Partition at Your Own Risk:

Evidence on Risk-Taking Prevalence and Motives from the Field

Saurabh Bhargava
Carnegie Mellon University
sbhar@andrew.cmu.edu

Timothy Hyde
University of Pennsylvania
thyde@upenn.edu

This Version: December 2023

Abstract

Clarifying the motives for risk taking is crucial for economic theory, welfare analyses, and policy formulation. In practice, attempts at such clarification in the field are frequently hindered by an inability to observe perceived risk, decision complexity, and limited generalizability. We provide new evidence on risk-taking prevalence and motives from unusually rich field data describing the decisions, productivity, and beliefs of 20,133 employees across 18 large firms participating in a simple, all-or-nothing, goal-rewards program with $9.4 million in incentives. We find nearly half of employees chose lower goals than predicted by an expected utility benchmark with plausible risk preferences, leading to an average loss of 46% in potential rewards. Conservative choice persisted across financial stakes ($69 to $4,500) and employee experience and was notably larger for women, resulting in a 22% gender gap in rewards. Prominent departures from EU such as biased beliefs, non-linear decision weights, and gain-loss utility do not meaningfully improve explanatory power and we replicate choice patterns in experiments where options are explicitly represented as financial lotteries. We advance and experimentally validate a novel, heuristic, explanation that presumes risk aversion emerges from partition-dependent inference in the context of approximate pairwise comparisons. The heuristic not only explains substantially more choices in the lab and field than other benchmarks but explains most of the gender gap in conservatism. A series of experiments demonstrate how the heuristic could help to resolve seemingly contradictory empirical puzzles in insurance involving excess demand in low-risk settings (e.g., home insurance) and insufficient demand in high-risk settings (e.g., prescription drug coverage). The findings imply ostensible anomalies in the level, heterogeneity, and gender disparity of observed risk-taking across economic settings may stem from variability in heuristic adoption rather than variation in risk preferences and perceptions alone.

Acknowledgements: We extend a special thanks to Tim Houlihan and George Loewenstein for facilitating access to these data. We additionally thank Ned Augenblick, Linda Babcock, Karna Basu, Daniel Benjamin, Ben Bushong, Lynn Conell-Price, Stefano DellaVigna, Russell Golman, Kareem Haggag, Ben Handel, David Huffman, Alex Imas, Botond Koszegi, Yucheng Liang, Ted O'Donoghue, Ricardo Perez-Truglia, Alex Reese-Jones, Silvia Saccardo, Emmanuel Saez, Peter Schwardmann, Justin Sydnor, Lowell Taylor, Richard Thaler, Oleg Urminsky and seminar participants at UC Berkeley and Carnegie Mellon University for constructive feedback. We also thank our partners at BI Worldwide for providing data and considerable program detail. We are especially appreciative of the generosity of Ray Harms, Jenn Kelby, Mark Hirschfeld, and Betsy Schneider. Finally, we thank Stephanie Rifai, Cassandra Taylor, and several research assistants for excellent project support. The authors had full editorial discretion and any errors are attributable to them.

1 INTRODUCTION

Economists have long sought to understand the motives for financial risk-taking. Clarifying such motives has profound implications for economic theory, consumer welfare analyses, and optimal design of programs/policies in domains such as insurance, contract design, financial markets, and consumer protection. From the perspective of Expected Utility Theory (hereafter, EU or the standard model), the dominant framework in economics for understanding risk-taking, risk aversion among fully informed, utility-maximizing, decision-makers (DMs) reflects the diminishing marginal utility of wealth generated by the concavity of the utility function (von Neumann and Morgenstern, 1947). However, empirical evidence on decisions under risk (and uncertainty) gives rise to ostensible puzzles relative to standard benchmark predictions.[1] For example, economic decisions often imply risk preferences inconsistent with the standard model. In the lab, researchers have documented a degree of risk aversion in small-to-medium sized gambles (e.g., Holt and Laury, 2002) suggesting implausibly high aversion to risk at larger scales (Rabin, 2000). In the field, researchers have documented inefficiently high insurance demand in certain low-risk settings (e.g., home insurance) and inefficiently low demand in high-risk settings (e.g., elderly prescription drug coverage).[2] As another example, observed variation in perceived risk and preferences for risk are often seen as insufficient to explain observed variation in decision outcomes.[3]

In recent decades, researchers have proposed several departures from the standard model to explain risk-taking through channels such as biased beliefs, non-linear decision weights (e.g., Kahneman and Tversky, 1979; Prelec, 1998), or gain-loss utility (e.g., Kahneman and Tversky, 1979; Koszegi and Rabin, 2006; Gul, 1991; Loomes and Sugden, 1986). For instance, risk averse choice could stem from the systematic overestimation of risk, disproportionate weighting of unlikely outcomes, or an aversion to unanticipated out-of-pocket expenses. It could also emerge from non-standard processes receiving less attention in economics such as heuristics, limited attention, affect, cognitive processes, or hormones (see Kusev et al., 2017; Fox et al., 2015). In practice, research aiming to elucidate risk-taking motives in the field often investigates settings like insurance, betting markets, or game shows. Such inquiry is, however, frequently hindered by limited insight into perceived risks (researchers cannot typically observe private beliefs), complexity of the choice environment (e.g., consumers may lack full understanding of how to

---

[1] For simplicity, we largely elide the distinction between risk (knowledge of the probability distribution over potential outcomes) and uncertainty (a lack of such knowledge) in the paper.

[2] Several papers have documented the inconsistency between consumer demand for insurance and predictions of EU benchmarks (for review see Barseghyan et al., 2018). For example, relative to standard benchmarks, Abaluck and Gruber (2011) and Heiss et al. (2013) find evidence of insufficient insurance demand for prescription drug coverage, while Sydnor (2010) finds evidence of excess demand for home insurance. Several studies have documented sub-optimal health insurance choice in employer-sponsored contexts (e.g., Handel, 2013, Handel and Kolstad, 2015; Bhargava et al., 2017).

[3] Researchers have asserted that neither heterogeneity in risk (Cohen and Einav, 2007) or risk and risk preferences (e.g., Cutler and Zeckhauser, 2004; Barseghyan et al., 2013) can explain variation in insurance demand. In extensive experimental analyses, Jaspersen, Ragin, Sydnor (2022) find only modest correlation between risk attitudes and insurance demand.

evaluate insurance contracts), and limited applicability to broader contexts (e.g., game show contestants may be swayed by their unique environment).

We address these challenges with data describing the risk-taking decisions—and beliefs—of 20,133 employees across 18 large North American firms in the context of an employee goal-reward program called GoalQuest© (GQ). Developed by a US consultancy specializing in the administration of engagement programs informed by behavioral science, GQ was primarily designed to increase employee productivity. Due to its distinctive economic structure, data transparency, and generalizability, however, we see GQ as offering unexpected insight into the motives underlying financial risk-taking. To clarify, at the onset of each one-to-three month GQ program, employee participants privately chose a productivity goal from a simple, standardized, menu of three options, personalized based on their prior performance. Critically, each goal was associated with an all-or-nothing reward (e.g., selecting Goal 3 but only attaining Goal 2 yielded no reward) denominated in points that could be exchanged for non-monetary prizes at a predetermined rate. To promote ambitious goal setting, menus featured linearly increasing goals (e.g., 100 units, 110 units, 120 units) and non-linearly increasing rewards (e.g., \$100, \$300, \$600), so that, for most well-informed employees, Goal 3 would have maximized expected value (EV). Given this structure, and significant variability in productivity across periods, one can interpret goal choice as a decision between financial lotteries varying only in their risk and reward. The setting was further distinguished by our access to data not only on goal choice and employee productivity but also contemporaneous perceptions of risk. That is, our partnership with the consultancy led to the temporary adoption of an enhanced enrollment module capturing employee beliefs of attaining each of the three goals immediately following goal choice. These data permitted us the opportunity to test several belief-based motives for risk-taking otherwise difficult to assess. Finally, the broad economic stakes involved (with rewards ranging from \$69 to \$4,500), the demographic and occupational diversity of decision-makers, and the near-universal participation (98+ percent), allude to the generalizability of the findings. Collectively, employees in our primary sample stood to earn \$9.4 million in rewards; an additional 15,345 employees with choice-only data stood to earn \$8.2 million.

Our analyses of GQ data, supplemented with evidence from experimental goal-reward paradigms, yield three insights into risk-taking prevalence and motives. A first insight is to document substantial risk aversion and heterogeneity in employee choice. Specifically, nearly half of employees selected conservatively relative to an expected value (EV) maximizing benchmark, despite a simple decision menu and significant economic stakes. For those attaining at least the low-goal threshold, this conservatism led to an average loss of \$139, or 46% of the potential reward given ex ante optimal choice. Only 45 percent of employees chose the EV-maximizing goal, a share that did not meaningfully vary across reward size, employee tenure, or employee salary and was not driven by outlier programs.

A second insight is that prominent explanations for risk aversion from the literature can only account for roughly half of employee decisions. For example, adopting an EU benchmark with CARA utility within the interval, $r \in [0.0003, 0.005]$—a range whose upper bound indicates a degree of risk aversion so severe as to imply the rejection of a 50/50 gamble of losing $175 or winning an infinite amount—does not improve explanatory power relative to the risk neutral benchmark. While assuming increasingly severe risk aversion moderately reduces the share of choice characterized as conservative, it increases the share of choice characterized as aggressive by a commensurate degree. Notably, over 40 percent of employee decisions cannot be rationalized by *any* value of *r* within the specified interval. Similar challenges arise when attempting to explain conservative choice through the assumption of CRRA utility across a range of plausible lifetime wealth levels. Other prominent behavioral explanations for risk aversion—such as biased beliefs, non-linear decision weights, and gain-loss utility—also fail to meaningfully improve explanatory power. For example, while systematic under-confidence in relative beliefs of low, versus high, goal attainment could theoretically explain conservative goal choice, our data suggests substantial *overconfidence* in both relative and absolute beliefs of high-goal attainment.

A third insight is that the propensity for conservative choice differed markedly across binary gender. Women selected a low goal 26 percent more frequently than men despite minimal differences in productivity, implying similarly sized deficits in conservativism under standard benchmarks. We estimate that this gender-based difference in conservatism accounts for just under half of the 22 percent gender deficit observed in realized rewards. Of the non-standard benchmark models we tested, none could account for this gender disparity including models incorporating biased beliefs—in contrast to consensus beliefs of economists (Bandiera et al., 2022), we observe comparable levels of overconfidence across men and women.

For additional evidence on mechanisms and to eliminate potential confounds, we designed an incentive-compatible rewards program resembling GQ in the context of an online effort task. The paradigm permitted us to observe six goal choices per participant in a setting where we could confirm understanding of program rules, explicitly denominate rewards in dollars, and minimize, or even eliminate, motives pertaining to reputation, signaling and perceived costs of effort. The experiment yielded a pattern of goal choice, overconfidence, and conservatism mirroring that observed in the field and provided an even more emphatic rejection of previously tested benchmarks. And no evidence was found to support additional theories involving sorting via contextual cues such as self-perceived ability (Kamenica, 2008) or taste for competition (Niederle and Vesterlund, 2007) or flexible models allowing for heterogeneity in parameters such as risk preferences or loss aversion. An additional study revealed similar choice patterns from a hypothetical scenario-based GQ context with explicit communication of attainment likelihoods and from a context where goal-rewards were explicitly recast as roulette-styled financial lotteries with known probabilities of success.

Given the limits of existing models for explaining conservative choice, we propose a novel heuristic explanation. The heuristic, which we refer to as Pairwise Partition Dependence (PPD), posits that employees engage risky decisions through a series of successive, approximate, pairwise comparisons. Crucially, due to the phenomenon of partition dependence—the propensity towards biased inference in decision contexts that impose a specific partitioning of the state space (Fox and Rottenstreich, 2003; Fox and Clemen, 2005; Tversky and Koehler, 1974)—and relative goal proximity, the heuristic predicts most pairwise comparisons will lead employees to systematically underestimate the relative likelihood of high-goal attainment, resulting in greater conservatism than predicted by standard benchmarks. For tractability, we adopt the approach of Fox and Clemen (2005) in representing partition-dependent beliefs as a convex combination of subjective beliefs and an ignorance prior that treats each partition as equally probable.

As a concrete example, consider an employee deciding between Goals 2 or 3 (having already ruled out Goal 1). The heuristic stipulates that the employee will evaluate the goal pair by considering three decision-relevant partitions of the state-space: attainment of neither goal, attainment of Goal 2 but not Goal 3, and attainment of Goal 3. Given the relative closeness of most GQ goals with respect to perceived attainment, partition dependence applied to pairwise comparison predicts overestimation of the middle partition, biasing employees towards Goal 2. Importantly, even modest variability in heuristic adoption or in bias severity predicts potentially extensive heterogeneity in observed choice. While not discussed previously, the heuristic draws on several well-established conjectures from the literature involving relative evaluation, partition-dependent inference, and decision noise.[4] Moreover, the phenomenon of partition dependence is itself consistent with non-standard decision processes recently contemplated by the literature such as contingency neglect (e.g., Martínez-Marquina, Niederle, and Vespa, 2019; Sunstein and Zeckhauser, 2010), selective attention and salience (e.g., Bordalo, Gennaioli, and Shleifer, 2012; 2013), and imperfect memory (Hilbert, 2012).

We sought evidence for the proposed heuristic from a scenario-based hypothetical choice experiment that additionally queried extensive decision-relevant beliefs and process details. Beyond again replicating the characterization of choice from the field, the experiment provided evidence supporting the process assumptions and descriptive predictions of the heuristic. For example, the experiment affirmed participant use of proximal pairwise comparisons and revealed systematic and, often substantial,

[4] The notion of relative evaluation is integral to economic theories of reference-dependent preferences (e.g., Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Koszegi and Rabin, 2007) and comparative evaluation (e.g., Koszegi and Szeidl 2013; Bushong, Rabin, and Schwartzstein, 2021). The allowance of decision noise or judgment error is a common feature of decision-making models such as those involving bounded rationality (e.g., Simon, 1955) or stochastic preferences (e.g., Loomes and Sugden, 1982) and recent frameworks explicitly recognizing the role of noise in decisions (e.g., Kahneman et al., 2021).

underestimation of pairwise conditional likelihoods, both in the context of goal attainment and entirely distinct context of weather forecasts. The magnitude of relative inferential bias strongly predicted the choices of participants even after controlling for beliefs elicited non-contingently. As additional evidence, when participants were randomized to a menu designed to discourage partition dependence—via the display of accurate likelihoods of relative attainment—participants were 48% more likely to select the EV-optimal option than from a baseline menu displaying the equivalent information framed non-contingently. Further, participant response from a menu displaying no explicit likelihood information was indistinguishable from that produced from a menu displaying relative likelihoods adjusted for the presumed bias associated with pairwise partition dependence. Perhaps even more diagnostic, we found that the heuristic accurately predicted a substantially larger share of choice than the dozens of previously tested benchmarks in both the lab and the field. With a moderate allowance for noise, the proposed heuristic explained 83 to 92 percent of employee choice in the field and 59 to 72 percent of choice in the field, a substantial improvement in predictive accuracy relative to standard benchmarks. The heuristic also explained the near entirety of the gender gap in conservative choice, implying apparent differences in risk aversion across genders, or other population sub-groups, may reflect differences in heuristic adoption or bias severity rather than inherent differences in risk preferences or perceptions alone.

We speculate that the proposed heuristic could offer insights into financial risk-taking far more broadly than employee reward programs. Specifically, we see the heuristic as potentially applicable to economic menus that can be conceptualized as offering a choice between "nested" financial lotteries—that is, lotteries of increasing risk and reward that draw from an identical risk distribution—of the sort routinely found in portfolio allocation or insurance plan choice. We illustrate the heuristic's specific applicability to insurance markets with a theoretical framework that describes heuristic plan choice from a menu of contracts varying only in cost and degree of actuarial cost-sharing. The heuristic predicts systematic bias in insurance demand, relative to standard benchmarks, of a direction and magnitude largely determined by structural features of the insurance market. For example, the heuristic predicts excessively low demand in markets with high baseline loss risk (e.g., elderly prescription drug coverage) and excessively high demand in markets with low baseline loss risk (e.g., home, medical, auto)—predictions aligned with actual empirical findings from analyses of consumer demand in Medicare Part D (Abaluck and Gruber, 2011; Heiss et al., 2013) and home insurance (Sydnor, 2010).

We conclude with a final set of experiments designed to assess whether the proposed heuristic might help resolve seemingly contradictory empirical puzzles in the insurance literature. The experiments asked participants to make hypothetical insurance choices from stylized menus adapted from Medicare Part D or the US home insurance market (Sydnor, 2010). The menus to which participants were randomized strategically varied the framing of prospective risk information to either encourage or discourage bias associated with partition dependence. At baseline, participants exhibited behavior consistent with the empirical literature—inefficiently low demand for prescription drug coverage and inefficiently high demand for home insurance. Consistent with heuristic predictions, however, participants engaging menus designed to discourage partition dependence selected EV-optimal plans 39% (prescription drugs) and 35% (home insurance) more frequently than baseline despite the economic similarity of menus. An additional experiment demonstrated how the bias towards over-insurance in low-risk settings declines across exogenous increases to baseline risk, a dynamic predicted by the heuristic. Beyond highlighting the potential role of inferential bias for understanding insurance choice, the experiments reveal stark violations of descriptive invariance, an axiomatic assumption of most economic analyses of insurance. The experiments also illustrate how variability in heuristic adoption or bias severity could help account for the extensive heterogeneity in demand often observed by insurance researchers.

Our research relates to multiple, disparate, literatures in economics. First, we contribute to prior work seeking to clarify the prevalence and motives for financial risk-taking in the field (see Barseghyan et al., 2018). In addition to documenting substantial risk aversion in an environment distinguished by its simplicity and generalizability, we leverage data on contemporaneous beliefs to test—and reject—for roughly half the sample, standard and prominent non-standard motives for risk-taking from the literature. Second, the explanation we do propose for employee choice aligns with previous research emphasizing the importance of heuristics for understanding behavior in domains such as insurance (e.g., Ericson and Starc 2012; Bhargava et al. 2017; Jaspersen et al., 2022) or asset allocation (e.g., Benartzi and Thaler 2007). Because the PPD heuristic implies potentially substantial inferential error, our findings raise the possibility of biased welfare analyses and sub-optimal policy design in settings where researchers mischaracterize underlying decision processes. Third, our finding of substantial gender differences in conservatism (and reward compensation) contributes to work cataloguing gender differences in risk-taking in the lab and field (see Niederle, 2017). Our findings suggest that such differences may be attributable to differences in heuristic adoption or bias severity rather than inherent differences in risk expectations or preferences. Finally, a growing (largely experimental) literature has catalogued the influence of partition bias on resource allocation and inference (see Benjamin, 2019) and outlined how partition bias can be axiomatically incorporated into an economics choice framework (Ahn and Ergin, 2010). We provide evidence as to the potentially first-order importance of partition-dependent beliefs in understanding economic risk taking. Our proposed heuristic, for which we provide novel experimental evidence, suggests that partition dependence may bias choice whenever individuals are likely to make relative comparisons involving risk and uncertainty.

2 BACKGROUND

2.1 Institutional Background

GoalQuest® (GQ) is an employee-rewards program conceived and administered by BI WORLDWIDE (BIW), a private global consulting firm. The firm, founded in 1950, specializes in the design and delivery of a suite of proprietary programs that leverage principles from behavioral science (e.g., non-monetary rewards, goal-setting, personalization, symbolic recognition, lotteries, contests, communication, and feedback) to improve employee, firm, and consumer engagement. As of 2021, BIW had engaged 6 million individuals across 144 countries through its various products. As of the same date, according to third-party estimates, the firm had approximately 1,500 employees and annual revenues between $500 million to $1 billion. Described as the world's only patented incentive-based sales program, GQ was designed to motivate employee productivity through self-selected performance goals tied to all-or-nothing non-monetary rewards.[5] As of 2018, BIW had administered over 1,000 GQ programs to over 1 million participants at firms primarily in the United States, Canada, and Europe since its 2001 inception. While marketed as a sales incentive program, our data indicate that the program has serviced a significant share of employees engaged in customer service and retention (e.g., call centers) across a diversity of sectors (e.g., communication, health care, manufacturing, financials, consumer discretionary).

2.2 GQ Program Overview

Although GQ programs are administered across many contexts, they share a standardized structure that permits comparability across programs. Participation entails three phases: an enrollment period (employees selected a goal), a performance period (employees logged performance towards their goal), and, for those achieving their selected goal, a reward redemption period. During the initial phase, employees were directed to an online portal where they proceed through a simple web-flow.[6] The web-flow itself consists of three parts: a program overview, an explanation of program rules, and goal selection. Employees select a goal from a menu of three personalized options (Goal 1, Goal 2, Goal 3) each associated with an all-or-nothing reward denominated in points (see Appendix for screenshot).[7] BIW promotes the program as having a 98 or 99 percent participation rate among eligible employees.[8]

In 2014, we asked BIW to implement an enhanced enrollment process to elicit additional data from employees, namely their beliefs regarding goal attainment. Enhanced enrollment added a fourth

[5] The World Intellectual Property Organization Publication Number associated with GQ is WO 01/13306 A2 (February 2001).
[6] While the design of the enrollment portal was standardized across programs at any point in time, its design evolved over time.
[7] In some programs, the goal selection period may have briefly overlapped the performance window.
[8] While we cannot directly verify participation statistics, high participation rates are plausible due to marketing and communication during the pre-period, the administrative ease of enrollment, and often-valuable rewards.

phase to the enrollment process: immediately after selecting their goal, respondents were prompted to complete a brief survey. The survey asked employees to estimate their perceived likelihood of goal attainment for each goal: "On a scale from 0% (no chance) to 100% (absolute certainty), how likely is it that you will meet or exceed each of the following achievement levels?" (the response scale was indexed in 10-point increments). Employees were additionally asked about their binary gender, age, and tenure with the firm. While the survey was optional and rewards did not depend on completion, survey participation across our sample was 60 percent.

Following goal selection, employees transitioned to a 30- to 90-day performance period during which they attempted to achieve their selected goal. In most programs, participants were able to log onto the website to check their progress or to remind themselves of their selected goal.[9] At the close of the performance period, employees who attained their goal exchanged reward points for a reward in the GQ marketplace. The non-monetary rewards included major electronics (e.g., flat-screen television), event packages, vacations, household items (e.g., luggage), or recreational items (e.g., golf clubs). Employees were educated as to the approximate conversion rate between points and the dollar value of the associated rewards during program marketing; for many programs, employees were familiar with the conversion rate through other BIW programs.

2.3 Goal and Reward Structure

We highlight three distinctive structural features of GQ relevant for the present research. First, based on the presumed importance of personal choice and personalization for engagement, GQ required employees to self-select a goal from a personalized menu of options. The personalization reflected the application of some uniform rule to an employee's productivity during an earlier baseline period (e.g., productivity during the prior quarter).[10] Most program menus featured additively linear goals of the form: $f(x_b)$, $f(x_b) + a$, $f(x_b) + 2a$, where $\mathrm{f}(x_b)$ is a function of baseline productivity, x_b, (e.g., $f(x_b) = 1.05x_b$) and a denotes some increment, potentially itself a function of baseline productivity (e.g., 10 or $0.10x_b$). Employees within a program were usually segregated into a few distinct groups based on comparability in factors such as baseline performance, experience, or job level. While personalization rules could vary across groups—permitting, for example, GQ to assign all new employees to a menu not informed by baseline data—menus within each group were personalized using the same rule. Employees were not given explicit encouragement to select any specific goal via recommendations, defaults, or persuasion.

[9] According to BIW, most programs provided data on intermediate performance to employees. In some programs, intermediate feedback was not technically feasible, or necessary, to track.

[10] The calculation of baseline performance was jointly determined by BIW and each firm on a program (and often group) specific basis based on considerations of data availability, employee tenure, and seasonal variation in productivity. For many programs, the baseline was calculated from employee performance over a recent period of similar duration to the program. New employees without historical performance were given a non-personalized menu.

Second, based on the presumed motivational potency of self-selecting high goals, GQ implicitly encouraged high goal choice by ensuring that high goals were financially attractive. Specifically, in contrast to the additively linear increase in goals, rewards typically increased in non-linear increments. For example, many menus followed the *k*, 3*k*, 6*k* reward structure, where *k* was approximately 1 percent of an average employee's salary over program duration. Due to the non-linearity of rewards, and their all-or-nothing nature, we estimate that, assuming rational expectations, Goal 3 maximized expected value for 84 percent of employees (Goal 2 maximized EV for 11 percent of employees). Lastly, the rewards associated with each goal were non-monetary, due to a belief that non-monetary rewards would be more motivating than monetary rewards of similar value.

3 THEORETICAL FRAMEWORK OF GOAL CHOICE

We now introduce a theoretical framework to organize our analysis of conservative goal choice. We represent the GQ goal choice as a decision between two simple lotteries and assume employees select the goal that maximizes their expected utility given their beliefs of goal attainment. We then amend the model to incorporate systematic departures from the standard framework, informed by the literature, such as the potential for biased beliefs, non-standard decision-weights, and reference-dependent utility.

3.1 Generalized Expected Utility Framework

We begin by outlining a generalized framework to describe how a utility-maximizing employee selects a productivity goal associated with an all-or-nothing reward from a menu of options. We represent goal choice as the decision to participate in one of two available lotteries, $G_n \in [G_h, G_l]$. Each lottery yields a reward x_n with some probability s_n and 0 with some probability $(1 - s_n)$. The high goal has a strictly higher reward and lower likelihood of attainment than the low goal, $x_h > x_l$ and $s_h < s_l$ and we assume an inter-temporal discount rate of 1, rendering the timing of reward receipt immaterial. We further assume that goal realization depends on employee-specific ability and employer-specific productivity shocks but not an employee's goal choice.

If we denote subjective probabilities of goal attainment by $\hat{s}_n$, then we can describe an employee's expected valuation of a goal-lottery as: $\mathrm{V}(G_n) = \pi(\hat{s}_n)\ v(x_n, \theta)$. Here, $\pi(.)$ denotes a decision-weighting function applied to some subjective probability of goal attainment and $v(.)$ is an always increasing function, potentially dependent on a reference point, θ, that represents an employee's preference for rewards. Given the employee is required to select a goal, the employee will select the low goal if $\pi(\hat{s}_l)v(x_l, .) \geq \pi(\hat{s}_h)v(x_h, .)$.

We first consider a baseline case in which an employee with rational expectations selects a goal that maximizes expected value. We define rational expectations, $\hat{s}_n^r$, regarding the attainment of each goal, as $\hat{s}_n^r = E(s_n \mid \Phi) = s_n + \varepsilon$. The parameter, Φ, denotes the information set available to an employee at the time of goal choice and ε is a normally distributed, mean-zero, error term with constant variance. At baseline, an employee will choose the low goal only if it maximizes expected value, $\hat{s}_l^r / \hat{s}_h^r > x_h / x_l$.

Assuming that the high goal maximizes expected value, $\hat{s}_l^r / \hat{s}_h^r < x_h / x_l$, then conservative choice—i.e., choice of a goal lower than the optimal goal—under the baseline benchmark can arise from utility-based risk aversion, non-standard beliefs, non-linear decision weights, or gain-loss utility.

3.2 Conservatism Due to Risk Aversion [$\pi(s) = s$, $v(x_n, .) = u(x_n)$]

An initial explanation we consider for conservative goal choice is risk aversion due to the diminishing marginal utility of wealth. We incorporate utility-based risk aversion via a parametric utility function from the constant absolute risk aversion (CARA) family, $v(x_n, .) = u(x_n, r)$, where r captures an employee's attitude towards risk (i.e., $r > 0$ implies risk aversion, $r = 0$ denotes risk neutrality; we ignore the possibility of $r < 0$):

$$u(x_n, r) = \begin{cases} -\dfrac{1}{r}\exp(-r x_n), & r > 0 \\ x_n, & r = 0 \end{cases}$$

The choice of a CARA function permits us to represent risk attitudes with a single parameter but implies the irrelevance of an employee's prior wealth for risk preferences.[11]

An employee averse to financial risks ($r > 0$) will choose conservatively if:

$$r > \frac{\ln\left(\frac{\hat{s}_l^r}{\hat{s}_h^r}\right)}{x_l - x_h}$$

The decision rule implies that conservative goal choice is positively increasing in the degree of risk aversion, relative expectations of goal attainment, and the gap between high and low goal rewards. Practically, we evaluate the descriptive accuracy of this benchmark for risk aversion parameters within some range of plausibility $r < r'$, where we determine r' by appealing to the literature and by assessing the degree of risk-taking the parameter implies in the context of simple lotteries involving financial stakes comparable to those found in GQ. We additionally evaluate the possibility that employees exhibit heterogeneity across risk preferences by evaluating the share of goal choice that can be rationalized by allowing employees to have *any* degree of risk aversion within an interval of plausibility, $r_i \in [0, r']$.

[11] In the Online Appendix, we consider arguably more realistic utility functions featuring constant relative risk aversion (CRRA) and show that the simplifying CARA assumption does not qualitatively affect model predictions.

3.3 Conservatism Due to Non-Standard Beliefs [$\hat{s}_n \neq E(s_n)$]

We next consider the possibility that conservative goal choice reflects the non-standard beliefs of an employee otherwise conforming to the standard expected utility framework. We denote non-standard beliefs with a goal-specific multiplicative constant, γ_n, such that $\hat{s}_n = \gamma_n s_n + \varepsilon$. Consequently, $\gamma_n > 1$ implies overconfidence while $\gamma_n < 1$ implies under-confidence. A risk averse employee with biased beliefs will select the low goal if:

$$r > \frac{\ln\left(\frac{s_l}{s_h}\right) + \ln\left(\frac{\gamma_l}{\gamma_h}\right)}{x_l - x_h}$$

The decision rule implies that conservative goal choice increases in the ratio of overconfidence regarding the low, relative to the high goal, γ_l/γ_h. This implies that low-goal overconfidence is not sufficient, by itself, to produce conservative choice since such overconfidence may be accompanied by equal, or even stronger, overconfidence about the high goal.

3.4 Conservatism Due to Non-Standard Decision weights [$\pi(s) \neq s$]

We proceed to consider whether the adoption of non-linear decision weights might help to explain conservative goal choice. Researchers have advanced several probability weighting functions to address violations of expected utility in which people appear to overweight highly improbable outcomes and underweight highly probably outcomes. We adopt arguably the most popular of these functions, the inverse-S shaped function proposed by Prelec (1998): $\pi_n = \exp\left(-(-\ln s_n)^{\alpha}\right)$.

The decision rule for a utility-maximizing employee with non-linear decision weights is given by:

$$r > \frac{\ln\left(\frac{\pi_l(s_l)}{\pi_h(s_h)}\right)}{x_l - x_h}$$

In theory, if employees were to systematically underweight moderate-probability outcomes (e.g., high goal) relative to higher-probability outcomes (e.g., low goal), a non-linear weighting function might help explain conservative goal choice.

3.5 Conservatism Due to Loss Aversion [$v(.) = v(x_n, \theta)$]

Finally, we consider the possibility that conservative goal choices may arise from prospective loss aversion in the context of gain-loss utility (Kahneman and Tversky 1979; Tversky and Kahneman 1992). Loss aversion has been advanced as a possible theoretical explanation for small- to moderate- scale risk aversion (Rabin (2000); Rabin and Thaler 2001) and has been cited as a possible explanation for risk aversion in the field across a range of contexts. While GQ employees engage prospective, as opposed to actual, losses, expectation-based approaches to gain-loss utility (e.g., Koszegi and Rabin, 2006; Gul,

1991) and research interpreting goals as reference points (Heath, Larrick, and Wu 1999) suggest the possible relevance of loss aversion in explaining conservative goal choice.

One practical challenge for assessing models of gain-loss preferences is the absence of clear theoretical guidance as to how to specify the reference point, the functional form of gain-loss utility, and the magnitude of loss aversion. In service of assessing the breadth of plausible representations, we appeal to the theoretical literature and practical considerations to identify, and then test, an exhaustive set of potential parameters.[12] With this in mind, we represent gain-loss utility, given some reference point, θ, with the following generalized function:

$$v(x_n, \theta) = \begin{cases} \eta m(x_n) + \ u^{+}(x_n - \theta), & \text{for } x \geq \theta \\ \eta m(x_n) + \lambda u^{-}(x_n - \theta), & \text{for } x < \theta \end{cases}$$

The term $m(x_n)$ is a strictly increasing function capturing consumption utility; u^{+} is a concave utility function capturing gain-loss utility for gains; u^{-} is a convex utility function capturing gain-loss utility for losses; and λ is the loss aversion parameter. The framework assumes consumption and gain-loss utilities are additively separable and we specify η as a scaling factor such that $\eta = 0$ collapses to a model with only gain-loss utility.[13] We evaluate this function for a range of potential reference points, loss aversion parameters, and scaling factor values.

4 DATA AND SAMPLE CONSTRUCTION

Our analysis of financial risk-taking relies on program- and employee-level administrative data from BIW. The employee-level data describes goal choice, productivity, beliefs of goal attainment, and demographic detail. The program-level data describes the identity of each firm, the dates of program administration, rules used to assign employees to distinct groups, and details of the GQ menu configuration. In this section, we describe the construction of the primary sample, summarize its key features, and define the variables central to the subsequent analysis.

4.1 Primary Sample

Our main evidence on employee behavior and beliefs draws from what we refer to as the *primary sample*. We constructed the primary sample—comprising 20,133 employees across 18 firms, 34

---

[12] For example, one resource for identifying candidate reference points is provided by recent work that evaluated the success of potential reference points, in the context of gain-loss utility, for explaining menu-based risky choice in the lab (Baillon, Bleichrodt, Spinu 2020). Specifically, the authors considered prospect-independent (e.g., status quo, the high outcome, the highest probability option, the highest option a person is certain to achieve) and prospect-dependent (e.g., the selected option, the expected value of the selected option) reference points.

[13] The composite functional representation of gain-loss utility draws on representations from prior work (Sugden 2003; Kobberling and Wakker 2005; Koszegi and Rabin 2006, 2007).

programs, and 232 distinct groups—by applying screening restrictions to an original dataset from BIW. This original dataset (n = 38,661) reflects the universe of data from GQ programs administered between 2014 to 2018 in the US or Canada with enhanced enrollment, at least 100 fully participating employees, and electronically archived data.[14] Before arriving at the primary sample, we generated an *expansive sample* (n = 35,478) that excluded roughly 8% of employees from the original dataset for whom a key data field was missing, the data was not consistent, or we inferred incomplete participation.[15] We then restricted the expansive sample to employees providing internally consistent beliefs during enhanced enrollment to create the primary sample.[16] In comparing the samples, employees completing enhanced enrollment were moderately more likely to select aggressive goals and modestly more likely to attain them, implying the conservatism and sub-optimal choice we subsequently document may, if anything, underestimate the actual degree of conservatism and sub-optimal choice in the broader employee population.[17] For robustness, we reproduce key analyses for the expansive sample in the Appendix.

Table 1 summarizes overall sample statistics as well as group-level (duration, financial stakes) and employee-level (age, gender, tenure, inferred income) characteristics for the primary sample.[18] On average, we observe data for 592 employees per program (IQR: 208 to 703) and 87 employees per group (IQR: 12 to 103). The groups varied with approximate uniformity across either 30, 60, and 90-day program durations (two outlier programs lasted 45 and 120 days). The distribution of potential reward values was asymmetric, such that 10 percent of employees engaged decisions with rewards averaging $2,150, despite a group-level average of $607 and an employee-level average was $466. Collectively, employees in the primary (expansive) sample had the opportunity to earn $9.4 ($17.5) million in possible rewards. The table also conveys the diversity of the sample across gender, age, and tenure.

4.2 Goal Choice, Employee Productivity, and Goal Attainment

Our analysis relies on measures of goal choice, productivity, and beliefs. We describe goal choice by summarizing choice shares and through indicators characterizing the optimality, aggressiveness, or conservativeness of each goal. For the latter, we compare the value of the selected goal with that of non-

---

[14] Data for a small number of programs was not archived by BIW. The size cutoff was necessitated by resource constraints.
[15] 5.2% of the original sample was missing critical data fields, 0.3% of the sample had contradictory award data, and 2.8% of the sample was identified as likely not participating or completing the program based on implausibly low performance reports.
[16] An employee was tagged as having inconsistent beliefs if such beliefs implied a strictly greater likelihood of attaining a higher, relative to a lower, performance threshold. We excluded 2,215 employees, or 9.5% of enhanced enrollees, for this reason.
[17] We compared the expansive and primary sample across observable factors through regressions of the following form: $y_{i,l} = \alpha + \theta enhance_i + \pi_l + \varepsilon$, where *y* indicates an observable factor, *enhance* indicates completion of enhanced enrollment and π_l denotes group-level dummy variables. The most notable difference is that enhanced enrollees were 0.091 more likely to select Goal 3 (baseline choice share of 0.34) and 0.031 more likely to attain Goal 3 (baseline attainment of 0.28) than counterparts. The comparison suggests that conservatism and sub-optimal choice documented in the primary sample not only exists but may be exaggerated in the expansive sample (we confirm this intuition in Section 5).
[18] Some firms participated in multiple GQ programs sequentially at the same location, so a small number of employees appear in the sample multiple times in different programs.

selected goals under the specified benchmarks. To facilitate comparisons across programs, we describe productivity with two normalized measures: (i) productivity relative to baseline (a figure determined by BIW, typically from prior performance) and (ii) productivity relative to the Goal 3 threshold.[19] Lastly, for each employee, we calculate goal attainment for both selected and non-selected goals.

Table 2 summarizes employee choice, productivity, and attainment. The table indicates 44 percent of employees selected the highest goal with an approximately even split across the other two goals. It also conveys a correlation between goal choice and productivity, consistent with more productive employees sorting themselves into higher goals (or alternatively, higher goal choice may lead to elevated performance). The table also shows that while only 29 percent of employees attained the highest goal, 66 percent of those attaining the lowest goal also attained the highest goal (i.e., 0.29/0.44). Figure 1, which depicts average shares of Goal 2 and 3 choice (residual choice is Goal 1) across programs and groups by entity size, provides insight into the distribution of choice shares across entities. Across meaningfully-sized entities, it both suggests the lack of outliers and non-trivial variation in choice patterns.

4.3 Employee Beliefs

We calculate two measures of employee beliefs: (1) subjective beliefs elicited through enhanced enrollment, (2) and econometric estimates of ex ante rational expectations. The subjective belief, $\hat{s}_{k,i}$, of employee, *i*, attaining goal *k*, reflects the employee's reported belief from the enhanced enrollment module (i.e., for tractability, we recoded beliefs of 0% to 1% and from 100% to 99%). To estimate an employee's rational expectation of goal attainment, $\hat{s}^{r}_{k,i}$, we appeal to a strategy routinely used in research on insurance. The strategy involved constructing employee sub-samples by program group and goal choice and then predicting each employee's ex ante likelihood of attaining each goal by adjusting the average sub-sample attainment rate (excluding the reference employee) by observable covariates. The exercise effectively assumes that one can proxy for rational expectations with the average realized attainment of similar others.[20]

There are two notable patterns from the summary of beliefs data provided by Table 3. First, the table reveals a correspondence between goal choice and attainment, again suggesting the non-randomness of employee choice. Second, comparing subjective and rational beliefs indicates substantial overconfidence among employees regarding future productivity. Specifically, employees were

---

[19] We did not have baseline data for 16 percent of employees. In most cases, this reflects the lack of past performance data for new employees or programs where performance goals were defined without reference to a baseline.

[20] More specifically, we initially estimated the following leave-out regressions for each employee *i* and goal $k \in [1,2,3]$: $\bar{s}_{k,l,-i} = \alpha + \boldsymbol{Z}\gamma + \pi_l + \varepsilon$. Each regression predicts average group-level attainment for each goal, $\bar{s}_{k,l,-i}$, leaving out employee *i*, as a function of employee characteristics included in vector **Z** (age, tenure, gender) and group fixed effects, π_l. (We estimated regressions at the program level to increase the precision of covariate estimates). We then calculated an employee's rational expectation of attaining goal *k*, as $\hat{s}^{r}_{k,i} = \hat{\alpha} + \boldsymbol{Z}\hat{\gamma} + \hat{\pi}$.

substantially overconfident, on average, with respect to every goal; such overconfidence was greater for higher, relative to lower, goals.

5 CHARACTERIZATION OF GOAL CHOICE BY BENCHMARK MODEL

We now characterize employee choice relative to predictions of the benchmark models outlined in the theoretical framework. Tables 4 and 5 report the share of optimal (goals matching benchmark predictions), conservative (goals lower than benchmark predictions), and, for completeness, aggressive (goals higher than benchmark predictions), choice for each benchmark. To clarify the economic consequences of sub-optimal choice, the tables also report measures of counterfactual losses associated with each benchmark. Finally, to better understand the moderating role of financial stakes and experience, we report optimal choice shares across reward size quartile and employee tenure.

5.1 Expected Utility with Risk Neutrality

We begin by characterizing choice under a baseline benchmark of risk-neutral expected utility with rational expectations (equivalent to an EV-maximizing benchmark). As presented in Table 4, under this benchmark, 45 percent of employees chose optimally with most remaining employees choosing conservatively. For employees engaging in conservative choice who attained at least the low goal, the average realized reward of $164 relative to the average benchmark-optimal reward of $303 implies a counterfactual loss of 85 percent. The share of optimal choice did not vary across reward size or employee tenure. Appendix Figure A2 shows the cumulative distribution of counterfactual loss overall and by goal choice and highlights the concentration of loss among low-goal choice.

Non-Standard Beliefs. It is possible that conservative choice could reflect widespread bias in perceptions of goal attainment. For example, systematic overconfidence about the relative likelihood of attaining low versus high goals, or systematic under-confidence about high goal attainment, could lead utility-maximizing employees to select conservatively. To assess such possibilities, we characterize choice under a modified benchmark incorporating subjective beliefs. While the transition from rational expectations to subjective beliefs led to a modest improvement in optimal choice share, it shifted the conservative choice share from 49 to only 48 percent (Table 4). Figure 2 provides more detail on the share of high- and EV-optimal goal choice across varying subjective values of Goal 3 and Goal 2. The figure suggests an underlying regularity in employee decisions with high-goal choice shares rising with the relative attractiveness of the high goal (lower left quadrant) and optimal choice shares with distance from the equal-value diagonal. Despite this regularity, however, the figure alludes to pervasive sub-optimal choice. At a minimum, roughly 15 to 20 percent of employees choose sub-optimally in nearly every cross-section with most cross-sections featuring sub-optimal choice majorities. Jointly, the panels point to a strong preference for Goal 2 over Goal 3 excepting high differences in value.

Table 3 provides insight as to why biases in beliefs do not meaningfully improve benchmark predictions. The table documents substantial employee overconfidence both with respect to absolute beliefs of goal attainment for each goal and with respect to beliefs about the relative likelihood of attaining Goal 3 compared to the other goals. The appendix provides additional clarity as to the relationship between beliefs and choice by depicting the sizable difference in the distribution of rational and subjective expectations by goal (Figure A3) and the relative unimportance of assumed information regime in the cumulative distribution of counterfactual reward losses (Figure A4).

5.2 Expected Utility with Risk Aversion

We now consider the possibility of attributing conservative choice under the baseline benchmark to utility-based risk aversion. We model risk aversion by assuming a CARA utility function with risk preferences within the interval, $r \in [0.0003, 0.005]$. To appreciate why one can interpret this interval (which was informed by our reading of the literature) as subsuming the range of plausible risk attitudes, we can translate what risk preferences within the interval imply for gambles involving potential losses comparable in magnitude to GQ rewards. For example, consider a simple lottery involving a 50% chance of losing \$175 (roughly the 25th percentile of GQ rewards) and a 50% chance of some unspecified gain. A risk aversion parameter of $r = 0.0003$ implies an employee would accept any such gamble if the potential gain exceeded \$184—a modest, but seemingly plausible, degree of risk aversion. The same employee would accept any 50/50 gamble involving a potential loss of \$350 (roughly the median GQ reward) so long as the potential gain exceeded \$391. The interval's upper bound, $r = 0.005$, implies the rejection of *any* 50/50 gamble involving a potential loss of \$175 or \$350, even if the potential gain was infinite.

As depicted in Table 4, the assumption of modest risk aversion ($r = 0.0003$) does little to shift the characterization of choice relative to a risk neutral benchmark, under rational expectations or subjective beliefs. The assumption of more severe, verging on implausible, risk aversion ($r = 0.005$) moderately shifts the characterization of choice from conservative to aggressive. It does not, however, meaningfully shift the ostensible share of optimal choice, the magnitude of counterfactual loss, nor the absence of moderation in optimal choice by reward size or tenure, relative to the prior benchmarks. Figure 3, which plots the share of optimal and conservative choice under the EU benchmark across different assumptions of beliefs, graphically depicts the insensitivity of choice characterizations to potential reward size.

Figure 4 provides intuition for how variation in assumed risk preferences affects choice characterization. The figure depicts the share of optimal goal choice for the expected utility benchmark given a "super-plausible" interval, $r \in [0.00, 0.10]$, under either rational expectations (Panel A) or subjective beliefs (Panel B). Across panels, the figure shows that the assumption of increased risk aversion, within the interval (shaded), modestly increases the optimality of lower-goal choice but reduces the optimality of Goal 3 by an offsetting degree. In the Appendix, we recharacterize goal choice for utility functions that assume constant relative risk aversion (CRRA) across a range of potential wealth. The analysis indicates that the assumption of CRRA utility, for plausible relative risk aversion, yields a characterization of choice nearly identical to CARA benchmarks (Appendix Table A1).

5.3 Empirical Tests of Potential Confounds and Robustness

Conceivably, conservative goal choice could reflect two additional motives consistent with the subjective EU framework—employee preference to avoid high goals in the context of convex effort costs and heterogeneity in employee preferences for risk. We address each potential confound through analyses of field data before revisiting them in the lab. With respect to effort costs and low-goal choice, while we interpreted the empirical setting as abstracting away from considerations of perceived effort given that we elicited beliefs of goal attainment after employees selected their goal, it is possible that employees systematically chose lower goals due to effort motives and that elicited beliefs reflect estimates of attainment conditioned on undesired effort provision rather than our favored interpretation.

We address this possibility through two empirical strategies. First, we characterize optimal choice—under a risk-neutral subjective EU benchmark—across plausible representations of hourly convex effort costs. Adopting a non-parametric approach, we specify a parameter, σ, to represent the costs of the incremental hourly effort required to achieve Goal 2, relative to Goal 1, as a percent of baseline hourly wage. Next, we represent effort cost convexity with a proportional scaler, $k \geq 1$, such that the incremental hourly cost of effort to achieve Goal 3, relative to Goal 2, is $k\sigma$. Appendix Table A2 conveys the improbability of explaining additional employee choice through the assumption of effort costs. Across extremely wide-ranging values of $\sigma \in [0, 1, 3, 5, 10, 25, 50]$ and $k \in [1.0, 1.1, 1.25, 1.5, 2.5, 5.0]$, incorporating effort costs into the benchmark model does not improve predictive accuracy. Intuitively, for effort costs to systematically explain Goal 2 choice, one must assume a baseline increment and convexity scaler within a highly (and perhaps implausibly) narrow interval (otherwise, effort costs would favor either the lowest or highest goal). One must additionally assume highly varying narrow intervals across employees given substantial variation in program length and employee beliefs of attainment. A second related strategy for assessing the explanatory power of effort costs is to observe whether longer programs lead employees to systematically favor low (or conservative) goals, under the assumption that high-goal effort costs accumulate over time. We find no evidence for such correlation.[21]

To evaluate the possibility of conservative choice due to heterogeneous risk preferences, we reassessed the optimality of choice after classifying any goal choice as optimal if it could be predicted by an EU benchmark given *any* value of *r* within the interval, $[0, 0.005]$, under either rational expectations or subjective beliefs. Under this more permissive standard, the share of optimal goal choice increases from 0.44 to 0.56 percent (rational expectations) and from 0.53 to 0.59 percent (subjective beliefs). Allowing for heterogeneous risk preferences also serves to increase the differential share of optimal choice across high and low reward size but not high and low employee experience. We revisit the predictive accuracy of models with heterogeneous risk in subsequent experimental analyses.

Robustness Analysis –Expansive Sample. To assess the robustness of the findings and their potential generalizability, we replicate the preceding analysis for the expansive sample (i.e., the sample inclusive of employees for whom beliefs were not observed). As summarized in Appendix Table A3, relative to a risk neutral benchmark, the characterization of choice in the expansive sample resembles that of the primary sample but for a modestly higher share of conservative choice and smaller share of optimal choice.[22] As with the primary sample, the assumption of moderate to severe risk aversion does not meaningfully improve the share of choice characterized as optimal relative to the EV-benchmark. The analysis suggests that to the extent the primary sample is not fully representative of the employee population, it modestly *underestimates* the share of sub-optimal and conservative choice.

5.4 Behavioral Departures from EU Framework

We proceed to consider whether one can explain conservative choice, and choice more broadly, through two prominent departures from the standard EU framework from the literature beyond non-standard beliefs: non-linear decision weights and gain-loss-utility. Table 5 summarizes the choice characterization under these behavioral benchmark models relative to a baseline benchmark that assumes subjective expected utility with moderate CARA risk aversion ($r = 0.0003$).

Non-Linear Decision Weights. A first behavioral departure we consider is the possibility that conservative choice emerges from the influence of non-linear decision weights. To test this possibility, we replace the linear decision weights with the weighting function of Prelec (1998; $\alpha = \beta = 0.65$). As the table indicates, the modified benchmark does not meaningfully shift characterization relative to baseline.

[21] A regression of an indicator for non-high goal choice on a linear index of program length in days with standard errors clustered at the program level yields a non-significant coefficient estimate, $b = 0.001$ ($p = 0.20$). The analogous regression for conservative goal choice (under a subjective EV benchmark) yields a near identical coefficient estimate, $b = 0.001$ ($p = 0.19$).
[22] To characterize choice under rational expectations in the expansive sample, we adhere to the previously adopted strategy but for excluding unobserved demographic variables in the regression estimates of beliefs.

The insensitivity of characterization to non-linear weights is perhaps unsurprising given that expected distortions from non-linear weighting functions are not typically large for moderate to high beliefs.

Gain-Loss Utility. We next consider the possibility that conservative choice may reflect prospective loss aversion in the context of gain-loss utility. Given an interest in testing all plausible representations of gain-loss utility, we evaluate a large combination of benchmark models reflecting varying candidate reference points, θ, functional scaling factors, η, and loss aversion parameters, λ, as discussed in the theoretical framework. Specifically, we considered five prospect-independent reference points: status quo (i.e., $0), the high probability goal (Goal 1), the high reward goal (Goal 3), the highest goal an employee felt they were certain to achieve as indicated by their survey response (or $0 if they were less than certain for all goals), and, for completeness, Goal 2. We additionally considered prospect-dependent reference points including the chosen goal, the expected value of the chosen goal, and in recognition of models of counterfactual regret, the nearest-goal either below or above the chosen goal. We assessed reference points in the context of composite utility specifications that assume a Kahneman-Tversky power function ($\alpha = 0.88$) for both consumption and gain-loss utility and, given the lack of empirical consensus in the literature, a wide range of scaling factors, $\eta \in [0, 5]$. In deference to the breadth of loss aversion parameters contemplated by the literature, we consider $\lambda = 1.5$, 2.25, and 3.0. Finally, we assume linear decision weights and subjective employee beliefs.

Across all tested benchmarks, the most successful—entailing a reference point set at the chosen goal reward ($\eta = 1$, $\lambda = 2.25$)— explained 59 percent of employee choices. Appendix Table A4, which reports the descriptive accuracy of gain-loss benchmarks, indicates that benchmarks with prospect-independent reference points explain approximately half of goal choices while benchmarks with prospect-dependent reference points produce more variation in descriptive accuracy. Table 5, which reproduces the full characterization of choice for the best-performing gain-loss parameterization, indicates that beyond improving the share of explained goal choice relative to baseline, the gain-loss benchmark does not generate moderation in descriptive accuracy by reward size or tenure. It is possible that the observed behavior reflects employees with gain-loss utility who vary in the severity of their loss aversion. To examine the possibility of heterogeneous loss aversion, we calculated the descriptive accuracy of a gain-loss utility benchmark ($\eta = 1$) that flexibly allows each employee any $\lambda = (1.0, 1.5, 2.0, 2.5, 3.0)$ that could rationalize their goal choice. The exercise yielded an optimal choice share of 0.71.

Collectively, attempts to explain employee goal choice with standard EU benchmark models, or common behavioral departures from such models, persistently failed to explain over 40 percent of employee choices. The most promising explanation we tested involves gain-loss utility with heterogeneous loss-aversion. Even allowing for non-trivial decision noise does not markedly improve the accuracy of standard benchmarks.[23] Employee conservatism resulted in an estimated 40 to 50 percent loss in rewards relative to the counterfactual associated with ex ante optimal choice.

5.5 Gender Differences in Risk-Taking

Given the emphasis on binary gender differences in risk attitudes in the literature (see Niederle, 2017), we report gender differences in goal choice and choice characterization in Table 6. Overall, we find that women selected a low goal—that is, Goals 1 or 2—26 percent more frequently than men (women: 0.63; men: 0.50). Due to similar levels of goal attainment, this choice pattern implies that, under risk-neutral SEU, women chose conservatively 28 percent more frequently than men (women: 0.55; men: 0.43). We can characterize the economic consequences of this difference by estimating the share of the gender gap in realized rewards—women earned 22% less rewards than men ($237 vs $304)—can be statistically attributed to gender differences in conservatism. The exercise suggests that had women chosen as aggressively as men, all else equal, they would have earned only 13% less than men, implying gender differences in risk aversion account for 40% of the gender gap in rewards.

Among explanations for greater risk aversion among women cited by the literature are the possibilities that women are less overconfident than women or that they have weaker utility-based preferences for risk than men. As conveyed in the table, our analysis suggests that the gender difference in conservative choice cannot be explained through differences in biased beliefs. That is, because women in the sample are as overconfident as men, a benchmark model allowing for biased beliefs does not reduce the implied gender gap in conservatism relative to rational expectations.[24] Comparing the conservative choice shares of women, under a risk-averse SEU benchmark with that of men under a risk-neutral SEU benchmark, further indicates that the gender gap in conservatism cannot be explained through systematic differences in risk preferences (similarly, it cannot be explained through differential adoption of non-linear decision weights). The table does, however, suggest the potential for explaining roughly one-half of the gender gap through disproportionate reliance on gain-loss utility. We return to potential explanations for gender differences in conservative choice in subsequent analyses.

[23] It is possible that conservative choice may reflect decision noise that, for unknown reasons, systematically favors low-goal choice. Such decision noise could reflect noisy beliefs—due to either actual uncertainty in such beliefs or uncertainty associated with our elicitation procedure—or computational imprecision. We find that even allowing for non-trivial decision error favoring low goals only slightly improves the accuracy of standard benchmarks. We implement an allowance for noise by evaluating whether the subjective EU model with moderate risk aversion can rationalize choice for any set of subjective beliefs within a +/-10 percent range of the self-reported figure. The allowance of a 20 percent error in subjective beliefs only moderately increases the share of choice deemed to be optimal, from 0.50 to 0.54.

[24] For example, if one defined overconfidence as the average difference in perceived and actual attainment, men and women were identically overconfident with respect to Goal 3 (both 0.32, $p = 0.73$).

6 ANALYSES OF POTENTIAL MECHANISMS

We continue to investigate the motives for conservative goal choice through three online experiments and accompanying analyses. In the first experiment, we re-examine prior benchmarks with greater statistical power, attempt to rule out potential confounds from the field, and assess alternative explanations from the literature including benchmark models with heterogeneous parameters. In a second experiment, we replicate findings from the field and the effort-task in an explicit lottery-choice scenario. In a third experiment, we evaluate a novel, heuristic, explanation of GQ goal choice. We provide additional evidence for the heuristic by assessing its descriptive accuracy in the field.

6.1 Additional Tests of Mechanisms via Online Goal-Reward Paradigm (Experiment A)

Overview and Implementation. We administered Experiment A in May 2019 on the Qualtrics platform to 407 employed US adults recruited from Amazon Mechanical Turk. The online instrument asked participants to complete a brief effort task in the context of an incentive-compatible goal-reward paradigm lasting a few minutes. The paradigm resembled GQ but for the inclusion of features intended to assess potential confounds from the field: Comprehension checks, dollar-denominated rewards, and multiple decisions per subject. We additionally asked multiple decision-relevant questions regarding goal attainment beliefs, risk preferences, and self-assessed relative ability and taste for competition.

We implemented the goal-reward paradigm by first explaining to consenting participants that they were about to participate in a timed effort task where they could earn financial rewards for "solving" a series of grids. To solve a grid, participants had to find the unique pair of numbers whose sum equaled 10 within a 3 x 3 matrix of single-digit numbers. After an opportunity for practice, we formally introduced participants to the all-or-nothing goal-reward paradigm named "GoalQuest" via an online web-flow resembling that used in field. The web-flow explained that participants would have four minutes to solve as many grids as they could and that they could earn a reward if they attained a self-selected performance goal. After questions to test comprehension of the paradigm, participants proceeded to select a goal.

To increase statistical power, we asked participants to select goals from each of six distinct menus, presented in succession. We explained that one menu would be selected at random to determine the participant's actual reward. The menus were strategically designed to facilitate tests of mechanisms. Specifically, a baseline menu resembled real-life GQ with three options, additively linear goals (6 grids, 8 grids, 10 grids) and non-linearly increasing rewards ($0.10, $0.20, $0.35). Four additional menus modified the baseline either by varying overall difficulty or the relative financial attractiveness of the high goal; two additional menus expanded the baseline menu to four choices by adding either a relatively unattractive high- or low-goal option. Following goal selection, we elicited beliefs of attaining several grid thresholds (we used these to estimate goal attainment beliefs across all menus).[25] We also presented participants with iterative hypothetical gambles to estimate a participant-specific loss aversion parameter and assessed self-perceived goal-solving ability and taste for competition relative to similar others.[26] Participants then completed the four-minute effort task so they could earn their reward. The online Appendix summarizes data on choice, beliefs, and attainment for 3-option menus from the experiment.

Baseline Comparison of Lab and Field. After auditing the data for incomplete and/or inconsistent beliefs, 277 remaining participants made 1,662 goal choices. Overall, the baseline menu generated a pattern of goal choice in the lab (Goal 1: 0.34, Goal 2: 0.28, Goal 3: 0.38) resembling that observed in the field (0.29, 0.27, 0.44). We also observed a similar correspondence between lab and field with respect to goal attainment beliefs (lab: 0.80, 0.66, 0.51; field: 0.78, 0.69, 0.63) and choice characterization under the subjective EV benchmark (lab: 0.50 optimal, 0.45 conservative; field: 0.50, 0.48). Lab participants exhibited overconfidence though not as severely as their employee counterparts from the field.

We interpret the similarity in choice, beliefs, and characterization across the lab and field as suggesting the improbability that documented patterns from the field reflect program confusion, reward value confusion, managerial signaling, or reputational concerns alone, as such factors were eliminated or diminished in the lab. Moreover, the latter two confounds, signaling and reputation, would presumably nudge employees towards less—not more—conservative goals. The baseline comparability of lab and field also validates the usefulness of the online paradigm for assessing mechanisms.

Benchmark Characterizations of Goal Choice. Table 7 summarizes multiple measures of optimal choice under a range of benchmark models considered in the field analyses. Beyond reporting the participant share whose full set of choices adhere to the predictions of a particular benchmark, the table also characterizes optimal choice after allowing for some decision error in the form of only five of six adherent choices. As indicated by the table, previously considered benchmarks explain, at most, 18 percent of participant choices, a rate that rises to 31 percent when allowing for decision error. Incorporating a participant-specific estimate of experimental loss aversion in the gain-loss benchmark yields an optimal choice rate of 0.19 (0.30 with decision error; not reported in table).[27] Overall, the table does not offer compelling evidence for any previously considered motive for financial risk-taking.

[25] We constructed a full distribution of beliefs through interpolation and extrapolation when required. For example, for all even values of n between 6 and 18, we estimated the belief of attaining exactly n-1 grids as the difference in beliefs associated with n and n-2 grids. We estimated beliefs of attaining below 4 and above 18 grids from using actual performance data.

[26] To assess confidence about their grid-solving ability, we asked participants to evaluate their ability to solve grids relative to others the study on a five-point scale ranging from well-below average to well-above average. We elicited participants' self-perceived relative taste for competition on a five-point scale from "much less" to "much more" competitive than others.

[27] We also assess the predictive accuracy of a flexible gain-loss benchmark that permits participants any λ = (1.0, 1.5, 2.0, 2.5, 3.0). The flexible benchmark yielded an optimal choice rate of 0.29 (0.44 with decision error), suggesting the high rate in the field may have been due to the low statistical power of the exercise. For roughly 20% of participants, we could not calculate a personalized loss aversion parameter. We ignored such participants from the optimal share rate calculation for that benchmark.

Contextual Sorting Heuristics. We also assess the predictive accuracy of two heuristic choice-strategies involving contextual sorting from the literature that we were unable to test in the field. The first presumes employees heuristically select the goal whose relative position in an ordered-menu corresponds to their perceived standing in some choice-relevant distribution such as ability. This heuristic would be a rational strategy for an employee unsure of what goal to select but who believes the menu was designed by an informed employer to sort employees in roughly equal shares. Contextual sorting of this sort was originally suggested as a potential explanation for uninformed consumer decisions from product menus (Kamenica 2008). A second heuristic reflects the related possibility that participants sort themselves into goals based on their relative taste for competition. The possibility that variation in apparent risk-taking might reflect differences in competitive taste was advanced by Niederle and Vesterlund (2007). As indicated in the table, we find no support for either heuristic.[28]

Benchmark Models with Heterogeneous Parameters. The multiple decisions collected for each experimental participant allows us to revisit flexible models with heterogeneous parameters with greater statistical power than the field. As such we reconsidered a subjective EU model that allows individuals any plausible set of CARA risk preferences, $r \in [0.00, 0.005]$, and a model with gain-loss utility ($\eta = 1$) that allows individuals any loss aversion parameter, $\lambda = (1.0, 1.5, 2.0, 2.5, 3.0)$. While the former provided no improvement in explanatory power, even allowing for decision error in the form of 5 of 6 correct decisions, the latter had more moderate success, explaining choices for 29 percent of participants (44 percent allowing for decision error). Nevertheless, the experiment suggests that the predictive accuracy of heterogeneous loss aversion in the field may overstate its case as a motive for risk-taking.

6.2 Replication of Conservatism with Explicit Financial Lotteries (Experiment B)

As a final test of potential confounds in the field (and to validate a paradigm for subsequent experiments), we administered Experiment B. The experiment was intended to test whether the observed conservatism from the field and Experiment A would persist in the context of hypothetical scenario-based GQ goal choice with explicit attainment likelihoods and choice in a modified scenario from a menu of explicit financial lotteries varying only in risk and reward. Specifically, we randomized 243 US adult participants from Amazon Mechanical Turk in December 2023 to one of two conditions, both of which asked participants to imagine they were new employees at a firm with an employee rewards program. The first condition introduced participants to the real-life GQ paradigm, tested their comprehension of the paradigm, and asked them to select a goal from a stylized GQ menu adapted. The menu featured three

[28] We tested the ability-sorting heuristic by asking participants to assess their grid-solving ability relative to other participants, mapping relative assessments to predicted goal choice by menu position (e.g., high relative ability predicts high-goal choice), and then comparing actual and predicted choices. We used a similar procedure to test the relative taste-for-competition heuristic.

sales goals (105 units, 110 units, 115 units) and dollar-denominated rewards ($150, $450, $900) representative of (percent-denominated) programs from the field.[29] The menu also explicitly conveyed the likelihood of attaining each goal with statistics roughly informed by the field (Goal 1: 83%, Goal 2: 74%, Goal 3: 65%) and strongly implying EV-optimality of Goal 3. The second condition introduced participants to a modified fictional paradigm named "RewardQuest" (hereafter, RQ). Participants were informed that RQ had been adopted to reward for departmental performance and that RQ participation entailed selecting a reward lottery from a menu of three risk-varying options. Each reward lottery was associated with a points threshold from 1 to 100 and participants were told that receipt of their selected reward would be determined by the random spin of an electronic wheel with 100 numbered slots, each corresponding to a point value. After a test of comprehension, participants were asked to select a reward from a menu with reward values and (explicitly displayed) likelihoods identical to the GQ menu.

Overall, experimental response affirmed the robustness of conservative choice to both scenario-based hypothetical GQ contexts and economically equivalent contexts involving financial lotteries. Specifically, as in the field, the experiment revealed a diverse pattern of choice across GQ goals (0.26, 0.44, 0.30) and RQ rewards (0.30, 0.45, 0.25). Such choices implied substantial conservatism relative to EV-optimal choice, with the lottery-based menu producing modestly greater, but not significantly so, conservatism (0.75) than the GQ menu (0.70; $p = 0.41$).

6.3. Pairwise Partition Dependence (PPD) - New Heuristic Explanation for (Conservative) Choice

What might explain (conservative) goal choice in the lab and the field? Given the challenges of explaining observed choice through existing benchmark models, we propose a novel heuristic explanation for financial conservatism informed by our reading of the literature on judgment and decision-making and exploratory pilot studies that asked participants to describe the phenomenology of their choice. The proposed heuristic, which we refer to as the Pairwise Partition Dependence (PPD) heuristic, broadly stipulates that a DM selects an option from a menu of simple lotteries through a succession of approximate, relative, pairwise comparisons between adjacent options. Critically, the heuristic presumes that such pairwise comparisons can lead to a potentially substantial inferential bias due to partitioning of the state space imposed by the pairwise comparison. When applied to the relatively modest differences that distinguish perceived attainment across goals in GQ menu for most employees, the heuristic predicts that employees underestimate the relative likelihood of high-goal attainment and select more conservatively than standard benchmarks would predict.

---

[29] The menu was representative of percent-denominated GQ programs (i.e., those with rewards expressed as a percent of baseline). To generate rewards we applied the modal rewards ratio (1-3-6) to the median Goal 1 reward ($150, after rounding). Goals reflect a 5-10-15 percent increase relative to a baseline of 100, reflecting the mean/median/ modal configuration of percent-denominated program. Average rewards in such programs were higher than the global program average.

Model Setup. We outline the PPD heuristic more formally, in the context of GQ, by returning to our earlier framework where we represent goal choice as a decision from a menu of two simple lotteries ordered from low to high risk (G_l, G_h), where G_j yields reward, $x = X_J$, with some probability and 0 otherwise. The heuristic specifies that the DM pairwise compares options by roughly evaluating whether the expected marginal gain from selecting the high relative to the low goal outweighs the potential loss of the low-goal reward. To carry out the pairwise comparison, the DM partitions potential outcomes into three states demarcated by two decision-relevant risk thresholds. A first threshold, X_L, specifies a value below which the two options similarly yield no payoffs or utility. A second threshold, X_H, specifies a value above which the payoff of the higher goal exceeds that of the lower goal. The two thresholds partition potential outcomes into the following three states: (i) A first, $s_L \equiv x < X_L$, specifies outputs of x for which goal choice is inconsequential, (ii) A second, $s_H \equiv x > X_H$, specifies outputs of x for which the high goal rewards exceeds that of the low goal, and (iii) A third, $s_M \equiv X_L \leq x \leq X_H$, describes outputs of x falling between the two thresholds and for which the low goal reward exceeds the zero payoff of the high goal. We denote a DM's perceived likelihood of state, s_K, by $\widehat{\varphi}_K$.

Pairwise Decision Rule. The heuristic specifies that a risk-neutral DM will select the high goal if the following decision rule is satisfied:

$$\widehat{\varphi}_{L+}\left[\widehat{\varphi}_{H|L+}\Delta x_{h,l} - (1 - \widehat{\varphi}_{H|L+})x_l\right] > |\omega|$$

In this rule, the parameter, $\widehat{\varphi}_{H|L+}$ denotes subjective belief of high-goal attainment conditioned on at least low-goal attainment, $(1 - \widehat{\varphi}_{H|L+})$ denotes subjective belief of the complement, $\Delta x_{h,l}$ is the difference between the high and low-goal reward, x_l is the low-goal reward, and ω denotes some decision error. While the formulation emphasizes the relative nature of the evaluation, we can simplify the rule to: $\widehat{\varphi}_H \Delta x_{h,l} - \widehat{\varphi}_M x_l > |\omega|$. Absent inferential bias, the rule simply restates the utility-maximizing proposition but for the allowance of decision noise.

The heuristic, however, presumes that pairwise comparisons are subject to partition dependence, a phenomenon describing the sensitivity of judgment to the potentially arbitrary partitions imposed by the decision context (Tversky and Koehler, 1994; Fox and Rottenstreich, 2003; Fox and Clemen, 2005). Specifically, the heuristic contends that individuals have partition-dependent beliefs in the context of the three partitions introduced by pairwise comparison. Following Fox and Clemen (2005), we represent partition dependence as a convex combination between actual beliefs and a naïve prior that assigns an equal probability to each partition. Accordingly, we define a parameter, $\theta \in [0, 1]$, as specifying the degree of partition bias ranging from an absence of bias ($\theta = 0$) to full bias ($\theta = 1$) so that: $\widehat{\varphi}_K = \theta/\pi + (1-\theta)\varphi_K$. Here, π indicates the number of decision-relevant partitions, which under pairwise evaluation—or in the case of a menu with only two options—is always equal to three. Under full partition bias, the heuristic implies that the DM will select the high goal if $(\Delta x_{h,l} - x_l) > |\omega|$, assuming renormalization of the noise parameter by $\widehat{\varphi}_H = \widehat{\varphi}_M$ = 1/3.

While we have restricted discussion thus far to a menu of two options, one can straightforwardly adapt the heuristic to larger menus, such as the ones featured in GQ, through one of several strategies. For tractability, we assume that when heuristically engage menus with more than two risk-ordered options, they successively compare proximal pairs of options beginning with the low-risk option and stopping any time a riskier option is rejected. As such, we assume GQ employees initially compare Goals 1 and 2, and either accept Goal 1 or proceed to compare Goals 2 and 3.[30]

To illustrate how the proposed heuristic might help to explain conservative goal choice in GQ, consider the stylized example in which an employee must select between a low goal with a $300 reward and 0.80 attainment likelihood, G_l = ($300; 0.80) and high goal with a $500 reward and 0.60 attainment likelihood, G_h = ($500; 0.60). An EV-maximizing employee would select the high goal since, assuming low-goal attainment, the expected marginal gain from high-goal attainment [$150 = $200 x 0.60/0.80] comfortably exceeds the expected marginal loss [$75 = $300 x (1- 0.60/0.80)]. In contrast, the proposed heuristic will predict low-goal choice for any θ exceeding 0.64 since such a theta marks the indifference point between the potential marginal gain [$120 = $200 x 0.60; $\widehat{\varphi}_{H|L+} = 0.60$] and loss [$120 = $300 x 0.40; $(1 - \widehat{\varphi}_{H|L+}) = 0.40$] associated with high-goal choice.

Figure 5 provides additional graphical intuition by depicting the distortion in perceived goal attainment likelihoods in the context of a pairwise comparison of two goals under heuristic and standard choice. The shaded region between the two CDFs depicts inferential bias under heuristic evaluation. It is worth noting that the heuristic does not inherently bias DMs toward less risk. Instead, the direction of the bias depends on the economic structure of the decision context. For example, because GQ features goals, particularly Goals 2 and 3, that are closely situated with respect to perceived attainment (as depicted in the figure), s_M will be typically small. In such circumstances, the heuristic predicts inflated estimates of not attaining the high goal, encouraging conservative choice. In other contexts, the heuristic may favor greater risk taking, a possibility we discuss subsequently in a discussion of insurance markets.

<u>Motivating Evidence for PPD Heuristic</u>. The PPD heuristic departs from standard models of decision-making through three non-standard assumptions: (i) reliance on pairwise evaluation, (ii) the influence of partition dependence within such evaluations, and (iii) an allowance for computational error. An extensive interdisciplinary literature supports each of these assumptions. First, the tendency for individuals to compare options is a well-documented psychological principle, supported by a wealth of

[30] Practically, in even larger menus, where decision-makers are unlikely to consider all options, we speculate that they may apply the heuristic to a focal subset of options.

interdisciplinary research including work in neuroscience, judgement and decision-making, and economics. This concept is integral to economic theories of reference-dependent preferences (e.g., Kahneman and Tversky 1979; Tversky and Kahneman 1992; Koszegi and Rabin 2007) and figures prominently in models of salience and attention (e.g., Bordalo, Gennaioli, and Shleifer 2012, 2013), and comparative evaluation (e.g., Koszegi and Szeidl 2013; Bushong, Rabin, and Schwartzstein, 2021).

Second, the possibility that judgment may be biased by the categorization imposed by the decision-making context, referred to as partition dependence, was introduced by Fox and Rottenstreich (2003). The theory builds upon the non-extensional support theory of Tversky and Koehler (1994), which posits the sensitivity of subjective beliefs to the granularity with which an event is described.[31] The effects of partition dependence have been documented across diverse contexts, often demonstrating the sub-additivity of probabilistic assessments (see Benjamin, 2019, for discussion). For example, partition dependence, with subadditivity, implies a heightened belief in the likelihood that tomorrow's temperature will exceed some threshold (e.g., 70 degrees) when elicited through specific partitions of the state space (e.g., 70 to 80 degrees, 80 to 90 degrees, etc.), rather than a general forecast free of partitions.

An implication of partition dependence in the context of financial risk is that the framing of options can influence evaluation, and choice, if such framing imposes a particular partitioning of the state space. Despite the strength of descriptive evidence, partition-dependent beliefs could conceivably emerge from several non-standard decision-making processes for which evidence exists. Such processes include the neglect or misunderstanding of contingencies (Martínez-Marquina, Niederle, and Vespa 2019; Sunstein and Zeckhauser 2010; Sunstein, 2002), noisy or imperfect memory (see Hilbert, 2012), or selective attention and informational salience (e.g., Bordalo, Gennaioli, and Shleifer 2012, 2013). Of relevance to economics, Ahn and Ergin (2010) demonstrate how to axiomatically incorporate partition-dependent beliefs into a decision-theoretic expected utility framework. While a straightforward application of partition dependence to GQ might suggest four partitions based on three performance goals (i.e., performance, x, is either less than goal 1, between goals 1 and 2, between goals 2 and 3, or greater than goal 3), our heuristic proposes dependence based on three decision-relevant partitions associated with pairwise comparison. Following Fox and Clemen (2005), we model the bias as a convex combination of genuine beliefs and a naïve prior that evenly distributes probability across partitions. While chosen for its analytic simplicity, this representation is consistent with the possibility that partition-dependent beliefs emerge from insufficient adjustment away from a naïve prior (Fox and Clemen, 2005).

[31] Fox and Rottenstreich (2003) discuss how partition dependence derives from both support theory and the pruning bias (Fischoff et al., 1978).

A third assumption of the proposed heuristic is an allowance for computational error. The allowance of decision noise or error in judgment is a common feature of decision-making models within and outside of economics. Decision noise may arise from sources such as bounded rationality (e.g., Simon, 1955), stochastic preferences (e.g., Loomes and Sugden, 1982), or more recent frameworks that explicitly recognize the role of noise in decision-making (e.g., Kahneman et al., 2021).

6.3. Evidence for the PPD Heuristic

We assess the plausibility of the proposed heuristic through three strategies. First, we present evidence from a new experiment designed to test whether individuals adopt two of the key decision-making precepts underlying the PPD heuristic—reliance on proximal pairwise comparisons and inferential errors in the context of relative comparisons; whether the magnitude of inferential bias, conditioned on partition independent beliefs, predicts goal choice; and whether a debiased menu intended to encourage partition independence increases the optimality of choice relative to standard benchmarks. Second, we revisit the prior experiment to assess the plausibility of partition dependence applied to pairwise comparisons as compared to a more straightforward interpretation. Last, we assess whether the proposed heuristic explains a greater share of employee choice in the field and lab than prior benchmarks.

Overview and Implementation - Experiment C. We administered Experiment C in July 2022 on the Qualtrics platform to 893 employed US adults recruited from Amazon Mechanical Turk. After describing the real-life GQ paradigm, we randomized the 82% of participants successfully completing comprehension checks to one of two experimental arms. Across arms, participants were asked to make a hypothetical decision from a GQ menu featuring sales goals (105 units, 110 units, 115 units) and dollar-denominated rewards ($150, $450, $900) identical to that adopted in Experiment B.

A first arm was designed to test whether participants adopted proximal pairwise comparisons, whether relative comparisons exhibited inferential bias consistent with partition dependence, and whether the degree of inferential bias, conditioned on native beliefs, predicted goal choice. Specifically, we provided participants randomized to the first arm a fictional distribution of sale figures to increase the realism of their goal choice. The distribution of prior sales reflected actual menu-specific averages from the field. After participants in this arm selected their goal, we asked them to introspect as to how they arrived at their goal choice by asking them to indicate which, if any, (two- or three-goal) comparisons they made during their deliberation (e.g., “At some point, I directly compared Goals 1 and 2”). We then asked participants to report their perceived likelihood of attaining each goal through both contingent and non-contingent elicitations, hypothesizing that the former would be subject to partition bias similar to that invoked by pairwise comparisons (e.g., we asked participants to estimate the likelihood of Goal 3 attainment given certain knowledge of attaining Goal 2).[32] Finally, to generate between-subject evidence for pairwise partition dependence, and to test the generalizability of the phenomenon, we elicited contingent and non-contingent weather forecasts across both experimental arms.[33] Across belief elicitations, validation rules ensured internal consistency.

A second arm was designed to test whether we would observe a greater share of EV-optimal choice from menus intended to discourage partition dependence. Specifically, we randomized participants to one of three reframed variations of the same representative menu. The first condition (*baseline*) communicated that a participant's likelihood of goal attainment was 83 (Goal 1), 74 (Goal 2), and 65 (Goal 3) percent (e.g., "You have an 83 percent chance of achieving Goal 1"). A second condition (*partition independent*) displayed the same likelihood for Goal 1 but contingently displayed the likelihood of Goal 3 given Goal 2 attainment implied by the non-contingent elicitation, $\widehat{\varphi}_{H|L+} = \varphi_{H|L+}$ ("If you achieve Goal 1, you have an 89 percent chance of also achieving Goal 2") and for Goal 3 ("If you achieve Goal 2, you have an 88 percent chance of also achieving Goal 3"). We hypothesized that this frame would lead to more accurate, and less partition dependent, inference. Finally, a third condition (*partition dependent*) displayed the non-contingent likelihood for Goal 1 and contingent likelihoods for Goals 2 and 3 modified to reflect bias of the form, $\widehat{\varphi}_{H|L+} = \varphi_{H}$ (i.e., we replaced the 89 and 88 percent from the prior condition with 74 and 65 percent, respectively).[34] We hypothesized that this frame would reaffirm the partition dependence naturally invoked from the baseline condition.

<u>Results - Experiment C</u>. The experiment yielded evidence consistent with widespread adoption of the proposed heuristic. First, regarding the process assumptions underlying the heuristic—86% of participants reported using pairwise comparisons in their deliberations; 93% of such participants made at least one proximal comparison. More critically, participants substantially underestimated conditional likelihoods compared to the Bayesian likelihood implied by non-contingent elicitations. Specifically, participants underestimated the likelihood of Goal 2 | Goal 1 attainment by 23% (0.64 relative to 0.82) and underestimated the likelihood of Goal 2 | Goal 1 attainment by 22% (0.59 relative to 0.76). Returning to our representation of bias, the pairwise comparisons imply θ = 1.01 (comparison of Goals 1 and 2) and

[32] Due to the hypothesized difficulty of credibly eliciting contingent expectations we piloted different communication strategies before arriving at the implementation used in the experiment. For example, to elicit contingent beliefs of weather we asked: "Suppose that you have a time-travelling friend who travels into the future. The friend returns and truthfully tells you that tomorrow's high temperature will be <u>at least</u> 80°F. Knowing for certain that the high temperature tomorrow will be <u>at least</u> 80°F, what are the chances that tomorrow's high will be <u>at least</u> 90°F?"
[33] We randomized participants to either forecast the likelihood that tomorrow's high temperature would be at least 70, 80, and 90 degrees Fahrenheit or to forecast the conditional likelihood of at least 90 degrees given certain knowledge of at least 80 degrees.
[34] The experiment was originally designed to test partition dependence defined as the neglect of the non-focal contingency of low-goal attainment ($\widehat{\varphi}_{L} = 0$ or $\widehat{\varphi}_{L+} = 1$). Defining bias as a convex combination between non-contingent beliefs and a naïve prior, the contingent probabilities displayed in the biased contingent condition imply a θ = .44 (comparison of Goals 1 and 2) and θ = .63 (comparison of Goals 2 and 3).

$\theta = 0.81$ (comparison of Goals 2 and 3). The between-subject estimates of weather implied a more severe, 38%, underestimation of contingent likelihoods. Second, we found that the absolute magnitude of the bias implied by within-subject estimates of contingent and non-contingent goal attainment strongly predicted optimal goal choice—even after controlling for non-contingent beliefs of attainment. Specifically, we estimated an additively linear model of each participant's EV-optimal goal choice, g^*, as a function of non-contingent beliefs of goal *k* attainment, $\hat{s}_k$, and the absolute bias in relative inference implied by contingent beliefs, $\hat{\lambda}_{k,k-1} = |\hat{s}_k/\hat{s}_{k-1} - \hat{s}_{k|k-1}|$:

$$g^* = \alpha + \pi_1\hat{s}_1 + \pi_2\hat{s}_2 + \pi_3\hat{s}_3 + \gamma_1\hat{\lambda}_{3,2} + \gamma_2\hat{\lambda}_{2,1} + \varepsilon$$

The regression (excluding six observations with non-unique optima) indicates that while the perceived likelihood of attaining Goal 3 (the EV-optimal goal for most participants based on their non-contingent beliefs) strongly, and expectedly, predicts EV-optimal choice ($\hat{\pi}_3 = 1.00$, $p < 0.001$), the magnitude of inferential bias strongly (negatively) predicts choice ($\hat{\gamma}_1 = -0.81$, $p < 0.001$). The estimates imply that eliminating the bias associated with relative beliefs of Goal 3 attainment would yield a 37% increase in the share of optimal choice (from 0.37 to 0.51). The substantial partial correlation between inferential bias and optimal goal choice is robust to a variety of alternative non-parametric specifications.

Lastly, we document a marked increase in EV-optimal choice from menus designed to discourage partition dependence. Specifically, 61 percent of participants in the *partition independent* condition selected the EV-optimal goal, a 48% increase relative to the *baseline* condition (from 0.41 to 0.61, p = 0.002). Beyond the increase in optimal choice, participants in the *partition independent* condition chose the EV-minimizing Goal 1 with significantly less frequency than those in the *baseline* condition. As further evidence for the proposed mechanism, the 39 percent optimal choice share of participants in the *partition dependent* condition, with biased display of contingent probability, was not statistically distinguishable from either the *baseline* menu (p = 0.82) or the menu from the first experimental arm with no explicit display of attainment likelihoods (p = 0.55).

<u>Pairwise versus Global Partition Bias - Experiment A</u>. As noted earlier, a more straightforward interpretation of partition dependence would be to assume inferential bias with respect to the partitions demarcated by each goal on the menu rather than to assume partitions associated with pairwise comparisons.[35] For additional insight into these competing interpretations we appeal to Experiment A, which elicited choice from menus varying in size. Experiment A offers little support for a global implementation of partition dependence. Assuming $\theta = 0.75$, a global implementation of partition dependence, accurately predicts the full set of decisions for 0 percent of participants and 5 of 6 decisions

[35] We interpret a non-pairwise implementation of partition bias as presuming n+1 partitions, if n is the number of choices.

for only 10 percent of participants. As we discuss in the next section, these figures signal a rate of descriptive accuracy far less than a pairwise implementation of the bias.

6.4. Descriptive Accuracy of the PPD Heuristic – Field and Lab

Perhaps the most instructive test of the pairwise heuristic is to gauge the accuracy with which it explains observed behavior. To assess the explanatory power of the heuristic in the field and lab, and to further evaluate the importance of each constituent assumption of the heuristic, Table 8 compares the share of accurately predicted heuristic choice, across a range of bias specifications and noise allowances, to a subjective EV baseline. Specifically, we consider heuristic choice with no inferential bias, with personalized bias inferred from the misalignment between self-reported conditional and non-conditional probabilities (Experiment A only), and a parameterized bias assuming $\theta = 0.75$ (a slightly more conservative degree of bias than suggested by the Goal 3 versus Goal 2 comparison in Experiment A). We additionally consider a range of noise allowances intended to subsume plausible levels of decision imprecision (i.e., \$0 to +/- \$50 for the field and Experiment C; \$0 to +/- \$0.02 for Experiment A).[36]

Across domains, the table documents substantially higher rates of predictive accuracy under the heuristic than the baseline, or previously tested, benchmarks. The increase in accuracy appears attributable both to the assumption of inferential bias and the inclusion of modest decision noise. For example, in the field, assuming a noise allowance of either +/- \$25 or +/- \$50, the parameterized heuristic explains 83 to 92 percent of employee choice, a 66 to 84 percent improvement over baseline. Analyses of Experiment C, where we can compare predictive accuracy of the heuristic under the personalized (59 to 72 percent) and parameterized (49 to 63 percent) bias under a moderate noise allowance, suggests that the parameterized estimates from the field may understate the true efficacy of the heuristic, subject to potential ceiling effects. Finally, analyses of Experiment A, where we are able to observe multiple decisions per participant, indicates a far more substantial degree of relative improvement in predictive accuracy under heuristic choice relative to very low baseline rates of accuracy.

<u>Revisiting Gender Differences in Risk-Taking</u>. The PPD heuristic offers an intriguing potential explanation for the previously discussed gender difference in conservatism. Returning to Table 6, the heuristic ($\theta = 0.75$; $\omega = [-25, 25]$) implies a gender difference in conservatism of 0.04—64 to 71 percent smaller in absolute terms than the gender gaps implied by EU benchmarks (0.11 to 0.14). While the diminished gender gap reduction is less pronounced when considered on a relative basis, this owes to the substantially lower baseline rates of conservative choice implied by the heuristic. We appeal to

---

[36] To provide context, a noise of allowance of \$25 in the context of the representative menu from the experiment is equivalent to 10 percent of the average difference in expected value between goals 2 and 3 and 12 percent of the average difference between goals 1 and 2. In terms of wage-based time-use, a \$25 noise allowance is equivalent to roughly one hour of wages. The noise allowances for Experiment A roughly correspond to the allowances for the field on a proportional basis.

Experiment C for additional insight as to how heuristic choice might help explain the gender divide in conservative choice. The experiment reveals small, insignificant, gender differences in the magnitude of inferential bias associated with either goal attainment or weather. This implies that the reduced gender gap, under heuristic choice, may be due to women adopting the heuristic more frequently than men. The evidence highlights the possibility that gender differences in financial risk-taking, or differences across other population sub-groups, may partially reflect heterogeneity in decision strategies, such as adoption of the PPD heuristic, rather than systematic differences in risk preferences and/or perceptions.

6.5. Explaining Residual Goal Choice – Narrow Application of Heuristic

We interpret evidence from the experiment and descriptive analyses as suggesting that a moderate to large share of DMs relied on a decision strategy resembling the proposed heuristic for goal choice. A simple interpretation of the lower bound of heuristic adoption is the difference in the share of explained choices between the standard benchmark and the heuristic. This interpretation, using field data, yields a lower bound of adoption between 33 and 42 percent (i.e., assuming noise allowances of either \$25 or \$50, the difference between 50 and 83 and 50 and 92 percent from Table 8). The table implies far higher upper bounds of adoption given that many employee choices consistent with standard benchmark predictions are also consistent with heuristic predictions. While we cannot precisely estimate the rate of adoption, it is notable that the heuristic explains a far higher share of choice also explained by the subjective EV benchmark than the converse. For example, in the field, given noise of +/- \$25, the heuristic explains 96 percent of choices also explained by the subjective EV benchmark and 69 percent of choices unexplained by the standard benchmark. Conversely, the baseline benchmark explains only 58 percent of choices also explained by the heuristic and 11 percent of choices unexplained by the heuristic.

A significant share of choice does not appear to adhere to any of the benchmark models we consider, including the proposed heuristic. While some decisions likely reflect confusion, inattention, or an otherwise random processes, we speculate that many employees may have narrowly applied a decision rule to a subset of the menu. For example, a GQ employee may have narrowly evaluated Goals 1 and 2 only, having ruled out Goal 3 for indeterminate reasons. Such narrow focus seems reasonable in the context of larger ordered menus where evaluating all options might be prohibitively effortful or involve consideration of clearly undesirable options. Appealing once more to Experiment C, we find that roughly one-quarter of participants whose choices were not consistent with the heuristic reported evaluating only Goals 1 and 2 (an indication of process consistent with their final goal choice). Applying heuristic choice to only the first two goals explained all but one of these decisions.

7 APPLYING PAIRWISE HEURISTIC TO INSURANCE PLAN CHOICE

We speculate that the proposed heuristic could offer insights into financial risk-taking extending beyond employee decisions in reward programs.[37] Specifically, we see the heuristic as applicable to economic contexts where DMs engage a menu of risky options that can be conceptualized as a choice between "nested" financial lotteries. We define a set of lotteries as "nested" if the lotteries draw from the same distribution of risk and uncertainty (e.g., an employee's productivity next month), if the lotteries are identical but for differences in risk and payoff (e.g., GQ goals varying only in their risk and reward), and if positive potential outcomes of a less-risky option entirely subsume that of a more-risky alternative (e.g., employee performance satisfying GQ Goal 3 will satisfy Goal 2). One can reasonably interpret economic menus involving contingent labor contracts, betting or gambling, (retirement) portfolio allocation, options investing, and insurance plan choice as constituting a choice between nested lotteries. While a comprehensive discussion of each of these domains is beyond the present scope, we briefly discuss the particularly promising application of the heuristic to plan choice in consumer insurance markets.[38]

7.1. PPD Heuristic and Insurance Plan Choice

Insurance plan choice offers a natural analogue to goal choice in GQ in that both can usually be approximated as a decision from a menu of nested financial lotteries (in the case of insurance, the analogy requires plans vary only in their cost and single index of cost-sharing; one distinction, of course, is that insurance plans vary in their non-contingent costs, or annual premia, whereas participation is costless for GQ employees).[39] We explore the applicability of the heuristic to plan choice by outlining a theoretical framework of plan choice under the heuristic, deriving the market conditions under which the heuristic predicts systematic biases (in either direction) in demand, positive or negative biases in demand, and then describing a series of experiments that test whether heuristic choice can help to resolve two seemingly contradictory puzzles in the empirical literature on insurance demand.

Model Notation and Setup. We begin by considering the stylized decision of a risk-neutral, utility-maximizing, DM required to purchase insurance to protect against the prospect of random loss, $x \geq 0$.[40] The DM selects from a menu of two plans, j ϵ $[l, h]$, varying only in annual premia, $p_h > p_l$, and cost-

---

[37] Employee reward programs such as GQ are arguably of independent economic interest given their popularity. GQ estimates that roughly 40 percent of Fortune 500 firms have adopted their program since its 2001 inception.

[38] As another example, consider an employee actively allocating their retirement savings via an ordered menu of fund options varying in risk (e.g., from fixed income products to small cap/growth equity). Heuristic choice could lead to inefficient conservatism if investors, in the context of pairwise comparisons, systematically underestimated the relative likelihood of more, versus less, positive market returns due to partition dependence.

[39] In practice, we believe that the analogy holds so long one can price non-financial price differences across plans and cost-sharing is largely achieved through a single channel such as a deductible or coinsurance.

[40] While the assumption of risk-neutrality departs from typical treatment of insurance choice, in contexts where plans differ sufficiently in their expected value, the assumption simplifies analysis without sacrificing generality.

sharing, $b_h(x) \geq b_l(x)$, where $b_j(.)$ is a function that specifies an indemnity payoff for a given loss. Both plans cover losses beyond some out-of-pocket maximum such that plan value, or utility, can be described by $u_j(x) = b(x) - x - p_j$. For any potential state of realized loss, *s*, we denote the difference between the high and low coverage plans in expected utility by Δu_s, indemnity payoff by Δb_s, and premia by Δp.

As before, we stipulate that the DM evaluates the plans via pairwise comparison. Such relative comparison causes the DM to partition potential loss outcomes into three decision-relevant states. A first state, $s_L \equiv x < X_L$, specifies realizations of loss for which neither plan provides coverage. As an example, given two plans whose cost-sharing takes the exclusive form of a deductible, $D_h < D_l$, the first partition threshold would be the deductible of the high coverage plan, $X_L = D_h$, since losses below this threshold would yield no indemnity payoffs for either plan (losses immediately above the threshold would yield payoffs for only the high coverage plan). As another example, for plans whose differential cost-sharing is actualized exclusively through different coinsurance rates, $X_L = 0$. A second state, $s_H \equiv x > X_H$, specifies realizations of loss for which the high coverage plan provides sufficient additional coverage, relative to the low coverage plan, to warrant its higher cost (i.e., $\Delta b_h > \Delta p$ or $\Delta u_h > 0$). Finally, a third state, $s_M \equiv X_L \leq x \leq X_H$, describes realizations of loss falling between the two thresholds.

<u>Pairwise Decision Rule</u>. A risk-neutral, utility-maximizing, DM, will enroll in the high coverage plan if: $\widehat{\varphi}_{L+}[\widehat{\varphi}_{H|L+}\Delta b_H + \widehat{\varphi}_{M|L+}\Delta b_M] > \Delta p$. While the expression reflects the presumed relativity of the comparison, it simplifies to: $\widehat{\varphi}_H \Delta b_H + \widehat{\varphi}_M \Delta b_M > \Delta p$. The rule implies that a DM will opt for greater coverage if the expected gain from such coverage in the event of severe, Δb_H, or moderate, Δb_M, loss outweighs the difference in plan costs, Δp. Absent inferential bias, $\widehat{\varphi}_K = \varphi_K$, and decision noise, this rule predicts EV-maximizing choice. Given heuristic choice with pairwise partition dependence, $\widehat{\varphi}_K = \theta/\pi + (1-\theta)\varphi_K$ and $\pi = 3$, plan choice depends not only on plan differences in cost and cost-sharing, and DM expectations of loss risk, but also the degree of bias, θ and the economic structure of the market, roughly proxied by $\varphi_H + \varphi_M$ (we ignore decision noise going forward for simplicity).

As discussed in the context of GQ, one can adapt the heuristic to larger, ordered, menus by specifying an initial pairwise comparison and a stopping rule. For example, one could assume that consumers successively compare proximal pairs of plans beginning with the low-coverage option and stopping their evaluation once higher coverage is rejected. In large menus where consideration of all options is impractical, it is reasonable to expect DMs to apply heuristic choice to a focal subset of options.

<u>Direction of Bias in Insurance Demand</u>. To better understand the direction and magnitude of bias implied by heuristic choice and the sensitivity of the bias across insurance markets and menu configurations, we can derive net bias as a function of model parameters. After differencing decision rules with and without bias, rearranging terms, and denoting bias with respect to φ_K as $\tau_K = \widehat{\varphi}_K - \varphi_K$, we arrive at the following expression for the increased marginal willingness to pay for increased plan

coverage (in terms of Δp) relative to EV-maximization: $\tau_H \Delta b_H + \tau_M \Delta b_M$. When positive, the heuristic predicts positive bias in demand; when negative, the heuristic predicts negative bias in demand.

Figure 6 provides graphical intuition for the expression by plotting demand bias under heuristic choice for varying market structures, menu configurations, and bias severity. The first panel shows how increasing bias severity amplifies the magnitude of demand bias for a given market structure and how perceptions of loss risk determine the direction of the bias. That is, the figure conveys the propensity towards excessively high demand, relative to standard benchmarks, under heuristic choice in markets characterized by relatively low perceived loss risk (e.g., home, vehicular, or health insurance) and excessively low demand in markets characterized by high perceived loss risk (e.g., dental insurance, vision care, or prescription drug coverage).[41] The second panel of the figure conveys how menu configuration can interact with market structure to shape demand bias. For example, given insurance menus in which plan options are closely situated with respect to cost sharing—and for which $\widehat{\varphi}_M / \widehat{\varphi}_H$ is relatively low—heuristic choice is more likely to predict a bias towards under-insurance than menus in which plan options are more distant in cost-sharing.

7.2. Experimental Evidence on Heuristic Insurance Choice across Insurance Markets

A large literature in economics has investigated the motives for consumer insurance demand across a variety of markets and has documented repeated instances in which utility-based risk preferences cannot account for the level and/or heterogeneity of demand (Barseghyan et al. 2018). We conclude with a series of experiments designed to test whether heuristic choice can help to explain two specific empirical puzzles, with opposite valence, from this literature. The first puzzle involves the propensity of many high-risk elderly Medicare Part D enrollees to undersubscribe, relative to standard benchmark predictions, high-coverage plans (e.g., Abaluck and Gruber, 2011; Heiss et al., 2013). Such a situation corresponds to an insurance market with a high baseline likelihood of at least some covered loss, practically implying a high $\widehat{\varphi}_M + \widehat{\varphi}_H$ for most pairwise comparisons. In this scenario, the heuristic predicts under-insurance, relative to standard benchmarks, due to systematic underestimation of high loss risk. The second puzzle involves the propensity to over-insure in the market for home insurance—a market typically characterized by low baseline loss risk (e.g., low-deductible enrollees had an average claim rate of less than 5% in the sample examined by Sydnor, 2010). In this scenario, the heuristic predicts excess insurance demand due to systematic overestimation of moderate and high loss risk.

[41]A potential exception to this taxonomy is catastrophic insurance, a setting long characterized by sup-optimally low insurance demand (e.g., Kunreuther 1973). While the literature has sought to explain low demand through non-standard decision motives and market failures, many of these explanations involve consumers who do not actively engage plan choice. We speculate that in catastrophic care settings with mandatory coverage requirements, the heuristic would lead to excess insurance demand.

Prescription Drug Coverage. A first experiment (Experiment C) investigated whether heuristic choice can help explain sub-optimally low demand for prescription drug coverage by examining the sensitivity of hypothetical plan choice to variation in the framing of loss risk. Specifically, we asked 432 US adults, recruited from Amazon Mechanical Turk, to imagine they were about to enter retirement and had to decide whether or not to purchase prescription drug insurance (participants were informed they had separate coverage for non-prescription medical expenses). After an educational module explained how drug bills mapped to out-of-pocket costs under various cost-sharing scenarios, participants were asked to select from a menu of two plan options (Silver, Gold) varying only in annual premia ($640, $1220) and cost-sharing (coinsurance rates: 50%, 15%). They also had the option of selecting no plan. The menu resembled actual Medicare Part D menus in that plans varied in cost-sharing primarily via coinsurance, covered all expenses beyond a fixed out-of-pocket threshold ($7,500), and had medal-color labels.

To facilitate plan choice, we provided participants with plausible information on the distribution of prospective drug bills costs for an actual high-risk Medicare Part D enrollee. We communicated the likelihood that drug bills would exceed three decision-relevant thresholds ($0, $1,280, $1,657), truthfully explaining that if drug bills exceeded the first threshold, the Silver or Gold Plan would minimize total costs and if drug bills exceeded the second threshold, the Gold Plan would minimize costs. For tractability, we additionally conveyed that drug bills would never exceed $10,000 in a year—even for those selecting no plan—and that they would follow a uniform distribution between conveyed thresholds. We assigned participants to one of two sets of likelihood thresholds (80%, 60%, 40%; 80%, 63%, 48%) deemed plausible and roughly resembling thresholds from an earlier experiment. Given the provided information, the Gold Plan comfortably minimized expected total costs (premium + out-of-pocket).

While all participants engaged the same set of plans, they were randomized to one of three experimentally varying menus. The first menu displayed cost information in non-contingent terms (*baseline*) (e.g., "You have a 60% chance of a drug bill exceeding $1,280"). The second menu displayed contingent costs after adjustment to reflect moderate to severe partition dependence with respect to the comparison between the Gold and Silver Plans (*partition dependent*) (e.g., "If your drug bill exceeds $1,280, you have a 40% chance of a drug bill exceeding $1,657").[42] The third menu, intended to discourage partition dependence, displayed accurate costs but in conditional terms (*partition independent*) (e.g., "If your drug bill exceeds $1,280, you have a 67% chance of a drug bill exceeding $1,657").

[42] The probabilities displayed in the partition dependent condition reflect an earlier implementation of partition dependence that assumed full neglect of the non-focal, low-state, partition, s_L. In the context of the present model, the displayed probabilities correspond to $\theta = 0.64$ (No Plan compared to Silver Plan) and $\theta = 1.71$ (Silver Plan compared to Gold Plan) for the first set of likelihoods (80%, 60%, 40%) and, analogously, $\theta = 0.59$ and $\theta = 1.08$ for the second set of likelihoods (80%, 63%, 48%).

The experiment, summarized in Appendix Table A5, produced two patterns of note. First, consistent with the literature, baseline participants exhibited substantial demand for non-EV-optimal plans, with only 30% selecting the cost-minimizing Gold Plan. The baseline share of optimal plan choice was nearly identical to that produced by the *partition dependent* menu which displayed contingent, but biased, cost information (29%; p = 0.82).[43] Second, relative to baseline, participants chose far more optimally from the *partition independent* menu (b = 0.12; p = 0.03). Said differently, the menu intended to discourage partition dependence led to 39% more EV-optimal choice than a menu with reframed, but otherwise equivalent, risk information. As with actual Medicare Part D, the economic consequences of inefficient plan choice were significant—those selecting the Silver Plan in lieu of the Gold Plan could have saved an estimated $452 each year in expectation, equivalent to 37% of the Gold Plan premium.

<u>Home Insurance</u>. A second experiment (Experiment E) investigated whether the proposed heuristic could help explain the over-insurance characterizing analyses of the US home insurance market. We implemented the test by comparing hypothetical plan choice across menus strategically varying the presence and salience of potential loss information. We hypothesized that increasing the salience of the likelihood of no loss would offset partition dependence, thereby increasing the efficiency of plan choice. Specifically, we asked 435 US adults, recruited from Amazon Mechanical Turk, to select lender-mandated insurance from a menu of three plan options adapted from Sydnor (2010). The options varied only in annual premium and deductible: (1) Basic Plan ($1,000 deductible, $616 premium) (2) Medium Plan ($500 deductible, $716 premium), (3) Premium Plan: ($250 deductible, $803 premium).

We randomized participants to plan choice from one of four experimentally varying menus. An initial baseline condition (*full information baseline*) accurately conveyed potential loss information roughly adapted from Sydnor (2010): a 4% likelihood of any loss; given some loss, a 75% likelihood of severe loss ($2500+) and a 25% likelihood of non-severe loss ($1 to $2500). After simplifying cost assumptions, the specified potential cost distribution implied the EV-optimality of the Basic Plan.[44] A second condition (*no information baseline*) conveyed no information regarding potential costs. A third condition (*partition dependent*), designed to encourage partition dependence, displayed the 75% relative likelihood of a severe versus non-severe loss but did not display the likelihood of any loss. A fourth condition (*partition independent*), designed to discourage partition dependence, conveyed the 96% likelihood of no loss—a framing deemed to be more salient than the 4% likelihood of any loss—without further displaying the relative likelihood of severe versus non-severe loss. Despite the variation of actual

[43] Treatment effect estimates generated from LPM regressions of the specified outcome on treatment and menu indicators.
[44] Expected cost calculations assume a 3% likelihood of severe losses with $2500 in damages and a 1% likelihood of non-severe losses with $500 in damages. The rank order of plans by expected value—Basic ($626), Medium ($707), Premium: ($781)—is not sensitive to the assumed distribution of costs within the specified range.

risk information across conditions, one can treat the *full information* baseline and *partition independent* conditions as economically equivalent since no partitioning of severe and non-severe loss, given the stated likelihood of any loss, could rationalize EV-optimal demand for anything but the Basic Plan.[45]

The experiment, summarized in Appendix Table A6, once again, corroborates findings from the empirical literature and predictions of the heuristic. First, across both baseline conditions, most participants chose a non-EV-maximizing plan, paralleling the over-insurance observed in the actual home insurance market. Second, when encouraged to actively attend to the high no-loss likelihood, participants demanded substantially less insurance than in either baseline condition—i.e., participants in the *partition independent* condition were 35 (54) percent more likely to select the EV-maximizing plan and 58 (69) percent less likely to select the EV-minimizing plan compared to the *no information* (*full-information*) *baseline*. Third, participants in the *partition dependent* condition, emphasizing the salience of the varying degrees of loss, demanded more coverage than any other condition, consistent with the hypothesized overestimation of baseline loss-risk. To rationalize plan choice under EV-maximization, the 77 percent of participants who did not select the Basic Plan would have had to assume an implausibly high annual loss rate of more than 27 percent. The possibility that these choice patterns simply reflect extremely inflated perceptions of baseline risk is inconsistent with the *lower* demand for insurance observed in the *no information* versus *full information* baseline conditions.

While Experiment E offers evidence consistent with heuristic predictions for a specific, low, level of baseline loss risk, the proposed heuristic also predicts diminished biases towards over-insurance across markets with increasing baseline risk (see Figure 6). We sought to test this prediction through a final experiment (Experiment F). Specifically, the experiment asked 348 US adults, recruited from Amazon Mechanical Turk, to select a hypothetical home insurance plan from the same plan menu as the prior experiment. To manipulate the likelihood of baseline loss, without drawing explicit attention to a specific numerical figure or changing markets, we randomized participants to one of three conditions varying the stated duration of plan coverage from one year, five years, and ten years (i.e., we explained that "The plan is different than most in that it provides coverage for the next [1, 5, 10] years [bolded] with a single deductible and one-time premium."). Within each plan duration condition, participants were cross-randomized to either a *baseline* condition displaying no loss information or a *partition independent* condition explicitly displaying the cumulative likelihood of no loss (e.g., "To help you make a decision, suppose that during your plan period, there's a roughly 77 percent [bolded] chance of no damage [underlined] to your home."), resulting in a 3 (baseline risk) x 2 (risk salience) between-subject design.

[45] While improbably high perceptions of baseline risk could rationalize demand for costlier plans in the no information baseline and focal contingency conditions, in conditions conveying baseline risk, the most pessimistic beliefs about relative severity can only rationalize the choice of the Basic Plan by an EV-maximizing consumer.

Assuming an annual loss likelihood of 5 percent and the specific cost assumptions from earlier (not communicated in this experiment), we estimated total expected costs would be minimized by the Basic Plan given a one year duration (5% loss risk), the Medium Plan given a five year duration (23% loss risk), and the Premium Plan given a ten year duration (40% loss risk).

As depicted in Figure 7, the experiment yielded choice patterns consistent with heuristic predictions. First, participants demanded more insurance, relative to standard predictions, with a loss risk of 0.05 compared to a loss risk of 0.23 ($p < 0.001$) (over-insurance was not possible with a loss risk of 0.40, as the high-coverage plan was EV-maximizing). Second, participants in the *partition independent* condition over-insured at a markedly lower rate than those in the *baseline* condition ($p = 0.006$). Finally, while the interaction was not statistically significant, *partition independent* participants were more likely to choose optimally, relative to baseline, in the lowest loss-risk setting and less likely to choose optimally in the highest loss-risk setting. Collectively, these patterns, while consistent with heuristic choice, cannot readily be explained through random plan choice, plausible utility-based risk preferences, or simple overestimation of baseline risk.

Beyond suggesting the importance of inferential error in insurance evaluation, the experiments allude to multiple violations of descriptive invariance, an axiomatic assumption of standard economic theory. They also demonstrate how heuristic choice could produce substantial heterogeneity in insurance demand, even in the context of plausible risk preferences and well-calibrated beliefs, due to the sensitivity of demand to heuristic adoption and bias severity.

8 CONCLUSION

We describe new evidence on the prevalence of financial risk-taking and its underlying motives. Our main evidence describes the decisions of several thousand employees in the context of a popular all-or-nothing goal-reward program. In contrast to other field settings, we see this setting as uniquely helpful for understanding risky choice given its simplicity, the diversity of employees and financial stakes, the high compliance and near-complete participation rate, and our capacity to observe contemporaneous beliefs. A central finding is to document substantial risk aversion—and heterogeneity—in the goal choices of employees, relative to a standard EU benchmark with plausible risk preferences, resulting in an average unrealized gain equivalent to 46% of potential rewards. The excess conservativism of employees was robust to reward size and employee tenure and was substantially higher for women than men. In investigating the motives for conservative choice, we found that it cannot be explained via prominent non-standard explanations (e.g., biased beliefs, non-linear decision weights, or gain-loss utility). We corroborated patterns of conservative choice and the limited explanatory power of previously tested benchmarks via an incentive-compatible, online goal-reward paradigm with verified comprehension, dollar-denominated rewards, and limited scope for signaling, reputation, or effort costs.

We propose a novel heuristic explanation for the observed patterns grounded in several conjectures from the literature. The heuristic stipulates that employees select goals through a series of approximate, proximal, pairwise comparisons. The heuristic posits that partition-dependent beliefs associated with pairwise comparison led many employees to substantially underestimate the relative likelihood of high-risk outcomes, thereby favoring conservative choice relative to standard benchmarks. We administered a second experiment that provided mechanistic evidence for the heuristic, revealed a significant correlation between the magnitude of inferential bias and optimal choice, and documented substantial improvement in choice efficiency from menus designed to discourage partition bias. Perhaps most compellingly, the heuristic explained a greater share of choice in the lab and in the field than any other tested benchmark; it also accounted for the near-entirety of the gender gap in conservatism, suggesting that gender differences in risk taking may reflect differences in decision strategy rather than differences in native risk preferences or beliefs.

Beyond providing new evidence on prevalence and motives, we speculate that the proposed heuristic may help understand a broad range of risk-taking in contexts where choice can be conceptualized as an ordered set of nested financial lotteries. Towards this end, we apply the heuristic to plan choice in consumer insurance and illustrate how it predicts, potentially substantial, bias in demand of a direction and magnitude shaped by structural features of the insurance market. A series of experiments implies that the heuristic offers a potentially unifying explanation for seemingly contradictory puzzles from empirical analyses of plan choice in Medicare Part D and home insurance markets. The heuristic also offers an explanation—variation in heuristic adoption or bias severity—for the ostensibly excessive heterogeneity routinely documented in analyses of insurance demand (e.g., Cutler and Zeckhauser, 2004; Cohen and Einav, 2007; Barseghyan et al. 2013). We hope future work will clarify the specific cognitive processes underlying the proposed heuristic, explore its applicability to other risk-taking domains, and work out its implications for welfare analyses in markets such as insurance.

9 REFERENCES

Abaluck, J. & J. Gruber. (2011). Choice inconsistencies among the elderly: Evidence from plan choice in the Medicare Part D program. *American Economic Review*, 101(4): 1180–1210.

Ahn, D., & H. Ergin. (2010). Framing Contingencies. *Econometrica*, 78(2): 655–695.

Baillon, A., Bleichrodt, H., & V. Spinu. (2020). Searching for the Reference Point. *Management Science*, 66(1): 93–112.

Bandiera, O., Parekh, N., Petrongolo, B. & Rao, M. (2022), Men are from Mars, and Women Too: A Bayesian Meta-analysis of Overconfidence Experiments. *Economica*, 89: S38-S70.

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

Barseghyan, L., F. Molinari, T. O'Donoghue, & J. C. Teitelbaum. (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review*, 103(6): 2499–2529.

Barseghyan, L., F. Molinari, T. O'Donoghue, & J. C. Teitelbaum. (2018). Estimating risk preferences in the field. *Journal of Economic Literature*, 56(2): 501–64.

Barseghyan, L., J. Prince, & J. C. Teitelbaum. (2011). Are risk preferences stable across contexts? Evidence from insurance data. *American Economic Review*, 101(2): 591–631.

Benartzi, S., & R. Thaler. (2007). Heuristics and Biases in Retirement Savings Behavior (2007). *Journal of Economic Perspectives*, 21(3): 81–104.

Benjamin, Daniel J. (2019). Chapter 2 - Errors in probabilistic reasoning and judgment biases. In B. Douglas Bernheim, Stefano DellaVigna, David Laibson (eds.), *Handbook of Behavioral Economics: Applications and Foundations 1, Volume 2* (pp. 69-186). North-Holland.

Bhargava, S., G. Loewenstein, & J. Sydnor. (2017). Choose to lose: Health plan choices from a menu with dominated option. *The Quarterly Journal of Economics*, 132(3): 1319–1372.

Bordalo, P., N. Gennaioli, & A. Shleifer. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3): 1243–1285.

Bordalo, P., Gennaioli, N. & Shleifer, A., 2013. Salience and consumer choice. *Journal of Political Economy*, 121(5): 803-843.

Bushong, B., M. Rabin, & J. Schwartzstein. (2021). A model of relative thinking. *The Review of Economic Studies*, 88(1): 162–191.

Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk and Uncertainty*, 2(1): 61-104.

Cohen, A. & L. Einav (2007). Estimating risk preferences from deductible choice. *American Economic Review*, 97(3): 745–788.

Cutler, D., & R. Zeckhauser. (2004). Extending the Theory to Meet the Practice of Insurance. In R. Litan & R. Herring (eds.), Brookings-Wharton Papers on Financial Service. Washington, D.C.: Brookings Institution.

Ericson, K., & A. Starc. (2012). Heuristics and heterogeneity in health insurance exchanges: Evidence from the Massachusetts Connector. *The American Economic Review*, 102 (3): 493-97.

Faccio, M., Marchica, M., & Mura, R. (2016). CEO gender, corporate risk-taking, and the efficiency of capital allocation. *Journal of Corporate Finance*, 39: 193-209.

Fox, C.R., & R.T. Clemen. (2005). Subjective Probability Assessment in Decision Analysis: Partition Dependence and Bias Toward the Ignorance Prior. *Management Science*, 51(9): 1417-1432.

Fox, C.R., Erner, C. & Walters, D.J., (2015). Decision under risk: From the field to the laboratory and back. The Wiley Blackwell handbook of judgment and decision making, 2: 41-88.

Fox, C.R., & Y. Rottenstreich. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, 14(3): 195-200.

Gomes, F., M. Haliassos, & T. Ramadorai. (2021). Household Finance. *Journal of Economic Literature*, 59(3), 919-1000.

Gul, F. (1991). A Theory of Disappointment Aversion. *Econometrica*, 59(3): 667–686.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7), 2643–2682.

Handel, B. R. & J. Kolstad. (2015). Health insurance for "humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8): 2449–2500.

Heath, C., Larrick, R. P., & G. Wu. (1999). Goals as Reference Points. *Cognitive Psychology*, 38(1): 79–109.

Heiss, F., Leive, A., McFadden, D., & Winter, J. (2013). Plan selection in Medicare Part D: evidence from administrative data. *Journal of Health Economics*, 32(6): 1325–1344.

Hey, J., & C. Orme. (1994). Investigating Generalizations of Expected Utility Theory Using Experimental Data. *Econometrica*. 62: 1291-1326.
Holt, C., & S. K. Laury. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92 (5): 1644-1655.
Jaspersen, J. G., Ragin, M. A., & J.R. Sydnor. (2022). Predicting insurance demand from risk attitudes. *Journal of Risk and Insurance*, 89: 63–96.
Kahneman, D., & A. Tversky. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263–291.
Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). Noise: a flaw in human judgment. First edition.
Kamenica, E. (2008). Contextual Inference in Markets: On the Informational Content of Product Lines. *The American Economic Review*, 98(5): 2127–2149.
Köbberling, V., & Wakker, P. P. (2005). An index of loss aversion. *Journal of Economic Theory*, 122(1): 119-131.
Kőszegi, Botond, & M. Rabin. (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, 121 (4): 1133–65.
Kőszegi, Botond, & M. Rabin. (2007). "Reference-Dependent Risk Attitudes." *The American Economic Review*, 97(4): 1047–73.
Kőszegi, B. & Szeidl, A., 2013. A model of focusing in economic choice. *The Quarterly Journal of Economics*, 128(1): 53-104.
Kunreuther, H. (1973). Recovery from natural disasters; insurance or Federal aid? United States: American Enterprise Institute for Public Policy Research.
Kusev, P., Purser, H., Heilman, R., Cooke, A.J., Van Schaik, P., Baranova, V., Martin, R. & Ayton, P., 2017. Understanding risky behavior: The influence of cognitive, emotional and hormonal factors on decision-making under risk. *Frontiers in Psychology*, 8: 102.
Lintner, J., 1965. Security prices, risk, and maximal gains from diversification. *The Journal of Finance*, 20(4): 587-615.
Loomes, G. & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, 53(2): 271-282.
Martínez-Marquina, A., Niederle, M., & E. Vespa. (2019). Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10): 3437-74.
Niederle, M. & L. Vesterlund. (2007). Do Women Shy Away from Competition? Do Men Compete Too Much?, *The Quarterly Journal of Economics*, 122(3): 1067-1101.
Niederle, M. (2017). Chapter 8. Gender. In J. Kagel & A. Roth (Ed.), The Handbook of Experimental Economics, Volume 2: The Handbook of Experimental Economics (pp. 481-562). Princeton: Princeton University Press.
Post, T., Assem, M., Baltussen, G., & R.H. Thaler. (2008). Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show. *American Economic Review*, 98(1): 38–71.
Prelec, Drazen. (1998). The Probability Weighting Function. *Econometrica*, 66(3): 497–527.
Rabin, M. (2000). Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica*, *68*(5): 1281–1292.
Rabin, M., & R.H. Thaler. (2001). Anomalies: Risk Aversion. *Journal of Economic Perspectives*, 15(1): 219-232.
Sugden, R. (2003). Reference-dependent subjective expected utility. Journal of Economic Theory, 111 (2), 172–191.
Sunstein, C.R. (2002). Probability Neglect: Emotions, Worst Cases, and Law. Yale Law Journal, 112, 61.
Sunstein, C.R., & R.J. Zeckhauser. (2010). "Dreadful possibilities, neglected probabilities." The Irrational Economist. Making Decisions in a Dangerous World. Eds. Erwann Michel-Kerjan & Paul Slovic, Public Affairs Books, NY: New York.
Sydnor, J. (2010). (Over)insuring Modest Risks. *American Economic Journal: Applied Economics*, 2(4): 177–99.

Tversky, A., & D. Kahneman. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5 (4): 297–323.
Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. Psychological review, 101(4), 547.
von Neumann, J., & O. Morgenstern. (1947). Theory of games and economic behavior (2nd ed.). Princeton, NJ: Princeton University Press.

Figure 1.
Average Program and Group Choice Shares for Goals 2 and 3

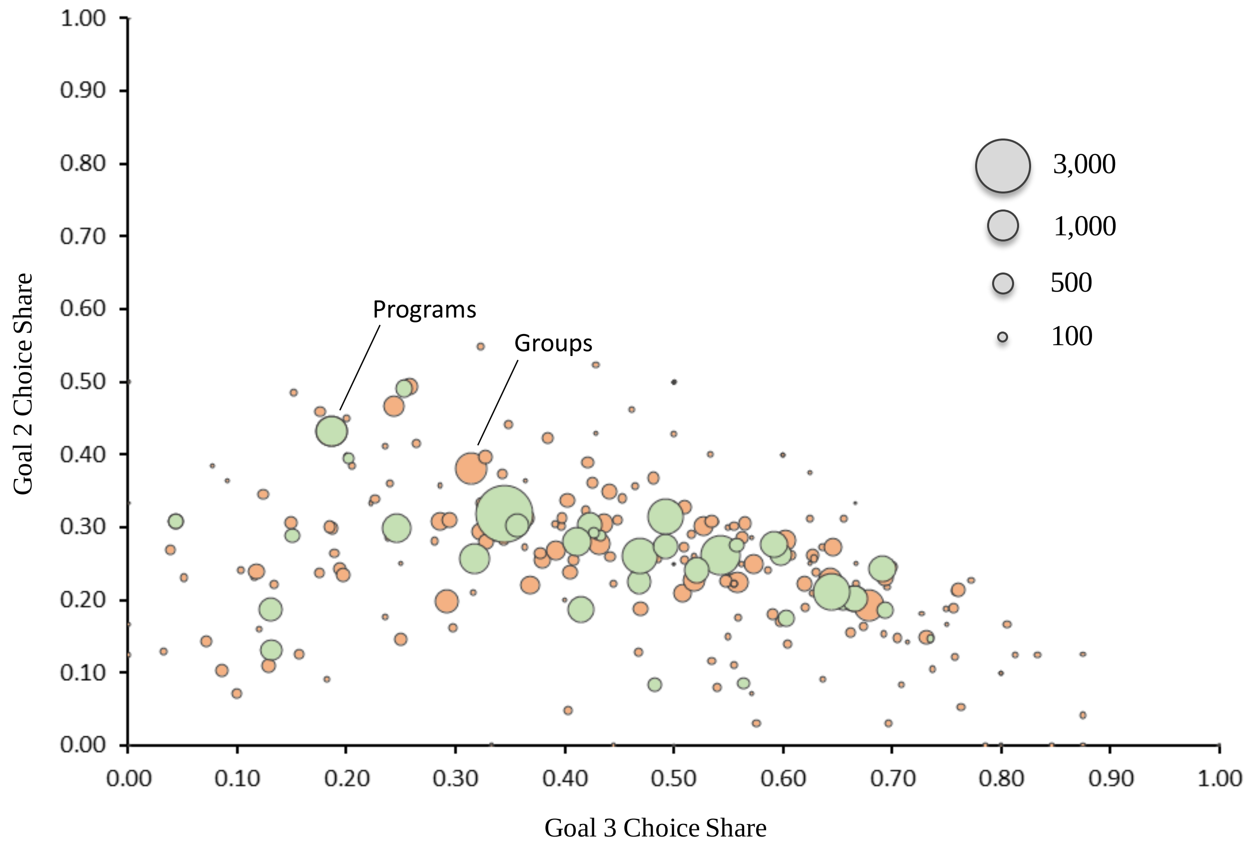

Notes: This figure depicts average choice shares for Goals 2 and 3 for each program (green) and group (orange). Groups with less than 10 employees excluded.

Figure 2.
Choice Characterization by Goal Expected Values

Panel A. Goal 3 Choice Share

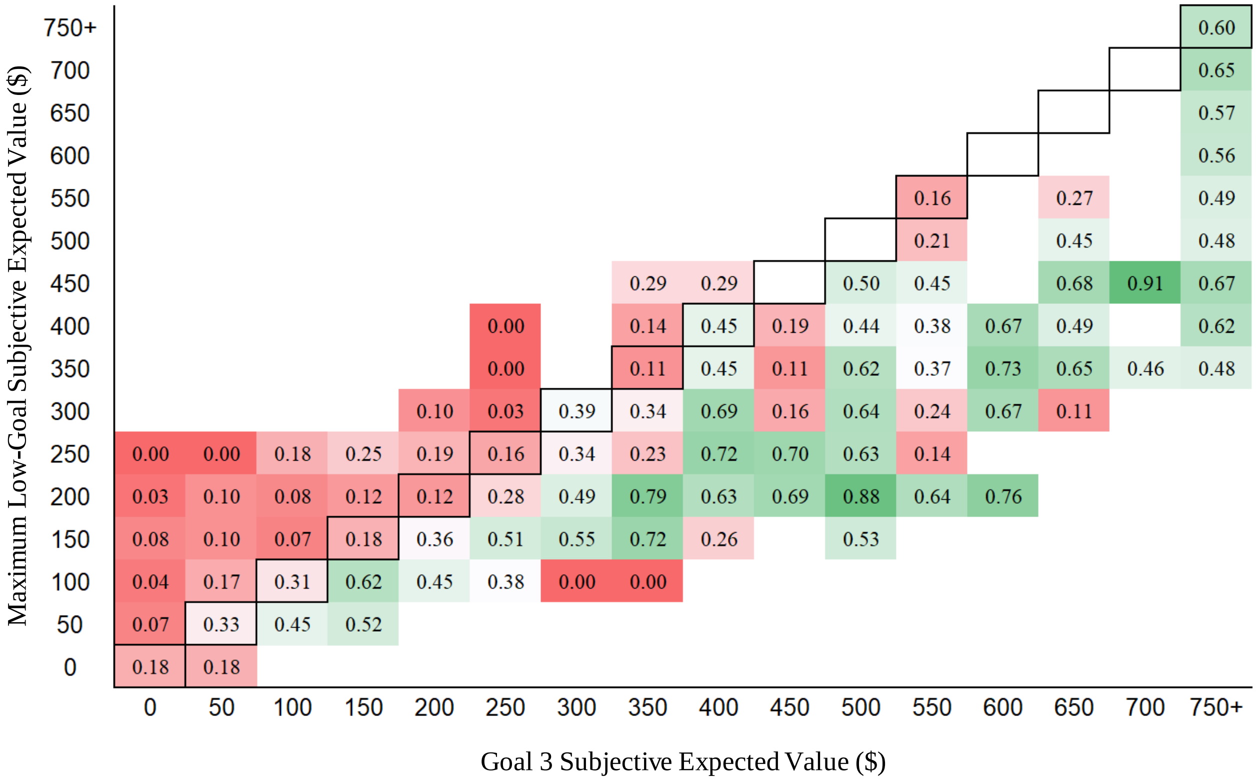

Panel B. Optimal Choice Share under Expected Value Benchmark

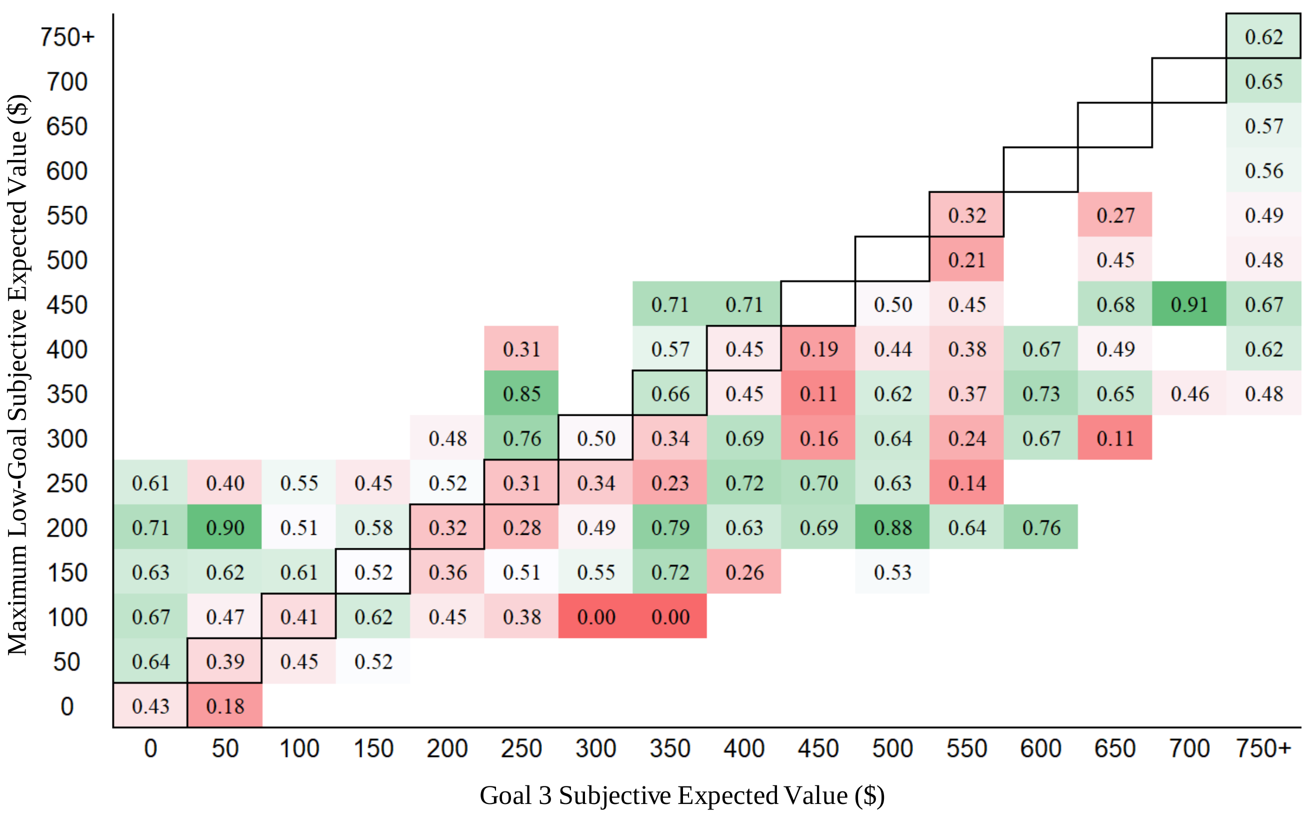

Notes: This figure depicts the Goal 3 (Panel A) and EV-optimal (Panel B) choice shares for varying low-goal and Goal 3 subjective expected values. Subjective expected values are censored at $750 and cells with less than 10 observations are omitted. Low-goal subjective expected values refer to the maximum subjective expected value of Goals 1 and 2.

Figure 3.
Optimal and Conservative Choice under Expected Utility Benchmark by Potential Reward and Information Regime

Panel A. Optimal Choice Share under Expected Utility (r = 0.0003)

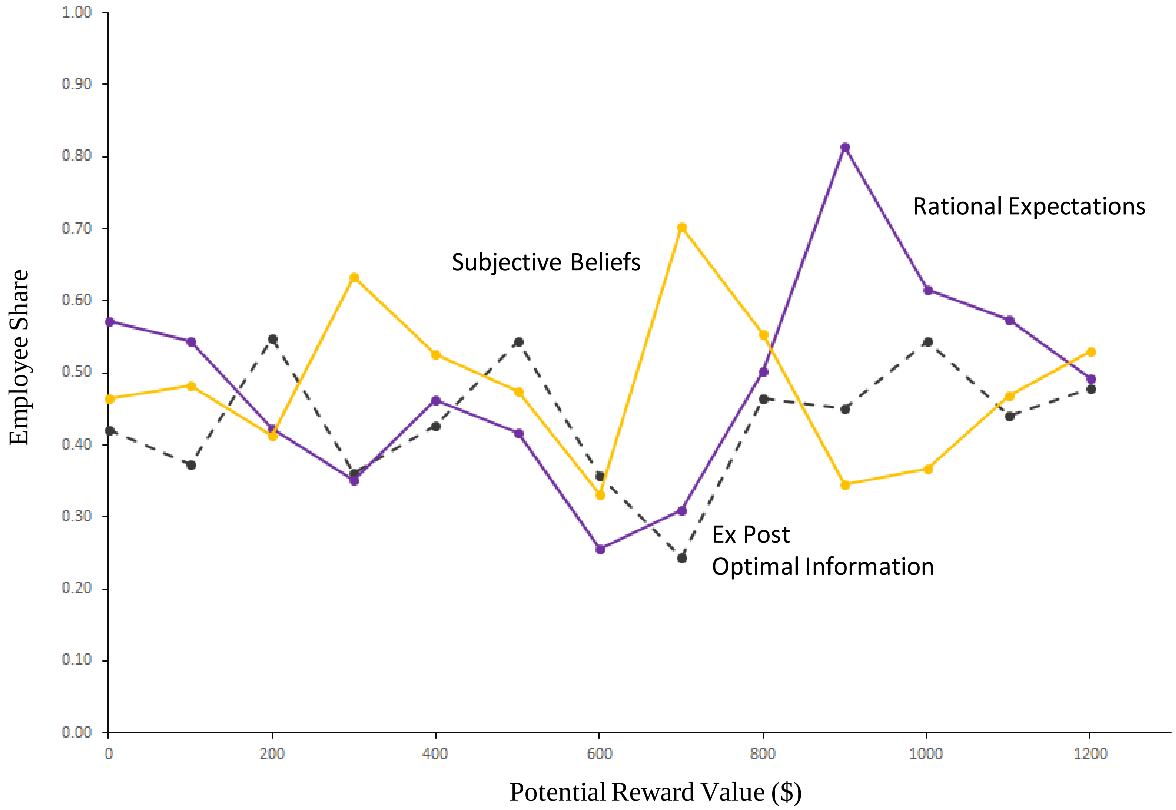

Panel B. Conservative Choice Share under Expected Utility (r = 0.0003)

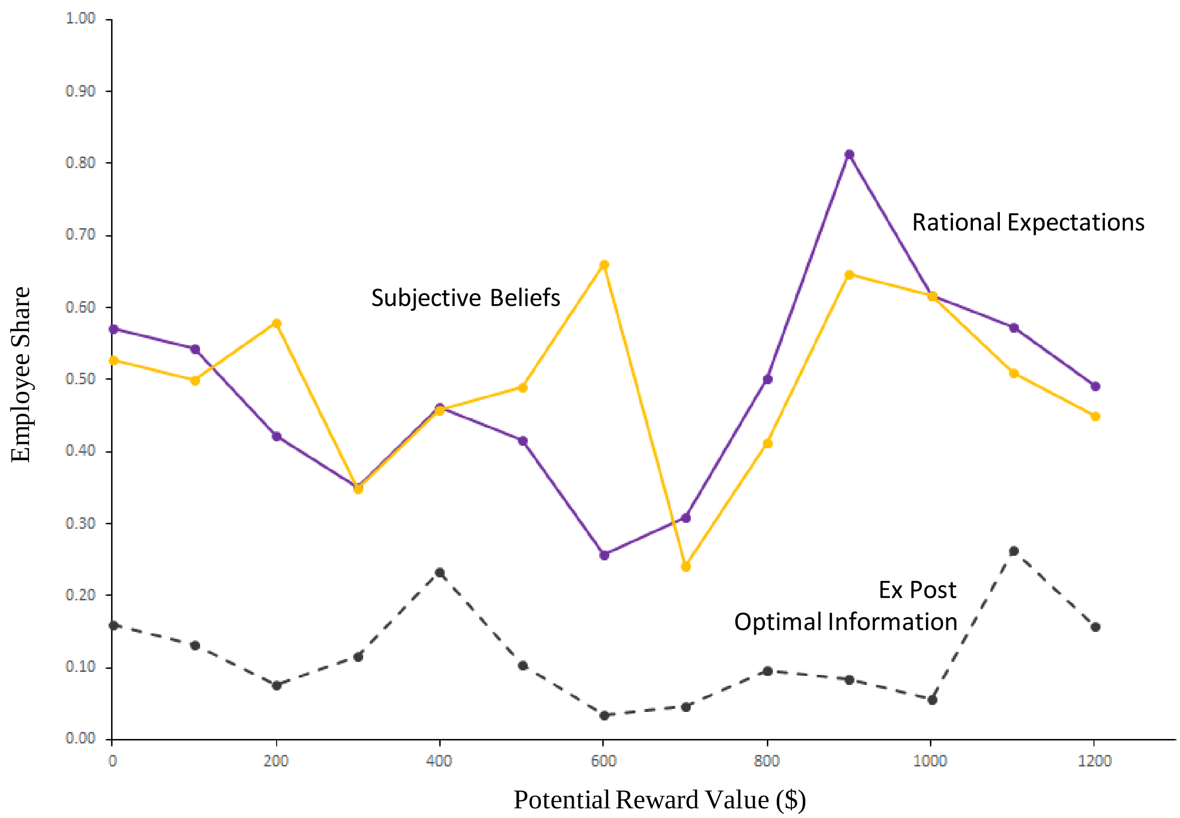

Notes: This figure reports the share of optimal (Panel A) and conservative (Panel B) choice by potential reward value under the expected utility (r = 0.0003) benchmark across varying assumptions regarding employee beliefs. Potential reward value refers to an employee's largest earnable reward (Goal 3) and is censored at $1,150.

Figure 4.
Optimal Choice under Expected Utility Benchmark by Risk Preference and Information Regime

Panel A. Expected Utility Benchmark assuming Rational Expectations

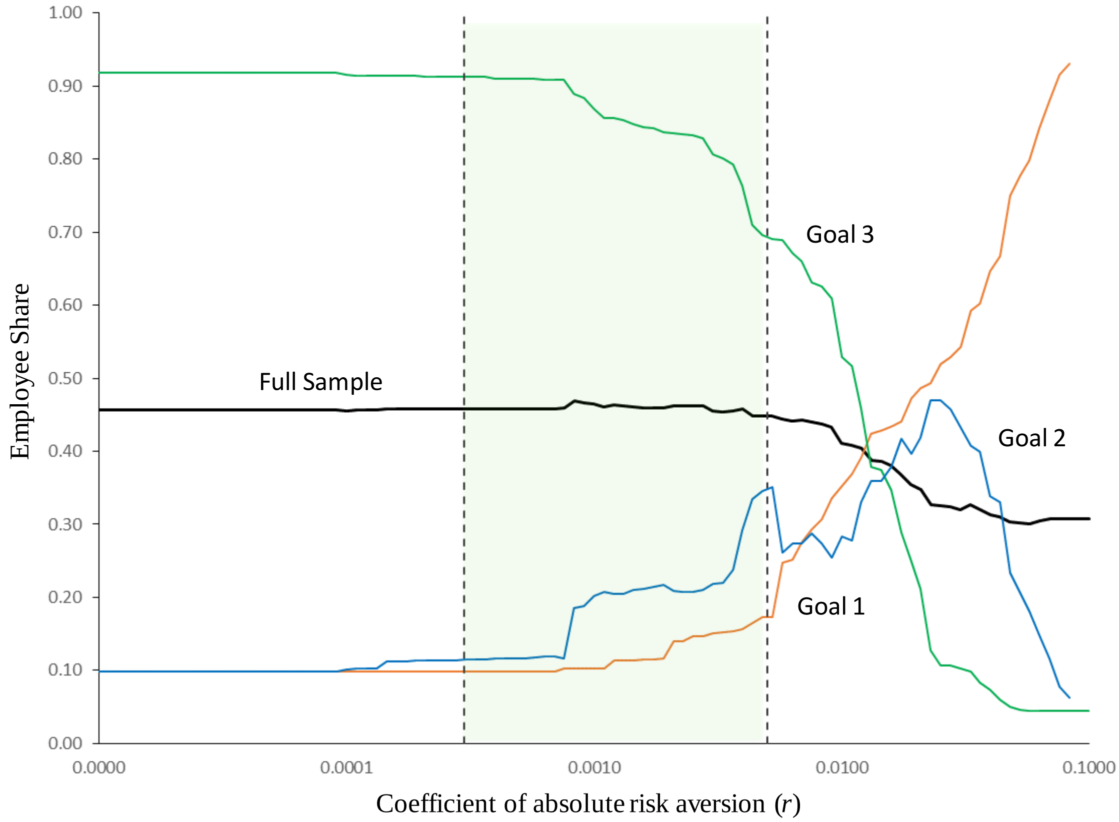

Panel B. Expected Utility Benchmark assuming Subjective Beliefs

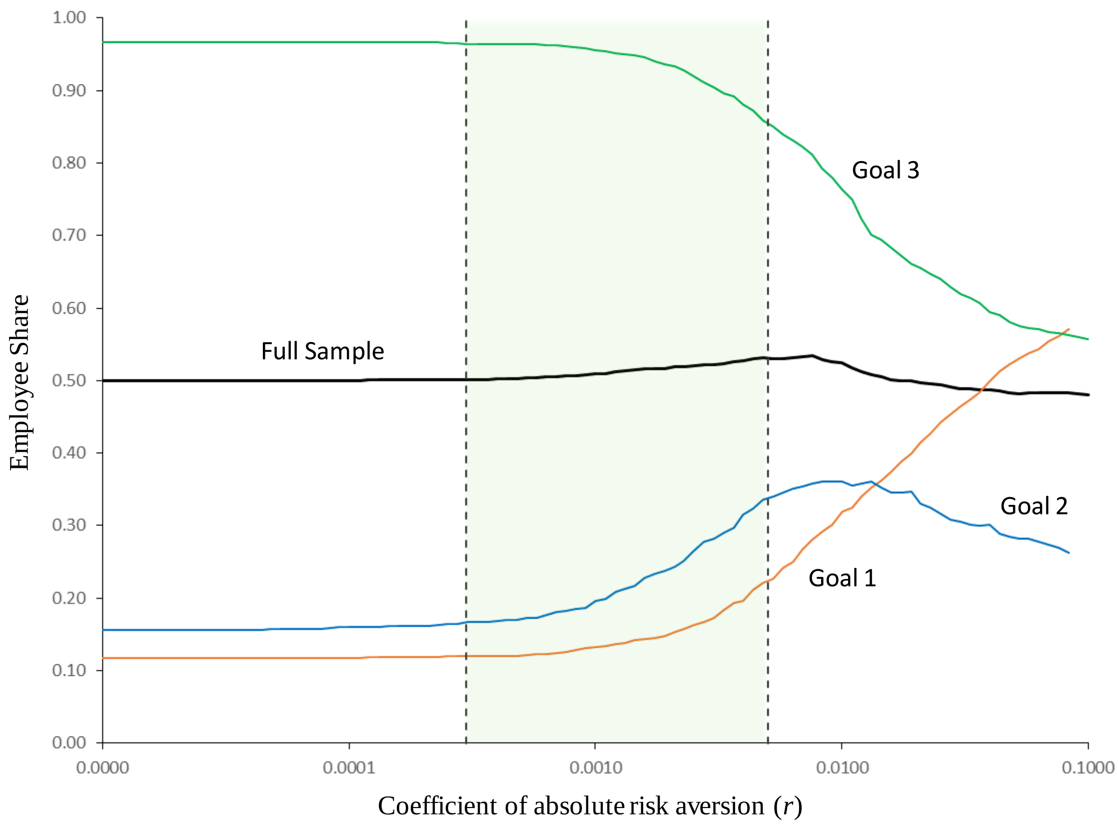

Notes: This figure depicts the share of optimal choice overall and by goal choice under expected utility across varying assumptions regarding CARA risk preferences and employee beliefs. Panel A depicts the share of optimal choice assuming rational expectations for an extended range of r on a logarithmic scale while Panel B depicts the analogous optimal choice share assuming subjective beliefs. The shaded region denotes the range of substantial but still plausible risk aversion, $r \in [0.0003, 0.005]$.

Figure 5.
Pairwise Partition Dependence and Perceived Goal Attainment CDFs

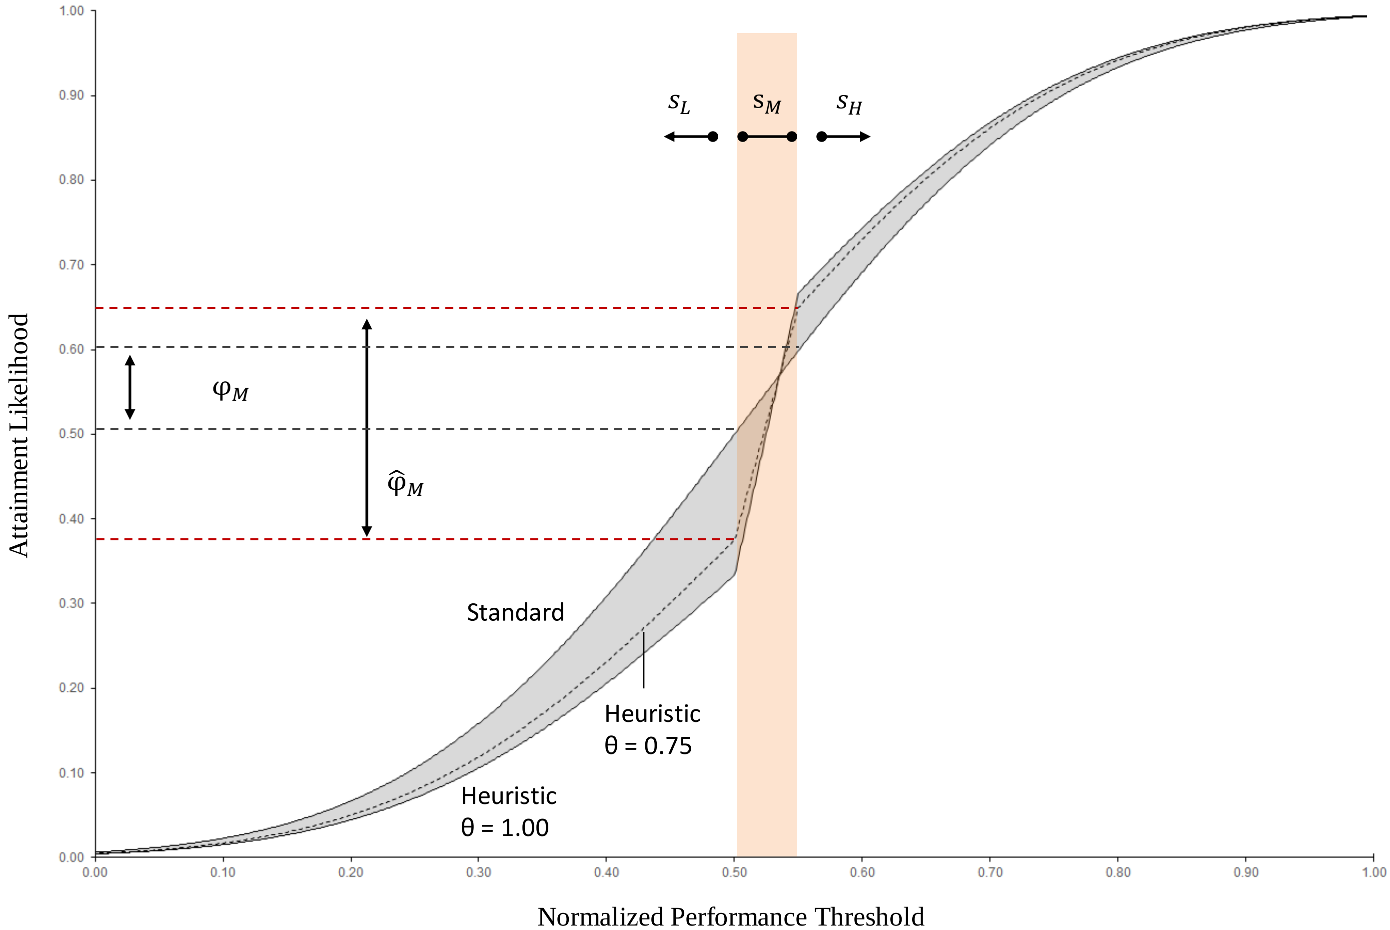

Notes: This figure depicts stylized goal attainment CDFs for GQ pairwise comparison under standard and heuristic choice (θ = [0.75, 1.00]). The shaded vertical region depicts the between-goal state space; the region to its left denotes the low-goal state space; and the region to its right denotes the high-goal state space. Shaded areas between CDFs depict inferential bias under heuristic evaluation. Arrows denote the actual and perceived likelihood of φ_M under standard and heuristic choice (θ = 0.75).

Figure 6.
Net Insurance Demand Bias under Heuristic Choice across Perceived Loss Likelihood

Panel A. Net Insurance Demand Bias and Bias Severity

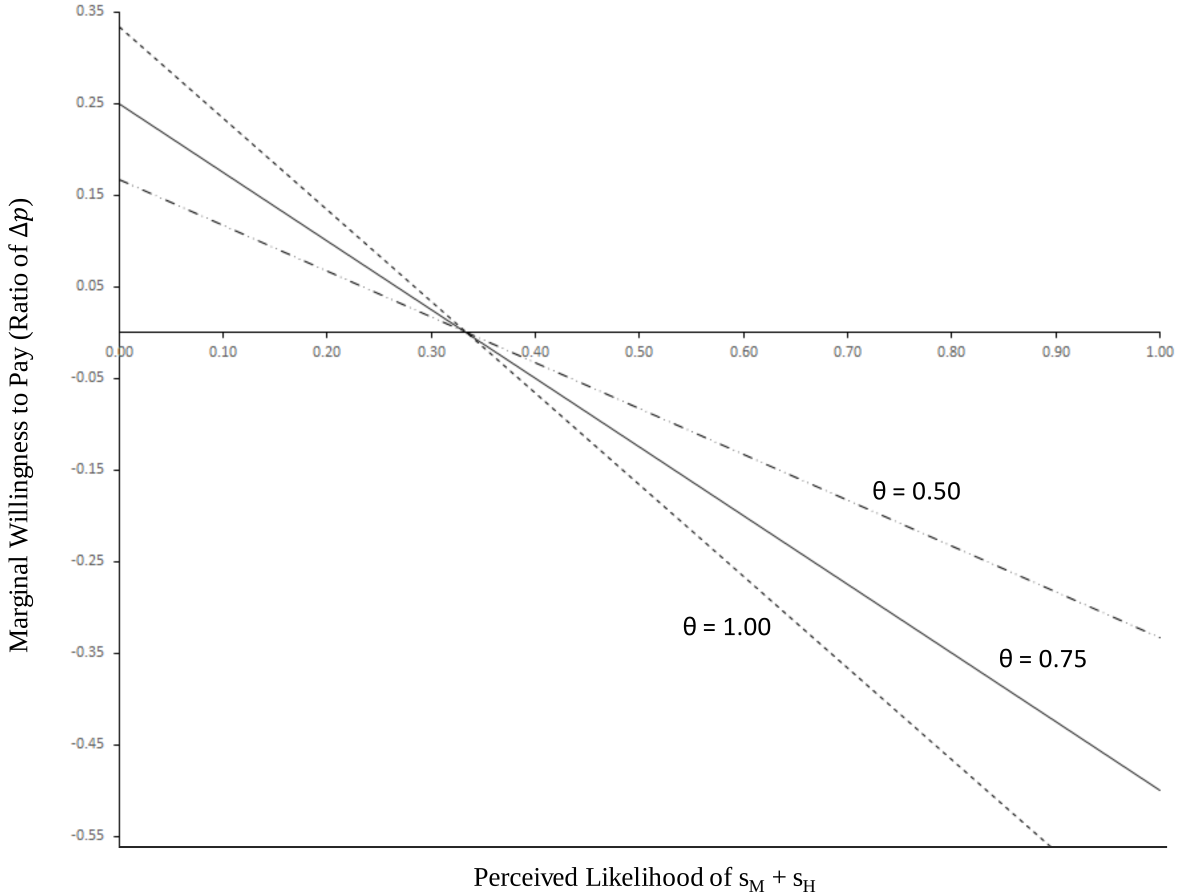

Panel B. Net Insurance Demand Bias and Relative Risk Ratio ($\frac{\hat{\varphi}_M}{\hat{\varphi}_H}$)

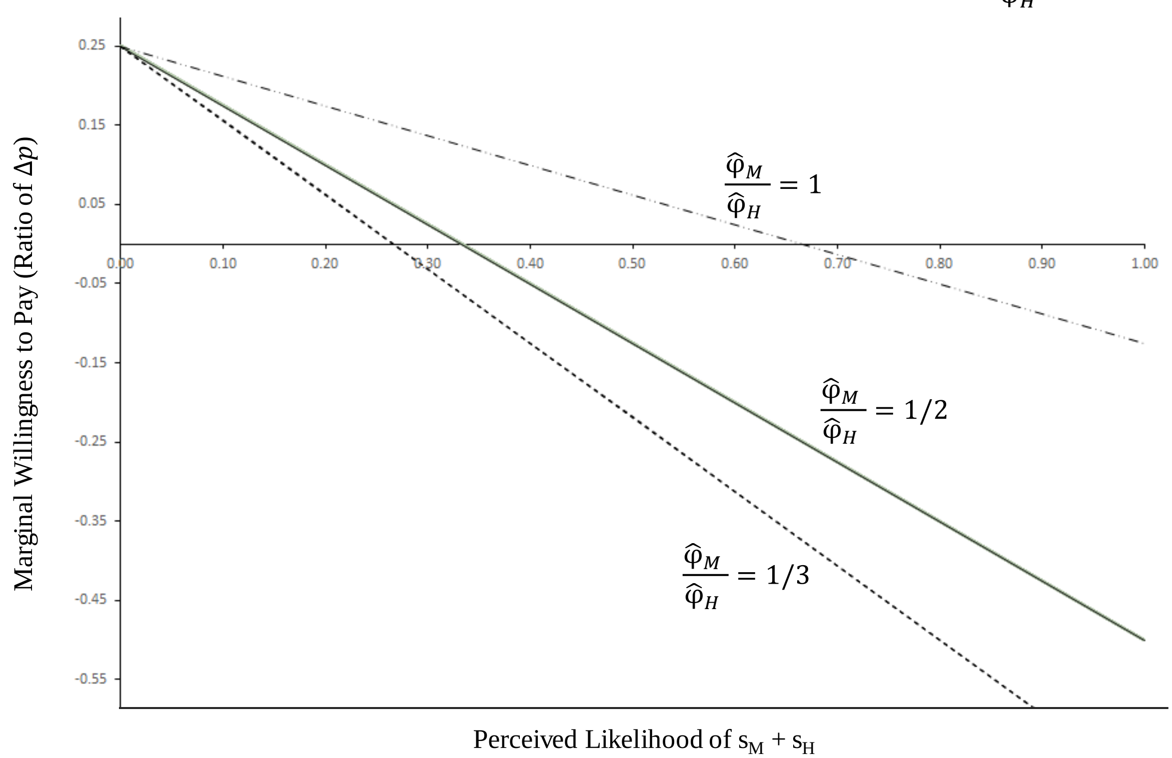

Notes: This figure depicts the net bias in insurance demand under heuristic choice across varying baseline levels of perceived risk loss for different bias severity (Panel B) and relative risk ratios (Panel B). Panel A assumes a relative risk ratio of 1:2 and Panel B assumes θ = 0.75. Net bias is expressed as the excess willingness to pay for a high versus low coverage plan, relative to a standard benchmark, as a ratio of the price difference between plans.

Figure 7.
Home Insurance Plan Choice across Loss Likelihood (Experiment F)

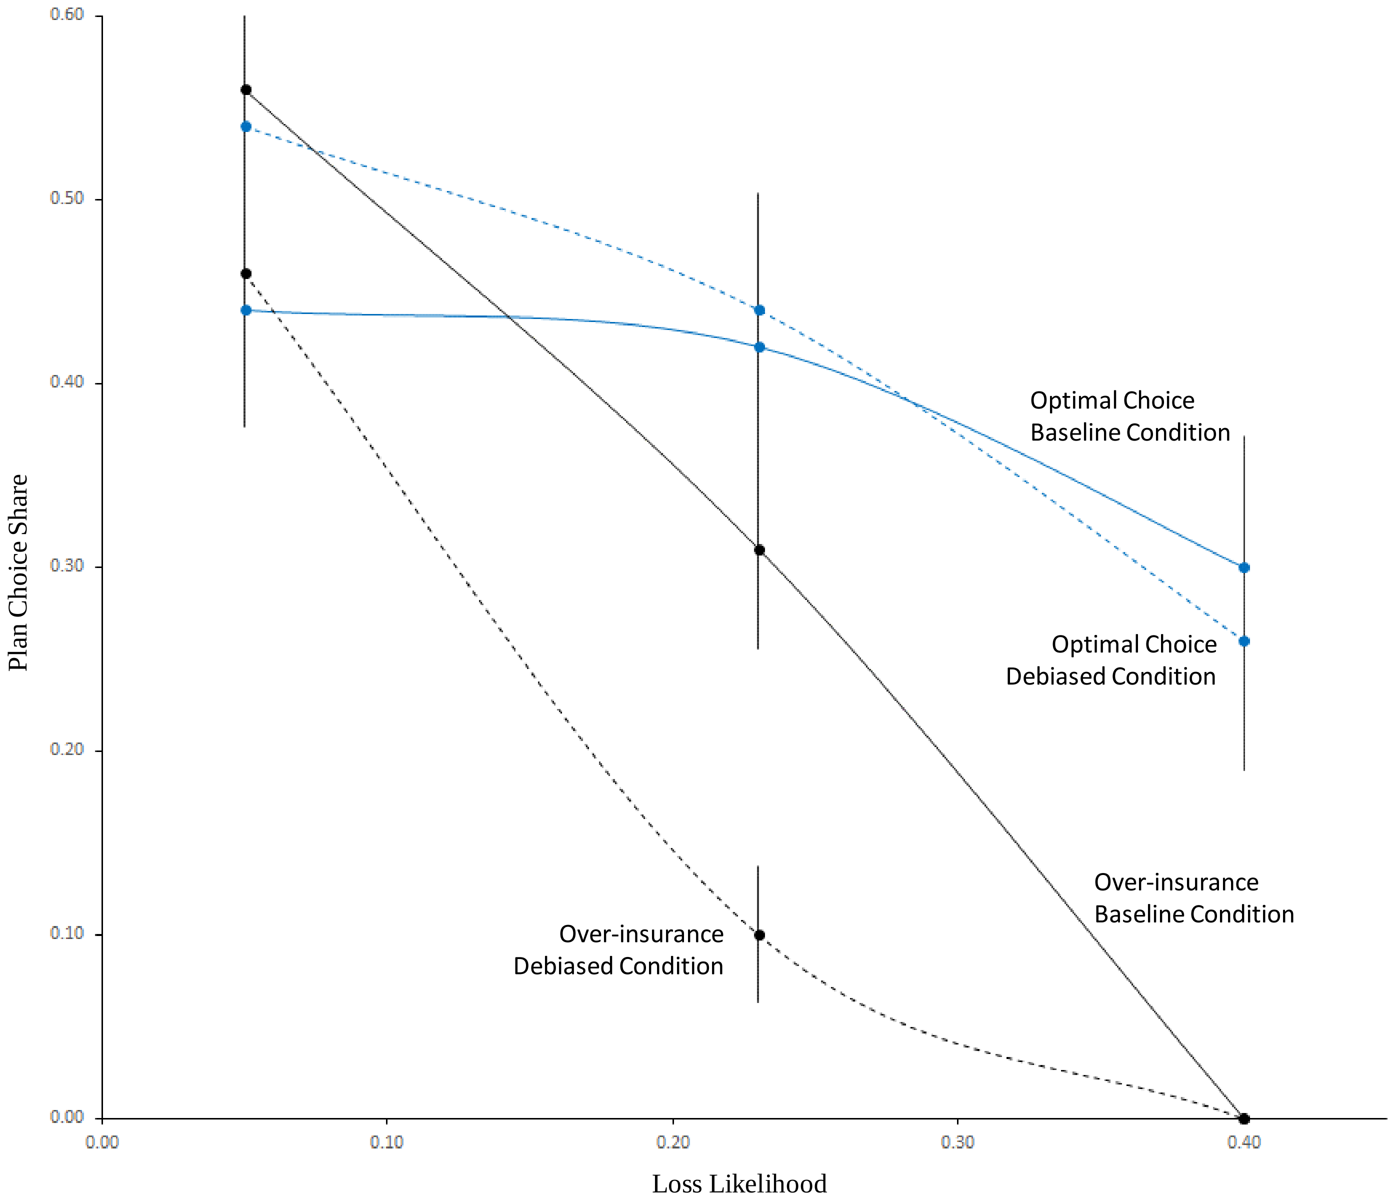

Notes: This figure depicts the share of over-insurance (black) and optimal choice (blue) for the baseline (solid) and debiased (dashed) conditions across varying loss likelihood from Experiment F. Loss likelihood refers to the percent likelihood of any loss during the plan coverage period (i.e., 0.05, 0.23, or 0.40). Error bars denote +/- 1 standard error.

Table 1.
Summary of Sample, Group and Employee Characteristics

	All	Potential Reward Value Below Median	Potential Reward Value Above Median
Panel A. Sample Overview			
Programs	34	-	-
Groups	232	-	-
Employees	20133	-	-
Firms	18	-	-
Employees per Group (Average)	87	-	-
	(139)		
Employees per Program (Average)	592	-	-
	(587.5)		
Panel B. Group Characteristics (Employee Shares)			
Program Duration			
≤ 30 days	0.39	0.51	0.28
45 to 60 days	0.28	0.12	0.42
≥ 90 days	0.33	0.38	0.29
Potential Reward Value (Estimated $)			
Average	467	150	746
	(482)	(58)	(517)
Median	350	168	525
25th Percentile	175	94	392
75th Percentile	525	175	914
Panel C. Employee Characteristics			
Age [Midpoint of 10-year bins]	36.9	36	37.6
Female	0.46	0.50	0.43
Tenure Category			
< 1 year	0.28	0.32	0.25
1 to 5 years	0.45	0.46	0.43
6 to 10 years	0.14	0.13	0.14
> 10 years	0.13	0.08	0.18
Program-Average Salary (Average) ($1,000s)	70.8	63.2	72.7
Data on Salary Available	0.25	0.10	0.38

Notes: This table summarizes observable detail on GQ programs and employees for the primary sample. Panel A describes the number and size of programs, while Panel B describes employee-level statistics regarding average program duration and potential rewards. Potential reward value refers to an employee's largest earnable reward (Goal 3 reward). Panel C summarizes employee demographic detail overall and by sub-groups distinguished by potential reward value. We impute age from self-reported 10-year bins, infer gender using a combination of self-reported data and inferences from first name, and approximate salary with program-level averages for programs with available data.

Table 2.
Goal Choice, Employee Productivity, and Goal Attainment

			Sample Restricted by Goal Choice	
	All	Goal 1	Goal 2	Goal 3
Panel A. Goal Choice				
Employees	20133	5866	5470	8797
Employee Share	1.00	0.29	0.27	0.44
Potential Reward Value (Average)	466	482	490	442
	(481.5)	(528)	(499)	(434.4)
Panel B. Employee Productivity				
Productivity Relative to Baseline				
Average	1.34	1.12	1.25	1.52
25th Percentile	0.88	0.78	0.89	0.91
50th Percentile	1.01	0.98	1.00	1.04
75th Percentile	1.20	1.11	1.15	1.27
Productivity Relative to Goal 3 Threshold				
Average	0.90	0.68	0.86	1.07
25th Percentile	0.60	0.30	0.63	0.77
50th Percentile	0.89	0.74	0.88	0.95
75th Percentile	1.02	0.95	1.00	1.09
Panel C. Goal Attainment				
Baseline	0.54	0.45	0.53	0.60
Goal 1	0.44	0.32	0.42	0.53
Goal 2	0.36	0.23	0.33	0.47
Goal 3	0.29	0.17	0.25	0.41
Earned Reward (Average)	121	33	92	197
Earned Reward (Average) | Goal Attainment	333	104	277	483

Notes: This table summarizes goal choice, productivity, and goal attainment for the primary sample overall and by employee goal choice. Panel A summarizes goal choice and average potential rewards, where potential reward value refers to an employee's largest earnable reward (Goal 3 reward). Panel B summarizes employee productivity both relative to baseline and to Goal 3 (the former measure excludes 18 percent of employees with no baseline data). Panel C summarizes goal attainment and average earned rewards.

Table 3.
Employee Beliefs and Confidence of Goal Attainment

	All	By Goal Choice		
		Goal 1	Goal 2	Goal 3
Panel A. Beliefs of Goal Attainment				
Rational Expectations				
Goal 1	0.44	0.41	0.44	0.46
Goal 2	0.37	0.32	0.36	0.39
Goal 3	0.30	0.25	0.28	0.33
Subjective Beliefs				
Goal 1	0.78	0.65	0.79	0.86
Goal 2	0.69	0.50	0.71	0.82
Goal 3	0.63	0.43	0.57	0.77
Panel B. Over/Under Confidence				
Ratio of Subjective/Rational Beliefs				
Goal 1	2.20	2.09	2.26	2.27
Goal 2	2.62	2.42	2.79	2.76
Goal 3	3.46	3.26	3.43	3.59
Relative Ratio of Over/Under Confidence				
Goal 3/Goal 1	1.45	1.41	1.42	1.48
Goal 3/Goal 2	1.22	1.24	1.15	1.22
Goal 2/Goal 1	1.13	1.08	1.18	1.17

Notes: This table summarizes employee beliefs and confidence associated with goal attainment for the primary sample overall and by employee goal choice. Panel A summarizes goal-attainment beliefs under rational expectations and subjective beliefs (see text for detail on estimates of rational expectations). Subjective beliefs reflect employee self-reports elicited during enhanced enrollment using a scale from 0 to 100 percent with 10-percent increments. For tractability, we adjust beliefs of 0 and 100 percent to 1 and 99 percent, respectively. Panel B summarizes average employee under/over confidence for each goal and relative under/over confidence for goal pairs. We represent confidence by the average ratio of subjective beliefs and rational expectations (> 1 indicates overconfidence), winsorized at the 5th and 95th percentiles.

Table 4.
Goal Choice Characterization under Expected Utility Benchmarks

	Risk Neutral EU		Expected Utility (CARA)					
			Rational Expectations			Subjective Beliefs		
	Rational	Subjective	r = 0.0003	r = 0.005	r [0, 0.005]	r= 0.0003	r = 0.005	r [0, 0.005]
Panel A. Characterization Overview								
Optimal Choice	0.45	0.50	0.45	0.44	0.56	0.50	0.53	0.59
Conservative Choice	0.49	0.48	0.49	0.38	--	0.48	0.39	--
Aggressive Choice	0.06	0.02	0.06	0.17	--	0.02	0.08	--
Panel B. Economic Consequences of Choice | Goal Attainment								
Counterfactual Loss								
Realized Reward	274	274	274	274	--	274	274	--
Counterfactual Reward | Ex Ante Optimal Choice	329	320	329	275	--	318	281	--
Counterfactual Loss (% of Counterfactual Reward)	0.17	0.14	0.17	0.00	--	0.14	0.02	--
Counterfactual Loss | Conservative Choice								
Realized Reward	164	162	162	122	--	159	118	--
Counterfactual Reward | Ex Ante Optimal Choice	303	281	302	244	--	276	222	--
Counterfactual Loss (% of Counterfactual Reward)	0.46	0.42	0.46	0.50	--	0.42	0.47	--
Panel C. Optimal Choice Share by Reward and Tenure								
Potential Reward Value								
Highest Quartile	0.42	0.48	0.42	0.39	--	0.49	0.55	--
Lowest Quartile	0.44	0.48	0.44	0.44	--	0.48	0.48	--
Employee Tenure								
Highest Category [10+ Years]	0.39	0.45	0.40	0.40	--	0.46	0.53	--
Lowest Category [< 1 Year]	0.44	0.47	0.44	0.44	--	0.47	0.50	--

Notes: This table characterizes the efficiency of goal choice for the primary sample under expected utility across varying assumptions regarding CARA risk preferences and employee beliefs. Panel A characterizes employee choices as either optimal, conservative, or aggressive relative to the prediction of the specified benchmark models. Panel B summarizes the economic consequence of goal choice via measures of average counterfactual loss. Panel C reports the share of optimal choice across employee sub-groups distinguished by potential reward value and employee tenure. Blank cells reflect an inability to uniquely characterize aggressive and conservative choices for benchmarks involving flexible values of r.

Table 5.

Goal Choice Characterization under Non-Standard Benchmarks

	SEU Baseline (CARA, r = 0.0003)	Non-Linear Weights [Prelec, α = β = 0.65]	Composite Gain-Loss [RP = g, η = 1, λ = 2.25]
Panel A. Characterization Overview			
Optimal Choice	0.50	0.47	0.59
Conservative Choice	0.48	0.52	0.24
Aggressive Choice	0.02	0.01	0.17
Panel B. Economic Consequences of Choice | Goal Attainment			
Counterfactual Loss			
Realized Reward	274	274	274
Counterfactual Reward | Ex Ante Optimal Choice	318	324	272
Counterfactual Loss (% of Counterfactual Reward)	0.14	0.15	-0.01
Counterfactual Loss | Conservative Choice			
Realized Reward	159	163	168
Counterfactual Reward | Ex Ante Optimal Choice	276	282	296
Counterfactual Loss (% of Counterfactual Reward)	0.42	0.42	0.43
Panel C. Optimal Choice Share by Reward and Tenure			
Potential Reward Value			
Highest Quartile	0.49	0.44	0.61
Lowest Quartile	0.48	0.46	0.55
Employee Tenure			
Highest Category [10+ Years]	0.46	0.42	0.59
Lowest Category [< 1 Year]	0.47	0.44	0.59

Notes: This table characterizes the efficiency of goal choice for the primary sample under two non-standard benchmark models. The first column provides a baseline characterization of choice under the SEU benchmark (CARA utility, r = 0.0003); the second column characterizes choice for a modified benchmark allowing for non-linear probability weights (Prelec 1998); the final column characterizes choice for the best-performing gain-loss utility benchmark (see text for details). Panel A characterizes employee choices as either optimal, conservative, or aggressive relative to the prediction of the specified benchmark models. Panel B summarizes the economic consequence of goal choice via measures of average counterfactual loss. Panel C reports the share of optimal choice across employee sub-groups distinguished by potential reward value and employee tenure.

Table 6.
Gender and Conservative Goal Choice in the Field under Different Benchmark Models

		Expected Utility (CARA)				Behavioral Departures from SEU		
		Rational		Subjective		Non-Linear Weights	Gain-Loss	PPD Heuristic
Sample	Lower Goal Share	r = 0	r = 0.0003	r = 0	r = 0.0003	[Prelec, α=β=0.65]	λ=2.25]	[θ = 0.75; w = $25]
All	0.56	0.49	0.49	0.48	0.48	0.52	0.24	0.15
Women	0.63	0.57	0.56	0.55	0.54	0.58	0.27	0.17
Men	0.50	0.43	0.42	0.43	0.43	0.46	0.21	0.13
Gender Difference [W-M]	0.13	0.14	0.14	0.12	0.11	0.12	0.06	0.04

Notes: This table compares conservative goal choice for women and men in the primary sample under various benchmark models. After describing the low-goal (i.e., Goals 1 or 2) choice share, a first set of columns summarizes conservative choice under the EU benchmark across varying assumptions regarding CARA risk preferences and employee beliefs. A subsequent set of columns summarizes conservative choice under benchmark models featuring behavioral departures from an SEU baseline (r = 0.0003). The final column summarizes conservative choice under the PPD heuristic given the specified parameters.

Table 7.
Goal Choice Characterization under Standard and Non-Standard Benchmarks — Experimental Paradigm (Experiment A)

						Contextual Sorting Heuristics	
	SEU Baseline (CARA, r = 0.0003)	Non-Linear Weights [Prelec, α = β = 0.5]	Composite Gain-Loss [RP = g; η = 1] [λ = 2.25]	Personal λ	Flexible λ	Ability	Taste for Competition
All Menus (6/6)	0.04	0.03	0.18	0.19	0.29	0.10	0.09
Nearly All Menus (5+/6)	0.16	0.13	0.31	0.30	0.44	0.12	0.10
All 3 Goal Menus (4/4)	0.16	0.09	0.25			0.13	0.12
All 4 Goal Menus (2/2)	0.10	0.10	0.32			0.15	0.15

Notes: This table characterizes multiple measures of optimal goal choice for experimental participants under a range of standard and non-standard benchmark models (Experiment A). The first column provides a baseline characterization of choice under the SEU benchmark (r = 0.0003), while the next few columns characterize choice for benchmarks modified to allow for non-linear decision weights and gain-loss utility. A final set of columns characterizes choice for heuristic models involving contextual sorting by self-reported ability or tastes for competition.

Table 8.
Goal Choice Characterization under PDD Heuristic in the Field and Lab

		PDD Heuristic - Model Parameters								
		No Bias			Personalized Bias			Parameterized Bias (θ = 0.75)		
Decision Sample	RN SEU	[w = $0]	[w = $25; 1¢]	[w = $50; 2¢]	[w = $0]	[w = $25]	[w = $50]	[w = $0]	[w = $25; 1¢]	[w = $50; 2¢]
Field Data	0.50	0.51	0.65	0.78	--	--	--	0.47	0.83	0.92
Experiment C	0.38	0.38	0.42	0.45	0.56	0.59	0.72	0.36	0.49	0.63
Experiment A										
All Menus (6/6)	0.04	0.14	0.27	0.33	--	--	--	0.11	0.38	0.59
Nearly All Menus (5/6)	0.16	0.29	0.45	0.56	--	--	--	0.30	0.71	0.86

Notes: This table characterizes the optimality of goal choice in the primary field sample and lab (Experiments C and A) under a baseline benchmark and various formulations of the PDD heuristic benchmark. The first column provides a baseline characterization of choice under a risk-neutral subjective EU benchmark. The subsequent columns characterize heuristic choice across varying specifications of bias and noise. In columns in which two noise allowances are listed, the first pertains to the allowance in the field and Experiment C, while the second pertains to Experiment A.

ONLINE APPENDIX

Characterizing Choice with CRRA Utility Benchmarks

Our primary analysis assessed employee goal choice for benchmark models featuring a utility function from the constant absolute risk aversion (CARA) family. The assumption of CARA utility, over the more common choice of constant relative risk aversion (CRRA) utility, was motivated by tractability given a lack of data on employee wealth. In this section we recharacterize choice for the core benchmark models assuming CRRA utility across a wide range of wealth and degrees of relative risk aversion.

Specifically, we assume employees are governed by CRRA utility of the form: $u(x) = \frac{x^{1-\rho}}{(1-\rho)}$ for $\rho \neq 1$ and $u(x) = \ln(x)$ for $\rho = 1$. We assess choice for initial lifetime wealth ranging from $1,000 to $1,000,000 and relative risk aversion, $\rho \in [0.10, 50]$. To appreciate the breadth of risk attitudes captured by the latter interval, we follow Post et al. (2008) in mapping risk parameters to the implied certainty coefficient—that is the certainty equivalent expressed as a fraction of expected value—associated with a 50/50 bet of ($0, $10k) assuming initial wealth of $25,000. This interval almost certainly subsumes the range of plausible relative risk aversion—asserted by Holt and Laury (2002) as bounded by 0 and 1.37.

Appendix Table A1 summarizes the choice characterization by reporting the optimal choice share for the EU benchmark across beliefs (rational, subjective), initial wealth, and relative risk aversion. The table indicates that within the (highlighted) range of plausible attitudes towards risk (spanning certainty coefficients from 0.87 to 0.99), the CRRA benchmarks explain a share of choice virtually identical to the CARA analogues from Table 4 assuming either risk neutrality or moderate risk aversion for rational expectations (0.45) and subjective beliefs (0.50). Overall, we interpret the table as suggesting that one cannot attribute the lack of descriptive accuracy of the benchmark models tested in the main analyses to the assumption of constant absolute, rather than constant relative, risk aversion.

ONLINE APPENDIX – REFERENCES

Holt, C.A., & S.K. Laury. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92 (5): 1644-1655**.**

Post, T., Van den Assem, M. J., Baltussen, G., & R.H. Thaler. (2008). Deal or no deal? decision making under risk in a large-payoff game show. *American Economic Review*, 98(1): 38-71.

Appendix Figure A1.
Sample Image of GQ Goal Selection Interface

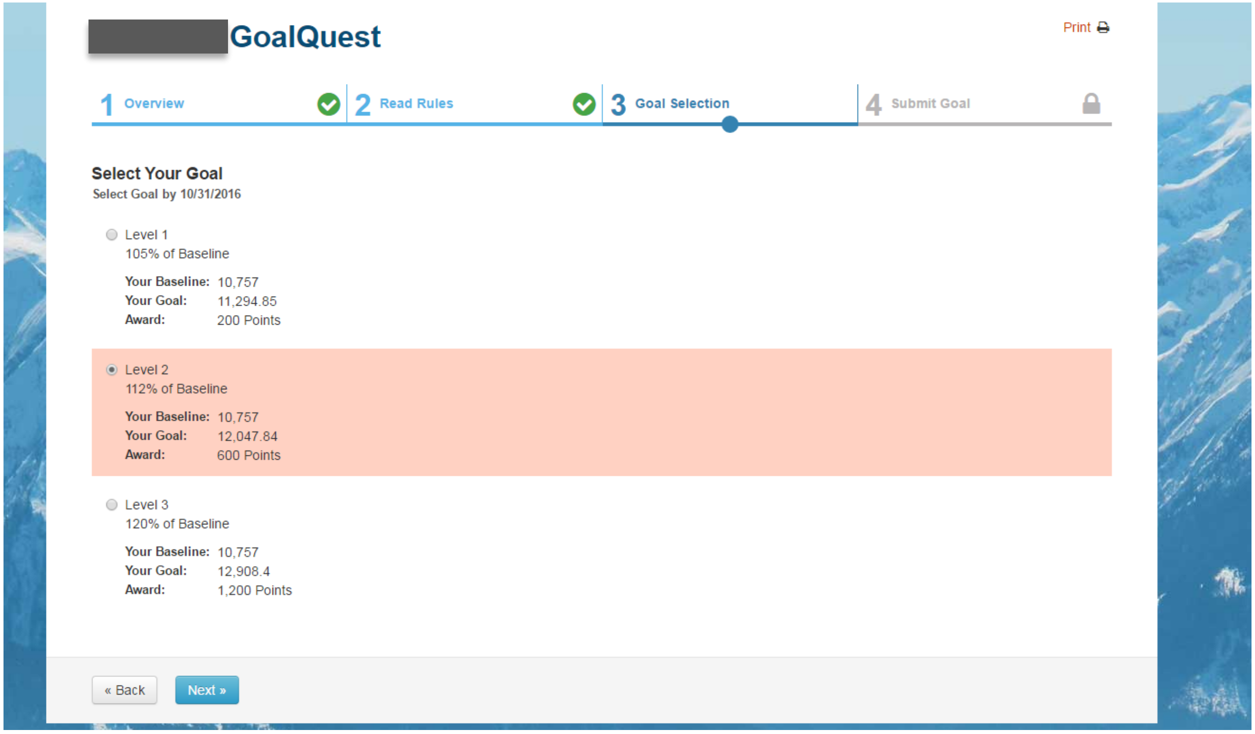

Appendix Figure A2.

Cumulative Distribution of Counterfactual Loss relative to Ex Post Optimal Choice | Goal Attainment

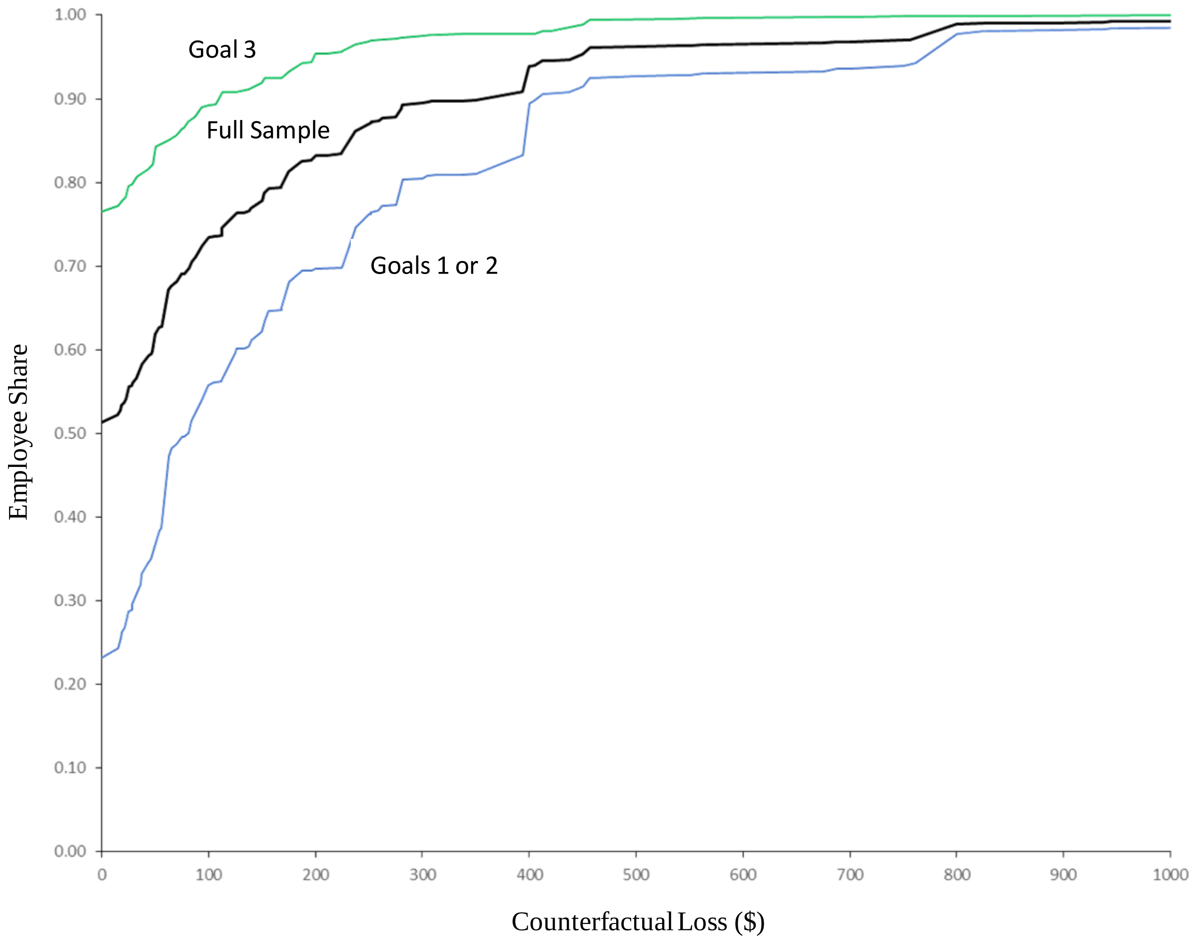

Notes: This figure depicts the cumulative distribution of counterfactual losses overall and by goal choice for employees whose productivity met or exceeded the Goal 1 threshold. Counterfactual loss refers to the difference between an employee's realized reward and the counterfactual reward an employee would have earned with ex post optimal choice (assuming no change in productivity). While counterfactual loss is censored at $1,000, a few employees had losses above $1,000 with a maximum loss of approximately $2,800.

Appendix Figure A3.
Distribution of Rational Expectations and Subjective Beliefs of Goal Attainment

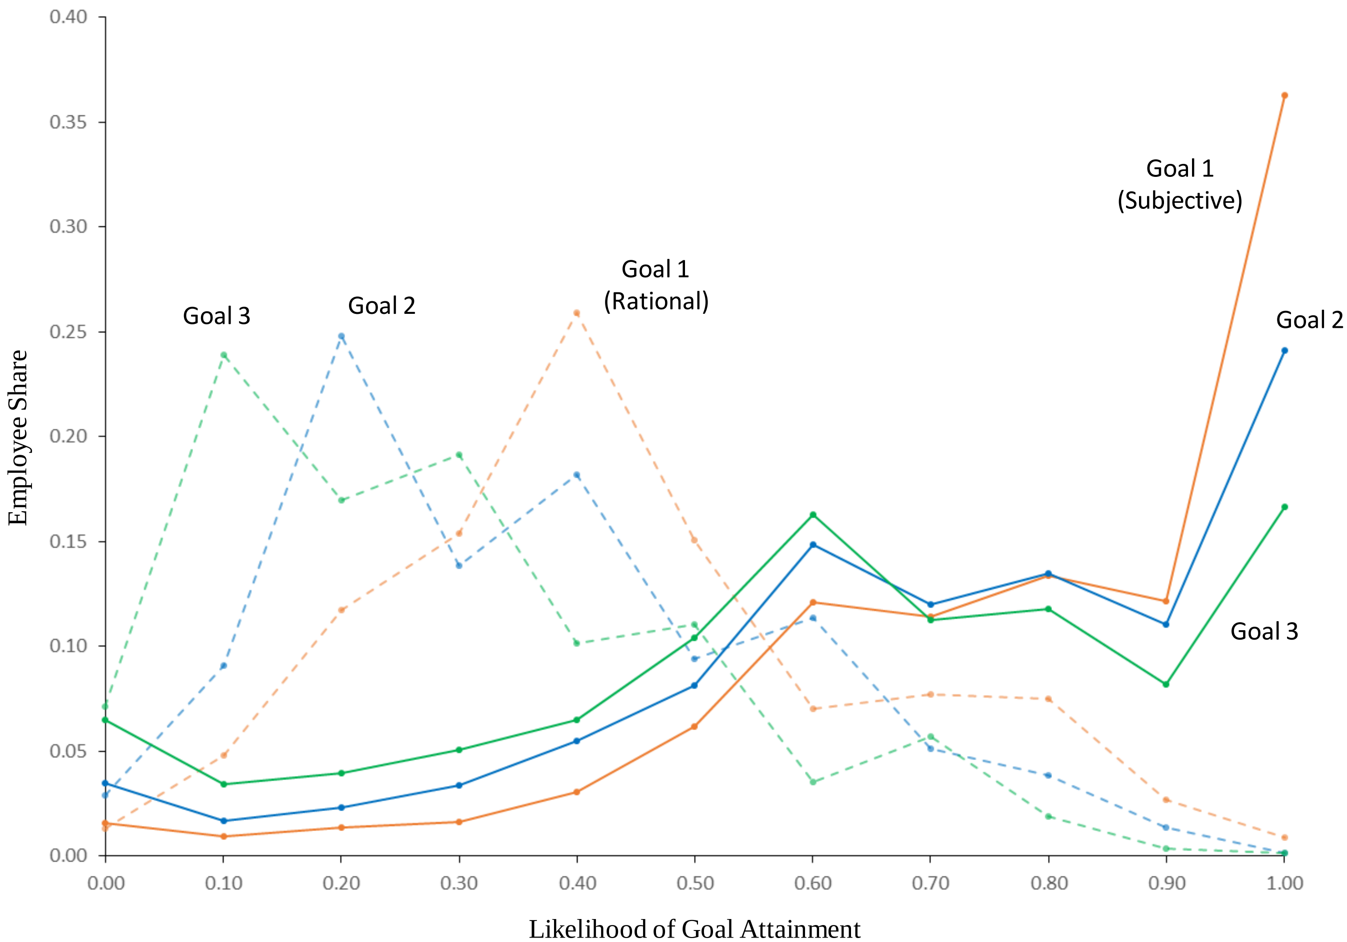

Notes: This figure compares the distributions of rational expectations and subjective beliefs of goal attainment for each goal. We assign employee- and goal-specific rational expectations by adjusting the ex post average rate of goal attainment at the group-level by employee age and gender, as estimated from a linear regression (for a small share of employees, for whom this strategy violated monotonicity, we adopted the unadjusted ex post average). Subjective beliefs for each goal reflect employee self-reports, elicited during enhanced enrollment, using an eleven-point scale (0, 10, 20,…, 100 percent). For ease of comparison, the figure groups rational expectations into bins that parallel the subjective belief data.

Appendix Figure A4.
Cumulative Distribution of Actual and Counterfactual Rewards under Risk Neutral Expected Utility Benchmark by Information Regime

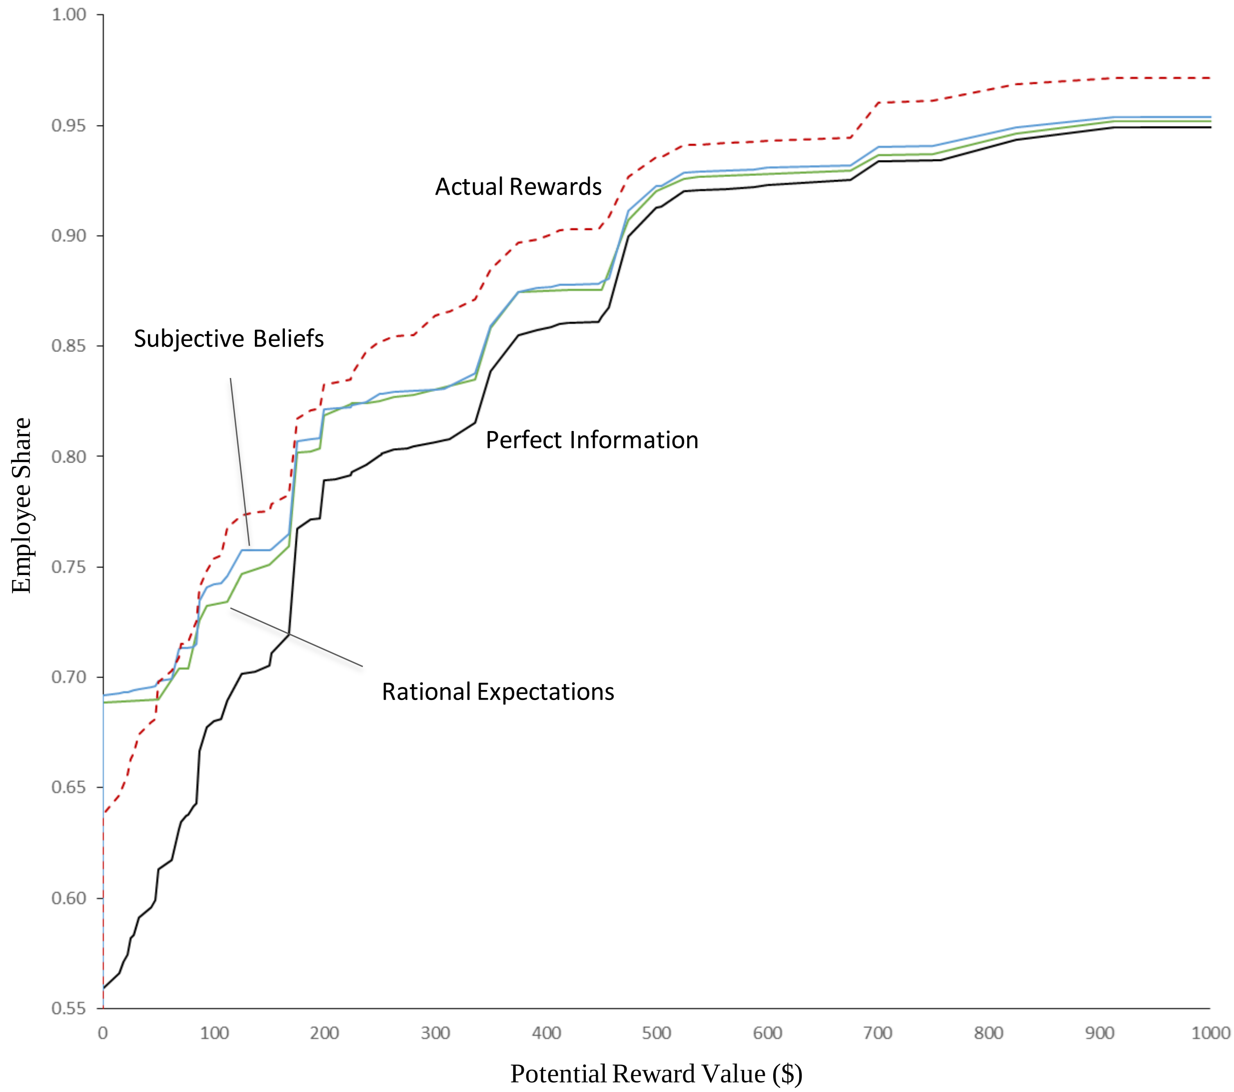

Notes: This figure depicts the cumulative distribution of actual rewards and counterfactual rewards under expected utility for risk neutral employees by information regime. Specifically, the dashed red line indicates the distribution of actual earned rewards, the black line indicates the distribution of counterfactual rewards given ex post optimal choice, the green line indicates the distribution of counterfactual rewards given ex ante optimal choice assuming rational expectations, while the blue line indicates the distribution of counterfactual rewards given ex ante optimal choice assuming subjective beliefs. The figure truncates the y axis at 0.55 to reflect the significant share of employees that did not attain any goal and, for clarity, truncates the x-axis at $1,000.

Appendix Table A1.
Goal Choice Characterization for CRRA Utiity Benchmarks

		Rational Expectations - Initial Lifetime Wealth						
ρ	CC(0/10k)	$1,000	$10,000	$25,000	$50,000	$100,000	$500,000	$1,000,000
0.10	0.99	0.45	0.45	0.45	0.45	0.45	0.45	0.45
0.25	0.98	0.45	0.45	0.45	0.45	0.45	0.45	0.45
0.50	0.96	0.45	0.45	0.45	0.45	0.45	0.45	0.45
0.75	0.94	0.45	0.45	0.45	0.45	0.45	0.45	0.45
1.00	0.92	0.45	0.45	0.45	0.45	0.45	0.45	0.45
1.50	0.87	0.45	0.45	0.45	0.45	0.45	0.45	0.45
2.50	0.79	0.46	0.45	0.45	0.45	0.45	0.45	0.45
5.00	0.61	0.46	0.45	0.45	0.45	0.45	0.45	0.45
10.00	0.37	0.42	0.46	0.45	0.45	0.45	0.45	0.45
50.00	0.07	0.30	0.45	0.45	0.46	0.45	0.45	0.45

		Subjective Expectations - Initial Lifetime Wealth						
ρ	CC(0/10k)	$1,000	$10,000	$25,000	$50,000	$100,000	$500,000	$1,000,000
0.10	0.99	0.50	0.50	0.50	0.50	0.50	0.50	0.50
0.25	0.98	0.50	0.50	0.50	0.50	0.50	0.50	0.50
0.50	0.96	0.50	0.50	0.50	0.50	0.50	0.50	0.50
0.75	0.94	0.50	0.50	0.50	0.50	0.50	0.50	0.50
1.00	0.92	0.51	0.50	0.50	0.50	0.50	0.50	0.50
1.50	0.87	0.51	0.50	0.50	0.50	0.50	0.50	0.50
2.50	0.79	0.52	0.50	0.50	0.50	0.50	0.50	0.50
5.00	0.61	0.53	0.50	0.50	0.50	0.50	0.50	0.50
10.00	0.37	0.53	0.51	0.50	0.50	0.50	0.50	0.50
50.00	0.07	0.48	0.53	0.52	0.51	0.50	0.50	0.50

Notes: This table characterizes the efficiency of goal choice with respect to benchmark models with CRRA utility across varying initial lifetime wealth and relative risk aversion. The second column reports the certainty coefficient (i.e., certainty equivalence as a share of expected value) assuming initial wealth of $25,000 for a fair bet of ($0, $10k). Highlighted region denotes interval of plausible relative risk aversion as indicated by Holt and Laury (2002). The first panel characterizes choice assuming rational expectations while the second panel characterizes choice assuming subjective expectations.

Appendix Table A2.
Optimal Goal Choice Shares under RN-SEU Benchmark with Effort Costs

	Baseline Effort Cost Increment as % of Wage						
Convexity (k)	0%	1%	3%	5%	10%	25%	50%
1.00	0.50	0.50	0.33	0.30	0.29	0.29	0.29
1.10	0.50	0.50	0.32	0.30	0.29	0.29	0.29
1.25	0.50	0.48	0.32	0.30	0.29	0.29	0.29
1.50	0.50	0.47	0.32	0.30	0.29	0.29	0.29
2.00	0.50	0.42	0.31	0.29	0.29	0.29	0.29
5.00	0.50	0.36	0.30	0.29	0.29	0.29	0.29

Notes: This table reports the share of optimal goal choice under a risk-neutral subjective EU benchmark model assuming varying specifications of effort costs. Baseline effort cost increment refers to the increase in hourly effort cost for Goal 2 versus Goal 1 as a % of wage. The convexity parameter refers to the proportional increase in effort costs for Goal 3 versus Goal 2 relative to the baseline increment. All calculations assume wage of $25/hour, 8 working hours per day, and the subjective beliefs elicited from employees. For example, a one-month program (~25 working days), baseline increment of 10%, and k = 1.5, implies total effort costs of $0, $500, and $1,250 for Goals 1, 2, 3, respectively.

Appendix Table A3.

Goal Choice Characterization for Expected Utility Benchmarks - Primary versus Expanded Sample

		Expected Utility (CARA)		
Characterization Overview	Risk Neutral	r = 0.0003	r = 0.005	r [0, 0.005]
Panel A. Primary Sample				
Optimal Choice	0.45	0.45	0.44	0.56
Conservative Choice	0.49	0.49	0.38	--
Aggressive Choice	0.06	0.06	0.17	--
Panel B. Expanded Sample				
Optimal Choice	0.41	0.41	0.40	0.54
Conservative Choice	0.56	0.55	0.48	--
Aggressive Choice	0.04	0.04	0.11	--

Notes: This table characterizes the efficiency of goal choice for the primary (Panel A) and expansive (Panel B) samples under expected utility across a range of assumptions regarding CARA risk preferences and assuming rational expectations.

Appendix Table A4.
Optimal Goal Choice Shares for Gain-Loss Utility Benchmarks by Candidate Reference Point

	Gain-Loss Utility (α = 0.88; η = 0)			Consumpton + Gain-Loss Utility (λ = 2.25)		
Candidate Reference Points	λ = 1.50	λ = 2.25	λ = 3.00	η = 1	η = 3	η = 5
<u>Panel A. Prospect Independent</u>						
Status Quo (0)	0.50	0.50	0.50	0.50	0.50	0.50
High Probability (Goal 1)	0.52	0.54	0.55	0.51	0.50	0.50
Compromise Goal (Goal 2)	0.50	0.52	0.52	0.50	0.50	0.50
Maximum Reward (Goal 3)	0.49	0.49	0.49	0.49	0.49	0.49
Maximum High Certainty	0.51	0.51	0.51	0.50	0.50	0.50
<u>Panel B. Prospect-Dependent</u>						
Reward of Chosen Goal	0.29	0.29	0.29	0.59	0.56	0.54
Expected Value of Chosen Goal	0.40	0.26	0.26	0.54	0.50	0.50
Reward of Chosen Goal + 1	0.55	0.55	0.55	0.53	0.53	0.52
Reward of Chosen Goal - 1	0.46	0.43	0.42	0.58	0.54	0.53
Regret (Expected Max Counterfactual)	0.50	0.50	0.50	0.50	0.50	0.50

Notes: This table assesses the descriptive accuracy of benchmark models involving gain-loss utility across several candidate reference points, functional forms, and parameter specifications. The first set of columns characterizes choice under benchmark models involving gain-loss utility following Kahneman and Tversky (1979) across potential values of the loss aversion parameter, λ. The second set of columns characterizes choice under benchmark models involving composite utility, an additively linear combination of consumption utility and gain-loss utility, across potential consumption utility scaling factors, n. (n = 0 therefore implies a model with gain-loss utility only). All benchmark models assume subjective beliefs. Panel A reports the share of optimal choice for prospect-independent candidate reference points while Panel B reports the analogous share of optimal choice for prospect-dependent candidate reference points. Please see text for additional detail on each of the benchmark models.

Appendix Table A5.
Demand for Prescription Drug Plans across Information Frames - Experiment D

	Menu Display		
	Baseline	Partition Dependent	Partition Independent
No Plan	0.11	0.18	0.13
Silver Plan [Coinsurance: 50%, Premium: $640]	0.59	0.53	0.44
Gold Plan [Coinsurance: 15%, Premium: $1220]	0.31	0.29	0.43
Expected Total Cost [Out-of-Pocket + Premium]	2076	2151	2041

Notes: This table reports average choice shares across conditions from Experiment D (N = 432). Participants were informed that coinsurance applies to all drug bills until the plan's out-of-pocket maximum of $7,500 (neither plan offered a deductible). They were also informed that annual drug bills could not exceed $10,000, even for those selecting no plan. Expected total cost refers to the estimated average total cost (premium + out-of-pocket costs) for participants in each condition based on their plan choices. Total cost estimates rely on an inferred distribution of potential drug bills (see text for details).

Appendix Table A6.
Demand for Home Insurance across Information Frames - Experiment E

	Menu Display			
	Full Information Baseline	No Information Baseline	Partition Dependent	Partition Independent
Basic Plan [Deductible: $1000, Premium: $616]	0.35	0.40	0.23	0.54
Medium Plan [Deductible: $500, Premium: $716]	0.39	0.41	0.50	0.38
Premium Plan [Deductible: $250, Premium: $803]	0.26	0.19	0.26	0.08
Expected Total Cost [Out-of-Pocket + Premium]	726	717	729	696

Notes: This table reports average plan choice shares across conditions from Experiment E (N = 435). Participants were informed that plans cover all expenses after the deductible has been met. Expected total cost refers to the estimated average total cost (premium + out-of-pocket costs) for participants in each condition based on their plan choice. Total cost estimates assume a 3 percent chance of damages exceeding $2500 and a 1 percent chance of damages of $500.