Journal of Financial Economics 113 (2014) 202–214

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Stock market returns and annuitization ☆

Alessandro Previtero *

Ivey Business School, Western University, 1255 Western Road, London, Ontario, Canada N6A3K7

A R T I C L E I N F O

Article history:
Received 17 December 2012
Received in revised form
30 September 2013
Accepted 30 October 2013
Available online 28 April 2014

JEL classification:
D14
G11
G22
H55

Keywords:
Annuities
Retirement income
Myopic extrapolation
Elderly

A B S T R A C T

I investigate the strong negative relation between recent stock returns and the annuitization of retirement savings using a novel data set with over 100,000 actual payout decisions. After controlling for several standard explanations (e.g., wealth effects), I present evidence supporting naïve beliefs and extrapolation from past returns. The effect of recent returns on annuitization dramatically increases with age, confirming that the elderly rely most heavily on recent information. My results provide insights into how beliefs are formed in old age and have implications for the design of public policies seeking to promote annuitization.

1. Introduction

What are the determinants of the decision to annuitize retirement savings? With over 31 million Americans expected to retire within the next ten years, this question is of great academic and practical interest. In defined benefit plans, employees traditionally receive at retirement lifetime income payments in the form of an annuity. In defined contribution plans, the most common type of retirement plan today, retirees have greater autonomy but directly bear the risk of outliving their retirement wealth.[1] By providing lifetime guaranteed income, annuities are a straightforward way to insure against this longevity risk.

Economists have been investigating annuitization decisions for almost 50 years, yet empirical evidence is still

[1] To understand the potential magnitude of this risk, we can look at the distribution of life expectancy at age 65 from the tables provided by the Social Security Administration. The difference between the 10th and 90th percentile of this distribution is equal to 22 years for men (dying at age 70 versus 92) and 23 years for women (dying at age 72 versus 95).

☆ An earlier version of this paper was titled "Stock Market Returns and Annuitization: a Case of Myopic Extrapolation." I would like to thank Shlomo Benartzi, Antonio Bernardo, Bruce Carlin and Mark Grinblatt for their useful comments. I am also grateful to Sebastien Betermier, Walid Busaba, Jeffrey Brown, James Choi, Timothy Conley, Dora Costa, Craig Dunbar, Amy Finkelstein, Steve Foerster, Andreas Fuster, Mark Garmaise, Hanno Lustig, Brian Melzer, Olivia Mitchell, Javier Perez-Estrada, Eduardo Schwartz, Albert Sheen, Stephan Siegel, Rossen Valkanov, Motohiro Yogo; participants at the Russel Sage Summer Institute in Behavioral Economics; seminar participants at UCLA, University of Washington, Notre Dame University, University of Western Ontario, HEC Montreal, University of Amsterdam, University of Mannheim; participants and discussants at 2010 Netspar International Pension Workshop, 2010 Northern Finance Association Meetings, 2011 SOBDR Conference at University of Toronto, 2011 Boulder Summer Conference on Consumer Financial Decision Making, 2011 Behavioral Finance Academy Annual Conference, 2011 CFS-EIEF Conference on Household Finance, 3rd Annual Behavioral Finance Conference at Queen's School of Business, 2012 ARIA Annual Meetings, University of Illinois at Urbana Champaign, 2013 AFA Meetings. For providing data used in this paper, I am indebted to an anonymous data provider; Federico Castellanos and Kathleen Roin from IBM; Matthew Drinkwater from Limra International; and Robert Shiller from Yale University.
* Tel.: +1 519 850 2417; fax: +1 519 661 3485.
E-mail address: aprevitero@ivey.ca

http://dx.doi.org/10.1016/j.jfineco.2014.04.006

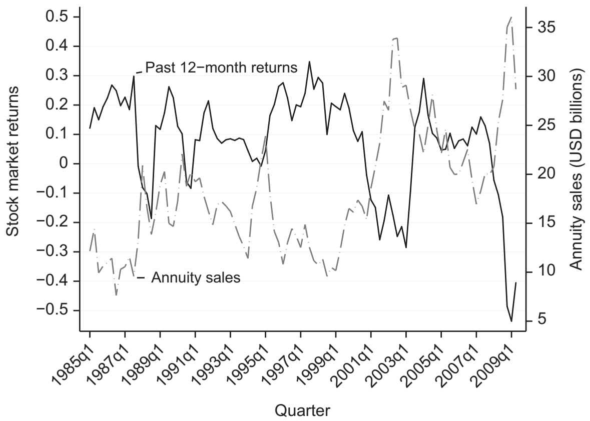

Fig. 1. Annuity sales and past stock market returns. This graph plots the past 12-month returns on the S&P 500 index (solid line) and the quarterly fixed annuity sales (dotted line). Annuity sales are deflated using the 2009 consumer price index (CPI).

limited (for a review, see Brown, 2007; Benartzi, Previtero, and Thaler, 2011). To help fill this void, I test the determinants of the relation between past stock returns and annuitization. In the main specification, I find that a one standard deviation variation in recent stock returns impacts the likelihood to annuitize by about 50% more than the effect of gender, the strongest cross-sectional predictor of annuitization in my sample. Further, I find that there is a strong negative correlation (-0.748) between fixed individual annuity sales and recent stock returns over the past 25 years (see Fig. 1). The robustness of this relation motivates efforts to understand the underlying causes. More specifically, I test several potential explanations for this relation between stock returns and annuitization. Given the worldwide trend of workers taking more and more responsibility in managing their retirement savings, understanding this particular relation is of utmost importance.

I analyze two novel data sets made available specifically for this study. First, I investigate the actual payout decisions of over 100,000 retirees enrolled in 112 different defined benefit plans across 63 different companies over seven years from 2002 to 2008. There is no enforced default option in any plan and each employee is required to make an explicit choice between a lump sum payment and a fixed annuity (i.e., an annuity whose payout is a fixed amount not contingent on stock market returns). Then, I study a defined benefit plan from IBM with over 18,000 actual retirement decisions between 2001 and 2009 and detailed information on the (financial) education of employees.

Three sets of results emerge from the analysis. First, I document that the strong negative relation between recent stock market returns and annuitization is robust to the inclusion of a host of different control variables: age, gender, tenure, benefit amount, interest rates, and fixed effects for retirement plans and metropolitan statistical areas (MSA) of residency. This effect is also economically significant; a change from the 25th to the 75th percentile of the past 12-month stock return distribution reduces the probability of selecting an annuity by about 10.4 percentage points (pp). Moreover, I find that the past 12-month stock returns largely drive the decision to annuitize and that returns older than two years have very limited influence.

Second, while several standard economic explanations might account for these findings, I find the most support for naïve beliefs and extrapolation from past returns. Using data on individual investors' beliefs on future stock market returns from the Yale Confidence Index, I find that a one standard deviation increase in the index implies a decrease in the probability of selecting an annuity by 9.8 pp. However, upon controlling for past returns, the effect of beliefs dramatically shrinks into statistical insignificance. These results suggest that past stock returns affect annuitization via beliefs. A falsification test using Confidence Index data on the beliefs of institutional investors rejects the hypothesis that more sophisticated beliefs about future returns affect annuitization.

Last, I document that the effect of past returns increases with age. The coefficient of past returns for the 60–69 and 70–75 age groups compared to the baseline group (age 50–59) is 2.0 and 5.2 times larger, respectively. These results are robust to the use of age quintiles instead of the previous cut-offs and to allowing for differences in the weighting of past returns across the three groups.

This paper connects to several strands of literature. My findings relate to the extensive theoretical literature on the annuity puzzle and portfolio choice with longevity products (see, e.g., Koijen, Van Nieuwerburgh, and Yogo, 2011; Cocco and Gomes, 2012). Empirical investigation of annuitization has been largely limited by the lack of actual micro-level data. Among the notable exceptions, Brown (2001) and Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) use intentions to annuitize collected through surveys. In the UK, where it is mandatory to annuitize retirement wealth, Finkelstein and Poterba (2004) analyze the selection among annuity types while Inkmann, Lopes, and Michaelides (2011) study the demand for (additional) voluntary annuities. These studies all share a focus on the cross-sectional determinants of annuitization such as life expectancy, marital status, bequest motives, or precautionary savings. I contribute to this literature by investigating the time-series determinants of annuitization. In particular, my focus is on identifying the specific mechanism by which past stock returns affect annuitization and, by extension, more accurately estimating its magnitude.

Studying the annuitization decision of Oregon state employees, Chalmers and Reuter (2012) also document a negative correlation between annuitization and the performance of the Standard & Poor's (S&P) 500 over the previous 12 months. My analyses reinforce the interpretation of their findings along two key dimensions. First, the main data set features annuities that are comparable with those available in the private annuity market (i.e., actuarially fair) and, consequently, my estimates are more likely to generalize to settings such as defined contribution plans where annuities are not subsidized. In contrast, Chalmers and Reuter find that the annuities offered to Oregon state employees are far more generous than those sold in the private annuity market (for the median retiree by 45%, for the average by 60%). This is one possible explanation for why their estimate of the effect of prior returns is only about 40% of my estimate[2] and why the annuity take up rate is higher (88% vs. 49% in my sample). Second, I attempt to control for potential omitted variables such as financial wealth. For example, positive stock returns can increase financial wealth (outside of retirement savings), reduce risk aversion, and make annuities less valuable [as theoretically shown in Mitchell, Poterba, Warshawsky, and Brown, 1999]. Therefore, failing to control for additional financial wealth can bias the estimates of the effect of past returns.

My results on extrapolation and age relate to the household finance literature on the effects of aging on financial decisions. Agarwal, Driscoll, Gabaix, and Laibson (2009) show that the ability to make financial decisions (in ten different types of credit behaviors such as credit card use) improves up to the early 50s before declining due to cognitive impairment. Analogously, Korniotis and Kumar (2011) find that older investors earn lower annual returns on average due to cognitive aging. A large psychology literature investigates the effects of aging on consumers' decision-making (for a review, see Yoon, Cole, and Lee, 2009; Drolet, Schwarz, and Yoon, 2010). Relevant to my results, different studies find that the elderly examine less information, consider fewer options when making choices (Cole and Balasubramanian, 1993; Besedes, Deck, Sarangi, and Shor, 2012), and experience a significant decline in explicit memory. My findings nicely complement these studies largely based on laboratory experiments by providing evidence that the elderly strongly rely on very limited and recent information (i.e., stock returns) to make actual retirement income decisions with serious welfare consequences.

This study also builds on research exploring the influence of past stock market returns in various settings: investors' beliefs and stockholdings (Vissing-Jorgensen, 2004); stock market participation and initial public offering (IPO) subscriptions (Kaustia and Knupfer, 2008, 2012); asset allocation (Benartzi, 2001; Benartzi and Thaler, 2007) and savings rates in 401(k) plans (Choi, Laibson, Madrian, and Metrick, 2009); mutual fund flows (Chevalier and Ellison, 1997; Sirri and Tufano, 1998; Bailey, Kumar, and Ng, 2011); and investments by young mutual fund managers (Greenwood and Nagel, 2009). My findings provide evidence that extrapolation from past returns can also influence a costly, irreversible decision. Moreover, I document a case of extreme myopic extrapolation in contrast to Benartzi (2001), who finds that employees would increase their company stock holdings after an increase in share price over the last ten years.

My stronger results for individuals at older ages suggest that beliefs might systematically change as people age. This latter finding provides a different perspective on the effects an aging population has on financial markets (Poterba, 2004). The traditional view, that individuals reduce their equity exposure as they age, can be only one side of the story if, for one example, the elderly also form their expectations in a systematically different way by giving higher weights to very recent returns. In addition to aiding our analysis of individual financial decisions, understanding these beliefs also has the potential to improve the formulation of asset pricing models based on micro-level evidence. As an example, Fuster, Hebert, and Laibson (2012) and Hirshleifer and Yu (2012) develop asset pricing models consistent with extrapolative expectations that are more strongly influenced by recent events.

The paper proceeds as follows. Section 2 describes the data and reports summary statistics. Section 3 introduces the empirical evidence on the relation between stock market returns and annuitization and documents its robustness. Section 4 is devoted to interpreting this evidence. Section 5 concludes.

2. Data summary and statistics

2.1. Retirement payout options in defined benefit plans

Defined benefit (DB) plans guarantee fixed benefits, typically based upon an employee's tenure at the company and pre-retirement level of income. While DB plans are compelled to offer participants the option to receive an annuity, some DB plans also offer a lump sum payout option. These plans are the focus of this paper.

The accrued benefits are usually defined in terms of an annuity beginning at the plan retirement age (typically age 65). The benefits are calculated by averaging the employee's earnings during the last few years of employment, taking a specified percentage of the average, and then multiplying it by the employee's number of years of service.[3] The lump sum distributions are determined by the present value of the future annuity payments to which the employee is entitled. The Internal Revenue Code (IRC) prescribes the interest rate and the unisex mortality table that the plan must use to determine the conversion from an annuity to a lump sum payment. A plan might decide to pay a larger lump sum, but is prohibited from paying less than the minimum amount derived under the IRC assumptions. In the empirical analyses, I include retirement plan fixed effects to account for this potential heterogeneity in the generosity of lump sums.

Note that while the annuity payments would only depend on pre-retirement income levels and tenure, the lump sum also depends on the interest rate used in the discount formula. For the majority of the sample period (2002–2007), the interest rate prescribed by the IRC was the rate on 30-year Treasury bonds. I include the long-term Treasury bond interest rate as a control variable in all analyses. The Pension Protection Act has revised the interest rate (changing it to a mix of short and long rates) and the mortality tables. As a consequence, starting in 2008, the value of lump sum payments will progressively

[2] A one standard deviation increase in the past stock returns reduces annuitization by 2.2 pp to 2.6 pp in Chalmers and Reuter's analyses vs. 6.4 pp in my baseline estimation in Table 2, Panel A.

[3] For example, let us consider an employee with average income before retirement of $3,000 per month and tenure of 40 years. If the plan promises 2.5% of the income for each year of tenure, the employee will be entitled to annuity payments of $3,000 per month.

decrease over the years. Year fixed effects are included in all empirical specifications to account for these changes.

Last, all employees in my sample work in private companies and the annuities are paid directly by those companies. In case of default, the defined benefits are guaranteed by the Pension Benefit Guaranty Corporation (PBGC) up to a maximum limit of roughly $50,000 during the sample period (e.g., $4,312.50 per month for plans terminating in 2008). Less than 12% of the employees in my sample will have benefits above the PBGC limits. Although I do not have information on the companies in the sample, the use of retirement plan (and retirement plan by year) fixed effects control for potential company heterogeneity in default risk.

2.2. Summary statistics

I investigate the relation between stock market returns and annuitization across two different samples: a large number of DB plans from an anonymous data provider (main sample) and a DB plan from IBM. In the Online Appendix 2.B, I also report results on individual annuity sales as collected by LIMRA International.

The main sample includes the actual payout decisions of over 103,000 employees enrolled in DB plans that offer the option to choose between an annuity and a lump sum. The payout decisions span seven years (2002–2008) and 112 different retirement plans offered by 63 different companies. While a company can offer more than one DB plan, the same plan cannot be offered by two different companies. Due to data collection issues and the addition of new plans over time, the panel of plans is unbalanced.[4] Therefore, all 112 plans are not observed for the entire seven-year period. At the employee level, I observe: (i) age, gender, tenure at the company, and zip code of residency; (ii) form of payout, benefit amount, and benefit start date; and (iii) identifiers for the retirement plan and company offering it.

The IBM data provide over 18,000 actual payout decisions from their DB plan. This data set is of particular interest for three reasons. First, these employees can choose partial annuitization (i.e., a mix of an annuity and lump sum) while options in the main data set are limited to one or the other. Second, this data set provides an additional downturn to analyze; it covers the nine years from 2000 to 2008 and thus includes decisions made after the Internet bubble. Last, I observe additional demographic information such as income before retirement and detailed information on education.

While retirees' health and life expectancy are very important determinants of the decision to annuitize (see, for example, Finkelstein and Poterba, 2004), data limitations prevent me from explicitly addressing them. Since the data are derived from payroll information, I only observe the longevity for the employees that take the annuity, as those that take the lump sum disappear from the data. My focus on the time-series determinants of annuitization (more than the cross-sectional determinants) make these data limitations less problematic. The results on the effect of stock market returns will be affected by health/mortality considerations only in the case that health and subjective life expectancies systematically vary with trends in the stock market. While this is an intriguing possibility, I believe that the identification of such a relation is beyond the scope of this paper.[5]

Table 1 contains some descriptive statistics. In the main sample, 49% of employees select an annuity. In the IBM plan, 88% of employees make a similar choice, 6% choose a lump sum, and the remaining 6% select a mix of the two. Among the 112 plans covered in the main sample, some plans present a high annuitization rate similar to the one found in the IBM plan. Considering the wide dispersion of annuitization rates across plans, I use retirement plan fixed effects to control for (non-time-varying) unobservable features of the plans that can drive the decision to annuitize.

The total wealth of employees is not observed, only the retirement benefits in the specific DB plan.[6] The average age and tenure of employees is, respectively, 60 and 25 years in the main data set and 58 and 29 years in the IBM data set. Substantial tenure at retirement suggests that the payout decision involves a significant fraction of their total retirement wealth. To support this claim, I use data from three waves (2001, 2004, and 2007) of the Survey of Consumer Finances (SCF) from the Board of Governors of the Federal Reserve System. Observations are weighed using SCF sample weights to represent the U.S. population. The mean and median values of retirement benefits in the main sample ($188,130 and $86,460, respectively) are of the same order of magnitude as the average and median net financial wealth for retirees in the SCF ($259,200 and $15,360, respectively). While the DB benefits in my sample are measured at the employee level, the data from the SCF are reported at the household level. All these considerations extend to the IBM sample, in which employees have significantly higher benefits.

To control for the omission of total wealth in the results, I use the median house prices at the MSA level. Looking at the data from the SCF, we can see how home equity represents a large fraction of total wealth of respondents: about 40% of the sum of home equity and net financial wealth ($167,460 plus $259,200). Therefore, including DB benefit amounts and median house prices in the regressions is likely to significantly reduce the omitted variable bias.

[4] To check if missing data for some plans or additional plans bias the results, I run all the analyses in the paper using only data from plans that are in the sample for at least four years. All my results are confirmed in this subsample.

[5] Nonetheless, I use year fixed effects in all main specifications to capture unobservable variables common to employees at the year level. As a consequence, my results are identified using within-year variations and not driven by any particular year where extreme stock returns might have caused a potential shift in life expectancies.

[6] I cannot exclude that some of these employees are also enrolled in a defined contribution (DC) plan offered by the same employer. However, when both plans are offered, the matching contributions of the employer to the DC plan are generally limited and so are the voluntary contributions made by employees.

Table 1
Summary statistics.

Annuity is a binary variable equal to one if the annuity is chosen. Age is equal to the age of the employee at the benefit start date. Female is an indicator variable equal to one if the employee is female. Tenure is the number of years the employee has worked for the company. DB benefits is the total amount of the benefits accrued to the employee. Median house price is the median house price in the Metropolitan Statistical Area (MSA) of residency at the moment of separation. I can match house prices only for 70,587 employees in the sample. Income is the total yearly income for the employee in the year of separation. Years of education represents years of education completed. Business education is an indicator variable equal to one if the employee obtained any financial education (either a bachelor's degree with a major in economics or accounting, or a master's degree in economics or in business administration). MBA is an indicator variable equal to one if the employee obtained a master's degree in business administration. Sample size is the total number of observations for each sample. The data from the Survey of Consumer Finances (SCF) represent averages across three waves of the survey (2001, 2004, and 2007). Observations are weighed using SCF sample weights, so as to be representative of the U.S. population. Data in columns 5 and 6 are from respondents that are retired and less than 75 years old. Data in columns 7 and 8 are from respondents with age between 50 and 75 years, regardless of their retirement status. Age and Female refer to the person answering the survey. Net financial wealth and Home equity are measured at the household level.

		Main sample		IBM		SCF (retirees)		SCF (age 50–75)	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Unit	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Annuity	0–1	0.49	0.00	0.88	1.00				
Age	1	59.83	60.00	58.33	57.86	70.20	70.13	60.63	59.60
Female	0–1	0.44	0.00	0.25	0.00	0.32	0.00	0.27	0.00
Tenure	1	24.52	25.66	28.92	30.59				
DB benefits	$1,000	188.13	86.46	413.04	387.1				
Net financial wealth	$1,000					259.2	15.36	262.79	6.72
Median house price	$1,000	213.33	166.10						
Home equity	$1,000					167.46	97.37	163.89	87.4
Income	$1,000			101.07	95.93				
Years of education	1			15.22	16				
Business education	0–1			0.12	0				
MBA	0–1			0.06	0				
Sample size	1	103,516	103,516	18,671	18,671	932	932	5,835	5,835

3. Empirical results

3.1. Stock market returns and annuitization

Table 2, Panel A reports the results of analyses using the defined benefit plans sample. In Panel B, I introduce similar estimations using data from the IBM retirement plan. Results from the latter are discussed in the following subsection. In both panels, I use the cumulative monthly stock returns R_{t-k} for different period lags (expressed in months) prior to the decision date. Specifically, I estimate:

$$Ann_{ijt} = \alpha + \beta \sum_{k=1}^{lag-1} R_{t-k} + \gamma' x_{it} + \delta' z_{jt} + \xi' t_t + \varepsilon_{it}. \quad (1)$$

The vector x_{it} of individual control variables includes age, gender, benefit amount, and tenure. The vector z_{jt} of time-varying plan control variables consists of the average of age, gender, and benefit amount, and the number of employees separating in a given year for each plan. The vector t_t of time-varying controls includes long-term interest rates and year and calendar-month fixed effects.

To proxy for interest rates, I use the long-term composite rate on Treasury bonds.[7] In practice, this rate is averaged over the six months before the separation date. As specified in Section 1, this is a good proxy for the discount rate that the employers are required to use in the conversion between the annuity and the lump sum. I use calendar-month fixed effects to account for potential seasonalities in payout forms (e.g., some plans might allow particular payout forms only in specific periods). The use of year fixed effects mitigates also the concern that the number of plans varies across years. The coefficients of the control variables have the signs that one would expect given the previous literature.

The coefficients on stock returns and on interest rates in Table 2 (and in the following tables) are standardized so as to compare the magnitude of the effect on annuitization across different variables. For example, in column 1 a one standard deviation increase in the past stock market reduces the probability of choosing an annuity by roughly 6.4 pp, a statistically significant result. To put the economic magnitude of this coefficient in perspective, note that one year of age increases the likelihood of annuitizing by 2.4 pp; being female increases the likelihood of selecting the annuity by about 4.1 pp; and an increase in the benefit amount of $100,000 increases the likelihood of annuitization by 3.3 pp. The effect of past returns appears stronger if we look at the past two years (col. 2). In untabulated evidence using non-overlapping windows (0–12 and 13–24 months), I find that the effect of past returns is largely driven by the past 12 months. Consistently, this effect becomes insignificant when the window is expanded to three years (col. 3). In the Online Appendix Table A.1, I refine the analysis using the weighting function of past returns adapted from Malmendier and Nagel (2011) and obtain similar results: almost no weight is assigned to returns older than two years.

In Table 2, the coefficient on interest rates is not significant in any of the specifications. If annuities are

[7] The Long-Term Composite Rate is the unweighted average of bid yields on all outstanding fixed-coupon Treasury bonds with maturity older than ten years.

Table 2
Stock market returns and annuitization.

This table reports results from OLS regressions. The dependent variable is a binary indicator which equals one if the employee chooses an annuity. Interest rates is the composite return on long-term Treasury bonds. In Panel A, additional and unreported controls include: plan controls (yearly average age, gender, benefits, and number of employees retiring in each plan); and calendar-month fixed effects. Standard errors are clustered across 15 company size/time groups. In Panel B, standard errors are clustered across eight geographical region/time groups. See Section 3.1 for more details on this methodology and the use of Linear Probability Models. All the coefficients can be interpreted as the percentage point variation in the probability of annuitization corresponding to a one-unit change in the independent variable. I use the following units: (i) one standard deviation for returns and interest rates; (ii) $10,000 for benefit amount and income; and (iii) one year for age, tenure and years of education.

Panel A: Defined benefit plans

	(1)	(2)	(3)	(4)	(5)
Past 12-month returns	−6.363***			−5.070***	−4.690***
	(1.723)			(1.429)	(1.488)
Past 24-month returns		−7.141**			
		(2.434)			
Past 36-month returns			−0.422		
			(0.479)		
Interest rates	−2.227	−1.511	−2.626	−0.120	0.076
	(2.344)	(2.410)	(2.032)	(1.915)	(1.835)
Female	4.134***	4.118***	4.119***	4.029***	3.955***
	(0.956)	(0.976)	(1.002)	(1.018)	(1.038)
Age	2.417***	2.436***	2.474***	2.418***	2.413***
	(0.250)	(0.263)	(0.288)	(0.256)	(0.258)
Benefits amount	0.330***	0.331***	0.333***	0.336***	0.329***
	(0.071)	(0.072)	(0.072)	(0.076)	(0.078)
Tenure	−0.192	−0.193	−0.198	−0.186	−0.179
	(0.143)	(0.144)	(0.145)	(0.150)	(0.152)
Plan controls	Yes	Yes	Yes	Yes	
Years F. E.	Yes	Yes	Yes	Yes	
Plan F.E.				Yes	
Plan/Year F.E.					Yes
Observations	103,516	103,516	103,516	103,516	103,516
R-squared	0.194	0.190	0.186	0.391	0.412

Panel B: IBM retirement plan

Sample:		All employees		Bus. education	MBA
	(1)	(2)	(3)	(4)	(5)
Past 12-month returns	−1.327***			−2.832**	−2.637*
	(0.310)			(0.922)	(1.244)
Past 24-month returns		−1.573**			
		(0.510)			
Past 36-month returns			−1.914**		
			(0.610)		
Interest rates	3.388***	3.702***	−4.291***	3.127***	4.696***
	(0.565)	(0.611)	(0.798)	(0.747)	(0.793)
Female	3.839***	3.864***	3.900***	1.073	−1.141
	(0.457)	(0.451)	(0.438)	(1.943)	(2.161)
Age	−1.697***	−1.683***	−1.667***	−1.594***	−1.976***
	(0.064)	(0.067)	(0.069)	(0.138)	(0.307)
Benefits amount	0.308***	0.309***	0.308***	0.323***	0.347***
	(0.058)	(0.059)	(0.058)	(0.056)	(0.073)
Tenure	0.550***	0.551***	0.555***	0.444**	0.219
	(0.101)	(0.100)	(0.099)	(0.170)	(0.211)
Income	−0.730***	−0.727***	−0.717***	−0.727***	−0.626***
	(0.108)	(0.108)	(0.104)	(0.126)	(0.102)
Years of education	−0.319*	−0.319**	−0.314**	−0.875	
	(0.135)	(0.130)	(0.127)	(0.537)	
Observations	18,671	18,671	18,671	2,271	1,062
R-squared	0.132	0.133	0.133	0.128	0.152

Standard errors in parentheses. Constant included.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

fairly priced, individuals can fully incorporate the impact of interest rates on the lump sum's value in their choice between annuities and lump sums. In practice, the value of cash holdings could also change in response to an increase in interest rates, netting out the effect of receiving a lower lump sum payment. An alternative interpretation is that employees might not respond to changes in interest rates because they do not have a good understanding of the factors determining the conversion between annuities and lump sums (Chalmers and Reuter, 2012).

I model the decision to annuitize with a linear probability model (i.e., ordinary least squares (OLS) estimation). Despite a binary dependent variable, I prefer this choice over the commonly used logit or probit models for three reasons: two substantial and one formal. First, I can use fixed effects while avoiding the incidental parameters problem. Although the data structure is pooled cross-sections and not a panel, payout decisions of different employees are likely to be dependent over time (e.g., peer effect within the same plan). Second, I can directly obtain unbiased coefficients for interaction terms (Ai and Norton, 2003). Last, this choice makes it easier to directly assess economic magnitudes by simply multiplying β with the variation in the past returns. As robustness check, I have also estimated Eq. (1) using a logit model. The results do not materially change.

To account for cross-sectional and intertemporal dependence in the data, I cluster the standard errors in Table 2 across 15 company size/time groups. More precisely, I partition the data into three 28-month periods and company size quintiles, proxied for by the number of employees separating from each company in the sample period. I derive this approach from Bester, Conley, and Hansen (2011).[8]

In columns 4 and 5, retirement plan fixed effects are added to control for unobservable (non-time-varying) plan characteristics that can influence employee decisions. For example, the inclusion of these fixed effects will help account for potential heterogeneity in terms of the generosity of lump sums. In other words, I can be sure that the results are not driven by those plans that offer more (or less) generous lump sums. These plan fixed effect specifications would not fully account for the case where plans would change how generous the lump sums are during the sample period. To account for this possibility, retirement plan by year fixed effects are added in column 5. Under this empirical specification, the effect of past returns is still economically and statistically significant. The results are also robust to the use of MSA fixed effects to control for unobservable variables such as MSA-level differences in ethnicity and demographic variables that can be correlated with life expectancy and hence affect annuitization.

3.2. Stock returns, annuitization, and financial education

In Table 2, Panel B, I report results for the IBM retirement plan. In this specification I add income and years of education to the vector of individual control variables x_{it}. The results in columns 1–3 show a strong effect of past stock returns. For example, a one standard deviation increase in the past 12-month returns reduces the probability of annuitizing by about 1.3 pp. There are two major differences between these results and the ones reported in Panel A. First, the effect of past returns is weaker and does not decline as quickly over time. Although weaker, this effect is still economically significant. For comparison, five additional years of education decreases the likelihood of selecting an annuity by about 1.6 pp. Second, the effect of interest rates is now economically and statistically significant. The estimates imply that a one standard deviation increase in interest rates increases the probability of annuitizing by 3.4 pp. Recall from Section 2.1 that while an increase in interest rates does not impact the annuity payments, it will instead reduce the value of the lump sum (by increasing the discount factor). This evidence indicates that employees not only consider current market conditions but also how the value of the lump sum changes over time. One possible explanation is that employees might use previous lump sum payouts as a reference point when making their decision (e.g., comparison with peers retiring in the past).

In columns 4–5, I investigate if financial education makes employees less sensitive to very recent stock market returns. In column 4, sample is restricted to employees that have obtained financial education in the form of a bachelor's degree in economics or accounting or a master's degree in economics or business administration. In column 5, I analyze only employees that have obtained an MBA. Results from both samples suggest that employees with any business education are just as likely to be influenced by past returns.[9] While stock returns are in general more salient to individuals with higher financial education, the most provocative interpretation of this evidence is that financial education can increase confidence about the ability to time and predict the stock market. Not having information on how employees invest their additional wealth limits the ability to provide more compelling evidence to interpret this result. Overall, the IBM results clearly demonstrate how heterogeneity across plans can change the estimates and highlight the need for the inclusion of retirement plan fixed effects in my specifications.

Following the methodology used in Panel A, I cluster here all the standard errors across eight geographical region/time groups. All the above results are also robust to the inclusion of working location fixed effects which account for potential co-worker effects or unobservable variables common at the working location (e.g., same job qualifications).

4. Interpretation of the evidence

4.1. Wealth effects: evidence from Hurricane Katrina

I do not observe the overall wealth of employees but only their retirement plan benefits. Therefore, the results

[8] Bester, Conley, and Hansen (2011) provide simulation evidence that cluster covariance estimators, as used in my analyses, outperform conventional inference procedures when the data exhibit cross-sectional and temporal dependence. For more details on this methodology, refer to the Online Appendix 1.B.

[9] I also run specifications for the entire sample with an indicator variable for business education or MBA and the interaction between these variables and past returns. In untabulated one-sided t-tests, I can reject the null that these interaction coefficients are greater than or equal to zero (i.e., employees with financial education are less likely to rely on past returns) with p-value$=$0.088 (for business education) and p-value$=$0.118 (for MBA).

potentially suffer from an omitted variable bias (see Online Appendix 1.C). If annuitization decreases as wealth rises, the estimates will be biased upward. In the opposite scenario, the estimates of past returns on annuitization will be, if anything, too conservative.

From a theoretical perspective, a wealth shock can either decrease or increase annuitization. Mitchell, Poterba, Warshawsky, and Brown (1999) show that more risk-averse people should be willing to pay more for annuities. With wealth-dependent risk aversion, as wealth increases (and risk aversion decreases), employees should value an annuity relatively less. However, bequests and precautionary motives can influence the decision to annuitize (Bernheim, 1991; Sinclair and Smetters, 2004; Ameriks, Caplin, Laufer, and Van Nieuwerburgh, 2011). If employees avoid annuitization to bequeath or to better handle liquidity needs such as health shocks, an increase in wealth might actually attenuate liquidity concerns and increase the likelihood of annuitization.

Which of these two effects prevails is an empirical question. I follow two different approaches to address this challenge. First, I use data on county-level damage from Hurricane Katrina as a proxy for exogenous shocks to wealth. Then, I use house price appreciation across different MSAs to proxy a wealth shock.

In August 2005, Hurricane Katrina caused more than 1,800 deaths and an estimated $81 billion in total property damage, primarily concentrated in Florida, Mississippi, Alabama, and Louisiana (Knabb, Rhome, and Brown, 2005). Even though the Gulf area has witnessed several hurricanes over the years, Katrina was unprecedented in terms of damages caused.[10] For this reason, I use this event as a proxy for an exogenous shock to the wealth of the employees living in that area at the time of retirement.

I use a differences-in-differences methodology to estimate the causal effect on annuitization of a shock to wealth due to the hurricane. Table 3 reports these results. In column 1, I estimate the same model in Eq. (1) with the addition of three explanatory variables: (i) "After Katrina," equal to one after the hurricane; (ii) "Katrina areas," equal to one for the counties afflicted by the hurricane[11]; and (iii) their interaction. This interaction represents the differential effect of an exogenous shock to wealth on annuitization between the "treated" group (i.e., employees living in the afflicted areas) and the control group (employees living in counties not afflicted). In column 1, this coefficient is economically and statistically significant; the hurricane decreases the likelihood of selecting an annuity by 9.6 pp.[12]

In column 2, I test if immediate liquidity needs are driving the results. In some areas, such as the entire state of Louisiana, the shock to wealth could have generated immediate and stringent liquidity needs and forced the employees into selecting the lump sum. I check for this possibility by excluding from the analysis all the employees (1,457) that are located in Louisiana. As expected, the coefficient of interest, the indicator for a hurricane-affected area, is lower than what was found earlier. Nonetheless, it remains economically and statistically significant; employees retiring after the hurricane in the counties affected outside of Louisiana are still 8.3 pp less likely to select an annuity.

The hurricane could have also impacted the companies operating in that area and increased their bankruptcy risk. Under this scenario, my results of higher demand for lump sums after the hurricane could be explained by employees perceiving annuities as riskier due to an increase in the default probability of their companies. Given that annuities in defined benefit plans are guaranteed up to a limit by the PBGC (see Section 2.1), the risk of insolvency of the annuity provider (the companies) is relevant only if we assume that employees: (i) have benefits higher than the PBGC limit (roughly 10% of the sample falls in this category); (ii) ignore the guarantee provided; or (iii) believe that the PBGC itself might become insolvent after the hurricane.

Since the identity of the companies is not known, I proxy for their exposure to the hurricane area by using the fraction of employees retiring in the areas afflicted by the hurricane. The rationale for this control variable is that the higher the proportion of employees in the Katrina area, the higher the business concentration of the companies in that area. Column 3 reports that the difference-in-difference estimates are unchanged if I control for this business exposure to the Katrina area. In the main specification I use the fraction of employees in the Katrina area during the entire sample period. As robustness checks, I also calculate this fraction before the event or only during the event period (2005–2006). My main results remain unchanged. In column 4, I show that the results are also robust to the exclusion of Louisiana employees from the sample to account for their immediate liquidity needs.

In a similar manner, the results could also be driven by employees negatively revising their life expectancy after the hurricane and, hence, finding the annuities less attractive. Even after controlling for this possibility (see Online Appendix 2.C), the results do not materially change. All this evidence from the exogenous shock to wealth caused by Hurricane Katrina suggests a positive relation between wealth shocks and annuitization (in this case, a negative shock reduces annuitization).

4.2. *Wealth effects: evidence from real estate prices*

In Table 4, I use real estates prices to proxy for wealth shocks. Both Hurst and Lusardi (2004) and Lusardi and Mitchell (2007) use variation in house prices as an instrument for wealth. After matching median house prices by MSA (from the National Association of Realtors) with payout data from the main sample, I obtain a final data set of 58,897 observations accounting for about 57% of the original observations. In column 1 of Table 4, I show that

[10] The second costliest Atlantic hurricane, Andrew in 1982, caused less than half of the total property damage (in 2005 US dollars) of Katrina.

[11] I use county damage estimates from the U.S. Department of Housing and Urban Development (HUD) for the 2006 hurricane season.

[12] Adapting the methodology previously used, the standard errors here are clustered across 12 region/time groups, obtained by combining the four US Census Regions and three 28-month periods. For additional details, refer to the Online Appendix 1.B.

Table 3
Wealth effects: evidence from hurricane Katrina.

This table reports results from OLS regressions, using only data from the defined benefit plans sample. The dependent variable is a binary indicator which equals one if the employee chooses an annuity. After Katrina is an indicator variable equal to one after 9/2005. Katrina areas is an indicator variable equal to one if the employee lived in a county afflicted by the hurricane at the moment of separation. Concentration in Katrina areas is the fraction of employees separating in the areas afflicted over the total number of separating employees in the same company. Additional controls include: (i) demographic controls (age, gender, tenure, and benefit amount); (ii) interest rates (the composite return on long-term Treasury bonds); (iii) calendar-month and year fixed effects; (iv) time-varying plan controls. Standard errors are clustered across 12 geographical region/time groups (see Section 3.1 for more details on this methodology and the use of Linear Probability Models). All the coefficients can be interpreted as the percentage point variation in the probability of annuitization corresponding to a one-unit change (for indicator variables) or a one standard deviation change in the corresponding independent variable.

Sample:	All states (1)	All states (without LA) (2)	All states (3)	All states (without LA) (4)
Past 12-month returns	−6.638***	−6.574***	−6.576***	−6.511***
	(1.937)	(1.959)	(1.907)	(1.931)
After Katrina	3.930	3.602	3.852	3.525
	(3.258)	(3.095)	(3.245)	(3.078)
Katrina areas	6.539**	5.354**	6.300**	5.126*
	(2.529)	(2.197)	(2.777)	(2.526)
After Katrina*Katrina areas	−9.618***	−8.291***	−9.562***	−8.228***
	(2.356)	(2.167)	(2.245)	(2.062)
Concentration in Katrina areas			0.335	0.339
			(2.897)	(2.905)
Additional controls	Yes	Yes	Yes	Yes
Observations	96,081	94,624	96,081	94,624
R-squared	0.199	0.196	0.199	0.196

Standard errors in parentheses. Constant included.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

Table 4
Wealth effects: evidence from real estate prices.

This table reports results from OLS regressions, using only data from the defined benefit plans sample. The dependent variable is a binary indicator which equals one if the employee chooses an annuity. The median house prices are collected at the MSA where the employee lived at separation. Variations in median house prices are computed from the lagged value until the time of retirement. Additional controls include: (i) demographic controls (age, gender, tenure, and benefit amount); (ii) interest rates (the composite return on long-term Treasury bonds); (iii) calendar-month and year fixed effects; (iv) time-varying plan controls. Standard errors are clustered across 15 company size/time groups (see Section 3.1 for more details on this methodology and the use of Linear Probability Models). All the coefficients can be interpreted as the percentage point variation in the probability of annuitization corresponding to a one standard deviation change in the corresponding independent variable.

Lag in median house prices:	(1)	1 Year (2)	2 Years (3)	3 Years (4)	1 Year (5)
Past 12-month returns	−7.302**	−7.316***	−7.109***	−7.288**	−5.883**
	(2.465)	(2.286)	(2.309)	(2.273)	(2.120)
Lagged median house price		−3.992***	−4.292***	−4.307***	−1.492**
		(0.881)	(0.876)	(0.806)	(0.560)
Δ Median house price		2.153*	1.784**	0.981	1.381**
		(1.022)	(0.728)	(0.854)	(0.539)
Additional controls	Yes	Yes	Yes	Yes	Yes
Plan F.E.					Yes
Observations	58,897	58,897	58,897	58,897	58,897
R-squared	0.169	0.176	0.176	0.176	0.378

Standard errors in parentheses. Constant included.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

the estimates of the past 12-month returns coefficient in this smaller sample are similar to the previous ones (see Table 2, Panel A, column 1).

Column 2 shows that the coefficient of past returns remains statistically and economically significant after controlling for levels and variations in median house prices. Since all the coefficients are standardized, I can directly compare across them. A one standard deviation increase in past returns implies a decrease in the probability of choosing an annuity by about 7.3 pp. Both variables related to real estate prices have a non-negligible effect on annuitization. A one standard deviation increase in the one-year lag of median house prices (about $125,600) reduces the likelihood of choosing the annuity by about 4.0 pp. A similar increase in the past 12-month appreciation of real estate values (about 11 pp) implies an increase in the likelihood of annuitizing by about 2.2 pp. As in Table 2, Panel A, standard errors are clustered across 15 company size/time groups by partitioning the data into company size quintiles and three 28-month periods.

These results highlight the importance of jointly controlling for levels and variations in real estate values. The coefficient of levels of house prices is driven by cross-sectional variations across MSAs and tells us the employees living in areas with higher prices — and therefore more likely to be wealthier — are less likely to choose an annuity. Using data from defined contribution plans from the Health and Retirement Study, Brown (2001) finds a similar (small) negative relation between annuitization and financial net worth. Two of the possible explanations for this relation presented by the author are that wealthier individuals might have less need for the insurance offered by the annuity or that they believe they can earn higher returns themselves. Both are plausible explanations for this result in my data. In the analysis of the IBM data, I indeed find a (small) statistically significant negative relation between annuitization and both income and education (see Table 2, Panel B).

The coefficient of variation of house prices is driven by time-series variation in prices: for a given level of real estate prices, employees that have experienced higher increases in prices are more likely to take an annuity. Therefore, precautionary motives also seem relevant in explaining the decision to annuitize. This positive relation between variation in wealth and annuitization provides evidence against the potential explanation that an increase in wealth — caused by stock market returns — is driving the main results.

In columns 3 and 4, I lengthen the time period used to control for levels and variation in median house prices, using two and three years respectively. Similar to the findings for stock returns, the effect of the variation in house prices on annuitization decreases going back in time and is not significant after three years.[13] In column 5, I confirm the results including retirement plan fixed effects. Overall, these results and previous evidence on the exogenous wealth shock caused by Hurricane Katrina highlight that wealth effects are not the likely driver of the relation between past returns and annuitization, and that failing to control for outside wealth is not biasing my estimates.

4.3. *Extrapolation from recent stock market returns*

Recent stock market returns can affect beliefs about future returns. After negative returns, employees might believe that this trend will continue into the future and consequently find the annuity, essentially a fixed-income financial product, more attractive. The opposite can happen after a positive trend in the market.

In Table 5, I test for this possibility using a measure of investor beliefs about future returns, the Confidence Index. This index corresponds to the percentage of individual investors expecting an increase in the Dow Jones (Industrial) in the coming year.[14] In column 1, I estimate the baseline model from Eq. (1), replacing past stock market returns with the six-month average of the Confidence Index. A one standard deviation increase in the index implies a 9.8 pp decrease in the probability of selecting an annuity. This result is not only statistically significant, but also comparable in magnitude with the effect found for past returns in Table 2, Panel A. Note that the standard errors in Table 5 are clustered across the 15 company size/time groups to allow for the same type of cross-sectional and serial correlation assumed earlier.

In column 2, I also control for the past 12-month returns. In this specification, the effect of beliefs about future returns dramatically shrinks and it is not statistically significant. This evidence suggests that the effect of beliefs on annuitization is mainly driven by previous returns. In column 3, using retirement plan fixed effects, I confirm that this finding is robust not only across but also within retirement plans.

In columns 4–6, a falsification (or placebo) test is conducted using data from a confidence index computed in the same manner but with answers from institutional investors. More sophisticated beliefs from institutional investors have a marginal effect on annuitization; this coefficient in column 4 is of the expected sign but smaller in magnitude and noisily estimated. Nonetheless, this effect is significant in explaining annuitization within retirement plans (see column 6 where plan fixed effects are added) and it is not reduced when I add past returns (see column 5). Beliefs from institutional investors contain information useful in understanding annuitization. Different than individual investor beliefs, this information is not easily captured by previous returns. As a simple test, I run a regression of the confidence index on past stock returns. The coefficient is significant for individual investors and not significant for institutional investors. The evidence in Table 5 suggests that past stock returns affect annuitization by changing beliefs. This result makes extrapolation a plausible interpretation of the evidence in my data. In a controlled laboratory environment, Agnew, Anderson, and Szykman (2012) also find evidence consistent with extrapolation from past returns influencing the decision to annuitize.

In the Online Appendix Sections 2.D, 2.E, and 2.F, I check the robustness of the results to several additional factors: (i) endogenous timing of retirement; (ii) stock market volatility; (iii) expectations about labor income; and (iv) expectations about inflation. These additional results do not change the main interpretation of the evidence that individuals extrapolate from past returns into the future.

[13] In additional analyses not tabulated, I find similar non-significant results for horizons of four and five years.

[14] The Confidence Index data are available on Robert Shiller's Web site (http://www.econ.yale.edu/~shiller/data.htm).

Table 5
Beliefs about stock market returns and annuitization.

This table reports results from OLS regressions, using only data from the defined benefit plans sample. The dependent variable is a binary indicator which equals one if the employee chooses an annuity. Confidence index is the (six-month average) percentage of investors expecting an increase in the Dow Jones Industrial in the coming year. Columns 1–3 report this percentage for individual investors; columns 4–6 for institutional investors. Additional controls include: (i) demographic controls (age, gender, tenure, and benefit amount); (ii) interest rates (the composite return on long-term Treasury bonds); (iii) calendar-month and year fixed effects; (iv) time-varying plan controls. Standard errors are clustered across 15 company size/time groups (see Section 3.1 for more details on this methodology and the use of Linear Probability Models). All the coefficients can be interpreted as the percentage point variation in the probability of annuitization corresponding to a one standard deviation change in the corresponding independent variable.

Sample:	Individual investors			Institutional investors		
	(1)	(2)	(3)	(4)	(5)	(6)
Confidence index	−9.803**	−2.518	−3.328	−3.217	−2.403	−2.974*
	(4.570)	(5.284)	(4.637)	(1.930)	(1.591)	(1.395)
Past 12-month returns		−4.858***	−3.871***		−5.386***	−4.577***
		(1.192)	(1.118)		(1.268)	(1.105)
Additional controls	Yes	Yes	Yes	Yes	Yes	Yes
Plan F.E.			Yes			Yes
Observations	103,516	103,516	103,516	103,516	103,516	103,516
R-squared	0.186	0.189	0.389	0.184	0.189	0.390

Standard errors in parentheses. Constant included.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

4.4. *Does the effect of extrapolation increase with age?*

This subsection looks more closely at the behavior of older retirees, the natural target of any policy intervention to provide retirement income solutions. The sample includes individuals aged 50–75. The effect of past returns can be uniform across age groups or vary with age. For example, individuals in their fifties could drive the results as they try to time the market and increase their retirement wealth. Alternatively, extrapolation might increase with age if the elderly are more prone to rely on less information due to memory loss and cognitive impairment (Cole and Balasubramanian, 1993; Besedes, Deck, Sarangi, and Shor, 2012). The interpretation of the results and their policy implications are substantially different between these two scenarios.

In Table 6, I investigate the effect of stock market returns for different age groups in the sample. Two age dummies are introduced for individuals in their sixties (with age between 60 and 69) and in their seventies (with age between 70 and 75, the sample cut-off). The data lack information on factors relevant to the decision to annuitize such as health status and longevity expectations that are likely to vary across these different age groups. The omission of these factors necessitates caution when interpreting the Age 60 and Age 70 indicator variables, but it is less problematic for the analyses on extrapolation and age (i.e., the interaction term).[15]

[15] See a general discussion about this issue in Section 2.1. An outstanding concern can be that individuals that decide to retire at age 50 can be systematically different from retirees in their 60s or 70s. I deal with this issue later in this section and in the Online Appendix where I

Table 6
Extrapolation from past returns and annuitization at older ages.

This table reports results from OLS regressions (col. 1 and 2) and non-linear least squares (col. 3–5), using only data from the defined benefit plans sample. The dependent variable is a binary indicator which equals one if the employee chooses an annuity. Additional controls include: (i) demographic controls (gender, tenure, and benefit amount); (ii) interest rates (the composite return on long-term Treasury bonds); (iii) calendar-month and year fixed effects; (iv) time-varying plan controls. Standard errors are clustered across 15 company size/time groups (see Section 3.1 for more details on this methodology and the use of Linear Probability Models). All the coefficients can be interpreted as the percentage point variation in the probability of annuitization corresponding to a one-unit change (for indicator variables) or a one standard deviation change in the corresponding independent variable.

	(1)	(2)
Past 12-month returns	−4.268**	−3.232**
	(1.513)	(1.412)
Age 60	20.916***	21.240***
	(2.416)	(2.279)
Age 70	11.969***	18.673***
	(3.945)	(3.683)
Age 60∗Past 12-month returns	−4.015*	−4.047**
	(2.209)	(1.820)
Age 70 ∗Past 12-month returns	−18.003***	−13.366***
	(3.866)	(1.850)
Additional controls	Yes	Yes
Plan F.E.		Yes
Observations	103,516	103,516
R-squared	0.186	0.384

Standard errors in parentheses. Constant included.
* Significant at 10%.
** Significant at 5%.
*** Significant at 1%.

The effect of stock market returns strongly increases with age. For individuals in their sixties, a one standard deviation increase in past 12-month returns decreases the likelihood of annuitizing by 8.3 pp (4.3+4.0), an effect twice as big as the baseline group. The effect becomes even more dramatic for individuals in their seventies; a similar variation in returns translates to a 22.3 pp (4.3+18.0) decrease in the probability to annuitize, equivalent to 5.2 times the effect in the baseline group. In column 2, I add plan fixed effects to control for time-invariant unobservables at the retirement plan level. In this specification, previous results are substantially confirmed with the effect of extrapolation increasing by a factor of 5.2 from individuals in their fifties to those in their seventies (3.2 pp vs. 16.6 pp).

Similar results are achieved if I divide the sample into age quintiles instead of using arbitrary age cut-offs. In untabulated results, I find that the effect of stock returns monotonically increases with age; a one standard deviation variation is 3.3 times larger moving from the first quintile, age 50–55, to the last quintile, age 65–75 (2.3 pp vs. 7.6 pp).

I identify the effect of age on the tendency to extrapolate from past returns by comparing cross-sections of employees retiring in their fifties, sixties, and seventies. People retiring later in life can be different along several dimensions. For example, they might need to work longer because they are less successful or because they have higher spending needs. My previous results from the IBM data (see Table 2, Panel B) document that financial education does not mitigate the tendency to extrapolate. This evidence is inconsistent with the alternative explanation that people working longer might be less financially sophisticated and more prone to extrapolate, independent of any effects aging may have.

As a robustness check, I run analogous analyses with the IBM data. In the Online Appendix 2.G, I estimate the baseline model separately for employees retiring before and after their 60th birthday. Controlling also for pre-retirement income, years of education, and if any formal business education has been obtained, I still find that the effect of extrapolation increases by a factor of 2.9 moving from individuals in their fifties to those in their sixties and beyond (2.0 pp vs. 5.7 pp). Furthermore, to account for any omitted variable that can jointly influence the decision to retire later and the tendency to extrapolate, I restrict the analysis to IBM employees that were laid-off. I still find an increase of 2.2 times in the extrapolation coefficient between the two age groups (2.0 pp. vs. 4.4 pp).

One last concern in interpreting these results is that older employees might have experienced significantly different stock returns during their lifetime such as the extremely negative returns during the Great Depression. To account for this possibility, Online Appendix 2.H reports results from specifications that include lifetime returns (i.e., since birth). While lifetime experiences are significant in explaining annuitization only when people have experienced unusually high returns (i.e., the highest quintile of lifetime returns), I still find that recent events are the ones largely driving annuitization.

(*footnote continued*)
report evidence from the IBM sample on employees terminated and, hence, less likely to time their retirement.

5. Conclusion

With the oldest cohort of the baby boomers recently reaching age 65, millions are expected to retire in the near future. Many of these retirees will soon face the challenge of managing their retirement wealth to make it last a lifetime while facing uncertainty regarding investment returns and longevity. In this paper, I show that recent stock returns have a substantial effect on the decision to annuitize retirement wealth.

More precisely, I document a robust and negative relation: after recent positive returns in the stock market, individuals are less likely to choose an annuity (negative returns have the opposite effect). My evidence supports extrapolation from past returns as a plausible explanation for these findings. Moreover, I find that the effect of past returns on annuitization increases at older ages, with a coefficient that is 5.2 times larger moving from the 50–59 age group to the 70–75 group. This result appears consistent with the literature on the effects of aging which shows how the elderly are prone to use more recent information due to cognitive aging and memory loss (Cole and Balasubramanian, 1993; Besedes, Deck, Sarangi, and Shor, 2012).

My results lead to three key implications. First, policy makers promoting annuitization as a retirement income solution should carefully consider the influence of recent stock market returns on the decision to annuitize. My finding that people later in life tend to rely more on recent events makes this argument even more compelling. In back-of-the-envelope calculations (see Online Appendix, 2.I), I estimate that the welfare consequences of annuitizing too early (or never) can be substantial and that the potential losses can range from 5% to 10% of retirement wealth, or the equivalent of working an additional two to five years. The Pension Protection Act has devoted a great deal of attention to helping ensure that employees have enough resources at retirement (by endorsing automatic enrollment, for example). Promoting adequate retirement income solutions should be the next step in policy makers' agendas. Carefully accounting for this tendency to extrapolate from recent returns will help to promote retirement income solutions such as annuities.

Second, I document how naïve beliefs about future stock market returns are likely to affect the attractiveness of annuities. This result highlights the potential demand for annuities — and retirement income solutions in general — that allow retirees to benefit from the potential upside of the stock market. As another example, Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011) estimate a potentially strong demand for annuities with long-term care insurance features. Given the general lack of payout solutions in defined contribution plans and individual retirement accounts, the size of this market [$9.4 trillion at year-end 2011, as estimated by the Investment Company Institute, 2012], and the large number of future retirees, progress in the design of new and more attractive retirement income solutions is of high interest for practitioners as well as policy makers.

Last, my findings on recent experiences contribute to the debate on how beliefs are formed and evolve over time. While the effects of recent market trends have been documented in several financial settings, recent studies (Malmendier and Nagel, 2011; Malmendier, Tate, and Yan, 2011) have provided evidence that lifetime events can also affect financial decisions, especially when they are extreme in nature such as the Great Depression. An interesting avenue for future research is to understand how and to what extent these two tendencies interact. Do extreme lifetime events mitigate or exacerbate the effect of recent trends? The answer to this question can improve our understanding about the long-run effects of recent dramatic events (e.g., the recent financial crisis) on investors' beliefs and simultaneously strengthen the microfoundations of our asset pricing models.

Appendix A. Supplementary data

Supplementary data associated with this paper can be found in the online version at http://dx.doi.org/10.1016/j.jfineco.2014.04.006.

References

Agarwal, S., Driscoll, J.C., Gabaix, X., Laibson, D., 2009. The age of reason: financial decisions over the life cycle and implications for regulation. Brookings Papers on Economic Activity Fall, 51–117.

Agnew, J.R., Anderson, L.R., Szykman, L.R., 2012. An experimental study of the effect of market performance on annuitization and equity allocations. Unpublished working paper. College of William and Mary.

Ai, D., Norton, E.C., 2003. Interaction terms in logit and probit Models. Economics Letters 80, 123–129.

Ameriks, J., Caplin, A., Laufer, S., Van Nieuwerburgh, S., 2011. The joy of giving or assisted living? Using strategic surveys to separate public care aversion from bequest motives. Journal of Finance 66, 519–561.

Bailey, W., Kumar, A., Ng, D., 2011. Behavioral biases of mutual fund investors. Journal of Financial Economics 102, 1–27.

Benartzi, S., 2001. Excessive extrapolation and the allocation of 401(k) accounts to company stock. Journal of Finance 56, 1747–1764.

Benartzi, S., Thaler, R.H., 2007. Heuristics and biases in retirement savings behavior. Journal of Economic Perspectives 21, 81–104.

Benartzi, S., Previtero, A., Thaler, R.H., 2011. Annuitization puzzles. Journal of Economic Perspectives 25, 143–164.

Bernheim, B.D., 1991. How strong are bequest motives? Evidence based on estimates of the demand for life insurance and annuities. Journal of Political Economy 99, 899–927.

Besedes, T., Deck, C., Sarangi, S., Shor, M., 2012. Age effects and heuristics in decision making. Review of Economics and Statistics 94, 580–595.

Bester, A.C., Conley, T.G., Hansen, C.B., 2011. Inference with dependent data using cluster covariance estimators. Journal of Econometrics 165, 137–151.

Brown, J.R., 2001. Private pensions, mortality risk, and the decision to annuitize. Journal of Public Economics 82, 29–62.

Brown, J.R., 2007. Rational and behavioral perspectives on the role of annuities in retirement planning. NBER Working Paper No. 13537.

Chalmers, J., Reuter, J., 2012. How do retirees value life annuities? Evidence from public employees. Review of Financial Studies 25, 2601–2634.

Chevalier, J., Ellison, G., 1997. Risk taking by mutual funds as a response to incentives. Journal of Political Economy 105, 1167–1200.

Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A., 2009. Reinforcement learning and investor behavior. Journal of Finance 64, 2515–2534.

Cocco, J.F., Gomes, F.J., 2012. Longevity risk, retirement savings, and financial innovation. Journal of Financial Economics 103, 507–529.

Cole, C.C., Balasubramanian, S.K., 1993. Age differences in consumers' search for information: public policy implications. Journal of Consumer Research 20, 157–169.

Drolet, A., Schwarz, N., Yoon, C. (Eds.), 2010. The Aging Consumer: Perspectives from Psychology and Economics, Routledge, Taylor & Francis, New York.

Finkelstein, A., Poterba, J., 2004. Adverse selection in insurance markets: policyholder evidence from the U.K. annuity markets. Journal of Political Economy 112, 183–208.

Fuster, A., Hebert, B., Laibson, D., 2012. Natural expectations, macroeconomic dynamics, and asset pricing. NBER Macroeconomics Annual 26, 1–48.

Greenwood, R., Nagel, S., 2009. Inexperienced investors and bubbles. Journal of Financial Economics 93, 239–258.

Hirshleifer, D., Yu, J., 2012. Asset pricing in production economies with extrapolative expectations. Unpublished working paper. University of California, Irvine.

Hurst, E., Lusardi, A., 2004. Liquidity constraints, household wealth, and entrepreneurship. Journal of Political Economy 112, 319–334.

Inkmann, J., Lopes, P., Michaelides, A., 2011. How deep is the annuity market participation puzzle? Review of Financial Studies 24, 279–319.

Investment Company Institute (ICI), 2012. A Review of Trends and Activity in the U.S. Investment Company Industry. Investment Company Fact Book, 52nd edition.

Kaustia, M., Knupfer, S., 2012. Peer performance and stock market entry. Journal of Financial Economics 104, 321–338.

Kaustia, M., Knupfer, S., 2008. Do investors overweight personal experience? Evidence from IPO subscriptions. Journal of Finance 63, 2679–2702.

Knabb, R.D., Rhome, J.R., Brown, D.P., 2005. Tropical cycle report, Hurricane Katrina, 23–30 August 2005. National Hurricane Center ⟨http://www.nhc.noaa.gov/pdf/TCR-AL122005_Katrina.pdf⟩ (Accessed in April, 2009).

Koijen, R., Van Nieuwerburgh, S., Yogo, M., 2011. Health and mortality delta: assessing the welfare cost of household insurance choice. NBER Working Paper No. 17325.

Korniotis, G.M., Kumar, A., 2011. Do older investors make better investment decisions? Review of Economics and Statistics 93, 244–265.

Lusardi, A., Mitchell, O.S., 2007. Baby boomer retirement security: the roles of planning, financial literacy, and housing wealth. Journal of Monetary Economics 54, 205–224.

Malmendier, U., Nagel, S., 2011. Depression babies: do macroeconomic experiences affect risk-taking?. Quarterly Journal of Economics 126, 373–416.

Malmendier, U., Tate, G., Yan, J., 2011. Overconfidence and early-life experiences: the effect of managerial traits on corporate financial policies. Journal of Finance 66, 1687–1733.

Mitchell, O.S., Poterba, J., Warshawsky, M., Brown, J., 1999. New evidence on the money's worth of individual annuities. American Economic Review 89, 1299–1318.

Poterba, J., 2004. The impact of population aging on financial markets. NBER Working Paper No. 10851.

Sinclair, S., Smetters, K., 2004. Health shocks and the demand for annuities. CBO Technical Working Paper. Congressional Budget Office, Washington, DC.

Sirri, E., Tufano, P., 1998. Costly search and mutual fund flows. Journal of Finance 53, 1589–1622.

Vissing-Jorgensen, A., 2004. Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions. In: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual, vol. 18; 2004, pp. 139–194.

Yoon, C., Cole, C.A., Lee, M.P., 2009. Consumer decision making and aging: current knowledge and future directions. Journal of Consumer Psychology 19, 2–16.