ERF WORKING PAPERS SERIES

Do Savers Learn from Experience? Evidence from Pension Contributions

Sadettin Haluk Çitçi, Halit Yanıkkaya and Yunis Dede

ECONOMIC RESEARCH FORUM
منتدى البحوث الاقتصادية

Working Paper No. 1799
October 2025

DO SAVERS LEARN FROM EXPERIENCE? EVIDENCE FROM PENSION CONTRIBUTIONS

Sadettin Haluk Çitçi,[1] Halit Yanıkkaya[2] and Yunis Dede[3]

Working Paper No. 1799

October 2025

We acknowledge the financial support of TÜBİTAK (The Scientific and Technological Research Council of Türkiye) under project number 222K023. We are also grateful to the Pension Monitoring Centre for providing valuable data and insights that contributed significantly to the completion of this research. We would like to express our sincere thanks to Ernst Fehr, Mustafa Akmaz, Hüseyin Kaya, and Murat Taşdemir for their valuable feedback and encouragement throughout the project. We also benefited greatly from discussions with numerous seminar participants at the University of Zurich and from comments received at the Economic Research Forum (ERF) conference. We further appreciate the insightful feedback provided during the Turkish Economic Association conference.

Send correspondence to:
Sadettin Haluk Çitçi
Gebze Technical University
hcitci@gtu.edu.tr

[1] Gebze Technical University & University of Zurich, Department of Economics. E-mail: hcitci@gtu.edu.tr & sadettinhaluk.citci@uzh.ch

[2] Gebze Technical University, Department of Economics, Türkiye. E-mail: halityanikkaya@gtu.edu.tr.

[3] Gebze Technical University, Department of Economics, Türkiye. E-mail: ydede@gtu.edu.tr.

First published in 2025 by
The Economic Research Forum (ERF)
21 Al-Sad Al-Aaly Street
Dokki, Giza
Egypt
www.erf.org.eg

Abstract

We examine whether households' voluntary retirement saving decisions are influenced by reinforcement learning (RL), a behavioral heuristic where recent outcomes disproportionately shape future choices. Using eight years of universe-wide administrative data from Türkiye's Individual Pension System, we show that savers over-weight recent return experiences. Specifically, individuals experiencing higher returns in one year substantially increase their voluntary contributions in the following year, and past returns continue to affect contributions with a diminished but persistent impact. The implied one-year learning weight is moderate, closely mirroring laboratory estimates. Alternative explanations such as inertia, skill learning, or asset rebalancing do not explain these observed behaviors.

Keywords: Reinforcement learning, heuristics, decision-making, household saving.
JEL Classifications: D80, G41, D70, J26, G51.

ملخص

ندرس في هذه الورقة ما إذا كانت قرارات الادخار التقاعدي الطوعي للأسر تتأثر بالتعلم التعزيزي (RL)، وهو أسلوب سلوكي تُشكل فيه النتائج الأخيرة خياراتها المستقبلية بشكل غير متناسب. باستخدام ثماني سنوات من البيانات الإدارية الشاملة لنظام التقاعد الفردي في تركيا، نُظهر أن المدخرين يُبالغون في تقييم تجارب العوائد الأخيرة. على وجه التحديد، يُزيد الأفراد الذين يحققون عوائد أعلى في عام ما مساهماتهم الطوعية بشكل كبير في العام التالي، بينما تستمر العوائد السابقة في التأثير على المساهمات بتأثير مُتناقص ولكنه مستمر. يُعدّ وزن التعلم المُفترض لمدة عام واحد متوسطًا، ويعكس بشكل وثيق التقديرات المختبرية. لا تُفسر التفسيرات البديلة، مثل القصور الذاتي، أو تعلم المهارات، أو إعادة توازن الأصول، هذه السلوكيات الملحوظة.

1. Introduction

Traditional life-cycle models assume that households update beliefs about portfolio returns in a fully Bayesian manner, weighting all past information efficiently. Laboratory work in psychology and game theory, however, documents a different algorithm—reinforcement learning (RL)—in which recent personal outcomes receive disproportionate weight and guide subsequent actions (Roth and Erev, 1995; Erev and Roth, 1998; Camerer and Ho, 1999). Given substantial policy efforts aimed at enhancing individual retirement savings, understanding whether RL heuristics govern saving behaviors over longer horizons, independently of institutional defaults or employer influences, is of critical economic importance.

Empirical evidence increasingly confirms RL heuristics as a parsimonious explanation for financial behaviors such as stock-market participation across cohorts (Kaustia and Knüpfer, 2008), savings decisions in U.S. 401(k) plans (Choi et al., 2009), and risk-taking (Malmendier and Nagel, 2011). Yet, direct individual-level evidence demonstrating that RL shapes voluntary saving decisions over substantial horizons, net of plan defaults and employer-driven influences, remains notably absent. By leveraging eight years of comprehensive administrative data from Türkiye's Individual Pension System (IPS), we directly address this gap and demonstrate that RL significantly influences households' choices to set aside additional resources for individual retirement accounts (IRA).

Reinforcement learning (RL) is a "repeat-what-just-worked" algorithm: each time an action pays off, its attractiveness is bumped up in proportion to the size of the payoff. In game theory experiments this simple rule has repeatedly out-predicted fully Bayesian, forward-looking models (Roth and Erev 1995; Erev and Roth 1998) and has even lured experimental subjects away from the Bayes-optimal choice when the two prescriptions conflict (Charness and Levin 2003). A salient property of the mechanism—sometimes called the power law of practice—is that reactions are strongest to the most recent outcomes and attenuate rapidly as those experiences recede into the past. Transposed to annual retirement savings, the theory yields clear empirical predictions: A good personal return this year should prompt a substantial rise in next year's voluntary contribution, while last year's return should still exert a positive but noticeably smaller pull. The ratio of these two effects reveals how quickly households discount older experiences and thus provides an estimate of the learning rate that governs their behavior.

To empirically test these predictions, we exploit a uniquely suitable dataset: the Pension Monitoring Centre's universe-wide administrative records for Türkiye's Individual Pension System (IPS) covering the period 2009–2016. For every participant and calendar year the database records the exact voluntary contribution, the net return earned on that saver's own fund mix, the equity–bond–cash split, and basic demographics. Restricting attention to accounts that contribute to in every year of the window yields a balanced panel that follows eight distinct return signals for each individual. Our empirical strategy involves estimating a linear savings model using panel data methods to determine how individuals respond to both contemporaneous and one-year-lagged returns. This design lets us isolate the imprint of personal performance on saving behavior.

Our findings confirm that savers over-extrapolate from their personal return experience in their savings decisions. A one-percentage-point improvement in a saver's own monthly average real return during the year is followed by roughly a 22 percent jump in the amount she voluntarily sets aside at year-end; the echo of the same shock a year earlier still adds about half as much again (approximately 11 percent). The gap between the two coefficients translates into one-year learning-rate of about 0.66. In lira terms the effect is sizable: a one–standard-deviation upside surprise raises the typical annual contribution by ₺936, a sum comparable to 2-week of median net pay.

We conduct extensive robustness analysis to validate these results. The identified recency gradient remains stable after controlling for demographic variables, portfolio allocations, pension company effects, and policy shifts such as government matching contributions. Alternative hypotheses—including inertia, skill-based learning, and portfolio rebalancing between pension and non-pension assets—fail to adequately explain the observed patterns. Thus, our analysis substantiates the presence and economic importance of reinforcement learning, widely documented in experimental contexts, in real-world financial decisions.

This paper makes two substantive contributions to the literature. First, it provides novel an eight-year, universe-wide administrative panel that follows the same Turkish savers through a full market cycle, allowing us to uncover medium-run reinforcement effects on the intensive margin of retirement saving. Second, the rich within-individual variation in those average return signals allows us to recover an estimate of the short-horizon reinforcement-learning weight. While this estimate is derived from a deliberately parsimonious one-lag specification and does not capture longer-horizon decay, it offers a calibrated benchmark that links laboratory parameters to real-world behaviour and can inform future structural models of household finance.

The remainder of the paper proceeds as follows. Section 2 details the institutional background of Türkiye's IPS. Section 3 describes the data, empirical design, and identification strategy. Section 4 presents the baseline estimates of the recency gradient. Section 5 probes alternative explanations and reports robustness checks. Finally, Section 6 concludes.

2. Institutional background

The IPS was introduced in Türkiye in October 2003 as a complementary scheme to the existing social security system. Designed to provide individuals with additional income during retirement, the IPS operates on a voluntary basis, granting participants flexibility in determining their contribution levels and allocating their savings among up to 230 different funds by the end of 2013 (PMC, 2014). Participants can transfer their savings across funds up to six times per year without incurring any fees.[4]

There are three distinct ways to participate in the IPS. First, individuals can open an IRA through individual contracts by selecting a pension company offering customized pension schemes. Eligibility is open to all Turkish citizens aged 18 and above, regardless of employment status. Second, institutional groups—such as professional associations, non-governmental organizations, and unions—can participate via group pension contracts. Third, employers can establish employer group contracts to enroll their employees, contributing a percentage of wages on their behalf. Participants may hold multiple contracts of any type simultaneously, paying contributions concurrently without legal restrictions. Unlike systems in countries such as the United States, Canada, and the United Kingdom, the IPS in Türkiye is not employer-sponsored. Individual contracts dominate the system, representing 74% of all contracts in terms of both fund size and the number of contracts by the end of 2013 (PMC, 2014).

Since its inception, the IPS has undergone two significant reforms. The first occurred in 2013, when the state introduced a matching contributions program, providing a 25% match on annual contributions, up to 25 percent of that year's gross minimum wage. This match, later increased to 30% in 2022. Matching funds are invested on the saver's behalf in a dedicated state contribution fund, which mainly holds government bonds. In the case of a participant has multiple contracts independent of contract type, the matching contributions is calculated from the total contributions in all contracts.

The second reform came in 2017 with the introduction of the Automatic Enrollment System (AES), aimed at increasing participation in the IPS. Under AES, employees under the age of 45 are automatically enrolled by their employers, contributing a minimum of 3% of their gross wage. Participants can opt out without penalty within the first two months; however, withdrawals after this period incur a 15% income tax on returns. For retirees, the tax rate is reduced to 10%. Importantly, enrollment in the AES does not affect existing individual contracts, and participants can continue contributing to or opening new individual retirement accounts while participating in AES.

3. Data and methodology

This study utilizes a unique administrative dataset from the Pension Monitoring Center, the central authority responsible for managing the IPS in Türkiye. This extensive dataset includes information from over 39 million pension contracts, capturing a wide array of financial, occupational, and demographic variables. For our analysis, we focus on the period from 2009 to 2016, isolating

[4] As of the second half of 2021, these allowable number of fund changes was increased 12 per a year.

potential effects AES policy, which is introduced in 2017, on saving behaviors of individuals (Yanıkkaya et al., 2023).

The dataset provides detailed information on participants' annual contributions, cumulative total assets, and portfolio allocations across 12 distinct fund groups. Additionally, it includes demographic data such as age, gender, and education level, alongside financial details such as contract types (individual, pension group, or employer group), pension company affiliations, and number of portfolio rebalancing. Income data, though available for individual contracts, is self-reported and therefore prone to potential biases. Given that income information is only available for 24% of the sample, we thus use education level as a proxy to control for income effects in our analysis[5].

To ensure a balanced evaluation of saving behavior over medium- and long-term horizons, we restrict the sample to participants with at least eight consecutive years of contributions. In the first stage of the analysis, we use data from 2009 to 2016 to explore general patterns RL heuristics in saving behavior. In the second stage, to exclude potential effects of matching contributions policy on RL heuristics and to check robustness of our results, we include only the sample in the pre-policy period (2009–2012).

Our analysis focuses exclusively on individual pension contracts. Group and employer-sponsored contracts are excluded, as their creation often depends on institutional or employer preferences, which may not reflect participants' independent saving choices. Approximately 30% of participants maintain multiple individual contracts; for these cases, we aggregate contributions and returns, weighting them proportionally.

To address potential distortions caused by extreme values, we exclude participants falling in the top and bottom 1% of the distributions for annual contributions and portfolio returns. This ensures that the empirical results are not unduly influenced by outliers.

The primary objective of this study is to investigate the role of past investment experiences in shaping participants' later saving decisions, with particular focus on whether these decisions align with RL heuristics. Specifically, we examine the sensitivity of annual contributions to both contemporaneous and lagged returns, testing whether participants over-extrapolate return experiences in their savings behavior. Our approach aligns with the empirical frameworks of prior RL studies (Kaustia and Knüpfer, 2008; Choi et al., 2009; Chiang et al., 2011).

[5] Numerous studies show a strong relationship between income and education level (Krueger and Ashenfelter, 1994; Harmon et al., 2003, etc.).

To model these relationships, we estimate the following baseline equation:

$$CP_{i,t} = \alpha + \beta_1 R_{i,t} + \beta_2 R_{i,t-1} + \gamma D_{i,t} + \delta P_{x,t} + \lambda S_{y,t} + \phi F_{z,t} + ¥T_t + \theta M_{ss} + \varepsilon_{i,t} \qquad (1)$$

$CP_{i,t}$ is the natural logarithm of the annual contributions for participant *i* at the end-of-year *t*. $R_{i,t}$ and $R_{i,t-1}$ are contemporaneous and lagged real returns, respectively (monthly arithmetic average of the annual percentage real return)[6]. $D_{i,t}$ is demographic control vector, including age, gender, and education. $P_{x,t}$ is the fourteen pension company dummies. $S_{y,t}$ denotes the portfolio allocation shares by fund groups and represents the investment preferences. $F_{z,t}$ is number of portfolio rebalancing and shows portfolio re-optimizing behavior in that year (change in portfolio allocation shares)[7]. T_t is the year fixed effects. M_{ss} is a binary variable indicating the matching contributions policy, which equals 1 for years 2013 onward and 0 otherwise.

To account for potential non-linear effects of age, we include both linear and quadratic terms for age in the demographic control vector $D_{i,t}$. Portfolio allocation shares are modeled to account for temporal variations in fund performance. Additionally, the proportions of assets allocated to equity and flexible funds are used as proxies for participants' risk tolerance (Malmendier and Nagel, 2011).

In some years and companies, financial advisors at pension companies may enable superior portfolio returns compared to market averages. Fisch et al. (2016) and Marsden et al. (2011) highlight that engaging with a financial advisor is associated with enhanced financial planning activities and outcomes. To account for the influence of differences in pension companies on investment performance—and consequently on savings behavior—our empirical framework incorporates pension company fixed effects. We use pension companies, portfolio allocation shares and number of portfolio rebalancing fixed effects separately and as interaction terms with year dummies.

Lastly, we recover the one-year learning weight, ψ, as

$$\hat{\psi} = \frac{\widehat{\beta_1}}{\widehat{\beta_1} + \widehat{\beta_2}} \qquad (2)$$

and compute standard error with delta method, using the joint covariance matrix of $(\widehat{\beta_1},\ \widehat{\beta_2})$.

[6] We utilize the monthly average portfolio returns, reflecting the predominant preference among individuals for the "monthly payment" option.

[7] Participants may choose to re-optimize and rebalance their portfolios in response to various factors, including shifts in risk tolerance, changes in market outlook and investment horizon. In practice, however, retirement accounts are notably inertial (Choi et al. 2002; Agnew et al. 2003; Ameriks and Zeldes 2011). Our data confirm this passivity: roughly 70 percent of savers never altered their portfolio weights, and only about 15 percent made more than one change over the entire sample period.

4. Estimation results

Table 1 provides summary statistics for the dataset, encompassing contributions, returns, demographic profiles, and portfolio characteristics. The average annual real contributions by participants during the sample period (2009–2016) are approximately 2,753 ₺ in 2013 prices, representing roughly 26% of the net minimum wage in that year. This ratio aligns with the total domestic savings-to-GDP ratio of 23.2%, as reported by the Ministry of Development (2024).

On average, participants' real portfolio returns approach zero. Peker (2016) links this outcome to the high management fees and administrative charges embedded in IRA assets, a pattern of underperformance that Ayaydın (2013) and Açıkgöz et al. (2015) likewise document.

Demographically, women slightly outnumber men in the sample, with females representing 53.7% of participants. The average age of participants is 35, with age distributions concentrated in the 26–35 and 36–45 age ranges. Approximately 48.6% of participants have attained an undergraduate degree or higher, indicating a strong preference for IPS participation among more educated individuals. Notably, 30% of participants maintain multiple individual contracts.

Table 1. Summary statistics

	Mean	Median	Std. Deviation	Minimum	Maximum	Observations
Annual Real Contributions	2,753	2,161	1,983	430	19,706	782,936
Monthly Average Real Return (%)	0.07	0.01	0.57	-1.64	2.42	782,936
Lagged Monthly Average Real Return (%)	0.02	-0.01	0.59	-1.64	2.42	685,069
			Gender (%)			
	Male			**Female**		
	46.27			53.73		
			Age (%)			
≤ 25	**26 - 35**		**36 - 45**		**46 - 55**	**55 >**
6.83	37.5		38.56		14.8	2.31
			Education Level (%)			
Less than High School			**High School – College**			**College Graduate or Higher**
19.93			31.47			48.6
			Average Portfolio Allocation Shares (%)			
Domestic Bonds	**Foreign Bonds**	**Equity**	**Balanced**	**Money Market**	**Standard**	**Flexible**
46.67	3.73	3.41	1.36	12.68	3.44	24.58
			Portfolio Rebalancers (%)			
None			1			> 1
68.31			17.00			14.68
			Number of Individual Contracts per Individual (%)			
1			2			≥ 3
69.69			17.99			12.32
			Having at Least One Pension Group or Employer Group Contract (%)			
	Yes			No		
	8.8			91.2		

Source: Pension Monitoring Center (2024)

Note: We categorize participants' age based on Pension Monitoring Center classification and education level.

Savers exhibit clear risk-averse preferences in their portfolio allocations. Government bond-dominated domestic fixed income funds constitute the largest share at 50% of total assets on average, while direct equity fund allocations remain modest. This conservative stance is partially offset by substantial allocations to flexible funds (25% of portfolios), which provide strategic diversification and equity-like exposure. Overall, this allocation pattern reflects participants' cautious approach to risk management within the IRA framework.

4.1. Estimation results

The regression results, presented in Table 2, examine the relationship between portfolio returns and savers' annual contributions, progressively incorporating controls to isolate the effects of return experiences. The dependent variable is the log of the end-of-year contribution, while the key covariates are the contemporaneous and one-year-lagged real returns, expressed as the arithmetic mean of monthly percentage returns over the calendar year.

Table 2. Regression of contributions on returns (2009 - 2016)

	(I)	(II)	(III)	(IV)	(V)	(VI)
Return	0.126***	0.122***	0.174***	0.208***	0.223***	0.222***
	(0.0015)	(0.0015)	(0.0025)	(0.0041)	(0.0041)	(0.0050)
Lagged Return	0.072***	0.0707***	0.135***	0.0915***	0.117***	0.109***
	(0.0012)	(0.0013)	(0.0018)	(0.0025)	(0.0025)	(0.0031)
Age		0.0341***	0.0304***	0.0297***	0.0291***	0.0335***
		(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0005)
Age^2		-0.0003***	-0.0002***	-0.0002***	-0.0002***	-0.0002***
		(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Gender		0.114***	0.0909***	0.0835***	0.0802***	0.0777***
		(0.0013)	(0.0013)	(0.0013)	(0.0013)	(0.0015)
High Sch. - College						0.0774***
						(0.0020)
College or higher						0.200***
						(0.0019)
Matching contributions dummy	0.139***	0.140***	0.209***	0.632***	0.511***	0.655***
	(0.0014)	(0.0014)	(0.0067)	(0.0361)	(0.0364)	(0.0619)
Constant	7.692***	6.834***	5.775***	5.604***	5.572***	6.344***
	(0.0010)	(0.0088)	(0.0695)	(0.0835)	(0.0814)	(0.0548)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Company x Year dummies	No	No	Yes	Yes	Yes	Yes
Portf. Alloc. Shares x Year dummies	No	No	No	Yes	Yes	Yes
Portf. Rebalancing x Year dummies	No	No	No	No	Yes	Yes
Observations	685,069	685,069	685,069	685,069	685,069	469,763

*Note: Clustered robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.*

In the simplest specification (column I), a one-unit increase in contemporaneous monthly average returns raises annual contributions by 12 percent, while lagged returns contribute an additional 7 percent. Including the age and gender in column II slightly reduces these the estimated coefficients. Columns III and IV add pension company fixed effects and portfolio allocation shares, and account for participant heterogeneity and provider-level differences. Finally, the fully specified model in column V interacts the number of within-year portfolio rebalancing with year dummies and column VI adds categorical controls for education. In this model, the combined effect of contemporaneous and lagged returns results in a 34% increase in annual contributions. Given a sample mean of ₺2753, this amounts to an additional ₺936 per saver, roughly one-third of the typical voluntary contribution and close to a 2-week of median net pay.

The ratio of contemporaneous to lagged return coefficients implies a short-horizon learning weight of $\psi \approx 0.66$ (95% CI [0.64, 0.67]). This estimate suggests moderate but rapid updating to new information, placing it at the upper bound of the 0.4–0.6 range typically observed in laboratory reinforcement-learning studies (Roth and Erev, 1995; Camerer and Ho, 1999; Charness and Levin, 2005). The result aligns with models of experience-weighted learning (Erev and Haruvy, 2013), which formalize how salient feedback amplifies recency effects. In real-world settings like pension contributions, where returns directly impact wealth and feedback is vivid and consequential, agents exhibit faster updating than in abstract, low-stakes laboratory tasks. This may explain why our estimate exceeds typical lab-based values, offering field validation for the contextual modulation of learning rates predicted by theory.

RL provides a natural framework for interpreting these findings. Savers appear to extrapolate from their own recent outcomes, implicitly assuming that investments that have paid off (or disappointed) will continue to do so. Rather than adopting fully forward-looking optimization, participants appear to rely on RL heuristics, adjusting their contributions upward following positive returns. This behavior likely serves as a cognitive shortcut in the face of investment complexity and uncertainty.

Our estimates extend earlier work. Choi et al. (2009) emphasize the short-term responsiveness of contributions to contemporaneous returns within the same cohort, and Malmendier and Nagel (2011) document the enduring impact of macroeconomic experiences across cohorts, our results bridge this gap by presenting evidence of medium-to-long-term RL effects within the same cohort. Choi et al. (2009) show strong sensitivity to contemporaneous returns but no lagged effect; we demonstrate that last year's personal performance still matters when an eight-year panel is used. Related evidence from Kaustia and Knüpfer (2008), Chiang et al. (2011), and Song et al. (2021) that returns predict future contributions is corroborated here, with the additional insight that the effect operates through a recency-weighted mechanism.

4.2. Pre-policy analysis

The 25 percent government match introduced in 2013 could magnify contribution elasticities. Table 3 therefore re-estimates the model for the pre-policy years 2009–2012.

Table 3. Regression of contributions on returns (2009 - 2012)

	(I)	(II)	(III)	(IV)	(V)
Return	0.0525***	0.0529***	0.0749***	0.0280***	0.0354***
	(0.0019)	(0.0019)	(0.0029)	(0.0051)	(0.0051)
Lagged Return	0.0016	0.00244*	0.0799***	0.0565***	0.0725***
	(0.0014)	(0.0014)	(0.0021)	(0.0029)	(0.0029)
Age		0.0330***	0.0287***	0.0281***	0.0275***
		(0.0006)	(0.0006)	(0.0006)	(0.0006)
Age^2		-0.00031***	-0.00026***	-0.00025***	-0.00024***
		(0.00008)	(0.00008)	(0.00008)	(0.00008)
Gender (Male = 1)		0.107***	0.0856***	0.0813***	0.0788***
		(0.0018)	(0.0017)	(0.0017)	(0.0017)
Constant	7.674***	6.871***	6.793***	6.609***	6.590***
	(0.00103)	(0.0123)	(0.0119)	(0.0303)	(0.0302)
Year dummies	Yes	Yes	Yes	Yes	Yes
Company x Year dummies	No	No	Yes	Yes	Yes
Portf. Alloc. Share x Year dummies	No	No	No	Yes	Yes
Portf. Rebalancing x Year dummies	No	No	No	No	Yes
Observations	293,601	293,601	293,601	293,601	293,601

*Note: Clustered robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.*

In all estimations, coefficient of the contemporaneous return is again significant at 1% level, suggesting presence of RL heuristics in saving decisions. However, the estimated coefficients of contemporaneous and lagged returns are smaller indicating either the matching contributions policy adds an upward bias in estimations or it amplifies RL heuristics (or both).

In the simplest specification (column I), a one-unit increase in contemporaneous monthly average returns is associated with a 5.2% increase in contributions. Including the age and gender in column II does not affect the coefficient. Columns III and IV add pension company fixed effects and portfolio allocation shares and account for participant heterogeneity and provider-level differences. Finally, the fully specified model in column V incorporates number of portfolio rebalancing as interaction terms with year dummies, producing the most precise estimates. In this model, the combined effect of contemporaneous and lagged returns results in a 10.8% increase in contributions implying that a one-unit increase in average monthly returns is associated with approximately ₺297 rise in annual contributions consistent with RL heuristics.

The positive, recency-weighted return elasticities are robust across a range of controls. We next test three alternative explanations—inertia, skill learning, and cross-account rebalancing—by estimating focused subsamples and interaction models. Section 5 reports the results of tests of alternative explanations.

5. Alternative scenarios

Some other stories that may explain the established positive impact of return on contributions. In this section, we test alternative explanations. First, we turn our attention to participants who make portfolio rebalancing, believing that the RL heuristic may reflect an existing inertia in IRA. Second, we assess the profitability of RL-driven portfolio adjustments by analyzing the persistence of individual portfolio alphas. Third, we investigate potential rebalancing effects between IRA assets and non-IRA assets by interacting RL heuristics with participants' age and education levels.

5.1. Inertia

There is extensive evidence that pension investors display marked inertia: most stick to their enrolment portfolio and rarely rebalance (Choi et al. 2002; Agnew et al. 2003; Ameriks and Zeldes 2011). Our data confirm this pattern: participants alter their asset mix only once in a total of 8 years, and 68% never make a single change over the eight-year window (see Table 1).

One might therefore argue that the positive return elasticities documented earlier are not the product of reinforcement learning (RL) but of a passive "set-and-forget" rule. The observed behavioral pattern could mimic RL heuristics (e.g., maintaining or increasing contributions following positive returns) while being driven by inertia.

The baseline regressions already include reallocation-count x year dummies, so the estimated return effects are net of contemporaneous trading intensity. Inertia could nevertheless matter if savers who never rebalance behave differently in ways those controls cannot fully absorb. To probe this possibility, we partition accounts by the total number of reallocations recorded between 2009 – 2016. Specifically, we divide the individuals in our sample into four subgroups based on the total number of reallocations they made during the eight-year period: zero, one, two and more than two. We then re-estimate equation (1) for these four subgroups respectively and report the results in Table 5.

Table 5. Regression results by total number of portfolio rebalancing

	= 0	= 1	= 2	> 2
	(I)	(II)	(III)	(IV)
Return	0.256***	0.201***	0.153***	0.131***
	(0.0057)	(0.0097)	(0.0140)	(0.0088)
Lagged Return	0.111***	0.118***	0.104***	0.0788***
	(0.0034)	(0.0060)	(0.0090)	(0.0063)
Age	0.0241***	0.0310***	0.0331***	0.0507***
	(0.0004)	(0.0012)	(0.0023)	(0.0023)
Age^2	-0.0002***	-0.0003***	-0.0002***	-0.0004***
	(0.00000)	(0.00001)	(0.00003)	(0.00003)
Gender (Male = 1)	0.0712***	0.0594***	0.0795***	0.100***
	(0.00142)	(0.00314)	(0.00549)	(0.00495)
Matching contributions dummy	0.877***	0.246***	0.401***	0.697***
	(0.0519)	(0.0874)	(0.139)	(0.111)
Constant	5.625***	5.636***	5.591***	6.111***
	(0.1020)	(0.1310)	(0.2430)	(0.1080)
Year dummies	Yes	Yes	Yes	Yes
Company x Year dummies	Yes	Yes	Yes	Yes
Portf. Alloc. Share x Year controls	Yes	Yes	Yes	Yes
Portf. Rebalancing x Year dummies	-	Yes	Yes	Yes
Observations	467,992	116,494	42,791	57,792

*Note: Clustered robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.*

Among the savers who never rebalance (column I), a one-unit increase in contemporaneous returns raises annual contributions by 25.6 percent, while last year's return adds another 11.1 percent. For savers who rebalance more than twice over the eight-year window (column IV) the corresponding responses decline to 13.1 percent (contemporaneous) and 7.9 percent (lagged). A plausible explanation is that active traders hold better-diversified portfolios, so any single-year gain or loss carries less weight. The persistence of sizeable return response across all activity levels—including those who actively manage their portfolios—shows that RL operates over and above both scheduled escalation and default stickiness.

In sum, we find no empirical support for the hypothesis that contribution inertia drives the observed return sensitivities and explains the results. Even when we isolate savers with minimal opportunity for inertia-driven drift, their contributions remain strongly and systematically tied to recent personal returns. The persistent, recency-weighted response aligns squarely with a reinforcement-learning mechanism.

5.2. Investment skills

One explanation for our finding that participants increase contributions in response to positive returns is that they may view high returns as evidence of superior investment skill. Over time, such participants might perceive themselves as more capable than average, leading them to devote additional resources to their individual retirement accounts. If this explanation holds, we expect persistence in individual portfolio alphas over time, as consistently high-performing participants maintain better-than-average results.

To test this hypothesis, we examine the persistence of participants' portfolio alphas between year *t* and year *t−1*, adapting the methodology of Choi et al. (2009). Formally, we define each participant's annual alpha as the risk-adjusted excess return.

$$\alpha_{\mathrm{i,t}} = \left[r_{i,t} - \left[r_{f,t} + \beta_{\mathrm{i,t}}\left(r_{m,t} - r_{f,t}\right)\right]\right] \qquad (3)$$

where $r_{i,t}$ is the saver's real portfolio return, $r_{m,t}$ the value-weighted market return for her fund group, and $r_{f,t}$ the year-*t* return on the ten-year Turkish government bond. Given the annual nature of our dataset, we estimate each saver's beta ($\beta_{\mathrm{i,t}}$) from the eight-year panel of her own annual returns. This coarse measure is inevitably noisy, betas can vary within fund groups and we see portfolio weights only at year end. Nevertheless, any economically meaningful skill should generate a positive unconditional relationship between $\alpha_{i,t}$ and $\alpha_{i,t-1}$.

We calculate a market return ($r_{m,t}$) for each fund group annually and assign it to individuals based on their portfolio allocations. Portfolio alphas are then calculated using the assigned market return, participants' annual returns, and the Türkiye 10-year treasury bond yield as the risk-free rate ($r_{f,t}$).

The average alpha across the 8-year period is near zero (0.05), suggesting that, as a group, participants fail to achieve consistent risk-adjusted outperformance. This outcome may reflect inherent challenges in outperforming benchmarks in an efficient market, compounded by administrative fees, suboptimal fund allocation strategies, or limited portfolio optimization. While these factors warrant further exploration, the lack of persistence provides a robust starting point for evaluating individual investment behavior and its implications for the IPS.

Our empirical analysis is summarized in the following regression model:

$$\alpha_{i,t} = \varphi + \beta_1 \alpha_{i,t-1} + \gamma D_{i,t} + \delta P_{x,t} + \mathrm{T} S_{y,t} + \phi F_{z,t} + \Upsilon T_t + \theta M_{ss} + \varepsilon_{i,t} \qquad (4)$$

where $\alpha_{i,t}$ represents participant *i*'s portfolio alpha in year *t*. The other variables are demographic control vector and pension company, portfolio allocation shares and number of portfolio rebalancing fixed effects, respectively, as in our basic model setup. We again control these vectors separately and as interactions with year dummies. To further explore persistence in alphas, we test the persistence of positive and negative lagged alphas, individually.

The regression results, reported in Table 6, provide evidence that portfolio alphas lack persistence over time. Across all specifications, lagged alpha coefficients are negative and statistically significant at the 1% level, indicating strong mean reversion rather than persistence. Participants with positive alphas in year $t-1$ tend to underperform in year *t*, while those with negative alphas in year $t-1$ tend to outperform subsequently. This pattern holds for the full sample and pre-policy (2009–2012) period.

Table 6. Persistence analysis of portfolio alphas

	2009 - 2016		2009 - 2012	
	(I)	(II)	(III)	(IV)
Lagged Alpha	-0.654***		-0.505***	
	(0.0023)		(0.0049)	
Lagged Alpha ≥ 0		-0.613***		-0.086***
		(0.0039)		(0.0079)
Lagged Alpha < 0		-0.692***		-0.848***
		(0.0028)		(0.0056)
Age	0.0320***	0.0040***	-0.0038***	0.0041***
	(0.0006)	(0.0006)	(0.0014)	(0.0013)
Age^2	-0.00043***	-0.00005***	-0.00006***	-0.00005***
	(0.00000)	(0.00000)	(0.00001)	(0.00001)
Gender (Male = 1)	0.00938***	0.00766***	0.0116***	-0.0036
	(0.0018)	(0.0018)	(0.0038)	(0.0036)
Matching contributions dummy	0.884***	0.903***		
	(0.1640)	(0.1640)		
Constant	-0.799***	-0.876***	-1.079***	-1.925***
	(0.1930)	(0.1960)	(0.1440)	(0.1460)
Year dummies	Yes	Yes	Yes	Yes
Company x Year dummies	Yes	Yes	Yes	Yes
Portf. Alloc. Share x Year controls	Yes	Yes	Yes	Yes
Portf. Rebalancing x Year dummies	Yes	Yes	Yes	Yes
Observations	685,069	685,069	293,601	293,601

*Note: Clustered robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.*

For the full sample, the lagged alpha coefficient is -0.654, consistent with reversals observed in both sub-periods. When splitting the sample by the sign of lagged alphas, the reversal is not more pronounced for participants with positive alphas (−0.613) compared to those with negative alphas (−0.692). These results show that past outperformance predicts subsequent underperformance rather than continued success.

Such pronounced mean reversion is inconsistent with a rational "learning my skill" story. Although measurement error in our annual alpha proxy could dampen persistence, the coefficients are so strongly negative that no plausible correction would turn them positive in an economically meaningful way. Reinforcement learning, not perceived investment ability, therefore, remains the most convincing explanation for the contribution patterns we observe.

5.3. Rebalancing

Another potential alternative explanation for the observed positive relationship between portfolio returns and contribution changes is the rebalancing hypothesis. This explanation posits that participants might adjust their contributions as part of a broader strategy to maintain a target allocation between IRA and non-IRA assets. If a participant holds a significant amount of non-IRA assets, a positive correlation between returns on IRA and non- IRA assets could create the appearance of return chasing due to rebalancing.

For example, consider a household aiming to maintain a stable buffer stock of non- IRA assets. When IRA asset returns are high, non-IRA asset returns may also increase (due to correlated market performance). To restore the non-IRA assets balance to its target level, participants might increase IRA contributions and withdraw or consume from non-IRA assets. Such behavior could mimic RL heuristics, where increased IRA contributions appear to follow higher returns.

This rebalancing explanation predicts two patterns: younger participants, who typically hold fewer financial assets, should exhibit weaker return-contribution correlations compared to older participants, who are more likely to hold substantial non-IRA assets. Also, more financially sophisticated participants, who tend to diversify investments and hold more non-IRA assets, should exhibit stronger return-contribution correlations compared to less financially sophisticated participants.

To evaluate these predictions, we examine participants' responsiveness to returns across different age groups and levels of financial sophistication. While financial sophistication is not directly observable, we use participants' education levels as a proxy. We include interaction terms for age, education level, and returns in our regression models to test whether these factors significantly influence the relationship between returns and contributions.

Table 7 presents the results of our analysis. In column I, we interact age and age-squared with both contemporaneous and lagged returns. Although the estimated coefficients for all four individual variables are statistically significant and similar to those in Table 2 (not reported here), the interaction terms coefficients are either not statistically significant or have a very small magnitude. We thus find no significant linear age effect, but a small significant quadratic effect. This suggests age explains minimal variation in reinforcement learning sensitivity.

Table 7. Impact of age and education level on RL

	(I)	(II)	(III)
Return x Age	-0.0004		-0.0003
	(0.0010)		(0.0011)
Return x Age^2	0.00004***		0.00004***
	(0.00001)		(0.00001)
Return x High Sch. – College		-0.00481	-0.00484
		(0.0042)	(0.0042)
Return x College Grad or higher		-0.00622	-0.00628
		(0.0040)	(0.0040)
Lagged Return x Age	-0.0006		-0.0003
	(0.000866)		(0.0010)
Lagged Return x Age^2	0.00003***		0.00003***
	(0.00001)		(0.00001)
Lagged Return x High Sch. – Undergrad		-0.00285	-0.00288
		(0.0036)	(0.0036)
Lagged Return x Undergrad and over		-0.0029	-0.0029
		(0.0034)	(0.0034)
Constant	5.591***	6.343***	6.343***
	(0.0822)	(0.0556)	(0.0556)
Year dummies	Yes	Yes	Yes
Company x Year dummies	Yes	Yes	Yes
Portf. Alloc. Shares x Year controls	Yes	Yes	Yes
Portf. Rebalancing x Year dummies	Yes	Yes	Yes
Observations	685,069	469,763	469,763

*Note: Clustered robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.*

In column II, we include participants' education levels as a proxy for financial sophistication, interacting them with contemporaneous and lagged returns. Again, the interaction terms are insignificant, suggesting no meaningful differences in responsiveness to returns across education levels. When age and education interactions are included together in column III, the results remain unchanged.

Taken as a whole, the results are inconsistent with the key predictions of the rebalancing hypothesis. If rebalancing were driving the observed behavior, we would expect return sensitivity to vary significantly across age and education levels, reflecting differences in non-IRA asset holdings and financial sophistication. However, the evidence offers no substantive heterogeneity across age or education levels. These results suggest that rebalancing is not the primary driver of the observed relationship between portfolio returns and contribution changes.

6. Conclusion

Using universe-wide administrative data from Türkiye's Individual Pension System, this paper shows that voluntary contributions react strongly to savers' own recent investment outcomes. A one-percentage-point increase in the monthly average real portfolio return raises the next year's contribution by 22 percent; the same shock realized a year earlier still adds 11 percent. The ratio of these coefficients implies a short-horizon reinforcement-learning weight of about one-half, providing the first field-based benchmark that matches laboratory estimates. In economic terms, a one-standard-deviation upside surprise (roughly four percentage points) adds ₺936 to annual saving, almost one-third of the average contribution and close to two weeks of median net pay.

Extensive robustness checks confirm that this recency gradient is not an artefact of policy changes, active rebalancing, or shifts in portfolio risk composition. The pattern predates the 25 percent matching program and survives controls for portfolio mix and provider-year conditions. Alternative explanations such as inertia, perceived investment skill, or cross-account rebalancing cannot reproduce the magnitude or pattern of the coefficients. Taken together, the results establish that a simple reinforcement-learning heuristic, long documented in laboratory settings, exerts a behaviorally potent and financially meaningful influence on real-world retirement saving.

Two limitations invite further work. First, our estimate captures the decay from year 0 to year 1; longer-horizon memory may follow a more complex path that richer panels or higher-frequency data could uncover. Second, we study voluntary contributions within a specific institutional setting; replicating the analysis in auto-enrolment systems or with opt-out features would test the generality of the mechanism. Addressing these questions would sharpen our understanding of how learning rules interact with retirement-saving architecture—and, ultimately, how policy can harness or temper the power of recent experience.

References

Açıkgöz, E., Uygurtürk, H., & Korkmaz, T. (2015). Analysis of factors affecting growth of pension mutual funds in Turkey. International Journal of Economics and Financial Issues, 5(2), 427-433.

Agnew, J., Balduzzi, P., & Sunden, A. (2003). Portfolio choice and trading in a large 401 (k) plan. American Economic Review, 93(1), 193-215.

Ameriks, J., & Zeldes, S. P. (2011). How do household portfolio shares vary with age? SSRN.

Ayaydın, H. (2013). Türkiye'deki Emeklilik Yatirim Fonlarinin Performanslarinin Analizi. Çukurova Üniversitesi Sosyal Bilimler Enstitüsü Dergisi, 22(2), 59-80.

Camerer, C., & Hua Ho, T. (1999). Experience‐weighted attraction learning in normal form games. Econometrica, 67(4), 827-874.

Charness, G., & Levin, D. (2003). Bayesian updating vs. reinforcement and affect: A laboratory study. Levine's Bibliography 666156000000000180, UCLA Department of Economics.

Charness, G., & Levin, D. (2005). When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect. American Economic Review, 95(4), 1300-1309.

Chiang, Y. M., Hirshleifer, D., Qian, Y., & Sherman, A. E. (2011). Do investors learn from experience? Evidence from frequent IPO investors. The Review of Financial Studies, 24(5), 1560-1589.

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2002). Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. Tax policy and the economy, 16, 67-113.

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2009). Reinforcement learning and savings behavior. The Journal of finance, 64(6), 2515-2534.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. American economic review, 848-881.

Erev, I. and Haruvy, E. (2013) Learning and the economics of small decisions. In the Handbook of Experimental Economics (Kagel, J.H. and Roth, A.E., eds.), Princeton University Press

Fisch, J. E., Wilkinson-Ryan, T., & Firth, K. (2016). The knowledge gap in workplace retirement investing and the role of professional advisors. Duke LJ, 66, 633.

Harmon, C., Oosterbeek, H., & Walker, I. (2003). The returns to education: Microeconomics. Journal of economic surveys, 17(2), 115-156.

Kaustia, M., & Knüpfer, S. (2008). Do investors overweight personal experience? Evidence from IPO subscriptions. The Journal of Finance, 63(6), 2679-2702.

Krueger, A. B., & Ashenfelter, O. C. (1992). Estimates of the economic return to schooling from a new sample of twins.

Malmendier, U., & Nagel, S. (2011). Depression babies: do macroeconomic experiences affect risk taking? The quarterly journal of economics, 126(1), 373-416.

Marsden, M., Zick, C. D., & Mayer, R. N. (2011). The value of seeking financial advice. Journal of family and economic issues, 32, 625-643.

Ministry of Developments. (2024). Economic and Social Indicators [data set]. Retrieved from https://www.sbb.gov.tr/wp-content/uploads/2019/01/Tab-2-6-Yurti%C3%A7i-Tasarruflar%C4%B1n-ve-SSYnin-GSYH-%C4%B0%C3%A7indeki-Pay%C4%B1-1998-2016.xls

Peker, I. (2016). Türkiye Bireysel Emeklilik Sistemi Üzerindeki Kesintiler ile Getirilerin Analizi (Analysing Turkish Individual Pension System Fees and Returns). Journal of Economics, Finance and Accounting–(JEFA), ISSN, 2148-6697.

PMC. (2014). Private Pension System Development Report 2013. PMC.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games and economic behavior, 8(1), 164-212.

Song, R., Jang, S., Wang, Y., Hanssens, D. M., & Suh, J. (2021). Reinforcement learning and risk preference in equity linked notes markets. Journal of Empirical Finance, 64, 224-246.

Yanıkkaya, H., Aktaş Koral, Z., & Çitçi, S. H. (2023). The Power of Financial Incentives versus the Power of Suggestion for Individual Pension: Are Financial Incentives or Automatic Enrollment Policies More Effective? Sustainability, 15(4), 3652.