

Salamanca, Nicolás; de Grip, Andries; Sleijpen, Olaf

Working Paper

How People React to Pension Risk

IZA Discussion Papers, No. 13077

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Salamanca, Nicolás; de Grip, Andries; Sleijpen, Olaf (2020) : How People React to Pension Risk, IZA Discussion Papers, No. 13077, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/216389>

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.

DISCUSSION PAPER SERIES

IZA DP No. 13077

How People React to Pension Risk

Nicolás Salamanca
Andries de Grip
Olaf Sleijpen

MARCH 2020

DISCUSSION PAPER SERIES

IZA DP No. 13077

How People React to Pension Risk

Nicolás Salamanca

University of Melbourne, ARC Centre of Excellence for Children and Families over the Life Course and IZA

Andries de Grip

ROA, IZA and Netspar

Olaf Sleijpen

Maastricht University and De Nederlandsche Bank

MARCH 2020

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

How People React to Pension Risk*

We show that people exposed to greater pension risk are less likely to invest in risky assets. We exploit a reform that links people's future pension benefits to their pension funds' funding ratio—a measure of the fund's financial health—making funding ratios a fund-specific measure of pension risk. The effect of pension risk is stronger for people who are better informed about their pensions, for retirees and pension-age non-retirees, and for wealthier people. The funding ratio does not affect investments in a pre-reform period, nor does it affect bequest intentions, (expected) retirement, or the motivations for saving.

JEL Classification: D14, J22

Keywords: individual portfolio choice, background risk, retirement planning, pension reform, The Netherlands

Corresponding author:

Nicolás Salamanca
Melbourne Institute: Applied Economic & Social Research
University of Melbourne
111 Barry St.
Parkville, VIC 3010
Australia
E-mail: n.salamanca@unimelb.edu.au

* We thank Alexandra de Gendre, Jan Feld, Dan Hamermesh, Erzo Luttmer, Raymond Montizaan, Joan Muysken, Paul Smeets, Mukul Tyagi, Ulf Zölitz, Maria Zumbühl, and participants at several seminars, workshops, and conferences for helpful suggestions. We also owe special thanks to Albertine Palsma (DNB) for her help in making data available to us. In this paper use is made of data of the DNB Household Survey and restricted administrative data on Dutch pension funds. This research was partly supported by the Australian Research Council Centre of Excellence for Children and Families over the Life Course (project number CE140100027). The Centre is administered by the Institute for Social Science Research at The University of Queensland with nodes at The University of Western Australia, The University of Melbourne, and The University of Sydney.

1. Introduction

People's pensions are an important source of retirement income. Although the risk associated with pensions varies widely depending on the pension scheme, all pension plans expose their participants' pension savings to some market risk. This is clearest in defined contribution plans, where people divert part of their savings into a long-term investment account.¹ Yet even defined benefit plans, where participants are guaranteed an annuity based on their wage history when they retire, also expose their participants to pension risk. Defined benefit plans traditionally take most of the market risk upon themselves since the terms of the annuities they offer are independent of market performance. However, since population aging puts pressure on the funding of traditional defined benefit plans, many of them have now transitioned to hybrid plans.² Hybrid defined benefit plans shift some market risk to their participants by conditioning pension benefits and contribution rates on the performance of the economy. This approach makes benefits conditional on the pension fund's financial position and is therefore linked to the performance of financial markets (Ponds and Van Riel 2009). The additional market risk is likely to affect the behavior and welfare of pension fund participants.

In this paper we quantify the effects of this pension risk on people's investment portfolios. To do this we exploit a 2007 pension reform in The Netherlands that suddenly exposed people's defined benefit pensions to stock market risk. Before the reform, retirement in The Netherlands was typically organized in standard defined benefit schemes. The reform linked pension funds' funding ratio—the ratio of the fund's assets to its liabilities—to the benefits people receive and the pension contributions they make. This effectively exposed people to market risk through their pensions. It also transformed each pension fund's funding ratio into a time-varying, fund-specific measure of the pension risk borne by each fund member. To quantify the effects of pension risk, we link data on pension fund performance from the Dutch National Bank (DNB) to the DNB Household Survey (DHS), a representative survey of Dutch households that includes detailed information on their savings and investments. We then estimate the effect of pension risk on people's investment behavior using the within-pension fund and within-year variation in pension risk while controlling for a rich set of observable characteristics. Since Dutch pension funds are organized by industry, this

¹ These pension plans include 401(k)s in the United States, private personal pension schemes in the United Kingdom, and mandatory superannuation funds in Australia.

² Defined benefit and hybrid defined benefit schemes are currently in place in several OECD and non-OECD countries and are the predominant retirement scheme in countries like Germany, The Netherlands, Luxembourg, and Switzerland. In the United States and Canada, 56 percent and 97.5 percent of occupational plans are (hybrid) defined benefit plans (OECD Pension Markets in Focus, 2016).

approach addresses the main concern of selection of workers into occupations or pension funds with different risk profiles. Our empirical strategy implies that the only remaining source of endogeneity in our findings would need to i) be unaccounted for by our extensive set of controls, ii) be time-varying within each industry or pension fund, iii) affect people in the same way pension risk would, and iv) have an effect large enough to drive our results. These types of confounders are very difficult to imagine.

Our results show that people who participate in pension funds with a one standard deviation higher funding ratio—and therefore have safer pensions—are 5.6 percentage points more likely to invest in risky assets and invest 2.7 percentage points more of their wealth in these assets, mostly taking this amount from their cash savings. This effect cannot be explained by time-varying characteristics such as people’s labor market status, wealth, income, or family composition, nor can it be explained by aggregate market movements or time-invariant characteristics of the pension funds themselves, such as systematic sorting of people into occupations (and thus pension funds). Consistent with these effects truly identifying a change in people’s investment behavior, we show that the impact of pension risk is stronger for people who are *a priori* more likely to exhibit a behavioral response, specifically: i) those who reported receiving an overview of their pension status in the last year, ii) those who report being aware of their pension fund’s indexation policy, and iii) retirees, whose main source of income is pension benefits. Moreover, the impact of funding ratios on people’s investment behavior only occurs in the postreform period. Finally, we provide evidence that the effect of pension risk is chiefly driven by a “pure risk” component and is inconsistent with a first-order “wealth” effect (i.e., by the implied drop in pension wealth).

Our findings contribute to the literature on background risk and portfolio choice by providing the first quasi-experimental evidence of the effects of pension risk using microlevel data. Most empirical evidence in this field comes from calibrated lifecycle models and does not use quasi-experimental identification. Heaton and Lucas (2000) provide an excellent review of this literature and tackle the issue of labor and housing background risk and portfolio choice by calibrating optimal individual decisions under different market frictions in the presence of tradable and non-tradable risk. Using a similar approach, Gomes and Michaelides (2005) provide additional evidence on the role of background risk in explaining stock market investment and participation rates by modeling fixed investment costs and

heterogeneity in people’s risk aversion.³ Our paper complements their findings by directly estimating the effects of background risk—which is only assumed and calibrated in these models—using quasi-experimental variation and individual-level data. Our approach is more similar to the studies that estimate the effects of personal and household characteristics on their investment behavior. In this literature, Curcuru et al. (2004) address pension risk by using cross-sectional data to show that participation in a defined benefit plan is uncorrelated to stock ownership. Guiso et al. (1996), Rosen and Wu (2004), Edwards (2008, 2010), Cardak and Wilkins (2009), and Angerer and Lam (2009) show strong correlations between portfolio choice and background risk from sources such as labor and health. Yet none of these studies can directly measure time-varying proxies of background pension risk, nor do they exploit quasi-experimental sources of variation to identify their estimates.

Conceptually, the most similar study to ours is Fagereng et al. (2016), which isolates the effect of background wage risk on people’s investment behavior by instrumenting wage risk with firm-level profit variability. They find a large marginal impact of background wage risk on people’s investment in risky assets yet argue that wage risk is economically unimportant for people’s investment behavior because firms provide substantial insurance on that risk. We similarly provide well-identified evidence of the negative effect of pension risk—another importance source of background risk—on people’s investment behavior. Our setting allows us to provide further evidence on the identification and behavioral interpretation of this effect via heterogeneous effects models on information and salience and to exploit a pre-reform period for placebo analyses. We also show that the distributional effects of pension risk are quite different from those of wage risk: they are driven disproportionately by wealthy people, retirees, and people at retirement age. Finally, we also provide evidence that pension risk does not affect people’s intended or actual labor supply, nor does it seem to affect their motivations for saving.

Finally, our estimates contribute much-needed empirical evidence to the theoretical literature on how people’s risk-taking behavior changes with background risk. This literature predicts that people exposed to an “unfair” background risk will find other risky assets less attractive and, based on these predictions, explores various fundamental aspects of people’s decision-making under risk and uncertainty (e.g., Pratt and Zeckhauser 1987; Kimball 1993; Gollier and Pratt 1996). Pension risk in our setting resembles in many ways this unfair background risk, and our results give unique empirical support for the predictions in this

³ See also Viceira (2001) Gomes and Michaelides (2003), Cocco et al. (2005) and French (2005) for further evidence using lifecycle models that incorporate background risk. Curcuru et al. (2004) provide further references on these models.

literature. We provide a theory-motivated exploration of the nature of pension risk and show that the effect of pension risk has a first-order pure risk component and negligible wealth effect component.

Overall, our findings suggest that background pension risk matters for people's investment behavior, and that pension reforms that increase exposure to this risk can have important distributional effects. The distributional aspects likely matter for welfare since, like uncertainty about the timing of retirement (Caliendo et al. 2016), pension risk can meaningfully decrease people's welfare and is more directly felt by people later in life, when adjusting labor supply is costly and income shocks can have stronger impacts. The disproportionate behavioral effects of pension risk on retirement-age people as well as the distributional implications of these effects across socioeconomic status should be considered when introducing new retirement schemes and considering alternatives for pension reform.

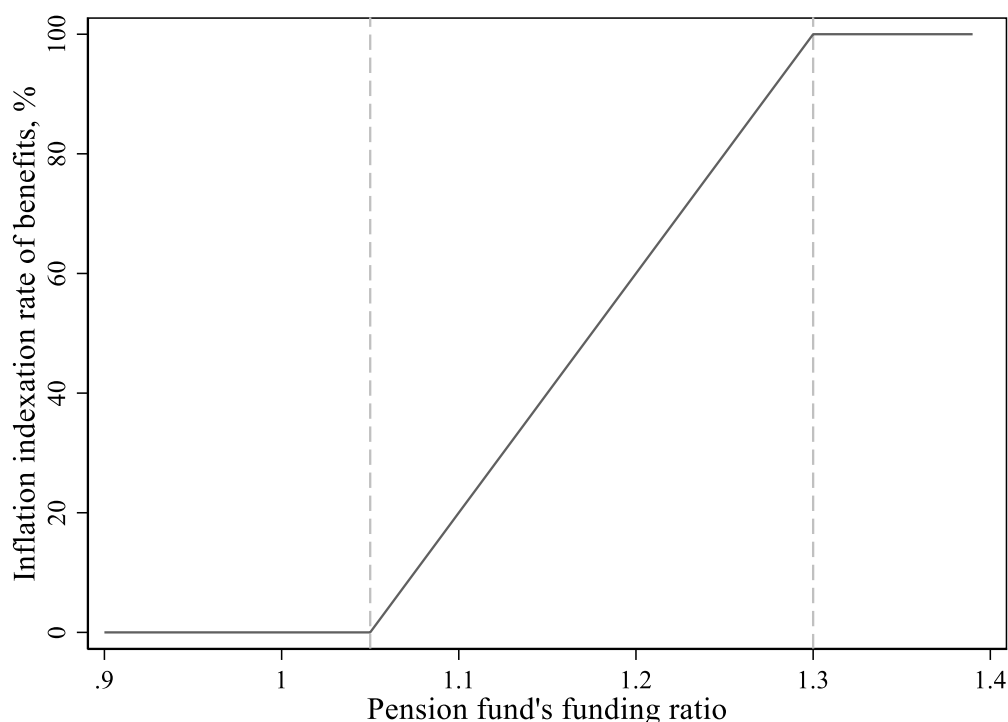
2. The Dutch Pension System

The pension system of The Netherlands is an ideal setting for studying people's reactions to background risk for three reasons. First, together with state-provided pensions that act as social security, firm-based retirement plans are the main pillar of retirement savings for the Dutch working population. In fact, the capitalization of Dutch occupational pension funds is among the world's highest (146 percent of gross domestic product in 2011, according to DNB statistics). Thus, pension fund savings are a crucial asset for most Dutch people.

Second, all employees in The Netherlands are obligated to participate in their employer's occupational pension plan, which is often shared by the entire industry sector. Each employer has a preassigned pension fund for all its employees, so even if employees are dissatisfied with their pension funds, they cannot change them without changing jobs. This contrasts with other settings in which employees are obligated to contribute to a pension fund yet are free to choose which one. The benefit of the more rigid system in The Netherlands is a high employee participation rate in Dutch pension funds from an international perspective, which allows employees to receive relatively generous pensions at retirement (OECD, 2011).

Third, in 2007 the Dutch government introduced a reform that effectively exposed all defined benefit pension holders to market risk through their pensions. Prior to the reform, employers and unions had been actively negotiating the terms of traditional defined benefit pension schemes to adapt them to the population aging (Sleijpen 2009). This process was catalyzed in the early 2000s by the dot-com crisis, which eroded the balance sheets of most pension funds, and culminated in the 2007 pension reform.

Figure 1. Typical indexation policy ladder for a Dutch pension fund after 2007



This figure illustrates a typical policy ladder regulating the inflation indexation policy of benefits for pension funds after 2007. In this example, the minimum required funding ratio to begin inflation indexation is 1.05, the funding ratio permitting full indexation is 1.3, and the taper rate for indexation is linear.

The key part of the reform was the introduction of legal requirements for pension funds to cut the inflation indexation of their benefits, nominally cut benefits, and/or increase mandatory contributions if the fund's financial health was compromised. Dutch authorities measure a fund's financial health using its *funding ratio*: the ratio of the market value of the fund's financial assets to the market value of its liabilities. A ratio of one indicates that the fund has just enough assets to cover its liabilities. Ratios higher than one indicate financially healthier funds and ratios lower than one indicate that the fund is in trouble. The 2007 reform requires pension funds to decrease retiree benefits and increase member contributions if i) their funding ratio drops below a fund-specific threshold and does not recover after three years, and ii) other recovery measures cannot be used as a last resort.⁴

⁴ Fund-specific thresholds are calculated by the Dutch Central Bank based on the fund's asset and liability mix and follow a general value-at-risk principle. Due to the financial crisis in 2008, the mandated recovery period was temporarily extended to five years. Nevertheless, several Dutch pension funds have had to reduce their benefits since then. The reform also introduced market valuation of assets and liabilities and determined that each pension fund should hold solvency reserves that depend on the riskiness of the pension fund's asset mix.

The link between funding ratios and the indexation and contribution policies of the pension funds are formalized by “policy ladders” (Ponds and Van Riel 2009). A typical indexation policy ladder is presented in Figure 1: it has no indexation below a funding ratio of 1.05, full indexation above 1.3, and partial indexation of benefits in between. Most pension funds would have similar policy ladders with slightly different thresholds for no benefit indexation. If the funding ratio remained over the 1.3 threshold for an extended period, then indexation that was not previously provided could be restored. If the funding ratio continued to be over 1.3, contribution rates could eventually be reduced, although this is an extremely unlikely event.

Conditional indexation of benefits and contributions to pension funds’ funding ratios effectively transferred some of the market risk of pension plans to the plans’ participants after 2007. Importantly, the rules that define benefit rates after retirement for these pension plans remain largely unchanged other than by their conditional indexation. This implies that any market risk taken by pension funds does not result in larger pensions for their participants; more pension risk is therefore largely unrewarded. Additionally, since participants do not know their fund’s precise asset mix, they are unable to hedge this risk by taking contrary investments. Together with participants’ inability to change their pension funds, this results in Dutch people being exposed to a background risk closely related to Gollier and Pratt’s (1996) “unfair” background risk—that which is undiversifiable and unrewarded—after 2007. Moreover, as we show in Section 3.2, the funding ratio of each pension fund is a good *inverse* measure of this background risk.

3. Data and Methodology

3.1 The Dutch National Bank Household Survey

Our main source of data is the Dutch National Bank Household Survey (DHS), an internet survey of a representative panel of Dutch households collected by CentERdata since 1993.⁵ We restrict our sample to household heads and their spouses, and we exclude respondents who are studying full time, looking for employment for the first time, or are solely living on disability benefits. Finally, we only include respondents who report more than €500 of financial wealth and some cash holdings.

⁵ The DHS currently surveys around 2,000 Dutch households each year. CentERdata provides participating households with an Internet connection, a television, and a set-top box for the television to facilitate survey responses. All persons aged 16 and over within each household are invited to take part in the survey. The response rate for a typical household questionnaire is around 86 percent. For a more detailed description of the DHS, see Teppa and Vis (2012) and <https://www.centerdata.nl/en/projects-by-centerdata/dnb-household-survey-dhs>.

We focus on equity investments as our measure of risky financial assets. These risky financial assets include stocks, mutual funds, growth funds, and options. The other components of financial wealth measured in the DHS are checking, savings, and deposit accounts; government and municipal bonds; mortgage bonds; illiquid saving certificates; and other unspecified investments and savings. We focus on the decision to hold *some* risky financial assets—the extensive margin in portfolio choice—by analyzing a dummy variable that takes the value of one if a person owns risky financial assets. In Section 4.1, we also analyze the share of wealth invested in risky assets.

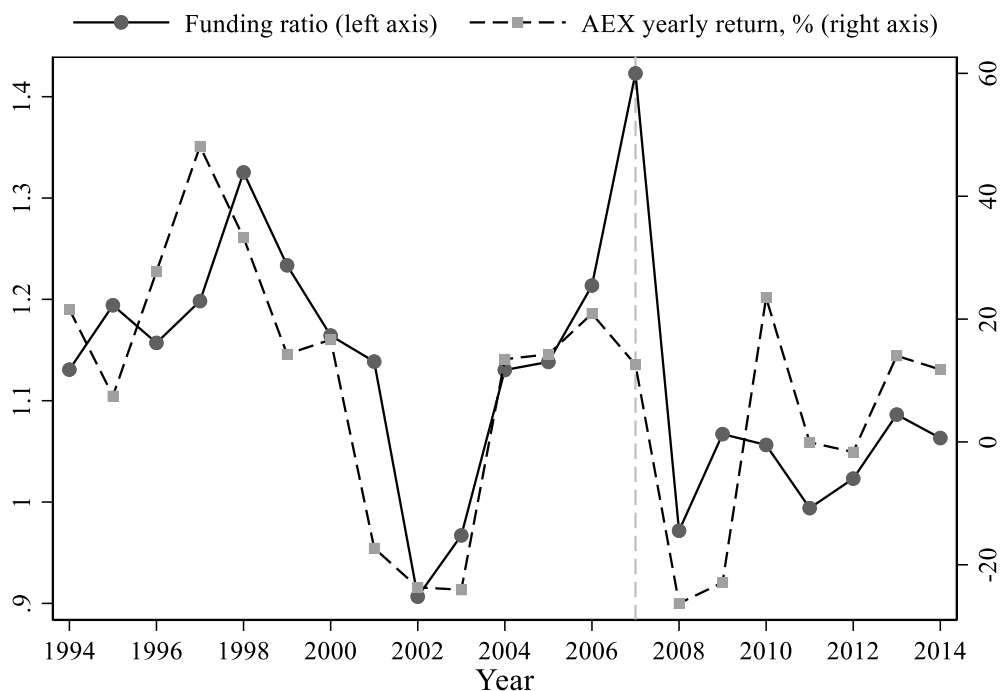
Our estimation sample excludes respondents for which we have incomplete financial asset information or that cannot be linked to past or current pension funds. We also exclude respondents without complete information on their sociodemographic characteristics (age, gender, occupation, and education) and key household characteristics (number of people in the household, number of children, household type, and net household income and wealth) since these are important covariates in our analyses. Our main estimation sample includes 1,937 people for a total of 5,268 person-year observations over the 2007–2014 period.

3.2 *Funding Ratio as a Measure of Background Risk*

To match the data on pension funds' funding ratios to the DHS, we use a survey question that asks the name of the pension fund to which each person contributes (for employees) or from which each person receives benefits (for retirees). We use the names of these pension funds to match the survey data with the funding ratios of each pension fund from 2007 through 2014, obtained directly from the Dutch Central Bank records. Our sample contains 30 different pension funds to which people contribute and can be matched to the 1994 through 2014 waves of the survey.

Funding ratios are a good measure of people's pension background risk since the ratios are highly correlated with the market and, after the 2007 reform, they are also linked to people's pension benefit inflation indexation and contribution rates. To illustrate this, Figure 2 shows that the average pension fund's funding ratio closely tracks the Amsterdam Exchange Index yearly returns over time, even after the 2007 reform. After 2007, the funding ratio is also closely related to the benefit inflation indexation rate (see Appendix Figure A1). Finally, the correlation between the funding ratio, the benefit inflation indexation rate, and

Figure 2. Average funding ratio and the Amsterdam Exchange Index (AEX) yearly returns, 1994–2014



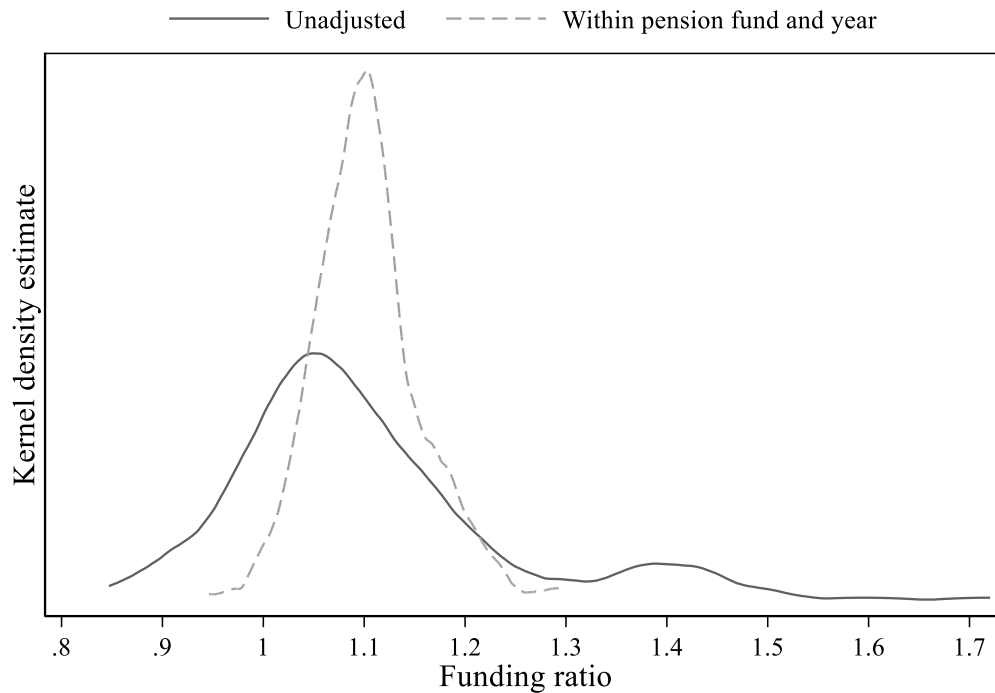
This figure illustrates the correlation between the funding ratio and the AEX market returns. The average funding ratio for the funds in our data is highly correlated ($\hat{\rho} = 0.62$) with the AEX yearly returns in the 1994–2014 period. This correlation remains after the 2007 pension reform (dashed line). The funding ratio averages funding ratios of 31 different pension funds each year. AEX yearly return averages 12 yearly returns based on the adjusted AEX closing price on the first day of each month, each year.

the AEX returns are high and positive for most of the pension funds in our data (see Appendix Figure A2).⁶

Figure 3 shows the variation in funding ratios as well as the within-pension-fund and within-year variation that identifies our estimates. There is substantial variation in unadjusted funding ratios (solid line), which is our inverse measure of pension risk, during 2007–2014. Overall, pension funds did well during this period—the median funding ratio was 1.07 and the average was 1.10. However, over 17 percent of the pension funds were underfunded at some point, and their funding ratio dropped as low as 0.86. The within-pension-fund and within-year variation in funding ratios (dashed line) is more compressed than the unadjusted one, and it is not right-skewed. However, the differences between adjusted funding ratios is substantial and this variation identifies our estimates.

⁶ There is no data available on the contribution rates, yet we expect them to be less tightly linked to funding ratios since changes in these rates are only implemented when finding ratios remain very high or very low for several consecutive years. Inflation indexation data is not available before 2007 since indexation was not linked to funding ratios then.

Figure 3. Variation in Dutch pension funds' funding ratios, 2007–2014



This figure shows the variation in funding ratios in our main estimation sample before (solid line) and after (dashed line) accounting for pension fund and year fixed effects. The dashed line is produced by subtracting the pension fund- and year-specific means of each funding ratio and then adding the overall mean to make both distributions comparable. The dashed line illustrates the identifying variation in our estimates. This figure is based on 220 pension-fund-by-year observations.

Table 1 presents the basic summary statistics of all key variables in our sample. About one quarter of respondents hold risky assets at any time, a share similar to the United States and other developed countries (see Guiso et al. 2008, p. 2561). Almost half of our sample live only with their partner, about 30 percent also live with their partner and children, and most of the remaining respondents live alone. There are relatively few women and almost everyone is either employed or retired. Our sample is relatively wealthy, although approximately 5 percent of people are net debtors at some point.⁷ The large majority earn between €22,000 and €75,000 in net household income each year. Overall, the average respondent in our

⁷ Appendix Figure A3 shows the distribution of net household wealth.

Table 1. Summary statistics of the main DHS estimation sample, 2007–2014

	Mean	Std dev.	Min	Max
Funding ratio (unstandardized)	1.09	0.17	0.85	1.72
Own risky assets	0.26	0.44	0	1
Share of wealth in risky assets [♦]	0.40	0.29	0	1
Female	0.36	0.48	0	1
High school degree	0.31	0.46	0	1
University degree	0.49	0.50	0	1
Occupation:				
Employed	0.63	0.48	0	1
Self-Employed	0.01	0.09	0	1
Unemployed	0.01	0.09	0	1
Retired	0.36	0.48	0	1
Household Composition:				
Living alone	0.20	0.40	0	1
Partner, no children	0.49	0.50	0	1
Partner, with children	0.28	0.45	0	1
No partner, with children	0.02	0.15	0	1
Other	0.02	0.12	0	1
People in the household	2.40	1.19	1	8
Children in the household	0.59	1.01	0	6
Household net wealth (€1,000)	73.2	165.9	-1,149.7	3,702.1
Household's net income (in Euro):				
less than €10,000	0.02	0.15	0	1
between €10,000 and €14,000	0.01	0.11	0	1
between €14,000 and €22,000	0.06	0.23	0	1
between €22,000 and €40,000	0.28	0.45	0	1
between €40,000 and €75,000	0.43	0.50	0	1
€75,000 or more	0.17	0.37	0	1
don't know	0.02	0.14	0	1

[♦]Share of wealth in risky financial assets is conditional on positive risky financial asset holdings.
Summary statistics using 5,268 person-year observations from 1,937 respondents.

sample is similar to the average DHS respondent, which is representative of the Dutch population (see Appendix Table A1).

3.3 Econometric Framework

We focus on estimating the effects of Funding ratio_{pt}, our inverse measure of pension risk for fund p at time t , on the probability of holding risky assets by using the following equation:

$$1[\text{Risky assets} > 0]_{ipt} = \beta \text{Funding ratio}_{pt} + \gamma' X_{it} + \delta_p + \theta_t + \varepsilon_{ipt}, \quad (1)$$

where $1[\text{Risky assets} > 0]_{ipt}$ is a dummy that marks whether person i participating in pension fund p at time t owns any risky assets. We rescale Funding ratio $_{pt}$ to have a zero mean and a standard deviation of one to facilitate the interpretation of β as the impact of a one standard deviation increase in the pension fund's funding ratio on the probability of investing in risky assets. The main covariates included in X_{it} are education, household size, partnership status, household income, household wealth, and age, gender, occupation, unemployment and retirement indicators. These covariates are common in the literature on portfolio choice (e.g., Rosen and Wu 2004; Hong et al. 2004; Guiso et al. 2008). The three error term components δ_p , θ_t , and ε_{ipt} capture time-invariant unobserved heterogeneity at the pension fund level, time-varying unobserved heterogeneity, and an independently distributed model error term. Both δ_p and θ_t can be correlated with observable and unobservable characteristics, and are accounted for via two-way pension-fund and year fixed effects. We use cluster-robust standard errors at the pension fund level. Throughout Section 4, we present our estimates of OLS linear probability models of Equation (1).

We argue that under plausible circumstances, and conditional on the covariates above, our estimates of Equation (1) identify the causal effect of pension risk on people's decisions to invest in risky assets. Individual time-varying confounders, such as single high-income men choosing riskier portfolios and also choosing jobs with more risk-taking pension funds, are accounted for by holding constant these characteristics in X_{it} . Any remaining self-selection of risk-seeking people into jobs (and thus pension funds) based on unobserved time-constant characteristics is accounted for by the pension fund fixed effects.⁸ These fixed effects also account for any other difference that is pension-fund specific and remain constant over time. Common market shocks that could drive funding ratios and individual investment decisions, such as the financial crisis market shock between 2007 and 2008, are accounted for by the year fixed effects, as well as anything else that is a year-specific common component. Finally, in our preferred specification we account for possible industry-specific time-varying shocks by expanding our set of year fixed effects to year-by-province fixed effects, noting that the industrial activity in The Netherlands is strongly segregated by province (see Section 4.1).

Our empirical strategy therefore accounts for most reasonable sources of potential omitted variable bias. Any remaining source of endogeneity needs to i) be unaccounted for by our extensive set of controls, ii) be time-varying within provinces, iii) affect people in the same

⁸ Also note that self-selection into riskier jobs based on risk preferences would not be enough to cause spurious effects in our setting unless the pension funds that manage riskier occupations are themselves more risk-taking in their portfolios.

way pension risk would, and iv) have an effect large enough to drive our results. These types of confounders are very difficult to imagine. And even if they exist, these confounders would also have to account for the pre-reform placebo analyses and heterogeneous effects addressed in Sections 4.2 and 4.3.

4. Results

4.1 *The Effect of Pension Risk on Investment in Risky Assets*

Table 2 presents our main results. Column (1) shows that a higher funding ratio, i.e., a lower risk in the pension fund’s portfolio, is positively related to people’s decisions to hold risky assets, as hypothesized. Column (2) shows that the effect of the funding ratio increases once we include pension fund fixed effects and year fixed effects (i.e., two-way fixed effects), which are the key components for our identification strategy. Column (3) shows that the size of this effect does not change when we include in the model various personal and household characteristics, including household income and wealth, although the precision in our estimate increases. Note also that the R-squared increases from 0.05 to 0.19 when moving from (2) to (3). Along with the very stable effect of pension risk, this suggests that unobserved heterogeneity correlated to the many observable characteristics we include in Column (3) is unlikely to drive the effect of the funding ratio. All these results are thus consistent with the idea that pension risk decreases people’s willingness to invest in risky assets.

Table 2. The effects of funding ratios on people’s decisions to hold risky assets, 2007–2014

Dependent Variable:	1 if person invests in risky assets			
	(1)	(2)	(3)	(4)
Funding ratio (standardized)	0.025** (0.011)	0.055* (0.030)	0.055** (0.026)	0.056** (0.026)
Person & household characteristics			✓	✓
Pension fund fixed effects		✓	✓	✓
Year fixed effects		✓	✓	
Province-by-year fixed effects				✓
R-squared	0.00	0.05	0.19	0.20

OLS regressions using 5,268 person-year observations from 1,937 respondents. The dependent variable is a risky asset investor dummy (stocks, mutual funds, growth funds, and options). Person & household characteristics include a linear age term and dummies for gender, occupation, and unemployment status, highest level of education, number of household members, children living in the household and partnership status, seven dummies for net household income bracket, and ten decile dummies for net household wealth. There are 30 pension fund fixed effects, eight year fixed effects, and 96 province-by-year fixed effects. Standard errors clustered at the pension fund level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

A remaining source of endogeneity in our estimates could be industry-specific shocks in unemployment, which can simultaneously drive the pension funds' funding ratio and other sources of income risk (through unemployment or earnings).⁹ However, this is unlikely to be driving our results for two reasons. First, we control for unemployment status, among other covariates, in column (3). Second, in our estimation period, industry-specific shocks would have been most likely to appear in 2008 when the financial crisis impacted The Netherlands. For these shocks to be a confounding factor, they would have had to affect labor markets across industries in heterogeneous ways, creating between-industry dispersion in labor markets. Yet data from the Dutch Labor Force Survey (LFS) reveal no such pattern.¹⁰

Nevertheless, to account for possible industry-specific labor market shocks, we exploit that industrial activity in The Netherlands is highly segregated between its 12 provinces.¹¹ A set of province-by-year fixed effects would thus capture a substantial part of this time-varying industry-specific heterogeneity. Column (4) of Table 2 presents estimates from a two-way fixed effect model with province-by-year and pension fund effects. This specification, which includes the most comprehensive set of fixed effects and controls, is our preferred one and we maintain it throughout the remainder of the paper. The estimates show that the effect of the funding ratio on investment in risky assets remains unchanged. Furthermore, the R-squared does not increase substantially between columns (3) and (4), indicating that unobserved province-specific shocks are not a source of endogeneity.¹²

The effect of the pension funding ratio on people's investment in risky assets is economically sizeable. People exposed to a one standard deviation higher funding ratio are 5.6 percentage points more likely to invest in risky assets. As only about a quarter of our sample holds risky assets, this implies an increase in risky asset participation of 23 percent over the sample mean. This effect is also larger than the investment gap between university

⁹ A set of industry-by-year fixed effects would account for this source of endogeneity. However, the DHS does not ask about respondents' sector of industry. More importantly, because most pension funds are industry-specific, industry-by-year fixed effects would eliminate almost all of our identifying variation.

¹⁰ Appendix Figure A4 shows that although the job separation rate did increase in The Netherlands in the first years of the economic crisis, the standard deviation of this rate across industries remained the same. Appendix Figure A5 presents a similar analysis on employment rates by sector of industry and leads to the same conclusion. This pattern is consistent with the crisis having limited heterogeneous effects on industry-specific labor markets.

¹¹ See <https://www.lisa.nl/data/gratis-data/overzicht-lisa-data-per-provincie>. Employment in manufacturing is, for example, three times as high in the provinces of Limburg and North Brabant than in the province of North Holland, whereas the financial sector is concentrated in North Holland, where share in total employment in that sector is three times higher than in the province of Limburg and five times higher than in the province of Flevoland.

¹² Specifications that control for people's investment experience and that include pension-fund-specific time trends reach very similar conclusions. We do not consider individual fixed effect models since they heavily restrict the identifying variation, yielding imprecise estimates, and can further introduce selectivity into identification (Miller et al. 2019). We also do not consider probit or logit models since including fixed effects does not remove unobserved heterogeneity from the likelihood function in these nonlinear models. Moreover, nonlinear models that can condition unobserved heterogeneity out of the likelihood function (e.g., conditional logit) cannot incorporate two-way fixed effects.

graduates and high school dropouts in our data (4.7 percentage points). Compared to other findings in the literature, the effect of the funding ratio on investing in risky financial assets is slightly larger than the effect of social interactions (Harrison et al. 2004) and similar to the effect of trust (Guiso et al. 2008) and right-wing political preferences (Kaustia and Torstila 2011), all of which are considered sizeable.¹³

Finally, the average funding ratio for funds in our data decreased by approximately 2 standard deviations between 2007 and 2014. In this same period, participation in risky assets decreased by approximately 10 percentage points, from 29 percent to 19 percent. This implies that the increase in pension fund risk could potentially explain the entire drop in equity participation in this period.

We then analyze the extent to which pension risk also affects people's portfolio risk-taking at the intensive margin. In these analyses we regress the share of people's wealth invested in risky assets, bonds, cash, and other assets, on their pension's funding ratio using our preferred specification. Column (1) of Table 3 shows that a one standard deviation increase in the funding ratio makes people increase the share of their wealth invested in risky assets by 2.9 percentage points. Since the unconditional share of wealth invested in risky assets is 10.1 percent for this sample, this effect implies a 29 percent increase over the sample mean. This increase is compensated by a 2.2 percentage point decrease in the share of their wealth in cash, and a smaller decrease in the share of their wealth invested in bonds and other assets.¹⁴ Compared to other findings in the literature, the effect of the funding ratio on the share of wealth invested in risky assets is similar to the effect of a 10 percentage point increase in the standard deviation of income risk (Heaton and Lucas 2000), a one-standard-deviation increase in health risk (Edwards 2008), or a 10 percent of a standard deviation increase in the uninsurable wage risk (Fagereng et al. 2016).

4.2 Placebo Estimates: Pre-reform Analyses

Since the mid-1990s, pension funds have been operating similarly in The Netherlands; their funding ratios have been used as a summary measure of their financial health, they have been rigidly organized by industry, and they have been communicating with their participants via

¹³ Regarding our unreported covariate coefficients: age is not related to investments in risky assets once we control for wealth and income; women and self-employed people are less likely to invest in risky assets; and wealthier, higher-income, and better-educated people are more likely to invest in risky assets. These findings generally match comparable estimates in the literature.

¹⁴ Specifications that control for previous investment experience and that include pension-fund-specific time trends, as well as specifications that exclude people who invest none or all of their wealth in one type of asset, reach very similar conclusions. As in our main results, we do not consider individual fixed effect models or nonlinear tobit or fractional regression models.

Table 4. Placebo analyses of the effects of funding ratios on people’s decisions to hold risky assets, 1994–2006

Dependent Variable:	1 if person invests in risky assets (1994-2006)			
	(1)	(2)	(3)	(4)
Funding ratio (standardized)	0.015 (0.013)	-0.009 (0.035)	0.005 (0.023)	-0.000 (0.023)
Person & household characteristics			✓	✓
Pension fund fixed effects		✓	✓	✓
Year fixed effects		✓	✓	
Province-by-year fixed effects				✓
R-squared	0.00	0.03	0.23	0.26

OLS regressions using 3,403 person-year observations from 1,463 respondents. The dependent variable is a risky asset investor dummy (stocks, mutual funds, growth funds, and options). Person & household characteristics include dummies for number of household members, children living in the household and partnership status; seven dummies for net household income brackets; and ten decile dummies for net household wealth. There are six pension fund fixed effects, 12 year fixed effects, and 156 province-by-year fixed effects. Standard errors clustered at the pension fund level in parentheses. *** p<0.01, ** p<0.05, * p<0.1

yearly statements. However, pension funds’ benefit indexation and contribution policies have been tied to their funding ratios only since the 2007 pension reform. The introduction of this reform allows us to estimate placebo analyses on the relationship between the pension fund’s funding ratio and people’s investments in risky assets for the years before 2007 to further support the causal interpretation of our results.¹⁵

Table 4 shows the effects of the funding ratio on people’s investment in risky assets for the pre-reform years of 1994–2006. The specifications mirror those of our main results in Table 2, yet in all of them, the coefficient of pension funds’ funding ratio is insignificant. This lack of significance is entirely due to the size of the coefficients and not due to a loss of statistical power in the pre-reform years; standard errors are also very similar to those in Table 2. Furthermore, the (unreported) coefficients of key covariates such as the female dummy and self-employment have a similar size and statistical significance as the analyses underlying Table 2, and the model fit is also comparable. This confirms that the effect of pension risk on investment behavior is only present after the 2007 reform, when the mechanism linking funding ratios to benefit indexation and contributions became effective.¹⁶

¹⁵ After the reform, the DHS also changed the names of many of the pension funds in their data. Because of this, we cannot estimate a joint model using the pre-2007 and post-2007 data akin to a difference-in-difference analyses with two-way fixed effects. However, our main and placebo results are robust to using different pension fund fixed effects based on various mappings of the pre-2007 and post-2007 pension fund categorization.

¹⁶ Appendix Figure A6 shows that the necessary identifying variation in funding ratios is also present in the prereform years.

4.3 *Heterogeneous Effects of Pension Risk*

We have shown that higher pension risk, as measured by a pension fund's lower funding ratio, makes people invest less in risky assets. To reaffirm our results, we explore the effects of pension risk for groups of people on whom we would expect to find stronger effects. We estimate variations of Equation (1) interacting subgroup dummies with the funding ratio and all other main covariates and always including pension fund and year-by-region fixed effects. We then calculate the average marginal effects of the funding ratio on investment in risky assets for each subgroup and test their differences. Table 5 presents these average marginal effects and their differences.¹⁷

We base our first two subgroup analyses on the intuition that pension risk should only affect the behavior of people who are aware of it. We test this using the answers to the following two DHS questions:

- (1) *Did you receive a pension overview last year?*
- (2) *A pension plan can include an arrangement for correcting the pension that can be claimed and/or the pension that is actually being paid according to a price index and/or to a salary index. Pensions that are corrected in this way are called indexed to inflation. Is your (future) retirement pension indexed to inflation?*

The first question refers to a report Dutch pension funds are required to send to their participants each year, which includes their pension fund entitlements and information about the pension fund's recent performance and current indexation conditions. The second question refers to the indexation policy of the pension fund, a widely debated and discussed topic in The Netherlands that draws considerable public attention. Answering the first question with "yes" instead of "no" indicates that respondents are keeping track of their pensions. Answering the second question with either "yes" or "no" instead of "I don't know" indicates that respondents are acquainted with the indexation conditions of their pension fund and are therefore more likely to be aware of the pension risk they face. The first two subgroups we analyze are the pension-risk-aware and the pension-risk-unaware people who are categorized based on these two questions.¹⁸

¹⁷ Prereform analyses show no effect of the funding ratio on investment in risky assets for any of these subgroups, reaffirming the causal interpretation of our results. See Appendix Table A2.

¹⁸ Eighty-one percent of our sample reports having received their pension overview, and 74 percent were aware of the indexation status of their pension funds. These questions are asked every year so the same respondent might belong to the aware or the unaware group depending on their response that year.

Table 5. Heterogeneous effects of funding ratios on people’s decisions to hold risky assets, 2007–2014

Dependent Variable: Effect of funding ratio (standardized) if:	1 if person invests in risky assets			
	Received overview (1)	Aware of indexation (2)	Retired (3)	Risk- Averse (4)
Yes	0.063** (0.027)	0.064** (0.027)	0.077** (0.029)	0.063** (0.028)
No	0.019 (0.030)	0.040 (0.027)	0.048* (0.027)	0.053** (0.025)
<i>Difference</i>	0.044*** (0.011)	0.024* (0.014)	0.028* (0.014)	0.010 (0.010)
Person & household characteristics	✓	✓	✓	✓
Pension fund fixed effects	✓	✓	✓	✓
Province-by-year fixed effects	✓	✓	✓	✓
R-squared	0.21	0.21	0.21	0.29
Observations	5,080	4,843	5,268	5,249

Average marginal effects of the standardized funding ratio in OLS regressions that the interact funding ratio with whether people report receiving their yearly pension overview (column (1)), are aware of their pension indexation policies (column (2)), are retired (column (3)), or score above median in risk aversion (column (4)). The dependent variable is a dummy that takes the value of one if the person owns risky assets (stocks, mutual funds, growth funds, and options). Marginal effects based on fully interacted models (except for fixed effects). Person & household characteristics include age, female, occupation, unemployed, retired, education, household members, children in household, partnership status, household income, and household wealth. Standard errors clustered at the pension fund level in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Column (1) of Table 5 shows that the negative effect of the funding ratio is only significant for people who reported receiving their yearly pension overview. Column (2) shows that the effect is also only significant for those who are aware of the indexation policy of their pension fund. Both group differences are statistically and economically significant, supporting the hypothesis that our main results reflect changes in people’s investment behavior due to changes in pension risk.

We can also verify the behavioral interpretation of our results by looking at the investment behavior of retirees. Since pension income is likely the primary source of income for retirees, and benefit indexation directly determines this income, we expect retirees to react more sharply to changes in their pension fund risk. Column (3) of Table 5 shows that pension fund risk has a stronger impact on retirees’ investments in risky assets. These results also highlight the disproportional effect of pension risk on retirees as an important distributional effect.

A final subgroup we can use for our tests is based on risk-aversion. Some findings from lifecycle portfolio theory suggest that the decrease in investment in risky assets in response

to background risk should be steeper for more risk-averse people.¹⁹ This prediction is somewhat less intuitive and is derived from models that impose stricter assumptions on people’s utility functions. Nevertheless, we can test whether risk-averse people react more strongly to pension risk with our data since we have a measure of people’s risk preferences in the financial domain. The DHS includes six self-reported statements measuring people’s willingness to take financial risks (i.e., investment risks with their own money). We construct an index of financial risk-aversion based on these questions, and classify people as risk-averse and non-risk-averse based on whether their score on this index is below or above the sample median.²⁰ Column (4) of Table 5 shows that the effects of the funding ratio for risk-averse people are slightly stronger than for non-risk-averse people. However, these differences are not statistically significant. This is the only result that is not strongly aligned with the interpretation of our findings as a behavioral response.

Overall, the evidence in this section supports the behavioral interpretation of our estimates as causal effects of pension risk on people’s investment decisions.

4.4 The Wealth and Risk Components of the Effects of Pension Risk

An intriguing question regarding our main results is whether they are not simply reflecting people’s reaction to a loss in their pension wealth rather than a pension risk effect. In this section we argue that one could, in principle, decompose the behavioral effects of pension funding ratios on risky investment into a “deterministic wealth” effect and a “pure risk” effect. We then highlight the conditions under which this decomposition can be accomplished, the limited way in which we can implement such decomposition in our data, and how those results can affect the interpretation of our main findings.

Gollier and Pratt (1996) show that the effect of an unrewarded undiversifiable background risk on risk aversion—as captured by the Arrow-Pratt absolute risk aversion measure—can be decomposed into two components: the effect of a deterministic downward shift in wealth and the effect of mean-zero uncorrelated pure risk (p. 1114). Through their shifts in the Arrow-Pratt measure, these two components can theoretically be conceived as the “wealth” and the “pure risk” effect of background risk on risk-taking. Applied to our analysis of

¹⁹ In Edwards (2008), this prediction is derived through the interaction between (health) background risk and the risk-aversion parameter in the model. However, this only holds for risk-averse people and relies on the CRRA formulation of his model. In Heaton and Lucas (2000), more risk-averse people decrease their investment in risky assets more substantially when faced with background risk.

²⁰ The index measures risk-aversion on a continuous scale between 1 and 7, with a median of 5.3. Further details on the specific questions included in this index and on the construction and validity of economic preference indices using the DHS data can be found in Salamanca (2016).

pension risk, one could thus think of separating the wealth and the pure risk effects of pension risk on the demand for risky assets. However, two issues must be highlighted. First, both the wealth effect and the risk effect are an integral part of the background risk effect, and considering either of them in isolation offers limited value. Second, the ultimate effect on the demand for risky assets operates through the Arrow-Pratt measure of risk aversion, complicating the comparative statics needed to disentangle these effects.

Nevertheless, if we wish to empirically disentangle the pure wealth and pure risk components of background risk, we could take two paths. The first one uses preference parameters described by three higher-order derivatives of the utility function—risk aversion, prudence, and temperance. Specifically, in Gollier and Pratt (1996) the pure risk component is a function of risk aversion and prudence, whereas the pure wealth component is, in addition, a function of temperance. One could therefore obtain empirical measures of these higher-order preferences and use them in a triangular system to isolate the wealth and pure risk components of pension risk. There are data and methodological issues with this approach. Data-wise, one would require measures of risk aversion, prudence, temperance, and the certainty equivalent and variance of retirement income attributable to pension risk. These measures are not easy to find in our data, or any other dataset for that matter. Methodologically, one would need to map prudence and temperance changes onto the Arrow-Pratt measure of absolute risk aversion and then map this measure to investment behavior, which is also not trivial.

The second approach, which we take in this paper, is to directly approximate the pure wealth effect of pension risk on investment behavior by estimating the effect of funding ratios on a measure of people's future retirement income. We then control for future retirement income in our main specification to approximate the pure risk effect as the remaining effect of pension funding ratios on investment behavior. This additional control does not improve our estimates of the effect of pension risk, but if this effect runs solely through changes in expected wealth, then controlling for expected wealth should result in large changes to the effect of pension risk. The benefit of this approach is that it is feasible with our data. The drawback is that it relies on our less-than-ideal measure of future retirement income. Because of this, we view its results as suggestive at best.

To implement this approach, we use a question that asks people what their pension will be after they retire if they were to make no further contributions. This amount—which is reported in a pension overview that Dutch people receive each year—is a measure of people's future pension wealth without any labor supply adjustments or further contributions, but

Table 6. The effects of funding ratios on people’s pension wealth and the decisions to hold risky assets, 2007–2014

Dependent Variables:	Pension wealth	1 if person invests	
	(1,000 Euro)	in risky assets	
	(1)	(2)	(3)
Funding ratio (standardized)	3.085* (1.622)	0.085** (0.034)	0.087** (0.034)
Pension wealth (1,000 Euro)			0.001 (0.000)
Person & household characteristics	✓	✓	✓
Pension fund fixed effects	✓	✓	✓
Province-by-year fixed effects	✓	✓	✓
R-squared	0.38	0.21	0.21

OLS regressions using 3,063 person-year observations from 1,377 respondents with nonmissing information on their pension wealth. The dependent variables are the yearly pension that respondents would receive if they retired without making any further contributions (columns (1) and (2)), and a risky asset investor dummy (stocks, mutual funds, growth funds, and options, columns (3) and (4)). A dummy for respondents who report not knowing their pension wealth (10.3% of sample) is always included. Person & household characteristics include age, female, occupation, unemployed, retired, education, household members, children in household, partnership status, household income, and household wealth. Fixed effects include 30 pension fund dummies and 96 province-by-year dummies. There are 30 pension fund fixed effects and 96 province-by-year fixed effects. Standard errors clustered at the pension fund level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

which should capture forecasted fund indexation.²¹ If the pure wealth component of pension risk is important, there should then be a positive effect of funding ratios on this measure of future pension wealth.

Column (1) of Table 6 shows that the funding ratio has a positive effect on the expected pension wealth of people, consistent with the existence of a pure wealth effect. However, the effect is not precisely estimated, reflecting the fact that expected pension wealth is measured with considerable error. Consistent with this, column (2) shows that the effect is smaller and not statistically significant once we control for person and household background characteristics. Column (3) and column (4) show that controlling for future pension wealth does not change the effect of the funding ratio on investment behavior; if anything, it increases it. Taken together, this evidence provides limited evidence that the main component of the effect of pension risk on people’s investment behavior is risk rather than wealth.

There is one final argument of why the pure wealth component of pension risk cannot be too large in our setting. As mentioned above, the wealth component effect of background risk

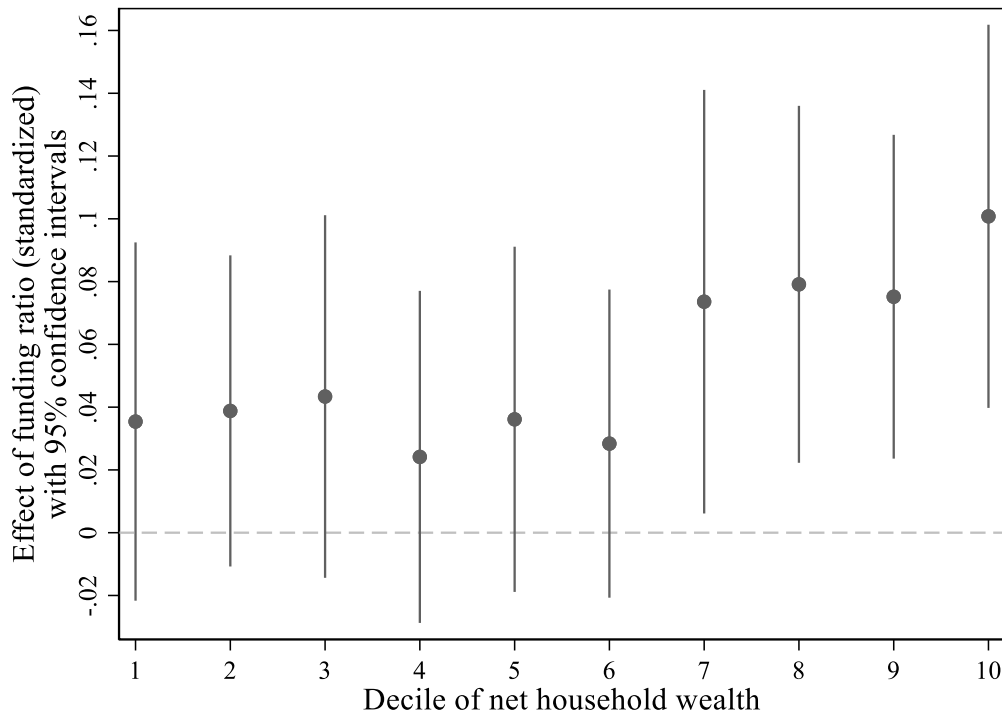
²¹ The question asks “According to [your pension] overview, how much pension had you built up until 1 January through your current/last employer? (i.e. the amount per year if you stopped working right then).”

is a function of prudence, a higher-order preference which by itself would make people increase their asset holdings when exposed to more background risk (Kimball 1990). This comparative static is further illustrated in Gomes and Michaelides (2005). A background risk effect with a dominant wealth component should therefore cause people to *increase* their risky asset holdings as a precautionary measure, not to decrease them as we observe in our data. Since our findings identify a negative effect of pension risk on the likelihood of investing in risky assets, it must be that the dominating component in our setting is the (negative) pure risk component and not the (positive) wealth component of pension risk.

4.5 The Distributional Effects of Pension Risk


It is also important to document the distributional effects of pension risk across several subgroups. People’s financial situations, exposure to other risks, and ability to adjust their portfolios can vary widely, and therefore so can their portfolio response to pension risk. Documenting these distributional effects can identify which population groups should receive attention when there are large changes in pension risk after pension reforms or market shocks.

Figure 4. The effect of funding ratios on people’s decisions to hold risky assets, by household net wealth decile, 2007–2014



This figure shows average marginal effects of the standardized funding ratio in OLS regressions that interact the funding ratio with net household wealth deciles using 5,268 person-year observations from 1,937 respondents. The dependent variable is a dummy that takes the value of one if the person owns risky assets (stocks, mutual funds, growth funds, and options). The rest of the model follows our preferred specification. The figure also shows 95 percent confidence intervals for each estimate based on clustered standard errors at the pension fund level.

Figure 5. The effect of funding ratios on people’s decisions to hold risky assets, by respondent age brackets and excluding retirees, 2007–2014



This figure shows average marginal effects of the standardized funding ratio in OLS regressions that interact the funding ratio with respondent age brackets of 20–29, 30–39, 40–49, 50–59, 60–69, and 70 or above using 5,268 person-year observations from 1,937 respondents. The dependent variable is a dummy that takes the value of one if the person owns risky assets (stocks, mutual funds, growth funds, and options). The rest of the model follows our preferred specification. The figure also shows 95 percent confidence intervals for each estimate based on clustered standard errors at the pension fund level.

Figure 4 shows that the effect of larger funding ratios is mostly present for people in the top four deciles of the net wealth distribution. One simple explanation for this finding is that the wealthiest people are more likely to own risky assets in the first place and can therefore adjust their investments more easily by selling them. Figure 5 illustrates a related finding: the effects of pension risk are largely driven by people of pension age, between 60 and 70 years old. This is not simply another illustration of the stronger effects of the funding ratio on retirees (Table 5), since the data in Figure 5 exclude retirees. Our findings are better explained by retirees and people around retirement age being uniquely affected by pension risk and having enough wealth and experience to be able to adjust their investments to this risk more easily.

Of course it could be that wealthier people also have more pension wealth and therefore pension risk affects them on a larger asset base. Yet we see there is no such gradient in the effect of funding ratios across our measure of future pension wealth. We also see little heterogeneity across the many other potential distributional effects we explore. There is some

suggestive evidence of larger effects of funding ratios for very high- and very low-income households, for unemployed people, for homeowners, and for people with very high subjective financial literacy. There is no evidence of heterogeneous effects across education or objective financial literacy, partnership status, or household size.

4.6 *The Effects of Pension Risk on Other Outcomes*

Finally, it is important to explore whether pension risk affects other outcomes. In our institutional setting, people cannot easily change jobs in response to pension risk, and aside from their investment portfolios, there are not many ways in which they can adjust their behavior in response to this risk. However, it is possible for them to adjust their labor supply by, for example, delaying retirement. They could also change their behavior in ways that would not be easy to detect like, for example, repurposing their savings to compensate for future pension shortcomings. These nonportfolio responses could be important to determine the potential reach of pension risk on people's overall welfare.

Table 7 shows that pension risk has no effect on other nonportfolio outcomes we can observe in our data. Columns (1) and (2) in Panel A show that funding ratios have no impact on people's retirement decisions or on their expectations for retirement before the age of 65. Column (3) also shows that funding ratios also have no effect on their intentions to leave a bequest to their families. Panel B further shows that pension risk does not affect the intentions behind people's saving behavior. Funding ratios do not have an effect on the likelihood of respondents' reporting that it is important or very important to save in order to supplement their pensions (column (4)), to not have to financially rely on others (column (5)), or to safeguard against unexpected future expenses (column (6)).

There are at least two reasons to interpret the evidence in Table 7 with caution. First, our measures of expected retirement and bequests suffer from relatively large nonresponse rates, and there is evidence of selective nonresponse for the bequest question (see Appendix Table A3). Second, the subjective nature of these outcome measures and the fact that many respondents may not have thought about these issues before can increase measurement error and decrease our power to detect effects of pension risk. We therefore view this evidence as only suggesting that pension risk does not affect these nonportfolio outcomes.

Table 7. The effects of funding ratios on other outcomes, 2007–2014

Panel A			
Dependent Variables:	1 if...		
	retired (1)	exp. retirement age<65 (2)	plans to leave bequest (3)
Funding ratio (standardized)	-0.010 (0.015)	0.031 (0.041)	-0.004 (0.042)
Outcome mean:	0.36	0.44	0.23
R-squared	0.67	0.15	0.10
Observations	5,268	2,626	3,772
Panel B			
Dependent Variables:	1 if (very) important to save to...		
	supplement pension (4)	not rely on others (5)	cover unexpect. expenses (6)
Funding ratio (standardized)	-0.017 (0.030)	-0.001 (0.034)	-0.011 (0.026)
Outcome mean:	0.72	0.62	0.48
R-squared	0.12	0.09	0.08
Observations	4,905	5,221	5,251
Person & household characteristics	✓	✓	✓
Pension fund fixed effects	✓	✓	✓
Province-by-year fixed effects	✓	✓	✓

OLS regressions. The dependent variables are dummies for whether the respondent: is retired (column (1)), expects to retire before the age of 65 (column (2)), plans to leave a bequest (column (3)); and whether the respondent believes it is important or very important (6 or 7 on a 1–7 scale) to save for: supplementing their pension (column (4)), not relying on others financially (column (5)), or to cover unexpected expenses (column (6)). Person & household characteristics include age, female, occupation, unemployed, retired, education, household members, children in household, partnership status, household income, and household wealth. In column (1), however, we exclude occupation as control to avoid perfectly predicting retirement. Standard errors clustered at the pension fund level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

5. Conclusions

In this paper, we analyze how an exogenous risk imposed on people's pension savings affects the way they invest their portfolios. We exploit a 2007 Dutch reform of its defined benefit pension system that tied people's pension benefits and contributions to their pension fund's market performance, as measured through their funding ratio. This effectively exposes people's retirement savings to background market risk. We show that a decrease in background pension risk, as measured by a one standard deviation increase in funding ratios

makes people 5.6 percentage points more likely to invest in risky assets and increases the share of wealth invested in risky assets by 2.9 percentage points.

People's behavior suggests that pension risk can substantially decrease people's welfare. It is difficult to assess the welfare loss imposed on people by pension risk, yet one way of doing so is by comparing the potential investment outcomes of the median Dutch person and a person who was exposed to the least amount of pension risk since the reform was enacted. In the 2007–2014 period, the funding ratio for the median person in our sample decreases by 1.9 standard deviations, yet for the best-performing fund, it increases by 0.15 standard deviations. A back-of-the-envelope calculation based on our estimates shows that the expected gap in wealth accumulation between the median and the least-exposed investor over these seven years is around €3,000. Our distributional effects further suggest that these losses could loom much larger for retirement-age and wealthier people. These estimates complement recent assessments that the lifetime welfare losses from retirement uncertainty is up to 5.7 percent of peoples' total lifetime consumption (Caliendo et al. 2016).

Moreover, there might be other important welfare losses even for people who manage to divest from other risky assets. Van Rooij et al. (2007) show that most Dutch people are reluctant to take control of their retirement savings investments, and, if given a choice, would prefer a managed pension scheme to avoid having to make complicated pension-related decisions. Therefore, even if people could generally offset the effect of additional background risk by divesting on other risky assets, they would still have to actively rebalance their portfolio to offset the risk, incurring direct and indirect costs. The good news, as we also show, is that there seem to be no effects of pension risk on people's retirement decisions or expectations, their bequest plans, or their motivations for saving. This, in many ways, simplifies the introduction of reforms and policies that explicitly account for pension risk as a behavioral response.

Our findings highlight a number of benefits of appropriately funded pension schemes that provide income certainty for people in their old age. There is a lot of value in pension certainty and we have shown that having more market risk on pensions can have important behavioral effects. Retirement plan designs should account for these nonnegligible risk effects, which can be quantified—as we have shown—and which could potentially be mitigated or outright eliminated by adopting well-designed pension insurance and risk-sharing schemes.

REFERENCES

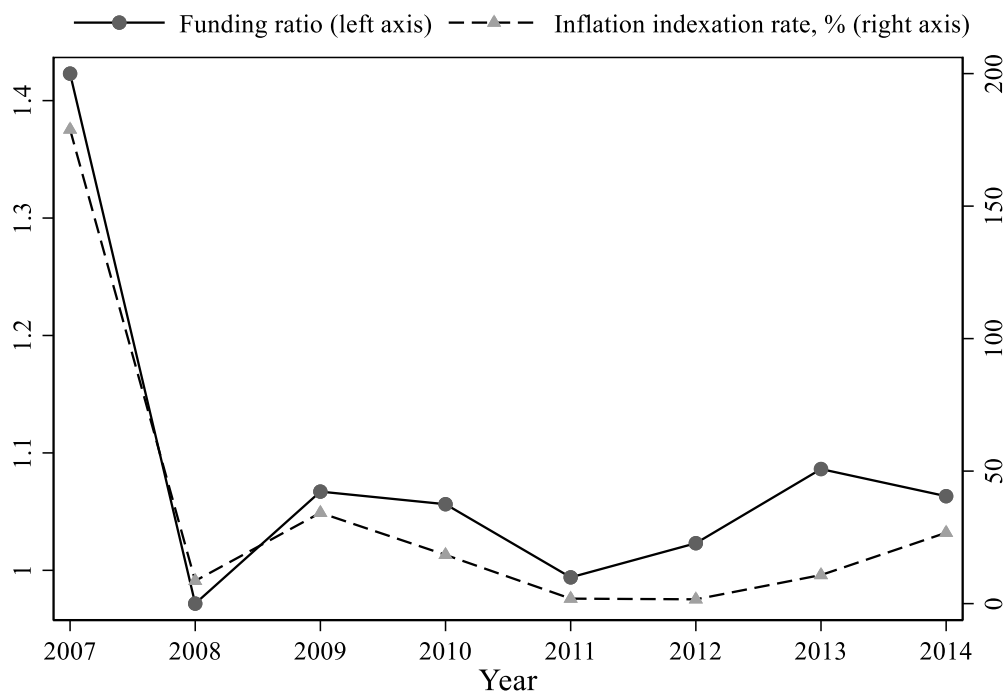
- X. Angerer, and P. Lam. Income risk and portfolio choice: An empirical study. *The Journal of Finance*, 64(2): 1037–1055, 2009.
- K. Arrow. *Aspects of the theory of risk-bearing*. Helsinki, Finland, Yrjö Jahnessonin Säätiö. 1965.
- S. Benartzi and R.H. Thaler. Heuristics and biases in retirement savings behavior. *The Journal of Economic Perspectives*, 21 (3): 81–104, 2007.
- Z. Bodie, R.C. Merton, and W.F. Samuelson. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control*, 16 (3-4): 427–449, 1992.
- F. Caliendo, M. Casanova, A. Gorry, and S. Slavov. The welfare cost of retirement uncertainty. National Bureau of Economic Research, No. w22609, 2016.
- C. Cameron, J.B. Gelbach, and D. L. Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90 (3): 414–427, 2008.
- B. Cardak and R. Wilkins. The determinants of household risky asset holdings: Australian evidence on background risk and other factors. *Journal of Banking & Finance*. 33(5), 850-860, 2009.
- L. Calvet, J.Y. Campbell, P. Sodini. Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, 115(5):707–47, 2007.
- G. Charness and U. Gneezy. Strong evidence for gender differences in investment. *Journal of Economic Behavior & Organization*, 83 (1): 50–58, 2012.
- J. Cocco, F. Gomes, and P. Maenhout. Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18(2): 491–533, 2005.
- S. Curcuru, J. Heaton, D. Lucas, D. Moore. Heterogeneity and portfolio choice: Theory and evidence. *Handbook of Financial Econometrics*, 1: 337–382, 2004.
- B. Donkers and A. Van Soest. Subjective measures of household preferences and financial decisions. *Journal of Economic Psychology*, 20 (6): 613–642, 1999.
- R. Edwards. Health risk and portfolio choice. *Journal of Business & Economic Statistics*, 26(4): 472–485, 2008.
- R. Edwards. Optimal portfolio choice when utility depends on health. *International Journal of Economic Theory*, 6(2): 205–225, 2010.
- A. Fagereng, L. Guiso, and L. Pistaferri. Portfolio choices, firm shocks and uninsurable wage risk. National Bureau of Economic Research, No. w22883, 2016.

- E. French. The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies*, 72(2): 395–427, 2005.
- C. Gollier and J.W. Pratt. Risk vulnerability and the tempering effect of background risk. *Econometrica*, 64 (5): 1109–1123, 1996.
- F. Gomes, and A. Michaelides. Portfolio choice with internal habit formation: A life-cycle model with uninsurable labor income risk. *Review of Economic Dynamics*, 6(4): 729–766, 2003.
- L. Guiso, T. Jappelli, D. Terlizzese. Income risk, borrowing constraints, and portfolio choice. *American Economic Review*, 86 (1): 158–72, 1996.
- L. Guiso, P. Sapienza, L. Zingales. Trusting the stock market. *Journal of Finance*, 63 (6): 2557–2600, 2008.
- J. Heaton and D. Lucas. Portfolio choice in the presence of background risk. *Economic Journal*, 110 (460): 1–26, 2000.
- H. Hong, J.D. Kubik, J.C. Stein. Social interaction and stock-market participation. *Journal of Finance*, 59 (1): 137–163, 2004.
- B. Honoré. Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica: Journal of the Econometric Society*, 60(3): 533–565, 1992.
- M. Kaustia and S. Torstila. Stock market aversion? Political preferences and stock market participation. *Journal of Financial Economics*, 100 (1): 98–112, 2011.
- M. Kimball. Precautionary Saving in the Small and in the Large. *Econometrica: Journal of the Econometric Society*, 58(1): 53-73, 1990.
- D. Miller, N. Shenhav, M. Grosz. Selection into identification in fixed effects models, with application to Head Start. National Bureau of Economic Research, No. w26174, 2019.
- B. Moulton. An illustration of a pitfall in estimating the effects of aggregate variables on micro units. *Review of Economics and Statistics*, 72 (2): 334–338, 1990.
- Organisation for Economic Co-operation and Development (OECD). *Pension Markets in Focus*. 2016.
- E. Ponds and B. van Riel. Sharing risk: The Netherlands’ new approach to pensions. *Journal of Pension Economics and Finance*, 8 (01): 91–105, 2009.
- J. Poterba, J. Rauh, S. Venti, D. Wise. Defined contribution plans, defined benefit plans, and the accumulation of retirement wealth. *Journal of Public Economics*, 91 (10): 2062–2086, 2007.
- J. Poterba and A.A. Samwick. Taxation and household portfolio composition: US evidence from the 1980s and 1990s. *Journal of Public Economics*, 87 (1): 5–38, 2003.

- J. Pratt and R.J. Zeckhauser. Proper risk aversion. *Econometrica*, 55 (1): 143–54, 1987.
- H. Rosen and S. Wu. Portfolio choice and health status. *Journal of Financial Economics*, 72 (3): 457–484, 2004.
- N. Salamanca. The Dynamic Properties of Economic Preferences. Unpublished mimeo, 2016.
- A. Samwick and J. Skinner. How will 401(k) pension plans affect retirement income? *American Economic Review*, 94(1), 329-343, 2004.
- O. Sleijpen. Exporteerbaarheid van het Nederlandse pensioenstelsel naar de Europese Unie [On the exportability of the Dutch pension system to the European Union] — Lecture at Maastricht University. 2009.
- F. Teppa, and C. Vis. The CentERpanel and the DNB household survey: Methodological aspects. *DNB Occasional Studies*, 10(4), 2012.
- R. Thaler and S. Benartzi. Naïve diversification strategies in defined contribution saving plans. *American Economic Review*, 91 (1): 79–98, 2001.
- M. Van Rooij, A. Lusardi, and R. Alessie. Financial literacy and stock market participation. *Journal of Financial Economics*, 60 (2–3): 401–448, 2011.
- M. Van Rooij, C.J.M. Kool, and H.M. Prast. Risk-return preferences in the pension domain: Are people able to choose? *Journal of Public Economics*, 91: 702–722, 2007.
- L. Viceira. Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance*, 56(2), 433–470, 2001.

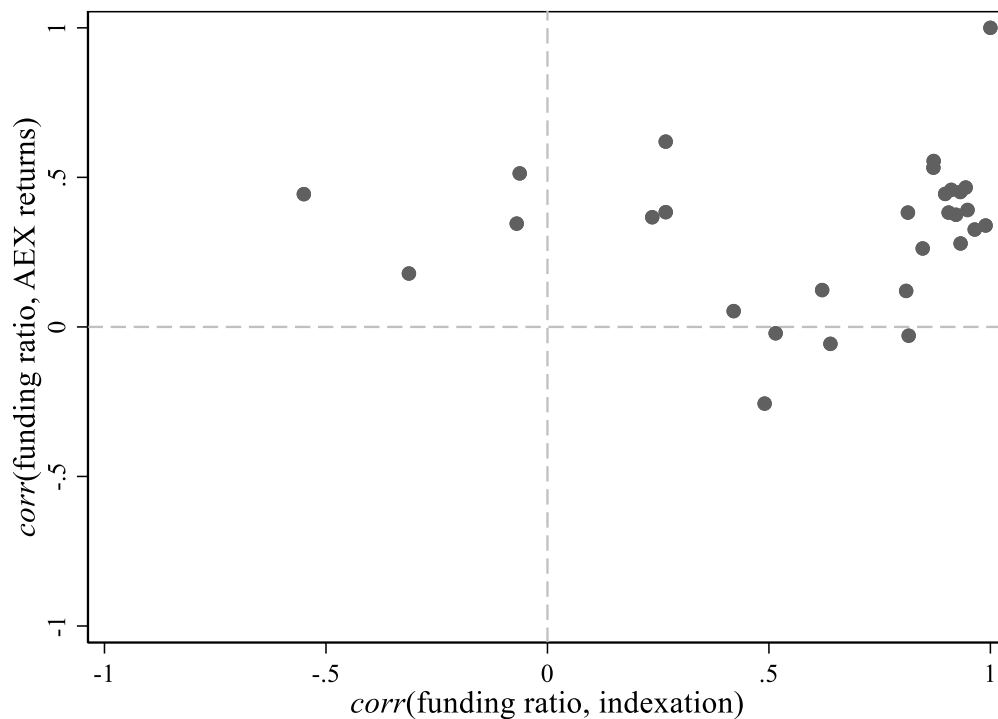
Appendix

Figure A1. Funding ratios and the inflation indexation rate of benefits, 2007–2014



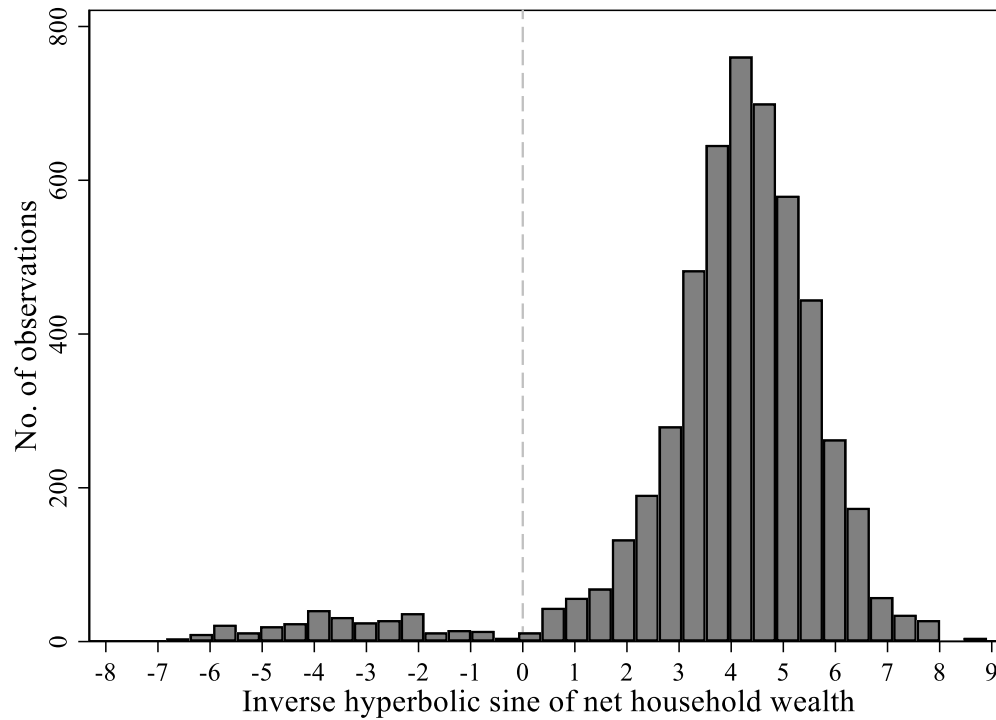
This figure illustrates the correlation between the funding ratio and the AEX market returns. The average funding ratio for the funds in our data is highly correlated ($\hat{\rho} = 0.97$) with the average inflation indexation rate of benefits in the 2007–2014 period. This correlation remains high ($\hat{\rho} = 0.60$) even after removing the high funding ratio and indexation rate observation in 2007. The funding ratio and benefit indexation rate observations average data from 31 different pension funds each year. Benefit indexation rate is measured in percentage points, where 0 percent means no indexation of pension benefits to inflation, 100 percent means full indexation of benefits to inflation, and values over 100 percent mean that benefits increased more than inflation that year.

Figure A2. Correlations between funding ratios, the Amsterdam Exchange Index (AEX) yearly returns, and the benefit indexation rate, by pension fund 2007–2014



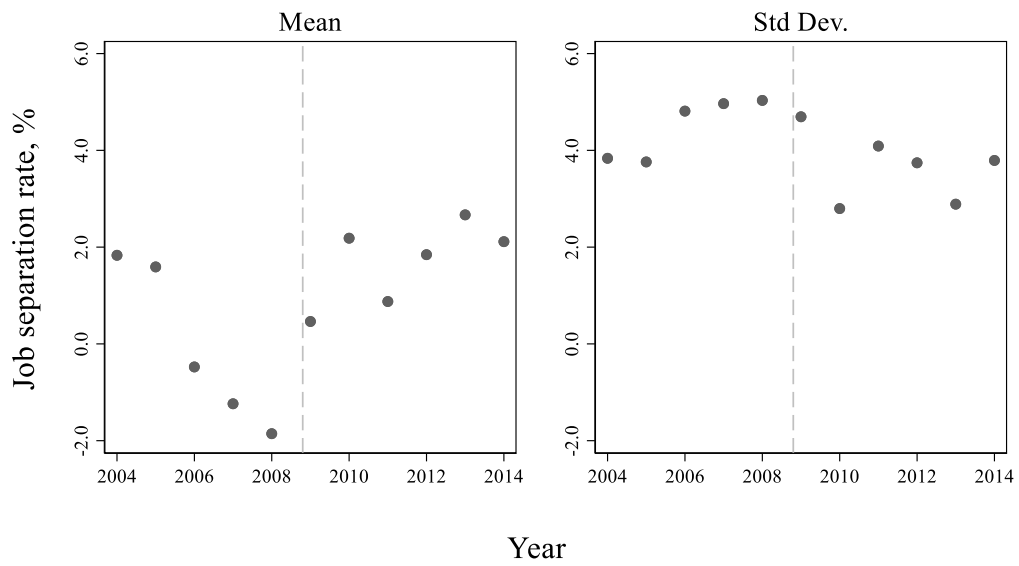
This figure plots the distribution of correlations between the funding ratio and the AEX market returns, and between the funding ratio and inflation indexation rate of benefits in the 2007–2014 period. Out of 29 pension funds with complete data for this analysis, almost all have a high positive correlation between their funding ratios and the AEX yearly returns (y axis), and between their funding ratios and their benefit indexation rate (x axis). Only four funds have negative correlations between their funding ratio and the AEX returns, and another four have a negative (though near zero) correlations between their funding ratio and their benefit indexation rate.

Figure A3. The inverse hyperbolic sine of net household wealth in the main estimation sample, 2007–2014




This figure illustrates the distribution of net household wealth in the 2007–2014 period. The inverse hyperbolic sine ($IHS = \sqrt{x + (x^2 + 1)}$) is similar to a natural logarithm transformation yet defined at zero (dashed line) and negative values. This transformation allows us to show the entire variation in net wealth in a simple manner. Net household wealth is the market value of all checking and savings accounts, single-premium annuity policies, owned real estate, business equity, and durables assets minus the market value of all loans, credits, credit card balances, debts, and outstanding mortgages. $N = 5,268$.

Figure A4. Job separation rates, means, and standard deviations in %, 2004–2014



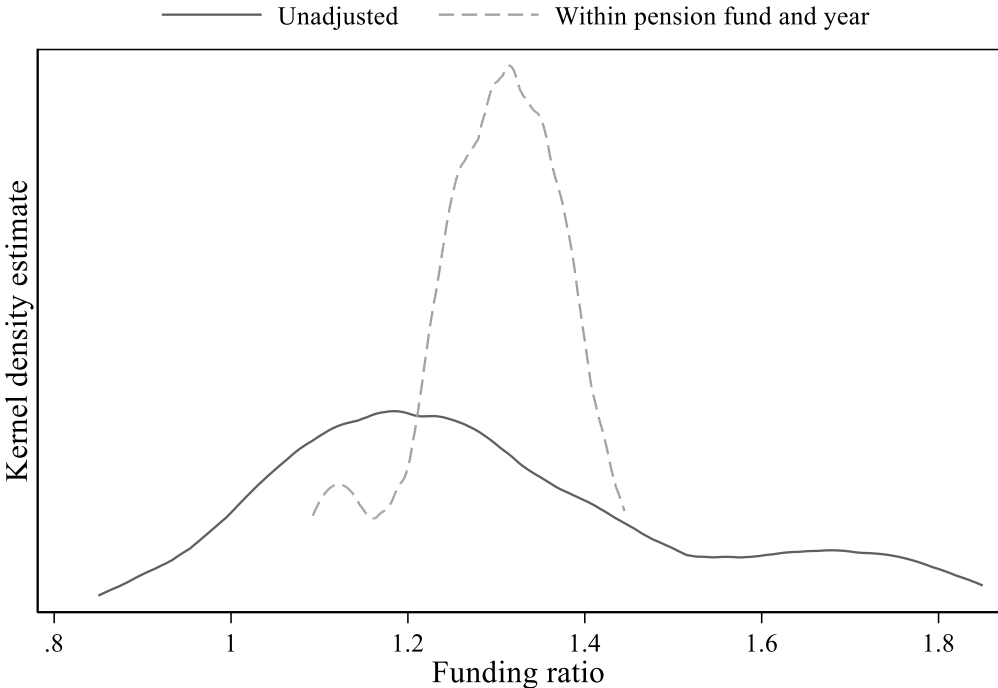
Source: Dutch Labor Force Survey (LFS) 2004–2015. The LFS is based on annual samples of more than 100,000 inhabitants, of which approximately 70,000 are employed in 21 industries. Annual percentages are those employed who were unemployed one year later. The financial crisis impacted The Netherlands in October 2008, which is indicated by the vertical dashed lines.

Figure A5. Means and standard deviations of employment rates, 1997–2014



Source: Dutch Labor Force Survey (LFS) 2004–2014. The LFS is based on annual samples of more than 100,000 inhabitants, of which approximately 70,000 are employed in 21 industries. The financial crisis impacted The Netherlands in October 2008, which is indicated by the vertical dashed lines.

Figure A6. Variation in Dutch pension funds' funding ratios in the placebo period, 1994–2006



The solid line shows the unadjusted funding ratios for the placebo period. The dashed line represents these same funding ratios adjusted by subtracting the pension fund- and year-specific means and then adding the overall funding ratio mean to make both distributions comparable. This figure is based on 249 pension-fund-by-year observations

Table A1. Comparison of main estimation sample with all available DHS data

	Estimation Sample		DHS 2007–2014		Difference
	Obs.	Mean	Obs.	Mean	
Funding ratio (unstandardized)	5,268	1.087	8,340	1.090	-0.004
Own risky assets	5,268	0.256	10,716	0.254	0.002
Share of wealth in risky assets [♦]	1,348	0.396	2,718	0.401	-0.005
Female	5,268	0.364	18,237	0.437	-0.073
High school degree	5,268	0.313	18,224	0.317	-0.003
University degree	5,268	0.493	18,224	0.450	0.043
Occupation:					
Employed	5,268	0.627	18,237	0.625	0.002
Self-Employed	5,268	0.009	18,237	0.078	-0.069
Unemployed	5,268	0.009	18,237	0.021	-0.012
Retired	5,268	0.356	18,237	0.276	0.079
Household Composition:					
Living alone	5,268	0.195	18,237	0.138	0.057
Partner, no children	5,268	0.490	18,237	0.462	0.028
Partner, with children	5,268	0.275	18,237	0.363	-0.088
No partner, with children	5,268	0.024	18,237	0.018	0.006
Other	5,268	0.016	18,237	0.018	-0.003
People in the household	5,268	2.396	18,237	2.612	-0.215
Children in the household	5,268	0.593	18,237	0.744	-0.150
Household net wealth (€1,000)	5,268	73.173	13,888	77.093	-3.920
Household's net income (in Euro):					
less than €10,000	5,268	0.022	11,284	0.029	-0.007
between €10,000 and €14,000	5,268	0.011	11,284	0.017	-0.006
between €14,000 and €22,000	5,268	0.058	11,284	0.067	-0.009
between €22,000 and €40,000	5,268	0.285	11,284	0.275	0.010
between €40,000 and €75,000	5,268	0.435	11,284	0.419	0.016
€75,000 or more	5,268	0.168	11,284	0.165	0.004
don't know	5,268	0.021	11,284	0.029	-0.008

[♦]Share of wealth in risky financial assets are conditional on positive risky financial asset holdings.

Table A2. Placebo analyses of the heterogeneous effects of funding ratios on people’s decisions to hold risky assets, 1994–2006

Dependent Variable: Effect of funding ratio (standardized) if:	1 if person invests in risky assets			
	Received overview	Aware of indexation	Retired	Risk- Averse
	(1)	(2)	(3)	(4)
Yes	0.033 (0.038)	0.021 (0.026)	-0.015 (0.024)	-0.017 (0.025)
No	0.063 (0.040)	-0.016 (0.020)	0.005 (0.024)	0.005 (0.023)
Individual & household characteristics	✓	✓	✓	✓
Pension fund fixed effects	✓	✓	✓	✓
Province-by-year fixed effects	✓	✓	✓	✓
R-squared	0.27	0.33	0.27	0.34
Observations	2,032	1,029	3,398	3,375

Average marginal effects of the standardized funding ratio in OLS regressions that interact the funding ratio with whether people report receiving their yearly pension overview (column (1)), are aware of their pension indexation policies (column (2)), are retired (column (3)), or score above the median in risk aversion (column (4)). The dependent variable is a dummy that takes the value of one if the person owns risky assets (stocks, mutual funds, growth funds, and options) in the pre-reform period of 1994–2006. Marginal effects based on fully interacted models (except for fixed effects). Person & household characteristics include age, female, occupation, unemployed, retired, education, household members, children in household, partnership status, household income, and household wealth. Standard errors clustered at the pension fund level in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A3. Selective nonresponse analyses for measures of other outcomes, 2007–2014

Dependent Variable:	1 if response is missing or "do not know" for...				
	Exp. retirement age	Plans to leave bequest	(very) important to save to...		
			supplement pension	not rely on others	cover unexp. expenses
Effect of funding ratio (standardized) if:	(1)	(2)	(3)	(4)	(5)
Funding ratio (standardized)	-0.024 (0.029)	0.037** (0.014)	-0.002 (0.004)	-0.007 (0.006)	-0.011 (0.021)
Mean of dependent variable	0.50	0.28	0.00	0.01	0.07
Person & household characteristics	✓	✓	✓	✓	✓
Pension fund fixed effects	✓	✓	✓	✓	✓
Province-by-year fixed effects	✓	✓	✓	✓	✓
R-squared	0.61	0.49	0.05	0.05	0.25
Observations	5,268	5,268	5,268	5,268	5,268

OLS regressions using 5,268 person-year observations from 1,937 respondents. The dependent variable is a dummy that takes the value of one if the respondent has missing data or responded “don’t know” for the dependent variable. Person & household characteristics include a linear age term and dummies for gender, occupation and unemployment status, highest level of education, number of household members, children living in the household and partnership status, seven dummies for net household income bracket, and ten decile dummies for net household wealth. There are 30 pension fund fixed effects, eight year fixed effects, and 96 province-by-year fixed effects. Standard errors clustered at the pension fund level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1