

Patki, Dhiren

Working Paper

Breaking the implicit contract: Using pension freezes to study lifetime labor supply

Working Papers, No. 21-7

Provided in Cooperation with:

Federal Reserve Bank of Boston

Suggested Citation: Patki, Dhiren (2021) : Breaking the implicit contract: Using pension freezes to study lifetime labor supply, Working Papers, No. 21-7, Federal Reserve Bank of Boston, Boston, MA,
<https://doi.org/10.29412/res.wp.2021.07>

This Version is available at:

<https://hdl.handle.net/10419/238089>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Breaking the Implicit Contract: Using Pension Freezes to Study Lifetime Labor Supply

Dhiren Patki

Abstract:

This paper studies the elimination of traditional pensions and subsequent adoption of 401(k) plans by U.S. employers. Using thousands of firm-level natural experiments, it shows that unexpected losses in future compensation engendered by pension plan transitions induce premature retirement for some workers and delayed retirement for others. Observed heterogeneity in retirement behavior is indicative of differences in wealth and in preferences for leisure. Using credibly identified treatment effects as estimation targets, it fits a structural model of retirement and uses the model to evaluate the effect of a counterfactual reform that eliminates Social Security payroll taxes for older workers.

JEL Classifications: J26, J32, J41, M52, D15, H24

Keywords: labor supply, pensions, retirement savings, Social Security payroll tax

Dhiren Patki is an economist in the Research Department of the Federal Reserve Bank of Boston. His email address is Dhiren.Patki@bos.frb.org.

The author is indebted to Matthew Shapiro, Charles Brown, Dana Muir, and Melvin Stephens for advice and support. He thanks Ariel Binder, John Bound, R. Andrew Butters, Chichun Fang, George Fenton, Jamie Fogel, Matthew Gross, Gábor Kézdi, Edward Lazear, Parag Mahajan, Stephanie Owen, Kathryn Shaw, Andrew Simon, Jonathan Skinner, Ellen Stuart, and seminar participants at the NBER Summer Institute, NBER Conference on Employer Challenges in Financing and Managing Pension Plans, Boston College Center for Retirement Research, Federal Reserve Bank of Dallas, Federal Reserve Bank of Boston, Michigan State University, University of Michigan, and the NBER Conference on Incentives and Limitations of Employment Policies on Retirement Transitions for valuable comments. The author is grateful to Kristin McCue for her help with the U.S. Census Bureau disclosure review process.

This research was supported in part 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. It was also supported through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor. This research uses data from the U.S. Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by 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.

This paper presents preliminary analysis and results intended to stimulate discussion and critical comment. The views expressed herein are those of the author and do not indicate concurrence by the Federal Reserve Bank of Boston, the principals of the Board of Governors, or the Federal Reserve System.

This paper, which may be revised, is available on the website of the Federal Reserve Bank of Boston at <https://www.bostonfed.org/publications/research-department-working-paper.aspx>.

1 Introduction

The aging of the baby boom generation poses stark challenges that affect not only the financial future of retirees but also the overall performance of the economy. As a large share of the population permanently withdraws from the labor force, economic growth will slow, more individuals will claim Social Security and Medicare benefits, and fewer workers will pay into these social insurance programs. To mitigate some of these adverse effects, researchers have proposed policies that incentivize delayed retirement by altering the tax code or by changing the structure of public and private pension benefits.

Predicting retirement responses to policy-induced changes in compensation structure is subject to two empirical challenges. The first challenge relates to defining the horizon that workers consider when making retirement decisions. [Krueger and Pischke \(1992\)](#), [Rogerson and Wallenius \(2013\)](#), [Brown \(2013\)](#), and other studies model retirement in the context of contemporaneous changes in compensation. In contrast, [Stock and Wise \(1990\)](#) emphasize that the retirement decision is determined not only by contemporaneous compensation but also by the entire path of expected future compensation.¹ The role of long horizons in determining retirement is important and has been validated empirically using both reduced-form and structural methods (see, for example, [Coile and Gruber \(2007\)](#) and references therein). Nevertheless, a limitation of studies that do adopt a forward-looking view on retirement behavior is that these studies use either cross-sectional variation or panel fixed effect designs. These empirical approaches can be problematic because within- and between-person differences in future compensation are likely correlated with unobserved determinants of lifetime labor supply. Closely related to this concern, the second challenge stems from the fact that unanticipated changes in compensation generate both wealth and substitution effects, and the identification of the structural parameters governing these effects is difficult when relying on panel data alone (see, for example, [MaCurdy \(1981\)](#)).²

In this paper, I address both challenges. I use a novel source of variation to identify a model of retirement behavior in which workers' retirement decisions depend on both current and expected future compensation. My identification strategy relies on the large-scale restructuring of U.S. private sector defined benefit (DB) pension plans that began in the early 2000s. Since then, many employers have reneged on long-standing promises to their workers by eliminating DB pension accruals and replacing them with defined contribution (DC) and cash balance (CB) pension plans.³ These actions, known as pension freezes, allow workers to keep previously earned pension benefits though

¹Implicit contracts that solve agency problems are a powerful force connecting current labor supply decisions to future compensation. See [Lazear \(1979\)](#), [Lazear \(1981\)](#), and [Akerlof and Katz \(1989\)](#) for important theoretical contributions.

²[Manoli and Weber \(2016\)](#) exploit quasi-experimental variation in anticipated retirement benefits and study the effects of these incentives on labor supply. They model retirement using the forward-looking framework of [Stock and Wise \(1990\)](#) but do not estimate the structural parameters of the model.

³The most common type of DC plan is a 401(k) plan. Unlike traditional DB pensions, which guarantee workers an income stream in retirement, DC plans allow workers to accumulate retirement savings in tax-preferred accounts. Employers typically provide incentives for participation in these plans by matching worker contributions. CB plans are functionally similar to DC plans. More details are provided in [Section 2](#).

unexpectedly change the present value of future compensation that workers have been promised but are yet to be paid.⁴ From a research perspective, pension freezes are useful because they generate unanticipated shifts in workers' age-compensation profiles that induce both wealth and substitution effects.⁵ As I describe below, this source of variation credibly identifies preference parameters governing retirement behavior, thereby providing a better way to predict how older workers will react to unexpected policy-induced changes in compensation structure.

To study why employers have increasingly restructured their DB plans and to examine how employees respond to these changes, I create a new data set that matches 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. My research design pools thousands of firm-level natural experiments and compares the labor supply behavior of workers whose employers freeze their DB plans with workers whose employers keep their DB plans intact.

On the firm side, I provide evidence that pension freezes are driven primarily by plan funding deficiencies and are unrelated to mass layoffs or employee age structure. On the worker side, I show that the retirement response to these shocks varies in two important dimensions. First, workers initially affected at different ages exhibit differing retirement responses because they experience compensation losses of varying magnitude. Second, (holding age 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 is therefore indicative of heterogeneity in wealth and/or preferences for leisure.

The administrative data that I use describe employment responses to pension freezes. These data do not include information on the dollar value of compensation changes induced by pension freezes. To estimate the pecuniary effects of pension freezes and to conduct counterfactual simulations, I develop and estimate a structural model of retirement timing and saving that relies on rich survey data from the Health and Retirement Study (HRS).⁶ The model incorporates heterogeneity in wealth and in preferences for leisure and allows work decisions to depend not only on current compensation but also on the option of higher earnings and increased pension benefits in the future, as in [Stock and Wise \(1990\)](#). Unlike [Stock and Wise \(1990\)](#), however, the model allows for saving in multiple asset types, includes Social Security benefits, and features nonlinear taxes.⁷ I fit the model using the method of simulated moments (MSM) by matching quasi-experimental employment responses

⁴The abrupt nature of these changes is reflected in the sentiment of one disgruntled Verizon employee who, after hearing of the company's decision to freeze his DB pension, said "Oh, it's outrageous that they ... *change the rules in the middle of the game.*" See "Verizon Unveils Major Changes to Retirement Benefits," National Public Radio, *All Things Considered*, December 6, 2005.

⁵[Krueger and Pischke \(1992\)](#), [Brown \(2013\)](#), and [Gelber et al. \(2016\)](#) also rely on unanticipated changes in pension compensation to study retirement behavior. These studies do not use variation in future compensation to identify a structural model or estimate the effect of counterfactual policies.

⁶The HRS is a large-scale longitudinal household survey of Americans over age 50. It tracks labor market transitions, household finances, and health characteristics, and it includes linkages to administrative data from the Social Security Administration (SSA).

⁷The model I use is a stochastic dynamic programming model of retirement and not an option value model as pioneered by [Stock and Wise \(1990\)](#). [Lumsdaine et al. \(1992\)](#) provide a comparison of the two approaches.

to pension freezes as observed in the administrative data.

The model that I estimate highlights the importance of forward-looking behavior in governing the retirement decision and lends itself to evaluating the impact of counterfactual policies aimed at incentivizing delayed retirement through a permanent shift in future compensation trajectories. I examine one such proposal through the lens of the model. In particular, I consider eliminating the Social Security, or Old Age and Survivors Insurance (OASI), component of the payroll tax for workers over the age of 60. This reform has been proposed as a way of lengthening lifetime labor supply by removing contribution requirements for workers who are fully vested in Social Security benefits (see, for example, [Burtless and Quinn \(2002\)](#) and [Goda et al. \(2009\)](#)). Simulations from the estimated model predict that an unexpected elimination of the OASI payroll tax on workers when they reach age 60 causes employment rates to rise by an average of approximately 4 percentage points over a 20-year horizon, which equates to a 10 month increase in the average retirement age. Longer careers allow workers to accumulate about \$5,700 more in retirement savings, and equivalent variation from the reform averages \$54,000 per worker.⁸

In this paper, I make three contributions to the literature. First, I provide the only available evidence of how older Americans have re-timed their retirement decisions in response to widespread DB pension freezes. These analyses, based on high-quality administrative data, shed new light on the ongoing transition from DB to DC pension provision. While studies examining the incentive effects of DB and DC pension plans emphasize that DC plans encourage longer careers (see, for example, [Friedberg and Webb \(2005\)](#), [Hurd and Rohwedder \(2011\)](#), and [Coile \(2021\)](#)), the findings I present show that unexpected transitions from DB to DC plans generate a mixed response: Some workers choose to shorten their careers, whereas others choose to lengthen them. Second, I estimate a model of retirement timing and saving that is identified by plausibly exogenous variation in future compensation arising from pension freezes. This natural-experiment-based identification strategy represents a departure from previous structural models of retirement that rely on cross-sectional or panel data to isolate variation in compensation (see, for example, [Stock and Wise \(1990\)](#), [Berkovic and Stern \(1991\)](#), [Rust and Phelan \(1997\)](#), [French \(2005\)](#), and [French and Jones \(2011\)](#)). Finally, I provide evidence about the effectiveness of a counterfactual reform to payroll taxes designed to lengthen careers. In contrast to previous efforts evaluating the effect of payroll tax sunsets using structural models ([Laitner and Silverman \(2012\)](#) and [Gustman and Steinmeier \(2015\)](#)), the model and identification strategy that I develop explicitly accounts for the long-term option value channel of continued employment and also quantifies the effect of the reform on retirement saving.

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 firms' decisions to freeze their plans. Section 3 outlines a model of retirement timing and saving and uses the model to evaluate how workers respond to pension freezes. Section 4 discusses three different sources of administrative data used in the analyses. Section 5 outlines the empirical framework and presents summary statistics. Section 6 provides empirical evidence on workers'

⁸Estimated in 2010 dollars with present values calculated at age 60.

labor supply responses to pension freezes. Section 7 explains the identification and estimation of the structural model and presents parameter estimates. Section 8 evaluates the counterfactual OASI payroll tax sunset. Section 9 concludes.

2 Why are employers freezing DB plans?

In 1980, DB plans covered 61 percent of pension-eligible private sector workers. By 2018, the same statistic had fallen to 14 percent.⁹ 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.¹⁰ Starting in the late 1990s, however, several prominent firms began converting their DB plans to CB plans. CB conversions switched pension accruals away from formulas based on years of service and earnings to account-based plans providing participants with employer contributions that were proportional to earnings and a market-linked rate of return on previous contributions.¹¹ In the early 2000s this shift was amplified, as many employers ended traditional DB pension accruals altogether in actions known as hard freezes.¹² The incidence of CB conversions and hard freezes from 1999 through 2019 is shown in Figure 1. By 2019, about half of all private sector single-employer DB plans had been either converted to CB plans or hard-frozen, affecting more than 6.3 million workers.¹³

2.1 Costs of DB plan provision have become increasingly volatile

Firms' 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 reduced the asset value of pension funds substantially. In 2000, the private sector DB system had \$1.44 in assets for each dollar of future liabilities. By 2004, the funding ratio had slipped to \$0.85 for each dollar. As the funding

⁹See Table E7, Private Pension Plan Bulletin Historical Tables and Graphs, U.S. Department of Labor, 2018.

¹⁰The Omnibus Budget Reconciliation Act of 1987 reduced the ability of DB-sponsoring employers to take tax deductions for pension contributions. This change led to a spike in non-distress terminations from 1987 through 1990 (see, for example, Table A-9 in Pension Insurance Data Book 1996, PBGC Single Employer Program <https://www.pbgc.gov/documents/1996databook.pdf>).

¹¹CB plans are functionally very similar to DC plans: They provide participants with individual accounts where annual accruals are tied to annual earnings. There are, however, two differences. First, CB plans allow for larger maximum pre-tax deferrals compared with DC plans. Second, it is common in CB plans for employers to bear some interest rate risk by promising a minimum rate of return on the value of the account. For these reasons, CB plans are subject to the same funding obligations required of DB plans and are legally treated as DB plans.

¹²Over the same period, employers also began closing their DB plans to new entrants in actions known as soft freezes. I do not consider soft freezes in this paper.

¹³Hard freezes are exceedingly rare in the public sector, affecting less than 0.5 percent of eligible state and local government workers. The main reason for this difference owes to differing legal protections in the private and public sectors. The Employee Retirement Income Security Act of 1974 (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. See, for example, Monahan (2012).

position of DB plans worsened, statutory provisions required firms to increase annual contributions; consequently, aggregate payments into DB pension funds rose nearly five-fold, from \$26 billion in 2000 to \$124 billion in 2003. Low interest rates and stock market losses during the Great Recession weakened firms' funding positions, prompting further increases in required pension contributions. These shocks disproportionately affected plans with marginal funding status, as they were not buffered against large contributions requirements in the same way that overfunded plans were. The role of worsening plan finances as a key predictor of subsequent pension freezes is demonstrated in greater detail in Appendix C, which relies on firm-level microdata.¹⁴

In 2006, at a time when the finances of DB pension funds were undergoing major changes, the Pension Protection Act (PPA) was signed into law. The PPA established more conservative standards for the interest rates that sponsors could use to discount future liabilities, reduced the length of the period during which sponsors needed to amortize funding deficits from 15 years to 7 years, and required that plans with funding ratios of 80 percent or lower make additional minimum contributions and pay higher insurance premiums to the Pension Benefit Guarantee Corporation (PBGC).¹⁵ 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 with marginal funding status.

2.2 Legal constraints to freezing plans have been alleviated

From 1998 through 2000, a handful of prominent employers converted traditional DB plans to CB plans. Older employees, who stood to lose substantial future pension accruals as a result of 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, for example, [Zelinsky \(2000\)](#)). As these cases played out in the court system, the legality of CB conversions remained uncertain. The threat 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.¹⁶ That same year, the new PPA law provided criteria 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.

¹⁴[Munnell and Soto \(2007\)](#), [Bovbjerg et al. \(2008\)](#), and [Rauh et al. \(2020\)](#) also find that cost savings and funding volatility are important determinants of firms' decisions to freeze their DB plans. A compliance check on frozen DB plans conducted by the IRS in 2012 and obtained through a Freedom of Information Act (FOIA) request for this paper further corroborates this conclusion, finding that sponsors typically froze their plans due to funding deficiencies.

¹⁵Some 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 from 2008 through 2011. The extensions were for 9 or 15 years rather than the 7-year requirement of the PPA.

¹⁶See, for example, Mary Williams Walsh, "Issues Left Unresolved on Pensions," *New York Times*, Jan. 17, 2007.

3 Model of retirement and saving

This section develops a model of retirement timing and saving. Using the model, I show that pension freezes generate both a substitution effect and a wealth effect and therefore have an ambiguous impact on retirement behavior. I explain how heterogeneity in wealth and in preferences for leisure implies differential retirement responses for workers faced with freeze-induced losses in future compensation. I summarize these comparative statics into three empirically testable implications.

3.1 Defined benefit pension incentives

Before characterizing the model, I illustrate key features of DB plans and their effect on total compensation using data from the Health and Retirement Study (HRS). The left panel of Figure 3 shows the age-compensation profile for the average HRS respondent aged 50 to 70 who is included in a non-frozen DB plan. The dashed line shows earnings from wage and salary compensation. The solid line adds pension accruals to wage and salary compensation. Note that about 40 percent of DB-eligible respondents have non-zero DC wealth; DC accruals for these respondents are included in the overall pension accrual estimates.¹⁷ Two types of nonlinearities typical of DB pensions are evident in the figure. First, DB accruals generate prominent spikes in compensation. At age 55, which is a common early retirement age (ERA) in many DB plans, workers can begin claiming benefits, thereby generating a substantial jump in pension wealth. After the ERA, accruals grow at a slower rate as workers approach the normal retirement age (NRA), which is typically 65. After the NRA, it is common for DB accruals to no longer grow in a manner that is actuarially fair and, in fact, become negative in real terms. This downturn occurs because workers who postpone retirement do not obtain sufficiently large increases in pension benefits even though they will obtain those benefits for fewer (mortality adjusted) years.¹⁸

The profiles shown in the left panel of Figure 3 illustrate how DB plans generate strong incentives for workers to remain employed through the ERA in order to collect generous retirement benefits. In contrast, DB plans generate strong incentives for workers to retire once they reach age 65. Underpinning these carefully designed incentives is an implicit contract between firms and workers: Firms make long-term unenforceable commitments to workers in which continued employment through retirement age will be rewarded with lucrative pension benefits. Workers, on the other hand, exert effort over the course of their careers and avoid early termination that would result in the loss of promised pension benefits (see, for example, Lazear (1981), Kotlikoff and Wise (1985), and Lazear and Moore (1988) for theory and evidence). By freezing DB pensions, firms are unilaterally renegeing on this implicit contract.

¹⁷Earnings 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.

¹⁸This implicit tax on continued work is a common feature of DB pensions specifically designed to encourage retirement (see, for example, Chapter 12 in Gustman et al. (2000)). In the figure, total pension compensation is positive for workers over 65 only because DC accruals for the 40 percent of respondents who have non-zero DC wealth offset negative DB accruals.

The following section outlines a model of retirement in the presence of DB pension promises and uses the model to understand how worker behavior changes in response to pension freezes.

3.2 Setup

Individuals are heterogeneous in age ($a \in \{50, 51, \dots, 80\}$), non-pension wealth $A_a \in [0, \infty)$, DC pension wealth $W_a^{DC} \in [0, \infty)$, and the extent to which they dislike working, g . For individual ι currently aged a , the disutility of working is modeled as

$$g_{\iota a} = \underbrace{\gamma + \phi a}_{\text{deterministic component}} + \underbrace{f_{\iota a}}_{\text{random component}} \quad (1)$$

$$f_{\iota a} = \rho f_{\iota, a-1} + v_{\iota a}, \quad (2)$$

where $v_{\iota a} \sim \mathcal{N}(0, \sigma_v^2)$ is an IID disturbance term. The deterministic component of g captures changes in preferences for leisure that are common to workers of the same age, whereas the random component captures individual specific differences in health, disability status, and non-pay aspects of employer-employee match quality. I assume that utility from consumption exhibits constant relative risk aversion (CRRA) and that consumption and leisure are separable. Under these assumptions, flow utility for an individual ι aged a is given by

$$\frac{c_{\iota a}^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - g_{\iota a} \times \mathbf{1}\{\iota \text{ is working at age } a\}. \quad (3)$$

Earnings at age a are given by e_a , which is taxable in the period that it is earned. Conditional on remaining employed, all workers share the same deterministic age-earnings profile, with earnings at age a given by e_a . DC wealth evolves according to

$$W_{\iota a}^{DC} = W_{\iota, a-1}^{DC}(1+r) + e_a(m_{\iota a}^w + m^e(m_{\iota a}^w)) \quad (4)$$

$$\bar{C} \geq e_a(m_{\iota a}^w + m^e(m_{\iota a}^w)), \quad (5)$$

where r is the interest rate, and $m_{\iota a}^w$ is the fraction of earnings that a worker defers toward their DC account on a pre-tax basis. $m^e(m_{\iota a}^w)$ is the fraction of earnings that the worker's employer contributes to their DC account. m^e is expressed as a function of $m_{\iota a}^w$, reflecting commonly used employer incentives for participation in DC plans. IRS rules limit total contributions to less than a threshold value \bar{C} .¹⁹

Workers earn DB pension wealth on the basis of a deterministic formula that is set by firms. The formula takes earnings history, tenure, and age as its arguments and returns the value of DB

¹⁹In 2010, which is the year to which all dollar values in the paper are inflated or deflated to, 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.

pension wealth:

$$W_a^{DB} = \mathcal{B}(\text{earnings, tenure, age}). \quad (6)$$

If a worker chooses to retire at age a , they obtain the annuity value of their DB wealth, which is b_a^{DB} . Workers also accrue Social Security benefits based on their earnings history in the form of an annuity whose value at worker's age a is given by b_a^{SS} . Like the DB annuity, Social Security can be claimed only in retirement. $b_a = b_a^{DB} + b_a^{SS}$ is the total annuity income for an individual who chooses to retire at age a .²⁰ I assume that individuals cannot hold debt, so $A_{\iota a} \geq 0$, and $W_{\iota a}^{DC} \geq 0$.

Total compensation for worker ι currently aged a is

$$\chi_{\iota a} = \underbrace{e_a(1 - m_{\iota a}^w)(1 - \tau)}_{\text{non-deferred compensation}} + \underbrace{\Delta W_{\iota a}^{DC}}_{\text{DC accrual}} + \underbrace{\Delta W_a^{DB}}_{\text{DB accrual}}, \quad (7)$$

where τ is the tax rate applied to income after pre-tax deferrals have been made.²¹ Equation (7) reflects the assumption that workers have passed the age where real earnings growth contributes to greater Social Security benefits. Consequently, there are no increments to Social Security wealth from additional years of employment.

3.3 Value of retirement

To conserve notation, denote age a variables by v and age $a + 1$ variables by v' . The value of retirement at age a is V_a^R , which is written recursively as

$$V_a^R(b, A, W^{DC}) = \max_{A', W^{DC'}} \left\{ u(c) + \beta(1 - p_a)V_{a+1}^R(b, A', W^{DC'}) \right\} \quad (8)$$

$$\text{s.t. } c \leq \left(b + W^{DC} - \frac{W^{DC'}}{1+r} \right) (1 - \tau) + A - \frac{A'}{1+r}. \quad (9)$$

In equation (8), $u(c)$ is utility from consumption, β is the annual discount factor, and p_a is the subjective probability of death within one year for an individual currently aged a . Annuitized pension income (b) is fixed at the age of retirement, which is why b rather than b' appears in the continuation value. Given the state variables and the budget constraint, the retirees consumption choice is recast as the choice of how to decumulate pension and non-pension wealth at each age. Pension income, which is the sum of annuitized income and liquidated DC pension wealth, is taxable, while liquidated non-pension wealth is not taxable.

Notably, the value function treats retirement as a self-absorbing state that excludes the possibility of short-term bridge employment. The assumption of self-absorbing retirement is important in the context of this paper because it makes future compensation contingent on current labor supply

²⁰The earliest age at which workers can claim Social Security benefits is 62. When estimating the model, I assume that workers who choose to retire before age 62 start claiming Social Security benefits once they turn 62.

²¹Income tax can be reduced by elective deferrals to DC accounts, but payroll taxes apply to all earned income up to the relevant earnings caps.

in a very stark way: If a worker retires, their future labor market options are entirely foreclosed. I provide empirical evidence in Section 6.1 that workers affected by pension freezes tend to follow a once-and-for-all retirement pattern, which is consistent with the modeling assumption I make here.

3.4 Value of working

Denote the age a vector of state variables for a working individual by $X = (e, g, b, W^{DC}, A)$. The value of working at age a is V_a^W , which is written recursively as

$$V_a^W(X) = \max_{A', m^w} \left\{ u(c) - g + \beta(1 - p_a)E_{g'} \left[\max \left\{ V_{a+1}^R(b', A', W^{DC'}), V_{a+1}^W(X') \right\} \right] \right\} \quad (10)$$

$$\text{s.t. } c \leq e(1 - m^w)(1 - \tau) + A - \frac{A'}{1 + r}. \quad (11)$$

In equation (10), flow utility from consumption is offset by the disutility of working g . The expectation operator for the continuation value term integrates over the random variable g' . By remaining employed at age a , workers preserve the possibility of obtaining higher earnings and higher pension accruals in the future, which is reflected in the continuation value. Given the state variables, the budget constraint (11), and the DC accumulation constraints (4) and (5), workers choose how much non-pension wealth to accumulate and what fraction of their earnings to defer toward DC pensions. I assume that DC pension wealth is illiquid until retirement.²²

3.5 Optimal retirement decision

A worker retires when $V_a^R \geq V_a^W$. This condition defines a cutoff value \bar{g}_a so that any draw of $g_a \geq \bar{g}_a$ will lead to retirement. Solving for \bar{g}_a yields

$$\bar{g}_a = u(c_a^W) - u(c_a^R) + \beta(1 - p_a) \times \quad (12)$$

$$\left[\underbrace{E_{g'} \left[\max \left\{ V_{a+1}^R(b', A', W^{DC'}), V_{a+1}^W(X') \right\} \right]}_{\text{1: Continuation value of working}} - \underbrace{V_{a+1}^R(b, A', W^{DC'})}_{\text{2: Continuation value of retiring}} \right]_{\text{Option value of working one more period}},$$

where c_a^W and c_a^R are the optimal consumption choices in the work and retirement states at age a . Consider terms 1 and 2 on the right side of equation (12). Higher future earnings and the potential for higher pension accruals raise term 1 and increase the incentive for continued work. On the other hand, higher levels of retirement wealth raise 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}_a .

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

²²IRS rules allow penalty-free distributions from DC accounts after age $59\frac{1}{2}$. Distributions taken before $59\frac{1}{2}$ are subject to a 10 percent tax. As of 2020, distributions must start at age 72. I do not incorporate these institutional features into the model.

the firm: By remaining employed in the current period, the worker obtains earnings growth and increased pension benefits in the future. Furthermore, because DB pension wealth is back-loaded and cannot be ported between employers, the option value of working has a particularly strong firm-specific component. Following [Stock and Wise \(1990\)](#), work decisions in this model vary not only due to period-by-period changes in compensation, but also due to changes in the option value of continued employment, which is embedded in term 1.

3.6 The effect of a pension freeze is theoretically ambiguous

When employers renege on long-term promises by freezing DB pensions, workers experience changes to the total value of compensation that they had previously expected to earn by retirement. The effect of these shocks is illustrated in the right-hand panel of [Figure 3](#). The solid blue line shows the path of total compensation in a world where DB plans are kept intact (that is, no pension freezes occur). When a DB plan is frozen, workers keep all previously earned benefits but no longer earn any new DB accruals. Instead, workers who do not already participate in DC plans are offered the chance to do so after the freeze. These changes are shown using dashed red lines that simulate the effect of a DB pension freeze for workers aged 52 and 66. The post-freeze simulations shown in the figure assume that respondents without DC wealth begin actively participating in a hypothetical new DC plan after the freeze.²³

The figure highlights that pension freezes shift the age-compensation profile 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 62. In contrast, the average 66-year-old experiences a gain. The main reason that pension freezes have positive effects on compensation over the long term is because workers with DC plans can continue to accumulate retirement wealth even at ages where traditional DB plans would penalize continued work.²⁴

Now consider terms 1 and 2 on the right side of [equation \(12\)](#) for a worker who is under age 65 when their plan is frozen. When everything else is held fixed, the loss in DB accruals generates two opposing effects. First, reductions to b^{DB} lower the return to working (term 1), thereby decreasing the option value of working and generating a substitution effect on labor supply. Second, reductions to b^{DB} make workers poorer in retirement (term 2), thereby increasing the option value of working and generating a wealth effect on labor supply. The same mechanisms work with opposite signs on a worker who is over age 65 when their plan is frozen: Freezes reduce the penalty for continued work, thereby increasing labor supply through the substitution effect. At the same time, greater pension wealth accumulation raises the continuation value of retirement, thereby lowering labor

²³Pension accruals drop sharply after a freeze because private sector DB plans do not require worker contributions, whereas DC plans almost always do. Thus, in order to offset lost accruals from a frozen DB plan and receive matching employer contributions, freeze-affected workers must defer some of their earnings.

²⁴In the post-65 age range, there is a second reason that freezes generate positive effects on total compensation. DB formulas often link benefits to average earnings over the last few years of a worker's career. Because earnings fall after age 60, a worker in a non-frozen plan will see their benefit shrink as their earnings decline, while a worker whose plan is frozen will see their benefit stay fixed. This positive effect on total compensation is not operative in plans that base benefits on years of highest earnings.

supply through the wealth effect. Since the substitution effect and wealth effect work against each other, the impact of freezes on retirement behavior cannot be signed.

Despite the overall theoretical ambiguity, two sources of heterogeneity allow me to characterize the following testable predictions about how pension freezes affect labor supply.

1. Between-age heterogeneity

- (a) *The relative strengths of the wealth effect and the substitution effect vary by age:* When workers under 55 have their plans frozen, they have large accruals still outstanding and experience substantial losses in compensation. When workers over 55 have their plans frozen, they have less left to earn from their DB pensions and therefore experience smaller losses in compensation. Therefore, wealth effects are more important for workers under 55. Conversely, substitution effects are more important for workers over 55.
- (b) *The effect of freezes on employment is reversed after age 65:* Because workers over the age of 65 experience increases rather than decreases in deferred compensation, their initial labor supply response should be reversed in sign relative to workers under the age of 65.

2. Within-age heterogeneity

- (a) *Short-term responses are dominated by the substitution effect:* Short ex ante work horizons reduce the role of the wealth effect since workers are about to retire anyway. Thus, the response of workers near the margin of retirement at the time of the freeze (that is, those with high values of g_{ia} , A_{ia} , and W_{ia}^{DC}) is dominated by the substitution effect. These workers will choose leisure over work when faced with a freeze.

4 Data

In this section, I describe three data sources that I use to estimate the effect of freezes on worker behavior. Pension plan data come from IRS Form 5500, employer characteristics come from the Census Longitudinal Business Database (LBD), and matched employer-employee data come from the Longitudinal Employer Household Dynamics (LEHD) data set. I am able to link 92 percent of plans from Form 5500 to employers in the LBD and 89 percent of those LBD employers to the LEHD data set.²⁵ Detailed descriptions of the data sets and the linking procedures that I use are provided in Appendix A.

4.1 Pension characteristics: Form 5500

Form 5500 (F5500) is an annual plan-specific filing that is collected jointly by the IRS, Department of Labor (DoL), and the Pension Benefit Guarantee Corporation (PBGC). These publicly available data contain rich information on the universe of privately sponsored pension plans. When DB plans

²⁵These data sets are used in conjunction to study the association between fringe benefits and employee mobility in [Decressin et al. \(2009\)](#).

are converted to CB plans or hard-frozen, this information is reported on F5500, thereby allowing me to identify plans whose participants are affected by freezes.²⁶ F5500 data also include important pension plan characteristics such as the number of plan participants, the present value of pension wealth for plan participants, the value of pension accruals earned by participants in the filing year, and the typical age at which participants claim benefits from the plan.

F5500 filings are identified by a combination of a Federal Employer Identification Number (EIN) and an employer designated plan number (PN) that remain consistent over time. While these identifiers are sufficient to match pension plans to single-unit firms (that is, operating only one establishment), they are not sufficient for matching to multi-unit firms (operating multiple establishments). This is because payroll tax filings for establishments that are part of a multi-unit firm may be recorded under EINs that are different from the one used in F5500. Attempting to match the F5500 data 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).

4.2 Firm characteristics: Census Business Register and Longitudinal Business Database

The Census Business Register (BR) is a database of the universe of establishments in the United States. It includes information on business location, organization, industry, revenue, payroll, and employment. The relationship between establishments belonging to multi-unit firms is determined using responses to the Census Bureau’s Company Organization Survey (now the Report of Organization), Economic Census, and 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.²⁷ Having matched F5500 records to the BR, I use the Census firm 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.

4.3 Worker characteristics: Longitudinal Employer Household Dynamics

To study outcomes at the individual level, I turn to the Longitudinal Employer Household Dynamics (LEHD), which is a quarterly matched employer-employee data set constructed from state-level unemployment insurance (UI) records.²⁸ 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 UI account number known as the state employer identification number (SEIN). I rely

²⁶Plans are considered frozen in F5500 if they meet the following condition: “As 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).”

²⁷In the absence of the Census firm identifier, I would be able to identify only those multi-unit establishments that share the same EIN as the one reported on F5500, thereby generating a false non-match problem alluded to earlier.

²⁸See [Abowd et al. \(2009\)](#) for details.

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.²⁹

An important feature of the LEHD is that states become part of the data set at different points in time. For example, Maryland enters in 1985:Q2, whereas Mississippi doesn't enter until 2003:Q3. Because of staggered entry, the scope of the data grows continuously over time.³⁰ As a consequence, when I link an employer from the LBD to the LEHD in a given year, I capture only those individuals who work in a state that has already entered the data set as of that year.

4.4 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 the F5500 database from 1996 through 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.³¹

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.³² 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. Consequently, 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 through 2014 as an experiment year, which is indexed by l .³³ This terminology reflects the research design wherein each experiment year yields a fresh sample of firm-level pension freezes. Workers employed at freezing firms five years before a given experiment year constitute the treated group, while workers employed at non-freezing firms five years before a given experiment year constitute the comparison group. I impose the restriction that firms file F5500 for the five calendar years preceding the experiment year, which I refer to as the pre-period. I then match the firm-experiment year data to the LEHD, the results of which are shown in the third row of Table

²⁹I 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.

³⁰Several populous states enter the data relatively early. Illinois enters in 1990, California and Pennsylvania in 1991, Florida in 1992, and New York and Texas in 1995.

³¹F5500 reports the count of active participants—that is, covered workers—in each plan.

³²Firms 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.

³³I 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 before 2001.

A1. I recover 89 percent of firm-experiment years and 93 percent of employee-experiment years. Finally, as shown in the fourth row of Table A1, I restrict the sample to firm-experiment years for which important pension plan data are not missing.³⁴

Regarding the implications of pension freezes on worker decisions, it is worth reiterating 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.³⁵ Within the high-coverage-rate firms, I select all workers employed at $l - 5$ who have at least two years of tenure as of $l - 5$ and who will be aged 50 to 70 in year l .³⁶

Having discussed the key sources of data and explained how I link and organize these data for analysis, I turn next to the regression framework and associated identification assumptions that I use to estimate causal effects.

5 Empirical framework for estimating treatment effects

This section describes the regression framework I adopt to analyze the impact of pension freezes on a variety of labor market outcomes. The identifying assumption is that pension freezes are independent of unobserved determinants of labor supply, conditional on a set of worker characteristics, firm characteristics, and fixed effects. Summary statistics lend credibility to the identifying assumptions.

5.1 Regression specification

Let i index firms, let j index cells that bin together workers in the same firm-state-gender-age-tenure-experiment year, and let t index calendar years. Let $k = t - l$ index years relative to the experiment year. Within each cell, the firm 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 experiment year, l :

$$y_{j(i)t}^l = \alpha_i + \gamma_t^l + \mathbf{x}'_{j(i)t} \boldsymbol{\beta}^l + \sum_{k=-5}^{m(l)} \delta_k^l T_{ik}^l + \varepsilon_{j(i)t}^l, \quad (13)$$

³⁴Important pension plan data include plan assets and liabilities, accruals earned during the plan year, and the typical benefit claim age for the plan.

³⁵The 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 $[l - 5, l - 2]$, that is five to two years preceding the experiment year. 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 and are not replaced with new, DB-eligible workers.

³⁶The two-year tenure restriction ensures that workers are fully vested in their pensions in the year when they may become subject to a freeze. This calculation is based on the seven-year maximum 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 when the state in which they work enters the data.

where $y_{j(i)t}^l$ measures a labor supply outcome of interest, α_i is a firm fixed effect, γ_t^l is a calendar-year fixed effect, and $\mathbf{x}_{j(i)t}^l$ is a vector of controls for age, gender, state, race, education, tenure, and previous earnings.³⁷ T_{ik}^l is an indicator variable that equals 1 if the firm freezes its plan, and the current period is k . $\varepsilon_{j(i)t}^l$ is the error term, which represents unobserved determinants of labor supply.³⁸ The parameters of interest are the δ_k^l coefficients, which capture the dynamic treatment effect of pension freezes on worker outcomes.

To maximize the precision of the estimates, I stack data from each of the experiment years together and estimate a version of equation (13) where

$$y_{j(i)t}^l = \alpha_i + \gamma_{lt} + \mathbf{x}_{j(i)t}^l \boldsymbol{\beta} + \sum_{k=-5}^{13} \delta_k T_{ik}^l + \varepsilon_{j(i)t}^l. \quad (14)$$

In equation (14), calendar-year fixed effects are replaced by experiment year-by-calendar-year fixed effects, which allow economy-wide shocks to differentially affect workers in each experiment year. In contrast, the effect of the $\mathbf{x}_{j(i)t}^l$ and T_{ik}^l variables are assumed to be constant across experiment years. This estimation strategy allows workers in the comparison group in a given experiment year to enter the treated group in a subsequent experiment year if their employer freezes pensions in the future. In this implementation, the δ_k coefficients are identified by within-experiment-year between-firm variation in worker outcomes as well as within-firm between-experiment-year variation in worker outcomes. Standard errors are clustered at the firm level.

5.2 Identification

The δ_k parameters in equation (14) represent causal effects of pension freezes on worker outcomes under the assumption that $E[\varepsilon_{j(i)t}^l | \alpha_i, \gamma_{lt}, \mathbf{x}_{j(i)t}^l, T_{it}^l] = 0$. Put differently, unobserved determinants of worker labor supply are assumed to have zero mean conditional on firm fixed effects, experiment year-by-calendar-year fixed effects, worker-level controls, and the freeze indicators. This assumption might be violated by 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.³⁹ Second, it is possible that freezing and non-freezing firms systematically differ in terms of pension generosity and benefit-claiming provisions 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

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

³⁸The upper limit of the sum, $m(l)$, represents the number of available post-period years for experiment year l . The maximum available post-period duration is 13 years (this happens when $l = 2001$, as the data run out in 2014).

³⁹Along with a broader set of statistics on firm dynamics around the freeze, I show in Appendix C that freezing and non-freezing firms do not differ in terms of their pre-period probability of experiencing distress.

the cell-level probability of experiencing a freeze expressed as a function of pre-period variables that influence both firms' decision to freeze and workers' labor supply responses. To mitigate concerns related to pre-freeze firm distress as 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, the age structure of employment at the firm, retirement rates, employment rates, and employer-to-employer transition (E-E) rates. Appendix D provides more details on the conditioning set and explains how the propensity scores are transformed into weights when equation (14) is estimated.

Beyond these key threats to the identification strategy, two other confounding effects are potentially at play. First, it is possible that funding deficiencies that led firms to freeze their DB plans also resulted in cutbacks to health insurance benefits. These unobserved changes could generate their own income and substitution effects on labor supply choices.⁴⁰ While these changes are not directly verifiable in the data that I use, indirect evidence from surveys suggests that firms have not altered health benefits as a consequence of freezes.⁴¹ Second, it is possible 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' retirement decisions are influenced not only by changes in compensation but also by changes to the retirement decisions of their peers. I ignore peer retirements as a first-order concern when interpreting the results because available evidence on the magnitude of peer effects of this type indicate that they are extremely small.⁴²

5.3 Summary Statistics

Before showing the impact of freezes on worker outcomes, I present a summary of raw data 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 C.

In Table 1 I show pre-period summary statistics for workers split into three different groups based on age as of the experiment year.⁴³ Workers in the sample are employed with DB-sponsoring firms

⁴⁰Employer-sponsored health insurance is reported in F5500 filings. However, the reporting requirement applies only to employers that cover more than 100 workers. More importantly, changes in health insurance plan characteristics cannot be ascertained from F5500.

⁴¹There are no reports of changes to employer-provided health insurance benefits in a sample of 17 large publicly traded firms that froze their plans during the 2004–2008 period, according to 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)).

⁴²Hamman 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 point and produces no detectable effect on women.

⁴³Using 55 and 65 as the modal ERA and NRA in DB plans, 50- to 55-year-old workers are below the ERA at the

as of $l - 5$, and the statistics are computed by averaging over 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 data set, pension and firm characteristics are worker weighted. For each age-specific panel, the first column shows the comparison group mean re-weighted using propensity scores, 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. Across the three sub-samples, workers are well educated and have relatively high earnings and the gender split is close to 50-50. Demographic characteristics, pre-period labor market outcomes, pension plan generosity, and firm-level characteristics are very similar in the treated and control groups.⁴⁴

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 likely to be white and male and more likely to have a college degree and higher earnings. Tenure, measured five years prior to the experiment year, is approximately equal across the three age groups.⁴⁵ 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. Appendix Table G1, which shows the same statistics without propensity score re-weighting, provides additional evidence of broad similarity between workers in the treated group and those in the control group.

6 How pension freezes affect labor supply and employer attachment

In this section, I use the regression framework and identification strategy presented earlier to investigate the causal impact of pension freezes on the labor supply and employer attachment. I show how freezes have heterogeneous labor supply effects based on the age at which workers experience them. In addition, holding age fixed, I show that freezes have heterogeneous effects. The treatment effects I present here are consistent with the testable implications developed in the theoretical model.

6.1 Employment, retirement, and earnings

Figure 4 plots the δ_k coefficients from specification (14) using the cell-level employment rate as the outcome variable. To investigate age-specific heterogeneity in labor supply responses, I split

time of the freeze, 56- to 64-year-old workers are between the ERA and the NRA at the time of the freeze, and 65- to 70-year-old workers are over the NRA at the time of the freeze.

⁴⁴It is important to note that the measures of pension wealth are based on plan-wide totals and should not be seen as representative of the workers in each age bin

⁴⁵The absence of variation in tenure between workers of different ages is an artifact of the way that states enter the LEHD data set. If a state enters the data set after a given employer-employee relationship is established, then the employer-employee history is left censored and tenure is understated.

the data by workers' ages as of the experiment year using the same age groups as in Table 1. I then estimate the regression model on each age group separately and show the coefficients in the respective panels of the figure.

Looking first at the far-left panel, it shows that treated workers in the 50- to 55-year-old age group exhibit a small reduction in employment rates in the first six years of the post-freeze period. Reductions in employment reflect substitution effects (that is, reduced lifetime labor supply), although the economically small magnitude of the coefficients and their statistical insignificance suggests that offsetting wealth effects (that is, increased lifetime labor supply) are equally important for workers in this age group. After about eight years post-freeze, wealth effects start to dominate the labor supply response, and workers in the treated group experience a 1.5 to 3.3 percentage point increase in employment relative to the comparison group. Another way of stating this finding is that the employment rates of both the treated group and the comparison group are declining in this age range, but the decline is markedly slower for treated workers. The muted substitution effect and substantial wealth effect that characterizes labor supply responses for 50- to 55-year-old workers aligns with the fact that large DB accruals are earned before age 55. Therefore, workers who are 55 or younger when first faced with a freeze experience large net losses in compensation, thereby eliciting strong wealth effects in favor of continued employment.

Moving next to the central panel of the figure, one sees that substitution effects play an important initial role for treated workers in the 56- to 64-year-old age group. In the first six years of the post-freeze period, employment rates for treated workers fall by 1.0 to 1.7 percentage points relative to the comparison group. Starting about eight years post-freeze, treated workers' labor supply diverges in the opposite direction of the comparison group's labor supply, as the employment rate differential rises by 1.2 to 2.1 percentage points. As with 50- to 55-year-old workers, the tendency of freeze-affected workers to extend their working lives relative to the comparison group is indicative of dominant wealth effects.⁴⁶

Most DB plans incentivize retirement by making the real value of DB accruals negative after age 65. When DB plans are frozen and replaced with DC plans, this implicit tax on continued employment is reduced, as workers can obtain offsetting DC accruals at ages when they would otherwise have been penalized. The far-right panel of Figure 4 shows labor supply behavior that is consistent with greater returns to work for the treatment group relative to the comparison group. Treated workers over age 65 increase their employment rates by 1.2 to 4.5 percentage points, which reflects dominant substitution effects. While these effects are not as precisely estimated, they are economically meaningful and align with theoretical predictions and the institutional design of DB plans.

Empirical evidence shown in Figure 4 lines up with the three theoretical predictions outlined in Section 3. Substitution effects are greater for 56- to 64-year-old workers relative to 50- to 55-year-old workers, which is a consequence of the former group losing less total compensation than the latter

⁴⁶Appendix Figure G1 shows that freeze-affected women are more likely to exhibit substitution-effect-dominant responses, whereas freeze-affected men are more likely to exhibit wealth-effect-dominant responses.

group (prediction 1[a]). Substitution effects are positive because of increased labor market returns for workers over 65, but they are negative for workers under 65 because of reduced labor market returns (prediction 1[b]). Finally, looking at the first two age groups, substitution effects play a dominant role in the short-term response of the most marginally attached workers, who reduce their employment rates (prediction 2[a]).

Figure 5 shows the impact of freezes on retirement, which is defined as a permanent departure from paid employment in the LEHD.⁴⁷ Freeze-induced changes in retirement rates are virtually mirror images of the employment effects shown in Figure 4, indicating that non-employment and retirement are essentially equivalent for workers affected by freezes. This finding indicates that bridge jobs appear not to be an important transition phase for freeze-affected workers who cut their careers short. It also substantiates the model’s assumption of self-absorbing retirement. On the whole, the treatment effects for retirement reinforce the employment-based findings.

Figure 6 shows the effect of freezes on the log of annual earnings. All the estimates shown here are conditioned on the sample of individuals with positive earnings in a given year. The far-left panel shows that there are no evident changes in earnings for 50- to 55-year-old workers in the first six to seven years of the post-freeze period. After seven years, workers in the treated group exhibit slower age-related earnings declines, which manifest as a positive earnings differential. In the final few years of the sample window, treated group workers’ earnings exceed those of control group workers by about 15 log points. This difference likely arises from the amount of continued full-time work for the treated group relative to the number of transitions into part-time or part-year work for the comparison group. It is consistent with wealth effects of the freeze inducing longer careers at full-time status.

For 56- to 64-year-old workers, who are shown in the central panel, the pattern is somewhat different. Treated workers experience a statistically significant earnings dip of 1.5 to 4.0 log points in the first five post-freeze years. This dip could arise for two reasons. First, treated workers who shorten their careers—that is, exhibit dominant substitution effects in the short term—may work less than full-time or less than full-year right after the freeze, thereby exhibiting reduced earnings relative to unaffected workers. Second, it is possible that short-term earnings declines are brought on by firm-specific factors.⁴⁸ The fact that 50- to 55-year-old workers do not experience earnings losses suggests that wage reductions are unlikely to be the only cause for the observed earnings dip that 56- to 64-year-old workers experience. In contrast, the long-term pattern shows a 20 to 30 log point increase in earnings relative to the comparison group and provides evidence of dominant wealth effects wherein freeze-affected workers delay retirement and continue in full-time or full-year employment while comparison group workers transition to part-time or part-year employment. As with the 50- to 55-year-old age group, it is important to reiterate that the positive long-term

⁴⁷See Appendix B for details on how retirement status is measured in the LEHD, along with estimates showing that it closely reflects the retirement pattern of a comparable sample of respondents in the HRS.

⁴⁸Analyses of firm-level payroll data shown in Appendix C indicate that freezes reduce firm-wide average earnings by approximately 2.5 log points. These earnings changes could stem from the changing post-freeze age composition of the firm’s workforce or from reductions in offered wages.

earnings differential reflects a slower decline in earnings levels rather than an increase.

Earnings differences for 65- to 70-year-old workers, shown in the far-right panel, are not precisely estimated over the sample window. However, the post-freeze coefficients indicate a fairly sustained drop in earnings of about 15 log points relative to the comparison group. Thus, while freezes induce workers over 65 to delay retirement, the pattern of earnings changes suggests that continued employment for these workers likely comes in the form of more part-time employment. That the oldest workers in the labor market are willing to extend their careers at less than full-time rates complements recent survey-based evidence on the importance of flexible hours in supporting longer working lives (Ameriks et al. (2020)).

6.2 Employer attachment

While the treatment effects shown thus far relate to the decisions about whether to work and how much to work, they do not address the decision about where to work. Employer-to-employer 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 the matched employer-employee structure of the LEHD to study differences in worker mobility between the treatment and comparison groups.

Figure 7 shows the percentage point change in the probability of leaving one’s DB-sponsoring employer.⁴⁹ 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 following a merger or acquisition. The large spike in transitions for treated workers that occurs four years before 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- to 55-year-old and 56- to 64-year-old age groups appear to respond with small but statistically significant increases in employer transitions. For 50- to 55-year-old workers, the transition rate increases by 1.4 percentage points from a baseline rate of 2.9 percent. For 56- to 64-year-old workers, the transition rate increases by 0.8 percentage point from 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 compared with workers closer to retirement age. After the first year, there is a statistically significant pattern of reduced employer-to-employer (E-E) transitions among workers in the treated group. Transition rates fall by 1.2 percentage points for 50- to 55-year-old workers and by about 0.6 to 0.9 percentage point for 56- to 64-year-old workers. When considered alongside the earnings results, reduced E-E mobility is consistent with continued employment in career jobs for treated-group workers as opposed to the counterfactual transition to part-time or bridge jobs for comparison-group workers. E-E mobility for workers over 65 is largely unaffected by freezes.⁵⁰

Reduced E-E transition rates for workers in the younger two age groups lines up with the time

⁴⁹That is, the employer that a worker is attached to in period $l - 5$.

⁵⁰Coefficient estimates are suppressed due to small sample sizes for the latter part of the estimation window.

window in which dominant wealth effects lead to increased employment and a 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 very 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. Therefore, robust demand for the labor services of older workers from their long-term employers is a key requirement to accommodate policies aimed at supporting longer working lives.

The results I have presented thus far characterize treatment effects of pension freezes on labor supply outcomes. These estimates provide insights about heterogeneity in workers’ preference for long careers and illustrate the importance of the right type of labor demand in sustaining those longer careers. In the next section, I use information about preferences revealed by these treatment effects to estimate parameters of the structural model.

7 Solving and estimating the structural model

In this section, I explain how I solve and estimate the structural model numerically using the method of simulated moments (McFadden (1989)). I identify the model’s parameters by matching model-based simulations of employment responses to pension freezes to the treatment effects observed in real-world administrative data.

7.1 Estimation

I estimate the model in two steps. In the first step, I calibrate several parameters (for example, earnings, pension accruals, DC match rates, Social Security benefits, tax rules, mortality rates, discount factors, and interest rates) using HRS survey data and other external data sources. Appendix E.1 provides a detailed description of the calibrations. In the second step, I estimate the model’s preference parameters $\theta = (\sigma, \gamma, \phi, \rho, \sigma_v)$ using MSM. The procedure is as follows:

1. For a given value of θ , I numerically solve the model using value function iteration under two different scenarios. In the first scenario, DB accruals progress normally. In the second scenario, I freeze DB accruals for each age from 56 through 64 and compute an alternative set of decision rules.
2. I simulate initial wealth and work disutility draws for 5,000 individuals who are initially aged 51 to 59 using information from the HRS. I apply decision rules from the no-freeze scenario to obtain work status and wealth accumulation paths from the initial age up through age 80 (the terminal age) to create a simulated control group.
3. Next, I apply the freeze decision rules for the same population of individuals starting five years after the initial age. Individuals in this exercise have the same work and wealth accumulation choices as the simulated control group for the first five years, but they have potentially different

work statuses and wealth accumulation choices once faced with a DB freeze. I call this sample the simulated treated group.

4. I compute two sets of moments using the simulated control group and treated group. The first set of moments is the difference in average employment rates between the two groups or the simulated treatment effect. I compute these differences for 12 periods, starting from the period of the freeze (12 moments). The second set of moments is the employment rate trend for the simulated control group. I compute the employment rate moments starting 4 periods before the freeze and lasting 12 periods after the freeze (17 moments). I then stack these 29 simulated moments together and compare them with the same moments estimated from LEHD data.
5. I use the Nelder-Meade algorithm to search for the value of θ that minimizes the distance between the simulated moments and the real-world moments.

Further details on the solution algorithm and the estimation procedure are provided in Appendix E.2.

7.2 Identification

Identification of the model’s parameters derives from two key sources of variation in real-world data. The first set of moments, which capture the age-based decline in employment rates in the comparison group, identify the constant (γ) and slope (ϕ) terms of the deterministic component of g .

The second set of moments—the treatment effect dynamics—are informative about within-age variation in individual preferences for continued work. Equation (12) helps to explain the identification argument for (σ, ρ, σ_v) . Notice that \bar{g} is composed of two separate terms: The first term is the change in utility from consumption at retirement, and the second term is the option value of working. Freezes affect both terms. Changes in \bar{g} driven by post-freeze re-optimization of consumption choices provide information about individuals’ preferences for smooth consumption profiles over the life cycle. This source of variation in \bar{g} aids in the identification of the intertemporal elasticity of substitution (IES), σ .⁵¹

The treatment-effect estimates reveal that freeze-induced changes in the option value of work induce early retirement for some workers and delayed retirement for others. These differences are informative about the persistence (ρ) and variance (σ_v) of the idiosyncratic component of g . In particular, ρ and σ_v need to be large enough so that workers of the same age and with the same level of wealth exhibit different retirement behavior when affected by a compensation shock of the same magnitude. For workers of a given age and wealth, those with high values of g have stronger preferences for leisure and choose to retire early. On the other hand, those with low values of g have weaker preferences for leisure and choose to delay retirement.

⁵¹Note that consumption data are not directly used to identify σ ; rather, variation in \bar{g} induced by freezes indirectly aids in identification.

7.3 Model fit and parameter estimates

Figure 8 shows how simulated moments compare with observed moments. The left panel shows the simulated and real-world employment rate trends for 56- to 64-year-old workers who are not subject to freezes, while the right panel compares simulated treatment effects with real-world treatment effects. These figures show that the model fits the trend rate of employment and, importantly, both the magnitude and the timing of employment rate fluctuations induced by pension freezes.

Table 2 shows the estimated parameters. The IES estimate is close to 1, which is equivalent to log utility. Because the 5 preference parameters in θ are identified using 29 moments, I am able to conduct a χ^2 overidentification test. The model is formally rejected on the basis of this test, which I report in the last row of the table. The main reason the test statistic is large is that the estimation procedure does not account for variance associated with first-step parameters obtained from the HRS. Consequently, the variance of second-step simulated moments is understated and the weight attached to each moment (which is inversely related to variance) is large.⁵²

8 Evaluating the effectiveness of a counterfactual payroll tax reform

In this section, I study the effect of eliminating the Old Age and Survivors Insurance (OASI) component of the payroll tax for older workers. I focus on transition cohorts for whom the policy change is unexpected. The main idea behind this reform proposal is to remove the disincentives for longer careers that exist under current law. While payroll taxes are applicable to all years of work, the Social Security benefit formula is based on the 35 years of highest paid work. For the typical worker over the age of 60 who no longer experiences real earnings growth, an additional year of work generates no substantial increase in Social Security benefits. Nevertheless, because all workers contribute OASI payroll taxes regardless of age, workers who remain employed beyond the 35-year vesting age experience tax burdens with no offsetting benefit increases. By relieving these fully “paid-up” workers from additional payroll tax contributions, the reform has the potential of lengthening careers. Understanding the costs and benefits of the reform requires the quantifying of behavioral responses, which I do here using the estimated model.

The flavor of reform that I consider involves an unexpected elimination of the OASI payroll tax for workers who are age 60 or older. To simplify the analysis, I assume that workers bear the full burden of the tax, which, under current law, is set at 5.3 percent each for workers and employers.⁵³ I assume that the disability insurance (DI) and hospital insurance/Medicare (HI) components of the payroll tax are unaffected, as is the Social Security benefit formula.⁵⁴ To more accurately portray incentives affecting the majority of employees in the current and future workforce, I remove DB pensions from the model.

⁵²Standard errors are understated for the same reason.

⁵³Full pass-through implies that the reform increases earnings subject to federal income tax by 10.6 percent. In the model, workers can choose to defer some of this extra income to DC accounts, thereby reducing their federal income tax liability.

⁵⁴DI and HI payroll tax components collectively amount to 4.7 percent.

Figure 9 shows model simulations of employment rate trends for workers aged 56 to 75 under current law and under the reform. The two trends diverge starting at age 60, with employment rate differences peaking at 12.4 percentage points at age 62. Averaging over the entire post-60 age window, the reform increases employment rates by 3.9 percentage points.

Table 3 provides additional statistics comparing the two different regimes. The left panel shows moments from the distribution of a variety of outcomes under current law, while the right panel shows the same moments under the reform. The third panel shows the mean difference in outcomes. Reflecting the overall increase in employment rates, the average retirement age under the reform rises by nearly 10 months. Extending a career by that much time allows workers to accumulate about \$5,700 more in their DC retirement accounts and pay \$12,400 more in federal income taxes. These gains accrue against a loss in payroll tax revenue per worker of \$25,300. Finally, as shown in the lower panel of the table, I find that equivalent variation from the reform averages about \$54,000 per worker, implying a substantial welfare gain.⁵⁵

The counterfactual analysis presented here shows that the relatively elastic labor supply of older workers can be harnessed to extend their careers with an OASI payroll tax sunset at age 60. I find that older workers would obtain substantial welfare gains from the reform and have more income in retirement, although net revenues collected from workers after age 60 decline by about 10 percent.

9 Conclusion

In this paper, I exploit widespread and under-examined shifts in employer-sponsored pension benefit programs to better understand how retirement behavior responds to changes in compensation. Over the last 20 years many employers have reneged on long-standing promises to continue supporting traditional retirement benefits in the face of rising costs of provision. Although these changes are substantial, there is no evidence on how they have affected workers' labor market outcomes. Creating a new data set that brings together detailed administrative information on pension plan characteristics with matched employer-employee data, I study the impact of these unexpected shocks on employment, retirement, earnings, and employer attachment for workers aged 50 to 70.

I find evidence of substantial heterogeneity among workers, even among those of the same age. When faced with freeze-induced compensation changes, some workers choose to retire early while others choose to delay retirement, thereby illustrating that differences in wealth and differences in preferences for leisure are important in explaining retirement decisions. Using these quasi-experimental treatment effects as targets, I estimate a structural model of retirement and saving that allows for heterogeneity in wealth and leisure preferences. I use the model to simulate the effect of a counterfactual policy that eliminates OASI payroll taxes for workers who are fully vested in their Social Security benefits. Simulations from the estimated model show that eliminating

⁵⁵Laitner and Silverman (2012) build revenue neutrality into their analysis of a similar reform by requiring a small increment to payroll tax contributions over the entire pre-vesting period. With revenue neutrality, Laitner and Silverman (2012) estimate smaller equivalent variations. I do not model the pre-vesting age lifespan because the simulation I develop relies heavily on the HRS, whose sample design covers only individuals over age 50.

the tax at age 60 induces a 10 month increase in the average retirement age, increases retirement saving by about \$5,700, and produces substantial welfare gains. However, these gains come at the cost of a 10 percent reduction in net revenues from workers over age 60.

The empirical setting and the model adopted in this paper highlight the importance of forward-looking behavior in determining labor supply decisions, particularly retirement decisions. It is likely that unobserved determinants of the relationship between employers and employees, such as implicit contracts with a rich set of contingencies, drive the association between current labor supply and expected future compensation (see, for example, Lazear (1981) and Akerlof and Katz (1989)). These contingencies create tight bonds between employers and employees that many standard models of labor supply do not incorporate. Notably, the durability of these relationships and the impact they have on forward-looking labor supply behavior are not unique to DB-style incentives, which are no longer common in the U.S. private sector. For instance, the median tenure for workers aged 55 to 64 has remained unchanged for two decades despite major economic transitions including the demise of DB pensions.⁵⁶ This fact suggests that a better understanding of the mechanisms driving long-term employer-employee relationships is critical in explaining lifetime labor supply behavior.

References

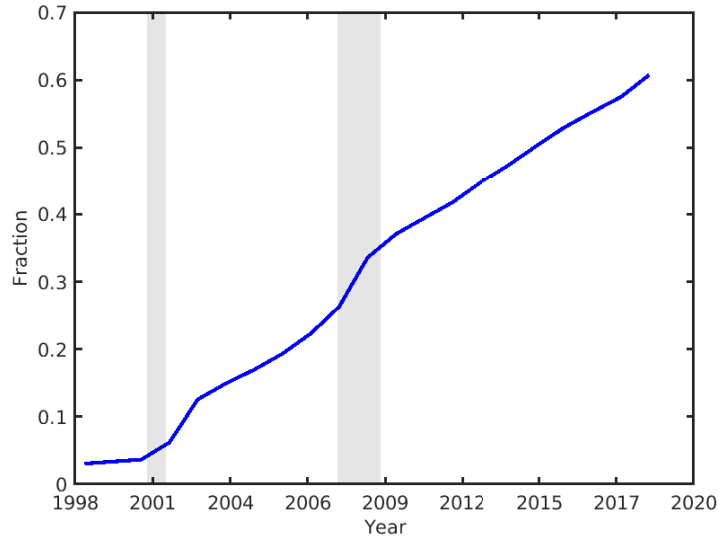
- Abowd, John, Bryce E. Stephens, Lars Vilhuber, Fredrik Anderson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock**, “The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators,” in Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press, 2009, pp. 149–230.
- Akerlof, George and Lawrence F. Katz**, “Workers’ Trust Funds and the Logic of Wage Profiles,” *Quarterly Journal of Economics*, 1989, 104 (3), 525–536.
- Ameriks, John, Joseph S. Briggs, Andrew Caplin, Minjoon Lee, Matthew D. Shapiro, and Christopher Tonetti**, “Older Americans Would Work Longer If Jobs Were Flexible,” *American Economic Journal: Macroeconomics*, 2020, 12, 174–209.
- Berkovic, James and Steven Stern**, “Job Exit Behavior of Older Men,” *Econometrica*, 1991, 59 (1), 189–210.
- 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.
- Burtless, Gary and Joseph P. Quinn**, “Is Working Longer the Answer for an Aging Workforce?,” Technical Report, Center for Retirement Research at Boston College December 2002.

⁵⁶The median tenure, as reported by the Bureau of Labor Statistics, for workers in the 55 to 64 age range in 1998 and in 2018 was 10.1 years.

- Coile, Courtney**, “The Evolution of Retirement Incentives in the U.S.,” in Axel Börsch-Supan and Courtney Coile, eds., *Social Security Programs and Retirement Around the World: Reforms and Retirement Incentives*, University of Chicago Press, 2021.
- **and Jonathan Gruber**, “Future Social Security Entitlements and the Retirement Decision,” *Review of Economics and Statistics*, 2007, 89 (2), 234–246.
- 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.
- French, Eric**, “The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behaviour,” *Review of Economic Studies*, 2005, 72 (2), 395–427.
- **and John Bailey Jones**, “The Effects of Health Insurance and Self-Insurance on Retirement Behavior,” *Econometrica*, 2011, 79 (3), 673–732.
- Friedberg, Leora and Anthony Webb**, “Retirement and the Evolution of Pension Structure,” *Journal of Human Resources*, 2005, 40 (2), 281–308.
- Gelber, Alexander M., Adam Isen, and Jae Song**, “The Effect of Pension Income on Elderly Earnings: Evidence from Social Security Full Population Data,” 2016.
- Goda, Gopi Shah, John B. Shoven, and Sita Nataraj Slavov**, “Removing the Disincentives in Social Security for Long Careers,” in Jeffrey Brown, Jeffrey Liebman, and David Wise, eds., *Social Security Policy in a Changing Environment*, University of Chicago Press, 2009, pp. 21–38.
- Gustman, Alan L. and Thomas L. Steinmeier**, “Effects of Social Security Policies on Benefit Claiming, Retirement, and Saving,” *Journal of Public Economics*, 2015, 129 (C), 412–437.
- 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.
- Hurd, Michael and Susan Rohwedder**, “Trends in Labor Force Participation: How Much is Due to Changes in Pensions?,” *Journal of Population Ageing*, 2011, 4, 81–96.
- 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.
- Krueger, Alan B. and Jorn-Steffen Pischke**, “The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation,” *Journal of Labor Economics*, 1992, 10 (4), 412–437.

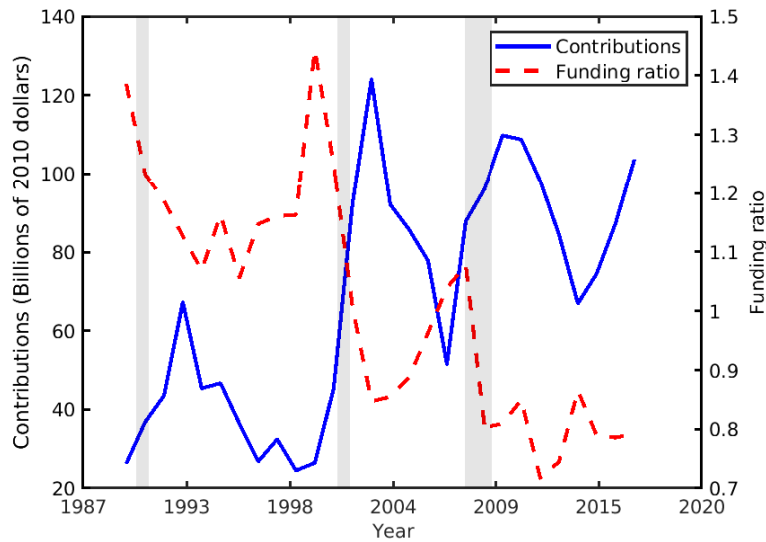
- Laitner, John and Daniel Silverman**, “Consumption, Retirement, and Social Security: Evaluating the Efficiency of Reform that Encourages Longer Careers,” *Journal of Public Economics*, 2012, 96 (7-8), 615–634.
- Lazear, Edward P.**, “Why Is There Mandatory Retirement?,” *Journal of Political Economy*, 1979, 87 (6), 1261–1284.
- , “Agency, Earnings Profiles, Productivity, and Hours Restrictions,” *American Economic Review*, 1981, 71 (4), 606–620.
- **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,” *American Economic Journal: Economic Policy*, 2016, 8 (4), 160–182.
- McFadden, Daniel**, “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration,” *Econometrica*, 1989, 57 (5), 995–1026.
- 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., Irina Stefanescu, and Stephen P. Zeldes**, “Cost Saving and the Freezing of Corporate Pension Plans,” 2020. NBER Working Paper No. 27251.
- Rogerson, Richard and Johanna Wallenius**, “Nonconvexities, Retirement, and the Elasticity of Labor Supply,” *American Economic Review*, 2013, 103 (4), 1445–1462.
- Rust, John and Christopher Phelan**, “How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets,” *Econometrica*, 1997, 65 (4), 781–831.
- Stock, James H. and David A. Wise**, “Pensions, the Option Value of Work, and Retirement,” *Econometrica*, 1990, 58 (5), 1151–1180.
- Zelinsky, Edward A.**, “The Cash Balance Controversy,” *Virginia Tax Review*, 2000, 19, 683–762.

Figure 1: DB plans hard-frozen or converted to cash balance



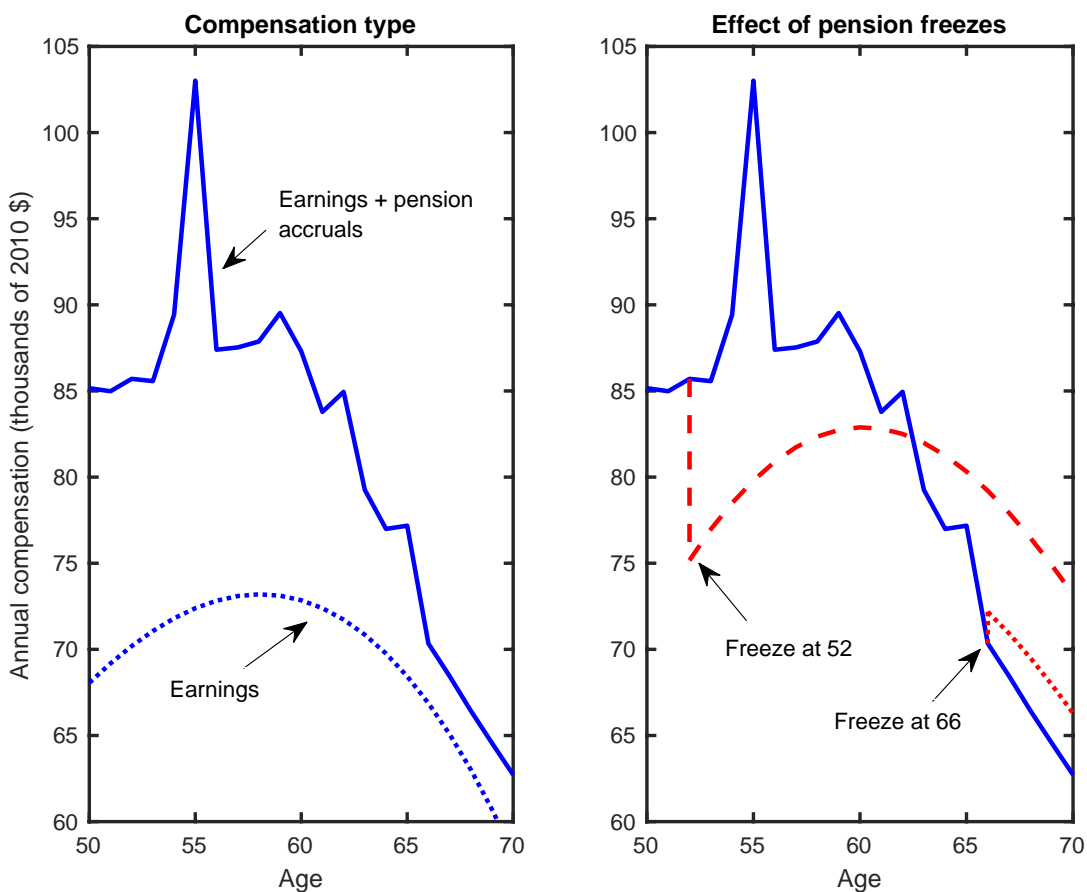
Notes: This graph shows the share of private sector single-employer DB pension plans that have been hard-frozen or converted to cash balance plans. Time series are based on 1999–2019 F5500 microdata. Gray shaded areas show recessions.

Figure 2: Pension costs and funding ratios




Notes: This figure shows total contributions and aggregate funding ratios for private sector single-employer DB sponsors. Total contributions (solid blue line) are measured on the left axis. Funding ratios (dashed red line) are measured on the right axis. The funding ratio is the ratio of the market value of assets to the present value of future pension liabilities. Contributions data are drawn from Table E13 of the Department of Labor’s Private Pension Plan Bulletin Historical Tables and Graphs 1975–2018. Funding ratios are drawn from Table S-44 of the PBGC Pension Insurance Data Book, 2018. Gray shaded areas show recessions.

Figure 3: Simulated effect of pension freezes on total compensation




Notes: This figure is based on data from DB-eligible HRS respondents in the 2010 wave who are employed in the private sector and whose DB plans have neither been frozen nor converted to cash balance plans. See Appendix F for details on the sample and simulation of post-freeze compensation.

Figure 4: Impact of freezes on employment




Notes: This figure shows the treatment effect of freezes on employment rates for workers in different age groups at the time of the freeze. The p-values for F-tests that the coefficients for periods -5 to -1 are jointly 0 are 0.41, 0.38, and 0.35 for each of the age groups, respectively. Dotted lines show 95 percent confidence intervals that are based on standard errors clustered at the firm level.

Figure 5: Impact of freezes on retirement




Notes: This figure shows the treatment effect freezes on retirement rates for workers in different groups at the time of the freeze. The p-values for F-tests that the coefficients for periods -5 to -1 are jointly 0 are 0.64, 0.77, and 0.35 for each of the age groups, respectively. Dotted lines show 95 percent confidence intervals that are based on standard errors clustered at the firm level.

Figure 6: Impact of freezes on log annual earnings



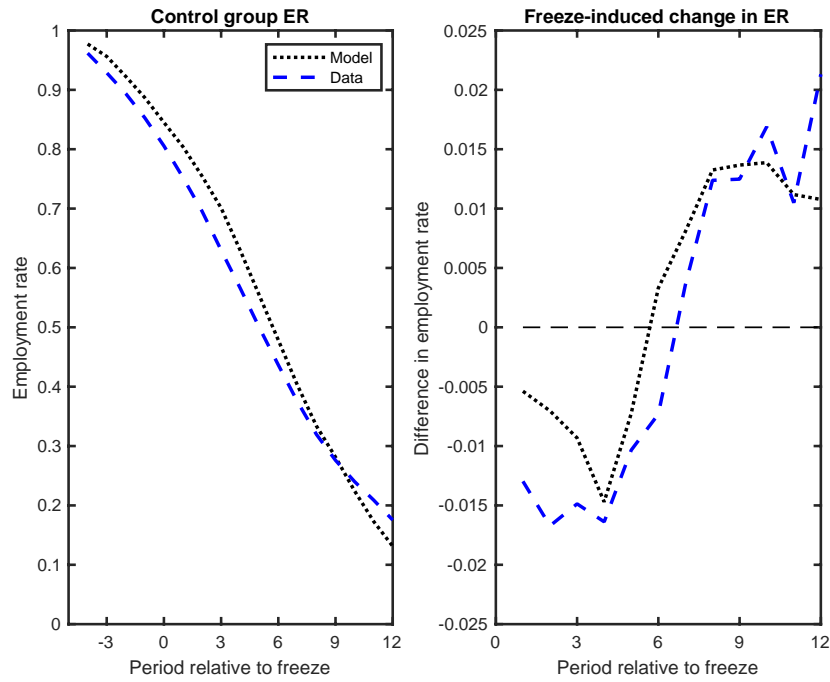
Notes: This figure shows the treatment effect of freezes on log annual earnings for workers in different age groups at the time of the freeze. The p-values for F-tests that the coefficients for periods -5 to -1 are jointly 0 are 0.17, 0.30, and 0.31 for each of the age groups, respectively. Dotted lines show 95 percent confidence intervals that are based on standard errors clustered at the firm level.

Figure 7: Impact of freezes on employer attachment




Notes: This figure shows the treatment effect of freezes on employer attachment for workers in different age groups at the time of the freeze. The p-values for F-tests that the coefficients for periods -5 to -1 are jointly 0 are 0.01, 0.06, and 0.58 for each of the age groups, respectively. Dotted lines show 95 percent confidence intervals that are based on standard errors clustered at the firm level.

Figure 8: Data moments compared with model simulated moments



Notes: The left panel shows the employment rate (ER) trend for 56- to 64-year-old workers, all of whom are employed five periods preceding the freeze. The right panel shows the difference between the treated group and control group employment rates. Dotted lines show moments from simulated data. Dashed lines show moments from real-world data.

Figure 9: Effect of OASI tax sunset at age 60 on employment rates



Notes: This figure shows the evolution of the employment rate under two scenarios simulated by the model. The dashed blue line shows the employment rate trend under current law. The dotted red line shows the counterfactual employment rate trend under a regime where the OASI payroll tax is unexpectedly eliminated at age 60. The trends are based on model simulations for workers who are employed at age 55.

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 (\$)	65,140	-1,477	0.75	65,280	-2,337	0.58	47,770	-3,562	0.13
Tenure at $l - 5$	7.8	-0.943	0.15	8.2	-0.977	0.11	8.0	-0.623	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 $l - 5$ employer	0.047	0.012	0.22	0.042	0.012	0.24	0.030	0.003	0.45
Pension and firm characteristics									
Log DB pension wealth/active participant	10.07	0.072	0.86	10.15	0.039	0.92	10.43	0.043	0.82
Log DB pension accrual/active participant	7.43	0.094	0.80	7.49	0.054	0.88	7.60	0.052	0.77
Pension plan claim age	62.8	-0.1	0.70	62.9	-0.1	0.72	64.4	0.0	0.88
Log firm size	8.55	-0.017	0.98	8.40	0.007	0.99	7.31	-0.162	0.74
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: Unless otherwise noted, statistics reported in the table average over the five-year period preceding any freeze activity. Pension wealth per active participant is computed as the present value of the liability owed to active participants divided by the number of active participants. All dollar values are expressed in 2010 terms. Tenure is understated because the LEHD does not capture the complete history of an employer-employee relationship when states enter the data set 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: Structural model parameter estimates

Parameter		Estimate	Standard error
IES	σ	1.02	0.005
Labor disutility persistence	ρ	0.858	0.01
Labor disutility standard deviation	σ_v	0.328	0.01
Labor disutility age slope	ϕ	0.067	0.001
Labor disutility constant	γ	-3.98	0.016
χ^2 statistic, 17 d.f.		712.7	

Notes: This table shows preference parameter estimates and standard errors for the structural model. See Appendix E for details.

Table 3: Counterfactual effects of OASI tax sunset at age 60

Variable	Current law				Reform				Difference
	p25	p50	p75	Mean	p25	p50	p75	Mean	Mean
Retirement age	62	64	66	64.1	63	65	67	64.9	0.8
DC wealth at retirement (\$)	140,045	207,370	342,035	292,051	142,485	212,405	350,010	297,715	5,663
PV federal income tax remitted after 60 (\$)	60,747	82,694	106,029	89,584	73,591	95,327	119,118	101,983	12,399
PV payroll tax remitted after 60 (\$)	21,886	42,110	60,599	41,072	10,419	16,683	22,378	15,792	-25,281
Equivalent variation (\$, PV at age 60)					p25	p50	p75	Mean	
					22,214	41,646	68,372	54,023	

Notes: Estimates are derived from model simulations under current law and under the reform. Simulated workers in the sample are all employed at age 55. Estimates are conditioned on workers not having retired before age 60. Monetary estimates are reported in 2010 dollars, and present values are computed as of age 60. See text for details.

Appendix for online publication

A Data

Form 5500

Form 5500 (F5500) is an annual plan-specific filing collected jointly by the Internal Revenue Service (IRS), Department of Labor (DoL), and the Pension Benefit Guaranty Corporation (PBGC) to ensure compliance with the Employee Retirement Income Security Act of 1974 (ERISA).⁵⁷ Each plan in the F5500 database is identified by a combination of a Federal Employer Identification Number (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 through 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 through 2009, the research data include all pension plans with more than 100 participants and a 5 percent sample of plans with fewer than 100 participants. I use records from the research data where possible and add back small plans (for example, those with fewer 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 Freedom of Information Act (FOIA) request to the DoL. The sample that I use covers plan years ending during the 1996–2014 period.

I focus primarily on defined benefit (DB) plans but also obtain data on defined contribution (DC) plans offered by employers that 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 or 3, and welfare benefit plans—such as employer-provided health insurance—have prefix 4.⁵⁸ 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 the string “supplemental” in the plan name field. I restrict my sample to single-employer plans, thereby eliminating multi-employer DB plans.⁵⁹

In addition to the main form, I include data taken from the actuarial information attachment. Prior to 2008 this attachment was labeled Schedule B. Since 2009 it has been labeled Schedule SB for single-employer plans. For 2008, I impute actuarial data by interpolating between the 2007 and 2009 values, because actuarial information is unavailable in electronic format. The actuarial attachment contains important plan-level data including detailed breakouts of plan assets and liabilities,

⁵⁷F5500 data for form years after 1999 are available at <https://www.dol.gov/agencies/ebsa/about-ebsa/our-activities/public-disclosure/foia/form-5500-datasets>.

⁵⁸The codes are different in data from 1999 and earlier. Welfare benefit plans need to be reported with an F5500 filing only when such plans cover more than 100 active participants.

⁵⁹Multi-employer plans are arrangements between a group of firms and/or unions to provide pension benefits to eligible employees within the group.

accruals earned during the plan year, the average benefit claim age, mortality and separation rate assumptions, etc.

Linking to the Census Business Register and Census Longitudinal Business Database

The Census Business Register (BR) is a database of the universe of establishments in the United States. It includes information on organization structure, location, industry, revenue, payroll, and employment. The relationship between establishments belonging to multi-unit firms is determined using responses to the Census Bureau’s Company Organization Survey (now the Report of Organization), Economic Census, and 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 of 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 an EIN plan 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.⁶⁰ 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. Consequently, when firms sponsor multiple plans, the data that I have do not allow me to determine 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.⁶¹ Using Census firm identifiers, I match these data with the Longitudinal Business Database (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 through 2014 as an experiment year, which is indexed by l .⁶² This terminology reflects the research design wherein each year yields a fresh sample of firm-level pension freezes. For a given experiment year, the panel data set of workers employed at freezing firms constitutes

⁶⁰Firms 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.

⁶¹When 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.

⁶²I start with 2001 because it is the first year in which pension freezes are reported in F5500. Cash balance (CB) conversions are reported in earlier years, but the number of firms making CB conversions before 2001 is small.

the treated group, while the panel data set of workers employed at firms that do not freeze their plans constitutes the comparison group. I impose the restriction that firms file F5500 for their DB plans for five calendar years preceding the experiment year, which I refer to as the pre-period. By requiring plan stability in the lead-up to the experiment 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.⁶³

Linking to the Longitudinal Employer Household Dynamics

The Longitudinal Employer Household Dynamics (LEHD) is a quarterly matched employer-employee data set 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 data set at different points in time. For example, Maryland enters in 1985:Q2 whereas Mississippi does not enter until 2003:Q3. Because of staggered entry, the scope of the data grows continuously over time.⁶⁴ I use the 2014 snapshot version of the LEHD, which provides 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 state employer identification number (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-level and employer-level characteristics from the F5500-LBD linked sample to state-level employers in the LEHD. I recover 89 percent of firm-experiment years, which corresponds to 93 percent of employee-experiment years from the

⁶³I 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.

⁶⁴Several populous states enter the data relatively early. Illinois enters in 1990, California and Pennsylvania in 1991, Florida in 1992, and New York and Texas in 1995.

F5500-LBD linked sample (see row 3 of Table A1). From this sample, I drop a small percentage of observations from which certain pension plan variables are missing (see row 4 of Table A1).⁶⁵

Worker sample

Regarding 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.⁶⁶ 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 $l - 5$ who have at least two years of tenure as of $l - 5$ and who will be aged 50 to 70 in the experiment year.⁶⁷

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-experiment years)	1419000	0.885	0.933
No missing pension data (firm-experiment years)	1256000	0.930	0.972

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

⁶⁵Critical pension plan information includes plan assets and liabilities, accruals earned in the filing year, and the average benefit claim age for the plan.

⁶⁶The 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 $[l - 5, l - 2]$ —that is, five to two years preceding the experiment year. 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 and are not replaced with new, eligible workers.

⁶⁷The two-year tenure restriction ensures that workers are fully vested in their pensions as of the experiment year when they may become subject to a freeze. This calculation is based on the seven-year maximum full vesting period allowed for DB plans by the Employee Retirement Income Security Act (ERISA).


B Imputing retirement in the LEHD

In the LEHD, a worker is employed if they have positive earnings in a given year. The definition of retirement is somewhat more involved. I classify an individual as having retired in year t if the last year in which they received non-zero earnings was $t - 1$. By this definition, retirement occurs only when an individual completely withdraws from paid employment. Recall, however, that both definitions exclude work in the form of self-employment because the LEHD data are based on UI-covered earnings. To examine the potential for misclassification of retirement in the 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 experiment year whose pensions have not been frozen.⁶⁸

These comparisons between HRS and LEHD data are shown in the three panels of Figure B1, which split the samples into three age categories as of 2004. The retirement rates align fairly well for all three age categories, with the best alignment for the 56- to 64-year-old age group.

⁶⁸Retirement 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 data set. 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.

Figure B1: 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 ≥ 80 percent. Sample splits are based on age as of 2004.

C 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, are 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.

C.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 C1. Comparing the left and right panels of the table shows that 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 approximately 70 and 35 percent, respectively. The lack of any meaningful differences 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 output is therefore also not a leading reason for a decision to freeze.

The 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 Pension Protection Act of 2006 (PPA) designates plans with funding ratios under 80 percent as being “at-risk” or distressed. According to the PPA’s threshold, 20 percent of freezing firms have distressed plans, whereas the distressed 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 toward underfunded plans therefore become larger and more volatile in the face of market risk.

Table C2 shows coefficients from a linear prediction model using a freeze in the experiment 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 reduces 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 reduces the likelihood of a future freeze by 0.2 to 0.5

percentage point. 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 1 percent improvement in plan funding. Industry fixed effects, which are included in columns 2 and 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.

C.2 Post-freeze changes

Figure C1 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 (14) 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 experiment year.

The upper row of Figure C1 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 that 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 the firm-level data alone does not clearly indicate that role for each channel.⁶⁹ The lower-left panel of Figure C1 shows that the fraction of DC-covered workers starts to rise about two years before the DB freeze, reflecting expanded DC plan eligibility, more generous DC match rates, or transitions to opt-out rather than opt-in DC enrollment.⁷⁰ In the period right around a DB freeze, DC coverage rates increase by about 5 percentage points from 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 toward 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 C1 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

⁶⁹The 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.

⁷⁰DC coverage is measured as the ratio of the maximum number of participants across a firm’s DC plans to total employment.

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. Therefore, it appears that distress is concentrated among smaller firms.

Table C1: Pre-period firm and plan characteristics

Variable	Non-freezing firms		Freezing firms	
	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-experiment years	428000		28500	
Firms	22500		6500	

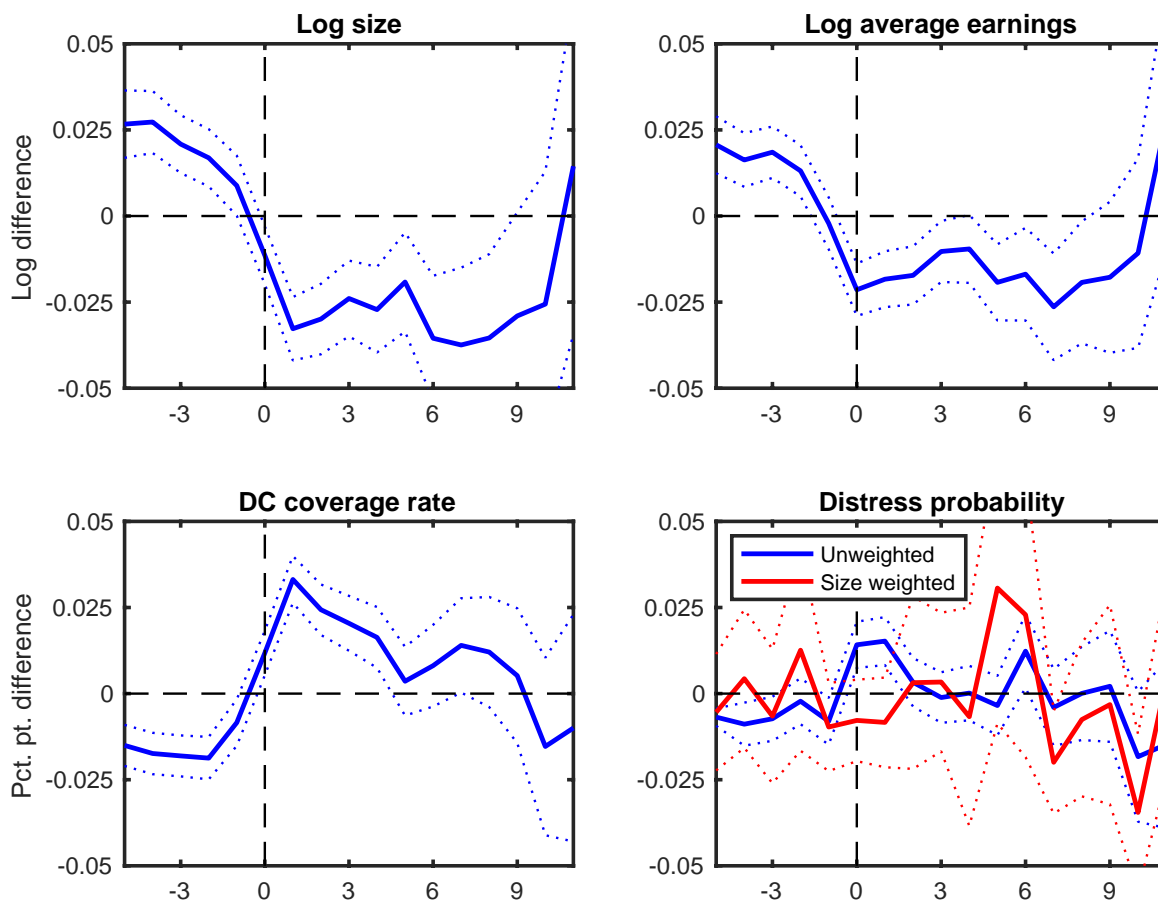
Notes: Statistics reported in the table average over the five-year period preceding any freeze activity. All dollar values are expressed in 2010 terms. Pension wealth per participant is computed as the present value of the liability owed to active participants divided by the number of active participants. 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 C2: Predictors of future freezes

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

Figure C1: 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 experiment year. Firm-years are coded as distressed if firm-wide year-on-year employment shrank by 30 percent or more.

D Propensity score re-weighting

For workers in cell $j(i)$, denote the probability of experiencing a freeze, or the propensity score, by $\hat{p}(\mathbf{z}_{j(i)})$. $\mathbf{z}_{j(i)}$ is a vector including all pre-period observations on log firm size, firm-level averages of log total pension wealth and log pension accruals per working participant, average benefit claim age from the pension plan, the age structure of the firm’s workforce, cell-level annual earnings, retirement rates, labor force participation rates, and employer-to-employer transition rates. I also condition on state, gender, tenure, and earnings averaged over the workers LEHD history prior to $l - 5$. $\hat{p}(\mathbf{z}_{j(i)})$ is estimated using logistic regression.

In the setting considered in this paper, the parameter of interest is the average treatment effect on the treated (ATET)—that is, 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 $\frac{\hat{p}(\mathbf{z}_{j(i)})}{1 - \hat{p}(\mathbf{z}_{j(i)})}$. 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 $\hat{p}(\mathbf{z}_{j(i)})$ —that is, 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 $\hat{p}(\mathbf{z}_{j(i)}) < 1$ 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 $\hat{p}(\mathbf{z}_{j(i)})$ is less than 1.

E Solution and estimation of the structural model

E.1 Calibrated parameters

DB pension wealth accruals

The Health and Retirement Study (HRS) provides high-quality estimates of DB pension wealth for DB-eligible survey respondents. These estimates are computed using summary plan descriptions (SPDs) obtained either directly from respondents’ employers or from attachments to F5500 administrative records. Information from these documents is coded along with relevant data on respondents’ past and future earnings projections and job tenure to calculate pension wealth at prospective retirement ages. Pension wealth projections are made using a software system known as the Pension Estimation Program (PEP).⁷¹

SPD-based pension wealth estimates are available for the 1992, 1998, 2004, and 2010 survey wave years of the HRS. While the 2004 and 2010 waves are most relevant because they coincide with the time period that I analyze, I rely exclusively on the 2010 wave for two reasons. First, unlike earlier waves, the 2010 wave 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 restricted-access version of the 2010 sample includes flags that allow me to remove frozen DB plans and CB plans from the sample and isolate the sub-sample of plans where the evolution of pension accruals occurs under the status quo—that is, where participants do not experience either a freeze or cash balance conversion.⁷²

DB pension wealth in the PEP is calculated using the following formula

$$W_q^{DB} = \sum_{t=q}^{119} p_t \left(\frac{1 + COLA}{1 + i} \right)^{t-T_0} B_{t|q}, \quad (15)$$

where W_q^{DB} is the present value of pension wealth at potential retirement date q , p_t is the probability of survival in period t conditional on being alive in period q , $COLA$ (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 $B_{t|q}$ is the annual pension benefit in period t conditional on retiring in period q .⁷³ DB pension wealth is based on the maximum amounts of wealth at the plan’s normal retirement age and early retirement age and the vested deferred value of benefits. When respondents have wealth in multiple plans, I sum pension wealth over each plan. I convert nominal values of

⁷¹See Fang et al. (2016) for more details.

⁷²I also exclude combination plans, which are plans that have both DB and DC features. While combination plans are generally uncommon, the HRS PEP Users Guide (April 2016) notes that “in 2010 many plans are coded as combination plans...many of these plans are frozen defined benefit plans that were followed by the establishment of a cash balance plan. Others are purely cash balance plans.”

⁷³ 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.

W_q^{DB} into 2010 dollars. I keep respondent quit dates where respondents have at least two years of service with their DB-sponsoring employer. This restriction ensures that simulated workers, when confronted with a pension freeze five years after the initial age of observation, have at least seven years of tenure, which is a restriction that I impose on the LEHD-based sample. It also guarantees that all workers are vested in their DB pension wealth using the seven-year maximum vesting period for DB plans allowed by ERISA.

I average across respondents in the sample to compute real DB wealth at each prospective quit age. I use the mortality and real interest rate assumptions from the PEP to convert these values into their annuity equivalents at each quit age, which is denoted by $\{b_a^{DB}\}_{a=51}^{80}$.

Earnings

In the PEP, earnings from the 2010 survey wave are projected forward and backward for different quit dates using the following formula

$$\log(e_q) = \log(e_{2010}) + \alpha_m(q - 2010) + \beta_1 \text{age}_q + \beta_2 \text{age}_q^2, \quad (16)$$

where α_m is nominal wage growth, and β_1 and β_2 adjust for age-based changes in earnings. α_m is set to 5 percent, which is the sum of 2.2 percent real wage growth and 2.8 percent inflation based on the intermediate term economic assumptions of the 2010 SSA Trustees report. I estimate β_1 and β_2 separately for men and women using the sample of HRS respondents with employer-sponsored retirement plans by regressing log earnings on a linear time trend, respondent age, and respondent age squared. I apply these values of α_m , β_1 , and β_2 to the set of respondents for whom I compute DB wealth estimates, and average the 2010 dollar value of projected earnings across respondents by quit age to compute an age-dependent earnings function $\{e_a\}_{a=51}^{80}$.

Social Security wealth

I compute household-level Social Security wealth estimates using the sample of respondents for whom I also compute DB pension wealth estimates, that is, respondents employed in the private sector who are covered by unfrozen traditional DB plans. For respondents who have yet to claim benefits, Social Security wealth in the HRS is computed by passing each respondent's administrative earnings record along with projected future earnings into the Social Security benefit formula. The benefit levels obtained from the formula are used to compute the present value of Social Security benefits at age 62, 65, and 70 (see Fang et al. (2016)). These values are reported in 2010 dollars. Because individuals cannot claim Social Security before age 62, I set Social Security wealth before that age as 0. I use the mortality and real interest rate assumptions from the PEP to convert household-level Social Security wealth estimates provided by the HRS into their annuity equivalents at each age, which is denoted by $\{b_a^{SS}\}_{a=51}^{80}$. I then linearly interpolate these values for ages 62, 65, and 70. Because Social Security benefits do not increase after age 70, I set $\{b_a^{SS}\}_{a=71}^{80} = b_{70}^{SS}$.

DC wealth and non-pension wealth

I compute household-level DC and non-pension wealth estimates using the sample of respondents for whom I also compute DB pension wealth estimates, that is, respondents employed in the private sector who are covered by unfrozen traditional DB plans.⁷⁴ These two forms of wealth holding are characterized by two salient facts. First, many surveyed households have no DC wealth at all. Second, conditional on having positive DC wealth, households with high levels of non-pension wealth tend to have high levels of DC wealth. I account for both empirical patterns when simulating wealth stocks for workers in the model. For each age from 51 through 59, I compute the share of HRS respondents who have no household-level DC wealth. For these respondents, I estimate an age-specific distribution of non-pension wealth (A), which I assume is log-normally distributed, and fit the parameters of the distribution using maximum likelihood. Next, for households with positive levels of DC wealth, I assume that non-pension wealth (A) and DC wealth (W^{DC}) are jointly log-normally distributed. For each age from 51 through 59, I fit the parameters of a joint log-normal wealth distribution using maximum likelihood.

For each initial age in the model (that is, 51 through 59), I assign zero DC wealth to the sample-implied fraction of individuals without any DC wealth. For these individuals, I simulate non-pension wealth by drawing values from the fitted univariate distribution for A . Next, for the sample-implied fraction of individuals with positive DC wealth I simulate the joint distribution of A and W^{DC} by drawing values from the fitted bivariate distribution for A and W^{DC} . This simulation accounts for the nontrivial share of individuals with zero DC wealth as well as the positive correlation between A and W^{DC} .

DC match function

The employer's DC match function, $m^e(m^w)$, is based on the sample of HRS respondents who either actively or passively participate in their employer-sponsored DC plans. Recall that m^e is the fraction of earnings that an employer contributes to its employees' DC account, and it is a function of the employees' own contribution rate m^w . To obtain the match function, I first bin worker contribution rates into 0.01-sized intervals and then compute the median employer contribution rate within each interval, which is denoted by \tilde{m}^e . I then regress binned worker contribution rates on median employer contribution rates for $m^w \in [.01, .065]$. Fitted values from the regression ($E[\tilde{m}^e|m^w]$) represent the estimated match function. Approximating the data for $m^w > .065$, I assume that $E[\tilde{m}^e|m^w > .065] = E[\tilde{m}^e|m^w = .065]$, which implies that employer matching contributions stop once workers contribute 6.5 percent of their earnings. The estimated match function is

$$\hat{m}^e(m^w) = \begin{cases} .011 + .39m^w & \text{if } m^w \leq .065 \\ .0364 & \text{if } m^w > .065 \end{cases} \quad (17)$$

⁷⁴Non-pension wealth includes the value of housing.

Additional parameters

The discount factor (β) and real interest rate (r) are based on values consistent with the PEP assumptions.⁷⁵ The maximum DC contribution limit, \bar{C} , is calibrated using 2010 IRS rules. Age-specific mortality rates expressed as the probability of dying within one year, $\{p_a\}_{a=51}^{80}$, are based on the 2010 actuarial life table published by the Social Security Administration (SSA). I average the gender-specific rates into a combined rate for each age and then divide by p_{81} to impose a maximum age of 80 for all simulated individuals. Finally, I use the NBER TAXSIM calculator to compute income and payroll tax liabilities when solving the model.

Table E1 provides a summary of all the calibrated parameters of the model and their sources.

E.2 Solution algorithm

The model does not have an analytical solution, so I solve it numerically through backward recursion. The state variables in the model are (a, e, g, A, b, W^{DC}) , of which g , A , and W^{DC} are continuous and exhibit both within- and between-age heterogeneity. I discretize the continuous state variables over finite dimensional grids of size 6 for g , 20 for A , and 20 for W^{DC} . Grid points for A and W^{DC} are narrowly spaced for low values and widely spaced for high values. To allow for the possibility that simulated individuals make no contributions or withdrawals from their wealth holdings, I increase the grid points for A and W^{DC} by a factor of $(1 + r)$ in every period. Then, iterating backward from the final age for a given value of $\theta = (\sigma, \gamma, \phi, \rho, \sigma_v)$:

1. I compute the probability distribution over the next period's work disutility, which arises from randomness in the AR(1) component f . I employ the Rouwenhorst (1995) approximation for the AR(1) term to obtain $P(f' = f_{l'} | f = f_l)$ for $l, l' = 1, \dots, 6$. I use these probabilities to estimate

$$E_{f'} \left[\max \left\{ V_{a+1}^R(a+1, b', A', W^{DC'}), V_{a+1}^W(X') \right\} \right]. \quad (18)$$

At age 80, $V^W = V^R = 0$, as no workers live beyond that age and there are no bequests.

2. For each value of the state variables (within the current age iteration), I compute V^W and V^R and the associated decision rules as follows:
 - (a) To obtain V^W , I compute the income and payroll tax rate faced by a worker for each potential $W^{DC'}$ choice given the earnings level for the current age. I then search over the $W^{DC'}$ and A' grid points to find the highest value of V^W . I impose the maximum contribution constraint (\bar{C}) and preclude withdrawals of DC wealth while working by setting the associated values of V^W to negative infinity. This grid search defines the decision rules for DC contributions and non-pension saving (or dissaving) while working.
 - (b) To obtain V^R , I compute the income tax rate faced by a retiree for each potential $W^{DC'}$ choice given the annuity income they would obtain by retiring at the current age.

⁷⁵The real interest rate is 2.9 percent, which is based on the 2010 SSA OASDI trust fund report assumption. I assume that the annual discount factor is 0.97 to be consistent with a real interest rate of approximately 3 percent.

I then search over the $W^{DC'}$ and A' grid points to find the highest value of V^R . I preclude contributions to DC wealth in retirement by setting the associated values of V^R to negative infinity. This grid search defines the decision rules for DC contributions and non-pension saving (or dissaving) while retired.

(c) I compute the retirement decision for each value of the state variables by comparing V^W with V^R . If $V^W \leq V^R$, the worker retires. If $V^W > V^R$, the worker remains employed.

3. The solution algorithm terminates when a value function and decision rule have been obtained for each age.

I compute a separate set of value functions and decision rules by imposing pension freezes at ages 56 through 64. A worker who experiences a pension freeze at age a has $b_{a+k}^{DB} = b_a^{DB}$ for all $k > 0$. I apply earnings reductions as estimated from LEHD data for workers who are affected by freezes from ages 56 through 64.⁷⁶ To the extent that career lengths change due to the freeze, I assume that they do not affect the value of Social Security wealth. This assumption reflects the fact that most workers over 55 already have long work histories and do not accrue substantial increases in Social Security benefits through continued work. Other than the change to b^{DB} and the earnings path, all other parameters remain the same.

E.3 Simulation estimator

Having solved the model for a given value of θ by obtaining decision rules for each age, I simulate data as follows.

1. I simulate initial wealth and work disutility draws for 5,000 individuals who are initially aged 51 to 59.⁷⁷ For a fixed θ , I apply decision rules under the no-freeze scenario to obtain work histories and wealth accumulation paths to create a simulated control group. I use linear interpolation to infer the optimal choices for simulated values of A , W^{DC} , and g that lie between grid points.
2. Next, for the same fixed value of θ , I apply decision rules under the freeze scenario for the same population of individuals—that is, individuals with the same initial wealth and work disutility draws—starting five years after the initial age. Individuals in this exercise have the same work status and wealth accumulation choices as the simulation control group for the first five years, but they have potentially different work statuses and wealth accumulation choices after they face a DB freeze. I call this sample the simulated treated group.
3. I compute two sets of moments using the simulated control and treated groups. The first set of moments is the age-specific employment rate for the simulated control group. I compute

⁷⁶Earnings losses occur in the year of the freeze and in the next four years. The estimates are -3.8 percent, -2.1 percent, -1.5 percent, -4.9 percent, and -1.6 percent. See the middle panel of Figure 6.

⁷⁷The share of individuals of each initial age is based on the sample of HRS respondents employed in the private sector who are covered by unfrozen traditional DB plans.

these rates for a 17-year horizon starting from the initial age. The second set of moments is the difference in average employment rates between the two groups or the simulated treatment effect. I compute these differences for 12 years starting from the period of the freeze. There are a total of 29 moments: 17 moments for the employment rate and 12 moments for the treatment effect.

Denote the 29×1 vector of simulated moments by $\hat{\mathbf{h}}^S(\boldsymbol{\theta})$. Denote the analogous 29×1 vector of observed moments from real-world data by $\hat{\mathbf{h}}^D$. The distance between the simulated and real-world moments is

$$\mathbf{m}(\boldsymbol{\theta}) = \hat{\mathbf{h}}^S(\boldsymbol{\theta}) - \hat{\mathbf{h}}^D. \quad (19)$$

The method of simulated moments (MSM) estimate of $\boldsymbol{\theta}$ is given by

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmin}} \mathbf{m}'(\boldsymbol{\theta}) \mathbf{W} \mathbf{m}(\boldsymbol{\theta}), \quad (20)$$

where \mathbf{W} is a weighting matrix.

I set $\mathbf{W} = \zeta \mathbf{I}$, where \mathbf{I} is the identity matrix, and ζ is a scaling vector. The scaling vector re-weights treatment effect moments by a factor of 10 to give them approximately the same numerical importance as the employment rate moments. I make this adjustment because the treatment effect moments are economically very informative but numerically an order of magnitude smaller than the employment rate moments. I first conduct a global search for an initial value by computing the objective function $J(\boldsymbol{\theta}) = \mathbf{m}'(\boldsymbol{\theta}) \mathbf{W} \mathbf{m}(\boldsymbol{\theta})$ for 5,000 points defined by a Sobol sequence over the feasible parameter space. I then use the value of $\boldsymbol{\theta}$ associated with the minimum objective function as the starting value to estimate $\hat{\boldsymbol{\theta}}$ using the Nelder-Meade algorithm.

As shown in Pakes and Pollard (1989), the asymptotic variance of $\hat{\boldsymbol{\theta}}$ is given by

$$\hat{\mathbf{V}} = (1 + N_s^{-1})(\hat{\mathbf{G}}' \mathbf{W} \hat{\mathbf{G}})^{-1} (\hat{\mathbf{G}}' \mathbf{W} \hat{\boldsymbol{\Omega}} \mathbf{W} \hat{\mathbf{G}}) (\hat{\mathbf{G}}' \mathbf{W} \hat{\mathbf{G}})^{-1}, \quad (21)$$

where N_s is the number of simulations, $\hat{\boldsymbol{\Omega}}$ is the variance-covariance matrix of the moment conditions, and \mathbf{G} is the Jacobian matrix:

$$\hat{\mathbf{G}} = \left. \frac{\partial \hat{\mathbf{h}}^S(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}'} \right|_{\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}}.$$

I obtain $\hat{\boldsymbol{\Omega}}$ using a parametric bootstrap procedure in which I estimate the simulated moments 500 times while holding $\hat{\boldsymbol{\theta}}$ fixed but re-drawing random components of the simulation that introduce sampling variability (that is, g and the initial values of A and W^{DC}). I obtain the Jacobian $\hat{\mathbf{G}}$ by computing derivatives numerically. Note that the expression for the variance of $\hat{\boldsymbol{\theta}}$ in equation (21) ignores uncertainty in the calibrated first-step parameters, which I treat as fixed values. I run the simulation once, so $N_s = 1$.

Table E1: Calibrated parameters

Parameter	Value	Description	Source
r	0.029	Real interest rate	HRS PEP/SSA
β	0.97	Annual discount factor	$r \approx .03$
π	0.028	Inflation rate	HRS PEP/SSA
α_m	0.05	Nominal wage growth	HRS PEP/SSA
\bar{C}	\$54,500	Maximum combined contribution limit to DC plans (2010 level)	IRS
$\{\beta_{1,g}, \beta_{2,g}\}_{g=m,f}$		Age-earnings profile curvature parameters by gender (m =male, f =female)	HRS
$\{s_a\}_{a=51}^{59}$		Share of individuals with zero DC wealth; initial ages 51-59	HRS
$\{\mu_a^A, \sigma_a^A\}_{a=51}^{59}$		Non-pension wealth distribution parameters for workers with zero DC wealth; initial ages 51-59	HRS
$\{\mu_a, \Sigma_a\}_{a=51}^{59}$		Bivariate wealth distribution parameters for workers with positive DC wealth; initial ages 51-59	HRS
$E[\tilde{m}^e m^w]$		Employer DC contribution rate expressed as a function of the worker's contribution rate	HRS
$\{b_a^{DB}\}_{a=51}^{80}$		Age-specific DB pension annuity	HRS
$\{b_a^{SS}\}_{a=51}^{80}$		Age-specific Social Security annuity	HRS
$\{p_a\}_{a=51}^{80}$		Age-specific mortality rate adjusted so that $p_{81} = 1$	SSA
$\{e_a\}_{a=51}^{80}$		Age-specific annual earnings	HRS
$\tau(\cdot)$		Federal income and payroll tax liabilities (2010 laws)	NBER TAXSIM

Notes: See Appendix E.1 for details.

F Simulated pension freezes in the HRS

This appendix describes how I construct the simulation underlying Figure 3. The simulations are based on the sample of DB-eligible HRS respondents from the 2010 survey wave who are employed in the private sector and whose DB plans have not been frozen or converted to CB as of the survey interview. Using these data, I consider three age-specific components of compensation: DB wealth, earnings, and DC wealth.

For survey respondents, the PEP provides estimates of DB wealth at each potential quit date on the basis of equation (15). In addition, the PEP provides earnings projections at each potential quit date for each respondent on the basis of equation (16). DC wealth is reported by respondents as of the survey year (2010). I use these measures along with self-reported own and employer contribution rates (m^w and m^e) to construct estimates of past and projected future values of DC wealth using the law of motion described in equation (4).⁷⁸ I then combine these estimates of DB wealth accruals, earnings, and DC wealth accruals to compute total compensation at each potential quit age using equation (7).⁷⁹ Finally, I average the data across respondents and quit dates to obtain the “no-freeze” path of compensation.

To simulate a hypothetical DB pension freeze for workers aged a^F , I make two changes. First, I assume that the annuity value of nominal DB wealth is frozen as of a^F . This is equivalent to receiving no new accruals either due to tenure increases or due to earnings growth.⁸⁰ Second, I assume that respondents with no DC wealth as of a^F —about 60 percent of the sample—begins contributing to a hypothetical new DC plan starting at age a^F . I assume that contribution rates for these workers are equal to the sample averages of m^w and m^e for respondents with non-zero DC wealth. Respondents who have non-zero DC wealth are assumed to continue contributing at the same rate they did before age a^F . Having defined post-freeze DB and DC wealth evolution, I combine the estimates of post-freeze DB wealth accruals, earnings, and DC wealth accruals to compute total compensation at each potential quit age using equation (7). I then average these data across respondents and quit dates to obtain the “Freeze at a^F ” path of compensation. I conduct these calculations for different values of a^F to obtain different post-freeze paths of total compensation.


⁷⁸I assume that respondents and their employers contribute at the same rate in all years. I rely on the RAND HRS files, which convert respondent reports of own and employer contributions to percentages of earnings.

⁷⁹All three components (DB wealth, earnings, and DC wealth) are converted to 2010 dollars. I do not account for taxes in these simulations because online tax calculators cannot be used within the restricted setting in which pension data are accessed.

⁸⁰By preventing future growth in the nominal value of benefits, DB freezes have the effect of lowering the present value of DB wealth in each year subsequent to the freeze, thereby generating negative future accruals. This is because benefits stay fixed, but the horizon over which they can be collected falls.

G Supplementary tables and figures

Figure G1: Impact of freezes on employment by gender (56- to 64-year-olds)



Notes: This figure shows the time path of the treatment effect of the freeze on employment rates for men and women separately. Workers are 56 to 64 years old at the time of the freeze.

Table G1: Pre-period summary statistics split by age group (without propensity score re-weighting)

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	0.002	0.90	59.6	0.034	0.29	67.0	-0.011	0.68
Male	0.466	0.034	0.30	0.478	0.032	0.38	0.478	0.030	0.50
High school	0.226	0.005	0.77	0.231	0.000	0.99	0.288	0.006	0.75
Some college	0.329	-0.011	0.24	0.314	-0.010	0.26	0.292	-0.007	0.42
College or more	0.391	0.002	0.93	0.394	0.007	0.80	0.315	0.004	0.87
White	0.793	0.020	0.28	0.819	0.027	0.16	0.824	0.042	0.05
Black	0.093	-0.009	0.24	0.081	-0.013	0.05	0.074	-0.016	0.05
Hispanic	0.065	0.001	0.95	0.052	0.000	0.99	0.051	-0.006	0.46
Other race	0.049	-0.013	0.11	0.048	-0.014	0.09	0.051	-0.020	0.06
Earnings (\$)	61,610	2051	0.61	62,240	1,243	0.73	48,130	-3915	0.10
Tenure at $l - 5$	7.6	-0.817	0.07	8.2	-0.898	0.04	8.2	-0.760	0.07
Retired	0.020	0.002	0.17	0.054	0.004	0.23	0.188	0.008	0.50
In labor force	0.966	-0.003	0.16	0.930	-0.005	0.23	0.783	-0.010	0.49
Switched $l - 5$ employer	0.034	0.025	0.01	0.030	0.024	0.02	0.027	0.006	0.07
Pension and firm characteristics									
Log DB pension wealth/active participant	10.18	-0.037	0.91	10.27	-0.077	0.82	10.5	-0.028	0.87
Log DB pension accrual/active participant	7.73	-0.205	0.52	7.80	-0.264	0.41	7.98	-0.325	0.05
Pension plan claim age	62.7	-0.015	0.96	62.7	0.036	0.89	64.2	0.137	0.41
Log firm size	8.46	0.066	0.90	8.32	0.079	0.88	7.42	-0.273	0.55
Fraction workforce ≤ 45	0.584	0.007	0.63	0.570	0.005	0.74	0.553	-0.015	0.27
Fraction workforce [46,50]	0.146	-0.003	0.50	0.142	-0.001	0.68	0.136	0.000	0.99
Fraction workforce [51,55]	0.124	-0.002	0.60	0.128	-0.002	0.69	0.122	0.002	0.64
Fraction workforce [56,60]	0.086	-0.001	0.89	0.094	0.000	0.94	0.095	0.005	0.21
Fraction workforce [61,65]	0.041	0.000	0.98	0.045	0.000	0.96	0.061	0.005	0.20
Fraction workforce [66,70]	0.011	-0.001	0.62	0.012	-0.001	0.66	0.021	0.001	0.55
Fraction workforce ≥ 71	0.007	-0.001	0.46	0.008	-0.001	0.59	0.012	0.002	0.38
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: Unless otherwise noted, statistics reported in the table average over the five-year period preceding any freeze activity. Pension wealth per active participant is computed as the present value of the liability owed to active participants divided by the number of active participants. All dollar values are expressed in 2010 terms. Tenure is understated because the LEHD does not capture the complete history of an employer-employee relationship when states enter the data set 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.

References

- 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.
- 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.
- Jarmin, Ron S. and Javier Miranda**, “The Longitudinal Business Database,” *Center for Economic Studies Working Paper 02-17*, 2002.
- Pakes, Ariel and David Pollard**, “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 1989, *57* (5), 1027–1057.
- Rouwenhorst, K. Geert**, “Asset Pricing Implications of Equilibrium Business Cycle Models,” in Timothy F. Conley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995, pp. 294–330.