NBER WORKING PAPER SERIES

DO EMPLOYER PENSION CONTRIBUTIONS REFLECT EMPLOYEE PREFERENCES? EVIDENCE FROM A RETIREMENT SAVINGS REFORM IN DENMARK

Itzik Fadlon
Jessica A. Laird
Torben Heien Nielsen

Working Paper 21665
http://www.nber.org/papers/w21665

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2015

We thank Raj Chetty, David Cutler, Matthew Gentzkow, Nathan Hendren, Lawrence Katz, David Laibson, Daniel LeMaire, Søren Leth-Petersen, Christian Schultz, and numerous participants at the Harvard Labor/Public Finance Lunch and the Zeuthen Workshop for helpful comments and discussions. This research was supported by the Danish Council for Independent Research (grant number 11-104234). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Do Employer Pension Contributions Reflect Employee Preferences? Evidence from a Retirement Savings Reform in Denmark

Itzik Fadlon, Jessica A. Laird, and Torben Heien Nielsen

NBER Working Paper No. 21665
October 2015
JEL No. J30,J32,J33

ABSTRACT

This paper studies how firms set contributions to employer-provided 401(k)-type pension plans. Using a reform that decreased the subsidy for contributions to capital pension accounts for Danish workers in the top income tax bracket, we provide strong evidence that employers' contributions are based on their employees' savings preferences. We find an immediate decrease in employer contributions to capital accounts, whose magnitude increased in the share of employees directly affected by the reform. This response was large relative to average employee responses within private IRA-type plans and was accompanied by a similar-magnitude shift of employer contributions to annuity accounts.

Itzik Fadlon
National Bureau of Economic Research
1050 Massachusetts Ave.
Cambridge, MA 02138
fadlon@fas.harvard.edu

Jessica A. Laird
Harvard University
Department of Economics
Littauer Center
1805 Cambridge Street
Cambridge, MA 02138
laird@fas.harvard.edu

Torben Heien Nielsen
University of Copenhagen
Department of Economics
Øster Farimagsgade 5, Building 25
DK-1353 Copenhagen K
torben.heien.nielsen@econ.ku.dk

1 Introduction

With the decline in the prevalence of defined-benefit pension plans, individual savings in defined-contribution accounts are becoming an increasingly important income source for post-retirement consumption. A large and growing portion of savings balances in defined-contribution accounts is within employer-sponsored pension plans, such as 401(k)s.[1] Recent research has underlined the important role that employers' decisions play in determining employees' actual savings within employer-sponsored accounts, since most workers do not actively deviate from the default options which are set by their employer (Madrian and Shea 2001; Choi et al. 2004; Choi et al. 2007; Beshears et al. 2009; Gelber 2011). The impact of employers' decisions has also been found to translate into large effects on individuals' overall level of savings, taking into account other savings vehicles outside of the workplace (Chetty et al. 2014). This research has led policymakers to consider introducing policies that encourage employer contributions to pension accounts in order to increase individuals' retirement savings (Beshears et al. 2010).

But how, in practice, do private firms set contributions to employer-sponsored 401(k)-like pension plans, and, specifically, are employers' choices based on their employees' savings preferences?

From a theoretical standpoint there is no one definite answer. Standard models of efficient compensation arrangements by competitive firms, all the way back to Rosen (1974), predict that employers will provide benefits when firms can purchase goods or services more cost-effectively than employees, and will provide the optimal package that their workers will value most highly.[2] In our application, provision of pension benefits may be less costly for the employer if there are economies of scale with respect to the purchasing and

[1] In 2014, 35% ($8.3 trillion) of retirement assets in the U.S. were held in defined-benefit plans, while a much larger share of 58% ($13.8 trillion) was held in defined-contribution accounts: 30% in Individual Retirement Accounts (IRAs) and 28% in employer-sponsored plans, mostly 401(k)s (Choi 2015).

[2] Both Rice (1966) and Phelps (1973) put a particular emphasis on the advantages of group purchases as a reason for the provision of benefits. For a review of the efficient compensation literature see Oyer (2005).

managing of the pension product (e.g., in terms of average fees to the financial service provider per saver), or if it is costly for individual households to optimally choose complex financial products.[3] However, the standard model assumes that individuals value their employer optimally setting pension contributions on their behalf. Given evidence that most employees are inattentive with regards to their own savings and may have only imperfect knowledge about their employer pension plans (Mitchell 1988; Gustman et al. 2009), it is not obvious that this assumption holds.

Some non-standard models, which assume that firms are rational benevolent planners but that some households are not fully rational, will have similar ambiguous predictions. Purely paternalistic firms will provide their employees with optimal retirement packages, taking into account their employees' misoptimization (Choi et al. 2003; Carroll et al. 2009, Cremer et al. 2008; Cremer and Pestieau 2011; Goda and Manchester 2013; Roeder 2014; Fadlon and Laibson 2015). However, if firms compete for employees, this prediction will hold only if the non-rational agents are aware of their biases and hence value the firm's paternalism. In our context, this would require employees to acknowledge they would save too little on their own, so that they will value ex-ante mandated savings by their firm. If some workers are naive with respect to their imperfect optimization, then they will not appreciate employer contributions to savings accounts at the cost of lower wages today. In that case, there will be reduced incentives for competitive firms to engage in the costly effort of providing pension plans and tailoring them to their employees' hypothesized preferences.[4]

[3]Both of these conditions are likely to be satisfied in our setting given the general prevalence of low levels of financial literacy (Brobeck 1990; Bernheim 1998; Hilgert et al. 2003; Lusardi and Mitchell 2007; Hastings et al. 2013) and given that in Denmark administrative and account managing costs paid to the financial service provider are on average lower in group accounts through the employer than in private accounts. In employer-sponsored accounts these costs are on the order of 1%-1.5% of the deposits annually, where the lowest costs are found in unionized mandatory funds (Økonomi- og Erhvervsministieriet et al. 2003). Reports by consumer organizations suggest that these costs are about 40% lower than those in private accounts (see, e.g., http://penge.dk/pension-skat/tjek-om-din-pension-bliver-rippet).

[4]This would be exacerbated by any direct costs of providing or updating pension plans

Beyond the question of whether firms tailor pension plans to employees' saving preferences, it is important for policy design to study the nature of firm responses to changes in these preferences. The increased reliance on employer-based savings as a source of consumption after retirement necessitates understanding empirically how fast and to what extent employers respond to shocks to the economic environment that can alter employees' savings incentives.

In this paper we empirically analyze how firms set characteristics of their contributions to employer-sponsored pension plans in practice.[5] To do so, we exploit a reform to the Danish retirement savings system. This reform differentially affected employees according to their exact location on the labor-income tax schedule and differentially changed tax deductions for contributions to two types of savings accounts: "capital" savings accounts, which are paid out in full at retirement, and "annuity" savings accounts, which are paid out as an annuity. Specifically, in 1999, the Danish government decreased the subsidy to contributions to capital pension accounts for workers in the top income tax bracket, while the subsidies to capital pension contributions for workers in lower tax brackets and for annuity pension contributions remained unchanged.

We find that immediately following the reform, employers significantly decreased their annual contributions to capital pension accounts. The average decrease was on the order of 27% – 0.76 percentage points (pp) on a baseline contribution rate of 2.81 pp. This decrease was entirely driven by firms in which some share of the workforce was directly affected by the reform, with no responses in workplaces in which all employees had earnings below the top income tax bracket. Moreover, the response strongly and continuously increased in the share of employees above the top tax threshold. We find that

by the employer.

[5]To our knowledge, this is the first paper that studies this specific topic. Some papers analyze firms' choice of whether to offer a pension plan (or the choice between defined-benefit or defined-contribution plans), but these papers usually focus on supply-side factors that affect firms' cost of providing the plan, e.g. the size of the firm (Aaronson and Coronado 2005; Dummann 2008; Hernæs et al. 2011). Papers that do analyze demand-side factors, relate these to the individual take-up of plans rather than the firms' decision to offer them (Aaronson and Coronado 2005; Dummann 2008).

an additional 10% of workers at the top bracket led to an additional decrease of more than 0.2 pp in employer contributions, so that workplaces in which all employees were at the top bracket experienced a significant drop of more than 2 pp on a base of 3.5 pp.

In order to put the employers' responses and their magnitude in context, the Danish setting allows us to compare contributions within employer-provided 401(k)-like accounts to individual contributions within private IRA-like accounts, which were equally affected by the reform. We find that for individuals at the top bracket almost the entire change in overall capital retirement savings was attributable to employer responses. We also show that the clear gradient of changes in employer-provided accounts with respect to the share of directly affected individuals in the workplace disappears in the analysis of changes in private accounts, suggesting that employer responses were not crowded out by individual responses in other closely substitutable accounts.

By changing the relative prices of contributions to capital and annuity accounts, the reform rendered contributions to annuity account more financially attractive through a substitution effect, but also led to an income effect that would push toward an overall decrease in pension contributions. Studying employer contributions to annuity accounts, we find that employers compensated for the decrease in capital contributions with an equally-large increase in annuity contributions, with no decrease in total pension contributions.[6] This suggests that at the employer-level the effect was driven by a substitution effect. In fact, annuity accounts serving as a close substitute is likely the reason there was such a large response in capital accounts. We also show that the potential income effect of the reform had no effect on other means of employee compensation, namely, labor income.

The paper proceeds as follows. In Section 2, we discuss the institutional setting of the policy change and the data that we use. In Section 3, we pro-

[6]These employer responses are in contrast to individual responses in private accounts, for which Chetty et al. (2014) find a shift of 57 cents to annuity accounts for each DKr that individuals would have contributed to capital accounts.

vide the empirical analysis of employers' responses to the reform and their heterogeneity with respect to workplace composition. We focus on changes in employers' contributions to capital pension accounts and their timing, but also analyze shifts to contributions to annuity accounts. To put the magnitude of the responses in context, we compare these changes in contributions to employer-provided accounts to individuals' responses within their private accounts. Section 4 concludes.

2 Institutional Details and Summary of Data

2.1 Institutions

This section provides the necessary background on Danish retirement institutions that is important for our empirical analysis.[7] In Denmark, there are two types of defined-contribution (DC) pension savings accounts similar to the US – employer-sponsored accounts, similar to 401(k)s, and private accounts, similar to IRAs. Employer and private DC accounts have equivalent tax properties but are completely independent, which makes them close substitutes. Within both the employer and the private DC pension plans, there are two types of tax-preferred accounts: capital pension accounts and annuity pension accounts. Capital pension accounts are paid out as a lump sum and taxed at 40% on payout, while annuity pension accounts are paid out over several years and are taxed as personal income. Balances in capital pension accounts can be converted to annuity pensions when they become liquid, but the reverse is not allowed. Contributions to both types of accounts are tax deductible at the time of contribution (as in traditional non-Roth 401(k)s and IRAs), and capital gains are taxed at 15%, compared to approximately 29% for assets in taxable accounts.

Our empirical research design exploits a 1999 tax reform, which aimed at reducing the generosity of capital accounts and incentivizing a shift to annuity accounts. To do so, the reform reduced the average deduction for contributions to capital pensions from 59 cents per DKr to 45 cents per DKr for individuals

[7]For additional information see OECD (2009) or Bingley et al. (2007).

in the top income tax bracket. The deduction for those in the lower tax bracket remained the same at 45 cents per DKr. While there were additional changes associated with this reform, they were orthogonal to whether a person was just above or below the top threshold.[8]

Most jobs in Denmark (roughly 80%) are covered by collective bargaining agreements between worker unions and employer associations. These agreements often have a pension plan in which a fixed proportion of an individual's earnings is paid into a retirement account. For the 20% of jobs that are outside the common agreements, employers set contribution rates to capital and/or annuity accounts for their workers.[9] While individuals cannot change the total contribution rate, they can choose a different allocation across capital and annuity accounts, but only if their pension fund allows both types of accounts.

2.2 Data Sources, Sample Selection, and Variable Definitions

We merge data from several administrative registers of the Danish population – the income tax register, the population register, and the Danish Integrated Database for Labor Market Research (IDA) – to obtain annual information on Danish employees and their matched firms from 1996 to 2001. These registers include data on taxable labor earnings, contributions to pension accounts, occupation, industry, and employees' demographics (such as age and educational attainment). All income and savings variables used in the analysis are based on third-party reports: earnings and pension contributions are reported directly by employers and pension funds to the tax authority.

Starting from the population data-set, we impose four restrictions to obtain our primary analysis sample. First, we exclude individuals under age 20

[8]Other changes that were associated with the reform include a reduction in the deduction value of negative capital income, the possibility to initiate a private rate pension plan (a special type of annuity pension plans) after age 55, a reduction in the bottom bracket tax rate, a move to equalize taxation on all liquid assets (i.e., stocks vs. bonds), and a decrease in the value of the Voluntary Early Retirement Plan.

[9]The contributions rates, default portfolio allocations, and administration fees are set by bargaining between the pension fund and the employer, which is usually represented by the heads of HR departments, CFOs, and pension brokers. Updates to the employees' plans, e.g., in response to taxation changes, are made by the pension fund and the employer.

or over age 60, at which the majority of the Danish workforce is eligible for early retirement benefits and retirement savings are eligible for withdrawal. Second, we focus our analysis on the 20% of workers who are outside collective bargaining agreements for which contribution rates are set by the firm (rather than within collective bargaining agreements).[10] To isolate the jobs that are not covered by collective bargaining, we exclude workers in the public sector or in blue-collar occupations, since they are likely covered by collective agreements.[11] Therefore, our analysis sample consists of workers in private firms with white-collar occupations.[12] Third, we exclude observations of workers with self-employment income because their "employer" contributions are not set by a firm. Finally, we exclude occupation-firm cells with fewer than five employees in order to decrease measurement error, as such small cells are unlikely to be treated as an independent unit by employers.[13] Overall, our sample choice allows us to study how private firms in competitive markets design their employees' pension plans.

Since firms often set contribution rates separately by occupation, we run our analysis at the occupation-firm level, and differentiate occupations at the 2-digit occupation code level.[14] We measure contribution rates to employer-sponsored accounts as contribution levels divided by taxable labor income. This measure of contribution rates may vary within an occupation-firm cell since employers may set pension contributions at a finer level within the firm

[10]While it would be interesting to additionally analyze how pension plans are designed within a collective bargaining setting, the data do not allow us to match workers to unions and firms to employer associations.

[11]See Appendix E for a complete description of occupations that we define as white- or blue-collar.

[12]Still, some white-collar jobs in the private sector are covered by collective bargaining. Therefore, in Appendix D, we assess how inadvertently including workers covered by collective agreements may affect the results, and show that it likely attenuates our estimates.

[13]Our results are not sensitive to this choice – see Appendix A for analyses that vary this minimal cell-size restriction.

[14]Due to measurement error in many-digit occupation codes, our choice for the analysis is the 2-digit code level. However, actual employer contributions may be set at a higher- (or lower-) digit occupation code level. This causes some measurement error in our identified decision unit. Our results stay similar if we aggregate occupations at the 1-digit level, or even at the firm level.

than the 2-digit code that we use, and since individuals can choose a different distribution of contributions between capital and annuity accounts than the employer's default when employers offer both accounts. Therefore, to identify the default contribution rates chosen by employers, we use the median contribution rate within an occupation-firm cell as our measure. In Appendix A, we assess the sensitivity of our results to other measures of defaults (namely, modes) and find very similar results. We also conduct various tests which verify that our results are driven by firm responses rather than by individual responses. For ease of discussion, we refer in the remainder of the paper to a 2-digit occupation-firm cell as a "workplace" and to the median contribution rate within a workplace (in a given year) as the "employer contribution rate" (or the "default").

2.3 Summary Statistics

Table 1 presents summary statistics for the sample of private white-collar wage earners between ages 20-60, in workplaces with at least five employees.[15] Our sample contains 2,020,705 worker-year observations from 1996-2001. These amount to 84,764 workplace-year observations with a total of 26,775 unique workplaces.

To provide an overview of contributions to retirement savings accounts in Denmark prior to the reform, Table 1 reports information on individual-level pension contributions for the year 1998.[16] Before the reform in 1999, contributions to employer-sponsored capital accounts were on average 3.6% of annual earnings, where 66% of workers had positive contributions to these accounts. The average of contributions to employer-sponsored annuity accounts was similarly at 3.7% of labor income, where 77% had positive contributions to these accounts. In contrast, individual contributions to both capital and annuity private pension accounts were much lower, with average contribution rates of 1.1% and 0.4%, respectively.

[15] During our sample period, 57% of wage earners were in the private sector, and 70% had white-collar occupations.

[16] All monetary values are reported in nominal Danish Kroner (DKr), where the exchange rate during this time period was approximately DKr 6.5 per US $1.

Importantly for our design, 49% of workers were above the top labor-income tax threshold, with a sizable standard deviation of the fraction of workers above the threshold across workplaces on the order of 35%. Given the restriction to at least five workers per occupation-firm cell, the average workforce size is 22, while the median cell size is 9.

3 Empirical Evidence

In this section, we analyze how employer contributions to pension accounts responded to the reform, namely, to the decrease in subsidies for capital pension contributions for workers with labor income in the top tax bracket. We begin by analyzing changes in employer contributions to capital accounts and their sensitivity to the share of workers who were directly affected by the reform. We then assess the magnitude of these responses to the reform by comparing employer responses to individual responses within private retirement accounts. Finally, we explore other potential margins of firm responses, in particular, whether changes in capital contributions translate into changes in overall savings or whether firms substitute contributions to annuity accounts.

3.1 Employer Responses in Contributions to Capital Pension Accounts

In the years preceding the reform, employer capital contributions steadily increased, such that they were on average 2.81% in 1998 (see Panel A of Appendix Figure 1).[17] In contrast, in 1999, when the capital subsidy decreased for workers above the top threshold, the average employer capital contribution rate decreased by 0.76 percentage points (pp). However, this average drop of 27% aggregates the responses of all the employers in our sample of private firms and white-collar occupations. Since the reform changed the savings incentives only for employees in the top labor-income tax bracket, our analysis focuses on the heterogeneity of firm responses with respect to the share of the workforce

[17]Note that these are employer contributions (measured by workplace-level medians) as opposed to individual-level contributions to employer-sponsored accounts that are reported in Table 1.

that was directly affected by the reform.

Graphical Analysis. To test whether and to what extent employers' capital contribution responses to the reform increased in the share of workers above the top income threshold, we divide workplaces into equal-sized groups by the fraction of employees above the threshold within a workplace. We begin by plotting in Figure 1 the mean employer capital contribution rate against the mean fraction of employees above the top threshold for each group in years 1996-2001. Panel A shows that before the reform, employer capital contributions were increasing in the fraction of workers above the threshold and that the slopes of this relationship were similar across years.[18] However, immediately following the reform – that took effect in 1999 – there is a significant change in this relationship, such that employer capital contribution rates became largely decreasing in the fraction of workers above the threshold. The decrease in employer capital contributions after the reform, i.e., the vertical distance between the lines of years 1998 and 1999, is noticeably larger for workplaces with a higher fraction of directly affected workers.

To clearly see these changes, Panel B of Figure 1 displays the year-to-year differences in employer capital contributions as a function of the fraction of employees above the threshold for each year from 1996 to 2001. This figure shows that annual changes in contributions were uniform across different shares of workers above the top tax threshold in the years prior to the reform. However, between years 1998 and 1999, the year of the reform, workplaces with no affected workers did not change their contributions, while those with a larger share of affected employees decreased their contribution rates in larger magnitudes. This change continuously increased in the share of the workforce at the top bracket, with about a 1.9 pp decrease for workplaces with the highest share of affected employees. This is a large response, since completely exiting capital accounts in workplaces with the highest share of affected employees – which is an upper bound to their response – would imply a 3.5pp reduction

[18]This is consistent with the fact that top-bracket workers enjoyed a larger subsidy for capital contributions on the margin, but as it is a cross-sectional relationship, there is a variety of other reasons for this pattern such as different preferences for savings across individuals with different labor income levels.

in contributions. In the subsequent years (2000 and 2001) there were some delayed responses to the reform, but the gradient with respect to the share of affected workers in those years is much smaller.

Overall, the graphical evidence clearly reveals that employers with a greater share of affected workers had larger capital contribution reductions in response to the reform. This suggests that employers are indeed responsive to changes in their employees' saving incentives, consistent with the hypothesis that employer-provided pension plans reflect the savings preferences of their particular workforce composition.

Regression Analysis. To quantify the firms' responses to the reform, we estimate regressions of the relationship between the change in employer capital contributions and the fraction of workers above the threshold. This also allows us to test the sensitivity of our results to a flexible set of controls. Our baseline estimating equation is of the form:

$$y_{ft} = \beta_0 + \beta_1 above_{ft} + \sum_{s=1996, s\neq 1998}^{2001} [\beta_s (I_{t=s} \times above_{ft}) + \mu_s] + X_{ft} + \varepsilon_{ft}. \quad (1)$$

The outcome variable y_{ft} is workplace f's behavior in time t , i.e., annual outcomes grouped at the occupation-firm level. Our first and main outcome variable is the change in employer capital contribution rates from year $t-1$ to year t. The right-hand side variables include the fraction of employees above the threshold in an occupation-firm-year cell ($above_{ft}$), year fixed effects (μ_s), and year dummies interacted with the fraction of employees above the threshold ($I_{t=s} \times above_{ft}$). In this specification, we omit 1998 as the baseline year, so that all the coefficients β_s are estimated relative to 1998. We choose a specification linear in the fraction of affected employees, since Panel B of Figure 1 revealed an approximately linear relationship between the change in employer capital contributions and the share of the workforce above the threshold. The main coefficient of interest is β_{1999}. This coefficient captures the relationship between the change in annual employer capital contributions

and the fraction of workers above the threshold in 1999 compared to that in 1998, thus estimating the effect of the reform on this relationship. The vector X_{ft} includes various sets of controls, which we add in order to verify the robustness of the estimated effect, as the share of employees above the threshold may be correlated with other characteristics of the firm that may affect the change in contribution rates.

Panel A of Table 2 reports the coefficients on the share of employees above the threshold and the interaction of this share with indicators for years 1996 through 2001 (omitting 1998) in regressions that include various sets of controls. In all columns, we include year fixed effects and cluster standard errors at the workplace level. We multiply the coefficients by 100 to convert them into percentage-point units. Column (1) estimates the baseline regression and in columns (2) to (4) we successively add controls to the vector X_{ft}. Importantly, we add high-order polynomials of the mean workplace-level income, separately for workers below and above the top tax threshold, as well as their interactions with the year dummies.[19] This allows us to further isolate the relationship between employer responses and whether employees are exactly above or below the threshold, by adding an underlying flexible continuous relationship between employer behavior and average labor income.[20] The additional controls that we include are the number of workers in a workplace and its square, as well as their interactions with year indicators, workplace (i.e., 2-digit occupation-firm) fixed effects, and 2-digit occupation-year fixed effects.

Across all specifications the results are very stable and are in accordance with the graphical results. There is no meaningful relationship between changes in employer capital contributions and the share of employees above the top

[19]The reported estimates are for polynomials of degree five, but the results are robust to higher- and lower-degree polynomials and are available from the authors on request. The decline in the number of observations from specification (1) to (2) is due to the inclusion of controls for average income separately for employees above and below the top bracket, which excludes observations in which all employees are either above or below.

[20]These controls alleviate concerns, for example, that "good" firms with higher wages may be more likely to respond to the reform and also have a higher fraction of workers above the top tax threshold. We estimated regressions that add controls for percentiles of the workplace's distribution of employee income and found similar results. The analysis is available from the authors on request.

threshold prior to the reform in years 1996-1998. However, the coefficient on the fraction of employees above the top threshold interacted with 1999 is approximately -2.2 pp and statistically significant at any conventional significance level. Focusing on the specification of column (4) with the full set of controls, this implies that in 1999 employers in workplaces with 100% of employees above the top-income tax threshold decreased their capital contribution rate by an average of 2.18 pp more than employers in workplaces with 0% of employees above the threshold. For years 2000 and 2001, the coefficients on the fraction of employees above the top threshold are -0.59 and -0.55, respectively, and statistically significant. These patterns are consistent with firms responding substantially just after the reform took place, with a small degree of delayed or gradual responses by some firms.[21]

The stability of the estimated effect across the different specifications suggests that the estimated relationship is not driven by omitted variables. Additionally, in Appendix A, we show that the results are robust to using modes instead of medians as our measure of defaults, and that the results are unlikely to be driven by individual responses. We also show that the results stay similar when we vary the minimal size of workplaces that we include in the analysis. In Appendix C, we demonstrate that the employer responses in contribution rates were attributable to changes in capital contributions (that is, the numerator) rather than changes in labor income (the denominator). Finally, in Appendix D, we show that inadvertently including workers who were covered by collective bargaining likely only attenuates our results.

In sum, our analysis is consistent with the notion that employers design pension plans to reflect the savings preferences of their workforce, and that they respond immediately to changes in their employees' incentives. In the next section, we gauge the magnitude of the employer responses that we estimated.

[21]We find that for workplaces with more than 50% of workers above the top threshold, approximately one half of the decrease in employer capital contributions after 1999 is delayed and attributable to firms that did not respond in 1999, and the other half is gradual and attributable to additional responses by firms that responded in 1999.

3.2 Employer vs. Individual Responses

In the analysis above, we showed that the average response of employers was large relative to their baseline contribution rates to capital accounts. In this section we assess the magnitudes of the employer responses to the reform by comparing them to the responses of individuals within their private accounts.

The ideal experiment that compares individuals' savings behavior and employers' savings behavior on the individuals' behalf would randomly assign savings decisions to either individuals or their respective employers. To mimic this experiment we exploit the Danish setting that provides us with administrative records of employee-level savings contributions to both employer-sponsored 401(k)-like accounts and private IRA-like accounts that are managed by the individuals themselves. We focus the analysis only on those who were directly affected by the reform – that is, employees at the top bracket of the labor income tax schedule – and compare their responses in private accounts to those of their employers in their employer-sponsored accounts.

In Figure 2, we divide the sample of affected workers into equal-sized groups according to the share of workers above the top threshold in their workplace. Panel A plots the change in the default contribution rate to employer-sponsored capital accounts, while Panel B plots the change in the average contribution rate to private capital accounts.

One key difference between employer and individual responses is that for any fraction of employees at the top bracket, the decrease in employer contributions to capital accounts was larger than the individuals' responses. The latter is at most a decrease of 0.75 pp, while the smallest decrease in employer contributions is more than 1 pp. This suggests that most of the overall decrease in capital contributions due to the reform was attributable to employer, rather than individual, responses. It is, in part, due to the fact that baseline contributions to capital accounts in the years prior to the reform were much smaller in individual accounts compared to employer-sponsored accounts (see Table 1). Another noticeable difference between the two panels of Figure 2 is that there is no gradient in private accounts with respect to the share of

employees above the top threshold, while there is a pronounced gradient in employer-sponsored accounts. This suggests that the response of employers in capital pension plans was not crowded out by individual responses in private plans.

To account for differential baseline contribution rates to private vs. employer-sponsored accounts and to understand better what underlies the aggregate responses, we analyze in Appendix B the changes in the two types of accounts at the employee level by studying their respective distributions.

Altogether, the comparison of responses between employer-sponsored accounts and private accounts reveals that, at the aggregate, the effect of the reform on overall savings for retirement was mostly driven by employer-level responses, underlining the large relative magnitude of employer responses. Next, we study whether the decrease in employer contributions to capital accounts translated into a decline in overall savings or into increased contributions to other substitutable accounts.

3.3 Substitution into Annuity Accounts

The empirical analysis of employer responses to the reform has focused so far on capital contributions. To understand the effects of the reform on the overall employer-sponsored savings portfolios of employees, we proceed by analyzing how employer contributions to annuity accounts may have changed.

The reform's decrease in subsidies to capital accounts had two main effects on employees' savings incentives for workers in the top income tax bracket. As it made contributions to savings accounts less attractive, the reform caused a negative income effect that pushed toward lower levels of total pension savings. At the same time, the reform created a substitution effect due to the decrease in the relative price of contributions to annuity accounts. The relative forces of these two effects determine whether and to what extent employers responded in their contributions to annuity accounts and in their employees' total compensation.

Figure 3 plots changes in employers' contributions to annuity accounts. We begin with Panel A, which plots changes in employer contribution rates to

annuity vs. capital accounts by year. This graph shows that the time series of changes in employer contributions to annuity accounts essentially mirrored the changes in capital contributions. Moreover, Panel B of Figure 3, which replicates Panel B of Figure 1 but with changes in employer annuity contributions as the outcome variable, reveals that employer responses in annuity accounts also mirrored the responses in capital accounts as a function of the share of employees above the top threshold. Before the reform, annual changes in both annuity and capital accounts were uniform across workplaces with different fractions of workers above the top tax threshold. In 1999, in response to the reform, employers decreased their capital contributions and increased their annuity contributions as a function of the share of their affected employees and in similar magnitudes. Panel B of Table 2 estimates equation (1) with a full set of controls and with changes in employer annuity contributions as the outcome variable. We find that in 1999 employers in workplaces with 100% of employees above the top-income tax threshold increased their annuity contribution rate by an average of 1.96 pp more than employers in workplaces with 0% of employees above the threshold, alongside the decrease of 2.18 pp in capital contributions. In fact, studying the sum of these responses in Panel C of Table 2, the evidence is consistent with full compensation of the decrease in capital contributions by an increase in annuity contributions, so that the change in overall employer contributions is not statistically different from zero.

In sum, the results suggest that the response at the employer-level was driven by a substitution effect, so that the decrease in capital contributions was almost-fully compensated for by an increase in annuity contributions, with no statistically significant effect on overall employer contributions. These employer responses are in contrast to individual responses in private accounts, for which Chetty et al. (2014) find a shift of 57 cents to annuity accounts for each DKr that individuals would have contributed to capital accounts. In Appendix C, we additionally show that there was no average effect on labor income, so that we do not find evidence that the potential income effect of the reform was offset by higher wages.

4 Conclusion

This paper provides evidence that employers set contributions to pension savings accounts in accordance with the savings preferences of their workforce, and that they respond immediately and significantly to changes in their employees' savings incentives. In particular, we find that the change in employer capital contributions in response to an increase in their relative price within the 1999 reform was strongly related to the fraction of workers who were above the top tax threshold and were directly affected by the reform. We also find that employers adjusted their employees' overall savings portfolios by significant shifts into the more subsidized annuity accounts, with almost no leakage of overall savings.

Since employer contributions and defaults are extremely effective at increasing individuals' total level of savings, some governments are considering implementing policies that incentivize employer-based savings accounts and default contribution rates. Given the increasing reliance of individual retirement savings on employers' contributions, our findings are promising and encouraging preliminary evidence that they are set carefully and in accordance with workers' savings preferences.

However, there are other important aspects of firm responses that we did not address in this paper. For example, our results do not reveal whether firms' behavior is attributable to benevolence or to competition, and, since our analysis is strictly positive, we do not know how far firms' responses were from the optimal changes in their contribution rates. We believe that addressing these issues is a fruitful direction for future research, as they have potentially important implications for the optimal design of employer-based retirement savings policies.

References

Aaronson, S., and J. L. Coronado (2005): "Are Firms or Workers Behind the Shift Away from DB Pension Plans?," *FEDS paper*, (2005-17).

Bernheim, B. D. (1998): "Financial Illiteracy, Education, and Retirement Saving," in *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement*, ed. by O. Mitchell, and S. Schieber, pp. 38–68. Philadelphia: Univ. Pennsylvania Press.

Beshears, J., J. Choi, D. Laibson, B. Madrian, and B. Weller (2010): "Public Policy and Saving for Retirement: The "Autosave" Features of the Pension Protection Act of 2006," in *Better Living Through Economics*, ed. by J. Siegried, pp. 274–290.

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2009): "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States," in *Social security policy in a changing environment*, pp. 167–195. University of Chicago Press.

Bingley, P., N. D. Gupta, and P. J. Pedersen (2007): "Fiscal Implications of Reforms in Retirement Systems in Denmark," in *Social Security Programs and Retirement around the World: Fiscal Implications of Reform*, pp. 119–154. University of Chicago Press.

Brobeck, S. (1990): *U.S. Consumer Knowledge: Results of a Nationwide Test*, Washington, DC: Consum. Fed. Am.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009): "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 124(4), 1639–1674.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014): "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 129(3), 1141–1219.

Choi, J. J. (2015): "Contributions to Defined Contribution Pension Plans," *Annual Review of Financial Economics*, 7(1).

Choi, J. J., D. Laibson, and B. C. Madrian (2007): "Mental Accounting in Portfolio Choice: Evidence from A Flypaper Effect," Discussion paper, National Bureau of Economic Research.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2003): "Optimal Defaults," *American Economic Review*, pp. 180–185.

——— (2004): "For Better or for Worse: Default Effects and 401 (k) Savings Behavior," in *Perspectives on the Economics of Aging*, pp. 81–126. University of Chicago Press.

Cremer, H., P. De Donder, D. Maldonado, and P. Pestieau (2008): "Designing a Linear Pension Scheme with Forced Savings and Wage Heterogeneity," *International Tax and Public Finance*, 15(5), 547–562.

Cremer, H., and P. Pestieau (2011): "Myopia, Redistribution and Pensions," *European Economic Review*, 55(2), 165–175.

Dummann, K. (2008): "What Determines Supply and Demand for Occupational Pensions in Germany?," *Journal of Pension Economics and Finance*, 7(02), 131–156.

Fadlon, I., and D. Laibson (2015): "Endogenous Paternalism and Pseudo-Rationality," Unpublished manuscript.

Gelber, A. (2011): "How Do 401(k)s Affect Saving? Evidence from Changes in 401(k) Eligibility," *American Economic Journal: Economic Policy*, 3(4), 103–122.

Goda, G. S., and C. F. Manchester (2013): "Incorporating Employee Heterogeneity into Default Rules for Retirement Plan Selection," *Journal of Human Resources*, 48(1), 198–235.

Gustman, A. L., T. L. Steinmeier, and N. Tabatabai (2009): "Do Workers Know about Their Pension Plan Type? Comparing Workers' and Employers' Pension Information," *Overcoming the Savings Slump: How to Increase the Effectiveness of Financial Education and Saving Programs*, pp. 47–81.

Hastings, J. S., B. C. Madrian, and W. L. Skimmyhorn (2013): "Financial Literacy, Financial Education, and Economic Outcomes," *Annual Review of Economics*, 5, 347–73.

Hernæs, E., J. Piggott, T. Zhang, and S. Strøm (2011): "Occupational Pensions, Tenure, and Taxes," *Journal of Pension Economics and Finance*, 10(03), 435–456.

Hilgert, M. A., J. M. Hogarth, and S. G. Beverly (2003): "Household Financial Management: The Connection Between Knowledge and Behavior," *Fed. Res. Bull.*, 89, 309.

Lusardi, A., and O. S. Mitchell (2007): "Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth," *Journal of Monetary Economics*, 54(1), 205–224.

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

Mitchell, O. S. (1988): "Worker Knowledge of Pension Provisions," *Journal of Labor Economics*, 6(1), 21–39.

OECD (2009): "Pension Country Profile: Denmark," in *OECD Private Pensions Outlook 2008*. Paris: OECD Publishing.

Økonomi-og erhversministeriet, B., and Finansministeriet (2003): "Større valgfrihed I pensionsopsparingen," Commisioned report.

Oyer, P. (2005): "Salary or Benefits?," Discussion paper, National Bureau of Economic Research.

Phelps, C. (1973): "The Demand for Health Insurance: a Theoretical and Empirical Investigation," .

Rice, R. G. (1966): "Skill, Earnings, and the Growth of Wage Supplements," *American Economic Review*, pp. 583–593.

Roeder, K. (2014): "Optimal Taxes and Pensions with Myopic Agents," *Social Choice and Welfare*, 42(3), 597–618.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, pp. 34–55.

Online Appendix A: Verification of Employer Responses

Recall that contribution rates within a workplace, defined by all employees with the same 2-digit occupation code in the same firm, can vary since employers may set pension contributions at a finer level within the firm (i.e., higher-digit occupation codes), and since individuals have some ability to choose a different distribution of contributions between capital and annuity accounts than the default set by the employer. Our main analysis identifies employer default contribution rates using workplace-level medians.

To validate our measure of choice and to verify that our results are not driven by employees' responses within their employer-sponsored accounts, but rather by the employers themselves, we augment the analysis in the following ways:

1. **Analysis of Modes:** Instead of workplace-level medians we calculate workplace-level modes for identifying employer contribution rates. Panel A.1 of Appendix Figure 2 replicates Panel B of Figure 1, but with changes in workplace-level modes of contribution rates instead of medians, and shows that our results persist with this choice of the outcome variable. To further show that the patterns are very similar to those provided by medians, Panel A.2 of Appendix Figure 2 replicates Panel B of Figure 1 that uses medians, but for the same sample that is included in Panel A.1 of Appendix Figure 2. This sample includes only workplaces with unique mode values. Due to rounding issues and multiplicity of possible modes we are left with 59,956 workplace-year observations, compared to 84,764 workplace-year observations in the main analysis of medians, which is the reason for our choice of medians over modes.

2. **Deviation from Aggregates:** Panels B.1 and B.2 of Appendix Figure 2 plot the distribution of the distance between employee-level capital contribution rates and workplace-level aggregates, for modes and medians, respectively. The analysis verifies that most employee contributions bunch exactly at these aggregates, supporting our design and choice of measures for identifying employer behavior.

3. **Workplace Size:** Statistically, medians are more likely to accurately identify default contribution rates in large workplaces. Conceptually, the firm decision making process is more likely to be tailoring defaults to groups of employees of similar occupations in larger workplaces. Therefore, we study the sensitivity of our analysis to the size of workplaces that are included in our sample. Appendix Figure 3 plots changes in medians for the samples of workplaces with more than ten, twenty, and fifty employees. While the sample size substantially decreases, the patterns of the results are very similar.

4. **Workplaces with No Contributions to Annuity Accounts:** To focus on a sample of employees with less discretion over contributions to employer-sponsored accounts, we repeat the analysis for employees who could not have reallocated contributions across different types of accounts. In particular, Panel C of Appendix Figure 2 constrains the sample to workplaces whose median annuity contribution rate in the years prior to the reform (1996-1998) was zero. In addition, we include only workplaces whose median capital contribution rate in 1998 was economically meaningful (here we choose having rates larger than 1.5%), to focus on employers that had non-negligible potential to reduce capital contributions in response to the reform. These restrictions substantially reduce the sample size (to a total of 4,056 observations), but reveal the same patterns.

Online Appendix B: Distribution of Employee-Level Changes in Contribution Rates

In section 3.2 we compared employer responses to the reform to the responses of individuals within their private accounts by using workplace-level aggregates. To account for differential baseline contribution rates to private vs. employer-sponsored accounts, we analyze percent changes in the two types of accounts at the employee level, by plotting the distribution of employee-level year-to-year percent changes in capital contribution rates. In addition,

to conduct a within-employee analysis when we compare responses within private vs. employer-sponsored accounts, we balance the sample such that we keep only employees that had positive contributions to both type of accounts in a previous year.

In Appendix Figure 4 we plot histograms of individual-level percent changes in contributions to capital accounts for employer-sponsored capital accounts (Panel A) and private capital accounts (Panel B). In each panel we plot the distributions for the two years prior to the reform (so that we average the changes from 1996 to 1997 and from 1997 to 1998) and the distribution of changes following the reform – which captures changes from 1998 to 1999.

Comparing the before and after distributions, we see that in both types of accounts most of the effect of the reform is driven by exiting capital accounts altogether (the large differences in the mass points at -100 in both panels). In fact, in both employer-sponsored and private accounts, there was an increase of 23 pp in the fraction of employees who exited capital accounts.[22]

Recall from Section 3.2 that employer responses were much larger than individual responses in the aggregate. The evidence here, which results from analyzing the contributions of the special sample of individuals with positive lagged contributions to both accounts, suggests that at the individual level percent change responses in private accounts were similar to those in employer accounts. However, this result is not entirely surprising since the analysis here diverges from the ideal experiment that we mentioned is Section 3.2 for two main reasons. First, we analyze only a small share of individuals. While among those that were at the top tax bracket in 1999 55% had positive contributions to employer-sponsored accounts and 30% had positive contributions to private accounts, less than 13% had positive contributions to both. Specifically, these workers are likely to be the most attentive to their pension accounts

[22]Note, however, that while the responses to the reform within private accounts were composed of opting out of capital accounts by almost exclusively those who would otherwise have small increases or no changes in contributions, opting out of employer-sponsored accounts was also equally driven by workplaces who would have otherwise decreased their contributions (the decrease in Panel A from 10 pp to 3 pp in the mass point at -4% that includes changes larger or equal to -4% and smaller than 0%).

than a randomly selected individual who does not typically contribute to private pension accounts. Second, the sample includes affected individuals in workplaces that also had to cater to employees who were not affected by the reform. Therefore, changes in employer-sponsored accounts do not capture the full potential extent of employer responses. The average fraction of workers at the top income tax bracket in these workplaces was 74% in 1999, and given the linear response we found in Section 3.1, there is a potential attenuation of the employer responses by approximately 26% compared to the response of employers with all workers at the top income tax bracket.

Online Appendix C: Capital Contributions vs. Labor Income Responses

Since we are analyzing employer capital contributions as a fraction of labor income, the heterogeneous responses of firms by the fraction of workers above the top threshold could be due to differential changes in the numerator, employer capital contributions, or the denominator, taxable labor income. Appendix Figure 5 breaks down the responses in capital contribution rates plotted in Panel A of Figure 1 into annual earnings and capital contribution levels.

Using different workplace-level measures for annual labor earnings, Panel A, B, and C show no noticeable differential patterns over the years across workplaces with different shares of employees above the top tax threshold. In contrast, Panel D, which plots workplace-level medians of contribution levels, exhibits exactly the same patterns of Panel A of Figure 1, confirming that our results are driven by changes in employer contributions to pension accounts.

This analysis also allows us to test an implication of the standard model of employee compensation. Recall that the decrease in subsidies for capital contributions potentially led to an income effect, as it reduced the overall compensation of workers above the top tax threshold. However, Appendix Figure 5 as well as the time series of the log of labor income in Panel B of Appendix Figure 1, display no detectable change in compensation through in-

creased annual earnings, either in the aggregate or as a function of the share of employees above the top tax threshold.[23]

Online Appendix D: Inclusion of Workers Covered by Collective Agreements

Recall that our analysis aims at isolating workplaces in which default contribution rates are set by the employer. To this end, we excluded workers in the public sector or with blue-collar occupations (as they are likely covered by collective agreements), so that we included only workers in private firms with white-collar occupations. However, some white-collar jobs in the private sector are covered by collective bargaining.

In this appendix, we assess how inadvertently including workers covered by collective agreements may have affected our results. Inclusion of such workers can alter the results in two different ways. First, given that there are union representatives within the collective bargaining process, employer contributions for these workers may be more closely related to workers' preferences than employer contributions that are set exclusively by firms. Thus, employer capital contributions for them may respond more to the reform. In this case, inadvertently including workplaces that are covered by collective bargaining could increase the magnitude of the coefficients. On the other hand, in such workplaces, employer contributions are set for groups of occupation-level worker unions and employer associations, not at the finer occupation-firm level. Since we analyze the occupation-firm cell as the decision cell in our analysis, variation in the share of employees above the top tax threshold across employers but within the same worker union-employer association unit will not exhibit differential responses. This will tend to flatten the gradient of responses with respect to the share of affected employees and to attenuate our results.

[23]This is consistent with the required compensation of the loss incurred at the initial bundle of contributions (on the order of no more than a few thousands of DKr) being negligible relative to annual labor earnings (on the order of hundred thousands of DKr). Empirically, such small changes are hard to detect, and, conceptually, some degree of wage rigidity due to, e.g., re-negotiation costs, would render these small changes non-profitable.

In order to empirically quantify how unintentionally including workplaces covered by collective bargaining affects our estimates, we analyze the relationship between employer capital contributions and the fraction of workers above the top tax threshold for different groups that are either more or less likely to be covered by collective bargaining agreements as compared to workers with white-collar occupations in the private sector.

Specifically, in column (3) of Appendix Table 1 we restrict the analysis to employees in the public sector or those with blue-collar occupations, who are more likely to be covered by collective bargaining. The coefficient of interest – on the interaction of the fraction of workers above the top threshold with the indicator for year 1999 – is -0.35, which is an order of magnitude smaller than (and statistically different from) the corresponding coefficient of -2.18 in our main sample analysis of employees with white-collar occupations in the private sector (replicated in column (1) of Appendix Table 1). In column (2) of Appendix Table 1 we restrict the analysis to private-sector workplaces of white-collar occupations with a highly-educated workforce. In particular, we focus on workplaces in which more than twenty percent of workers had more than sixteen years of education, which are even less likely to be covered by collective bargaining (as compared to all private-sector workplaces of white-collar occupations). The point estimate of the coefficient of interest in this case is large, -3.68 pp, and is significantly different from the coefficient for the remainder sample of private-sector white-collar occupations.

Put together, the results suggest that inadvertently including some workers who are covered by collective bargaining in our main specification may have attenuated our results.

Online Appendix E: Occupation Codes

In this appendix, we list 2-digit occupation codes used by Statistics Denmark. In our analysis, white-collar occupations are defined as occupations with codes whose first digit is between 1 and 5.

Occupation Codes:

1 Leadership at the top level of companies, organizations and the public sector

11 Legislative work and leadership in public administration and interest groups

12 Senior management of the company

2 Work that requires knowledge at the highest level in the area

21 Work within the non-biological branches of science and computer science, statistics, architecture and engineering sciences

22 Work in medicine, pharmacy, and biological branches of science, as well as midwives, general nursing work etc.

23 Teaching in primary schools, vocational schools, colleges, universities, and research organizations

24 Work within the social sciences and humanities

3 Work that requires Intermediate knowledge

31 Technicians in non-biological topics

32 Technicians and other work in biological topics

33 Caring and educational work

34 Work in sales, finance, business, administration, etc.

4 Office jobs

41 Internal office work

42 Office work with customer service

5 Retail sales, service and care work

51 Service and care work

52 Retail and models

6 Work in agriculture, horticulture, forestry, hunting and fishing

61 Work in agriculture, horticulture, forestry, hunting and fishing

7 Craft and related trades workers

71 Work in mining and construction

72 Metal and engineering work

73 Precision craftsmanship, graphic work, etc.

74 Other craft and work in related trades

8 Machine operating and work in transportation and civil engineering

81 Stationary plant

82 Operation of industrial machinery

83 Transportation

9 Other work

91 Cleaning and renovation work, messenger and security services, telephone canvassing

92 Assisting in agriculture, horticulture, fisheries and forestry

93 Manual work in the construction sector, manufacturing and transportation

FIGURE 1
Employer Contributions to Capital Pension Accounts
by the Share of Workers above the Top Tax Threshold

(a) Contribution Rates

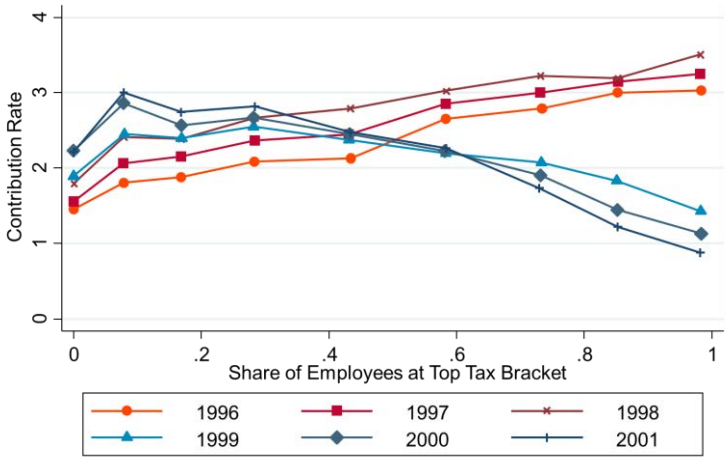

(b) Changes in Contribution Rates

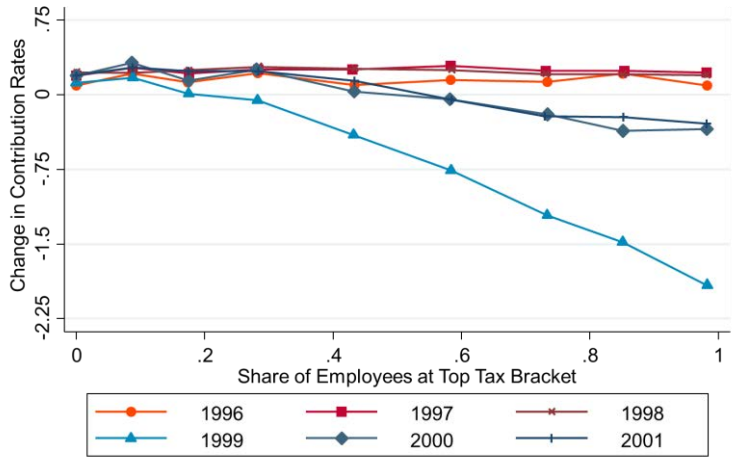

Notes: These figures plot employers' contributions to capital pension accounts as a function of the share of their employees whose earnings were above the top labor-income tax threshold, for years 1996-2001. Panel A plots employer capital contribution rates (as a fraction of labor income), and Panel B plots changes in employer capital contribution rates from the previous year. The observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and then plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

FIGURE 2
Employer vs. Individual Contributions to Capital Pension Accounts of Workers above the Top Tax Threshold

(a) Changes in Employer Contributions by the Share of Workers above the Top Tax Threshold

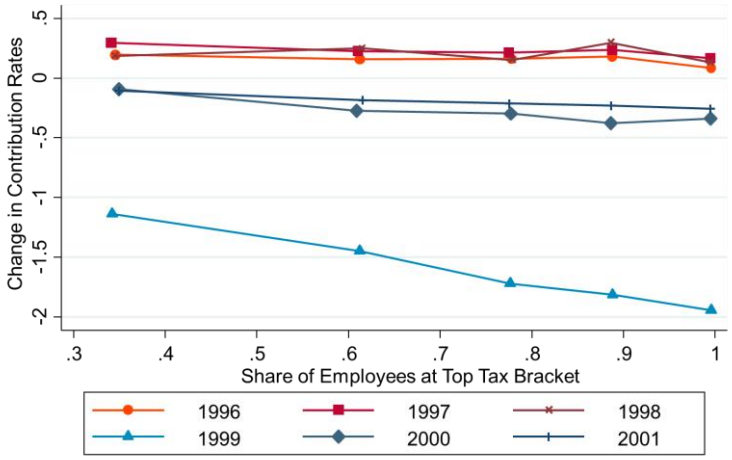

(b) Changes in Individual Contributions to Private Accounts by the Share of Workers above the Top Tax Threshold

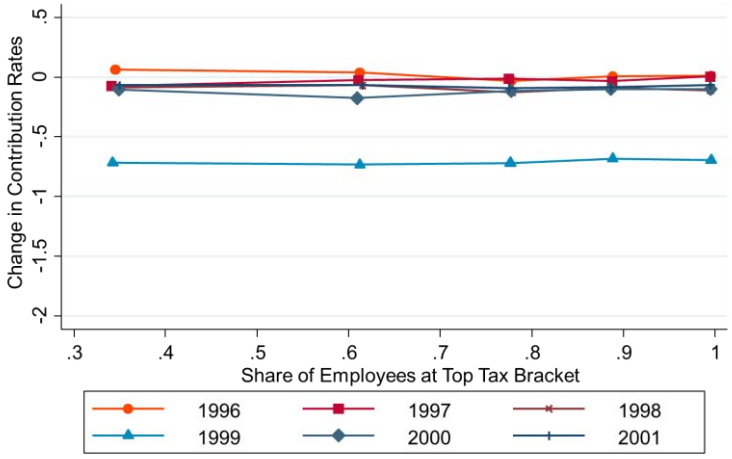

Notes: These figures plot changes in workplace-level contribution rates to capital pension accounts for employees with labor income at the top tax bracket as a function of a workplace's share of employees above the top tax threshold, for years 1996-2001. Panel A plots changes in median capital contribution rates to employer-sponsored (401(k)-like) accounts, and Panel B plots changes in average capital contribution rates to private (IRA-like) accounts. The observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals, workplaces with less than five employees, and employees with earnings below the top income threshold.

FIGURE 3
Employer Contributions to Annuity Pension Accounts

(a) Changes in Contribution Rates by Year

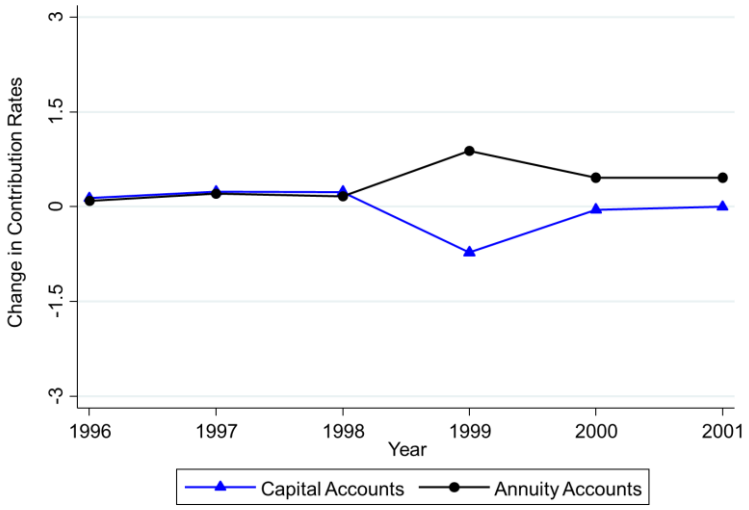

(b) Changes in Annuity Contributions by the Share of Workers above the Top Tax Threshold

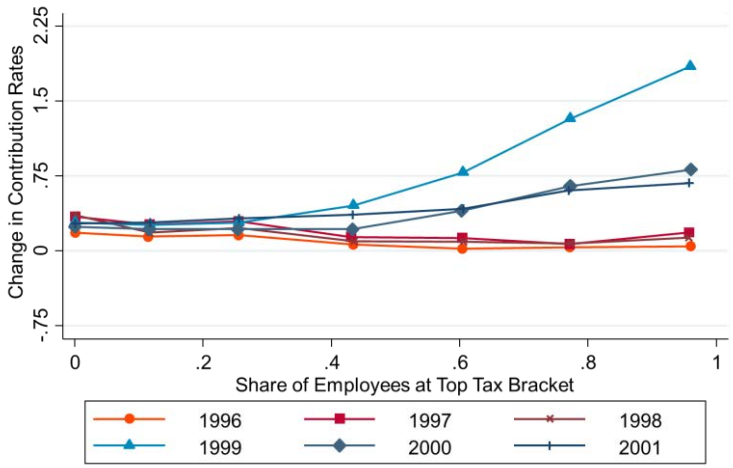

Notes: These figures plot changes in employers' contributions to annuity pension accounts, for years 1996-2001. Panel A plots changes in employer contributions by year, comparing annuity contributions (black line and circles) to capital contributions (blue line and triangles). Panel B plots changes in employers' contribution rates to annuity pension accounts as a function of the share of their employees whose earnings were above the top labor-income tax threshold, for years 1996-2001. The observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

APPENDIX FIGURE 1
Employer Capital Contributions and Labor Income by Year

(a) Workplace-Level Medians of Capital Contribution Rates

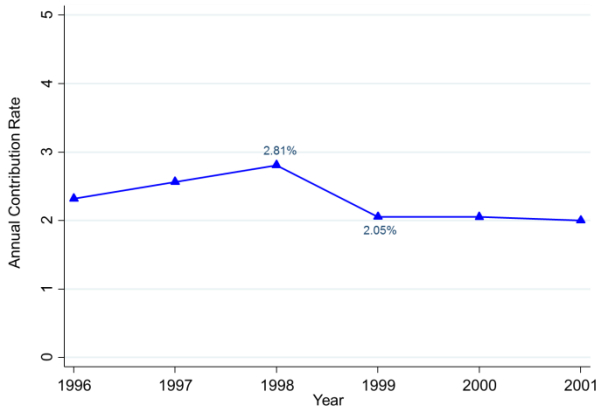

(b) Workplace-Level Medians of the Log of Labor Income

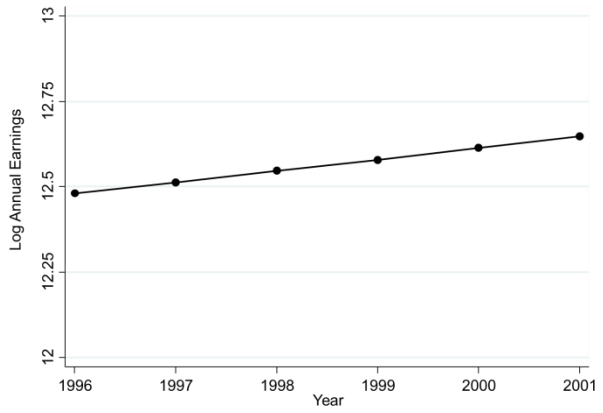

Notes: These figures plot means of different workplace-level outcomes by year, for years 1996-2001. The outcome in Panel A is the median contribution rate to employer-sponsored capital accounts, and the outcome in Panel B is the median of the log of annual labor income. A workplace is defined as the group of all employees with the same 2-digit occupation code in the same firm. The sample includes employees in private-sector firms with white-collar occupations, and excludes self-employed individuals and employees in workplaces with less than five employees.

APPENDIX FIGURE 2
Robustness of Employer Responses in Capital Contributions

(a) Changes in Contribution Rates by the Share of Workers above the Top Tax Threshold

(1) Modes (2) Medians for Same Sample

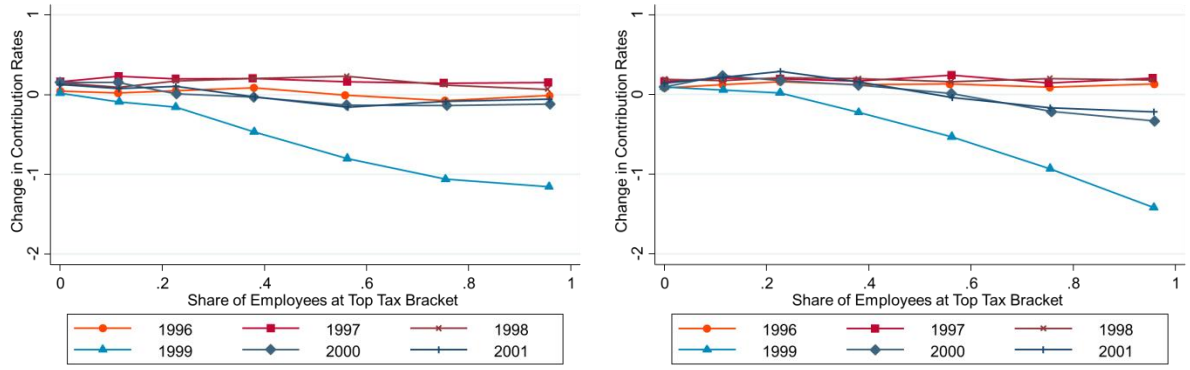

(b) Distance of Individual-Level Measures from Workplace-Level Aggregates

(1) Modes (2) Medians

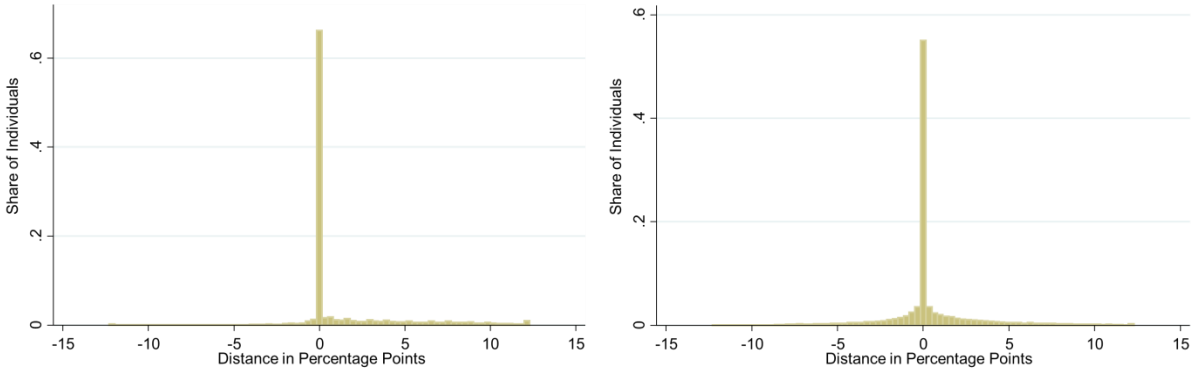

(c) Workplaces with No Annuity Contributions

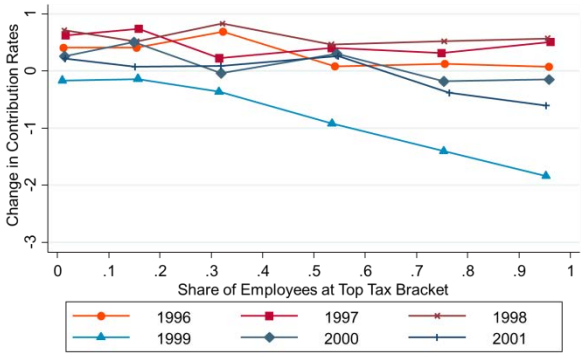

Notes: The sample for these figures includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees. Panel A.1 plots changes in workplace-level modes of contribution rates to capital accounts as a function of a workplace's share of employees above the top tax threshold, for years 1996-2001, including only workplaces with unique mode values. Panel A.2 plots workplace-level medians, including only the sample that is included in Panel A.1. These figures are plotted in the same way as the figure in Panel B of Figure 1. Panels B.1 and B.2 plot the distribution of the distance between employee-level capital contribution rates and workplace-level modes and medians, respectively. Panel C replicates Panel A.2 for all private-sector firms and white-collar occupations but constrains the sample to workplaces whose median annuity contribution rate in the years prior to the reform (1996-1998) was zero and whose median capital contribution rate in 1998 was larger than 1.5% (so that they had non-negligible potential to reduce capital contributions in response to the reform), with a total of 4,056 observations.

APPENDIX FIGURE 3
Employer Capital Contributions Responses by Size of the Workforce

(a) Workforce Larger than 10 Employees

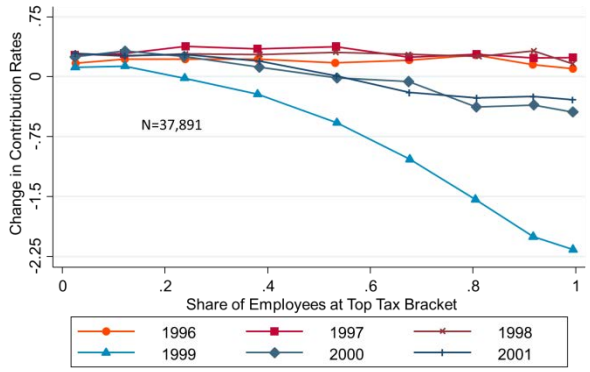

(b) Workforce Larger than 20 Employees

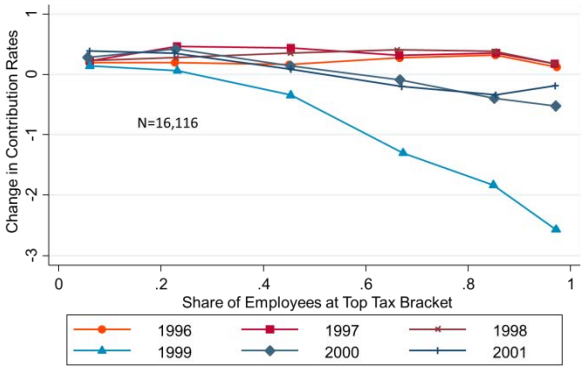

(c) Workforce Larger than 50 Employees

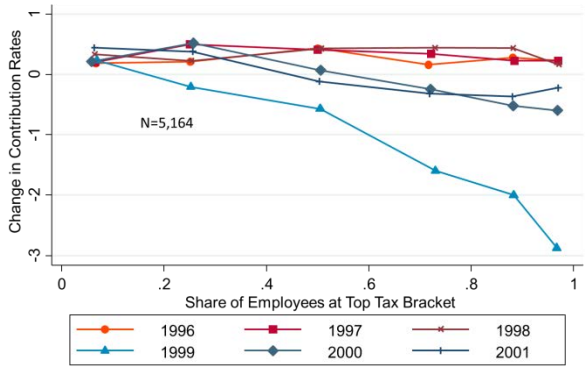

Notes: These figures plot changes in employers' contribution rates to capital pension accounts as a function of the share of their employees whose earnings placed them above the top labor-income tax threshold, for years 1996-2001. Panel A excludes workplaces with less than ten employees, Panel B excludes workplaces with less than twenty employees, and Panel C excludes workplaces with less than fifty employees. In each panel we indicate the number of included observations, denoted by N. The observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations and excludes self-employed individuals.

APPENDIX FIGURE 4
Changes in Employer vs. Individual Contributions to Capital Pension Accounts of Workers above the Top Tax Threshold

(a) Distribution of Employee-Level Changes in Contributions to Employer-Sponsored Accounts

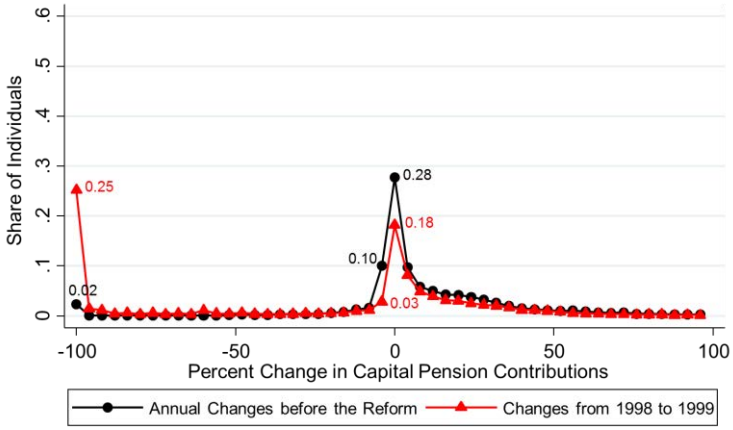

(b) Distribution of Employee-Level Changes in Contributions to Private Accounts

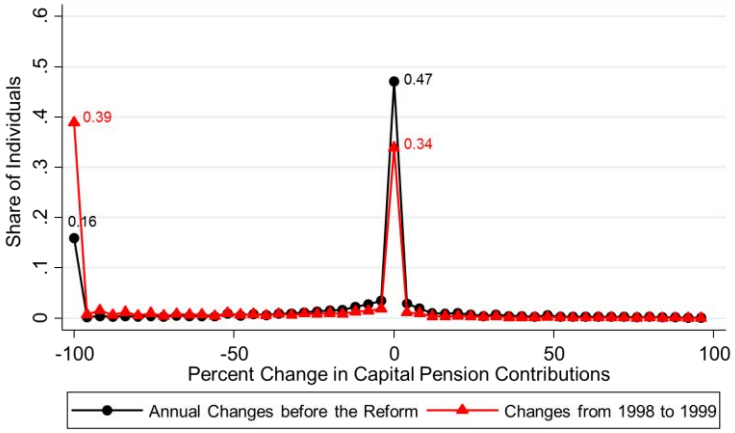

Notes: These figures plot the distribution of changes in employee-level capital pension contributions in percents relative to lagged contributions for individuals above the top tax cutoff for the years before the reform (black lines and circles) and from 1998 to 1999 (red lines and triangles). Panel A plots changes in contributions to employer-sponsored (401(k)-like) accounts, and Panel B plots changes in contributions to private (IRA-like) accounts. The figures include only individuals with positive lagged contributions to both types of accounts. Each point represents the floor of bins of 4% width, so that the point at 0% represents individuals with changes in the range [0%,4%). The curves for the distribution of annual changes in the years prior to the reform (in black lines and circles) include changes from 1996 to 1997 and 1997 to 1998. The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals, workplaces with less than five employees, and employees with earnings below the top income threshold.

APPENDIX FIGURE 5
Employer Capital Contributions and Labor Income
by the Share of Workers above the Top Tax Threshold

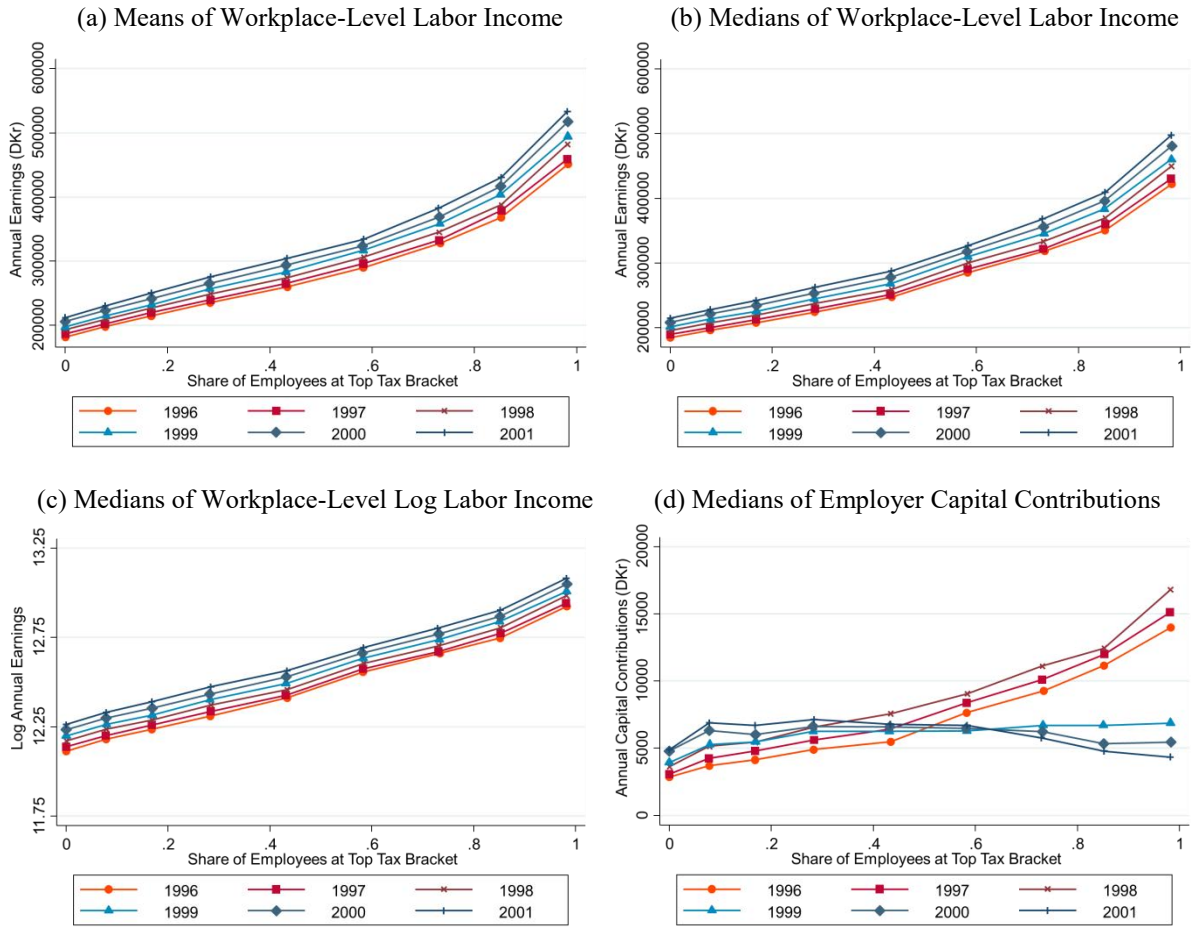

Notes: These figures plot different workplace-level outcomes as a function of the share of employees whose earnings placed them above the top labor-income tax threshold, for years 1996-2001. Panel A plots means of annual labor income, Panel B plots medians of annual labor income, Panel C plots medians of log annual labor income, and Panel D plots median employer contributions to capital pension accounts. The observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm. We plot these figures by dividing the sample into equal-sized groups according to the share of employees above the top tax threshold, and plotting for each group the mean outcome (on the y-axis) against the mean share of employees above the top tax threshold (on the x-axis). The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees.

TABLE 1
Summary Statistics of Analysis Sample

	Mean (1)	Median (2)	SD (3)
Individual-Level Variables:			
Labor Earnings (DKr)	285,740	253,893	214,500
Pension Contributions before the Reform			
Employer-Sponsored Accounts			
Capital Contributions (DKr)	11,665	6,679	12,490
Capital Contribution Rate	3.6%	2.8%	3.8%
Fraction with Capital Contributions	66%		
Annuity Contributions (DKr)	12,604	5,047	17,830
Annuity Contribution Rate	3.7%	1.8%	4.5%
Fraction with Annuity Contributions	77%		
Private Accounts			
Capital Contributions (DKr)	3,125	0	7,627
Capital Contribution Rate	1.1%	0%	2.6%
Fraction with Capital Contributions	28%		
Annuity Contributions (DKr)	1,470	0	5,859
Annuity Contribution Rate	0.4%	0%	1.7%
Fraction with Annuity Contributions	15%		
Workplace-Level Variables:			
Fraction of Employees above Top Tax Threshold	49%	50%	35%
Number of Employees	22	9	97
Number of Worker-Year Observations		2,020,705	
Number of Workplace-Year Observations		84,764	
Number of Workplaces		26,775	

Notes: This table presents means, medians, and standard deviations of key variables in our analysis sample of white-collar workers in private-sector firms from 1996 to 2001. The classification of white-collar occupations is described in detail in Appendix E. All monetary values are reported in nominal Danish Kroner (DKr), where the exchange rate during this time period was approximately DKr 6.5 per US $1. Labor income is calculated as total pre-tax wage earnings plus employer pension contributions. The values reported in the table for pension contributions before the reform are based on data from 1998. Pension contribution levels are winsorized at their 99th percentile.

TABLE 2
Changes in Employer Contribution Rates to Pension Accounts
by the Share of Workers above the Top Tax Threshold

Panel A					
Dependent Variable:		Δ Capital Contributions			
		(1)	(2)	(3)	(4)
Fraction of Employees Above Top Tax Threshold (Baseline Year 1998)		-0.052*	-0.113*	0.122	0.132
		(0.032)	(0.0670)	(0.137)	(0.145)
Fraction of Employees Above Top Tax Threshold Interacted with:					
	Year 1996	0.035	0.101	0.202	0.273
		(0.046)	(0.098)	(0.139)	(0.167)
	Year 1997	0.061	0.075	0.071	0.096
		(0.046)	(0.096)	(0.126)	(0.147)
	Year 1999	-2.126***	-2.167***	-2.181***	-2.182***
		(0.061)	(0.129)	(0.176)	(0.203)
	Year 2000	-0.606***	-0.569***	-0.556***	-0.593***
		(0.048)	(0.104)	(0.149)	(0.173)
	Year 2001	-0.558***	-0.598***	-0.590***	-0.553***
		(0.046)	(0.101)	(0.152)	(0.176)
Year Fixed Effects		X	X	X	X
Income and Workforce Size Controls			X	X	X
2-Digit Occupation-Firm Fixed Effects				X	X
2-Digit Occupation-Year Fixed Effects					X
Number of Observations		84,764	60,643	60,643	60,643
Number of Clusters		26,775	20,642	20,642	20,642

Panel B	
Dependent Variable:	Δ Annuity Contributions
Fraction of Employees Above Top Tax Threshold Interacted with Year 1999	1.962***
	(0.195)

Panel C	
Dependent Variable:	Δ Overall Employer Contributions
Fraction of Employees Above Top Tax Threshold Interacted with Year 1999	-0.220
	(0.217)

Notes: This table reports estimates of employers' responses to the reform as a function of the share of their employees whose earnings were above the top labor-income tax threshold (equation (1)). In Panel A, the outcome variable is the change in employer capital contribution rates from the previous year. We regress this outcome on the fraction of workers above the top tax threshold, year fixed effects, the fraction of workers above the top tax threshold interacted with year fixed effects, and different sets of controls as indicated in the table. The baseline year is 1998, so that the coefficient on the fraction of employees above the top tax threshold refers to that year. The coefficient on the fraction of employees above the top tax threshold interacted with other year indicators estimates this relationship relative to the relationship of the baseline year. Income controls include a fifth-order polynomial of the mean workplace-level labor income, separately for workers above and below the top tax threshold, as well as their interactions with year indicators. Workforce size controls include the number of workers in a workplace and its square, as well as their interactions with year indicators. In Panels B and C, we replicate the specification with the full set of controls in column (4) of Panel A, but where the outcome variables are changes in employer annuity contribution rates and changes in overall employer contribution rates to both capital and annuity accounts, respectively. In all regressions the observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. The sample includes private-sector firms and white-collar occupations, and excludes self-employed individuals and workplaces with less than five employees. Standard errors are clustered at the workplace level. Coefficients are multiplied by 100 so that they are converted to percentage point units.
*** p<0.01, ** p<0.05, * p<0.1

APPENDIX TABLE 1
Changes in Employer Contribution Rates to Capital Pension Accounts for Different Samples

Dependent Variable:	Δ Capital Contributions		
Sample:	White-Collar in Private Sector		Blue-Collar or in Public Sector
	All	Highly-Educated	
	(1)	(2)	(3)
Fraction of Employees Above Top Tax Threshold (Baseline Year 1998)	0.132	0.303	0.138**
	(0.145)	(0.611)	(0.074)
Fraction of Employees Above Top Tax Threshold Interacted With:			
Year 1996	0.273	-0.159	-0.093
	(0.167)	(0.752)	(0.081)
Year 1997	0.096	-0.234	0.063
	(0.147)	(0.675)	(0.076)
Year 1999	-2.182***	-3.681***	-0.347***
	(0.203)	(0.861)	(0.084)
Year 2000	-0.593***	-1.467**	-0.129*
	(0.173)	(0.747)	(0.077)
Year 2001	-0.553***	-0.178	-0.080
	(0.176)	(0.915)	(0.078)
Year Fixed Effects	X	X	X
Income and Workforce Size Controls	X	X	X
2-Digit Occupation-Firm Fixed Effects	X	X	X
2-Digit Occupation-Year Fixed Effects	X	X	X
Number of Observations	60,643	7,989	78,533
Number of Clusters	20,642	3,742	25,932

Notes: This table reports estimates of employers' responses to the reform, for different samples, as a function of the share of their employees whose earnings were above the top labor-income tax threshold. The outcome variable is the change in employer capital contribution rates from the previous year. We regress this outcome on the fraction of workers above the top tax threshold, year fixed effects, the fraction of workers above the top tax threshold interacted with year fixed effects, and a set of controls as indicated in the table. The baseline year is 1998, so that the coefficient on the fraction of employees above the top tax threshold refers to that year. The coefficient on the fraction of employees above the top tax threshold interacted with other year indicators estimates this relationship relative to the relationship of the baseline year. Income controls include a fifth-order polynomial of the mean workplace-level labor income, separately for workers above and below the top tax threshold, as well as their interactions with year indicators. Workforce size controls include the number of workers in a workplace and its square, as well as their interactions with year indicators. Column (1) replicates exactly column (4) of Table 2, which restricts the sample to private-sector firms and white-collar occupations. Column (2) repeats the analysis of column (1) for private-sector firms and white-collar occupations, but restricts the sample to workplaces with a highly-educated workforce. Specifically, we focus on workplaces in which more than twenty percent of workers had more than sixteen years of education. Column (3) restricts the estimation of employer responses to employees in the public sector or those with blue-collar occupations. All the regressions exclude self-employed individuals and workplaces with less than five employees. The observation units are workplaces, defined by all employees with the same 2-digit occupation code in the same firm, where employer contribution rates are calculated as the median annual contribution rate within each workplace in a given year. Standard errors are clustered at the workplace level. Coefficients are multiplied by 100 so that they are converted to percentage point units.
*** p<0.01, ** p<0.05, * p<0.1