Private Equity for Everyone?
Measuring Retail Investor Literacy and Behavioral Patterns

Nora Bulava* & Hloriia Pishko†
Master's Thesis in Finance

Stockholm School of Economics
May, 2026

Abstract

As private equity (PE) becomes increasingly accessible to retail investors through retirement plans and semi-liquid investment vehicles, the question of whether retail investors understand the asset class has become more pressing. This thesis examines retail investor understanding of private equity by comparing general financial literacy with PE-specific literacy, and by analyzing how overconfidence and narrative framing shape willingness to invest in PE. Using survey data from 1,032 U.S. retail investors, we find that respondents perform substantially worse on PE-specific questions than on standard financial-literacy measures, suggesting that knowledge of traditional financial concepts does not readily translate into understanding of private markets. PE literacy is concentrated primarily among respondents with the highest levels of general financial literacy, while overconfidence independently predicts willingness to invest in PE. Positive framing also increases stated investment intent, particularly among respondents with moderate levels of PE literacy. Overall, the findings suggest that retail access to private equity is expanding faster than retail investors' ability to evaluate the structures and risks embedded in the asset class, with important implications for investor protection and financial education.

Keywords: Financial literacy; private equity; overconfidence; retail investors; framing; retirement.

Supervisor: David T. Robinson.

Acknowledgements: We would like to thank our supervisor, David Robinson, for his guidance, patience, and continuous encouragement throughout this process. His insights and support during the survey data collection phase were invaluable and helped shape the direction of our work. We are also very grateful to Anders Anderson for his generosity with both his time and expertise, particularly in helping us refine the methodology and strengthen the literature review. Their support, thoughtful feedback, and encouragement made this thesis both possible and deeply rewarding to complete. Any mistakes or questionable decisions are entirely our own.

*42880@student.hhs.se
†42887@student.hhs.se

Contents

1. Introduction

> *"Assets that will define the future – data centers, ports, power grids, the world's fastest-growing private companies – aren't available to most investors. They're in private markets, locked behind high walls, with gates that open only for the wealthiest or largest market participants. But nothing in finance is immutable."*
>
> – LARRY FINK, BLACKROCK CHAIRMAN'S LETTER

> *"All they're [PE firms] doing is lying a little bit to make the money come in."*
>
> – CHARLIE MUNGER, BERKSHIRE HATHAWAY ANNUAL MEETING

The gates are opening. In August 2025, an Executive Order directed the Department of Labor to facilitate alternative assets including private equity in defined-contribution retirement plans (Trump, 2025). BlackRock, Blackstone, Apollo, and KKR are racing to build retail-accessible structures, and ELTIF 2.0 is doing the same in Europe (European Parliament and Council of the European Union, 2023). The argument for access is compelling, and the political momentum is real. There is just one question nobody has properly answered: do retail investors know what they are walking into? This is the question this thesis seeks to answer. And the answer is mostly no.

The financial literacy literature has documented for two decades that general knowledge is low and consequential. Only around a third of U.S. adults answer Lusardi and Mitchell's (2014) benchmark questions on compound interest, inflation, and diversification correctly – and that share is declining: FINRA Foundation (2025) reports that the share answering at least 4 of 5 correctly fell from 42% in 2009 to 32% in 2024. Yet, the standard Big Five questions were built around publicly traded instruments. They do not ask whether an investor can interpret a carried-interest waterfall, evaluate a multi-year lock-up, or recognize the leverage embedded in a buyout. As the domain-specific literacy literature has established for debt (Lusardi & Tufano, 2015) and emerging for crypto (Jones et al., 2024), general literacy is a weak proxy for product-specific knowledge. No validated instrument measuring PE literacy has ever been fielded with a retail sample. This thesis introduces one.

We surveyed 1,032 U.S. retail investors on their knowledge of private equity, its structure, fees, liquidity, leverage, and returns, alongside the Big Five section as a within-person benchmark. The gap is large and structurally revealing. Respondents answered 59% of Big Five questions correctly and 31% of PE questions: a within-person gap of 28 percentage points. Meaningful PE knowledge is almost entirely confined to respondents with a perfect Big Five score, roughly 18% of the sample, and even they answered barely half the PE questions correctly.

Just over half the sample, 50%, is overconfident: they believe they know more than they do. As Kruger and Dunning (1999) showed, low performers lack the metacognitive ability to recognize their own errors; as Anderson et al. (2017) demonstrated, perceived literacy can drive financial behavior as forcefully as actual literacy. We find that both channels are at play. Overconfidence predicts PE investment intent independently of actual knowledge: moving from the bottom to the top quartile of the overconfidence distribution raises predicted intent by 15 pp, similar to the effect of holding an investment account. Investors are not being pushed toward PE by knowledge – they are being pulled by the illusion of financial competence.

The third piece is framing. Tversky and Kahneman (1981) established that equivalent choices generate different preferences depending on presentation; Shiller (2017) extended the insight to contagious economic narratives. We test this directly: respondents were randomly assigned to either a positive PE description – Hilton Hotels and its success story under PE ownership – or a negative one: Toys "R" Us, bought by PE and bankrupted. Positive framing raises intent by 7 percentage points on average, uniformly across the overconfidence distribution, but the effect is concentrated among respondents with superficial understanding of PE and is essentially absent among those who do not know anything about this asset class.

Taken together, the thesis contributes to three strands of the literature. To the financial literacy literature, it offers the first PE-specific knowledge instrument fielded at scale, showing that PE literacy is not a smooth extension of general financial competence but a distinct layer of understanding that many retail investors do not possess. To the overconfidence literature, it shows that miscalibration drives participation in an unfamiliar asset class independently of actual knowledge, with an effect size comparable to prior investment experience. To the policy debate on retail PE access, it provides direct evidence that the sophistication argument underpinning current regulatory liberalization rests on knowledge that, for many retail investors, remains limited.

The thesis proceeds as follows. Chapter 2 reviews the literature on financial literacy, overconfidence, framing, and private equity. Chapter 3 describes the survey, variable construction, and empirical approach. The remaining chapters present results and discussion.

2. Literature review

2.1 Financial literacy

2.1.1 Measuring financial literacy

Lusardi and Mitchell (2014) refer to financial literacy as "people's ability to process economic information and make informed decisions about financial planning, wealth accumulation, debt, and pensions" – a definition that has created the field's working standard. The OECD/INFE framework expands this definition to include attitudes and behaviors (Atkinson & Messy, 2012), and other scholars decompose financial literacy into knowledge, application, and self-assessment components (Hung et al., 2009; Huston, 2010; Remund, 2010). In practice, however, the empirical literature has largely converged on the cognitive dimension of the concept and measures it through a relatively small set of instruments.

The "Big Three" questions, introduced by Lusardi and Mitchell (2006) and refined in subsequent work (Lusardi & Mitchell, 2008, 2011), serve as the dominant instrument, measuring understanding of compound interest, inflation, and risk diversification. These items are favored for their simplicity, relevance, and suitability for large-scale surveys (Lusardi & Mitchell, 2014). The 2009 FINRA National Financial Capability Study extended this battery into the "Big Five" by adding mortgage interest and bond pricing (FINRA Investor Education Foundation, 2009; Hastings et al., 2013). Beyond these benchmarks, the S&P Global Financial Literacy Survey (Klapper et al., 2015) applies a closely related four-concept instrument across more than 140 economies, while more specialized measures include van Rooij et al. (2011)'s basic and advanced literacy modules, debt-literacy scales (Lusardi & Tufano, 2015), and refinements based on item response theory (Houts & Knoll, 2020; Knoll & Houts, 2012).

Financial literacy measurement is sensitive to question wording (van Rooij et al., 2011), while short financial knowledge batteries cover a narrow range of topics and difficulty levels, limiting their ability to discriminate across the full distribution of financial knowledge ability (Knoll & Houts, 2012). Overall, the Big Three remains the most commonly used benchmark for cross-study comparison, while the Big Five is mostly used in U.S. survey research.

2.1.2 Empirical patterns

Research demonstrates that financial literacy among the population is not uniform. Globally, 33 percent of adults meet the threshold introduced by the S&P Global FinLit survey – the largest measurement of financial literacy, conducted among 150,000 respondents in 140 countries (Klapper et al., 2015). Financial literacy ranged from 71% in Scandinavia to 13% in Yemen, with the U.S. at 58% in 2014. FINRA Foundation (2025) shows that after a long decline, U.S. financial literacy plateaued: 46% answered at least four of seven questions correctly in 2024 versus 45% in 2021, while the comparable five-question score was 32%, down from 42% in 2009. Part of this decline has been attributed to a rise in "Don't know" responses rather than to an increase in incorrect answers (FINRA Foundation, 2025).

Demographic patterns are strongly associated with financial literacy. Literacy follows an inverted-U over the life cycle, peaking around age 50 (Agarwal et al., 2009) and declining by roughly 1 percent annually after age 60 (Boyle et al., 2025; Finke et al., 2017). Education and income are also positively associated with literacy (FINRA Foundation, 2025; Klapper et al., 2015; Lusardi & Mitchell, 2014), though the income relationship is harder to interpret causally. The gender gap is widely discussed but often misinterpreted: women score below men on general-literacy surveys (Bucher-Koenen et al., 2017; Lusardi & Mitchell, 2014), yet a substantial portion of this gap reflects greater use of "Don't know" responses and lower confidence rather than lower underlying knowledge. Investment experience matters too: holders of savings, investment, or retirement products consistently outperform non-holders on the OECD/INFE battery (OECD, 2023).

Evidence on financial education remains contested. Fernandes et al. (2014) report that financial-education interventions explain only about 0.1 percent of the variance in downstream financial behavior, with effects becoming negligible after roughly twenty months. By contrast, Kaiser et al. (2022) re-examine the evidence using 76 randomized experiments and find materially larger treatment effects without the signs of rapid decay, though they also caution against assuming long-run sustainability.

2.1.3 Overconfidence and the perception-knowledge gap

The distinction between actual and perceived knowledge is central to behavioral finance. Kruger and Dunning (1999) provide the foundational mechanism: low performers lack the metacognitive ability to recognize their own errors. Subsequent work documents that perceived and actual knowledge predict financial behavior independently. Allgood and Walstad (2016) find that perceived literacy is a stronger predictor of costly credit-card behavior than actual literacy, while Parker et al. (2012) show that confidence predicts investment action conditional on objective knowledge. Anderson et al. (2017) similarly find that perceived, rather than actual, literacy drives precautionary saving and retirement planning, and that low-literacy respondents who see themselves as financially competent are less likely to seek professional advice.

Confidence also shapes vulnerability to fraud. While low financial literacy predicts susceptibility to lottery and sweepstakes scams, investment fraud victims are disproportionately college-educated, financially literate, and optimistic men (AARP, 2011; Consumer Fraud Research Group, 2006; Graham, 2014; Kieffer & Mottola, 2017). The mechanism most commonly advanced is overconfidence: investors with some knowledge overestimate their ability to evaluate complex offers (Gamble et al., 2014; Xiao et al., 2022).

Whether men are more overconfident than women is contested. Barber and Odean (2001) find that men trade roughly 45 percent more than women and earn lower net returns, consistent with greater overconfidence among men. Bucher-Koenen et al. (2017) document a literacy gap driven partly by disproportionate female use of the "Don't know" option, and Bucher-Koenen et al. (2021) attribute roughly one-third of that gap to lower confidence rather than lower knowledge. Lawrence et al. (2024) reach the opposite conclusion in a broader FINRA sample, finding women more overconfident once miscalibration is measured as the gap between self-perceived and objective knowledge. Bannier and Neubert (2016) show that confidence increases with education among men but decreases among women, while Kim et al. (2022) find women exhibit greater overconfidence among older adults specifically – suggesting the direction of the gender gap varies across samples and life stages.

The difference between self-assessed and actual score is an overestimation measure in the Moore and Healy (2008) taxonomy and does not capture overplacement or overprecision. A related concern is that guessing noise can read as miscalibration even absent systematic upward bias. The standard remediation, which this thesis follows, is to enter the actual score as a control alongside the overconfidence measure, so that residual variation in the gap captures miscalibration net of ability (Anderson et al., 2017).

2.1.4 Framing effects in financial decisions

Framing effects in finance trace back to prospect theory. Tversky and Kahneman (1981), building on Kahneman and Tversky (1979), show that mathematically equivalent choices generate different preferences depending on presentation. The literature on household finance framing effects can be organized into two broad approaches: structural and narrative framing.

The first concerns defaults, contribution rules, and disclosure. Benartzi and Thaler (2007) argue that households rely on heuristics rather than optimization. The Save More Tomorrow program (Thaler & Benartzi, 2004) reframes higher saving as a commitment out of future raises, raising contribution rates from 3.5 to 13.6 percent over 40 months in its first implementation. Beshears et al. (2011) and Choi et al. (2010) show that even simplified disclosure and fully transparent fee information fail to alter allocations or minimize costs – even among MBA students and Harvard staff – suggesting that how a choice is structured matters more than what information is provided.

The second channel is narrative. Tversky and Kahneman (1973) introduce the availability heuristic – the tendency to assess frequency or likelihood based on the ease with which relevant examples are recalled; Shiller (2017, 2019) extends it into a theory of contagious economic narratives. Appelbaum and Batt (2014) document the polarized discourse around private equity, in which firms appear alternately as rescuers or predators. Framing effects are not expected to operate uniformly across investors: Agnew and Szykman (2005) show that low-knowledge respondents accept default allocations at far higher rates than high-knowledge ones, providing direct evidence that framing susceptibility can vary with financial literacy.

2.2 Private equity

2.2.1 Defining private equity

Private equity (PE) generally refers to investing in non-public companies, often through closed-end limited partnerships in which the PE firm acts as the general partner (GP) and investors provide capital as limited partners (LPs) (Metrick & Yasuda, 2010). The term can also include taking public companies private, so it is broader than only buying privately held firms (Kaplan & Strömberg, 2009; Metrick & Yasuda, 2010). The main strategy families are buyouts, growth equity, and venture capital, with buyouts being the dominant and most studied segment in the literature (Kaplan & Strömberg, 2009).

These funds are typically closed-end vehicles, so investors usually cannot redeem their money at will as they could in a mutual fund or ETF (Kaplan & Strömberg, 2009). Instead, capital is committed for a long fund life, often around 5–10 years, with money returned only as portfolio companies are exited (Kaplan & Strömberg, 2009). PE funds commonly use a "2 and 20" fee structure: a management fee is charged annually, and carried interest gives the manager a share of profits above a contractual threshold (Metrick & Yasuda, 2010). Carried interest is the GP's share of profits after investors have received back their capital and, in many funds, a preferred return or hurdle; an 8 percent hurdle is commonly cited in buyout funds (Metrick & Yasuda, 2010; Phalippou & Gottschalg, 2009).

Leverage is central to leveraged buyouts, as the PE sponsor borrows money to finance part of the acquisition, and the acquired company carries that debt after the transaction. This structure can amplify equity returns when performance is strong, but it also increases financial risk and bankruptcy exposure (Kaplan & Strömberg, 2009).

On performance, the literature does not converge on a univocal answer. Kaplan and Schoar (2005) find that average PE returns are roughly in line with public markets, while Harris et al. (2014) report buyout outperformance of more than 3 percent per year using Burgiss data. Robinson and Sensoy (2016) show that conclusions depend materially on the performance measure used, and Phalippou (2020) argues that headline outperformance often disappears once benchmarks and fees are adjusted. Against this backdrop, the present thesis uses 3 percent as a practical benchmark drawn from one influential estimate in the literature, while recognizing that expected outperformance remains contested.

2.2.2 Democratization of private equity

The question of whether retail investors can competently evaluate PE has become more pressing as access has widened rapidly. In the United States, the SEC's 2020 amendment to the accredited-investor definition marked a partial shift away from wealth as the sole proxy for sophistication, introducing qualification pathways based on credentials such as the Series 7, 65, and 82 licenses (U.S. Securities and Exchange Commission, 2020). Around the same time, private-markets firms expanded distribution through vehicles more accessible than the traditional closed-end partnership – interval funds, non-traded business development companies, and semi-liquid or feeder structures targeted at high-net-worth and retail channels (Aramonte & Avalos, 2021; SEC Investor Advisory Committee, 2025).

This democratization has been reinforced by the structural shift in retirement saving. As the U.S. system has moved from employer-directed defined-benefit provision toward participant-directed defined-contribution saving, allocation decisions have fallen increasingly on households rather than sponsors (Hacker, 2011; Poterba et al., 2007). The extension of private-market products into retail and retirement channels means households are now deciding not only whether to save, but which alternative vehicles will receive those savings (Aramonte & Avalos, 2021; SEC Investor Advisory Committee, 2025).

The shift accelerated further in 2025, when President Trump signed an Executive Order directing the Department of Labor to facilitate alternative-asset access in defined-contribution plans and to reconsider prior guidance limiting PE inclusion (Trump, 2025). In Europe, ELTIF 2.0 points in the same direction by broadening retail access to long-term illiquid assets, indicating that the liberalization of private-markets access is international even if the present empirical setting is U.S.-focused (European Parliament and Council of the European Union, 2023). Historically, PE access was constrained primarily by wealth thresholds. Increasingly, it is being justified by claimed sophistication, which presupposes that such sophistication can be measured.

2.3 Hypotheses development

Existing financial literacy measures were developed around publicly traded instruments. The standard Big Five instrument is general rather than context-specific (Lusardi & Mitchell, 2014), while PE is an alternative, structurally distinct asset class.

The domain-specific literacy literature suggests this distinction matters empirically. Debt literacy predicts debt outcomes more strongly than general financial knowledge (Lusardi & Tufano, 2015), and crypto literacy predicts crypto-specific behavior beyond what general measures capture (Jones et al., 2024). Yet despite these parallels, no validated instrument measuring PE literacy currently exists.

This thesis addresses that gap. By measuring PE-relevant knowledge in a retail sample and linking it to investment intent, it offers a first attempt to understand what the general public actually knows, and what it thinks it knows, about private equity. A further question is whether overconfidence compounds susceptibility to narrative framing, exploring whether the gap between perceived and actual knowledge is what makes retail investors responsive to how PE is presented rather than what it is.

The structural gap implies that PE literacy should be substantially lower than general literacy within the same investors. No validated PE-specific instrument exists in the published literature; the seven-item battery here is, to our knowledge, the first publicly fielded attempt, making the hypothesis below both a directional prediction and a foundational measurement test. The additional claim – that PE knowledge becomes meaningful only at the top of the general-literacy distribution – follows from the precondition logic of the domain-specific literature but is tested empirically.

Hypothesis 1. Within the same respondents, PE literacy is significantly lower than general financial literacy, and meaningful PE knowledge is concentrated among respondents with a perfect Big Five score; below this threshold, PE literacy is limited.

The first hypothesis concerns measurement; the second concerns behavior. General literacy predicts stock-market participation (van Rooij et al., 2011) and retirement savings, wealth accumulation, and mortgage choice (Lusardi & Mitchell, 2014), and the domain-specific analog is equally established (Jones et al., 2024; Lusardi & Tufano, 2015). If PE literacy captures genuine structural understanding, it should likewise translate into willingness to hold PE – a prediction not previously tested.

Hypothesis 2. Higher PE literacy is associated with greater willingness to include PE in a retirement plan, controlling for overconfidence and demographics.

While the second hypothesis concerns what investors know, the third concerns the gap between what they know and what they think they know. Kruger and Dunning (1999) establish the foundational mechanism: low-performing individuals not only make poor judgments but lack the metacognitive ability to recognize their errors. In financial settings, perceived and actual knowledge frequently predict behavior independently. The closest precedent is Anderson et al. (2017), who find that perceived rather than actual literacy drives precautionary saving and retirement planning in a financially active sample. Whether this dominance of perception over knowledge intensifies in a less familiar asset class is a direct and testable extension.

Hypothesis 3. The gap between self-assessed and actual scores is a significant positive predictor of willingness to invest in PE, independently of literacy.

Tversky and Kahneman (1981) show that mathematically equivalent choices elicit different preferences depending on presentation, and narrative framing is particularly powerful when investors cannot evaluate a product technically and must rely on vivid stories (Shiller, 2017). The PE context is especially susceptible: Appelbaum and Batt (2014) document polarized public discourse in which firms are depicted alternately as rescuers or destroyers, so retail investors encounter PE primarily through simplified narratives rather than clean performance data. Framing effects are heterogeneous in knowledge – Agnew and Szykman (2005) find that low-knowledge respondents accept default allocations at ten times the rate of high-knowledge ones. Overconfident investors lack both the self-awareness to dismiss optimistic narratives and the domain knowledge to evaluate them critically, so positive PE framing should land harder on them. This extends the overconfidence-participation channel into a narrative domain and has not been tested previously in PE.

Hypothesis 4. Positive framing increases investment willingness more for overconfident respondents than for calibrated ones.

3. Research settings and methodology

3.1 Survey design

The instrument is organized in four blocks of items administered in fixed order. The first collects demographics: age (Q1), gender (Q2), education (Q3), employment status (Q4), annual gross income (Q5), investment experience (Q6), and prior finance or economics education (Q7).

The second replicates the Big Five items verbatim: compound interest (Q8), inflation (Q9), risk diversification (Q10), mortgage amortization (Q11), and bond pricing (Q12). General-literacy performance can therefore be benchmarked directly against published U.S. norms. Each item carries a “Don’t know” option, in line with Lusardi and Mitchell (2014) and for the calibration reasons set out in Section 2.1.3.

The third block contains the framing experiment and the PE literacy module. Each respondent is randomly assigned, with equal probability, to one of two short framing descriptions. The positive frame (Q13) describes Hilton’s 2007 leveraged buyout and its success; the negative frame (Q14) describes the 2005 leveraged buyout and eventual bankruptcy of Toys “R” Us. Each closes with a yes/no awareness question. Assignment is between-subjects, following Tversky and Kahneman (1981); the achieved split is Hilton $n = 501$ (48.5 %) and Toys “R” Us $n = 531$ (51.5 %). Respondents then answer seven custom PE literacy items (Q15-Q21) covering fund structure, fee structure, liquidity and lock-up, leverage, long-run returns relative to public markets, the “2 and 20” convention, and carried interest. Wording follows the Lusardi-Mitchell format, and the seven concepts span the structural features any retail investor offered PE exposure would need to evaluate, drawing on Phalippou (2009, 2020) and Metrick and Yasuda (2010). To our knowledge, this is the first publicly fielded PE-specific literacy instrument.

The fourth block measures investment intent. Q22 asks how likely respondents would be to include PE in their retirement plan if such products were available, on a six-point scale running from “Definitely no” to “Definitely yes”. Q22 is the source of the binary outcome variable defined in Section 3.3.3. Q23 is an optional slider asking respondents to estimate how many of the preceding 12 knowledge questions they answered correctly. Its 0–12 scope, rather than separate PE-only and Big Five-only sliders, has consequences we return to in Section 5.6.

3.2 Sample

The data come from a Qualtrics (2025) online panel of 1,032 U.S. retail investors fielded in November 2025. Probit specifications use 961 respondents after listwise deletion on the controls and exclusion of “Other” or “Prefer not to say” on gender. Overconfidence analyses use 963, since sixty-nine respondents skipped the optional self-assessment slider.

Table 3.1: Sample Demographics

Variable	Category	N	%
Age	18–24	210	20.3
	25–34	192	18.6
	35–44	210	20.3
	45–54	210	20.3
	55–64	210	20.3
Gender	Male	505	48.9
	Female	525	50.9
	Other	1	0.1
	Prefer not to say	1	0.1
Education	High school/GED	320	31.0
	Vocational/technical school	151	14.6
	Bachelor's degree	312	30.2
	Master's degree or above	249	24.1
Employment status	Employed full-time	591	57.3
	Employed part-time	126	12.2
	Self-employed/entrepreneur	64	6.2
	Unemployed	134	13.0
	Student/retired	117	11.3
Annual gross income (USD)	Less than $34,999	257	24.9
	$35,000–$49,999	109	10.6
	$50,000–$74,999	104	10.1
	$75,000–$99,999	78	7.6
	$100,000–$149,999	194	18.8
	$150,000–$199,999	173	16.8
	$200,000–$499,999	100	9.7
	$500,000 or more	17	1.6
Investment account	Yes – actively invests	444	43.0
	Yes – rarely makes changes	213	20.6
	No – has invested in the past	99	9.6
	No – has never invested	276	26.7
Finance/economics education	Yes, in high school	230	22.3
	Yes, in college/university	347	33.6
	Yes, through self-study	100	9.7
	No	355	34.4

Notes: Sample summary statistics for the survey sample ($N = 1{,}032$). Each panel reports the unweighted count of respondents in each category and the corresponding share of the full sample. Income bands are reported in U.S. dollars as elicited in the survey instrument. Columns within each panel sum to 1,032 (± 0.1 pp due to rounding).

Demographics are shown in Table 3.1. The sample is 50.9 % female and 48.9 % male, roughly uniform across the 18–24 through 55–64 age bands. Education is concentrated at high school (31.0 %), bachelor's (30.2 %), and master's or above (24.1 %). Income is bimodal: 24.9 % below \$35,000 and 46.9 % above \$100,000. Investment-account ownership stands at 63.6 % and turns out to be the largest single predictor in the probit reported in Section 4.4.

External validation comes from the Big Three subscore: compound interest (Q8), inflation (Q9), and risk diversification (Q10). 35.7 % of the sample answer all three correctly, within rounding of the U.S. benchmark of approximately 34 % in Lusardi and Mitchell (2014). The Qualtrics panel is non-probabilistic and reweighted only through internal quota mechanics, so the Big Three match is reassurance rather than formal external validity.

3.3 Variable construction

3.3.1 Literacy scores

Two literacy scores are constructed per respondent, since H1 needs the within-person comparison and Specification 1 in Section 3.4.3 needs them as separate regressors. The Big Five score sums correct answers to Q8-Q12 (range 0 to 5); the PE literacy score sums correct answers to Q15-Q21 (range 0 to 7). Following Lusardi-Mitchell, "Don't know" is coded as incorrect for scoring but retained as a separate descriptive category. For the H1 paired test, each score is expressed as a proportion correct:

$$\text{big5\_prop} = \frac{\text{big5\_score}}{5}, \qquad \text{pe\_prop} = \frac{\text{pe\_score}}{7}. \tag{3.1}$$

Both then sit on the unit interval and are directly comparable.

Cronbach's α for the seven-item PE literacy scale is 0.61 (k=7,N=1,032), which is within the acceptable range for short knowledge scales scored as correct/incorrect in exploratory research. Because the instrument is designed to capture multiple dimensions of PE understanding rather than a single highly homogeneous construct, extremely high internal consistency is neither expected nor necessarily desirable. No individual item materially increases α when removed, so the full scale is retained.

3.3.2 Framing treatment indicator

The framing indicator equals 1 for respondents assigned the Hilton description and 0 for those assigned Toys "R" Us. Since framing is administered between-subjects, the framing coefficient in Section 4.4 captures the average treatment effect of positive versus negative narrative framing on stated PE investment intent.

3.3.3 Dependent variable: investment intent

The binary outcome variable is built from Q22 using two thresholds. Threshold A is primary, and Threshold B is reported as a robustness check.

Primary outcome (Threshold A) = positive intent – invest $= 1$ if Q22 $\in$ {"Definitely yes", "Probably yes"}. We use "positive intent" throughout to denote the union of the two affirmative Likert categories on the PE-investment intent question. A conservative commitment threshold (Threshold B = "Definitely yes" only) is reported as robustness. The baseline rate at Threshold A is 35.1 % (362 of 1,032) – the combination of the two affirmative Likert categories.

Threshold B is the commitment-only measure: invest $= 1$ only for respondents selecting "Definitely yes." The baseline rate is 16.6 % (171 of 1,032). Qualitative probit conclusions are unchanged across thresholds (Section 4.7.1).

3.3.4 Overconfidence gap

Following the Moore and Healy (2008) overestimation taxonomy, and operationalized as in Anderson et al. (2017), overconfidence is the difference between self-assessed and actual number of correct answers:

$$\text{overconfidence\_gap}_i = \text{Q23}_i - \text{total\_score}_i \tag{3.2}$$

where Q23 is the respondent's estimate of correct answers across the combined Big Five plus PE pool. Q23 is empirically already on a 0–12 integer scale (mean 6.12, median 6, range 0–12), so no rescaling applies. Positive values indicate overconfidence, zero calibration, negative values underconfidence. The continuous gap enters Section 3.4.3 directly.

A categorical version follows Kim et al. (2020) (hereafter KLH) building on Robb et al. (2015) and Xia et al. (2014), with median splits on actual score (median $= 5$) and self-assessment (median $= 6$) yielding four cells: *Correctly Confident* (high actual, high self), *Overconfident* (low actual, high self), *Underconfident* (high actual, low self), and *Correctly Unconfident* (low actual, low self). The KLH binary variable used in the robustness check of Section 4.7.3 equals 1 only for the Overconfident cell:

$$\text{oc\_klh}_i = \mathbb{1}[\text{actual}_i < \tilde{m}_{\text{actual}}] \cdot \mathbb{1}[\text{self}_i \geq \tilde{m}_{\text{self}}] \tag{3.3}$$

3.3.5 Control variables

Age enters as a continuous variable using the midpoint of each reported band (21.5, 29.5, 39.5, 49.5, 59.5; the 65+ band is coded at 70). The female indicator equals 1 for female respondents and 0 otherwise; respondents selecting "Other" or "Prefer not to say" are excluded from gender-specific subgroup tests but retained in pooled regressions with an imputed value of 0. Income is ordinal, 1 to 8, from less than $34,999 through $500,000 or more; education is ordinal, 1 to 4, from high school/GED through master's degree or above. The investment-account indicator (`has_invested`) equals 1 for respondents reporting active, infrequent, or former investment activity and 0 for those who have never invested.

3.4 Empirical approach

The four hypotheses divide naturally between a within-person comparison of literacy scales (H1) and a joint probit specification covering literacy, overconfidence, and framing (H2, H3, H4).

3.4.1 Within-person comparison of literacy scores (H1)

Every respondent completes both modules, so the two scores can be compared within a person and unobserved respondent-level heterogeneity absorbed. The main test is a paired-samples t-test on big5_prop and pe_prop, reported alongside Cohen's d. Both distributions are non-normal, with the PE score sharply right-skewed, so a Wilcoxon signed-rank test is reported as a non-parametric robustness check.

The respondents are categorized into three Big Five tiers (*Low*: 0–2; *Medium*: 3–4; *High*: 5), and we plot the mean PE score per tier and per unit Big Five score. A linear regression would impose the wrong functional form if PE knowledge emerges only above a threshold of general literacy, so a step-function visualization lets the non-linearity speak for itself. The tier comparison is tested with one-way ANOVA, and Kruskal-Wallis as a non-parametric robustness check.

Q15, the foundational PE-fund definition item, supplements the threshold-tier analysis as a single-item necessary-condition test. We compare the PE and Big Five scores of respondents who answered Q15 correctly against those who did not, using Welch t-tests and reporting Cohen's d.

3.4.2 Overconfidence descriptives

Section 4.3 documents the distribution of the overconfidence gap and its variation across demographic subgroups. We report mean and median of the gap between self-assessed and objective literacy across five demographic dimensions – gender, age, education, income, and investment experience – together with the share of each subgroup we term *Overconfident*: those whose self-assessed literacy exceeds their objective score. Group assignment follows the four-category calibration scheme of Kim et al. (2020). Welch t-tests handle the gender comparison, and one-way ANOVA handles the others.

The gender comparison tests the asymmetry that men are more confident aligning with Barber and Odean (2001) rather than with Lawrence et al. (2024). Whether it holds in an unfamiliar asset class such as PE, where "Don't know" is the most chosen option, is the empirical question.

3.4.3 Probit specifications (H2, H3, H4)

Probit is preferred over OLS since the outcome is binary, and over logit by convention in the financial-economics literature we draw on, following Anderson et al. (2017) and Kim et al. (2020); logit results are nearly indistinguishable. We estimate two specifications and one exploratory variant.

Specification 1 estimates the main effects of PE literacy, overconfidence, and framing, with general financial literacy partialled out alongside PE literacy:

$$\begin{aligned}\Pr(\text{invest}_i = 1) = \Phi\big(&\alpha + \beta_1 \text{pe\_score}_i + \beta_2 \text{big5\_score}_i \\ &+ \beta_3 \text{oc\_gap}_i + \beta_4 \text{framing}_i + \gamma' \mathbf{X}_i\big)\end{aligned} \tag{3.4}$$

where Φ is the standard-normal cumulative distribution function and $\mathbf{X}$ stacks age, female, income, education, and investment experience. H2 predicts $\beta_1 > 0$; H3 predicts $\beta_3 > 0$; the framing coefficient β_4 is the average treatment effect of positive versus negative narrative framing on stated intent.

Specification 2 adds the framing $\times$ overconfidence interaction:

$$\begin{aligned}\Pr(\text{invest}_i = 1) = \Phi\big(&\alpha + \beta_1 \text{pe\_score}_i + \beta_2 \text{big5\_score}_i + \beta_3 \text{oc\_gap}_i \\ &+ \beta_4 \text{framing}_i + \beta_5 (\text{framing}_i \times \text{oc\_gap}_i) + \gamma' \mathbf{X}_i\big)\end{aligned} \tag{3.5}$$

H4 predicts $\beta_5 > 0$, since positive framing should amplify the overconfidence effect when overconfident respondents lack the PE-specific knowledge to evaluate optimistic narratives critically. The interaction is tested with a likelihood-ratio test of Specification 2 against Specification 1:

$$\text{LR} = 2(\ell_2 - \ell_1) \sim \chi^2_1 \tag{3.6}$$

Predicted probabilities for a 2×2 grid of framing arm and overconfidence level (25th and 75th percentiles of the gap distribution) are reported in Section 4.6.

Specification 2b is exploratory. Following the H4 null in Section 4.6.1, we estimate a not-pre-registered specification that swaps the H4 interaction for framing $\times$ PE-literacy, with PE literacy entered as a three-level categorical moderator rather than the continuous score. Respondents are assigned to one of three mutually exclusive buckets:

- **Literate**: pe_score ≥ 4 – anchored to the mean PE score of 3.60 in the subgroup with perfect Big Five (Section 4.2.2), the only tier in which PE knowledge departs meaningfully from the guessing baseline.
- **Moderate**: pe_score < 4 *and* Q15 correct – respondents who recognize the foundational definition of a private-equity fund but lack the mechanics.
- **Illiterate**: the residual (Q15 wrong) – respondents who fail the gating item that Section 4.2.2 establishes as the cleanest single-item summary of the precondition story.

With Illiterate as the reference category, Specification 2b is

$$\begin{aligned}\Pr(\text{invest}_i = 1) = \Phi\big(\alpha + \beta_1 \text{big5\_score}_i + \beta_2 \text{oc\_gap}_i + \beta_3 \text{framing}_i \\ + \beta_4 \text{Moderate}_i + \beta_5 \text{Literate}_i \\ + \beta_6(\text{framing}_i \times \text{Moderate}_i) + \beta_7(\text{framing}_i \times \text{Literate}_i) + \gamma' \mathbf{X}_i\big)\end{aligned} \quad (3.7)$$

The continuous pe_score is omitted as it is collinear with the bucket dummies. The mechanism, H4b, is comprehension-gating: the framing description only updates priors for respondents with enough PE literacy to interpret the implied causal claim. H4b therefore predicts $\beta_6, \beta_7 > 0$, tested jointly with a likelihood-ratio test of Specification 2b against the restricted model without interactions (Equation 3.6, distributed χ^2_2).

Standard errors are unclustered, since the sample is drawn independently. Marginal effects are reported as averages over the sample distribution rather than evaluated at sample means.

4. Results

The central finding concerns the gap between U.S. retail investors' actual knowledge of private equity and their perceived knowledge of it, and how this gap influences their reported willingness to include PE in their retirement portfolios. The main outcome measure used throughout is Threshold A – a positive intent to invest in PE, captured by "Definitely yes" and "Probably yes" responses to Q22 (willingness to invest in PE), which 35.1% of the sample chose. Threshold B, a stricter commitment-only category limited to "Definitely yes" responses (16.6%), is reported as a robustness check in Section 4.7.1.

4.1 Knowledge scores: distributions and per-question correctness

4.1.1 Score distributions

Respondents answer 59.0 % of Big Five items correctly on average (2.95 of five, SD = 1.48, median = 3), and 61.7 % of the Big Three subset (1.85 of three, SD = 1.05, median = 2), against 31.3 % of PE items (2.19 of seven, SD = 1.74, median = 2). The shapes of the two distributions preview H1. Big Five scores are approximately bell-shaped and peak around 3 and 4. The PE distribution is sharply right-skewed: 20.8 % of respondents score zero and only 0.2 % reach a perfect seven. Figure 4.1 plots both distributions and Table 4.1 reports the summary statistics.

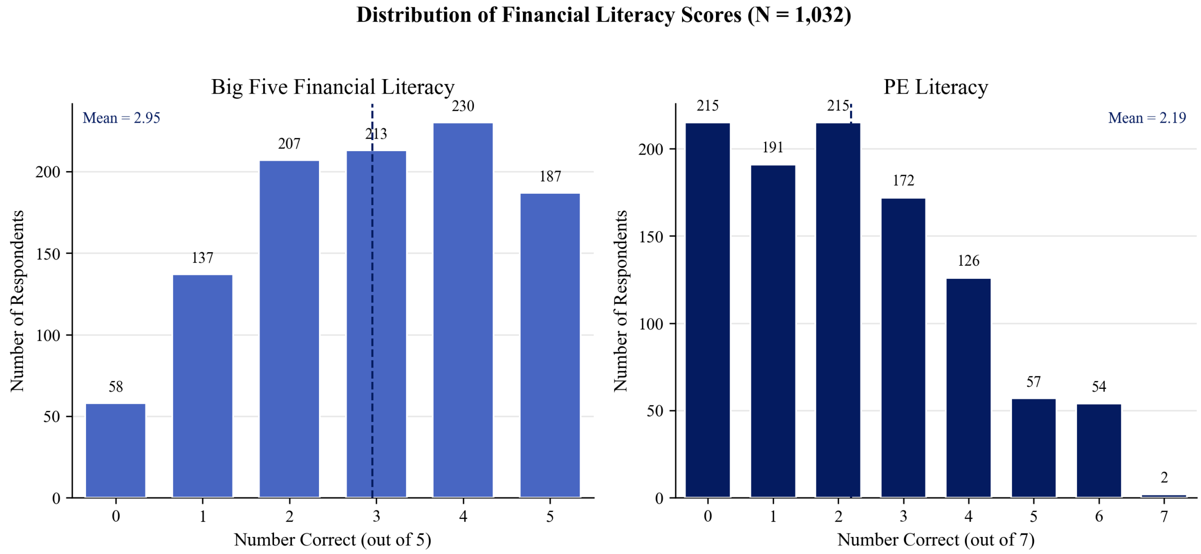

Figure 4.1: Distribution of Five and PE literacy scores ($N = 1{,}032$).

Mean PE literacy varies systematically across these cells. It rises with education from 1.72 of seven (24.6 %) at high school to 2.81 (40.2 %) at master's- degree or above, and with income from 1.61 (23.0 %) below $35,000 to 2.70 (38.6 %) in the $150,000–$199,999 band, plateauing thereafter. Prior investment experience is the sharpest single cut: 2.71 (38.7 %) for active investors against 1.46 (20.9 %) for those who have never invested. Men score 2.42 against 1.98 for women, the gender pattern returned to in Section 4.4.2. Age is non-monotonic and modest, peaking at 2.74 in the 35–44 band and dipping to 1.80 in the 55–64 band. Big Five literacy shows broadly the same ordering across education, income, and investment experience but with a narrower spread (2.41 to 3.36 across education; 2.29 to 3.30 across investment experience).

Table 4.1: Summary statistics for Big Three, Big Five, and PE literacy scores

Measure	N	Mean	Median	SD	Min	Max	Range	% Perfect	% Zero
Big Three	1,032	1.85	2.0	1.05	0	3	0–3	35.7	13.0
Big Five	1,032	2.95	3.0	1.48	0	5	0–5	18.1	5.6
PE Literacy	1,032	2.19	2.0	1.74	0	7	0–7	0.2	20.8

Notes: Score distributions for the three knowledge measures used in the analysis ($N = 1{,}032$). Big Three is the Lusardi and Mitchell (2014) subset of compound interest (Q8), inflation (Q9), and risk diversification (Q10). Big Five sums correct answers to Q8-Q12 and PE Literacy sums correct answers to Q15-Q21. "% Perfect" is the share scoring at the maximum of each scale; "% Zero" is the share answering none correctly.

4.1.2 Per-question correctness and "Don't know" rates

Figure 4.2 plots per-question correct rates across both modules. On the Big Five, compound interest (Q8) at 77.6 % and mortgage amortization (Q11) at 77.4 % are the best-answered questions; bond pricing (Q12) at 32.8 % is the hardest. PE per-item rates range from 17.4 % on Q19, the PE-versus-public returns item, to 47.1 % on Q16, the fee-structure item. The leverage item (Q18) is the second-hardest at 18.8 %, reflecting the sophistication needed to define where in a leveraged buyout the acquisition debt actually sits. Q19 is also the most often skipped: 38.5 % select “Don’t know,” roughly double the rate of any Big Five item.

A more revealing pattern sits inside the PE module. Q15, the foundational PE-fund definition question asking which of three statements best describes a private-equity fund, is correct for only 31.2 % of respondents – substantially fewer than the 47.1 % who answer the fee item correctly. More respondents know that PE charges fees than know what PE is. It resembles surface-attribute recognition not backed by structural understanding.

Across the PE module, “Don’t know” rates average 32 % against approximately 15 % on the Big Five. The pattern is evidence of genuine unfamiliarity rather than guesswork: respondents know they do not know PE. The full per-question correct, incorrect, and “Don’t know” breakdown is in Figure 4.2 and Figure 6.1 of Appendix B.

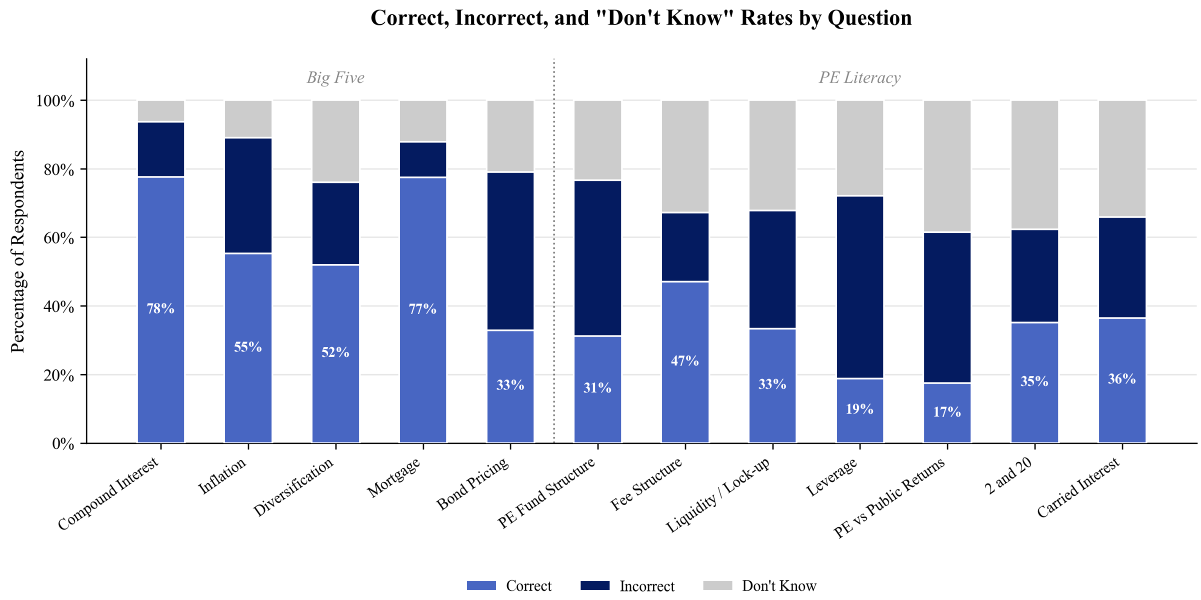

Figure 4.2: Per-item correct rates across the Big Five and PE modules.

4.2 Hypothesis 1: PE literacy is lower than general financial literacy

PE literacy sits well below general financial literacy in the same respondents, and the gap is structured rather than uniform. Knowledge concentrates among the small subgroup that scores perfectly on the Big Five, and the foundational PE question on what is PE (Q15) functions as a prerequisite to the rest of the scale.

4.2.1 *Within-person paired test*

Table 4.2: H1 paired tests: within-person Big Five vs. PE literacy proportions

Test	Statistic	p-value (one-tailed)	Effect size
Paired t-test	$t(1031) = 28.418$	5.81×10^{-132}	Cohen's $d = 0.885$
Wilcoxon signed-rank	$W = 453{,}918$	6.69×10^{-108}	–

Notes: Within-person tests of H_0 : big5_prop = pe_prop against H_1 : big5_prop > pe_prop, N = 1,032. The Wilcoxon signed-rank test is reported as a non-parametric robustness check given the right-skew of the PE distribution.

Table 4.2 reports the paired-test results. The within-person Big Five minus PE proportion-correct gap averages 27.7 percentage points (59.0 % vs. 31.3 %). The paired-samples test returns $t(1031) = 28.42$, $p < 10^{-131}$, Cohen's $d = 0.88$, large by conventions. The non-parametric Wilcoxon signed-rank test returns $W = 453{,}918$, $p < 10^{-107}$. The result is not tail-driven: 79.2 % of respondents have a positive gap, 2.8 % are tied, and 18.0 % score higher on PE in proportion-correct terms (Figure 4.3).

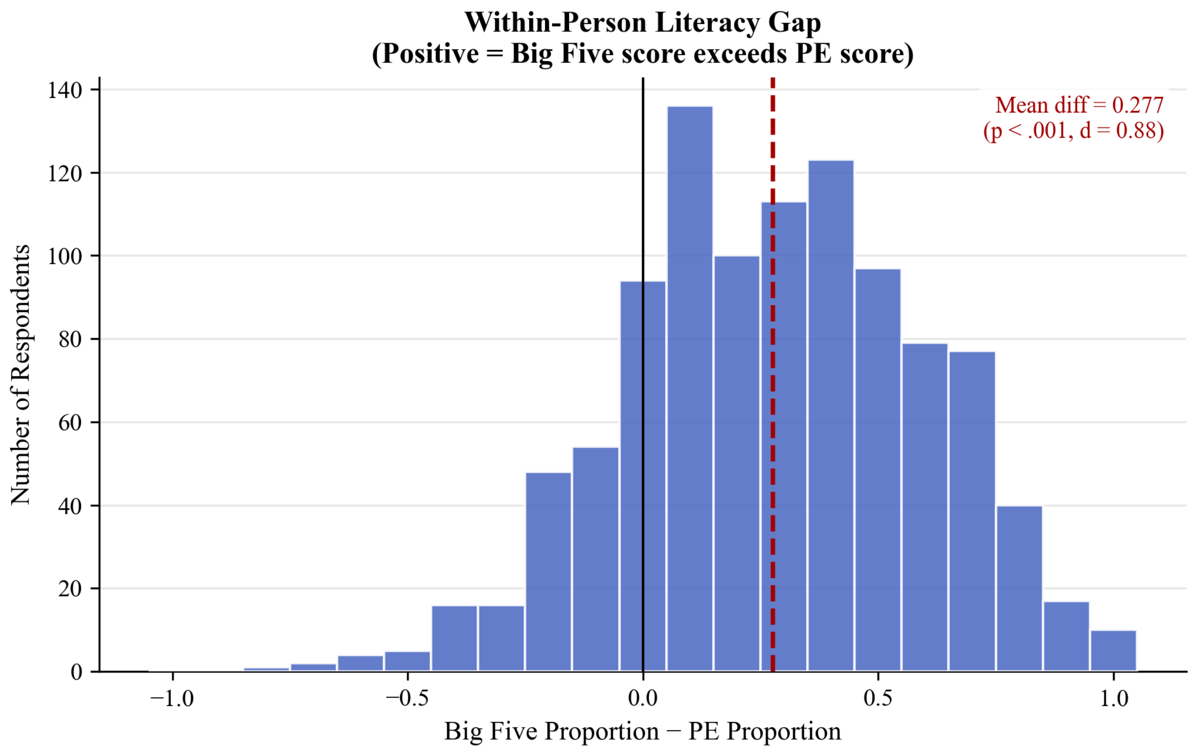

Figure 4.3: Within-person Big Five minus PE proportion-correct distribution.

4.2.2 *Threshold-tier analysis: general literacy as a precondition*

The within-person gap establishes the divergence; the structure of that divergence is the more interesting finding. Figure 4.4 plots mean PE score against Big Five score at each of the six possible Big Five values and reveals a strong threshold pattern. Respondents scoring 0 through 4 achieve mean PE scores of 1.14, 1.69, 1.81, 1.96, and 2.17, an almost flat profile. Respondents with a perfect Big Five of five reach a mean PE of 3.60, around 1.7 times the benchmark for the group who scored 4 on Big Five. The first four increments of general literacy add on average 0.26 correct PE answers; the fifth increment alone adds 1.43.

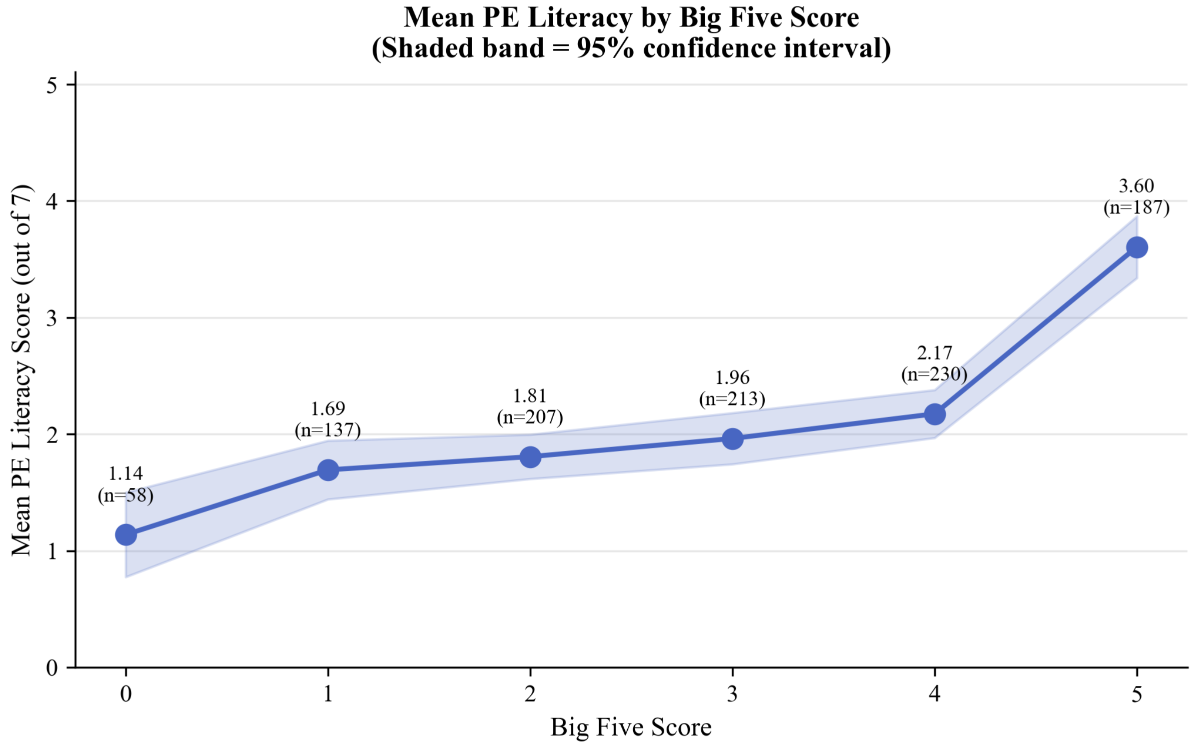

Figure 4.4: Mean PE literacy score by Big Five score, with 95% confidence band.

Table 4.3: Mean PE literacy score by Big Five tier

Big Five tier	N	Mean PE score	Median PE score	SD
Low (0–2)	402	1.67	1.5	1.45
Medium (3–4)	443	2.07	2.0	1.61
High (5)	187	3.60	4.0	1.85

Notes: Mean PE literacy score (range 0–7) by Big Five tier, $N = 1{,}032$. Tiers are constructed on the Big Five score: Low (0–2), Medium (3–4), High (5 of 5). One-way ANOVA across the three tiers gives $F(2, 1029) = 95.61$, $p < 10^{-38}$; the non-parametric Kruskal-Wallis test gives $H = 133.30$, $p < 10^{-29}$.

General financial literacy is a necessary but not sufficient condition for PE-specific knowledge. Below perfect Big Five performance, PE literacy is statistically indistinguishable from a guessing baseline; even perfect general literacy does not lift PE understanding above 52 %. The structural gap will not close without differentiated PE education.

Q15 – the foundational question asking what a PE fund is, prior to any item about fees, lock-up, leverage, or returns – functions empirically as a gate. Respondents who fail it score 1.65 of seven on the PE questions versus 3.40 for those who answer correctly (Welch $t = 16.29$, $p = 4.1 \times 10^{-49}$, Cohen's $d = 1.15$), much larger than the within-person H1 gap on a single binary partition. The same split runs symmetrically on the Big Five: Q15-correct respondents score 3.69 of five against 2.61 for those who fail it (Welch $t = 11.79$, $p < 10^{-28}$, Cohen's $d = 0.78$). The gating effect at $d = 1.15$ and the necessary-condition effect at $d = 0.78$ reaffirm that Q15 is the prerequisite question (Table 4.4).

Table 4.4: Welch t-tests of PE and Big Five scores by Q15 correctness

Outcome	Group	N	Mean	SD
Big Five score	Q15 correct	322	3.69	1.34
Big Five score	Q15 wrong	710	2.61	1.42
PE score	Q15 correct	322	3.40	1.66
PE score	Q15 wrong	710	1.65	1.48

Notes: Group means and standard deviations of the Big Five score (0–5) and PE literacy score (0–7) for respondents who answered Q15 (the foundational PE-fund definition item) correctly versus incorrectly.

4.3 Overconfidence: distribution and demographic patterns

4.3.1 Calibration and the gap distribution

Of 1,032 respondents, 963 (93.3 %) gave a valid response to the optional self-assessment item Q23 and we retain them for the overconfidence analysis. Mean self-assessed score is 6.12 of twelve and mean actual score (Big Five plus PE combined) is 5.19, giving an overconfidence gap with mean $+0.93$, SD $= 3.32$, and median $= +1$. 50.4 % of respondents are overconfident, 13.6 % calibrated, and 36.0 % underconfident. The distribution is right-skewed beyond the median: at the 75th, 90th, and 95th percentiles the gap is $+3$, $+6$, and $+8$ (Figure 4.5).

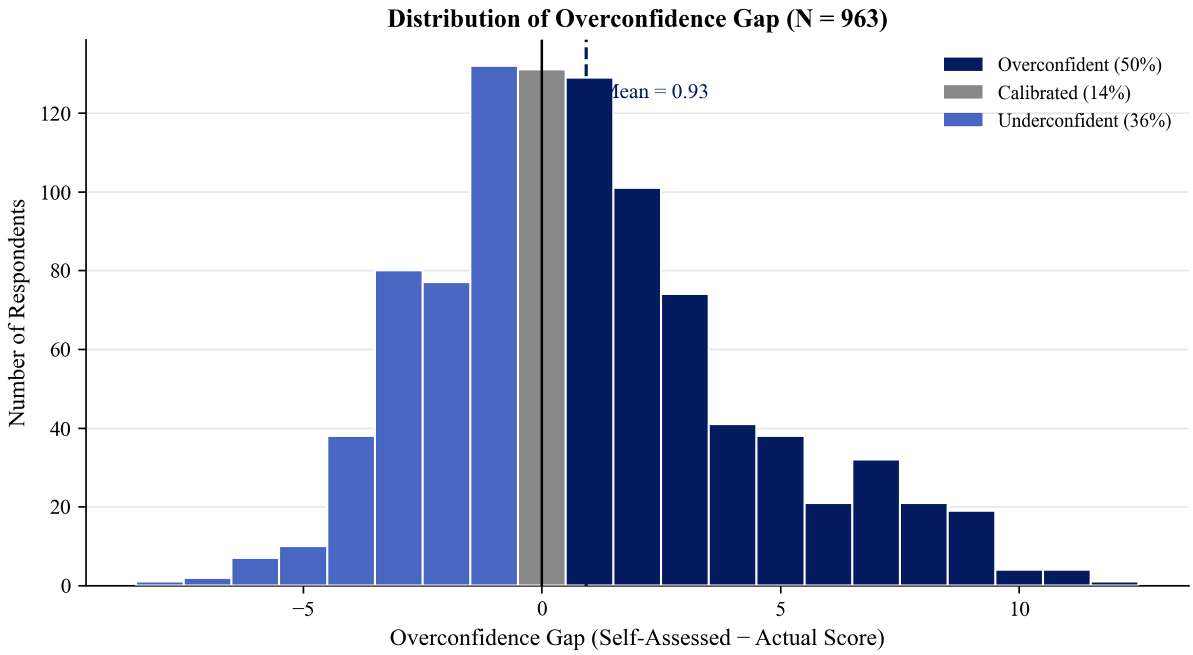

Figure 4.5: Distribution of the overconfidence gap (self-assessment minus actual score).

4.3.2 KLH 2×2 typology

The Kim et al. (2020) four-group classification cuts the sample on actual knowledge crossed with self-assessment. *Correctly Confident* respondents, above the sample median on both actual knowledge (mean 7.36) and self-assessment (mean 8.40), form the largest cell at 37.3 % ($n = 359$). *Correctly Unconfident* respondents, low on both at 2.64 and 2.82, are 25.4 % ($n = 245$). *Overconfident* respondents, low on knowledge but high on self-belief, are 17.5 % ($n = 169$), with mean actual score 3.07 against self-assessment 8.78, an average gap of +5.72. *Underconfident* respondents mirror the pattern at 19.7 % ($n = 190$), with actual 6.25 against self-assessment 3.73. Among the four groups, the Overconfident stand out as the most policy-relevant – they are disproportionately likely to express PE intent (Section 4.5, with the KLH robustness check in Section 4.7.3). Figure 6.3 in Appendix C shows the sample composition and Table 4.5 reports the cell-level means.

Table 4.5: KLH four-group classification of respondents

Group	N	%	Mean actual	Mean self	Mean gap
Correctly Confident	359	37.3	7.36	8.40	1.04
Overconfident	169	17.5	3.07	8.78	5.72
Underconfident	190	19.7	6.25	3.73	−2.53
Correctly Unconfident	245	25.4	2.64	2.82	0.17

Notes: Four-group classification following Kim et al. (2020), $N = 963$. Cells are defined by independent median splits on actual total score (median = 5) and self-assessment Q23 (median = 6). The Overconfident cell defines oc_klh = 1 in the robustness check of Section 4.7.3.

4.3.3 Overconfidence by gender, age, education, income and investment experience

Age, education, income, and investment experience each generate significant between-group variation in the gap, one-way ANOVA $p < 0.001$ in every case. The gender result sits in the main body given its centrality to the residual gender penalty in Section 4.4.2; the four-panel demographic figure (Figure 6.2), the KLH-by-gender breakdown (Figure 6.4), and the full demographic table (Table 6.2), together with overall summary statistics (Table 6.1), are in Appendix C.

Gender

Male respondents have a mean overconfidence gap of +1.69 against +0.21 for women: Welch $t = 7.07$, $p < 0.001$, Cohen's $d = 0.45$, a medium effect (Table 4.6). Men outperform women on actual knowledge by 0.69 points (5.54 vs. 4.85) but inflate self-assessment by 2.17 points (7.23 vs. 5.06). The gender gap in overconfidence is driven by self-perception, not knowledge. The pattern matches the asymmetry that men are more overconfident suggested by Barber and Odean (2001), where women score lower on objective literacy but are also less likely to overestimate.

Table 4.6: Overconfidence components by gender

Measure	Male	Female	Diff (M−F)	Test	p-value
N	470	491	–	–	–
Mean actual score	5.54	4.85	0.69	$t = 4.14$	< 0.0001
Mean self-assessment	7.23	5.06	2.17	$t = 10.95$	< 0.0001
Mean overconfidence gap	1.69	0.21	1.48	$t = 7.07$	< 0.0001
Mean PE "Don't Know" count	1.65	2.80	−1.15	–	–
% Overconfident (KLH)	21.7	13.4	8.3 pp	–	–

Notes: Gender comparison of the components of the overconfidence gap and the share classified as Overconfident under the Kim et al. (2020) four-group scheme. Sample restricted to respondents reporting Male or Female ($N = 961$); "Other" or "Prefer not to say" (and respondents with missing Q23) are excluded. Welch t-tests are reported for the continuous comparisons. "% Overconfident (KLH)" is the within-gender share falling in the low-actual/high-self cell of the four-group classification.

Women select "Don't know" on PE questions much more often than men: 2.80 against 1.65 per respondent, consistent with Bucher-Koenen et al. (2017). Under the KLH classification, 48.1 % of men are Correctly Confident against 27.1 % of women, while 34.0 % of women are Correctly Unconfident against 16.6 % of men. Women are systematically more willing to acknowledge knowledge gaps, accounting for a meaningful share of the measured gender difference in both overconfidence and literacy. The epistemic-humility interpretation reappears in the residual gender penalty in Section 4.4.2 and the discussion in Section 5.4.

Age, education, income, and investment experience

Age follows an inverted U: those aged 18–24 are mildly underconfident at -0.19, the gap peaks at 35–44 at $+1.82$, and falls back to $+0.39$ at 55–64. Bachelor's and vocational respondents sit around $+0.6$ while the group with master's or above level of education stands out at $+1.94$. The pattern is striking, since if any subgroup might recognize the limits of its own PE knowledge it would be the most generally educated. Section 4.4.1 shows that education does not independently predict intent once literacy and overconfidence are controlled, suggesting that its association with PE intent is largely captured by these channels. Higher-income respondents show larger gaps. Investment experience produces the sharpest sorting: active investors $+1.66$, former investors $+1.24$, passive investors essentially calibrated at $+0.02$, never-invested only mildly overconfident at $+0.36$.

4.4 Hypothesis 2: PE literacy raises PE-investment intent

Specification 1 estimates Threshold A intent on PE literacy, overconfidence, framing, and demographic controls (Equation 3.4) and identifies the H2 and H3 channels jointly.

4.4.1 *Probit Specification 1*

Maximum-likelihood estimation on $N = 961$ complete-case observations after listwise deletion of missing controls returns a Pseudo R^2 of 0.24 at Threshold A, a strong fit for a binary behavioral outcome of this kind (McFadden, 1974). Table 4.7 reports regression results of Threshold A and Threshold B side by side. Following Equation 3.4, big5_score enters alongside pe_score, identifying the overconfidence coefficient as pure self-assessment error net of both literacy levels.

Table 4.7: Probit Specification 1 at Thresholds A and B

	Threshold A (positive intent)		Threshold B (definitely yes)	
Variable	Coef. (SE)	AME	Coef. (SE)	AME
PE literacy score	0.217^{***} (0.031)	0.060	0.177^{***} (0.039)	0.032
Big Five score	−0.005 (0.039)	−0.001	−0.037 (0.046)	−0.007
Overconfidence gap	0.132^{***} (0.017)	0.037	0.171^{***} (0.021)	0.031
Positive framing	0.252^{***} (0.094)	0.070	0.130 (0.113)	0.024
Age	-0.020^{***} (0.004)	−0.006	-0.010^{*} (0.005)	−0.002
Female	-0.399^{***} (0.101)	−0.111	-0.386^{***} (0.126)	−0.070
Income (ordinal)	0.052^{*} (0.029)	0.014	0.018 (0.036)	0.003
Education (ordinal)	0.138^{**} (0.054)	0.038	0.090 (0.067)	0.016
Has investment account	0.453^{***} (0.127)	0.126	0.765^{***} (0.181)	0.140
Constant	-0.869^{***} (0.208)	–	-1.925^{***} (0.288)	–
N	961		961	
Log-likelihood	−472.82		−314.92	
Pseudo R^2	0.238		0.264	
Baseline rate	35.1%		16.6%	

Notes: Probit estimates of Specification 1 (Equation 3.4) at Threshold A (positive intent on Q22, primary) and Threshold B (commitment-only, robustness). Coefficients with standard errors in parentheses; AME is the average marginal effect computed over the sample distribution. $N = 961$ respondents with complete data on all controls. Stars denote significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 4.6 plots the average marginal effects of Specification 1 at Threshold A with 95% confidence intervals. PE literacy raises positive intent strongly: $\beta = +0.217$, SE = 0.031, $z = 6.92$, $p < 0.001$, with an average marginal effect of +6.03 percentage points per correctly answered PE question. Over the full range from 0 to 7, this implies a predicted-probability difference of approximately 42.2 percentage points between a respondent who answers all PE questions correctly and one who answers none, with other variables held at sample means. **H2 is supported.**

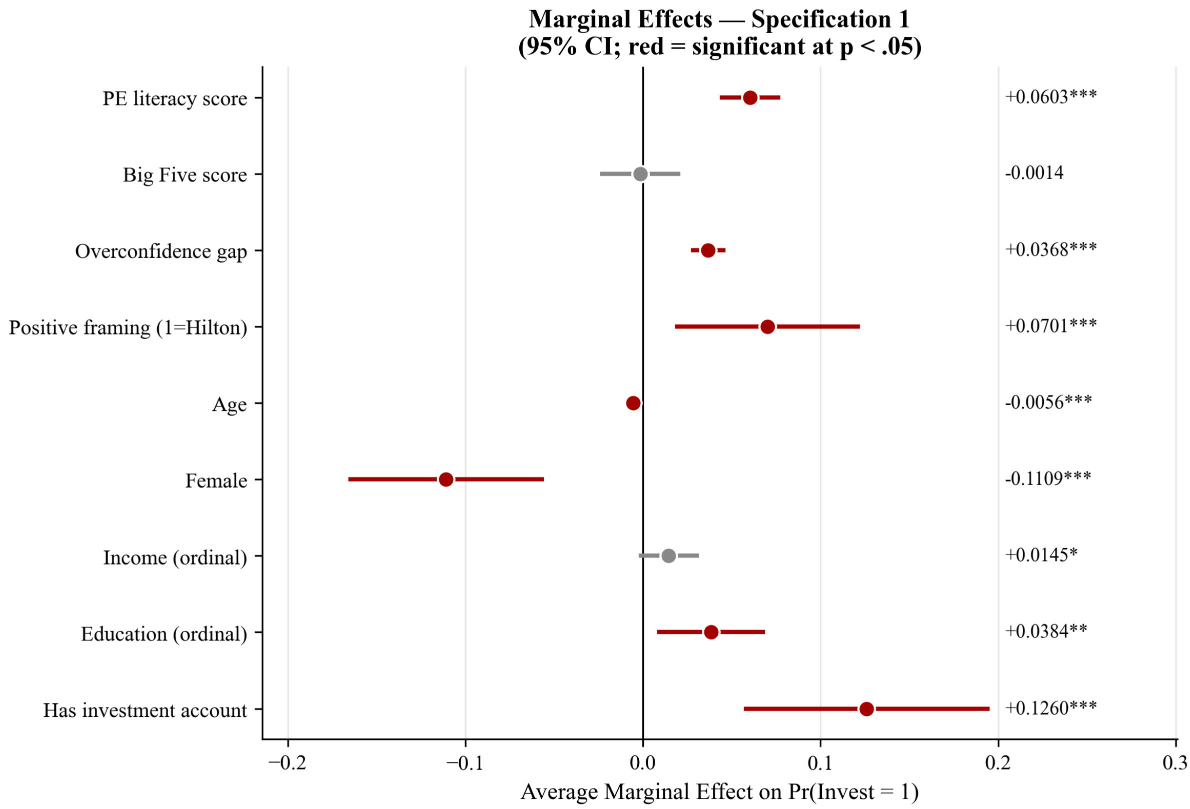

Figure 4.6: Average marginal effects of Specification 1 regressors at Threshold A, with 95% confidence intervals.

The big5_score coefficient is essentially null at the positive-intent threshold: $\beta = -0.005$, $\mathrm{SE} = 0.039$, $z = -0.13$, $p = 0.90$. General financial literacy does not independently predict willingness to add PE to retirement once PE knowledge and overconfidence are held fixed. General literacy gates PE-specific knowledge (Section 4.2), and once PE-specific knowledge is held constant, residual general-literacy variation carries no additional signal for PE intent. At Threshold B the coefficient remains null at $\beta = -0.037$, $p = 0.42$.

Anderson et al. (2017) find that perceived literacy dominates objective literacy in driving general financial behavior. In PE, both matter, and they matter separately. Each is significant at $p < 0.001$ in the joint specification, and their correlation ($\rho = -0.29$) is the expected sign – overconfident respondents score lower on the objective scale – and well within bounds that pose no threat to identification.

4.4.2 Other covariates

Holding an investment account is the largest single predictor in the model, raising the probability of positive PE intent by 12.60 percentage points, approximately double the AME of a one-unit increase in PE literacy. Whatever the channel, prior portfolio inclusion is the strongest predictor of future portfolio inclusion.

Age reduces intent at AME $= -0.56$ pp per year, $p < 0.001$, consistent with rising risk aversion across the lifecycle. Income (AME $= +1.45$ pp per band, $p = 0.072$) and education (AME $= +3.84$ pp per band, $p = 0.011$) are positive and at or near conventional significance. The framing main effect is positive: $\beta = +0.252$, AME $= +7.01$ pp, $p = 0.008$, and is discussed in Section 4.6.

The female coefficient is large and negative: women are 11.09 percentage points less likely to express positive PE intent than men, after controlling for knowledge, overconfidence, age, income, education, and investment experience. The penalty is larger at the positive-intent threshold than at the conservative threshold (Section 4.7.1). Combined with the much higher PE "Don't know" rate among women (2.80 vs. 1.65 items per respondent) and the much smaller female overconfidence gap ($+1.69$ vs. $+0.21$), we read this not as a missed knowledge deficit but as epistemic humility: women are more willing to acknowledge unfamiliarity by selecting weaker intent categories.

4.5 Hypothesis 3: Overconfidence raises PE-investment intent

4.5.1 Overconfidence as an independent intent driver

The overconfidence-gap coefficient is positive and highly significant: $\beta = +0.132$, SE $= 0.017$, $z = 7.70$, $p < 0.001$, with an AME of $+3.68$ percentage points per unit of gap. **H3 is supported.**

The result reproduces the Barber and Odean (2001) mechanism in a PE context: respondents who believe they know more than they do are substantially more willing to commit to an unfamiliar asset class. We follow Anderson et al. (2017) in entering objective and perceived literacy as separate channels. Because oc_gap is constructed as Q23 minus the combined Big Five and PE objective score, including big5_score and pe_score in the regression is what isolates the OC coefficient as pure miscalibration rather than partialling out additional controls.

4.5.2 Effect size across the gap distribution

To put the overconfidence effect in context, consider two respondents who differ only in how much they overestimate their knowledge. A respondent at the 75th percentile of the gap distribution (overestimating by 3 points) is 14.7 percentage points more likely to express PE investment intent than one at the 25th percentile (underestimating by 1 point). This interquartile effect is roughly comparable to the female effect (-11.1 pp) and the investment-experience effect (+12.6 pp), and about 2.5 times larger than the effect of answering one additional PE question correctly. A categorical robustness check replacing the continuous gap with the binary KLH Overconfident indicator yields an even larger effect of +15.9 pp.

4.6 Framing: registered (H4) and exploratory (H4b) interactions

If positive framing uniquely amplifies willingness to invest among respondents who already overestimate their PE knowledge, behavioral targeting would be a clear driver of retail PE uptake. Specification 2 (Equation 3.5) adds the framing $\times$ overconfidence interaction to Specification 1 and tests this registered hypothesis. Specification 2b (Equation 3.7) tests the exploratory framing $\times$ literacy alternative.

4.6.1 Probit Specification 2

The interaction coefficient is not statistically distinguishable from zero: $\beta(\text{framing} \times \text{oc\_gap}) = -0.0377$, SE $= 0.0285$, $z = -1.32$, $p = 0.185$ (Table 4.8). The point estimate is negative, opposite the sign predicted by H4, but the imprecision rules out substantive inference either way. Main effects of PE literacy, female, and those with investment experience are essentially unchanged from Specification 1.

Table 4.8: Probit Specifications 1 and 2 at Threshold A

Variable	Spec. 1	Spec. 2
Constant	-0.8691^{***} (0.2082)	-0.8935^{***} (0.2095)
PE literacy score	0.2169^{***} (0.0314)	0.2164^{***} (0.0314)
Big Five score	-0.0051 (0.0391)	-0.0060 (0.0392)
Overconfidence gap	0.1322^{***} (0.0172)	0.1519^{***} (0.0228)
Positive framing (1=Hilton)	0.2521^{***} (0.0942)	0.2943^{***} (0.0996)
Framing $\times$ Overconfidence	–	-0.0377 (0.0285)
Age	-0.0202^{***} (0.0040)	-0.0200^{***} (0.0040)
Female	-0.3988^{***} (0.1013)	-0.3965^{***} (0.1014)
Income (ordinal)	0.0521^{*} (0.0289)	0.0501^{*} (0.0290)
Education (ordinal)	0.1381^{**} (0.0542)	0.1384^{**} (0.0541)
Has investment account	0.4531^{***} (0.1269)	0.4591^{***} (0.1272)

Notes: Probit estimates at Threshold A (positive intent on Q22), $N = 961$. Specification 1 is Equation 3.4; Specification 2 (Equation 3.5) adds the framing $\times$ overconfidence interaction. Coefficients with standard errors in parentheses. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The likelihood-ratio test of Specification 2 against Specification 1 returns $\chi^2(1) = 1.76$, $p = 0.185$. The interaction adds essentially nothing to model fit. **H4 is not supported.**

4.6.2 Interpreting the null

Predicted probabilities of positive PE intent across a 2×2 grid of framing arm and overconfidence level, with other covariates at sample means, make the null concrete. Low overconfidence is the 25th percentile at gap $= -1$; high is the 75th at gap $= +3$. The overconfidence main effect dominates: high-OC respondents have predicted intent well above low-OC respondents in both framing arms. The framing main effect is visible too, since positive-framing respondents have higher intent in both OC strata, but it is roughly constant across OC levels and slightly smaller at high OC than at low OC. The two effects sit additively rather than multiplicatively. The framing main effect is real, but it operates on the population at large rather than uniquely amplifying the response of overconfident retail investors.

4.6.3 H4b (exploratory): framing × PE literacy

The H4 null prompted a diagnostic check on whether framing depends on PE literacy itself. The Hilton and Toys "R" Us vignettes only update priors for respondents with enough PE literacy to interpret the implied causal claim; below that comprehension threshold the vignette is just a story about a hotel chain or a toy retailer. The specification was not pre-registered. Categorizing PE literacy according to the buckets defined in Section 3.4.3 – Illiterate ($n = 617$, Q15 wrong and pe_score < 4), Moderate ($n = 176$, Q15 correct but pe_score < 4), Literate ($n = 239$, pe_score ≥ 4) – and estimating Specification 2b at Threshold A returns a jointly significant pair of interaction terms: $\beta(\text{framing} \times \text{Moderate}) = +0.710$, SE $= 0.278$, $z = 2.55$, $p = 0.011$, and $\beta(\text{framing} \times \text{Literate}) = +0.354$, SE $= 0.221$, $z = 1.60$, $p = 0.109$ (Table 6.4). The joint likelihood-ratio test against the no-interaction restricted model returns $\chi^2(2) = 7.70$, $p = 0.021$. Both interaction coefficients are in the predicted direction; the Moderate interaction is significant at conventional levels and the Literate interaction is positive but imprecise.

The full coefficient set for Specification 2b is reported in Table 6.4 of Appendix D.

The descriptive evidence on Q22 willingness, measured on the 1–6 Likert scale, lines up with the probit. The overall positive-vs-negative difference is $+0.15$ (3.35 vs. 3.20, Welch $t = 2.26$, $p = 0.024$). Within buckets, Illiterate respondents move by $+0.12$ ($p = 0.18$, n.s.), Moderate by $+0.34$ ($p = 0.032$), and Literate by $+0.23$ ($p = 0.080$). The strongest movement sits in the Moderate band.

The full descriptive breakdown is shown in Figure 6.5 and Table 6.3 in Appendix D.

The pattern is consistent with comprehension-gating. Illiterate respondents – who fail the foundational Q15 definition item – are essentially unmoved, as expected if the vignette cannot register without the floor of knowing what a PE fund is. Moderate respondents, who recognize the definition but lack the mechanics to discount it, show the strongest framing response. Literate respondents move directionally with the frame but the contrast is no longer significant: they have enough domain knowledge to evaluate a single anecdote critically. This inverts the conventional H4 prediction: the financially illiterate are the least moved by optimistic PE framing, not the most. We treat the result as an empirically supported pattern rather than a confirmed registered finding, since H4b was not pre-registered. The joint LR test at $p = 0.021$ and the Moderate within-bucket contrast at $p = 0.032$ are jointly significant at conventional levels; replication with a stronger or richer manipulation would be required to upgrade H4b to registered status.

4.7 Robustness

Five robustness exercises leave the qualitative conclusions of Sections 4.4–4.6 intact. The exercises test robustness to outcome-threshold choice, PE-item construction, overconfidence operationalization, sample quality, and calibration coding. Four of the five specifications are anchored at the conservative commitment threshold (Threshold B), since B is the most demanding outcome and the hardest place to retain significance; any specification that survives at B automatically survives at the positive-intent primary outcome, which the threshold-sensitivity exercise establishes directly.

4.7.1 Outcome-threshold sensitivity

The exercise re-estimates Specifications 1 and 2 at Threshold B. The two thresholds produce qualitatively identical results for the literacy and overconfidence channels; the framing main effect is significant at A but null at B. The PE-literacy coefficient is +0.217 at A and +0.177 at B; both are highly significant. The overconfidence coefficient is +0.132 at A and +0.171 at B; the larger coefficient at the conservative threshold supports the secondary interpretation: overconfidence drives commitment in particular. The big5_score control is null at both thresholds (-0.005, $p = 0.90$ at A; -0.037, $p = 0.42$ at B). The framing main effect appears at A (+0.252, $p = 0.008$) but is null at B (+0.130, $p = 0.25$): framing moves intent in the soft probable region but rarely all the way to firm commitment. The framing $\times$ overconfidence interaction is null at both thresholds. The coefficient comparison across thresholds is plotted in Figure 6.6 of Appendix E.

4.7.2 Drop the two interpretation-sensitive PE items (Q17, Q19)

Q17, the liquidity/lock-up question, and Q19, the question asking about the PE vs. public returns, are the two PE questions whose keyed answers most depend on benchmark or structural assumptions (Section 5.6). Q17's keyed answer ("Only after the fund ends and sells its investments") is correct for the whole-fund, rather than the deal-by-deal structure, as well as understates the partial liquidity available through secondaries markets and the quarterly-redemption windows in newer interval and evergreen structures. Q19's keyed 3 % premium tracks the influential Harris et al. (2014) strand but is benchmark- and period-sensitive (Phalippou, 2020). Removing them and reconstructing PE literacy as a 0–5 score leaves the headline H2 result intact: $\beta(\text{pe\_score, reduced}) = +0.251$, $p < 0.001$. The overconfidence coefficient is essentially unchanged at $\beta_{\text{oc}} = +0.168$, $p < 0.001$, and the big5_score coefficient remains insignificant at $\beta = -0.034$. The headline H2 result does not depend on the borderline items. The full coefficient set is in Table 6.5 of Appendix E.

4.7.3 KLH-binary overconfidence

Replacing the continuous overconfidence gap with the binary KLH "Overconfident" indicator ($n = 169$) yields $\beta = +0.819$, $p < 0.001$, AME $= +15.85$ percentage points. The PE-literacy coefficient retains its sign and significance at $\beta_{\text{pe}} = +0.136$, AME $= +2.63$ pp. The big5_score coefficient becomes mildly significant at $\beta = -0.094$, $p = 0.04$. We do not interpret this shift substantively. The coefficient is significant in only one of the five robustness specifications and may reflect overlap between `big5_score` and the binary KLH classifier, whose actual-score component is constructed from the combined `big5_score` and `pe_score` measures. An Overconfident-cell respondent is 15.85 percentage points more likely to commit to PE intent than a respondent in the pooled reference group: a useful categorical summary of H3 that complements the continuous-gap estimate. The full coefficient set is in Table 6.6 of Appendix E.

4.7.4 Drop Qualtrics-flagged respondents (gc = 4)

Qualtrics flagged 21 respondents (gc = 4) for quality concerns and retained them in our delivered sample. Re-estimating Specification 1 on the 1,011-respondent gc = 1 subset yields β(pe_score) = +0.183, β(big5_score) = −0.049 ($p = 0.30$), and β(oc_gap) = +0.171, with female and has_invested coefficients within one standard error of headline. The H1 paired t on the gc = 1 subset is $t = 28.42$, Cohen's $d = 0.894$ against 0.885 in the full sample. Only seven of the 21 flagged respondents would have entered the original probit; the other 14 were already excluded for missing demographic controls. The flagged cases drive no headline result. The full coefficient set is in Table 6.7 of Appendix E.

4.7.5 "Don't know" as correct answer

The exercise treating the "Don't know" as correct is best read as a stress test on H3 rather than on H2. The "Don't know" rate is itself a trait, since women, less-educated respondents, and the strictly knowledgeable all use DK at higher rates (Section 4.3.3). The relevant question is whether overconfidence still predicts intent when the willingness to admit ignorance is removed from the calibration measure.

We retain the strict pe_score and big5_score (DK coded as wrong) but rebuild overconfidence on a total, which treats DK as correct, oc_gap_dk = Q23 − total_score_dk. This reconstructs overconfidence while leaving the literacy scores unchanged. The OC coefficient on the DK-rebuilt gap is +0.183 ($p < 0.001$), slightly larger than the headline +0.171 at the conservative threshold. The PE-literacy coefficient attenuates substantially to +0.050 ($p = 0.21$) and is no longer significant. The Pseudo R^2 rises to 0.333 at Threshold B, against 0.264 in the headline at the same threshold. The overconfidence channel survives the recoding intact and, if anything, strengthens; this is the cleanest takeaway from the exercise. We do not interpret the literacy attenuation as a verdict on H2, since the DK-tolerant total is constructed from the same items that enter the literacy scores, so the two channels cannot be separated cleanly under this coding. The exercise speaks to H3's robustness rather than to a relative ranking between H2 and H3. Full coefficient table in Appendix E, Table 6.8.

5. Discussion

5.1 The threshold character of PE literacy

The results suggest that PE literacy is not a smooth extension of general financial literacy. Instead, it represents a qualitatively different form of financial understanding that emerges only at the upper end of the general-literacy distribution. Respondents who score between 0 and 4 on the Big Five have relatively little variation in PE performance and cluster around mean PE scores close to random-guessing. Respondents with a perfect Big Five score display a discontinuous increase in PE literacy. Yet, even among this financially literate subgroup, average PE performance reaches only 51.5 % of the questions. This indicates that conventional financial literacy remains an incomplete proxy for understanding institutionally complex assets. The threshold structure is consistent with the broader domain-specific literacy literature, which indicates that specialized financial products require forms of knowledge not captured by general literacy batteries (Jones et al., 2024; Lusardi & Tufano, 2015).

The structure of the PE responses further suggests that familiarity with PE is often superficial rather than conceptual. Respondents who correctly answer Q15 – PE-definition item – score nearly two additional points on the PE questions relative to those who fail it. Only 31.2 % of respondents correctly identify what a PE fund is, while substantially more respondents (47.1 %) correctly recognize the "2 and 20" fee structure. Therefore, more respondents know that PE charges fees than can define the asset class itself. The distinction is economically meaningful because it separates recognition of widely publicized attributes from understanding of the institutional structure underlying the product. The pattern aligns with Hastings et al. (2013), who argue that financially complex products with multidimensional attributes are particularly difficult for households to evaluate.

Another notable feature of the data is the consistently high use of "Don't know" responses throughout the PE module. Relative to the Big Five, respondents are substantially more willing to acknowledge uncertainty when confronted with PE questions. This suggests that the issue is not simply random guessing or misinformation but recognition of the limits of one's own financial competence. Interestingly, the response distribution on the PE-definition item resembles that of the bond-pricing question from the Big Five module, where respondents may recognize the topic itself while lacking a coherent conceptual framework for evaluating it.

5.2 *Overconfidence as an independent driver of investment intent*

The results show that overconfidence predicts PE-investment intent independently of actual literacy and does so with a degree of consistency that exceeds every other behavioral variable in the model. This suggests that willingness to participate in private equity depends not only on what respondents know but also on what they believe they know. This extends prior findings that perceived financial competence can predict economic behavior independently of objective literacy (Anderson et al., 2017) into a structurally unfamiliar asset-class setting. Just over half of the sample (50.4%) overestimates its own financial performance, and the magnitude of the relationship is economically meaningful: moving from the 25th percentile of the overconfidence-gap distribution (gap = -1) to the 75th percentile (gap = $+3$) raises predicted PE-investment intent by approximately 14.7 percentage points, an effect comparable in size to already holding an investment account. The KLH-binary specification sharpens the interpretation further because respondents in the low-knowledge/high-self-assessment cell are 15.85 percentage points more likely to express commitment than the pooled reference group despite an average of 3.07 correct answers out of 12 while rating themselves at 8.78. Therefore, the respondents most willing to engage with PE are not necessarily the most informed, but those whose perceived competence most exceeds their demonstrated understanding (Kruger & Dunning, 1999). The overconfidence coefficient is robust across specifications, including the recalibration in which Don't know responses are treated as correct only within the calibration measure.

The demographic structure of the overconfidence results is particularly notable because the highest gaps appear not among financially disengaged respondents, but among financially active and highly educated groups. Respondents aged 35–44 display the largest average gap ($+1.82$), master's-degree holders exhibit substantially higher overconfidence than less-educated groups ($+1.94$), and active investors show the highest overconfidence among all investment-experience categories ($+1.66$), while passive investors are essentially calibrated ($+0.02$). In the PE setting, financial engagement and formal education appear to increase confidence substantially more than actual understanding. Notably, this demographic profile resembles the group vulnerable to complex investment fraud (AARP, 2011; Consumer Fraud Research Group, 2006; Graham, 2014; Kieffer & Mottola, 2017).

5.3 *Framing: a real main effect, with no interaction*

Our framing results point to a clear main effect but little evidence of behavioral targeting through overconfidence. Respondents exposed to the positive "Hilton" framing are 7.0 percentage points more likely to express positive PE intent than respondents exposed to the negative "Toys 'R' Us" framing, when PE literacy, overconfidence, and demographics are held constant. The framing $\times$ overconfidence interaction, however, is null and directionally negative ($\beta = -0.038$, $p = 0.185$), while the likelihood-ratio test indicates that the interaction adds essentially no explanatory power to the model. The data therefore do not support the hypothesis that optimistic PE narratives amplify willingness among overconfident respondents.

The threshold analysis further clarifies the mechanism – framing increases "positive intent" responses but loses significance at the conservative "definitely yes" threshold. This indicates that narrative framing primarily shifts respondents out of hesitation rather than into firm commitment.

The exploratory framing $\times$ PE-literacy interaction produces the more theoretically interesting result. The frame effect is the smallest among the PE-illiterate respondents ($+0.12$, n.s.), the strongest among the moderate-literacy group ($+0.34$, $p = 0.032$), and positive but less precise among the PE-literate respondents ($+0.23$, $p = 0.080$). The joint interaction test is statistically significant, with the moderate-literacy interaction driving the result ($\beta = +0.710$, $p = 0.011$). Framing effects remain positive among PE-literate respondents but are strongest among the moderate-literacy group. The pattern suggests that respondents with partial familiarity with PE may be the most responsive to narrative framing: knowledgeable enough for the vignette to register, but not sufficiently informed for the framing effect to weaken substantially.

5.4 The residual gender gap in investment intent

Women remain less likely than men to express positive PE-investment intent after controlling for literacy, calibration, and demographic characteristics. This confirms that women are less willing than men to engage with unfamiliar financial products. Taken at face value, the residual may appear consistent with a preference-based interpretation, in which gender-linked differences in risk tolerance, financial confidence, or attitudes toward complexity explain variation in PE investment intent beyond observable characteristics. The financial-literacy literature, however, has attributed observed gender gaps in literacy scores to lower confidence rather than lower knowledge (Bucher-Koenen et al., 2021), and the calibration evidence in our sample extends the same logic to PE-investment intent.

Men exhibit a substantially larger overconfidence gap than women (+1.69 versus +0.21), and the difference is driven primarily by self-assessment rather than objective knowledge. Men outperform women on the combined literacy scale by only 0.69 points (5.54 versus 4.85), yet rate their own performance 2.17 points higher (7.23 versus 5.06). Women also select "Don't know" more frequently on the PE module, averaging 2.80 responses per respondent compared with 1.65 for men. The KLH classification reinforces the same pattern: 34.0 % of women fall into the Correctly Unconfident category against 16.6 % of men, while 48.1 % of men are classified as Correctly Confident compared with 27.1 % of women. Therefore, two established asymmetries from the financial-literacy literature appear simultaneously in the sample: the male overconfidence pattern documented by Barber and Odean (2001) and the higher female use of "Don't know" responses documented by Bucher-Koenen et al. (2017). In the PE setting, the two patterns describe the same underlying phenomenon from opposite directions: unfamiliarity is translated by men into stated willingness to invest and by women into acknowledged uncertainty.

The PE setting aligns with Barber and Odean (2001) and Bucher-Koenen et al. (2017) rather than Lawrence et al. (2024), plausibly because women's high DK rate on PE-battery depresses their self-assessment more than their objective performance, amplifying the perceived-knowledge gap underlying the gender pattern.

The policy implication follows directly from the calibration evidence rather than from a preference-based interpretation. If the residual female coefficient partly reflects differences in calibration rather than solely gender-specific preferences against PE, then reducing the participation gap depends less on accommodating heterogeneous preferences and more on improving comprehension and calibration. Expanding retail access to PE without improving investor understanding may therefore widen participation asymmetries, since overconfident male respondents exhibit greater willingness to invest despite being poorly equipped to evaluate the structure and risks embedded in the asset class.

5.5 Policy implications

The policy context is set by the August 2025 U.S. Executive Order (Trump, 2025) instructing the Department of Labor to facilitate the inclusion of alternative assets, including private equity, in defined-contribution retirement plans. Retail PE access is therefore imminent rather than hypothetical, and several implications follow.

Disclosure regimes face a binding comprehension floor: Because respondents are far more likely to recognize superficial features (like fee structures) than to actually define the asset class itself, current disclosure frameworks likely assume a baseline conceptual understanding that many retail investors simply do not possess.

General financial literacy is insufficient: General financial literacy appears necessary but insufficient for PE-specific comprehension. Even among the most financially literate respondents, mastery of PE concepts remains remarkably low. This demonstrates that standard financial education does not smoothly transfer to complex alternative assets.

Access may target the most miscalibrated investors: Expanded retail access may disproportionately draw from respondents whose perceived competence exceeds their demonstrated understanding, rather than those with the strongest PE-specific knowledge. Because overconfidence is most pronounced among highly educated and active retail investors, these groups are most at risk of participation driven by miscalibration rather than genuine comprehension.

Narrative framing exploits partial familiarity: Positive narrative framing meaningfully increases PE-investment intent, but the effect is concentrated among respondents who have only partial conceptual familiarity. This suggests that marketing narratives are most effective, and potentially most hazardous, for investors who know just enough to recognize the product but not enough to critically evaluate it.

Unaddressed comprehension gaps will widen gender asymmetries: The data suggest that lower female participation reflects calibration under uncertainty more than lower financial capability itself. Women exhibit a greater willingness to acknowledge incomplete understanding in this highly complex domain. Therefore, expanding retail access without first improving investor comprehension may inadvertently widen gender-based participation asymmetries.

Our findings invert the sophistication argument: the investors most likely to gain early access to retail PE, credentialed, financially active, and confident, are the most miscalibrated, while meaningful PE understanding remains concentrated in a small minority even within the most financially literate.

5.6 Limitations

Several design choices and scope conditions bound the claims that follow, and we set them out here so the reader can weigh them against the results.

First, the overconfidence measure is imperfectly targeted. Q23 asks for combined performance across all twelve items rather than the PE items alone, so the overconfidence gap blends general-literacy and PE-specific self-assessment without clean decomposition. The Dunning-Kruger prediction motivating the variable is PE-specific cognitive failure driving willingness to accept PE exposure, which the instrument cannot cleanly isolate. PE items account for seven of twelve scored questions, so PE-specific cognition contributes the larger share. The robustness check treating DK answer as correct (Section 4.7.5) provides partial reassurance: the overconfidence coefficient survives manipulation with magnitudes within 0.02 of headline, suggesting the combined-measure variable is not an artefact of differential ignorance-acknowledgment. A re-fielded survey with separate PE-only and Big Five-only self-assessments could decompose the measure cleanly. Apart from that, our overconfidence variable is an overestimation measure in the Moore and Healy (2008) taxonomy – the gap between self-assessed and actual score – and does not capture overplacement or overprecision, which the same authors show are conceptually and empirically distinct. The interpretation throughout refers to overestimation, not the stronger claim that overconfident respondents consider themselves more knowledgeable than their peers.

Second, the seven PE items differ in how much their keyed answers depend on benchmark or structural assumptions. The Q17/Q19 robustness check (Section 4.7.2) re-estimates the probit on a reduced 0–5 PE scale and leaves the headline H2 result intact, which alleviates concerns surrounding the two items whose keyed answers are most defensibly contestable.

Third, the outcome variable is stated investment intent rather than realized behavior, a standard survey limitation shared with the broader literature on retail investor decision-making (Anderson et al., 2017; Lusardi & Mitchell, 2014). Stated intent is not a substitute for direct behavioral data, and the gap between intent and realized choice can be large. In retirement saving specifically, Madrian and Shea (2001) show that defaults dominate stated preferences under auto-enrollment, regardless of what employees previously said they intended. How well stated PE intent tracks realized PE-allocation behavior under any specific choice architecture is, as far as we are aware, undocumented.

Fourth, we document association, not causation, outside the framing manipulation. The structural link between literacy, overconfidence, and intent admits multiple causal orderings, and our cross-sectional design cannot distinguish them. Fernandes et al.'s (2014) meta-analytic caution on literacy effects applies with particular force in an unfamiliar asset class: the association we report between PE literacy and intent in Section 4.4 may reflect both causal and selection effects.

Fifth, our hypothesis set is not the only one the data could support. The instrument collected information on employment, finance education, and risk tolerance that we do not fully exploit. The relationship between prior formal finance education and PE literacy, the role of employment status in moderating intent, and the interaction of overconfidence with explicit risk-preference elicitation are all productive extensions, which we leave to subsequent work.

Sixth, the framing manipulation uses two opposing vignettes, Hilton (positive) and Toys "R" Us (negative), without a neutral no-framing control. The estimated framing coefficient therefore captures the difference between the two arms rather than the absolute effect of either vignette relative to no exposure. Throughout the paper, "positive framing raises intent by 7 pp" should be read as a between-arm contrast: we cannot decompose how much of the effect reflects the positive vignette lifting intent versus the negative vignette suppressing it.

Finally, our results are specific to the U.S. retail population in November 2025 and should not be extrapolated to European or Asian retail markets without separate fielding. The regulatory architecture of retail PE access differs materially across jurisdictions, ELTIF 2.0 in the EU, the Investment Trusts Act in Japan, and the 401(k)/IRA framework in the U.S., and corporate-governance and cultural differences could further shift the empirical patterns. A European extension is a natural next step given the parallel ELTIF 2.0 retail opening, but the conditions under which our U.S. results would translate are uncertain.

Despite these limitations, the four-channel structure we document – PE literacy, overconfidence, framing, and prior investment-account ownership – is robust across thresholds, item-set redefinitions, and DK recodings, and the policy implications of Section 5.5 follow conditional on the scope conditions above.

Conclusion

Private equity is becoming more accessible to retail investors on the premise that sophistication, rather than wealth, should determine access. This thesis asks what that sophistication actually looks like in a general retail population. The evidence suggests that it is substantially narrower, more domain-specific, and more unevenly distributed than current retail-access debates implicitly assume.

Across 1,032 U.S. respondents, we document a large and highly structured gap between general financial literacy and private-equity literacy. PE literacy remains relatively flat across most of the Big Five distribution before rising sharply among respondents with perfect general-literacy performance. Even within this highly literate subgroup, respondents answer only around half of the PE questions correctly, indicating that general financial competence does not readily translate into understanding of private markets. The results therefore support a threshold interpretation of PE literacy: general financial knowledge appears to function less as a smooth predictor of PE understanding than as a precondition for acquiring it at all.

Knowledge alone, however, does not determine willingness to participate. Overconfidence predicts PE-investment intent independently of actual literacy and remains one of the most robust predictors across specifications and robustness checks. The respondents most willing to allocate to PE are not necessarily the best informed, but those whose perceived competence exceeds their demonstrated understanding. The results therefore extend the literature on perception versus knowledge into a structurally unfamiliar asset class, suggesting that confidence can substitute for comprehension even where objective knowledge remains weak.

The framing results qualify this picture in an important way. Positive narrative framing increases PE-investment intent, but the effect is not concentrated among the most overconfident respondents. Instead, framing appears strongest among investors who recognize the basic concept of PE but lack the structural understanding required to evaluate narratives critically. Respondents with very low literacy are comparatively unresponsive to framing, while more literate respondents appear less sensitive to it. The evidence therefore supports a comprehension-gating interpretation rather than a pure manipulation story.

Taken together, the findings suggest that current retail-access debates rest on an overly generalized view of financial sophistication. Existing frameworks implicitly assume that the knowledge required to evaluate publicly traded securities transfers relatively smoothly to complex private-market products. Our evidence instead suggests that PE literacy is threshold-based, domain-specific, and only partially captured by standard measures of financial literacy. The central policy question is therefore not simply whether retail investors should have access to private equity, but whether current concepts of investor sophistication adequately capture the type of understanding that participation in private markets requires.

References

AARP. (2011). *Investment fraud vulnerability study* (tech. rep.). AARP Foundation.

Agarwal, S., Driscoll, J. C., Gabaix, X., and Laibson, D. (2009). The age of reason: Financial decisions over the life cycle and implications for regulation. *Brookings Papers on Economic Activity*, *2009*(2), 51–117.

Agnew, J. R., and Szykman, L. R. (2005). Asset allocation and information overload: The influence of information display, asset choice, and investor experience. *Journal of Behavioral Finance*, *6*(2), 57–70.

Allgood, S., and Walstad, W. B. (2016). The effects of perceived and actual financial literacy on financial behaviors. *Economic Inquiry*, *54*(1), 675–697.

Anderson, A., Baker, F., and Robinson, D. T. (2017). Precautionary savings, retirement planning, and misperceptions of financial literacy. *Journal of Financial Economics*, *126*(2), 383–398.

Appelbaum, E., and Batt, R. (2014). *Private equity at work: When Wall Street manages Main Street*. Russell Sage Foundation.

Aramonte, S., and Avalos, F. (2021). The rise of private markets. *BIS Quarterly Review*, 69–82.

Atkinson, A., and Messy, F.-A. (2012). *Measuring financial literacy: Results of the OECD/INFE pilot study* (OECD Working Papers on Finance, Insurance and Private Pensions No. 15). Organisation for Economic Co-operation and Development.

Bannier, C. E., and Neubert, M. (2016). Gender differences in financial risk taking: The role of financial literacy and risk tolerance. *Economics Letters*, *145*, 130–135.

Barber, B. M., and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics*, *116*(1), 261–292.

Benartzi, S., and Thaler, R. H. (2007). Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives*, *21*(3), 81–104.

Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2011). How does simplified disclosure affect individuals' mutual fund choices? *Explorations in the Economics of Aging*, 75–96.

Boyle, P. A., Yu, L., Wilson, R. S., and Bennett, D. A. (2025). Cognitive aging and the decline in financial literacy. *Journals of Gerontology: Series B*.

Bucher-Koenen, T., Alessie, R., Lusardi, A., and van Rooij, M. (2021). Fearless woman: Financial literacy and stock market participation. *National Bureau of Economic Research Working Paper*, (28723).

Bucher-Koenen, T., Lusardi, A., Alessie, R., and van Rooij, M. (2017). How financially literate are women? An overview and new insights. *Journal of Consumer Affairs*, *51*(2), 255–283.

Choi, J. J., Laibson, D., and Madrian, B. C. (2010). Why does the law of one price fail? An experiment on index mutual funds. *Review of Financial Studies*, *23*(4), 1405–1432.

Consumer Fraud Research Group. (2006). *Investor fraud study final report* (tech. rep.). NASD Investor Education Foundation.

European Parliament and Council of the European Union. (2023). Regulation (eu) 2023/606 amending regulation (eu) 2015/760 (ELTIF 2.0).

Fernandes, D., Lynch Jr., J. G., and Netemeyer, R. G. (2014). Financial literacy, financial education, and downstream financial behaviors. *Management Science*, *60*(8), 1861–1883.

Finke, M. S., Howe, J. S., and Huston, S. J. (2017). Old age and the decline in financial literacy. *Management Science*, *63*(1), 213–230.

FINRA Foundation. (2025). *The NFCS 2024 state-by-state survey: Highlights* (tech. rep.). FINRA Investor Education Foundation.

FINRA Investor Education Foundation. (2009). *Financial capability in the United States: National survey – executive summary* (tech. rep.). FINRA Investor Education Foundation.

Gamble, K. J., Boyle, P. A., Yu, L., and Bennett, D. A. (2014). Aging and financial decision making. *Management Science*, *61*(11), 2603–2610.

Graham, E. J. (2014). Investor vulnerability and investment fraud victimization. *Working paper*.

Hacker, J. S. (2011). The institutional foundations of Middle-Class insecurity. *Policy Network Essay*.

Harris, R. S., Jenkinson, T., and Kaplan, S. N. (2014). Private equity performance: What do we know? *Journal of Finance*, *69*(5), 1851–1882.

Hastings, J. S., Madrian, B. C., and Skimmyhorn, W. L. (2013). Financial literacy, financial education, and economic outcomes. In *Annual review of economics* (pp. 347–373, Vol. 5).

Houts, C. R., and Knoll, M. A. Z. (2020). The financial knowledge scale: New analyses, findings, and development of a short form. *Journal of Consumer Affairs*, *54*(2), 775–800.

Hung, A. A., Parker, A. M., and Yoong, J. (2009). *Defining and measuring financial literacy* (RAND Working Paper No. WR-708). RAND Corporation.

Huston, S. J. (2010). Measuring financial literacy. *Journal of Consumer Affairs*, *44*(2), 296–316.

Jones, M., Luu, T., and Samuel, A. (2024). Crypto literacy and cryptocurrency investment behavior. *Working paper*.

Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Kaiser, T., Lusardi, A., Menkhoff, L., and Urban, C. (2022). Financial education affects financial knowledge and downstream behaviors. *Journal of Financial Economics*, *145*(2A), 255–272.

Kaplan, S. N., and Schoar, A. (2005). Private equity performance: Returns, persistence, and capital flows. *Journal of Finance*, *60*(4), 1791–1823.

Kaplan, S. N., and Strömberg, P. (2009). Leveraged buyouts and private equity. *Journal of Economic Perspectives*, *23*(1), 121–146.

Kieffer, C. N., and Mottola, G. R. (2017). *Understanding and combating investment fraud* (tech. rep.). FINRA Investor Education Foundation.

Kim, K. T., Lee, J., and Hanna, S. D. (2020). The effects of financial literacy overconfidence on the mortgage delinquency of US households. *Journal of Consumer Affairs*, *54*(2), 517–540.

Kim, K. T., Lee, J., and Kim, H. C. (2022). Gender differences in financial knowledge overconfidence among older adults. *International Journal of Consumer Studies*.

Klapper, L., Lusardi, A., and van Oudheusden, P. (2015). *Financial literacy around the world: Insights from the S&P global FinLit survey* (tech. rep.). Standard & Poor's Ratings Services Global Financial Literacy Survey.

Knoll, M. A. Z., and Houts, C. R. (2012). The financial knowledge scale: An application of item response theory to the assessment of financial literacy. *Journal of Consumer Affairs*, *46*(3), 381–410.

Kruger, J., and Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*(6), 1121–1134.

Lawrence, E. R., Nguyen, T. D., and Wick, B. (2024). Gender differences in financial literacy overconfidence. *Working paper*.

Lusardi, A., and Mitchell, O. S. (2006). *Financial literacy and planning: Implications for retirement wellbeing* (Working Paper). Pension Research Council, Wharton School, University of Pennsylvania.

Lusardi, A., and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, *52*(1), 5–44.

Lusardi, A., and Tufano, P. (2015). Debt literacy, financial experiences, and overindebtedness. *Journal of Pension Economics and Finance*, *14*(4), 332–368.

Madrian, B. C., and Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, *116*(4), 1149–1187.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics* (pp. 105–142). Academic Press.

Metrick, A., and Yasuda, A. (2010). The economics of private equity funds. *Review of Financial Studies*, *23*(6), 2303–2341.

Moore, D. A., and Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, *115*(2), 502–517.

OECD. (2023). *OECD/INFE 2023 international survey of adult financial literacy* (tech. rep.). Organisation for Economic Co-operation and Development. Paris.

Parker, A. M., Bruine de Bruin, W., Yoong, J., and Willis, R. (2012). Inappropriate confidence and retirement planning: Four studies with a national sample. *Journal of Behavioral Decision Making*, *25*(4), 382–389.

Phalippou, L. (2020). An inconvenient fact: Private equity returns and the billionaire factory. *Journal of Investing*, *30*(1), 11–39.

Poterba, J., Rauh, J., Venti, S., and Wise, D. (2007). Defined contribution plans, defined benefit plans, and the accumulation of retirement wealth. *Journal of Public Economics*, *91*(10), 2062–2086.

Qualtrics. (2025). Qualtrics XM platform (2024 ed.). https://www.qualtrics.com

Remund, D. L. (2010). Financial literacy explicated: The case for a clearer definition in an increasingly complex economy. *Journal of Consumer Affairs*, *44*(2), 276–295.

Robinson, D. T., and Sensoy, B. A. (2016). Cyclicality, performance measurement, and cash flow liquidity in private equity. *Journal of Financial Economics*, *122*(3), 521–543.

SEC Investor Advisory Committee. (2025). *Recommendations regarding retail investor access to private markets* (tech. rep.). U.S. Securities and Exchange Commission.

Shiller, R. J. (2017). Narrative economics. *American Economic Review*, *107*(4), 967–1004.

Thaler, R. H., and Benartzi, S. (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy*, *112*(S1), S164–S187.

Trump, D. J. (2025, August). Democratizing access to alternative assets for 401(k) investors, exec. order no. 14,330.

Tversky, A., and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*(2), 207–232.

Tversky, A., and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458.

U.S. Securities and Exchange Commission. (2020). Accredited investor definition: Final rule [Effective December 8, 2020].

van Rooij, M., Lusardi, A., and Alessie, R. (2011). Financial literacy and stock market participation. *Journal of Financial Economics*, *101*(2), 449–472.

Xia, T., Wang, Z., and Li, K. (2014). Financial literacy overconfidence and stock market participation. *Social Indicators Research*, *119*(3), 1233–1245.

Xiao, J. J., Li, Y., and Zhou, C. (2022). Financial literacy overconfidence and investment fraud susceptibility. *Working paper*.

6. Appendices

A Survey instrument

The survey below was administered to all participants via Qualtrics. Question wording, response options, and ordering are reproduced exactly as shown to respondents. Each respondent saw only one of the two framing vignettes in Section A.3 (Q13 or Q14), randomly assigned.

A.1 Demographics and Socio-Economic Background

Q1. What is your age?

- ○ 18–24
- ○ 25–34
- ○ 35–44
- ○ 45–54
- ○ 55–64
- ○ 65+

Q2. What is your gender?

- ○ Male
- ○ Female
- ○ Other
- ○ Prefer not to say

Q3. What is the highest level of education you have completed?

- ○ High School/GED
- ○ Vocational/Technical School
- ○ Bachelor's degree
- ○ Master's degree or above

Q4. What is your employment status?

- ○ Employed full-time
- ○ Employed part-time
- ○ Self-employed/Entrepreneur

- Unemployed
- Student/Retired

Q5. What is your annual gross income (before tax)?

- Less than $34,999
- $35,000 to $49,999
- $50,000 to $74,999
- $75,000 to $99,999
- $100,000 to $149,999
- $150,000 to $199,999
- $200,000 to $499,999
- $500,000 or more

Q6. Do you currently have an investment account?

- Yes – I actively invest in stocks, funds, or other instruments
- Yes – but I rarely make changes to my investments
- No – but I have invested in the past
- No – I have never invested

Q7. Have you ever studied finance or economics?

- Yes, in high school
- Yes, in college/university
- Yes, through self-study
- No

A.2 Financial Literacy (Big Five)

Q8. Suppose you had $100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money there?

- **More than $102**
- Exactly $102
- Less than $102
- Do not know

Q9. Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, how much would you be able to buy with the money in this account?

- More than today
- Exactly the same
- **Less than today**

- ○ Do not know

Q10. True or false: buying a single company's stock usually provides a safer return than a stock mutual fund.

- ○ True
- ○ **False**
- ○ Do not know

Q11. True or false: a 15-year mortgage typically requires higher monthly payments than a 30-year mortgage, but the total interest paid over the life of the loan will be less.

- ○ **True**
- ○ False
- ○ Do not know

Q12. If interest rates rise, what will typically happen to bond prices?

- ○ They will rise
- ○ **They will fall**
- ○ They will stay the same
- ○ There is no relationship
- ○ Do not know

A.3 Private Equity Perception (Randomized Framing Vignettes)

Each respondent was randomly assigned to one of the two vignettes below (Q13, positive frame, or Q14, negative frame) and saw only that vignette.

Q13. Did you know that Hilton was owned by a private-equity firm?

The next few questions ask you about private equity. Please read this text carefully. One of the most famous examples of how private equity works is Hilton Hotels. The company's value more than doubled while a private-equity firm owned it. During that period, private equity investors helped Hilton expand into new markets, modernize its booking systems, and strengthen its brand after the financial crisis. The company emerged stronger than ever and has since been recognized by Great Place to Work and Fortune magazine as one of the world's best workplaces, an achievement many attribute to that successful partnership.

- ○ Yes, I already knew Hilton was owned by a private-equity firm
- ○ No, this is new information to me

Q14. Did you know that Toys "R" Us was owned by a private-equity firm?

The next few questions ask you about private equity. Please read this text carefully. One of the most famous examples of how private equity works is Toys "R" Us. Once every child's favorite store, it collapsed after years under private-equity ownership. The company was loaded with *debt, couldn't keep up with online competition, and eventually shut down, costing more than 30,000 jobs. Many former employees said they felt the owners cared more about debt repayments than about saving the business.*

- ○ Yes, I already knew Toys "R" Us was owned by a private-equity firm
- ○ No, this is new information to me

A.4 Private Equity Literacy

Q15. Which statement best describes a private-equity fund?

- ○ It offers short-term savings products to private individuals
- ○ It pools money from private investors to buy shares of companies traded on the stock market
- ○ **It pools money to invest directly in private companies or to take public companies private**
- ○ Do not know

Q16. When people invest in private-equity funds, what is usually true about the fees?

- ○ **The fund charges a yearly management fee and takes a share of the profits**
- ○ The fund only charges a one-time entry fee
- ○ The fund does not charge any fees unless it makes a profit
- ○ Do not know

Q17. In a traditional private-equity fund, when can investors usually get their money back?

- ○ Anytime, like in a mutual fund
- ○ **Only after the fund ends and sells its investments**
- ○ Every quarter, provided the investor gives notice in advance
- ○ Do not know

Q18. Private-equity firms often use leverage when buying companies. What does this mean?

- ○ Private-equity firms borrow money to fund part of the purchase
- ○ **Companies take on debt when they are acquired by private-equity firms**
- ○ Private-equity firms lend money to the companies they buy
- ○ Private-equity firms ask their investors to borrow money to help finance the acquisition
- ○ Do not know

Q19. On average, how have private-equity funds performed compared with public stock markets?

- ○ Private-equity funds have earned about 10% higher returns per year
- ○ Private-equity funds' returns are fixed and do not depend on public markets
- ○ **Private-equity funds have earned about 3% higher returns per year**
- ○ Do not know

Q20. Traditional Private Equity funds often charge fees according to a "2 and 20" scheme. What does the "20" refer to?

- ○ That 20% of the total capital is paid in fees over the life of the fund
- ○ **A 20% performance fee on profits over a specified threshold**
- ○ A 20-month period during which withdrawals are not allowed
- ○ Do not know

Q21. What is "carried interest" in a Private Equity fund?

- ○ A fixed yearly payment to the fund manager
- ○ **A share of the profits the manager receives when returns exceed an agreed minimum**
- ○ The tax investors pay on investment profits
- ○ Do not know

A.5 Private Equity Investment Intention

Q22. How likely are you to include Private Equity investments in your retirement plan if they were available?

- ○ Definitely yes
- ○ Probably yes
- ○ Not sure, I would like to learn more first
- ○ Neutral
- ○ Probably no
- ○ Definitely no

A.6 Self-Perception

Q23. How many questions do you think you answered correctly?
Open-ended numerical response (0–12).

B Per-item correctness and "Don't know" rates

This appendix collects the per-question descriptive material that supports Section 4.1.2 in the main text. The figure plots per-item "Don't know" rates across both literacy modules; the table reports the full per-item correct, incorrect, and "Don't know" share alongside the Lusardi and Mitchell (2014) U.S. benchmarks for the Big Five items.

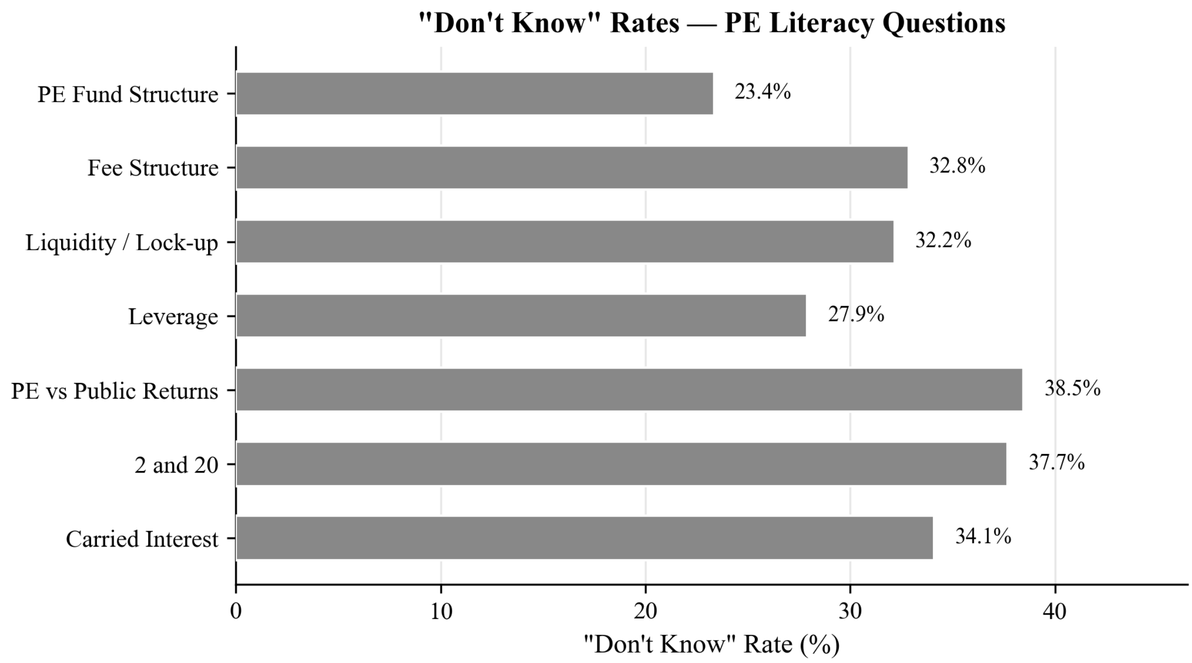

Figure 6.1: “Don’t know” rates by PE item.

C Demographic patterns of overconfidence

This appendix supports Section 4.3.3 in the main text by presenting the full demographic breakdown of the overconfidence gap. The summary table reports the overall calibration distribution. The demographic-breakdown figure shows the gap by age, education, income, and investment experience. The two KLH figures display the four-group sample composition and its breakdown by gender. The final table reports the underlying numbers for all demographic cuts.

Table 6.1: Overconfidence summary statistics

Statistic	Value
N (valid Q23)	963
N (missing Q23)	69
Mean self-assessed score	6.12
Mean actual score	5.19
Mean overconfidence gap	0.93
SD overconfidence gap	3.32
Median overconfidence gap	1.0
% Overconfident (gap $>$ 0)	50.4
% Calibrated (gap $=$ 0)	13.6
% Underconfident (gap $<$ 0)	36.0
Median actual score (KLH split)	5.0
Median self-assessment (KLH split)	6.0

Notes: Distribution of the overconfidence gap, defined as Q23 − total_score over the combined 12-item Big Five plus PE pool. $N = 963$ excludes respondents who skipped the optional Q23 self-assessment slider. Calibrated, overconfident, and underconfident shares partition the 963 valid respondents. Median actual and self-assessment scores are the splits used in the Kim et al. (2020) four-group classification reported in Table 4.5.

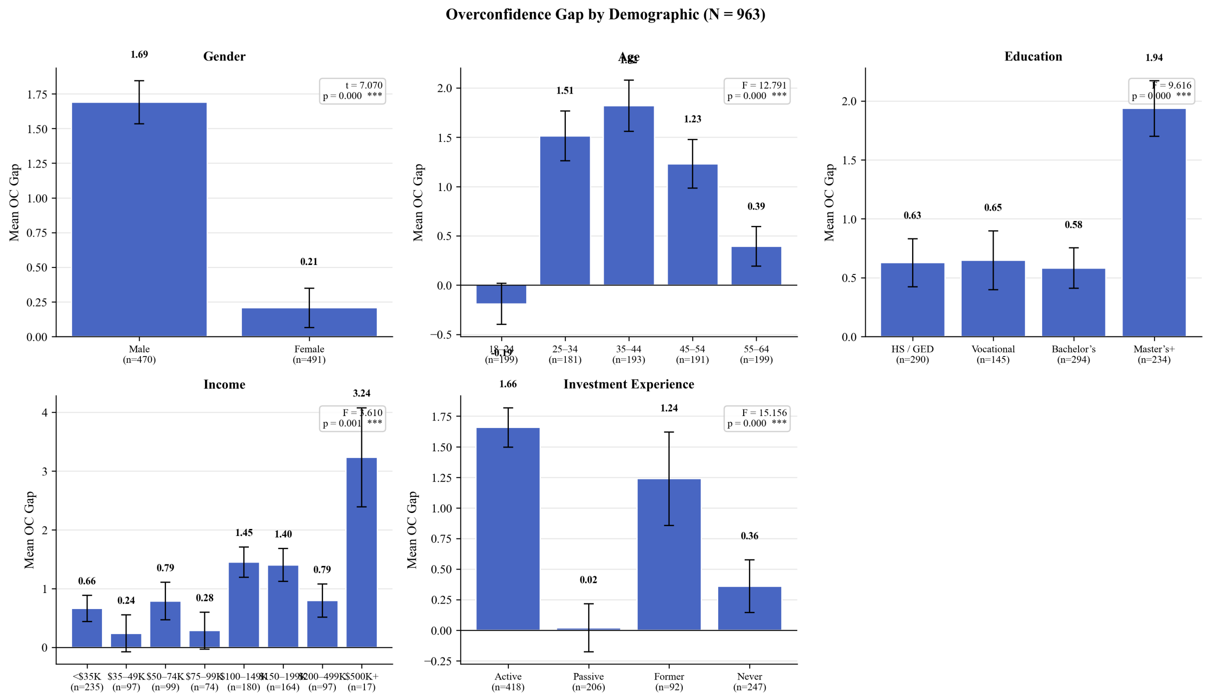

Figure 6.2: Overconfidence gap by age, education, income, and investment experience, four-panel view.

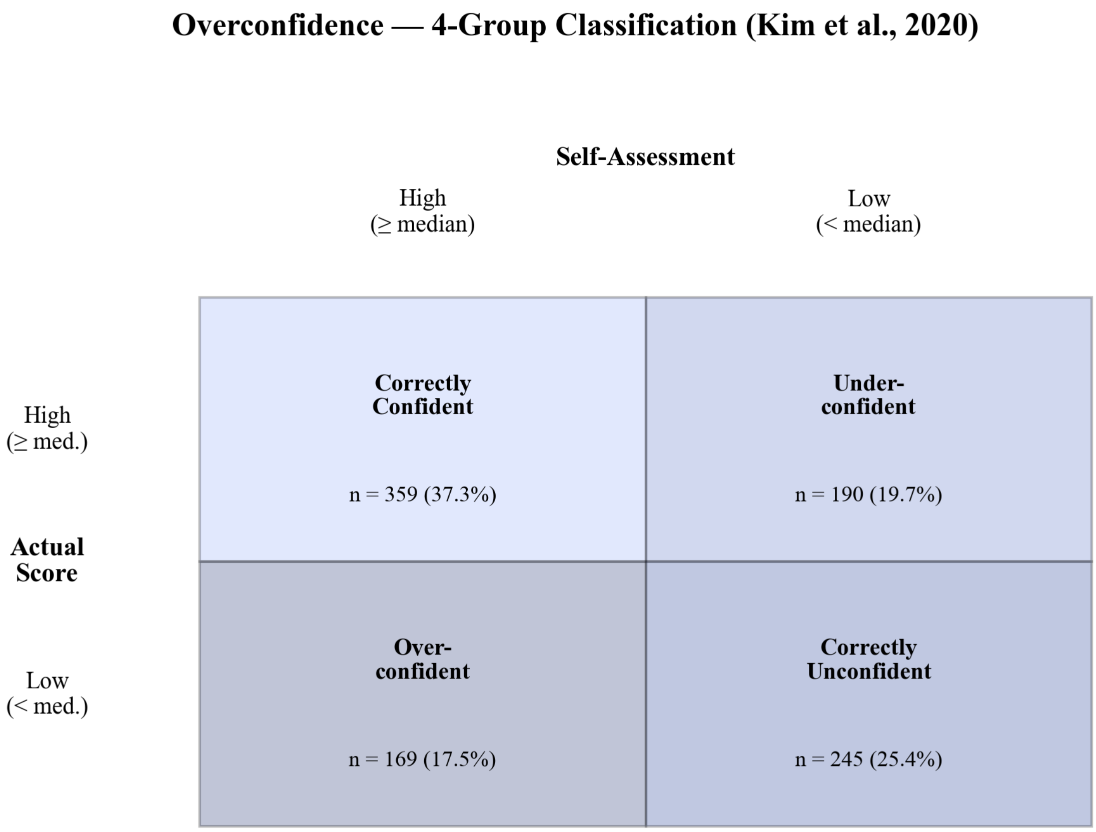

Figure 6.3: Sample composition under the KLH (2020) four-group classification.

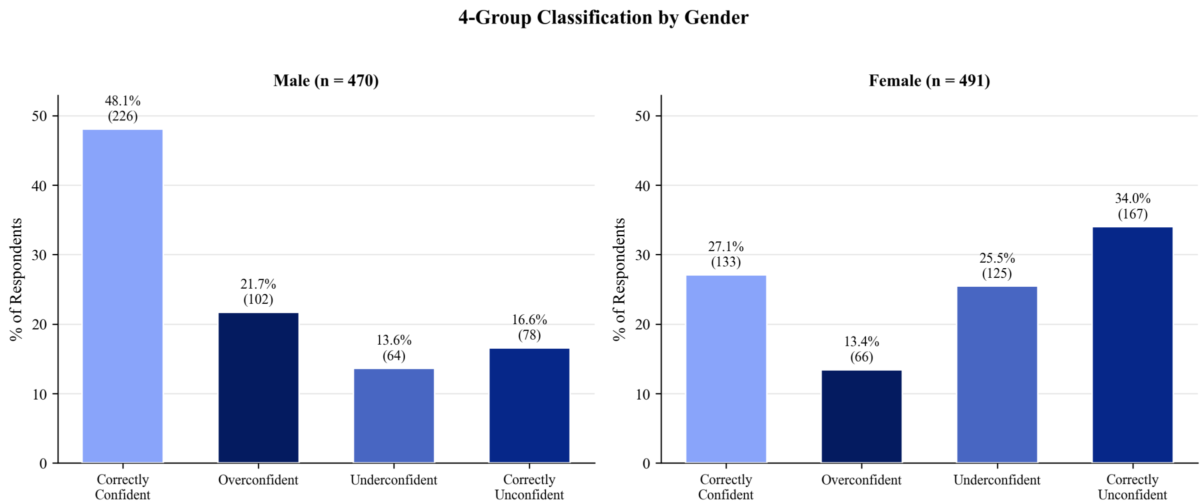

Figure 6.4: KLH four-group breakdown by gender.

Table 6.2: Mean and median overconfidence gap, n, and Overconfident-share by demographic subgroup.

D Framing supplementary material

This appendix collects the supplementary visualizations and cell-level descriptives for the exploratory H4b analysis (Section 4.6.3). H4b was not pre-registered; these materials report the bucket-level descriptive evidence and the full Specification 2b coefficient set.

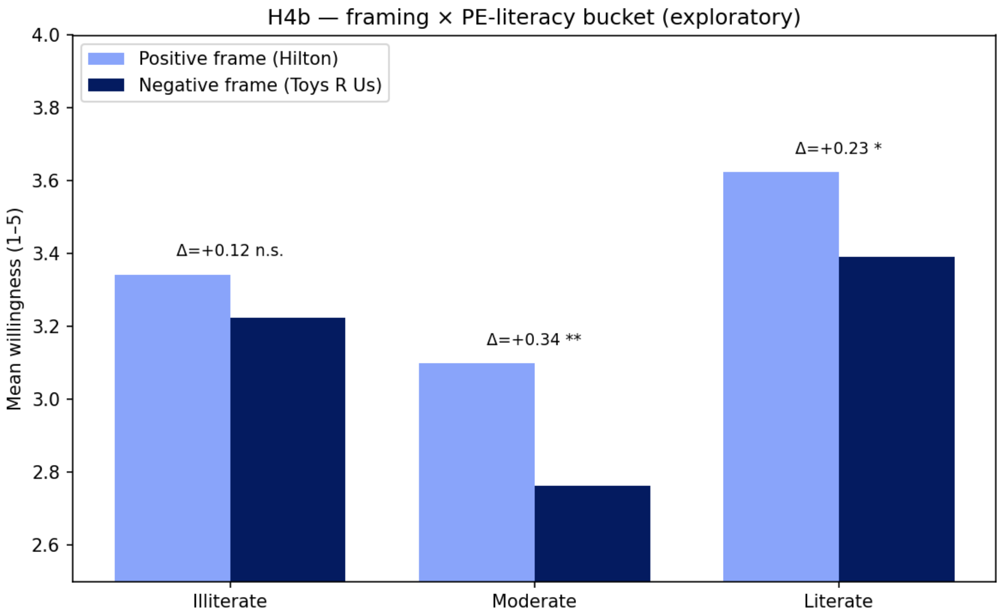

Figure 6.5: Mean Q22 willingness (1–6) by framing arm and PE-literacy bucket. Illiterate: Q15 wrong and pe_score < 4 ($n = 617$); Moderate: Q15 correct and pe_score < 4 ($n = 176$); Literate: pe_score ≥ 4 ($n = 239$). Δ values represent the positive-minus-negative frame difference within each bucket; stars denote significance from within-bucket Welch t-tests ($^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$). Total $N = 1{,}032$ (descriptive $N = 1{,}032$, Probit $N = 961$ after listwise deletion of demographic controls).

PE-literacy bucket	Framing	n	Mean Q22	SD
Illiterate	Positive (Hilton)	297	3.34	1.10
Illiterate	Negative (Toys “R” Us)	320	3.23	1.09
Moderate	Positive (Hilton)	100	3.10	0.94
Moderate	Negative (Toys “R” Us)	76	2.76	1.08
Literate	Positive (Hilton)	104	3.63	1.01
Literate	Negative (Toys “R” Us)	135	3.39	1.02

Table 6.3: Full 3×2 cell means and standard deviations of Q22 willingness (1–6) by framing arm and PE-literacy bucket. Bucket definitions as in Figure 6.5. Supports the H4b descriptive analysis in Section 4.6.3.

Table 6.4: Probit Specification 2b at Threshold A: framing × PE-literacy bucket interactions

Variable	Coef.	SE	p-value	AME
Constant	-0.442^{**}	0.196	0.024	–
Big Five score	0.030	0.039	0.450	0.008
Overconfidence gap	0.115^{***}	0.017	< 0.001	0.033
Positive framing (1=Hilton)	0.051	0.122	0.680	0.014
Moderate (vs Illiterate)	-0.606^{***}	0.225	0.007	−0.174
Literate (vs Illiterate)	0.233	0.156	0.136	0.067
Framing × Moderate	0.710^{**}	0.278	0.011	0.204
Framing × Literate	0.354	0.221	0.109	0.101
Age	-0.022^{***}	0.004	< 0.001	−0.006
Female	-0.402^{***}	0.100	< 0.001	−0.115
Income (ordinal)	0.056^{**}	0.028	0.050	0.016
Education (ordinal)	0.177^{***}	0.053	0.001	0.051
Has investment account	0.506^{***}	0.125	< 0.001	0.145
N		961		
Log-likelihood		−485.12		
Pseudo R^2		0.218		
Joint LR (both interactions = 0)		$\chi^2(2) = 7.70$, $p = 0.021$		

Notes: Probit estimates of the exploratory Specification 2b (Equation 3.7) at Threshold A (positive intent on Q22), $N = 961$. PE literacy is entered as a three-bucket categorical moderator (Illiterate = reference). The specification swaps the H4 framing × overconfidence interaction for framing × PE-literacy-bucket interactions, motivated by the comprehension-gating mechanism (H4b) discussed in Section 4.6.3. The joint likelihood-ratio test compares Specification 2b against a restricted model that retains the bucket dummies but drops both framing × bucket interactions. AME is the average marginal effect. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E Robustness regression tables

This appendix collects the full coefficient tables for the robustness exercises summarized in Sections 4.7.2 through 4.7.5. The headline numbers from each specification appear in the main-text body; the tables here reproduce the full coefficient set, standard errors, z-statistics, p-values and average marginal effects for completeness. The threshold-sensitivity coefficient comparison referenced in Section 4.7.1 is plotted in Figure 6.6 below.

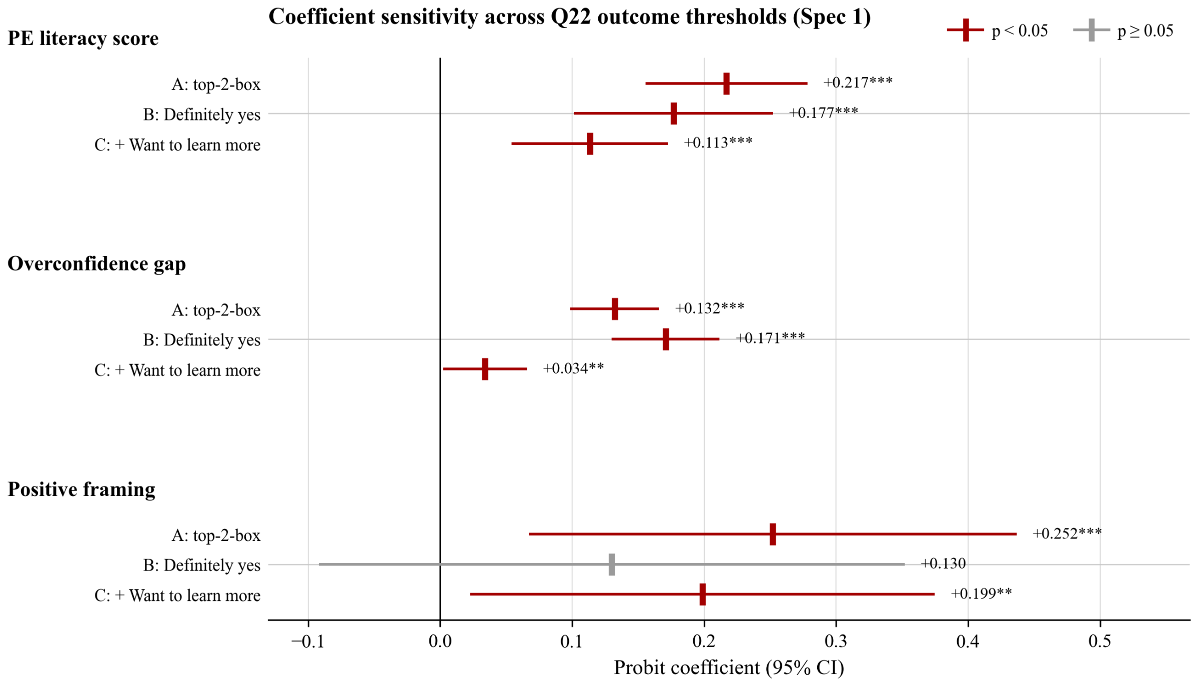

Figure 6.6: Coefficient forest plot for pe_score, oc_gap, framing, and framing × oc across Thresholds A and B.

Variable	Coef. (SE)	AME
PE score (reduced, 0–5)	0.251*** (0.050)	0.045***
Big Five score	−0.034 (0.046)	−0.006
Overconfidence gap	0.168*** (0.021)	0.030***
Framing (positive)	0.119 (0.114)	0.022
Age	−0.011** (0.005)	−0.002**
Female	−0.377*** (0.126)	−0.068***
Income	0.015 (0.036)	0.003
Education	0.090 (0.067)	0.016
Has invested	0.756*** (0.181)	0.137***
Constant	−1.921*** (0.286)	–
N	961	
Log-likelihood	−312.70	
Pseudo R^2	0.269	

Table 6.5: Specification 1 with PE literacy reconstructed as a 0–5 score excluding Q17 and Q19, estimated at Threshold B. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Variable	Coef. (SE)	AME
PE score	0.136^{***} (0.037)	0.026^{***}
Big Five score	-0.094^{**} (0.046)	-0.018^{**}
KLH overconfident (1/0)	0.819^{***} (0.155)	0.159^{***}
Framing (positive)	0.136 (0.110)	0.026
Age	-0.009^{*} (0.005)	-0.002^{*}
Female	-0.531^{***} (0.120)	-0.103^{***}
Income	0.030 (0.035)	0.006
Education	0.133^{**} (0.064)	0.026^{**}
Has invested	0.703^{***} (0.170)	0.136^{***}
Constant	-1.568^{***} (0.272)	–
N	961	
Log-likelihood	−336.95	
Pseudo R^2	0.212	

Table 6.6: Specification 1 with the continuous overconfidence gap replaced by the binary KLH "Overconfident" indicator. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Variable	Coef. (SE)	AME
PE score	0.183^{***} (0.039)	0.033^{***}
Big Five score	−0.049 (0.047)	−0.009
Overconfidence gap	0.171^{***} (0.021)	0.031^{***}
Framing (positive)	0.132 (0.114)	0.024
Age	-0.010^{*} (0.005)	-0.002^{*}
Female	-0.384^{***} (0.127)	-0.070^{***}
Income	0.020 (0.036)	0.004
Education	0.075 (0.067)	0.014
Has invested	0.822^{***} (0.186)	0.149^{***}
Constant	-1.937^{***} (0.292)	–
N	954	
Log-likelihood	−310.45	
Pseudo R^2	0.267	

Table 6.7: Specification 1 estimated on the gc = 1 subset (dropping the 21 Qualtrics-flagged respondents). The 1,011 raw gc = 1 respondents reduce to $N = 954$ after complete-case constraints across all controls. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table 6.8: Specification 1 with strict `pe_score` and DK-rebuilt `oc_gap_dk`

Variable	Coef.	SE	p-value	AME
Constant	-1.384	0.293	< 0.001	–
PE literacy score (strict)	0.050	0.040	0.208	0.008
Big Five score	-0.005	0.048	0.923	-0.001
Overconfidence gap (DK-rebuilt)	0.183	0.018	< 0.001	0.030
Positive framing (1=Hilton)	0.130	0.119	0.276	0.021
Age	-0.007	0.005	0.185	-0.001
Female	-0.248	0.134	0.065	-0.041
Income (ordinal)	0.009	0.038	0.806	0.002
Education (ordinal)	0.062	0.071	0.383	0.010
Has investment account	0.605	0.191	0.002	0.100
N			961	
Log-likelihood			-285.29	
Pseudo R^2			0.333	

Notes: Probit estimates of Specification 1 (Equation 3.4) at Threshold B (commitment-only on Q22), $N = 961$, with the *pure-H3* "Don't know" recoding described in Section 4.7.5. The PE literacy score retains the strict coding (DK = wrong); only the overconfidence gap is rebuilt on a DK-as-correct total, oc_gap_dk = Q23 − total_score_dk. Coefficients with standard errors and AMEs.

A note on threshold anchoring. The four robustness specifications in this appendix are anchored at the conservative commitment threshold (Threshold B, "Definitely yes" only) by design. Threshold B is the most demanding outcome and the hardest place to retain coefficient significance, so any specification that survives at B survives *a fortiori* at the more inclusive positive-intent primary outcome (Threshold A). The threshold-sensitivity exercise in Section 4.7.1 confirms this directly: every coefficient of interest produces a qualitatively identical result across A and B, with the OC coefficient slightly larger at B and the framing main effect appearing only at A.

F Note on the use of generative AI

This appendix discloses the use of generative-AI tools during the preparation of the thesis, in accordance with the thesis guidelines.

Tools and purpose. We used Anthropic's Claude Code and OpenAI's ChatGPT as auxiliary tools. Claude Code assisted with code refactoring, debugging, LaTeX formatting, writing and cross-reference checks. ChatGPT assisted with editorial support: clarity, concision, and consistency. Neither tool was used to formulate the research question or choose the empirical strategy.

Contribution to thesis quality. Where AI helped most was in catching the small things easy to miss: inconsistent labels, unclear captions, and bugs in the Python scripts that clean the survey responses, construct the literacy and overconfidence variables, and run the paired tests, probit regressions, and robustness reruns. It also tightened prose and verified that every cross-reference in the main text points where it should.

Risks and mitigation. We identify three principal risks: AI fabricates claims, AI misreads statistics, and AI prose drifts toward generic polish. Every suggestion was reviewed before inclusion; every numerical result came from our own Python scripts, re-run and checked against source data; interpretive claims were accepted only when they matched the empirical evidence; and any prose that drifted from an academic register was rewritten.

Insights. The tools made the revision process more systematic, revealing repetition, missing caveats, and places where the interpretive language around the framing manipulation needed clearer statement. The process highlighted the distinction between code assistance, editorial support, and substantive research judgment.

Author responsibility. All survey data were collected by the authors via Qualtrics, with respondents recruited from a U.S. retail-investor panel ($N = 1{,}032$). All empirical results were generated by the authors' Python scripts and reviewed by the authors. The authors retain full responsibility for the content, analysis, and conclusions of this thesis.