants a market determined rate of return on a fixed proportion of earnings.[6] In the early 2000's this shift was amplified as many employers altogether ended traditional DB pension accruals in actions known as hard freezes.[7] The incidence of CB conversions and hard freezes between 1999 and 2015 is shown in Figure 1. The most recent data indicate that about half of all private sector single employer DB plans have either been converted to CB plans or hard frozen, affecting about 40 percent of active participants.

[5]The Omnibus Budget Reconciliation Act of 1987 (OBRA-87) reduced the ability of DB sponsoring employers to take tax deductions for pension contributions. This change led to a spike in non-distress terminations between 1987 and 1990 (see, e.g., Table A-9 in Pension Insurance Data Book 1996, PBGC Single Employer Program `https://www.pbgc.gov/documents/1996databook.pdf`).

[6]CB plans are functionally very similar to DC plans: they provide participants with individual accounts whose value is tied to market rates of return. Individual CB accounts are hypothetical rather than real, however, so CB plans are subject to the same funding obligations required of DB plans and are legally treated as DB plans under ERISA.

[7]Over the same period, employers also began closing their DB plans to new entrants in actions known as soft-freezes.

2.1 Costs of DB plan provision become increasingly volatile

Firm's decisions to renege on DB promises through CB conversions and hard freezes occurred in a deteriorating financial environment that increased the volatility of DB pension costs. As shown in Figure 2, the wake of the dot-com bubble and ensuing recession lowered the asset value of pension funds substantially. In 2000, the private sector DB system had $1.87 in assets for each dollar of future liabilities. By 2004, the funding ratio had slipped to just 0.62. As the funding position of DB plans worsened, statutory provisions required firms to increase annual contributions; consequently, aggregate payments into DB pension funds rose five-fold from $26 billion in 2000 to $126 billion in 2003. Low interest rates and stock market losses during the Great Recession weakened firms' funding positions drawing further increases in required pension contributions. The role of worsening plan finances as a key predictor of subsequent pension freezes is demonstrated greater detail in Appendix B, which relies on firm-level microdata.[8]

In the midst of major changes to the finances of DB pension funds, the Pension Protection Act (PPA) was signed into law in 2006. The PPA established more conservative standards on the interest rates that sponsors could use to discount future liabilities, reduced the period over which sponsors needed to amortize funding deficits from 15 years to 7 years, and required that plans with funding ratios of 80 percent or below make additional minimum contributions and pay higher insurance premiums to the Pension Benefit Guarantee Corporation (PBGC).[9] These key provisions of the PPA, which were phased-in as of 2008, raised statutory minimum pension contributions and imposed greater cost pressure on DB sponsors.

2.2 Legal constraints alleviated

Between 1998 and 2000, a handful of prominent employers had converted traditional DB plans to CB plans. Older employees, who stood to lose substantial pension accruals as a result these transitions, brought class action lawsuits against their employers claiming that CB conversions violated the age discrimination provisions of the Employee Retirement Income Security Act of 1974 (ERISA) (see, e.g., Zelinksy (2000)). As these cases played out in the court system, the legality of CB conversions remained uncertain. The threat

[8] Munnell and Soto (2007), Bovbjerg et al. (2008), and Rauh et al. (2017) also find that cost savings and funding volatility are important determinants of firms' decisions to freeze their DB plans. An internal compliance investigation on frozen DB plans conducted by the IRS in 2012 and obtained through a Freedom of Information Act (FOIA) request for this research project further corroborates this conclusion finding that sponsors typically froze their plans due to funding deficiencies.

[9] Some of the PPA's requirements were relaxed during the Great Recession as pension sponsors sought relief from strict funding targets. For instance, The Preservation of Access to Care for Medicare Beneficiaries and Pension Relief Act of 2010 allowed firms to elect extended amortization periods for any two plan years between 2008 and 2011. The extensions were for 9 or 15 years rather than the 7 year requirement of the PPA.

of litigation along with large potential settlement costs for class action lawsuits likely constrained other employers from converting traditional DB plans. In 2006, however, a Federal appeals court ruled that IBM's CB conversion was age-neutral, thereby ending uncertainty surrounding the legality of what was seen by many as an important test for restructuring traditional DB plans.[10] In the same year, and perhaps more importantly, the new PPA law provided guidelines that CB conversions needed to meet in order to be considered age-neutral. Together, the appeals court ruling and the PPA's new provisions gave willing employers the legal cover they needed to restructure their DB plans.

2.3 The public sector remains largely unaffected

The economic forces affecting private sector DB sponsors also affected sponsors in the public sector where DB pensions are more common. Indeed, as of 2017, unfunded pension liabilities in the public sector stood at \$3.84 trillion, substantially larger than the \$0.76 trillion funding gap in the private sector.[11] Despite this stark difference in funding, however, hard freezes are exceedingly rare in the public sector affecting less than 0.5 percent of eligible state and local government workers.[12] The reason for the difference in the incidence of hard frozen pensions between the private and public sectors lies in differing legal protections. ERISA, which governs private pensions, explicitly allows sponsors to modify future pension rules provided that earned pension benefits are not reduced. In the public sector, which is not subject to the provisions of ERISA, several state governments protect both earned and unearned future pension benefits.[13]

During recent periods of fiscal strain, prominent city and state governments including Detroit and New Jersey have successfully challenged some of the protections afforded to public sector pension accruals.[14] Perhaps the most relevant challenge stems from California's Public Employee Pension Reform Act of 2012 which would allow public pension sponsors to modify unearned future accruals. The law has been chal-

[10] See, e.g., Mary Williams Walsh, "Issues Left Unresolved on Pensions," *New York Times*, Jan. 17, 2007.

[11] See, e.g., Rauh (2017) for public sector estimates and Pension Benefit Guarantee Corporation (PBGC) for private sector estimates `https://www.pbgc.gov/prac/data-books`. Multiemployer pension plans (i.e. plans that are formed between multiple employers and potentially multiple labor unions) account for the bulk of funding deficiencies in the private DB system as of 2017.

[12] See Bureau of Labor Statistics (BLS), National Compensation Survey data from 2017. These data indicate that 56 percent of state and local government workers were in DB plans closed to new entrants (i.e. soft frozen).

[13] The protection of unearned future benefits is based on a set of legal precedents collectively referred to as the "California rule." The rule prevails not only in California but also, with varying degrees of stringency in Alaska, Colorado, Idaho, Kansas, Massachusetts, Nebraska, Nevada, Oklahoma, Oregon, Pennsylvania, Vermont, and Washington. See, e.g., Monahan (2012).

[14] A hard freeze of legacy pension accruals was one among the many provisions that Detroit negotiated after it filed for bankruptcy in 2013 (see, e.g., Urahn et al. (2018)). Notably, unearned pension accruals are not protected by Michigan law. In a 2016, the New Jersey Supreme Court upheld a 2011 state law that froze cost of living adjustments (COLAs) for retirees covered by the state's pension plan. Plaintiff's had argued that COLAs were subject to the same protections as pension accruals themselves.

lenged in two separate cases which are currently being reviewed by the California Supreme Court.[15]

In the remainder of this paper, I focus my attention on CB conversions and hard freezes in the private sector. For the sake of brevity, I jointly refer to both types of changes as "freezes" unless otherwise noted. Having laid out the institutional context in which freezes occur, I turn next to their effect on individual labor supply decisions using a option value model of labor supply and asset accumulation.

3 Theoretical Framework

3.1 Defined benefit pensions as implicit contracts

Starting from Becker (1962), a number of studies have emphasized that total compensation paid to a worker may not equal her marginal product in every period. By paying a worker less than her marginal product when young and more than her marginal product when old, employers can elicit higher levels of effort (or obtain the same level of effort with less monitoring), benefit from greater retention, and obtain more investment in firm specific human capital throughout the worker's career (see, e.g., Lazear (1979), Lazear (1984) and Lazear and Moore (1988)).

This pattern in lifetime compensation is particularly evident with DB pension plans which defer the largest pension accruals to the last phase of a worker's career — a practice commonly known as backloading. The left panel of Figure 3 shows evidence of backloading using detailed information from pension plans linked to DB eligible workers in the HRS. The dashed blue line shows the age progression of non-deferred compensation, while the solid blue line shows compensation gross of pension accruals. Accrual spikes at salient ages (55, 59, 62, and 65) contribute substantially to the backloaded nature of total compensation.

A corollary of this view of the employer-employee relationship is that DB pensions function like implicit contracts: workers give up compensation when young in return for the unenforceable promise that their employer will pay them the full value of their marginal product by the time they retire. In the implicit contract view of DB pensions, workers *expect* to be retained until retirement and they *expect* to earn the full value of pension benefits that they would be entitled to at retirement.[16]

In the following subsection I outline a model of saving and retirement timing that draws on Stock and

[15] The relevant cases are Marin Assoc. of Public Employees v. MCERA and Cal Fire Local 2881 v. CalPERS.

[16] See, e.g., Ippolito (1985) and Kotlikoff and Wise (1985) for theory and evidence on the implicit contract view of DB pensions. An alternative view of DB pensions is the legal theory wherein workers do not form long term expectations of pension promises (see, e.g., Bulow (1982)).

Wise (1990a), Stock and Wise (1990b), and Lumsdaine et al. (1992) and emphasizes the implicit contract view of DB pensions. I use the model to characterize the effect of pension freezes on worker behavior and enumerate four empirically testable implications.

3.2 Model of labor supply and asset accumulation

Worker earnings in period t are given by y_t which is taxable in the period that it is earned. Workers may elect to defer a fraction m_t^w of their earnings towards a DC account on a pre-tax basis. Elective deferrals to DC accounts and DB pension accruals are not subject to taxes until that income is claimed in retirement. Total compensation earned in each period can be written as follows

$$\mathcal{C}_t = \underbrace{y_t(1 - m_t^w)(1 - \tau(y_t, m_t^w))}_{\text{non-deferred compensation}} + \underbrace{\Delta W_t^{DC}}_{\text{DC deferral}} + \underbrace{\Delta W_t^{DB}}_{\text{DB deferral}} \tag{1}$$

where $\tau(y_t, m_t^w)$ is tax rate applied to non-deferred income.[17] DC wealth evolves according to

$$W_t^{DC} = W_{t-1}^{DC}(1 + r) + y_t(m_t^w + m^e(m_t^w)) \tag{2}$$

$$\bar{C} \geq y_t(m_t^w + m^e(m_t^w)) \tag{3}$$

where r is the interest rate and $m^e(m_t^w)$ is the fraction of earnings that the employer contributes to the DC plan. m^e is expressed as a function of m_t^w reflecting commonly used employer incentives for participation in DC plans. IRS rules limit total contributions to be less than a threshold value $\bar{C}$.[18]

DB pension wealth is earned passively as a deterministic function of earnings and tenure and is paid out upon retirement as an annuity whose period t value is given by b_t^{DB}. DB wealth, W_t^{DB}, represents the present value of future annuity payments. Workers accumulate Social Security wealth based on their earnings history in the form of an annuity whose period t value is given by b_t^{SS}. Like the DB annuity, Social Security can only be claimed in retirement. $b_t = b_t^{DB} + b_t^{SS}$ is the total annuity income for an individual who retires in period t. In addition to pension assets, workers can accumulate non-pension wealth which is denoted by A_t.

[17] Income tax can be reduced by elective deferrals to DC accounts but payroll taxes apply to all earned income up to the relevant earnings caps.

[18] In 2010 the total contribution limit was $49,000 with workers over age 50 being allowed to make an extra $5,500 in catch-up contributions. These values are updated annually by the IRS based on cost of living adjustments. See https://www.irs.gov/pub/irs-tege/cola_table.pdf

Individuals face stochastic disutility from work, g, which is drawn from a distribution G. This random variable reflects individual preferences for labor supply and is influenced by health, disability status, and non-pay aspects of employer-employee match quality. Casting the workers problem this way abstracts away from the intensive margin labor supply choice and focuses purely on the participation margin.

3.2.1 Value of retirement

Denote period t variables by v and period $t+1$ variables by v'. The value of retirement is given by

$$V^R(b, A, W^{DC}) = \max_{c, W^{DC'}} \left\{ u(c) + \beta V^R(b, A', W^{DC'}) \right\} \tag{4}$$

$$\text{s.t. } c \leq \left(b + W^{DC} - \frac{W^{DC'}}{1+r} \right)(1-\tau) + A - \frac{A'}{1+r(1-\tau_k)} \tag{5}$$

where $u(\cdot)$ is a concave function and c is consumption. τ reflects the relevant tax rate on retirement income, whereas τ_k reflects the tax rate on income earned from non-pension wealth. In this framework, retirement is a self-absorbing state. While this assumption excludes the possibility of short-term bridge employment, I provide empirical evidence in Section 6.1, that workers affected by pension freezes tend to follow a once-and-for-all retirement pattern.

3.2.2 Value of working

The value of working is

$$V^W(y, g, A) = \max_{c, m^w} \left\{ u(c) - g + \beta E \left[\max \left\{ V^R(b', A', W^{DC'}), V^W(y', g', A') \right\} \right] \right\} \tag{6}$$

$$\text{s.t. } c \leq y(1-m^w)(1-\tau) + A - \frac{A'}{1+r(1-\tau_k)}. \tag{7}$$

The expectation operator integrates over the random variable g which is the only source of uncertainty in the model. The budget constraint imposes the assumption that pension wealth is illiquid until retirement. τ reflects the relevant tax rate on non-deferred income, whereas τ_k reflects the tax rate on income earned from non-pension wealth.

3.2.3 Optimal retirement decision

A worker retires when $V^R(b, A, W^{DC}) \geq V^W(y, g, A)$. This condition defines a cutoff value $\bar{g}$ so that any draw of $g \geq \bar{g}$ will lead to retirement. Solving for $\bar{g}$ yields

$$\bar{g} = u(c^W) - u(c^R) + \beta \left[\underbrace{\overbrace{E\left[\max\left\{V^R(b', A', W^{DC'}), V^W(y', g', A')\right\}\right]}^{\text{1: Continuation value of working}} - \overbrace{V^R(b, A, W^{DC})}^{\text{2: Continuation value of retiring}}}_{\text{Option value of working one more period}} \right] \tag{8}$$

where c^W and c^R are the optimal consumption choices in the work and retirement states respectively. Consider terms 1 and 2 on the right side of equation (8). Higher future earnings and pension accruals raise term 1 and increase the incentive for continued work. On the other hand, higher retirement wealth (b and W^{DC}) raises term 2, thereby generating an incentive to retire. These two offsetting incentives are key determinants of the optimal retirement decision, which is summarized by the cutoff value $\bar{g}$.

In this framework, the decision about whether to continue working or to retire is determined by the option value channel which embodies the implicit contract of continued employment with the firm. Because DB pension wealth is backloaded and cannot be ported between employers, the option value of working has a particularly strong firm-specific component. As a result, labor supply varies not due to period-by-period changes in flow compensation, as in a spot market, but due to changes to the long-term contractual value of employment. This distinction is of substantive importance in the definition and estimation of labor supply elasticities which I take up in Section 7.

3.2.4 The effect of a pension freeze is theoretically ambiguous

Treating DB pensions as implicit contracts, workers expect to be paid the full value of their marginal product over the duration of their career. If employers renege on this implicit bargain by freezing DB pensions, workers face changes to the total value of compensation that they had previously expected to earn by retirement. These within-employer changes in compensation are simulated using data from the HRS in the right-hand panel of Figure 3. The solid blue line shows how total compensation (earnings plus deferred compensation) varies with age.[19] The dashed red lines simulate the effect of a DB pension freeze followed by an immediate transition to a DC plan for workers aged 52, 58, and 64 respectively. As the graphs show, pension freezes

[19] Earnings and pension benefits are computed for each individual at each age, so the age-based variation in compensation reflects plan-specific formulas for accruals and not changes in sample composition due to retirement.

alter the path of total compensation in different ways based on the age at which workers experience them. The average 52 year-old experiences large initial losses in compensation followed by gains after age 65; the average 58 year-old experiences smaller initial losses followed by gains after age 65; finally, the average worker 65 or older experiences unambiguous gains. The reason that pension freezes have positive effects on compensation over the long-term is because they eliminate negative real accruals that are common in DB plans after age 65.[20]

Having illustrated how pension freezes alter the age-compensation path, I use the model to investigate their impact on retirement behavior. Consider terms 1 and 2 on the right side of equation (8) for a worker under 65 when their DB plan is frozen. If the plan is frozen in period t_F then $b_t^{DB} = b_{t_F}^{DB}\ \forall\, t > t_F$. This change causes term 1 to drop and generates a substitution effect away from work toward retirement. The substitution effect translates to a lower value of $\bar{g}$. Nevertheless, the fact that b^{DB} is frozen causes term 2 to fall in subsequent periods. The reduction in term 2 induces a wealth effect that raises $\bar{g}$ and lowers the value of retirement relative to the value of working. The opposite signs are at play for workers over the age of 65: freezes eliminate the penalty for continued work thereby increasing $\bar{g}$ through the substitution effect. At the same time, greater pension wealth accumulation raises the continuation value of retirement thereby lowering $\bar{g}$ via the wealth effect. In both the under-65 and over-65 cases, substitution and wealth effects work against each other so the impact of the freezes on retirement behavior cannot be signed.

Despite the overall theoretical ambiguity, two sources of heterogeneity yield the following testable predictions about how these changes in compensation affect labor supply.

1. Between-age heterogeneity

   (a) *1A Substitution effects are larger the closer workers are to age 65:* For workers under 65 when first faced with a freeze, the amount of lost compensation gets progressively smaller in the age at which workers experience the shock (see the right-hand panel of Figure 3). As such, wealth effects on labor supply get progressively weaker and substitution effects get progressively stronger with proximity to age 65.

   (b) *1B The sign of substitution effect is reversed after age 65:* Because workers over the age of 65 experience increases rather than decreases in total compensation, their initial labor supply

[20] Growth in DB pension wealth flattens out after age 65 because the actuarially adjusted present value of DB wealth becomes negative after that age. This pattern in traditional DB accruals is ubiquitous and is designed specifically to induce retirement; see, e.g., Chapter 12 in Gustman et al. (2000). Further details on the sample and the pension freeze simulation are discussed in Appendix D.

response should be opposite in sign to workers under the age of 65.

2. Within-age heterogeneity (under 65)

   (a) *2A Short-term responses are dominated by the substitution effect:* Workers with relatively high values of g, A, and W^{DC} are closer to the margin of retirement. The substitution effect will dominate the behavioral response of these workers and they will be induced to retire early.

   (b) *2B Long-term responses are dominated by the wealth effect:* Workers with relatively low values of g, A, and W^{DC} are farther away from the retirement margin. The wealth effect will dominate the behavioral responses of these workers and they will be induced to retire later.

4 Data

To shed light on the response of individual behavior to pension freezes, I link three distinct sources of administrative information: pension plan characteristics from IRS Form 5500 (F5500) records, employer characteristics from the Census Longitudinal Business Database (LBD), and worker characteristics from the Census Bureau's matched employer-employee Longitudinal Employer Household Dynamics (LEHD) dataset.[21] In this section, I first discuss each administrative data source separately and then explain how I use these sources together to analyze the impact of pension freezes on labor supply. Detailed descriptions of the data and the linking methodology are provided in Appendix A.

Because the pecuniary costs of pension freezes cannot be observed directly in the administrative data sources that I employ, I turn to restricted data on pension eligible respondents in the HRS. Based on DB pension plan parameters obtained from the employers of HRS respondents, these data provide precise measures of DB pension wealth at hypothetical retirement ages. I use these measures to simulate the effects of pension freezes on total compensation. In the latter part of this section, I provide details about the HRS linked pension sample. It is worth re-stating that survey data cannot be used to make inferences about behavior because they are only reported by a handful of HRS respondents.

[21] The F5500-LBD-LEHD link has been used to study the role of fringe benefits on employee mobility in Decressin et al. (2009).

4.1 Administrative Data

4.1.1 Form 5500

F5500 is an annual, plan-specific, filing that is collected jointly by the IRS, Department of Labor (DoL), and the PBGC to ensure compliance with ERISA. These publicly available data contain rich information on the universe of privately sponsored pension plans. When DB plans are converted to cash balance plans or hard frozen, this information is reported on F5500, thereby allowing me to identify plans whose participants were affected by freezes.[22]

F5500 filings are identified by a combination of a Federal Employer Identification Number (EIN) and a employer designated Plan Number (PN) that remain consistent over time. While these identifiers are sufficient to match pension plans to single-unit (i.e. operating only one establishment) firms, they are not sufficient for matching to multi-unit (i.e. operating multiple establishments) firms. This is because payroll tax filings for establishments that are part of a multi-unit firm may be recorded under different EINs than the one used in F5500. Attempting to match the F5500 to establishment level data on EIN alone would therefore generate many false non-matches. To overcome this issue, I turn to the Census Business Register (BR).

4.1.2 Census Business Register and Census Longitudinal Business Database

The BR is a database of the universe of establishments in the United States.[23] It includes information on business location, organization, industry, and information on revenue, payroll, and employment that is collected from administrative tax records as well as survey data. The relationship between establishments belonging to multi-unit firms are determined using responses to the company organization survey, the economic census, and the annual survey of manufactures. Establishments that are part of the same multi-unit firm share the same Census assigned firm identification number even if they have different EINs.

I rely on the presence of EINs in both the F5500 and the BR to create an initial link between the two files. Secondary to this link, I use the Census firm identifier to identify *all* the establishments associated with a multi-unit pension plan sponsor.[24] Having matched F5500 records to the BR, I use the Census firm

[22] Plans are considered frozen in F5500 if they meet the following condition: "[a]s of the last day of the plan year, the plan provides that no participant will get any new benefit accrual (whether because of service or compensation)."

[23] Information about the BR is confidential and protected by Title 13 and Title 26, U.S. Code. Information in the following paragraph is drawn from https://www.census.gov/econ/overview/mu0600.html

[24] In the absence of the Census firm identifier, I would only be able to identify those multi-unit establishments that shared the same EIN as the one reported on F5500, thereby generating a false non-match problem alluded to earlier.

identifier to further match those records to the LBD. The LBD is a cleaned and research ready version of the BR that is restricted to active employers in the private sector.[25]

4.1.3 Longitudinal Employer Household Dynamics

To study outcomes at the individual level, I turn to the LEHD which is a quarterly matched employer-employee dataset constructed from state-level unemployment insurance (UI) records.[26] These data cover almost all wage and salary workers in the United States but exclude individuals who are self-employed. In the LEHD, employers are identified using a state unemployment insurance account number known as the SEIN. I rely on the crosswalk between the SEIN and the Census firm identifier developed in Haltiwanger et al. (2014) to link pension plans in the matched F5500-LBD data to employers in the LEHD.[27]

An important feature of the LEHD is that states become part of the dataset at different points in time. For example, Maryland enters in 1985:Q2 whereas Mississippi enters only in 2003:Q3. Because of staggered entry, the scope of the data grows continuously over time.[28] As a consequence, when I link an employer from the LBD to the LEHD in a given year, I only capture those individuals who work in a state that has already entered the data as of that year.

4.2 Sample restrictions and data structure

Appendix Table A1 describes the results of four data linking and sample restriction procedures. The first row shows the match rate between the universe of DB plans extracted from F5500 database between 1996 and 2014 and the BR. The massive scope of the BR allows for a 92 percent match rate at the plan-year level and a 95 percent match rate at the participant-year level.[29]

The set of plan-years represented in the F5500-BR merge contains a mix of firms that sponsor just one DB plan and firms that sponsor multiple DB plans.[30] I limit my sample to firms that have a single plan within the 1996-2014 window for which I have F5500 data. When firms have multiple plans, I retain only

[25] See Jarmin and Miranda (2002) for details.
[26] See Abowd et al. (2009) for details.
[27] I rely on the 2014 snapshot of the LEHD, which incorporates UI data from 49 states and the District of Columbia through the first quarter of 2015. Alabama is not included in the version of the data that I use.
[28] A number of populous states enter the data relatively early. Illinois enters in 1990, California and Pennsylvania in 1991, Florida in 1992, and New York and Texas enter in 1995.
[29] F5500 reports the count of active participants in each plan. Active participants refers to workers who are accruing benefits under the plan.
[30] Firms that sponsor multiple DB plans typically do so to cover different types of workers. For example, a firm may sponsor different DB plans for salaried and hourly workers or unionized and non-unionized workers.

those employers who choose either to never freeze their plans, or freeze them all at the same time. The principle driver of this restriction is that I cannot observe individual pension plan coverage. When firms sponsor multiple plans, there is no way of knowing — using the F5500, LBD, or LEHD data — which plan a worker may be covered by. By imposing this restriction, however, I can ascertain whether workers at a given firm have been affected by a freeze in a given year. This sample restriction allows me to retain 94 percent of firm-years but only about 40 percent of worker-years. The discordance between these two rates reflects the fact that only the very largest employers sponsor multiple DB plans.

Having matched F5500 records to the BR and the LBD, I structure the data as follows. I treat each year from 2001-2014 as an experimental cohort year, which is indexed by c.[31] This terminology reflects the research design wherein each year yields a fresh sample of firm-level pension freezes. For a given cohort-year, the panel dataset of workers employed at freezing firms constitutes the treated group while the panel dataset of workers employed firms that do not freeze their plans constitutes the comparison group. I impose the restriction that firms file F5500 for 5 calendar years prior to the cohort-year, which I refer to as the pre-period. I then match the firm-cohort-year panels to the LEHD, the results of which are shown in the third row of Table A1. Because of the staggered nature of state level data coverage in the LEHD, I recover 89 percent of firm-cohort-years and 93 percent of employee-cohort-years. Finally, as shown in the fourth row of Table A1, I restrict the sample to firm-cohort-years for which important pension plan data is not missing.[32]

When considering the implications of pension freezes on worker decisions, it is important to reiterate that I do not observe individual information on pension plan coverage. To study worker responses in a way that limits the potential for misclassification error, I restrict the sample to firms where DB eligibility is near universal. I impose this restriction by retaining firms where the DB coverage rate is 80 percent or greater in the pre-period.[33] Within the high-coverage rate firms, I select all workers employed at $c - 5$, who have at least two years of tenure as of $c - 5$, and who will be between the age of 50 and 70 in year c.[34] I focus on

[31]I start with 2001 because it is the first year in which pension freezes are reported in F5500. Only a handful of firms engaged in CB conversions prior to 2001.

[32]Critical pension plan information includes plan assets and liabilities, accruals earned in the filing year, the typical benefit claim age for the plan.

[33]The firm-wide DB coverage rate is the ratio of active participants in the plan as reported in F5500 to the count of total employees in the LBD. The 80 percent average coverage rate requirement is based on years $[c - 5, c - 2]$. Restricting the sample this way likely eliminates soft freezes in which the firm's plan is closed to new workers. A firm that imposes a soft freeze is likely to see its DB coverage rate decline as workers quit or retire but are not replaced with new eligible workers.

[34]The two year tenure restriction ensures that workers are fully vested in their pensions as of the cohort year when they may become subject to a freeze. This calculation is based on the 7 year maximum full vesting period allowed for DB plans by ERISA. Note that tenure measurements are right censored in the LEHD when a worker's employment spell begins prior to the year in which

older workers to study labor supply and retirement behavior among individuals most affected by DB freezes.

4.3 Imputing retirement in the LEHD

I infer LFP in the LEHD based on the earnings history of an individual. Workers who receive positive earnings in a calendar year are deemed to be in the labor force, whereas those that receive zero earnings in a calendar year are deemed to be out of the labor force. I classify an individual as having retired in year t if the last year in which she received non-zero earnings was $t-1$. In this definition, retirement only occurs when an individual completely withdraws from the labor force. Recall, however, that both definitions exclude labor force participation through self-employment because the LEHD data are based only on UI covered earnings which exclude self-employment. To examine the potential for misclassification of retirement in administrative data, I compare the retirement rate of employed, DB eligible, respondents from the 2004 wave of the HRS with a comparable sample of individuals in the LEHD drawn from the 2004 cohort-year whose pensions have not been frozen.[35]

These comparisons between HRS and LEHD data are shown in the three panels of Figure 4, which split the samples into three age categories as of 2004. The retirement rates align very well for the 56-64 year-old age group but diverge somewhat for the 50-55 and 65-70 year-old age groups. For 65-70 year old individuals, the HRS-based rates are lower than the LEHD-based rates which potentially reflects the fact that self-employment is not covered in the LEHD. For the 50-55 year old age group, the HRS-based rates are higher than the LEHD-based rates. This discordance exists even when HRS retirement rates are constructed from a question asking respondents if they have zero earnings. As such, it is unlikely to reflect younger respondents reporting their labor force status as retired even when they have non-zero labor market earnings. In subsequent analyses that rely on the HRS for the estimation of two-sample labor supply elasticities, I focus on the 56-64 year-old age group because labor supply outcomes in the LEHD align most closely to the HRS.

the state that they work in enters the data.

[35] Retirement in the HRS is inferred from a respondent's labor force status report. To align with the LEHD-based definition of permanent departure from paid employment, I consider an HRS respondent as retired only if they continually report their labor force status as retired. In this definition, a respondent who reports being retired in 2006 but re-enters the labor force in 2010 will not be counted as a retiree in the dataset. The labor-force-status-based retirement statistics are very similar to those constructed from a question asking respondents if they have zero earnings from employment.

4.4 Pension wealth in the Health and Retirement Study

The HRS obtains pension summary plan descriptions (SPDs) that provide details on the pension plan provisions of survey respondents. Information from these documents is coded along with relevant data on the earnings histories and job tenure of linked respondents to calculate pension wealth at prospective retirement ages.[36] Because these calculations are developed using employer provided pension plan benefit formulas, they provide high quality measures of DB pension accruals for a representative sample of older workers.

Linked pension data are available for the 1992, 1998, 2004, and 2010 survey years of the HRS. While the 2004 and 2010 samples are most relevant because they coincide with the time period that I analyze, I rely exclusively on the 2010 linked pension sample when estimating labor supply elasticities for two reasons. First, unlike the prior wave samples, the 2010 sample explicitly separates public and private sector plans. By focusing on respondents working in the private sector, I am able to align the survey data to reflect private sector pension provisions which constitute the relevant subset for the analysis. Second, the 2010 sample separates CB plans from plans that continue to follow legacy DB formulas. By isolating the subsample where the evolution of pension accruals occurs under the status-quo — i.e. where participants do not experience a freeze — I can accurately characterize the difference in pension accruals in the no-freeze and freeze states. Appendix D provides additional detail on these data and explains how I use them to simulate the difference in total compensation under frozen and non-frozen scenarios. I refer to the linked HRS pension sample of DB eligible respondents from 2010 as the "HRS sample" in subsequent sections.

5 Empirical framework

5.1 Regression specification

Let i index firms, let j index cells that bin together workers in the same firm-state-gender-age-tenure-cohort, and let t index calendar years. Let $k = t - c$ index years relative to the cohort year. Within each cell, the firm (i) represents a worker's employer as of $k = -5$, and tenure represents the duration of employment at firm i as of $k = -5$. Consider the following regression framework for a given cohort, c,

$$y^c_{j(i)t} = \alpha_i + \gamma^c_t + \boldsymbol{x}^{c'}{}_{j(i)t}\boldsymbol{\beta}^c + \sum_{k=-5}^{m(c)} \delta^c_k T^c_{ik} + \varepsilon^c_{j(i)t} \tag{9}$$

[36] See Fang et al. (2016) for more details.

where $y^c_{j(i)t}$ measures a labor supply outcome of interest, α_i is a firm fixed effect, γ^c_t is a calendar year fixed effect, and $\boldsymbol{x}^c_{j(i)t}$ is a vector of controls for age, gender, state, race, education, tenure, and prior earnings.[37] T^c_{ik} is an indicator variable that equals 1 if the firm freezes its plan and the current period is k. $\varepsilon^c_{j(i)t}$ is the error term which represents unobserved determinants of labor supply.[38] The parameters of interest are the δ^c_k coefficients which capture the dynamic treatment effect of the freeze on worker outcomes.

Because freezes are relatively rare events, I stack data from each of the cohorts together and estimate a version of equation (9) where

$$y^c_{j(i)t} = \alpha_i + \gamma_{ct} + \boldsymbol{x}^{c'}{}_{j(i)t}\boldsymbol{\beta} + \sum_{k=-5}^{13} \delta_k T^c_{ik} + \varepsilon^c_{j(i)t}. \tag{10}$$

In equation (10), calendar year fixed effects are replaced by cohort-by-calendar year fixed effects which allow economy-wide shocks to affect each each cohort differently. In contrast, the effect of the $\boldsymbol{x}^c_{j(i)t}$ and T^c_{ik} are assumed to be constant across cohorts. This estimation strategy allows workers in the comparison group in a given cohort to enter the treated group in a subsequent cohort if their employer freezes pensions in the future.[39] In this implementation, the δ_k coefficients are identified by within-cohort between-firm variation in worker outcomes as well as within-firm between-cohort variation in worker outcomes. Standard errors are clustered at the firm level.

5.2 Identification

The δ_k parameters in Equation (10) represent causal effects of pension freezes on worker outcomes under the assumption that $E[\varepsilon^c_{j(i)t}|\alpha_i, \gamma_{ct}, \boldsymbol{x}^c_{j(i)t}, T^c_{it}] = 0$. Put differently, unobserved determinants of worker labor supply are assumed to have zero mean conditional on firm fixed effects, cohort-by-calendar year fixed effects, worker level controls, and the freeze indicators. This assumption is not robust to two important sources of bias. First, firm-specific economic distress in the pre-period may result in a subsequent freeze as well as a reduction in firm-specific labor demand through downsizing. Second, it is possible that freezing and non-freezing firms systematically differ in terms of pension generosity and benefit claiming provisions

[37] State is defined based on the location of the workplace in $k = -5$. I control for prior earnings using two variables: the log of average annual earnings prior to $k = -5$ and the average growth rate of earnings prior to $k = -5$.

[38] The upper limit of the sum, $m(c)$, represents the number of available post-period years for cohort c. The maximum available post-period duration is 13 years (this happens when c =2001 as the data runs out in 2014).

[39] This approach is similar to recent literature that studies the impact of job loss on earnings. See, e.g., Krolikowski (Forthcoming) and Flaeen et al. (2019).

in the pre-period which can influence post-period differences in labor supply behavior.

To account for time-varying pre-period confounders, I rely on propensity score re-weighting. The main idea behind the use of propensity scores is to make the treatment and comparison group units more comparable in terms of observed pre-period characteristics thereby mitigating concerns that post-period differences in behavior are subject to bias. In this setting, the propensity score is the cell-level probability of experiencing a freeze expressed as a function of pre-period variables which influence both firms' decision to freeze and workers' labor supply responses. To mitigate concerns related to firm distress an omitted variable, the propensity score model includes the pre-period trend in firm size and in worker compensation. To mitigate concerns that differences in plan- and firm-level characteristics between treatment and comparison groups are responsible for post-period labor supply decisions, the propensity score model also includes pre-period trends in pension wealth, pension accruals, benefit claim ages, and age structure of employment at the firm, retirement, labor force participation, and E-E rates. Appendix C provides more details on the conditioning set and explains how the propensity scores are transformed into weights when estimating Equation (10).

Beyond these key threats to the identification strategy, two other confounding effects are potentially at play. First, it is possible that cost concerns that lead firm's to freeze their DB plans also resulted in cut backs to health insurance benefits. These unobserved changes could generate their own income and substitution effects on labor supply choices. While these changes are not directly verifiable in the firm-level sample that I study, indirect evidence from survey's suggests that firms have not altered health benefits as a consequence of freezes.[40] Second, it is feasible that observed changes in labor supply are influenced by network effects within the firm. This breakdown of the so called stable unit treatment value assumption (SUTVA) could occur if freeze-affected workers' labor supply choices are influenced not only by changes in compensation but also changes in the labor supply of their peers. Available evidence on the magnitude of peer effects of this variety indicate that they are extremely small, I therefore ignore peer retirements as a first order concern when interpreting the results[41]

[40]There are no reports of changes to employer provided health insurance benefits in a sample of 17 large publicly traded firms that froze their plans between 2004 and 2008 as compiled by the Boston College Retirement Research Center (see `http://crr.bc.edu/uncategorized/fact-sheets/`). Similarly, a Government Accountability Office (GAO) survey of freezing employers indicates no reported changes to health benefits (see Bovbjerg et al. (2008)).

[41]Hamman et al. (2016), who use large-scale linked employer-employee data from Germany to investigate these spillovers on retirement behavior, find that one additional peer retirement (at the establishment level) increases the probability of retirement for men by 0.01 percentage points and produces no detectable effect on women.

5.3 Summary Statistics

Before showing the impact of freezes on worker outcomes, I present a summary of raw data garnered from the F5500, the LBD, and the LEHD on pre-period characteristics of workers and their employers. These statistics are based on workers employed at the sample of firms where the pre-period coverage rate is in excess of 80 percent. A table showing firm characteristics for the full sample of DB sponsoring employers is provided in Appendix B.

In Table 1 I show pre-period summary statistics for workers split into three different age groups as of the cohort year. Workers in the sample are employed with DB sponsoring firms as of $c-5$, and the statistics are computed by pooling together data for five pre-period years. The top panel shows worker characteristics while the lower panel shows pension plan and firm characteristics. Because the statistics are computed from a worker-level dataset, pension and firm characteristics are worker weighted. For each age-specific panel, the first column shows the propensity score re-weighted comparison group mean, the second column shows the difference between the treatment and the comparison group, and the third column shows the p-value for the null hypothesis that there is no difference between the two groups.

Although there are no statistically significant differences between the treatment and comparison groups for any variable, there are economically meaningful differences in pension wealth per participant and firm size. Statistics on pension wealth are derived from plan liability reports on F5500 which are calculated using a variety of different actuarial methods. Because not all firms use the same actuarial methods, or even the same interest rates and mortality assumptions, it is likely that observed differences in wealth are influenced by firms' use of different actuarial standards. Average DB pension accruals earned in each year and the typical claim age from the plan, shown in the subsequent two rows of the table, are very similar between the treatment and comparison groups. Unlike the wealth statistics, accruals and claim age statistics are not influenced by actuarial methods thereby providing evidence that the incentive structure of DB plans between the treatment and comparison groups is, in fact, very similar. It is important to add that DB pension wealth and accruals statistics represent firm-level averages rather than averages for workers in each age bin.[42]

Secondary to the differences between treatment and comparison groups within each age bin, there are also several notable differences between the three age bins. Workers in the younger two age bins are less

[42] Younger workers constitute the majority of the typical firm's workforce. Furthermore, because of backloading, young workers' DB wealth and accruals are small relative to the levels of older workers. Pension wealth and accrual statistics are therefore indicative only of overall plan characteristics and should not be seen as representative of the workers in each age bin.

likely to be white and male, more likely to have a college degree, and have higher earnings. Tenure, measured five years prior to the cohort-year, is approximately equal across the three age groups which reflects left censoring of employer-employee histories in the LEHD.[43] Finally, the oldest workers are employed at substantially smaller firms with a higher proportion of workers over age 60 and a lower proportion of workers under age 45. Differences in firm-wide age structure across the three age groups could reflect differences in DB pension formulas or other unobserved workplace characteristics. Some of these differences are reflected in higher average pension wealth and delayed retirement claim ages at firms that employ the oldest workers in the sample.

6 Labor supply and employer attachment

In this section, I use the regression framework and identification strategy developed earlier to investigate the causal impact of pension freezes on the labor supply and employer attachment. These results shed light on the testable implications developed in the theoretical model and form the basis for the estimation of the extensive margin labor supply elasticity.

6.1 Labor force participation, retirement, and earnings

Figure 5 plots the δ_k coefficients from specification (10) using the labor force participation rate (LFPR) as the outcome variable. To investigate age-specific heterogeneity in labor supply responses, I split the data into three different age groups and estimate the regression model on each age group separately. The three panels in the figure show coefficients for workers aged 50-55, 56-64, and 65-70 years-old as of the cohort year.

Treated workers in the 50-55 year-old age group exhibit a small but insignificant reduction in LFP rates in the first six years of the post-freeze period. Reductions in participation reflect substitution effects, although the economically small magnitude of the coefficients and their statistical insignificance suggests that offsetting wealth effects are important for workers in this age group. After about 8 years post-freeze, wealth effects start to dominate the labor supply response and workers in the treated group experience a 1.5 - 3.3 percentage point increase in participation relative to the control group. The muted substitution effect and substantial wealth effect that characterizes labor supply responses for 50-55 year-old workers aligns with

[43] States can enter the LEHD data after a given employer-employee relationship is established, thereby leading to left censoring of employer-employee history and an understatement of true tenure.

the fact that large DB accrual spikes occur after age 55. As such, workers at or under 55 when first faced with a freeze experience large net losses in compensation thereby eliciting strong wealth effects in favor of continued labor force participation.

Substitution effects are more pronounced and are statistically significant for treated workers in the 56-64-old age group reflecting the fact that they experience smaller losses in total compensation compared with 50-55 year-old workers. In the first six years of the post-freeze period, LFP rates for treated workers fall by 1.0 - 1.7 percentage points. Starting about 8 years post-freeze, treated workers' labor supply diverges in the opposite direction from the comparison group as participation rates rise by 1.2 - 2.1 percentage points. As with 50-55 year-old workers, the tendency of freeze affected workers to lengthen their working lives relative to the comparison group is indicative of dominant wealth effects.

Most DB plans incentivize retirement by making the actuarially adjusted value of DB accruals negative after age 65.[44] Pension freezes therefore have the effect of eliminating the penalty associated with continued labor force participation for people over age 65. The right most panel of Figure 5 shows labor supply behavior that is consistent with higher returns to work for the treatment group relative to the comparison group. Treated workers over age 65 increase their participation rates by 1.2 - 4.5 percentage points as a consequence of substitution effects. Although these effects are not as precisely estimated, they are economically meaningful and align with theoretical predictions.

Empirical evidence shown in Figure 5 lines up with the four theoretical predictions outlined in Section 3. Substitution effects are larger for 56-64 year-old workers relative to 50-55 year-old workers which is a consequence of the former group losing less total compensation than the latter group (prediction 1A). Substitution effect are positive because of increased labor market returns for workers over 65 but negative for workers under 65 because of reduced labor market returns (prediction 1B). Looking within the first two age groups, substitution effects play a dominant role in the short-term response of workers by lowering LFP rates (prediction 2A) while wealth effects play a dominant role in the long-term response of workers by raising LFP rates (prediction 2B).

Figure 6 shows the impact of freezes on retirement, which is defined as a permanent departure from paid employment in the LEHD. Freeze induced changes in retirement rates are virtually mirror images of the LFPR effects shown in Figure 5, indicating that non-participation and retirement are equivalent for workers affected by freezes. This finding substantiates the model's assumption of self-absorbing retirement. Changes

[44] See, e.g., the left panel of Figure 3 which is based on HRS data or Chapter 12 in Gustman et al. (2000).

in retirement patterns reinforce LFPR-based findings and provide further evidence supporting the between- and within-age labor supply implications of pension freezes predicted by the model.

Figure 7 shows the effect of freezes on the log of annual earnings. In contrast to the LFPR and retirement rate changes, earnings are informative about intensive margin adjustments and therefore provide additional tests for the labor supply implications of pension freezes. For 50-55 year-old workers, the first 6-7 years of the post-freeze period shows no significant change in earnings between groups which is consistent with wealth effects offsetting substitution effects. In the final few years of the sample window, treated group workers earnings exceed that of control group workers by about 15 percent. This difference likely arises from continued full-time work for the treated group relative to part-time work for the comparison group. It is consistent with the wealth effects of the freeze inducing longer full-time labor force participation.

For 56-64 year-old workers the pattern is very similar, although there is a statistically significant earnings dip of 1.5 - 4 percent in the first five post-freeze years. The dip likely reflects intensive margin adjustments including less than full year employment, which is consistent with dominant substitution effects in the short-term.[45] The long-term pattern, which shows a 20-30 percent increase in earnings, provides evidence for freeze induced wealth effects of longer full-time labor force participation. Earnings differences for 65-70 year-old workers are not precisely estimated over the sample window. However, the post-freeze coefficients indicate a fairly sustained drop in earnings of about 15 percent. Thus, while freezes induce workers over 65 to stay in the labor force, the intensive margin estimates suggest that continued participation comes in the form of more part-time work. That the oldest workers in the labor market are willing extend their careers at less than full-time rates complements recent survey-based evidence on the importance of flexible hours in supporting longer labor force participation (Ameriks et al. (2018)).

6.2 Employer attachment

While the behavioral responses discussed thus far relate to the decision about whether to work and how much to work, they do not address the decision about where to work. Employer-to-employer (E-E) transitions are a potentially important margin of adjustment particularly given that pension freezes induce employer-specific rather than worker-specific or market-wide changes in compensation. In this subsection, I exploit

[45] Analyses of firm-level payroll data shown in Appendix B indicate that freezes lower firm-wide average earnings by approximately 2.5 percent. It is unclear whether the firm-wide average earnings losses stem from changing age composition of the firm's workforce or from reductions in offered wages. The fact that 50-55 year-old workers do not experience earnings losses suggests that wage reductions are unlikely to be the only cause for the earnings dip for 56-64 year-old workers.

the matched employer-employee structure of the LEHD to study differences in worker mobility between the treatment and comparison groups.

Figure 8 shows the percentage point change in the probability of leaving one's $c - 5$ employer — i.e. the DB sponsoring employer to which workers are attached five years prior to the cohort-year. In the data, a worker is coded as having experienced an employer change if the EIN associated with their UI record in the LEHD changes. It is worth noting that not all EIN changes reflect employee mobility as some firms change their EINs in the course of a merger or acquisition. The large spike in transitions for treated workers that occurs four years prior to the freeze year in the left panel and the center panel is likely an artifact of firm-level EIN recoding. One period after the freeze, treated workers in the 50-55 and 56-64 year-old age group appear to respond with small but statistically significant increases in employer transitions. For 50-55 year-old workers the transition rate increases by 1.4 percentage points off a baseline rate of 2.9 percent.

For 56-64 year-old workers the transition rate increases by 0.8 percentage points off a baseline rate of 2.8 percent. The magnitude of these effects indicates that relatively younger workers faced with compensation losses from a pension freeze have better outside options to exercise than workers closer to retirement age. After the first year, there is a statistically significant pattern of reduced E-E transitions among workers in the treated group. Transition rates fall by 1.2 percentage points for 50-55 year-old workers and about 0.6 to 0.9 percentage points for 56-64 year old workers. When considered alongside the earnings results, reduced E-E mobility indicates continued employment in career jobs for treated group workers as opposed to the counterfactual transition to part-time bridge jobs for comparison group workers. E-E mobility for workers over 65 is largely unaffected by freezes.[46]

Reduced E-E transition rates for workers in the younger two age groups lines up with the time window where dominant wealth effects lead to increased LFP and higher propensity for full-time work. Ironically, the ability of workers to extend their careers in order to make up for lost compensation comes from extended attachment to the employers responsible for the pension freeze. Reduced E-E mobility, even in the face of substantial employer-specific compensation shock, indicates that full-time work opportunities are limited outside of workers' long-term employers. As such, robust demand for the labor services of older workers within their long-term employers is a key requirement to accommodate policies aimed at supporting longer working lives.

[46]Coefficient estimates are suppressed due to small sample sizes for latter part of the estimation window.

7 Two-sample estimates of the extensive margin labor supply elasticity

Although they provide granular information on large samples of workers affected by pension freezes, the administrative data that I rely on do not contain information DB pension accruals or DC contribution and match rates. As such, I cannot use these data to estimate labor supply elasticities directly. In this section, I describe a two-sample procedure to estimate labor supply elasticities that divides behavioral responses observed in the administrative data by freeze-induced changes in total compensation simulated using the HRS sample.

I consider the Marshallian (or uncompensated) labor supply elasticity (η^M) because I cannot parse wealth and substitution effects in my setting. Under the assumption that leisure is a normal good, the Marshallian elasticity is smaller than the Hicksian (or compensated) elasticity (η^H) which holds wealth constant. As shown, for example, in MaCurdy (1981), the Hicksian elasticity is smaller than the intertemporal (Frisch) labor supply elasticity (η^F) which holds the marginal utility of wealth constant. The three elasticities can therefore be ordered as follows

$$\eta^M < \eta^H < \eta^F. \tag{11}$$

The Frisch elasticity governs intertemporal labor supply responsiveness to temporary and predictable changes in compensation. In contrast, the Hicksian elasticity is relevant for evaluating the welfare effects of permanent changes in compensation; it is typically used to understand the steady state impact of permanent changes in tax rates that are rebated as lump-sums. Due to the inequalities described in (11), the estimates I present are lower bounds for the Hicksian and Frisch elasticities.

7.1 Labor supply elasticity with long-term implicit contracts

Define R^* as the optimal retirement age for a worker whose labor supply is governed by a long-term implicit contract with her employer. R^* can vary not only due to financial considerations such as earnings, pension accruals, and wealth but also because of differences in health capacity, disability status, and non-pay characteristics of employer-employee match quality. As such, R^* measures the horizon over which the worker's implicit contract with her employer is defined.

In the context of long-term implicit contracts, define the extensive margin Marshallian labor supply elasticity by

$$\eta_a^M := \frac{\partial \log(\text{LFPR}_a)}{\partial \log(\omega_{a,R^*})} \tag{12}$$

where LFPR_a is the labor force participation rate of workers aged a. ω_{a,R^*} measures the flow value of an implicit contract which the worker expects to last until R^*. Assuming away job-destruction, job-switching, or mortality risk prior to R^*, the flow value of the contract for a worker currently aged a can be expressed by the annuitized equivalent of total future compensation:

$$\omega_{a,R^*} = \frac{1}{R^* - a} \sum_{t=a}^{R^*} \left(\frac{1}{1+r}\right)^{t-a} \mathcal{C}_t \tag{13}$$

where $\mathcal{C}_t$ is flow compensation as defined as in equation (1). The underlying reason for the formulation in equation (13) is that flow compensation does not necessarily reflect the worker's marginal product in each period as it would in a spot market. When workers are paid the full value of their marginal product over long horizons, it is the present value of total compensation over the relevant horizon rather than flow compensation that determines labor supply.[47] By changing the path of $\mathcal{C}_t$ over the existing contract period, freezes induce plausibly exogenous changes in ω_{a,R^*} thereby aiding in the identification of η^M.

7.2 Two-sample estimates

The treatment effect estimates based on administrative data presented in Figure 6 describe changes in the numerator of expression (12). To obtain estimates of the denominator, I exploit rich information on compensation and expectations derived from the HRS sample. A critical input from survey data is self-reported expected retirement age which, by serving as a measure of R^*, defines career length for each respondent.[48] In the following discussion, I focus on workers in the 56-64 year-old age group at the time of the freeze. I restrict attention to 56-64 year-old workers for two reasons. First, wealth effects approximately cancel out substitution effects for the 50-55 year-old sub-sample. Second, LFPR estimates are imprecisely estimated in the 65-70 year old sub-sample and retirement patterns obtained from the LEHD diverge substantially from

[47] With spot markets, the relevant elasticity would be $\frac{\partial \log(\text{LFPR}_a)}{\partial \log(\mathcal{C}_a)}$ which is the definition adopted in almost all of the literature estimating labor supply elasticities.

[48] This proxy for contract duration assumes that expected retirement ages elicited from the HRS represent the final year of work with career employers not the terminal age of labor force participation after post-career bridge job or post-career self employment. To the extent that post-career work is factored into expected retirement age, R^* overestimates the duration of the contract.

the HRS-based estimates for this age group (see Section 4.3).

Pooling together respondents in the HRS sample with initial ages in the 56-64 range with 7 or more years of tenure (i.e. to match the LEHD sample screen), I estimate the path of $\mathcal{C}_t$ under two scenarios: the first assumes no changes to the status-quo; i.e. non-deferred compensation, DC accruals, and DB accruals evolve under normal projections. The second scenario assumes that a freeze occurs at the initial age. Under this scenario, all future DB accruals are set to 0. When respondents report positive DC balances, I assume that both the respondent and their employer would continue to contribute at the same rate in the post-freeze period. When respondents who report no DC balances, I impute the post-freeze contribution rate using the average in-sample employer and respondent contribution rate rate for respondents with positive balances.[49] Further details on these calculations are presented in Appendix D.

Armed with estimates of behavioral responses from the LEHD and simulated compensation changes from the HRS, I estimate (12) using analog's from both samples as

$$\hat{\eta}_k^M = \frac{\left(\frac{\Delta \text{LF}\hat{\text{P}}\text{R}_k}{\text{LF}\hat{\text{P}}\text{R}_k}\right)^{\text{LEHD}}}{\Delta \log(\hat{\omega}_{a,R^*})^{\text{HRS}}}. \tag{14}$$

In equation (14), k denotes periods since the freeze and a is the age at which workers experience freezes. In the numerator, $\Delta\text{LF}\hat{\text{P}}\text{R}_k$ is the average freeze-induced percentage point change in the LFPR in period k while $\text{LF}\hat{\text{P}}\text{R}_k$ is the average control group LFPR in the same period. The denominator measures the freeze induced loss in the value of the implicit contract holding its original duration fixed. To isolate substitution effects as far as possible, I focus on workers who are near the participation margin at the time of the freeze and are therefore induced to leave the labor force (as highlighted in prediction 2A in Section 3). I do this by pooling over those k where $\Delta\text{LF}\hat{\text{P}}\text{R}_k$ is negative. This procedure does not account for heterogeneity in non-pension wealth or DC wealth at the time of the freeze, so the estimates do not account for differences in initial wealth.

Table 2 shows estimates of participation elasticities in the top panel and retirement elasticities in the bottom panel. The retirement elasticity is analogous to the participation elasticity but uses retirement as the behavioral outcome of interest. The first column pools together data from both genders and the second and third columns provide gender-specific estimates. In the rows of each panel, I separately show the numerator

[49] This assumption imputes DC saving choices rather than using the model to solve for them. In ongoing work, I rely on the structural model to solve for optimal consumption and saving choices in the post-freeze environment.

and denominator of equation (7) and then show the estimated elasticity. I compute standard errors of the two-sample estimates using the delta method.

The participation elasticity estimates are 0.47 for the sample as a whole, 0.34 for men, and 0.63 for women. That women have larger elasticities than men potentially reflects a combination of labor supply preferences and higher levels of non-labor income through spousal earnings and is consistent with existing estimates. The pooled estimate I obtain is two times larger than an average of prime-age-based compensated estimates of 0.25 in the meta-analysis of Chetty et al. (2012). The difference in magnitudes between the estimates is a product of two different factors. First, even after netting out differences in household wealth levels, prime-age workers may have smaller labor supply elasticities than older workers as consequence of differing preferences. More importantly, the existing estimates are based on the spot market perspective of employer-employee associations which may not adequately capture the relevant incentives affecting labor supply behavior for a large swath of the labor force.

Two related studies provide benchmarks on the relative magnitude of the pension freeze-based behavioral elasticities reported here. Importantly, both studies examine the behavior of older workers and both are based on the spot market perspective of employer-employee associations. First, the participation elasticity estimates I obtain are similar to Gelber et al. (2017) who exploit budget set non-linearities to estimate extensive margin responses of 63-64 year-olds to the Social Security annual earnings test (AET). They obtain lower bound estimates of 0.49 for both genders pooled together, 0.25 for men, and 0.49 for women. Second, the retirement elasticity estimates I obtain are substantially larger than Manoli and Weber (2011) who exploit tenure-specific severance pay discontinuities for workers over age 55 in Austria.[50] They report estimates of 0.12 for men and 0.38 for women, which are about one-third as large as the retirement elasticities shown in Table 2.

If the AET penalty for older Americans who have already claimed Social Security is relevant primarily for self-employment or work in short-term bridge jobs as opposed to long-term career jobs, then implicit contracts are less important in the setting Gelber et al. (2017) study. As such, behavioral elasticities are accurately captured by the AET-based estimates. On the other hand, tenure-specific severance pay embodies the idea of a long-term job by its very design. Interpreting severance pay using the lens of a spot market, as in Manoli and Weber (2011), mechanically leads to smaller elasticity estimates because severance pay

[50] The estimates in Manoli and Weber (2011) are interpreted as Frisch elasticities because they involve temporary and anticipated change in compensation. Their estimates therefore represent the upper bound of the scale of labor supply elasticities (see equation 11).

constitutes a large share of flow compensation but a smaller share of the total value of compensation earned over the expected duration of the contract.[51]

The argument presented here underscores the idea that the labor market context in which behavioral elasticities are estimated is critical for inference. Furthermore, it indicates that relatively large elasticity estimates obtained from pension freeze natural experiments are, in fact, accurate depictions of labor supply preferences.

8 Conclusion

In this paper, I exploit important and understudied changes in employer sponsored pension benefit programs to better understand labor supply behavior in the context of long-term jobs rather than spot markets. In particular, I examine the private sector DB environment where employers have reneged on promises to continue supporting traditional retirement benefits in the face of rising costs of provision. Creating a new dataset that brings together detailed administrative information on plan characteristics with matched employer-employee data, I study the impact of these unexpected changes on labor force participation, retirement, earnings, and employer attachment for a sample of workers between age 50 and 70.

I combine estimates of the extensive margin labor supply response to freezes with estimates of their impact on total compensation simulated among HRS respondents. Within the long-term implicit contract model, freeze induced changes in compensation are measured as changes in the annuitized present value of compensation holding fixed the existing horizon of the contract. I estimate extensive margin labor supply elasticities of 0.47 for 56-64 year-old workers. For men in the same age range, the elasticity is 0.34. For women it is 0.63. The elasticity estimates are uncompensated because they do not account for wealth effects. Nevertheless, they are two times larger than estimates reported in prior studies. I argue that these differences are a product of estimating elasticities within the framework of long-term jobs rather than spot markets.

[51] In the spot market perspective adopted in Manoli and Weber (2011), severance pay is viewed as a part in flow compensation in the year that it is paid out. Since severance pay is large relative to annual earnings, it constitutes a large share of flow compensation. In the spot market perspective, the change in labor supply behavior induced by severance pay is divided by a large percentage change in compensation thereby yielding a small elasticity. In the long-term contract perspective, severance pay is part of the present discounted value of compensation to be earned over the remaining duration of the contract. In this perspective, severance pay constitutes a smaller percentage of total compensation. Dividing the same change in labor supply behavior by a smaller percentage change in compensation would yield a larger elasticity. A similar reliance on the spot market framework leads to a retirement elasticity estimate of 0.18 on the basis of changes to the pension accrual rate for public sector workers near retirement age in California. See Brown (2013) and the related calculations in Chetty et al. (2012).

References

Abowd, John, Bryce E. Stephens, Lars Vilhuber, Fredrik Anderson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock, "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," in Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press, 2009, pp. 149–230.

Abowd, John M. and David Card, "Intertemporal Labor Supply and Long-Term Employment Contracts," *American Economic Review*, 1987, *77* (1), 50–68.

_ and _ , "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 1989, *57* (2), 411–445.

Ameriks, John, Joseph S. Briggs, Andrew Caplin, Minjoon Lee, Matthew D. Shapiro, and Christopher Tonetti, "Older Americans Would Work Longer If Jobs Were Flexible," 2018. NBER Working Paper No. 24008.

Beaudry, Paul and John DiNardo, "Is the Behavior of Hours Worked Consistent with Implicit Contract Theory?," *Quarterly Journal of Economics*, 1995, *110* (3), 743–768.

Becker, Gary S., "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, 1962, *70* (5), 9–49.

Bianchi, Marco, Bjorn R. Gudmundsson, and Gylfi Zoega, "Iceland's Natural Experiment in Supply-Side Economics," *American Economic Review*, 2001, *91* (5), 1564–79.

Blundell, Richard, Antonie Bozio, and Guy Laroque, "Labor Supply and the Extensive Margin," *American Economic Review Papers and Proceedings*, 2011, *101* (3), 482–486.

Bovbjerg, Barbara D., Charles A. Jeszeck, Charles Ford, Isabella Johnson, Luann Moy, Mark Ramage, Joe Applebaum, Craig Winslow, Gene Kuehneman, Brian Friedman, Melissa Swearingen, Marietta Mayfield, Sue Bernstein, and Walter Vance, "Defined Benefit Pensions: Freezes Affect Millions of Participants and May Pose Retirement Income Challenges," Technical Report, United States Government Accountability Office 2008.

Brown, Kristine M., "The Link Between Pensions and Retirement Timing: Lessons from California Teachers," *Journal of Public Economics*, 2013, *98*, 1–14.

Bulow, Jeremy, "What are Corporate Pension Liabilities," *Quarterly Journal of Economics*, 1982, *97* (3), 435–452.

Busso, Matias, John DiNardo, and Justin McCrary, "New Evidence on the Finite Sample Properties of Propensity Score Reweighting and Matching Estimators," *Review of Economics and Statistics*, 2014, *96* (5), 885–897.

Card, David and Dean R. Hyslop, "Estimating the Effects of a Time Limited Earnings Subsidy for Welfare Leavers," *Econometrica*, 1995, *73* (6), 1723–70.

Carrington, William J., "The Alaskan Labor Market During the Pipeline Era," *Journal of Political Economy*, 1996, *104* (1), 186–218.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," in Daron Acemoglu, Jonathan Parker, and Michael Woodford, eds., *NBER Macroeconomics Annual 2012*, Vol. 27, University of Chicago Press, 2012, pp. 1–56.

Decressin, Anja, Tomeka Hill, Kristin McCue, and Martha Stinson, "The Role of Fringe Benefits in Employer and Workforce Dynamics," in Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press, 2009, pp. 473–505.

Devereux, Paul J., "Changes in Relative Wages and Family Labor Supply," *Journal of Human Resources*, 2004, *39* (3), 696–722.

Eissa, Nada and Hilary Hoynes, "Taxes and the Labor Market Participation of Married Couples," *Journal of Public Economics*, 2004, *88* (9-10), 1931–1958.

_ **and Jeffrey B. Liebman**, "Labor Supply Response to the Earned Income Tax Credit," *Quarterly Journal of Economics*, 1996, *111* (2), 605–637.

Fang, Chichun, Charles Brown, and David Weir, "Cohort Changes in Social Security Benefits and Pension Wealth," 2016. Michigan Retirement Research Center Working Paper WP 2016-350.

Flaeen, Aaron, Matthew D. Shapiro, and Isaac Sorkin, "Reconsidering the Consequences of Worker Displacements: Firm versus Worker Perspective," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 193–227.

Gelber, Alexander M., Damon Jones, Daniel W. Sacks, and Jae Song, "Using Non-Linear Budget Sets to Estimate Extensive Margin Responses: Method and Evidence from the Social Security Earnings Test," 2017. NBER Working Paper No. 23362.

Gustman, Alan, Olivia Mitchell, Andrew Samwick, and Thomas L. Steinmeier, "Evaluating Pension Entitlements," in Olivia Mitchell, P. Brett Hammond, and Anna M. Rappaport, eds., *Forecasting Retirement Needs and Retirement Wealth*, University of Pennsylvania Press, 2000, pp. 310–326.

Haltiwanger, John, Henry Hyatt, Erika McEntarfer, Liliana Sousa, and Stephen Tibbets, "Firm Age and Size in the Longitudinal Employer-Household Dynamics Data," *Center for Economic Studies Working Paper 14-16*, 2014.

Hamman, Mary K., Daniela Hochfellner, John M. Nunley, and Christopher Ruhm, "Peer Effects and Retirement Decisions: Evidence from Pension Reform in Germany," 2016. Unpublished working paper.

Ippolito, Richard, "The Labor Contract and True Economic Pension Liabilities," *American Economic Review*, 1985, *75* (5), 1031–1043.

Jarmin, Ron S. and Javier Miranda, "The Longitudinal Business Database," *Center for Economic Studies Working Paper 02-17*, 2002.

Kotlikoff, Laurence J. and David A. Wise, "Labor Compensation and the Structure of Private Pension Plans: Evidence for Contractual versus Spot Labor Markets," in David A. Wise, ed., *Pensions, Labor, and Individual Choice*, University of Chicago Press, 1985, pp. 55–88.

Krolikowski, Pawel, "Choosing a Control Group for Displaced Workers," *Industrial and Labor Relations Review*, Forthcoming.

Lazear, Edward P., "Why Is There Mandatory Retirement?," *Journal of Political Economy*, 1979, *87* (6), 1261–1284.

—, "Pensions as Severance Pay," in Zvi Bodie and John B. Shoven, eds., *Financial Aspects of the U.S. Pension System*, University of Chicago Press, 1984, pp. 57–85.

— **and Robert L. Moore**, "Pensions and Turnover," in Zvi Bodie, John B. Shoven, and David A. Wise, eds., *Pensions in the U.S. Economy*, University of Chicago Press, 1988, pp. 163–190.

Lumsdaine, Robin L., James H. Stock, and David A. Wise, "Three Models of Retirement: Computational Complexity versus Predictive Validity," in David A. Wise, ed., *Topics in the Economics of Aging*, University of Chicago Press, 1992, pp. 21–60.

MaCurdy, Thomas E., "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 1981, *89* (6), 1059–1085.

Manoli, Dayanand S. and Andrea Weber, "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions," 2011. NBER Working Paper No. 17320.

Meyer, Bruce and Dan T. Rosenbaum, "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers," *Quarterly Journal of Economics*, 2001, *116* (3), 1063–1114.

Monahan, Amy, "Statutes as Contracts - The California Rule and its Impact on Public Pension Reform," *Iowa Law Review*, 2012, *97*, 963–1029.

Munnell, Alicia H. and Mauricio Soto, "Why are Companies Freezing their Pensions?," 2007. Working Paper 2007-22. Center for Retirement Research, Boston College.

Rauh, Joshua D., "Hidden Debt, Hidden Deficits: 2017 Edition: How Pension Promises are Consuming State and Local Budgets," Technical Report, Hoover Institution 2017.

—, Irina Stefanescu, and Stephen P. Zeldes, "Cost Saving and the Freezing of Corporate Pension Plans," 2017. Columbia Business School Research Paper No. 16-4.

Rosen, Sherwin, "Implicit Contracts: A Survey," *Journal of Economic Literature*, 1985, *23* (3), 1144–1175.

Stock, James H. and David A. Wise, "The Pension Inducement to Retire: An Option Value Analysis," in David A. Wise, ed., *Issues in the Economics of Aging*, University of Chicago Press, 1990, pp. 205–229.

— **and** —, "Pensions, the Option Value of Work, and Retirement," *Econometrica*, 1990, *58* (5), 1151–1180.

Urahn, Susan K., Greg Mennis, Susan Banta, David Draine, Aleena Oberthur, Nicholas Heitger, and Fatima Yousofi, “The Challenge of Meeting Detroit’s Pension Promises,” Technical Report, The Pew Charitable Trusts March 2018.

Zelinksy, Edward A., “The Cash Balance Controversy,” *Virginia Tax Review*, 2000, *19*, 683–762.

Figure 1: Hard frozen and cash balance DB Plans

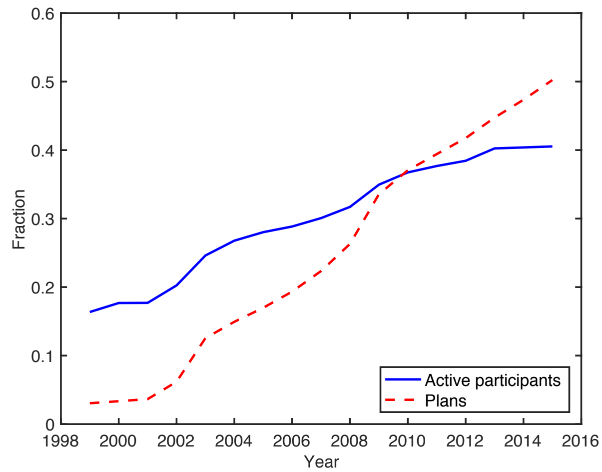

Notes: Time series are based on 1999-2015 F5500 microdata.

Figure 2: Pension Costs and Funding Ratios

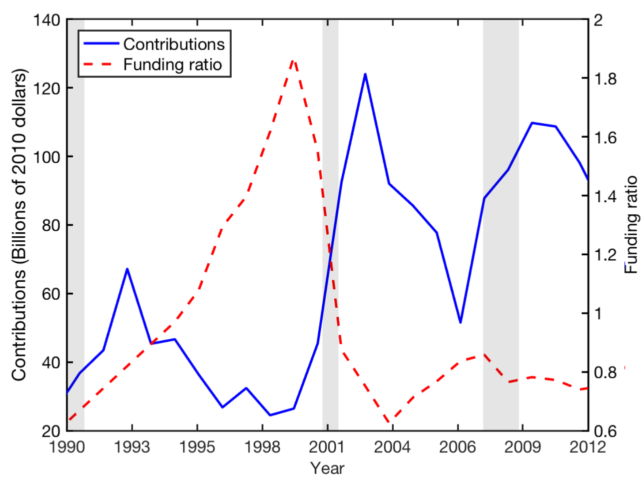

Notes: Contributions data are drawn from the Department of Labor's Private Pension Plan Bulletin Historical Tables and Graphs 1975-2015. Funding ratios are drawn from the PBGC Pension Insurance Data Book, 2016. Funding ratios between 1990 and 1995 are interpolated. Contributions and funding data are based on single employer DB plans. Shaded areas indicate NBER recessions.

Figure 3: Simulated effect of pension freezes on deferred compensation

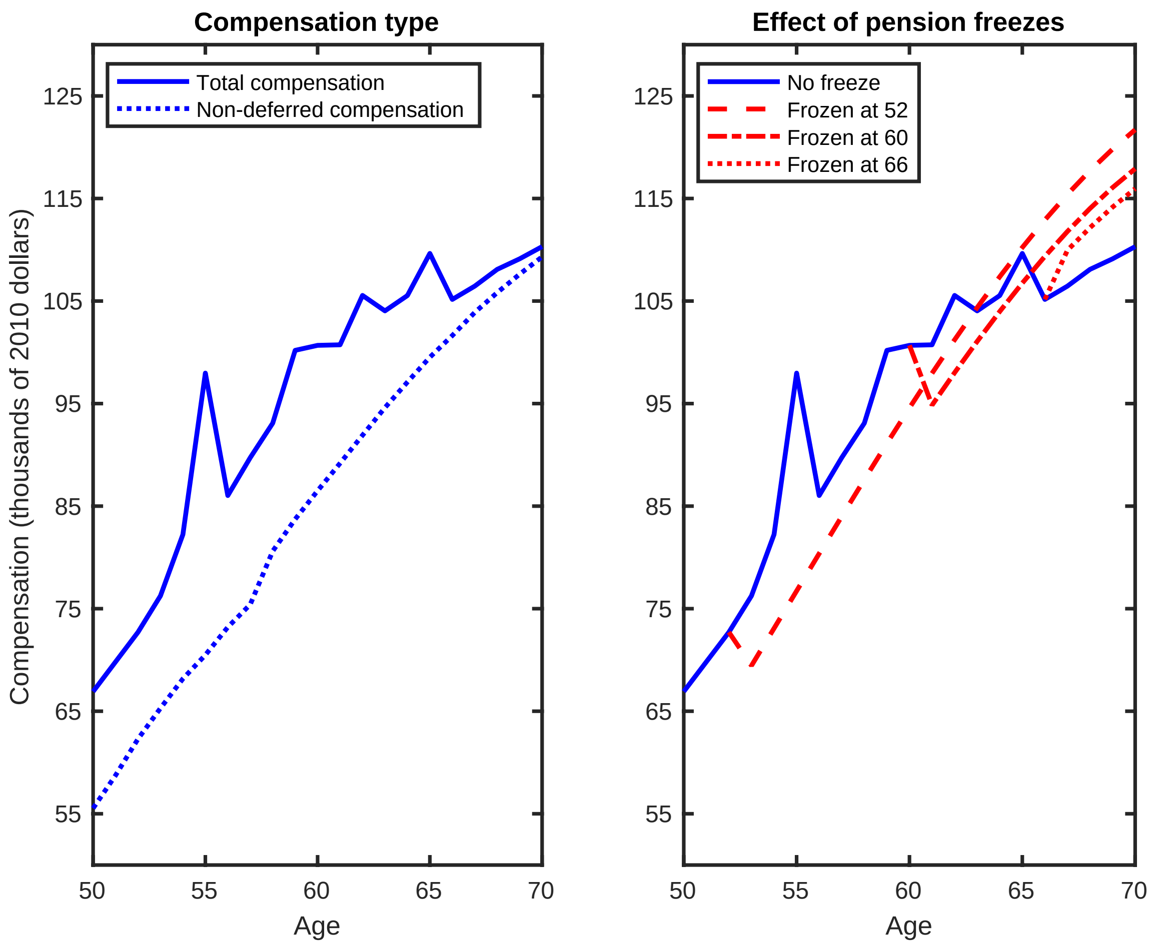

Notes: This figure is based on data from HRS respondents in the 2010 wave who are employed in the private sector. For simulated freeze compensation paths, post-freeze DB accruals are set to zero. Respondents with DC plans are assumed to continue participating in them at the same rate; post-freeze contibution rates for respondents without DC plans are imputed using the sample average rate for participants. See Appendix D for details on the sample and simulation of post-freeze compensation.

Figure 4: HRS versus LEHD retirements

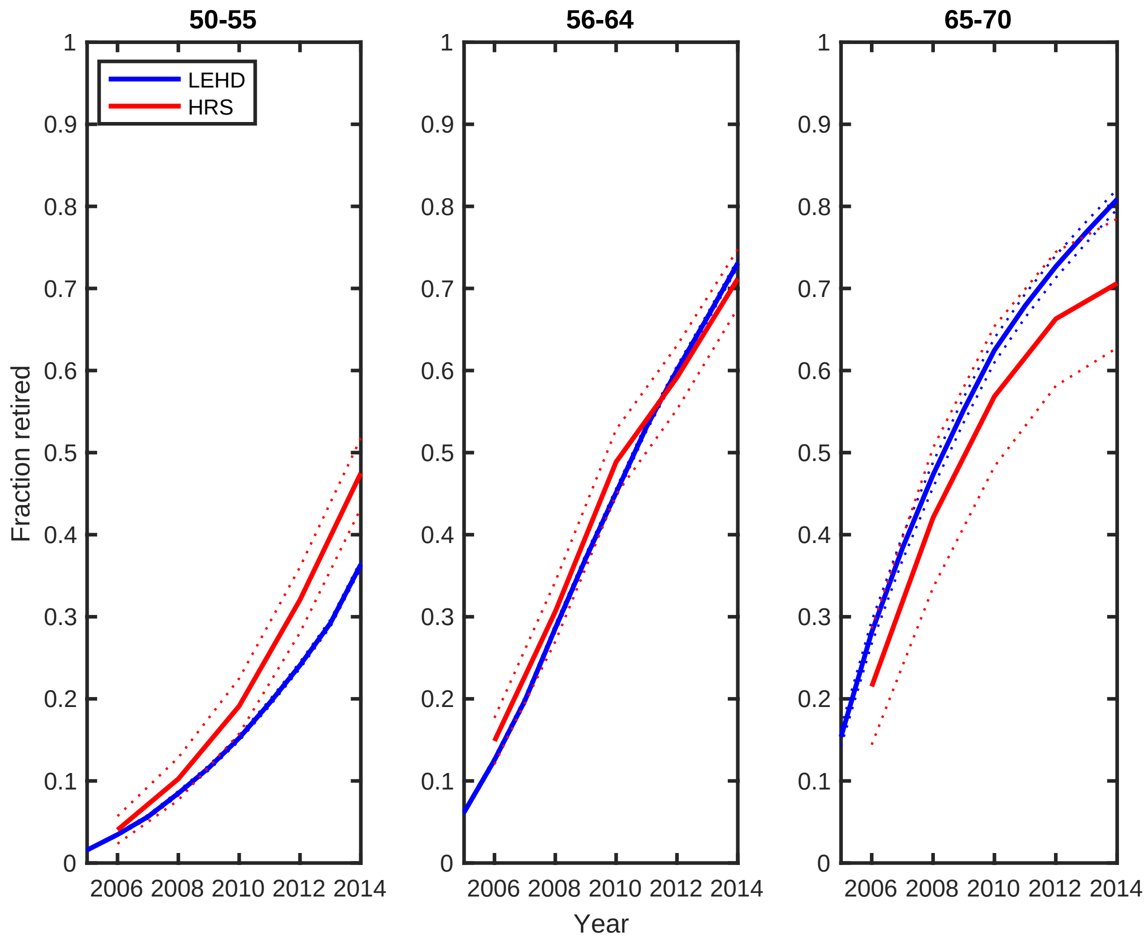

Notes: HRS data are based on respondents who are working as of 2004 and are eligible for employer sponsored DB pension plans. LEHD data are based on the sample of individuals employed at DB sponsoring firms as of 2004 where the firm-wide coverage rate is $\geq$ 80 percent. Sample splits are based on age as of 2004.

Figure 5: Impact of freezes on LFPR

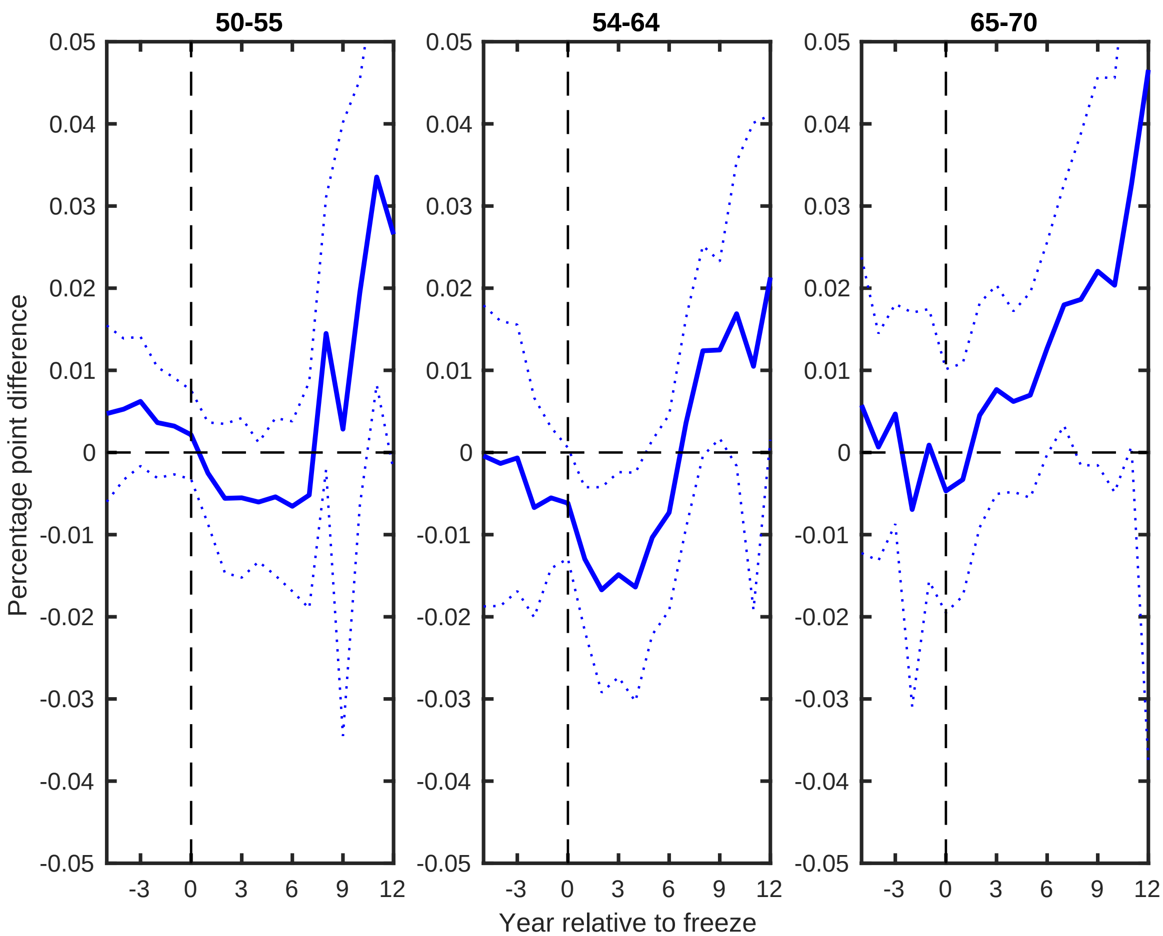

Notes: Dotted lines show 95 percent confidence intervals which are based on standard errors clustered at the firm-level.

Figure 6: Impact of freezes on retirement

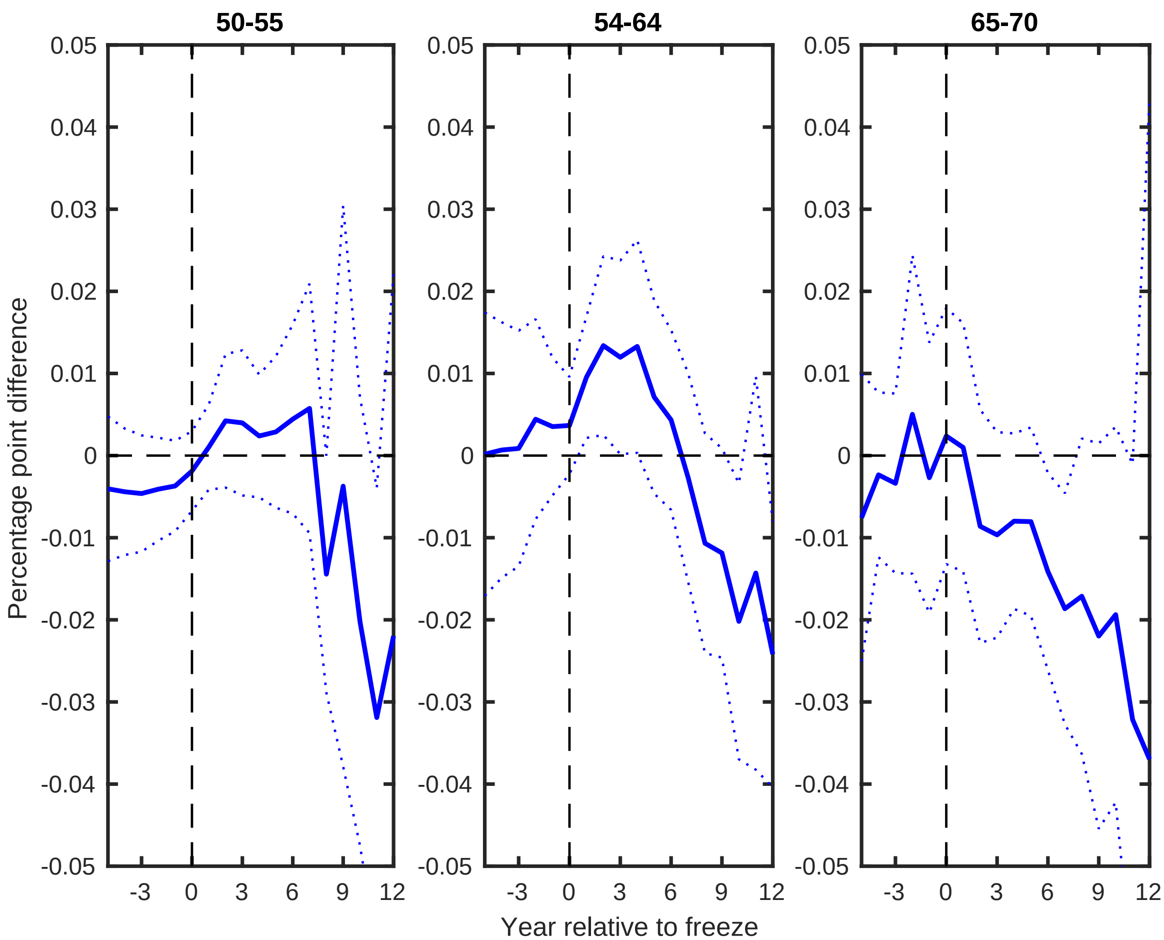

Notes: Dotted lines show 95 percent confidence intervals which are based on standard errors clustered at the firm-level.

Figure 7: Impact of freezes on log annual earnings

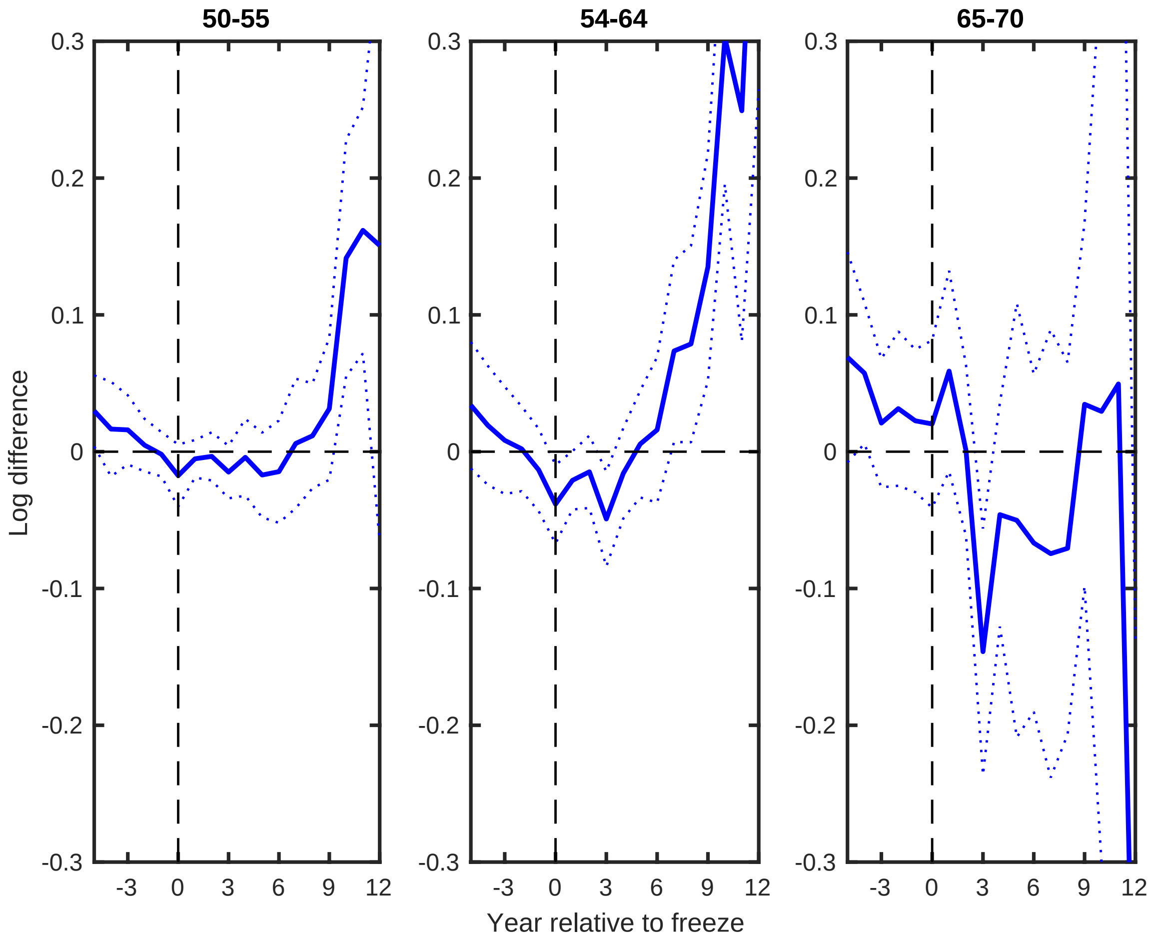

Notes: Dotted lines show 95 percent confidence intervals which are based on standard errors clustered at the firm-level.

Figure 8: Impact of freezes on employer attachment

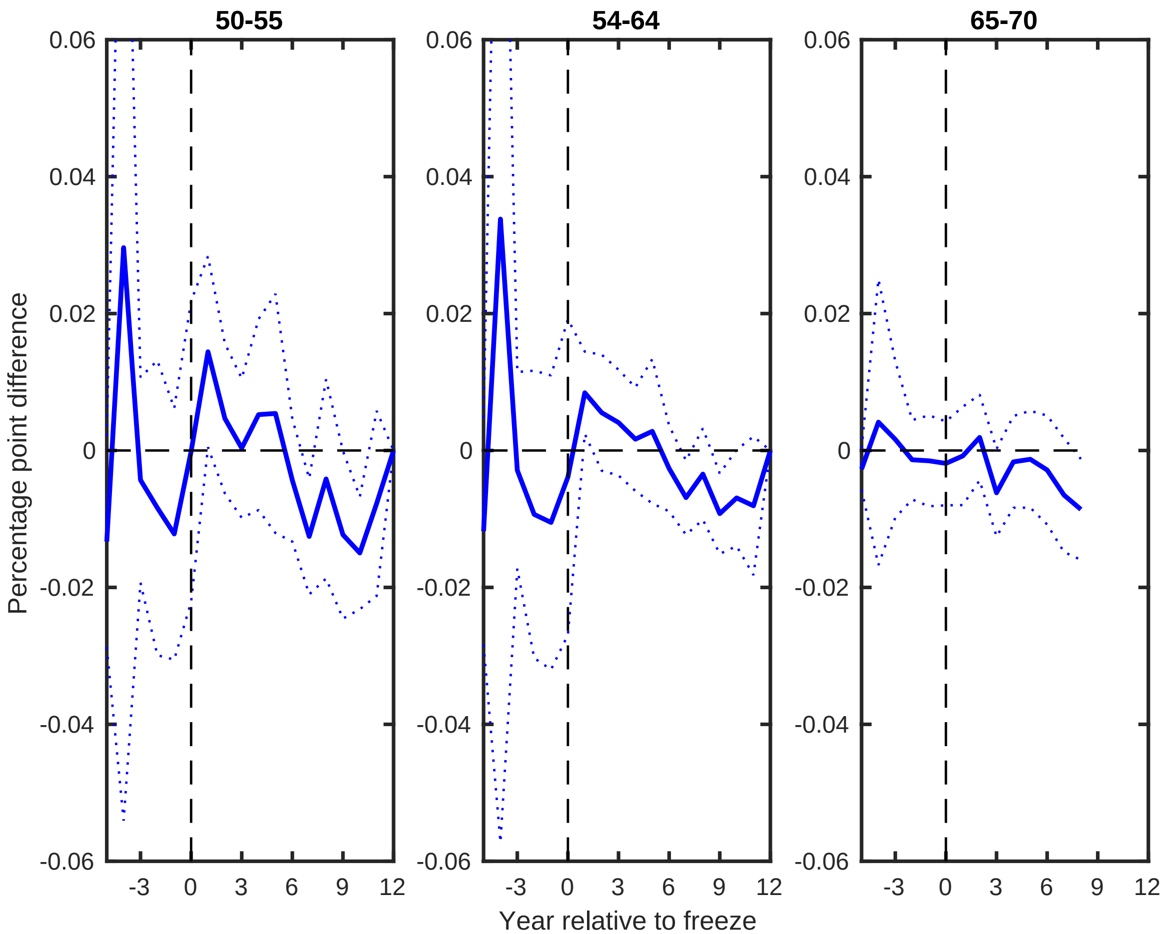

Notes: Dotted lines show 95 percent confidence intervals which are based on standard errors clustered at the firm-level.

Table 1: Pre-period summary statistics split by age group

Variable	50-55			56-64			65-70		
	Comp. mean	Diff.	p-value	Comp. mean	Diff.	p-value	Comp. mean	Diff.	p-value
				Worker characteristics					
Age	52.5	0.001	0.93	59.6	-0.031	0.32	67.0	0.016	0.56
Male	0.475	0.025	0.46	0.487	0.023	0.55	0.494	0.014	0.76
High school	0.232	-0.001	0.93	0.232	-0.001	0.93	0.294	0.001	0.98
Some college	0.324	-0.006	0.55	0.312	-0.007	0.42	0.290	-0.005	0.58
College or more	0.386	0.008	0.81	0.392	0.010	0.77	0.311	0.008	0.76
White	0.791	0.022	0.20	0.826	0.020	0.25	0.838	0.027	0.25
Black	0.096	-0.011	0.15	0.079	-0.010	0.12	0.069	-0.010	0.21
Hispanic	0.068	-0.003	0.81	0.053	0.000	0.97	0.049	-0.005	0.61
Other race	0.045	-0.008	0.27	0.043	-0.009	0.23	0.044	-0.013	0.31
Earnings ($)	65140	-1477	0.75	65820	-2337	0.58	47770	-3562	0.13
Tenure at $c-5$	7.8	-0.9	0.15	8.2	-1.0	0.11	8.0	-0.6	0.17
Retired	0.022	0.000	0.99	0.056	0.002	0.57	0.189	0.006	0.59
In labor force	0.964	-0.001	0.83	0.927	-0.003	0.47	0.779	-0.006	0.66
Switched $c-5$ employer	0.047	0.012	0.22	0.042	0.012	0.24	0.030	0.003	0.45
				Pension and firm characteristics					
DB pension wealth/ptcp. ($)	24290	-4044	0.36	27150	-4640	0.34	30010	-649.7	0.85
DB pension accrual/ptcp. ($)	2179	-230	0.59	2373	-262	0.55	2560	21	0.95
Avg. benefit claim age	62.8	-0.1	0.70	62.9	-0.1	0.72	64.4	0.0	0.88
Firm size	15200	-6353	0.48	14320	-5624	0.51	4503	-562	0.82
Fraction workforce ≤ 45	0.580	0.010	0.44	0.568	0.007	0.61	0.542	-0.003	0.80
Fraction workforce [46,50]	0.146	-0.002	0.62	0.141	0.000	0.91	0.135	0.001	0.79
Fraction workforce [51,55]	0.124	-0.002	0.61	0.127	-0.001	0.78	0.124	0.001	0.89
Fraction workforce [56,60]	0.087	-0.002	0.63	0.095	-0.002	0.67	0.098	0.001	0.74
Fraction workforce [61,65]	0.044	-0.002	0.40	0.047	-0.002	0.52	0.066	0.000	0.99
Fraction workforce [66,70]	0.012	-0.001	0.38	0.013	-0.001	0.54	0.022	0.000	0.87
Fraction workforce ≥ 71	0.007	-0.001	0.33	0.008	-0.001	0.56	0.013	0.001	0.70
Comparison group workers		383000			373000			77000	
Treated group workers		60000			66500			11000	
Comparison group firms		7700			8600			4600	
Treated group firms		1500			1700			900	

Notes: Statistics reported in the table average over years [c-5,c-1]. All dollar values are expressed in 2010 terms. Pension wealth per participant is computed as the present value of the liability owed to active participants divided by the number of active participants. Tenure is understated because the LEHD does not capture the complete history of an employer-employee relationship when states enter the dataset after a given employee-employer relationship is established. P-values for the difference between treatment and control groups are obtained by regressing the statistic of interest on a indicator variable for treatment status and clustering standard errors at the firm-level.

Table 2: Two-sample elasticity estimates

	Overall	Men	Women
	Participation elasticity		
$(\Delta \hat{\text{LFPR}}/\hat{\text{LFPR}})^{\text{LEHD}}$	-0.022	-0.017	-0.027
	(0.008)	(0.009)	(0.009)
$\Delta \log(\hat{\omega}_{a,R^*})^{\text{HRS}}$	-0.047	-0.050	-0.043
	(0.002)	(0.003)	(0.003)
Elasticity	0.467***	0.341*	0.628***
	(0.166)	(0.189)	(0.207)
	Retirement elasticity		
$(\Delta \hat{\text{RR}}/\hat{\text{RR}})^{\text{LEHD}}$	0.029	0.020	0.039
	(0.013)	(0.015)	(0.015)
$\Delta \log(\hat{\omega}_{a,R^*})^{\text{HRS}}$	-0.047	-0.050	-0.043
	(0.002)	(0.003)	(0.003)
Elasticity	0.619**	0.403	0.905**
	(0.271)	(0.293)	(0.361)

Notes: Standard errors for the elasticities are estimated using the delta method. *** p<0.01, ** p<0.05, * p<0.1. LEHD-based estimates are obtained by pooling estimates for $k \in$ 1-6. All estimates are based on samples of 56-64 year-old workers at the time of real or simulated DB pension freeze.

A Data Appendix

Form 5500

F5500 is an annual filing collected jointly by the IRS, DoL, and the PBGC to ensure compliance with ERISA.[52] Each plan in the F5500 database is identified by a combination of an EIN and plan number (PN). The PN is assigned by the plan's sponsor and stays fixed over the life of the plan. For form years 2000-2015, the DoL has prepared an edited research sample of the data in which logical and arithmetic errors are corrected and multiple filings for the same plan are de-duplicated. From 2000-2009, the research data include all pension plans with more than 100 participants and a 5 percent sample of plans with less than 100 participants. I use records from the research data where possible and add back small plans (i.e. those with less than 100 participants) from the raw data if they are excluded from the research sample. I de-duplicate multiple filings for the same plan in the raw data files by retaining the most recent filing in a given year. I obtained pre-1999 data through a FOIA request to the DoL. The sample that I use covers plan years ending 1996-2014.

I focus primarily on DB plans, but also obtain data on DC plans offered by employers who sponsor DB plans. Plan characteristics are coded using a set of numbers and letters. In post-1999 F5500 data, DB plans have prefix 1, DC plans have prefix 2 and 3, and welfare benefit plans — such as employer provided health insurance — have prefix 4.[53] Hard frozen DB plans are recorded using code 1I. Cash balance plans are recorded using code 1F. I eliminate supplemental plans which I identify by searching for any case versions of the string "supplemental" in the plan name field. I restrict my sample to single employer plans, thereby eliminating multi-employer DB plans.[54]

In addition to the main form, I include data taken from the actuarial information attachment. Prior to 2007 this attachment was labeled Schedule B. After 2009 it was labeled Schedule SB for single employer plans. For 2008, I impute actuarial data by interpolating between the 2007 and 2009 values because actuarial

[52] F5500 data are publicly available at https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/public-disclosure/foia/form-5500-datasets.

[53] The codes are different prior to 1999. Welfare benefit plans only need to be reported with a F5500 filing when such plans cover more than 100 active participants.

[54] Multi-employer plans are arrangements between a group of firms and/or unions to provide pension benefits to eligible employees within the group.

information is unavailable in electronic format. The actuarial attachment contains important plan level data including detailed breakouts of plan assets and liabilities, accruals earned in the current year, the average retirement age/benefit claim age, mortality and separation rate assumptions, etc.

Linking to the Census Business Register and Census Longitudinal Business Database

The BR is a database of the universe of establishments in the United States.[55] It includes information on business location, organization, industry, and information on revenue, payroll, and employment that is collected from administrative tax records as well as survey information. The relationship between establishments belonging to multi-unit firms are determined using responses to the company organization survey, the economic census, and the annual survey of manufactures. Establishments that are part of the same multi-unit firm share the same Census assigned firm identification number even if they have unique EINs.

To link the F5500 files to the BR, I match EIN-plan-end-years in F5500 to EIN-years in the BR. Because the massive scope of the BR, I am able to match approximately 92 percent of DB plan-years in the F5500 files to specific establishments in the BR (see row 1 of Table A1). Non-matches occur when a plan EIN does not map to any establishment with positive payroll in the BR which could happen, for example, when a plan is sponsored by a union or an employer-association.

The set of plan-years represented in the F5500-BR merge contains a mix of firms that sponsor just one DB plan and firms that sponsor multiple DB plans.[56] I limit my sample to firms that have a single plan within the 1996-2014 window for which I have F5500 data. When firms have multiple plans, I retain only those employers who choose either to never freeze their plans, or freeze them all at the same time. The principle driver of this restriction is that I cannot observe individual pension plan coverage. When firms sponsor multiple plans, there is no way of knowing — using the F5500, LBD, or LEHD data — which plan a worker may be covered by. By imposing this restriction, however, I can ascertain whether workers at a given firm have been affected by a freeze in a given year. This sample restriction allows me to retain 94 percent of firm-years but only about 40 percent of worker-years (see row 2 of Table A1). The discordance between these two rates reflects the fact that only the very largest employers sponsor multiple DB plans.[57]

[55]Information about the BR is confidential and protected by Title 13 and Title 26, US Code. The following information is drawn from `https://www.census.gov/econ/overview/mu0600.html`.

[56]Firms that sponsor multiple DB plans typically do so to cover different types of workers. For example, a firm may sponsor different DB plans for salaried and hourly workers or unionized and non-unionized workers.

[57]When firms sponsor multiple plans and pass the sample screen, I sum plan-level variables such as assets, liabilities, and participant counts across all plans sponsored by the firm. I compute the weighted average of the retirement age reported on F5500 using the number of participants in each plan as weights.

Using Census firm identifiers, I match these data with the LBD which is a cleaned and research ready version of the BR. The LBD covers private sector establishments with non-zero payroll but excludes some industrial sectors (see p.4 of Jarmin and Miranda (2002) for details).

Having matched F5500 records to the BR and the LBD, I structure the data as follows. I treat each year from 2001-2014 as an experimental cohort year, which is indexed by c.[58] This terminology reflects the research design wherein each year yields a fresh sample of firm-level pension freezes. For a given cohort-year, the panel dataset of workers employed at freezing firms constitutes the treated group while the panel dataset of workers employed firms that do not freeze their plans constitutes the comparison group. I impose the restriction that firms file F5500 for their DB plans 5 calendar years prior to the cohort-year, which I refer to as the pre-period. By requiring plan stability in the lead up to the cohort-year, I implicitly follow a specific set of workers covered by a DB plan regardless of whether their employer merges, grows from single-unit to multi-unit, or vice versa. I match information on DC plans offered by the set of DB sponsoring employers to these data using the same EIN-based linking procedure described above. The key DC-related variable is the number of workers covered by DC plan(s). When a firm offers multiple DC plans, I pick the maximum number of active participants across plans and use that count to estimate the DC coverage rate at the firm.[59]

Linking to the Longitudinal Employer Household Dynamics

The LEHD is a quarterly matched employer-employee dataset constructed from state-level unemployment insurance (UI) records. The UI system covers 96 percent of wage and salary employment nationally, although the data exclude independent contractors, the unincorporated self-employed, railroad workers covered by railroad unemployment insurance, and some other minor categories of workers who are not covered by state-level UI laws. State and local government employees are included in the data but elected officials, members of the judiciary, and some emergency employees are excluded. Federal government workers and workers employed in Alabama are excluded from the version of the data that I use in this paper.

An important feature of the LEHD is that states enter the dataset at different points in time. For example, Maryland enters in 1985:Q2 whereas Mississippi enters only in 2003:Q3. Because of staggered entry, the scope of the data grows continuously over time.[60] I use the 2014 snapshot version of the LEHD, which pro-

[58] I start with 2001 because it is the first year in which pension freezes are reported in F5500. Cash balance conversions are reported in earlier years but the number of firms making CB conversions before 2001 is small.

[59] I use the maximum across plans rather than the sum across plans because workers can participate in multiple DC plans. The DC coverage rate is the ratio of the maximum number of DC participants from F5500 to the count of employees from the LBD.

[60] A number of populous states enter the data relatively early. Illinois enters in 1990, California and Pennsylvania in 1991, Florida

vides matched employer-employee histories from each state's entry quarter up through 2015:Q1. I eliminate the single quarter of 2015 from these data as it represents partial year information on earnings and is not representative of the annual data structure that I employ.

In the LEHD, employers are identified using a state UI account number known as the SEIN while workers are identified using a variable known as a protected identification key (PIK). I begin by matching firm-level data from the F5500-LBD linked sample to the T26 Employer Characteristics File (ECFT26) in the LEHD. The ECFT26 is a SEIN-quarter-year level file that contains the Census firm identifier associated with each SEIN. Using this common unique identifier, I can match national plan- and employer- level characteristics from the F5500-LBD linked sample to state level employers in the LEHD. Because of the staggered nature of state level data coverage in the LEHD, I recover 89 percent of firm cohort-years which corresponds to 93 percent of employee-cohort-years from the F5500-LBD linked sample (see row 3 of Table A1). From this sample, I drop a small percentage of observations where certain pension plan variables are missing (see row 4 of Table A1).[61]

Worker sample

When considering the implications of pension freezes on worker decisions, it is important to reiterate that I do not observe individual information on pension plan coverage. To study worker responses in a way that limits the potential for misclassification error, I restrict the sample to firms where DB eligibility is near universal. I impose this restriction by retaining firms where the DB coverage rate is 80 percent or greater in the pre-period.[62] Within the sample of high-coverage rate firms, I use the LEHD Employment History File (EHF) to obtain matched employer-employee data. The EHF is a SEIN-PIK-year level file that provides the earnings history associated with each employer-employee combination. To these data, I add information on date of birth, race and ethnicity, and education from the Individual Characteristics File (ICF). I then select all workers employed at a DB sponsoring firm in $c - 5$, who have at least two years of tenure as of $c - 5$, and who will be between the age of 50 and 70 in year c.[63]

in 1992, and New York and Texas enter in 1995.

[61]Critical pension plan information includes plan assets and liabilities, accruals earned in the filing year, the average benefit claim age for the plan.

[62]The firm-wide DB coverage rate is the ratio of active participants in the plan as reported in F5500 to the count of total employees in the LBD. The 80 percent average coverage rate requirement is based on years $[c - 5, c - 2]$. Restricting the sample this way likely eliminates soft freezes in which the firm's plan is closed to new workers. A firm that imposes a soft freeze is likely to see its DB coverage rate decline as workers quit or retire but are not replaced with new eligible workers.

[63]The two year tenure restriction ensures that workers are fully vested in their pensions as of the cohort year when they may become subject to a freeze. This calculation is based on the 7 year maximum full vesting period allowed for DB plans by ERISA.

Table A1: Data linkage and sample restrictions

Match/restriction type	N	Match rate	Person weighted rate
F5500-BR (plan-years)	852000	0.922	0.947
Multi-plan restriction (firm-years)	699000	0.942	0.378
LBD-LEHD (firm-cohort-years)	1419000	0.885	0.933
No missing pension data (firm-cohort-years)	1256000	0.930	0.972

Notes: F5500 data are based on years ending 1996-2014. Pension data is treated as missing if plan liabilities, assets, accrual amounts, and claim ages are either missing or unreadable in electronic format and cannot be interpolated.

B Firm and plan characteristics around the freeze

This appendix describes how firm and plan characteristics evolve around the freeze. I establish three facts. First, I show that worsening plan finances rather than worsening firm performance is the main predictor of DB pension freezes. Second, I show that the aftermath of a freeze leads to small but persistent reductions in firm size and average pay. Third, I show that DB freezes generate an immediate transition towards DC plan participation.

B.1 Pre-freeze environment

To describe the environment prevailing prior to the firms freeze decision, I show several characteristics of firms and their pension plans averaged over a five year pre-period in Table B1. Comparing the left and right panels of the table shows the freezing firms are very similar in terms of size, average pay, employee age structure, and DB pension plan characteristics. DB and DC coverage rates within the two sets of firms are both approximately 70 and 35 percent respectively. The lack of any meaningful difference in employee age structure, pension liability, pension accrual rates, and claim ages indicates that the freezing firms are not disproportionately staffed by older workers at the threshold of retirement. Put differently, firms that freeze their plans are not on the brink of a large liability cliff. The likelihood of experiencing a mass layoff, which is recorded as a 30 percent reduction in employment and labeled firm distress, is about 5.5 percent in both groups. Economic distress driven by a large negative shock in the output market is therefore not a leading reason for freeze decision either.

They key distinction between freezing and non-freezing firms lies in the financial health of their DB plans. For every dollar in future liabilities, non-freezing firms have $1.10 in assets. The same ratio — referred to as the funding ratio — is 1.05 for firms that ultimately freeze their plans. The PPA designates plans with funding ratios under 80 percent as being "at-risk" or distressed. Using the PPA's threshold, 20 percent of freezing firms are have distressed plans whereas the same rate is 15 percent for non-freezing firms. Funding deficiencies are particularly important from a cost management perspective because gaps must be closed to meet statutory requirements. Furthermore, once underfunded, plans are no longer buffered against financial market shocks the way overfunded plans are. Required contributions towards underfunded plans therefore become larger and more volatile in the face of market risk.

Table B2 shows coefficients from a linear prediction model using freeze in the cohort year as the outcome and a variety of pre-event characteristics as predictors. The regressions are estimated on data pooling over a five year pre-period, thereby allowing for the inclusion of firm fixed effects. Columns 1 and 2 do not include firm fixed effects, while columns 3 and 4 do. The regressions show that a 1 percent improvement in the funding ratio lowers the likelihood of a future freeze by 2.5 percentage points. This partial effect is stable and statistically significant across all four specifications. Firm size is negatively correlated with future freezes, but the magnitude of the effect is small: a 1 percent increase in firm size lowers the likelihood of a future freeze by 0.2 to 0.5 percentage points. In specifications with firm fixed effects, employee age structure has no statistically significant impact on freezes and the magnitudes of the partial effects are negligible when expressed in proportional terms. DB plans that are collectively bargained are about 2.5 percentage points less likely to experience a freeze which implies that unions offer approximately the same protective effect as a one percent improvement in plan funding. Industry fixed effects, which are included in columns 2 an 4 have no appreciable impact on the estimated coefficients indicating that industry-specific factors are not important, conditional on the other predictors in the model.

B.2 Post-freeze changes

Figure B1 compares the evolution of four variables between freezing and non-freezing firms before and after the freeze. Each panel plots coefficients from an event study regression using the specification described in equation (10) with firm-cohort-year level data. Note that the estimated coefficients are net of firm fixed effects and therefore remove time invariant unobserved heterogeneity between firms. In each panel, the horizontal axis represents the calendar year relative to the cohort-year.

The upper row of Figure B1 shows the difference in log of total employment and log of average pay between freezing and non-freezing firms. The estimated coefficients show that freezes lead to a persistent 2.5 percent reduction for both outcome variables — a gap which closes only after about 10 years. Post-freeze differences in size and pay between the two types of firms represent some combination of changes to the age or seniority composition of the firm's workforce after a freeze either through labor supply responses or changes in labor demand, although it is not clear from firm-level data alone what the role for each channel is.[64] The lower left panel of Figure B1 shows that the fraction of DC covered workers starts to rise about 2 years prior to the DB freeze reflecting expanded DC plan eligibility, more generous DC match rates,

[64] The person-level analyses presented in Section 6 isolate labor supply factors by using propensity score methods to condition on the pre-freeze path of firm size which serves as a proxy for latent changes in output demand for the firm.

or transitions to opt-out rather than opt-in DC enrollment.[65] In the period right around a DB freeze, DC coverage rates increase by about 5 percentage points off a baseline coverage rate of about 35 percent. In subsequent years, non-freezing firms gradually increase their DC coverage and catch up to the DC coverage rate prevailing at freezing firms. Whether the catch up occurs through soft-freezes that close existing DB plans to younger workers, or through more generous incentives for DC participation, the results shown here provide evidence that DB freezes accelerate the inevitable transition towards DC pension coverage within firms. Evidence for increased DC participation is important in explaining the extended labor force participation of some freeze-affected workers as it allows them to offset DB losses.

The lower right panel of Figure B1 shows the change in the likelihood of a freezing firm to experience economic distress, which is defined as a reduction in employment of 30 percent or greater. The coefficient estimates from an unweighted regression (in blue) show that the immediate aftermath of a freeze induces a 1.5 percentage point increase in the probability of distress which lasts for two years. When the same regression is weighted by firm size (in red), thereby representing the change in the probability of freeze affected workers experiencing large employment contractions, the point estimates are economically and statistically insignificant. As such, it appears that distress is concentrated among smaller firms.

[65]DC coverage is measured as the ratio of the maximum number of participants across a firm's DC plans to total employment.

Table B1: Pre-period firm and plan characteristics

	Non-freezing firms		Freezing firms	
Variable	Mean	Std. error	Mean	Std. error
Size	273.7	3.3	254.6	19.5
Average earnings ($)	69300	252	68010	520
Firm age	20.9	0.0	21.2	0.1
Multi-unit	0.273	0.001	0.250	0.003
Fraction workforce ≤ 45	0.579	<0.001	0.578	0.001
Fraction workforce [46,50]	0.117	<0.001	0.115	0.001
Fraction workforce [51,55]	0.110	<0.001	0.109	0.001
Fraction workforce [56,60]	0.092	<0.001	0.095	0.001
Fraction workforce [61,65]	0.057	<0.001	0.060	0.001
Fraction workforce [66,70]	0.024	<0.001	0.023	<0.001
Fraction workforce ≥ 71	0.022	<0.001	0.020	<0.001
Distressed firm	0.056	<0.001	0.057	0.001
DC plan offered	0.502	0.001	0.521	0.003
DC plan coverage rate	0.348	0.001	0.362	0.002
DB plan coverage rate	0.723	0.000	0.708	0.001
DB pension wealth/ptcp ($)	98910	372	101600	1612
DB pension accrual/ptcp ($)	14200	41	14690	145
Average benefit claim age	63.2	0.0	63.3	0.0
Collectively bargained plan	0.040	<0.001	0.030	0.001
Funding ratio	1.11	0.00	1.05	0.00
Distressed plan	0.159	0.001	0.193	0.002
Firm-cohort-years	428000		28500	
Firms	22500		6500	

Notes: Statistics reported in the table average over the five year period preceeding any freeze activity (i.e. cohort-years [c-5,c-1]). All dollar values are expressed in 2010 terms. Pension wealth per participant is computed as the present value of the liability owed to active participants divided by the number of active participants. Plan-years are coded as distressed if their ratio of assets to liabilities is under 80 percent — the threshold below which DB plans are considered "at risk" in the Pension Protection Act of 2006. Firm-years are coded as distressed if firm-wide year-on-year employment shrank by 30 percent or more.

Table B2: Predictors of future freezes

Variables	(1)	(2)	(3)	(4)
Log funding ratio	-0.0264***	-0.02545***	-0.0258***	-0.0258***
	(0.0020)	(0.0020)	(0.0022)	(0.0022)
DB coverage rate	-0.0075**	-0.0029	0.0008	0.0009
	(0.003)	(0.0030)	(0.0036)	(0.0036)
Firm age	-0.00009	0.00001	0.00027	0.00034
	(1.04e-04)	(1.07e-04)	(5.49e-04)	(5.40e-04)
Log size	0.0026***	0.0021***	-0.0055**	-0.0056**
	(6.32e-04)	(6.51e-04)	(0.0024)	(0.0024)
Log average pay	0.00004	-0.00063	-0.00591***	-0.00598***
	(0.00114)	(0.00117)	(0.00182)	(0.00181)
Fraction workforce ≤ 45	0.0168*	0.0178*	-0.0141	-0.0143
	(0.0098)	(0.0099)	(0.0124)	(0.0124)
Fraction workforce [46,50]	0.0097	0.0119	-0.0115	-0.0117
	(0.01083)	(0.0109)	(0.013)	(0.013)
Fraction workforce [51,55]	0.0085	0.0114	-0.0217*	-0.022*
	(0.01068)	(0.0107)	(0.0131)	(0.0131)
Fraction workforce [56,60]	0.0223**	0.0251**	-0.0106	-0.0108
	(0.01087)	(0.0109)	(0.0129)	(0.0129)
Fraction workforce [61,65]	0.0253**	0.0272**	0.0037	0.0036
	(0.0116)	(0.0116)	(0.0131)	(0.0131)
Fraction workforce [66,70]	0.0024	0.0036	0.0069	0.0069
	(0.0124)	(0.0124)	(0.0123)	(0.0123)
DB plan collectively bargained	-0.0215***	-0.0257***		
	(0.0035)	(0.0036)		
Firm offers DC plan	0.0016	0.0030*		
	(0.0016)	(0.0016)		
Multi-unit firm	-0.0119***	-0.0069***		
	(0.0023)	(0.0024)		
Observations	456000	456000	456000	456000
Adjusted R-squared	0.015	0.016	0.362	0.362
Firm FE	No	No	Yes	Yes
Cohort-period FE	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	Yes
Number of firm clusters	23500	23500	23500	23500

Notes: Robust standard errors, clustered at the firm level, in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Regressions are estimated on a panel dataset that pools the five year period preceeding any freeze activity (i.e. cohort-years [c-5,c-1]).

Figure B1: Firm characteristics around freezes

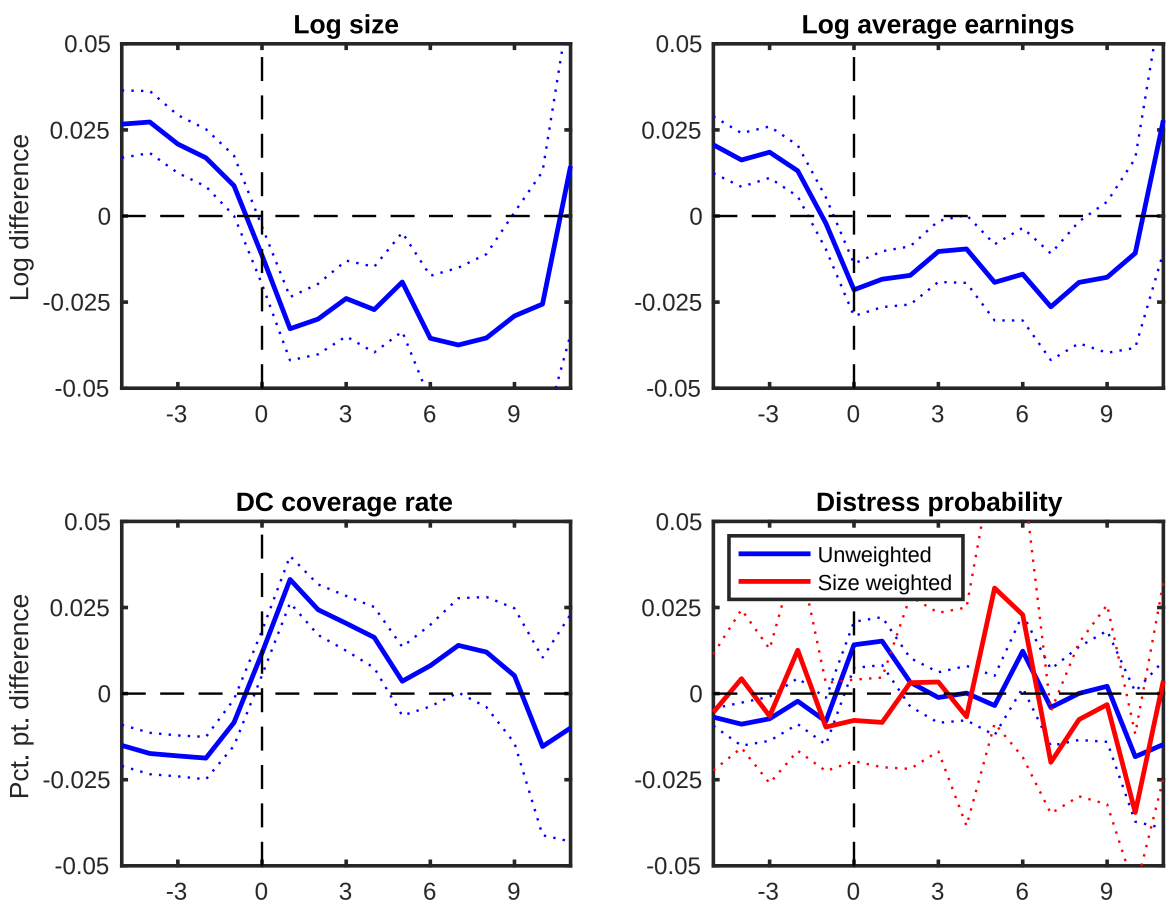

Notes: Dotted lines show 95 percent confidence intervals which are based on standard errors clustered at the firm level. Horizontal axes show years relative to the cohort-year. Firm-years are coded as distressed if firm-wide year-on-year employment shrank by 30 percent or more.

C Propensity score re-weighting

For workers in cell $j(i)$, denote the probability of experiencing a freeze, or the propensity score, by $\hat{p}(\boldsymbol{z}_{j(i)})$. $\boldsymbol{z}_{j(i)}$ is a vector including all pre-period observations on firm size, firm-level averages of total pension wealth and pension accruals per working participant, average benefit claim age from the pension plan, the age structure of the firm's workforce, cell-level earnings, retirement rates, labor force participation rates, and employer-to-employer transition rates. I also condition on state, gender, tenure, and prior earnings. $\hat{p}(\boldsymbol{z}_{j(i)})$ is estimated using logistic regression.

In the setting being considered in this paper, the parameter of interest is the average treatment effect on the treated (ATET) — i.e. the impact of pension freeze shocks on the labor supply of workers affected by those shocks. To estimate the ATET, each comparison group unit is re-weighted by $\frac{\hat{p}(\boldsymbol{z}_{j(i)})}{1-\hat{p}(\boldsymbol{z}_{j(i)})}$. Following Busso et al. (2014), the weights are first normalized to sum to 1 so that the number of weighted units in the comparison group is unaffected by the re-weighting procedure.

Constructing good counterfactuals for treated group units requires that comparison group units with the same $\hat{p}(\boldsymbol{z}_{j(i)})$ — i.e. the same ex-ante probability of experiencing the treatment — can be found in the sample. This requirement is referred to as the "common support condition" or the "overlap condition." Formally, the common support condition for the ATET parameter requires that $\hat{p}(\boldsymbol{z}_{j(i)}) < 1$ for all $j(i)$. In practice, the treatment and comparison groups in the data I use share a large region of common support and the maximum $\hat{p}(\boldsymbol{z}_{j(i)})$ is lower than 1.

D Counterfactual compensation after a freeze

DB pension wealth in the HRS sample is calculated using the following formula

$$W_{T_0}^{DB} = \sum_{t=T_0}^{119} P_t \left(\frac{1+COLA}{1+i}\right)^{t-T_0} B_{t|T_0} \tag{15}$$

where $W_{T_0}^{DB}$ is the present value of pension wealth at retirement age or quit date T_0, P_t is the probability of survival at age t conditional on being alive in T_0, $COLA$ (cost of living adjustment) is the plan specific annual growth rate of payments (for most DB plans in the HRS $COLA = 0$), i is the nominal interest rate and $B_{t|T_0}$ is the annual pension benefit at age t conditional on retiring at T_0.[66] In the HRS, these values are adjusted from T_0 to 2010 to facilitate comparisons across a variety of hypothetical retirement ages. Earnings are a critical component of the DB formulas used to compute $B_{t|T_0}$. In the HRS sample, earnings for years prior to 2010 reflect actual survey measures whereas earnings for years after 2010 are projected using a variety of assumptions detailed in Section I-B of Fang et al. (2016).[67] Tenure is the other major component of the DB formula. HRS DB wealth estimates rely both on self-reports of tenure and on linked Social Security earnings histories to ascertain tenure.

DC accounts constitute the second part of deferred compensation. Recall from equation (2) that DC wealth at time t can be expressed as

$$W_t^{DC} = W_{t-1}^{DC}(1+r) + y_t(m_t^w + m^e(m_t^w))$$

where W_{t-1}^{DC} is the stock of DC wealth carried forward from the prior period, r is the real interest rate, y_t is earnings in the current period, and m_t^w and m_t^e are the respective proportion of earnings contributed to the DC plan by the worker and the firm. In the HRS, respondents report the value of their DC account balances along with the contribution rates, m^w and m^e.[68] Recall from equation (1) that total compensation is the sum

[66] P_t are based on the 2010 version of gender-specific cohort mortality tables published by the Social Security Administration (SSA). The 2010 pension wealth calculations assume a nominal interest rate of 5.7 percent and an inflation rate of 2.8 percent according to economic assumptions detailed in the 2010 Annual Report of the Board of Trustees of the old age, survivors, and disability insurance (OASDI) trust funds of the SSA.

[67] The basic HRS linked pension sample provided to researchers does not incorporate age-based curvature in the projected earnings profile. I adjust projected earnings for age induced changes using estimates from a regression of log annual earnings on age and age-squared using data from all working respondents across all HRS waves.

[68] Individual and employer contributions are expressed either as a percentage or as a dollar value in the HRS. When expressed as a dollar value, I divide them by earnings to convert them to a percentage. Some respondents provide intervals rather than point measures of contribution rates. I use the mid-point of the interval to impute these contribution rates.

of cash compensation and accruals in deferred compensation is given by

$$\mathcal{C}_t = y_t(1 - m^w) + \Delta W_t^{DB} + \Delta W_t^{DC}$$

Define R^* as the optimal retirement age for a worker currently aged a. Recall from equation (13) that the annuity value of a worker's long-term employment with the firm is given by

$$\omega_a = \frac{1}{R^* - a} \sum_{t=a}^{R^*} \left(\frac{1}{1+r}\right)^{t-a} \mathcal{C}_t$$

where all values are expressed in constant dollar terms. If a freeze occurs in period k, then $\Delta W^{DB} = 0$ in period k and in all subsequent periods. I assume that the post-freeze substitution to DC or transition to CB are equivalent. Define ω^F as the annuity value of a worker's long-term employment in the post-freeze environment. Then, for a worker aged a when faced with a freeze:

$$\omega_a^F = \frac{1}{R^* - a} \sum_{t=a}^{R^*} \left(\frac{1}{1+r}\right)^{t-a} \Bigg(y_t + \underbrace{\Delta \tilde{W}_t^{DC}}_{\text{Post-freeze DC accruals}}\Bigg) \tag{16}$$

Simulating changes in compensation due to freezes

I begin by selecting DB eligible respondents in the HRS sample, all of whom were working when surveyed. I retain respondents employed in the private sector whose plans are coded as traditional DB plans (i.e. their plans have neither been frozen nor converted to cash balance). Note that these data have a psuedo-panel structure: they provide pension wealth and earnings estimates for each respondent hypothetical retirement ages. Let i index respondents in the HRS and let t index time periods in the psuedo-panel; denote $\tilde{a}_{it}$ as the worker's hypothetical retirement age. I select the sample of observations where $\tilde{a}_{it} \in [56, 64]$ and respondents have 7 or more years of tenure as of age $\tilde{a}_{it}$ in order to match the LEHD sample screen. [69]

With these data I simulate the effect of a pension freeze on total compensation as follows:

1. Earnings (y_{it}) either reflects actual prior earnings or is based on an age specific projection provided by the HRS. ΔW_{it}^{DB} is obtained by differencing DB wealth computed using the HRS pension estimation program (PEP). DB pension wealth is based on the maximum of wealth at the plan's normal retirement

[69] I do not simulate pension freeze induced changes in compensation for 50-55 year-olds because wealth and substitution effects largely offset each other in the reduced form estimates from the LEHD. I exclude 65-70 year old workers from the simulation because the retirement patterns obtained from the LEHD diverge substantially from the HRS-based estimates (see Section 4.3).

age, early retirement age, and vested deferred value of benefits. When respondents have wealth in multiple plans, I sum pension wealth over each plan.

2. When respondents report a positive DC account balance, I use the law of motion described in equation (2) to estimate the respondent's DC balance at each hypothetical retirement age. In making this calculation, I use earnings (y_{it}) , 2010 survey data on the respondent's report of m^w and m^e, and the nominal interest rate and inflation rate assumed in the PEP to complete the calculations. When respondent's make contributions to multiple plans or receive employer contributions to multiple plans, I sum the fraction of earnings contributed across all plans to determine the total contribution rate.

3. Using y_{it}, ΔW_{it}^{DB}, and ΔW_{it}^{DC} and R_i^* , I compute $\mathcal{C}_{it}$ and ω_i for each respondent.

4. Next, I simulate pension freezes by setting $\Delta W_{it}^{DB} = 0$ for all $k \geq 0$. When respondents have zero DC account balances at $k = 0$, I initiate DC contributions in the post-freeze period using the in-sample average rates. When respondents report positive balances at $k = 0$, DC accruals continue at the self-reported contribution rates. Using these estimates, I compute ω_i^F for each respondent.

5. I compute $\log(\omega_i^F) - \log(\omega_i)$ as the freeze-induced change in annuitized compensation for each respondent as of $k = 0$.