The Risk of Caution: Evidence from an R&D Experiment

Richard T. Carson,[a] Joshua Graff Zivin,[a,b] Jordan J. Louviere,[c] Sally Sadoff,[d] and Jeffrey G. Shrader[e]

Affiliations

[a] Department of Economics, University of California, San Diego, 9500 Gilman Dr. #0508, La Jolla, CA 92093-0508, USA.

[b] School of Global Policy and Strategy and Department of Economics, University of California, San Diego, 9500 Gilman Dr. #0508, La Jolla, CA 92093-0508, USA.

[c] School of Marketing, University of South Australia, 101 Currie Street, Adelaide, South Australia 5000, Australia.

[d] Rady School of Management, University of California, San Diego, 9500 Gilman Drive #0553, La Jolla, CA 92093-0553, USA.

[e] School of International and Public Affairs, Columbia University, 420 W 118th St., New York, NY 10027, USA.

Abstract

Innovation is important for firm performance and broader economic growth. But breakthrough innovations necessarily require greater risk-taking than more incremental approaches. To understand how managers respond to uncertainty when making research and development decisions, we conducted experiments with master's degree students in a program focused on the intersection of business and technology. Study participants were asked to choose whether to fund hypothetical research projects using a process that mirrors real-world research and development funding decisions. The experiments provided financial rewards that disproportionately encouraged the choice of higher-risk projects. Despite these incentives, most participants chose lower-risk projects at the expense of projects more likely to generate a large payoff. Heterogeneity analysis and additional experimental treatments show that individual risk preferences predict greater tolerance of high-risk projects and suggest that more appropriate decision making can be learned. Thus, for firms seeking to fund breakthrough R&D, appropriate screening and training of employees may play important roles in increasing the likelihood of success.

1. Introduction

Research and development (R&D) is an important determinant of firm growth and performance (Porter 1985; Amit and Zott 2001; Stephan 2010; Teece 2010; Keupp, Palmie, and Gassmann 2012). Innovation is also thought to be a fundamental driver of long-run economic growth (for instance in the Schumpeterian growth model of Aghion and Howitt 1992). But while R&D is important for the success of companies in many sectors, it is generally an expensive and complex undertaking. Deciding which elements of prior knowledge are important for current projects, what knowledge should be drawn from, and the particular form in which knowledge should be combined is often shrouded in uncertainty (Boudreau et al. 2016). Appropriate risk-taking is important because projects with greater uncertainty have a lower probability of bearing fruit but may also generate more path-breaking innovations if successful (Azoulay, Graff Zivin, and Manso 2011). In this paper, we study the effect of uncertainty on research funding decisions by asking how research project risk affects project choice.

One ingredient to a successful R&D program is its ability to encourage appropriate risk taking—tolerating failure in pursuit of reward (March 1991, Manso 2011). This is consistent with recent empirical evidence on research grants (Azoulay, Graff Zivin, and Manso 2011) as well as for venture backed funding of start-up firms (Tian and Wang, 2011). Although the importance of appropriate risk taking may be widely recognized, it is often challenging in practice. For example, the decline in new drugs and breakthrough therapeutics—despite increased R&D spending—has been attributed in part to lack of risk taking by pharmaceutical and biotech companies (Munos and Chin 2011, Krieger et al. 2019). Similar concerns exist in private sector areas including semiconductor manufacturing (Bloom et al. 2017) as well as in academic research. For example, Marks (2011) writes that "everyone familiar with

NIH operations knows that it is extremely difficult to obtain funding for groundbreaking, high-risk research."

To inform our understanding of risk taking in R&D, we focus on individual decision-makers who often serve as gatekeepers in selecting which ideas to invest in and commercialize. Such individuals include R&D managers, external review board members, and investor analysts. Prior work has examined how incentive structures affect risk taking. In this study, we highlight that even if incentive structures are aligned with risk taking, the way individuals respond to uncertainty may shape R&D investment decisions. We examine several potential barriers to risk taking that stem from individual decision-makers, including cognitive limitations in processing variance, a desire for diversification, loss aversion, sensitivity to ambiguous payoffs, and personal risk preferences.

We do so using a discrete choice experiment designed to uncover the role of uncertainty in shaping a manager's decision to fund R&D projects. Experimental participants were asked to rank a series of uncertain research projects. The choice scenarios were designed to be similar to the investment decisions R&D managers make in the real world. We instructed participants to assume the role of the director of the R&D group at a private company, and they were asked to choose their preferred research projects from a series of hypothetical proposals that had been judged and scored by an objective, third-party science advisory panel. Similar ratings procedures are commonly used as inputs to allocate internal funding at firms, attract external investors, and award government research grants.[1]

Compensation was determined by a competitive "tournament" structure. Participants were compensated for the performance of the R&D projects that they chose to fund relative to the choices of their peers in the experiment. The highest scoring participants received a substantially larger monetary reward than their peers. There was no penalty for low performance: the bottom 75% of scorers all

[1] Personal communication with Hanneke Schuitemaker, PhD, VP, Head Viral Vaccine Discovery and Translational Medicine, Janssen Vaccines and Prevention B.V., Johnson and Johnson, 3 February, 2020.

received the same compensation. Because there were large rewards for high performance and no downside risk for poor performance, the incentive structure disproportionately rewarded participants for choosing higher-variance (i.e. riskier) projects. That is, projects with greater disagreement in ratings (i.e., some high ratings and some low ratings) had a higher chance of success than projects with the same average rating but greater agreement (e.g., project ratings of 5,5,3,1,1 vs. 3,3,3,3,3 which has the same mean but lower variance).

The experiment was conducted with 290 Master of Business Administration (MBA) and Master of Finance (MFin) students at a major research university in a program focused on the intersection of business and technology. Many of these students come from an R&D background and will go on to work at investment firms or serve as managers making R&D decisions at companies in the health and technology sectors.

The experiment took place in two phases. In the first phase, 150 research subjects were asked to evaluate projects under three distinct sets of choice scenarios, described below. The second phase of the study was designed to explore some of the potential mechanisms driving our first phase results. It included three experimental treatments, where 140 subjects were randomized into (1) a replication of the baseline experiment from phase 1, (2) a version of the experiment created to test the role of loss aversion, or (3) a version designed to test for the effects of ambiguity. After the choice experiments, we elicited participants' personal risk preference parameters (both phases) and loss aversion parameters (second phase only).

In phase 1 of the experiment, the first set of choice scenarios assessed whether the incentives to choose high-variance projects in fact led to such choices among participants. Each participant was presented with ten scenarios where they were asked to rank four potential projects based on their preferences for funding. For each project, the participant was shown the individual scores from the advisory panel members and the average of those scores.

We find that most participants acted in an excessively risk-averse manner when selecting projects. Because of the competitive incentives, when offered two otherwise identical options, choosing a higher-variance project first-order stochastically dominated choosing a lower variance project. Despite this, participants were more likely than not to choose dominated projects. In other words, holding average score constant, participants were, on average, significantly *less* likely to choose a project as variance in ratings increased. Even in ideal cases where the participants were choosing between two projects that had identical mean scores, they chose the dominated project—the one with lower variance—three-quarters of the time. Because no risk aversion parameter can rationalize this behavior, we refer to the strong distaste for high variance projects exhibited by the participants as variance aversion.

Why did the participants behave this way? And what might a manager do to overcome or circumvent this behavior in employees overseeing R&D funding? Our subsequent choice scenarios, preference parameter elicitations, and second phase of the experiment examined potential mechanisms for variance aversion and tested interventions to address it.

We first examine heterogeneity in behavior across elicited preference parameters and participant demographics. The analysis reveals strong correlation between variance aversion and multiple dimensions of heterogeneity. Participants who were more risk loving and had more R&D experience exhibited a greater taste for variance. We also find less robust evidence of greater taste for variance among participants who faced more competition for a high reward payment (because of discreteness in the number of higher rewards issued) and participants in the MBA program (compared to the MFin program). In contrast, college coursework, elicited discount rates, and a measure that checked for understanding of the experimental tasks instructions were not correlated with variance preferences.

Second, we assess whether cognitive limitations in effort or attention drive the response to uncertainty. To do so, the second set of choice scenarios tested a simple informational intervention in which we additionally showed participants the variance of project scores. The variance is straightforward to infer from the individual scores in the first set of choice scenarios, so the second set of choice scenarios measures the effect of lowering the cognitive calculation costs and increasing the salience of score variance. Showing the variance backfired, leading subjects to engage in even more variance-averse behavior than in the first set of choice scenarios.

Third, we examine whether variance aversion can be explained by a desire to diversify risk. The third set of choice scenarios asked participants to construct portfolios of research projects rather than pick single projects to assess whether allowing for diverse portfolios would encourage more risk taking. We find, consistent with the first two sets of choice scenarios, that individuals continued to make variance-averse choices. By randomizing the budget that was provided to the subjects for each portfolio problem, we are also able to assess the effect of budgetary pressure. We find that tighter budgets exacerbated the problem, leading to more variance aversion.

Fourth, we examine answers that participants gave during a debriefing that followed the choice scenarios in which they were given free space to tell us why they had made their choices. The majority of participants stated that they treated the choice as a simple mean-variance tradeoff. These participants were, unsurprisingly, more likely to exhibit variance aversion in their project choices. A minority of participants were more sophisticated in their decision-making process and looked at individual project scores, explicitly discussed the idea that successful R&D requires one to embrace uncertainty, or gave other answers that suggested they were willing to be variance loving in this setting. Such answers were significantly, positively correlated with a greater preference for high-variance projects, being risk loving in the risk preference elicitation, and with prior work experience in the R&D sector.

Finally, in the second phase of the study, we investigated whether variance aversion is due in part to loss aversion—i.e., avoiding high variance projects that could result in losses—or ambiguity aversion. To address the former, we replicated the first choice scenario of the baseline, phase 1 experiment but removed all mentions of potential losses. Participant behavior was largely indistinguishable from the behavior in the baseline experiment. In addition, elicited loss aversion preferences were associated, if anything, with greater taste for variance, and considerable variance aversion remained after accounting for loss averse preferences.

To examine whether participants' response to variance in ratings partially reflects a distaste for ambiguity—in particular, ambiguity in how ratings map onto expected financial returns—we tested a version of the baseline experiments in which the ratings of projects were explicitly denominated in financial terms. The average subject was slightly more averse to choosing high-variance projects than in the baseline experiment, indicating that ambiguity aversion does not explain our core findings.

Our results suggest that explicit risk-taking incentives might not be enough to encourage optimal R&D within a firm, and that excessive risk aversion could lead to suboptimal R&D investment. To be more concrete, consider an example based on the empirical results that highlights the effect this behavior could have on breakthrough advances. In the experiments, subjects were shown hypothetical projects with ratings on a 1 to 5 scale. Consider two stylized examples of projects with identical average ratings but different variances. The first project is rated a 4 out of 5 by all seven panelists on the advisory committee. The second project is more divisive: receiving three ratings of 3, one rating of 4, and three ratings of 5. The first project has a variance of 0 while the second project has a variance of 1. Based on the findings from our experiment, subjects would be 6 percentage points less likely to choose the second project, despite the fact that the first project has no chance of producing an outcome of the highest possible quality and the second project has a 43% chance of doing so.

Our examination of participant characteristics points to potential solutions. Risk-loving participants performed better, on average, on the experimental tasks and chose projects more in line with optimal theory. Performance was hampered by treating the choices as "standard" portfolio optimization problems, an impulse that appears to have been tempered, in part, by training and work experience. These findings suggest that firms aiming to encourage more innovation may want to include the risk preferences of those workers in charge of research and development as a factor in their hiring and promotion decisions as well as emphasize the convex nature of returns to R&D as part of their training.

2. Literature

Theoretical models of optimal R&D argue that both firms and public funders should invest in high-variance research projects. An important early contribution to this literature, Dasgupta and Maskin (1987) argue that from the perspectives of the individual scientist, competitive firms, and society as a whole, the spoils from R&D are skewed toward novel, high-quality discoveries. Given the disproportionate benefits from producing the highest quality discoveries, investing in riskier R&D projects is optimal from both a social and private perspective.[2]

More recently, in a theoretical setting similar to our experiment, Tishler (2008) shows that competition among firms or research groups should lead them to adopt high-variance R&D portfolios. Given two projects with the same expected discovery quality, a firm should choose the higher variance project to capture convex returns. The incentives in our experiment are meant to replicate the competitive compensation scheme modeled by Tishler (2008) and observed in real-world R&D. Participants were paid substantially more if their research projects and portfolios performed well relative to the other participants.

Despite models showing that optimal R&D entails investment in high variance research, many observers have documented low rates of risk-taking by agencies that disburse research funds (Azoulay, Graff Zivin, and Manso 2011; Marks 2011) and firms that conduct R&D (Munos and Chin 2011). These papers leave open the question of how the preferences of individual decision-makers help drive suboptimal risk taking in R&D even when explicit incentives for innovation are in place.

A separate strand of research highlights the potential link between individual preferences and innovation. Prior work demonstrates that scientists have important

[2] Theoretical work in this area continues to develop. For instance, Hopenhayn and Squintani (2021) study the role that R&D fads can play in diverting researchers away from potentially more valuable but riskier R&D areas.

nonpecuniary motivations (Dasgupta and David, 1994; Merton. 1973). In particular, scientists are willing to accept a lower salary to work in organizations that allow them to pursue independent research (Stern, 2004); and such preferences are positively correlated with innovative performance, as measured by patent applications (Sauermann and Cohen, 2010). Related work across a range of industries finds that willingness to take risks is positively correlated with assessments of innovative creativity in the workplace at both the organizational (Amabile et al., 2017) and individual levels (Madjar et al., 2011). Bringing these strands together, recent evidence suggests that less risk-averse individuals generate more novel inventions by pursuing riskier innovation strategies (Graff Zivin and Lyons, 2020).

Finally, related to our suggestion that firms may want to take into account the risk preferences of their R&D managers, prior work explores the relationship between preferences and selection into innovative sectors. A large literature examines the relationship between risk preferences and both selection and performance of entrepreneurs with mixed findings (Astebro et al. 2014 provide a review). Goel and Thakor (2008) show theoretically that firms might value overconfident Chief Executive Officers (CEOs) if that overconfidence helps counteract risk aversion. Overconfident CEOs are also more likely to invest in risky projects, leading to higher innovation if the firm is in an innovative sector (Hirshleifer, Low, and Teoh 2012). Related work by Kagan, Leider, and Lovejoy (2019) makes a similar point about equity contracts in entrepreneurial teams. Traditionally, researchers have argued that contract structure matters for team performance, but Kagan, Leider, and Lovejoy (2019) show that individual preferences determine which types of contracts are taken up by workers. This selection confounds estimates of the effects of contract type on firm performance and means that individuals in charge of hiring should pay close attention to the preferences of potential employees.

3. Experimental Design

3.1. Experimental Setup

The experiments were implemented among master's degree students enrolled in a program focused on the intersection of business and technology. The typical student has three to four years of work experience with a background either in research-intensive firms in science and technology sectors, or in finance, banking and economics. All have formal academic training in assessing risky tradeoffs and portfolio analysis. Many of the graduates will work for investment firms or will assume management positions within research divisions of corporations across a wide spectrum of science and technology spaces. Thus, studying the decisions of this group is particularly germane for our understanding of R&D investment choices within the private sector. Summary statistics for the study participants are discussed in Section 3.1.4.

Participants were asked to assume the role of the head of a research division at an organization considering whether to fund project proposals based on ratings from a third-party scientific advisory panel (see Appendix A for the instructions). They were then tasked with ranking research projects in a series of choice scenarios. Participants were allowed to take as much time as they wanted to complete the experiment. Empirically, the average participant spent 58 minutes on the experiment.

The experiment was conducted in two phases. Participants in the baseline phase 1 experiment were asked to rank research projects in three sets of choice scenarios, described below. In the second phase, participants were randomized into one of three distinct experimental treatments to help elucidate the mechanisms driving our baseline results. The two phases of the experiment were distinct. Participants were not explicitly randomized across the first and second phases, but they were drawn from a demographically similar subject pool.

3.1.1. Baseline, Phase 1 Experiment

In the baseline experiment, each participant took part in three sets of choice scenarios. In the first set, they were presented with a list of four research projects rated by seven reviewers (on a scale of 1 to 5) along with the average reviewer score for each of the projects. The subjects ranked projects based on the likelihood that they would fund them. The ranking was carried out by first choosing the most and least preferred project, then by ranking the remaining two projects. This process was repeated for ten different groups of research projects, with each group characterized by different reviewer score profiles.

In the second set of choice scenarios, the same procedure was repeated for ten more groups of projects, but the subjects were also shown the variance of reviewer scores. Because participants could calculate the variance themselves based on the individual ratings, the second set of choice scenarios did not provide more information than the first one. It was designed to address concerns about cognitive calculation costs, computation errors or misunderstandings, but also made that feature more prominent. An example of the initial project choice screen is shown in the Appendix.

The third set of choice scenarios presented each subject with eight portfolio choices. For each portfolio choice, subjects were presented with ten different projects rated by seven reviewers. As in the second set of scenarios, each project was rated by seven reviewers, and participants saw the individual ratings as well as each project's average rating and variance of ratings. In addition, each project was assigned a cost of either \$1, \$4, \$7, or \$10 million. Subjects were provided a randomized budget that they could use to fund the projects in the portfolio. One of eight possible budgets (\$12, \$13, \$14, \$15, \$16, \$17, \$18, or \$19 million) was chosen without replacement for each portfolio choice, so each subject saw the full set of possible budgets. Participants could select and deselect projects from their portfolio. We displayed the remaining funds in their budget for their chosen portfolio until they finalized their choices. An example portfolio choice question is shown in the Appendix.

At the end of the experiment—after participants made their decisions but before learning of their performance—subjects completed a debriefing about why they had made their decisions as well as a survey that included questions about demographics and their risk preferences. We utilized a multiple price list to elicit risk preferences, a standard technique in the experimental economics literature (Charness, Gneezy, and Imas 2013). Subjects were provided with a list comparing a guaranteed payment to gambles with progressively lower variance and expected values. The subjects were then asked to make hypothetical choices between the gambles and the guaranteed payment. Based on their choices, we classified participants as risk-averse, risk-neutral, or risk-loving, and we calculated each subject's coefficient of relative risk aversion (details on this calculation can be found in Appendix Section A.5).

3.1.2. Phase II Experiment

In order to tease out the mechanisms underlying the results from our baseline experiment, we conducted a second round of experiments with a new cohort of students drawn from the same academic program. Due to the COVID-19 pandemic, these experiments were administered to students taking their classes remotely. In this second phase, participants were randomized into one of three distinct experimental arms: (1) one that replicated the first and second sets of choice scenarios in the baseline experiment; (2) an identical experiment that removed all loss-framing language; and (3) an identical experiment that replaced reviewer rating scores with an objective payoff matrix. In addition, we also gathered information on loss aversion preferences from the subjects using the elicitation from Imas et al. (2017), based on the design of Abdellaoui et al. (2008). Details on the calculation of the loss aversion parameter can be found in Appendix Section A.6.

3.1.2.1 Arm 1: Replication Experiment

The replication experiment was designed to create a bridge between the two experimental phases by allowing for a direct comparison between the behavior of subjects in each. The randomization across experiments within the second phase also ensures that comparisons across these follow up experiments can be interpreted in a causal framework.

3.1.2.2. Arm 2: Experiment Without Loss Framing

Participants randomized into the experiment without loss framing were presented with the same choice scenarios as in the replication experiment (choice scenario sets 1 and 2). The only difference was in the instructions that introduced the experiment and choice scenarios. The no-loss framing experiment removed all language that stated or implied that losses were possible when investing in R&D projects. The complete instructions are available in Appendix Section A.2. In all of the experiments, subjects could only gain money, so the change in the instructions only affected the framing of the experiment rather than the true, underlying incentives.

3.1.2.3. Arm 3: Experiment with Objective Returns

Participants randomized into the objective returns version of the experiment also engaged in the same choice scenarios as in the replication experiment but with a different framing. In this case, the possible value of a project was presented, not in terms of scores generated by an outside scientific advisory panel, but as objective financial returns. The returns were consistent with the no-loss experiment in that all of the projects were shown to have strictly positive net returns. The returns were displayed in units of millions of dollars. The values for the returns had the same distribution as the advisory panel ratings from the other three experiments. The instructions are shown in Appendix Section A.3, and example choice scenarios can be seen in Figure S6.

3.1.3. Incentives

By design, participants in all experimental treatments and phases were incentivized to choose riskier (i.e., higher-variance) projects. At the beginning of the experiment, subjects were told that they would receive a score based on the projects and portfolios that they chose. The realized value for each project was generated by an independent draw from a normal distribution with mean and variance of the reviewer scores. To maintain incentive compatibility throughout the ranking, final scores were affected by the full ranking of all project choices that the subject made. For each project choice question in the first and second experiments, the final score for each individual project was equal to the full realized value for the first-choice project, 0.75 times the value drawn for the second-choice project, 0.5 times the value drawn for the third-choice project, and 0.25 times the value drawn for the fourth-choice project. The value of the portfolio questions was similarly drawn from a normal distribution with mean equal to the sum of each individual project's mean weighted by cost; and with variance equal to the sum of each project's variance weighted by cost. The project and portfolio scores (where applicable) were summed to create the total score for the participant.

We then publicly awarded prizes to the top performers in each session: the top 10-25% of scores received $25 and the top 10% of subjects received $100. All subjects received a $15 participation fee. Because we offered large rewards for performance in the right tail of the distribution and offered no additional rewards for performance in the bottom three-fourths of the distribution, there was a large potential upside and no downside risk from choosing higher variance projects. Thus, subjects had a strong incentive to choose higher variance projects to maximize their probability of winning the largest prizes. For two projects with the same average rating, choosing the higher variance project first-order stochastically dominated choosing a lower variance project, meaning that all subjects, regardless of risk preferences, should have chosen higher variance projects on the margin. We assessed participant understanding of the incentives through actions taken during the portfolio choice section of the experiment and by debriefing the participants after they had completed all project ranking scenarios. We report results for these two assessments below.

3.1.4. Recruitment, Sample Size, and Sample Summary Statistics

All baseline experimental sessions were implemented during regularly scheduled class sessions of the MBA and MFin programs. Participants in the other three experiments were also recruited through their classes. Each professor chose whether to field the experiment during class time or outside of class. The randomization was stratified by class for all experiments.

All students in the class were eligible to take part, and participation was voluntary. After obtaining informed consent from all participants, they completed the experiment on their own computers. On average, the experiment took subjects about 1 hour to complete. The baseline experimental sessions were conducted in-person, and subjects were paid at the end of the session. The three other experiments were conducted remotely, and subjects were paid after all participants in their class had completed the experiment—at most one week after the initial distribution of the experiment.

[Table 1 about here]

For the phase 1 experiment, a total of 196 students were recruited in six experimental sessions. One subject started the experiment but had to leave before completing it, and four subjects failed to provide us with answers sufficient to calculate risk preferences. They were excluded from the analysis. In the first session, the order of projects was not randomized due to a coding error, so we exclude the 36 students from that session in the baseline analysis (results with all participants are shown in Appendix Table B4). Five additional participants exhibited multiple switching on the risk preference elicitation and were also excluded from the base sample. All results are robust to the inclusion or exclusion of these participants. The final sample therefore contains 150 subjects.

Each subject faced 10 choice scenarios in choice scenario sets 1 and 2. Each scenario involved choosing between 9 potential options, yielding 13,500 total observations for each set. The options in choice scenario set 3 varied by budget, which was randomized across subject. The average subject had 1,399 options, leading to a total sample size for set 3 of 219,310 observations.[3] The standard errors for all analyses are clustered at the subject level to account for correlation within subject across choice scenario.

For the three additional experimental arms, we recruited 140 subjects across 4 sessions. We randomized within session at the individual level: 46 subjects

---

[3] More specifically, in both choice scenario sets 1 and 2, subjects engaged in 10 choice scenarios. For each choice scenario, they first selected their top and bottom choice from a set of 4 options. They then selected their second favorite choice from the remaining two options. We model this as three choice occasions per scenario, so there were four observations for the first choice occasion, three options in the second occasion, and two in the final occasion. For choice scenario set 3, the set of feasible portfolios determined the choice set faced by the subject in each of the 8 choice scenarios. Feasible portfolios were those that had total cost less than or equal to the budget. Because budget was randomized, the size of the choice set varied by subject and choice scenario.

completed experimental arm 1 (phase 1 replication), 47 subjects completed arm 2 (no loss framing), and 47 subjects completed arm 3 (objective payoffs). Each of the additional experiments involved two sets of 10 choice scenarios. As in the first and second set of choice scenarios from the phase 1 experiment, each scenario involved ranking four potential projects.

Summary statistics for all study participants in the main estimation sample are shown in Table 1. The statistics show that the typical participant has multiple years of work experience, and across the two phases of the experiment, 41% of the participants reported R&D sector work experience. The average participant was risk averse, and in phase 2, the typical subject was also loss averse according to the preference elicitation. The summary statistics provide an initial indication that there are no gross imbalances across observables, which is further validated by formal tests of balance across the treatment arms in phase 2, as shown in Appendix Table B3.

3.2. Design of the Discrete Choice Experiments

The design for the choice scenarios presented to the subjects builds on models of random utility theory to estimate discrete choice models using decisions from discrete choice experiments.[4] These designs allow discrete choice models to be applied to situations where individuals are making choices that are not currently observed in real markets. We followed this tradition by developing experiments to simulate hypothetical but potentially real proposals and projects and asking individuals to evaluate them and make choices. The design allows us to estimate statistical models using the experimental choices as data to approximate the individuals' choice processes.

Discrete choice experiments (DCEs) are based on traditional experimental design concepts for fractional factorial designs widely used in applied statistical

[4] Random utility theory was developed by Thurstone (1927) and underlies applications of the Method of Paired Comparisons (e.g., David 1988). Models for multiple choices were proposed by Luce (1959) and random utility theory was extended to statistical models for multiple discrete choices by McFadden (1974). Louviere and Woodworth (1983) proposed discrete choice experimental designs consistent with random utility theory.

work.[5] To construct the choice sets in our experiment, we first enumerated all possible combinations of seven hypothetical raters using a 5-category rating scale. We then calculated the mean and variance of each combination and sorted them from highest to lowest and identified 16 orthogonal combinations of means and associated variances. Using these combinations, we constructed the choice sets for the twenty individual project ranking questions and then constructed the choice sets for the eight portfolio questions.

To construct the choice sets for the project ranking task, we used a Balanced Incomplete Block Design (BIBD)—see Louviere, Flynn, and Marley (2015)—to create 20 sets of four project proposals. Each proposal was described by seven ratings. The mean and the variance of these ratings were the two primary attributes associated with each proposal. The 20 sets of projects were divided into two groups of 10 to create sets of choice scenarios.

To ensure that the models we estimated were not saturated and to enhance the degrees of freedom, we made two versions of the DCE by randomly rearranging the original DCE attributes (mean and variance) and again making 20 sets of four proposals using the same BIBD. Again, these 20 sets of proposals were divided into two subsets of 10.

We then randomly blocked each of the two versions of the DCE—Version I and Version II—and the two subsets within version—Subset A and Subset B. This produced four treatment groups: Version I.AB, Version I.BA, Version II.AB and Version II.BA – where the first letter refers to the subset used in choice set 1 and the second letter refers to the subset used in choice set 2 . By showing identical choice scenarios to different participants in the same treatment group, we can identify the effect of changes in attributes (score mean and variance) conditional on choice

[5] Basically, a DCE is a sparse, incomplete contingency (crosstab) table, one side of which represents the observed discrete choice options presented in the DCE. Thus, DCEs use experimental designs from the factorial family of combinatorics designs to create sets of choice options called choice sets. The experimental design provides the basis for creating the choice options and the choice sets to which they are assigned.

scenario and participant fixed or random effects. Balance across the experiment versions is shown in Appendix Table B1.

To construct the choice sets for the portfolio selection task, we used the complement of the BIBD used to construct the choice sets for the project ranking task (the complement contains all combinations not included in the first BIBD). Costs were also added as an additional attribute for the proposals, with costs randomly assigned following the same procedure for mean and variance used in the project selection tasks. Costs were blocked so that subjects would routinely face choices between two projects with identical expected value (same cost and same mean) but different variance. We exploit this feature to study risk taking behavior as a function of portfolio budget in the results section. We arrayed the 16 combinations into 16 sets of ten proposals. We then created four blocks of eight choice sets using the method discussed above to make two versions of the DCE and two subsets within each DCE. We randomly assigned each block of eight portfolio selection questions—Block 1, Block 2, Block 3, Block 4—to one of the four experimental versions discussed above (i.e., Version I.AB.1, Version I.BA.2, Version II.AB.3, Version II.BA.4).

In the Phase I experiment, we randomized individuals to one of the four versions, stratified on session. In the Phase 2 experiment, within each of the three experimental arms, participants were randomly assigned to one of the four versions, stratified on session. The order in which projects were presented within each version was also randomized across sessions. The experimental instrument was programed and delivered using the Sawtooth Software platform.

4. Empirical Specification

We estimate the relationship between project attributes and subject choice using a generalized multinomial logit (G-MNL) model. The estimating equation models the probability that subject i chose alternative j in choice scenario t as

$$\Pr(\text{choice}_{it} = j|\beta_i) = \frac{\exp(\beta_i' x_{itj})}{\sum_{k=1}^{J} \exp(\beta_i' x_{itk})} \tag{1}$$

where x_{itj} is a vector of attributes (mean and variance of the projects in the baseline models and interactions with subject demographics in the models exploring heterogeneity); and β_i is the vector of individual-specific coefficients on the vector of attributes. These coefficients can be interpreted as utility weights placed on the attributes by each individual and are defined by

$$\beta_i = \sigma_i \beta + \eta_i \tag{2}$$

The coefficients in Equation (2) are a vector β that is constant across individuals and measures the average utility weights across the sample for the different variables in x; a single parameter for the scale of the individual-level idiosyncratic error σ_i, which captures overall scaling of an individual's tastes; and, a random vector η_i distributed multivariate normal with mean 0 and variance-covariance matrix Σ, which captures taste heterogeneity. We follow Fiebig et al. (2010) and assume that σ_i is distributed lognormal with mean $\bar{\sigma} + \theta z_i$ and standard deviation τ. The parameter $\bar{\sigma}$ is a normalizing constant and z_i is a vector of subject characteristics that explain differences in σ_i across individuals. In our application, we focus on project and portfolio attributes and limit our attention to subject indicators in z_i.[6]

The workhorse model in applied microeconomic studies of discrete choice is the multinomial (conditional) logit model. We prefer estimates based on the more flexible G-MNL model due to the strong restrictions imposed by the standard

[6] This is a G-MNL type I model in the terminology of Fiebig et al. (2010) because the standard deviation of η_i is assumed to be independent of the scaling of β. We make this assumption to speed convergence of the model and based on analyses that showed this constraint led to superior model fit relative to the other choices of constraints commonly used in the literature (including not constraining the relationship between the standard deviation of η_i and the scaling of β). These alternative results are available upon request.

conditional logit. The restrictions effectively rule out many kinds of heterogeneity that are of potential interest when studying behavior and which can lead to confounding. The random coefficient (mixed) logit model relaxes assumptions about preference heterogeneity but restricts all agents to have their error component drawn from the same distribution, such that differences in these "scale" parameters could easily be misidentified as differences in preference parameters. This occurs because preference and scale parameters are not separately identified in choice models, behavior is governed by their ratio, so that an upward shift in the scale parameter must shift the magnitude of the vector of preference parameters upward to maintain the same ratio. The GMNL model nests both the mixed logit model that allows heterogeneity in preference parameters and models that allow for scale heterogeneity by allowing both to vary in a reasonably flexible, but in a statistically identified way.

For ease of interpretation, however, we also present corresponding conditional logit models for the main results in the paper. The analysis of the third set of choice scenarios (portfolio choices) is also carried out using standard conditional logit and fixed effects linear regression specifications. We estimate these specifications because we are interested in the effect of budget constraints on choice, and budget was randomly varied within subject, across choice scenario. As such, we rely on between-subject comparisons that preclude the use of individual and choice scenario-specific heterogeneity parameters. We verify that the budget randomization was balanced on observable characteristics in Appendix Table B2.

5. Results

5.1. Initial Evidence for Variance Aversion

Our primary question of interest is whether subjects responded to the incentives we gave them by choosing higher variance projects when faced with a choice between two otherwise similar research proposals. We formally test this by estimating statistical models that control for the average score, allowing us to isolate the effect of variance on the likelihood that a subject would choose a given project. As discussed above, the repeated, within-subject sampling of the experimental design allows us to estimate generalized multinomial logit (G-MNL) models that further account for latent subject-specific heterogeneity while relaxing strong assumptions that underly the estimation of conditional logit models.

Table 2 shows results from the first set of choice scenarios in the phase 1 experiment. In all columns, the dependent variable is an indicator equal to one if the subject chose the project.[7] The explanatory variables are the project mean and variance, and they are standardized to have an average value of zero and standard deviation of one.

[Table 2 about here]

The coefficients in the top portion of the table (labelled "Average Utility Weight") are the estimates of the utility weight that subjects placed on average project score and score variance (the β terms in Equation 2). The second section of the table (labelled "Utility Weight Heterogeneity") reports estimates of the heterogeneity in preference (the σ_i terms in Equation 2). The third section of the table reports the estimate of the standard deviation of individual-level scale heterogeneity, which we estimate to be small in this case.

The results show that, on average, participants had strong preference for projects with higher average scores and lower score variance. This behavior is at odds with the incentives the participants faced and provides our first evidence of variance aversion. The result holds both in the G-MNL model, a traditional conditional logit

[7] For the project choice questions, subjects ranked all projects by first choosing their first and fourth favorite projects, then choosing their second favorite project from the remaining two choices. In the analysis in Table 2, we treat these decisions as three separate choice scenarios. In the first scenario, the choice set is all four projects, and the subject's choice is their top ranked project. In the second scenario, the choice set is the three remaining projects after excluding the top ranked project and the choice is their second ranked project. The third scenario's choice set is the remaining two projects, and the choice is the third ranked project. Results using just the first choice (of the most preferred project) are similar and available from the authors. Rank-based multinomial logit results are shown in Appendix Table B6.

model (Column 2), and across each choice occasion when analyzed with a rank-based multinomial logit model (Appendix Table B6).

The standard deviation of those preferences is also large, indicating that there was substantial individual heterogeneity that would be overlooked by more traditional conditional logit models. The full distributions of estimated preference parameters are shown in Appendix Figure C1 and demonstrate that 1/6th of the sample exhibited variance loving (rather than variance averse) preferences. It is this heterogeneity—the fact that the average subject, despite the incentives, exhibited variance aversion while some subjects behaved in accordance with the incentives to seek variance—that we look to explain in the next section.

The marginal effect of an increase in average project score is 0.23 for the G-MNL estimates from Column 1 and 0.34 for the conditional logit estimates in Column 2. The marginal effect of an increase in project score variance is -0.032 for the G-MNL estimates and -0.06 for the conditional logit estimates (with all estimates significant at the 1% level). For both estimation methods, higher average scores raised the probability that a participant would select that project while increases in score variance decreased the probability. The difference in marginal effects shows that the conditional logit results overstate both effects, however, by not accounting for subject-specific heterogeneity.

[Figure 1 about here]

Figure 1 summarizes the G-MNL results from Table 2 for a range of different project attributes. The figure shows the estimated probability that the average subject would choose a project with a high, medium, or low average score,[8] and score variances that span the set shown to study participants. The figure again shows that

[8] A score of 3 out of 5 (corresponding to the 25th percentile of scores shown to subjects), a score of 3.5 out of 5 (the median), and a score of 4 out of 5 (the 75th percentile) respectively.

subjects strongly preferred to choose projects with higher mean scores and lower score variance. The figure also reveals that the effect of variance was stronger for projects with a higher mean. For a project with a high mean score, an increase in variance by 1 reduced the probability that a typical subject chose that project by 11.9 percentage points (95% confidence interval of 9.1 to 14.7). In contrast, for a low mean score project, a 1 unit increase in variance decreased the probability of selection by only 2.2 percentage points (95% confidence interval of -1.0 to 5.3).

In other words, participants were particularly variance averse when choosing between high mean projects—so much so that they were willing to frequently forgo selection of projects with the very highest average scores if it meant reducing their exposure to variance. When assessing projects with middling average scores, participants were less reluctant to choose high variance projects, potentially because neither project looked particularly attractive. We report results from a debriefing with participants that shed further light on this behavior in the next section. Irrespective of the underlying drivers of this behavior, it has potentially important implications for research project selection if otherwise stellar research projects are being rejected due to disagreement in reviewer scores while more mediocre, if less controversial, projects are being favored instead.

As described in Section 3.1, we replicated the phase 1 experiment with a new set of subjects. The participants were drawn from a similar population as those who completed the baseline experiment. We find that the participants in the replication experiment behaved similarly to those in the baseline: both groups preferred projects with higher mean scores and lower score variance, on average. There is no significant difference in the variance preferences across the two groups, as shown in Appendix Table B5.

5.2. Mechanisms Underlying Project Choice and Variance Aversion

We now investigate the factors that might explain why subjects preferred projects with lower variance. We first assess whether a variety of subject characteristics, demographics, and other covariates do or do not correlate with variance aversion in the first set of choice scenarios from the phase 1 experiment. We then present results from four additional sets of choice scenarios and experimental interventions to determine whether differences in the experiment itself could lead participants to choose higher variance projects.

5.2.1 Heterogeneity in Variance Aversion

The results from the G-MNL estimates in Table 2 showed that there was substantial subject-specific heterogeneity in variance preferences. In Table 3 and Figure 2, we examine correlation between variance aversion and observable, subject-level covariates. Table 3 presents G-MNL estimates that include interactions between project attributes and 8 dimensions of heterogeneity. Figure 2 gives the marginal effect of project score variance along these different dimensions. The measures are all based on observational data rather than explicit randomization, so we emphasize that the results are suggestive and correlational.

[Table 3 about here]

The figure shows a wide range in the correlation between observables and variance preferences. On one side, there is essentially no correlation between variance preferences and the number of college math courses taken, the number of decision science courses taken, and the elicited discount rate. The fourth measure, titled budget spending, is an indicator equal to 1 if the participant left any of their allotted budget unspent during the third set of choice scenarios. The incentives were to spend all of the budget, so this indicator is a simple measure of whether the participants understood the instructions and incentives they faced. Although many participants left some budget unspent, this behavior also does not strongly correlate with variance aversion.

[Figure 2 about here]

In contrast, the next four measures do correlate with variance preferences. Participants who were risk averse were substantially and significantly more variance averse than participants who were risk loving. A risk averse subject responded to a 1 unit increase in project score variance by reducing the probability of choosing that project by 6 percentage points. A risk loving participant, on the other hand, reduced their probability by only 2 percentage points.[9]

Participants with prior work experience in the R&D sector (30% of this sample) were also substantially less variance averse. The difference in marginal effects of variance between participants with R&D experience versus those without was 4.8 percentage points (95% confidence interval of 2.3 to 7.3). This result also holds when simultaneously estimating the effect of all of these dimensions of heterogeneity (see Figure C2), suggesting that it is not simply selection into prior work experience due to coursework, risk or time preferences, or other factors.

Similarly, participants who were pursuing an MBA were, on average, variance loving while those pursing an MFin were substantially more variance averse, potentially because the MFin students were more likely to treat the choice scenarios as typical financial portfolio problems. The difference in marginal effects of variance between the two groups was 6.3 percentage points (95% confidence interval of 3.7 to 8.9).[10]

The strength of incentives faced by participants also affected preferences. Based on the number of other participants in a session, the probability that a subject would receive an extra $25, for example, varied between 20% and 25%. For a session

[9] Although, as noted above, these results could be driven by true risk preferences or correlates including cognitive ability (Frederick, 2005; Benjamin, Brown, and Shapiro, 2013).

[10] Figure C2 shows that this effect is attenuated (though still significant at the 5% level) when all dimensions of heterogeneity are included simultaneously. See the discussion at the end of this section.

with 36 people, for instance, exactly nine people would receive a $25 or $100 award. In a session with 35 people, however, only 8 (or 22.9%) would receive a larger reward. Participants in the more competitive (lower probability) sessions were substantially less variance averse.[11] This result, however, is not robust to simultaneously including all dimensions of heterogeneity simultaneously, again suggesting that the incentive structure faced by the participants played a weaker role in determining behavior.

Appendix Table B9 and Figure C2 show the correlation of these measures when included simultaneously in the same regression. As discussed above, elicited risk aversion and prior R&D experience continue to be two of the strongest predictors of variance aversion, suggesting that these two measures (or their correlates) independently predict the preferences of participants. The discount rate also becomes a strong predictor, with higher discount rates being associated with lower variance aversion. Appendix Table B9 also shows that these dimensions of observable heterogeneity explain a substantial fraction (74%) of the subject-level heterogeneity estimated in the baseline G-MNL regression reported in Table 2.

5.2.2 Assessing Cognitive Limitations in Effort or Attention

Second, we assess whether cognitive limitations in effort or attention help explain the response to uncertainty. The second set of choice scenarios tested a simple informational intervention in which, in addition to the individual project scores and mean of the scores, participants were shown the variance of project scores. The variance is straightforward to infer from the individual scores in the first

[11] The marginal effect of score variance for participants with a 20% chance of winning was 5 percentage points more negative than the effect for other participants (95% confidence interval of 0.6 to 8.8).

set of choice scenarios, so the second set of choice scenarios measures the effect of lowering the cognitive calculation costs and increasing the salience of score variance.

[Table 4 about here]

Table 4 compares behavior between the first and second choice scenarios. The results show that participants were substantially and significantly more variance averse when shown the variance. The G-MNL model in Column 2 that nests responses from both sets of choice scenarios indicates that participants were twice as variance averse when the variance was displayed. The marginal effect of project variance in choice scenario set 2 is -0.061, versus -0.032 in set 1. The difference is significant at the 1% level. That the act of reporting variance, which should have made it easier for subjects to respond to the incentives of the contest, was associated with more risk-averse choices is quite surprising. We note, however, that participants were not randomized into being treated with choice scenario set 1 or 2 (all participants took part in both), so although the estimates come from a within-subject comparison, they still fall short of the ideal experiment and should be treated with some circumspection.

5.2.3 Effect of Risk Diversification and Budgetary Pressure

Next, we examine behavior in a portfolio choice setting—the third set of choice scenarios—to determine whether the ability to diversify the set of projects affects risk taking. Consistent with the first two sets of choice scenarios, individuals continued to make variance-averse choices. We can see this result from a simple analysis of choices over similar portfolios. For a given budget, subjects could often construct two portfolios with identical expected values but different variances. For instance, with a budget of $12 million, there were two different portfolios with the highest possible mean score (53.48), one with higher variance (15.14) and one with lower variance (8.44). Most subjects (71%) chose the lower variance project when faced with this choice. Appendix Table B7 reports regression results that further corroborate this analysis. The point estimates suggest that individuals were slightly more variance averse when choosing portfolios than in the first set of choice scenarios, although the difference is not statistically significant at conventional levels.

[Figure 3 about here]

By randomizing the budget given to participants in each choice scenario, we are also able to assess the effect of budgetary pressure. Figure 3 shows that participants were more variance averse when they faced a smaller budget. For ease of presentation, we show the relationship for budgets that were less than the average (between $12 and $15 million) and for budgets that were greater than average (between $16 and $19 million). The difference in slope between the two fitted lines indicates that the effect of variance was smaller for choices made with larger budgets. This difference is statistically significant at the 1% level and shows that for two otherwise similar portfolios (same mean score, same cost), subjects were roughly twice as reluctant to choose a portfolio with a higher variance if they had a smaller budget than if they had a larger budget.[12]

5.2.4 What Participants Said About Their Own Behavior

After the third set of choice scenarios, participants were asked an open-ended question about their decision-making process.[13] The answers can help shed light on the results and heterogeneity reported above. The most common answers indicated that the participant was treating the choices as a typical mean-variance tradeoff, with a typical respondent saying that they "[f]irst sought lowest variance with the highest average value." More than 70% of participants gave an answer along these lines (see Appendix Table B11).

One-fifth of participants indicated that they were variance loving. For example, one respondent wrote, "I went with the highest number and then often the highest

[12] Regression estimates corresponding to the differences shown in the figure are in Appendix Table B7.

[13] The exact question was, "Briefly describe how you went about deciding which projects you put in the portfolios you wanted to fund."

number with the highest variance. Some people can have a difference of opinion." Some were even more explicit about how the R&D process is improved by judicious risk taking: "Consensus was not applied to my choices. When the respondents had a unanimous consensus, I took that to mean that the work was not groundbreaking. Therefore, I chose proposals that would challenge the experts and thus may drive at better outcomes, positive or negative." Most participants who indicated that they were variance loving also said that they looked at individual project scores and not just the mean or variance. Finally, participants who stated in their answers that they sought out variance were significantly more likely to exhibit variance loving preferences in the first and second sets of choice scenarios, corroborating the self-reported statements.

Reassuringly, only about 1% of participants gave answers that were directly contradictory to the experimental instructions. A larger fraction—12% of participants—expressed some form of loss aversion, indicating that they left budget unspent out of concern that the remaining projects would result in losses or saying that projects with a high proportion of low scores might generate losses.[14] Loss aversion is a potential alternative explanation for the behavior we observe, and we designed the phase 2 experiment to test for the effect of loss aversion directly.

5.2.4 Phase 2 Experiment: Assessing Loss and Ambiguity Aversion

The phase 2 experiment consisted of a replication of the first set of choice scenarios from phase 1 (treatment arm 1), a replication of choice scenario set 1 from the phase 1 experiment but without any mention of potential losses (treatment arm 2), and a version of treatment arm 2 that further presented the project attributes as objective financial returns rather than review scores (treatment arm 3). The similar

[14] For example, one subject wrote, "Firstly, I will rank by the average score from high to low and prefer those with a significantly small variance. After that, I will check whether there is a possibility of extreme loss in this project. If so, I would like to not fund the project."

results between the phase 1 experiment and the arm 1 replication were reported above in Section 5.1.

The second and third arms allow us to assess the effects of loss aversion and ambiguity aversion. Table 5 compares behavior across the arms and shows that preferences were largely the same in all arms. Participants had a statistically significantly lower preference for average project score in arm 2 compared to arm 1, but there was not a significant difference in variance preferences.

[Table 5 about here]

In addition, Appendix Table B10 adds elicited loss aversion and risk aversion parameters to the estimation model and shows that a larger loss aversion preference parameter is associated, if anything, with greater taste for variance. Variance aversion is also relatively unaffected by the inclusion of controls for elicited loss aversion. Together, these results indicate that loss aversion is not driving the variance aversion we find.

Arm 3 tested for the effect of ambiguity aversion. Table 5 shows that presenting the project attributes using explicit financial returns led participants to be more variance averse, if anything.[15] This result is consistent with the debriefing results and heterogeneity analysis showing that participants who treated their choices as financial portfolio problems sought to maximize mean while minimizing variance.

6. Discussion

Anemic research pipelines and the apparent slowdown of paradigm-shifting discoveries over the past quarter century have drawn considerable ire from both the research and investor communities. This has led to episodic concerns of policymakers regarding national scientific competitiveness and its role in shaping economic

[15] The difference is statistically significant for the G-MNL model but not for the conditional logit in Column 2. The point estimates from both models indicate substantially greater variance aversion on average.

growth. If a small number of breakthrough research projects are responsible for a disproportionate amount of scientific progress, then research funders should target projects with greater uncertainty in order to have any chance of hitting upon rare but important results (Lotka 1926, Helpman 1998).

In our setting, participants did not behave this way. They consistently chose lower variance projects despite incentives that expressly rewarded risk taking and that mirrored the risk-reward trade-off laid out above. The results suggest that one possible reason for the lack of scientific breakthroughs is the risk appetite of R&D managers. We found that subjects routinely made dominated decisions—choosing lower variance projects even when projects with higher variance and the same mean score were available. These decisions caused excessively risk-averse subjects to leave money on the table. Comparing participants by the variance of the projects they actually chose, those subjects in the top quartile of variance were three times more likely to earn a reward than subjects in the bottom quartile.[16]

To highlight the effect of variance aversion and personal preferences, consider a typical project in the experiment which had an average rating of 3.5 out of 5. The variance of ratings ranged between 0.28 and 2.28. A participant with an elicited coefficient of relative risk aversion of 1.185 (the 25th percentile in the sample) would be just over 2 percentage points less likely to choose the project with higher variance. In contrast, a participant with a preference parameter of 1.45 (the 75th percentile) would be more than 5 percentage points less likely to choose the project. On the other extreme, a participant with an elicited parameter of 0.93 (the 10th percentile) would be indifferent to project variance in this case.

To put this into context, we can construct a stylized example of how these differences could affect broader R&D. Patent citations and other measures of R&D performance have highly skewed distributions, with the top 10 to 20% of inventions typically capturing more than 90% of returns (Scherer and Harhoff, 2000; Silverberg

[16] Because incentives were competitive, raising the variance of project choices would only have led to a larger expected payment, conditional on unchanged choices by other participants.

and Verspagen, 2007). If we imagine that scores of 5 in the experiment represent the 80th percentile of quality, then a project with mean of 3.5 and variance of 0.28 has only a 3% chance of yielding a score of 5. The higher variance project has a much higher 25% chance. Given the convexity in returns found in citation data, the expected value of a project that could yield $1 million if successful is just $65,000 for the low variance project and $250,000 for the high variance project. Individuals with 75th percentile risk aversion, by preferentially choosing the lower variance project, would lose out on $10,000 in expectation (or 5.5% of the difference in expected value between the two projects). Individuals with 25th percentile risk aversion would lose 2.2% of the difference in value, while 10th percentile individuals would lose nothing in expectation.

These findings suggest that the personal preferences of those individuals in charge of research investments may exert an oversized influence on investment decisions within the firms in which they are employed. That is, who is placed in charge of research investment decisions may be as important as the incentives that firms provide them to make those decisions. As such, efforts to integrate risk preferences into the criteria driving hiring and promotion decisions could yield increases in the productivity of the R&D divisions in which they are employed. Whether this is best achieved through the importation of one of the many assessment approaches developed in economics and psychology, through new AI tools that are beginning to be integrated into HR departments (McKinsey, 2020; Li, Raymond, and Bergman, 2020), or some combination of the two will likely depend on firm characteristics and goals as well as the continued evolution of these tools over time (Cowgill, 2018).

At the same time, our finding that individuals with more R&D experience perform better, suggests that the relevant decision-making skills can be learned. A critical question for managers is determining which employees are best positioned to acquire this knowledge and the best approaches to accelerate this learning. The recent evidence that suggests that entrepreneurship training, as distinct from standard business school training, can greatly improve decision making in highly uncertain domains (Camuffo et al., 2020; Lyons and Zhang, 2018) offer some reason for optimism. How to optimize those for R&D choices has important implications for the advancement of science, and the fate of research-intensive firms.

We conclude by noting that, as with all experiments, it is unclear how the experience in the 'laboratory' generalizes to more realistic work settings. Two important considerations merit particular attention in our context. First, and perhaps most importantly, we examined individual decision making but many R&D decisions are made by teams. While team risk preferences do appear to influence how teams respond to incentives to innovate (Graff Zivin and Lyons, 2020), precisely how risk preferences are aggregated within a team and what that implies about group decision making remains an open question. Second, decisions in our experiment were one-shot, while many firm level decisions are sequential with opportunities for mid-course corrections and early project determination. Whether the tradeoff of risk and reward in a more options-oriented framework differs significantly from those in our setting is equally unknown. Finally, it is important to recognize that not all additional risk taking is equally valuable to the firm. Efforts to encourage more radical forms of exploration must balance concerns regarding moral hazard and the relative values of more incremental forms of innovation, with the optimal mix likely to vary across firms and sectors. Together, these questions represent an area ripe for future research.

References

Abdellaoui M, Bleichrodt H, L'Haridon O (2008) A Tractable Method to Measure Utility and Loss Aversion under Prospect Theory. *Journal of Risk and Uncertainty.* 36(3):245–266.

Aghion P, Howitt P (1992) A Model of Growth Through Creative Destruction. *Econometrica.* 60(2): 323–51. https://doi.org/10.2307/2951599.

Amit R, Zott C (2001) Value Creation in E-Business. *Strategic Management Journal.* 22(6–7): 493–520. https://doi.org/10.1002/smj.187.

Azoulay P, Graff Zivin JS, Manso G (2011) Incentives and Creativity: Evidence from the Academic Life Sciences. *The RAND Journal of Economics.* 42(3): 527-554.

Benjamin DJ, Brown SA, Shapiro JM (2013) Who is 'Behavioral'? Cognitive Ability and Anomalous Preferences. *Journal of the European Economic Association.* 11(6): 1231-1255.

Bloom N, Jones CI, Van Reenen J, Webb M (2017) Are Ideas Getting Harder to Find? *National Bureau of Economic Research.* 23782: 1-53. http://www.nber.org/papers/w23782.

Blom, G (1961) Statistical Estimates and Transformed Beta-variables. *Biometrische Zeitschrift.* 3(4): 285.

Boudreau KJ, Guinan EC, Lakhani KR, Riedl C (2016) Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance, Novelty, and Resource Allocation in Science. *Management Science* 62(10): 2765-2783.

Cabral L (1994) Bias in Market R&D Portfolios. *International Journal of Industrial Organization.* 12(4): 533–47. https://doi.org/10.1016/0167-7187(94)90006-X.

Cabral L (2003) R&D Competition When Firms Choose Variance. *Journal of Economics and Management Strategy.* 12(1): 139–50. https://doi.org/10.1162/105864003321220760.

Camuffo A, Cordova A, Gambardella A, Spina C (2020) A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial. *Management Science.* 66(2): 564-586.

Charness G, Gneezy U, Imas A (2013) Experimental Methods: Eliciting Risk Preferences. *Journal of Economic Behavior and Organization.* 87: 43–51.

Cowgill B (2018) Bias and Productivity in Humans and Algorithms: Theory and Evidence from Resume Screening. *Working Paper, Columbia Business School*. 29.

Dasgupta P, Maskin E (1987) The Simple Economics of Research Portfolios. *The Economic Journal* 97: 581–95. https://doi.org/10.2307/2231709.

Dasgupta P, David PA (1994) Toward a New Economics of Science. *Research Policy* 23(5): 487-521.

David, HA (1988) *The Method of Paired Comparisons.* (Griffin, London).

Fiebig DG, Keane MP, Louviere JJ, Wasi N (2010) The Generalized Multinomial Logit Model: Accounting for Scale and Coefficient Heterogeneity. *Marketing Science.* 29(3): 393-421.

Frederick S (2005) Cognitive Reflection and Decision Making. *Journal of Economic Perspectives.* 19(4): 25-42.

Goel AM, Thankor AV (2008) Overconfidence, CEO Selection, and Corporate Governance. *The Journal of Finance.* 63(6): 2737–84. https://doi.org/10.1111/j.1540-6261.2008.01412.x.

Graff Zivin J, Lyons E (2020) The Effects of Prize Structures on Innovative Performance. *National Bureau of Economic Research*. 26737.

Helpman E (1998) *General Purpose Technologies and Economic Growth*. (MIT Press).

Hirshleifer D, Low A, Teoh SH (2012) Are Overconfident CEOs Better Innovators? *The Journal of Finance.* 67(4): 1457–98. https://doi.org/10.1111/j.1540-6261.2012.01753.x.

Hopenhayn H, Squintani F (2021) On the Direction of Innovation. *Journal of Political Economy.* 129(7): 000-000.

Imas A, Sadoff S, Samek A (2017) Do People Anticipate Loss Aversion? *Management Science*. 63(5): 1271-1284.

Kagan E, Leider S, Lovejoy WS (2019) Equity Contracts and Incentive Design in Start-Up Teams. *Management Science*. mnsc.2019.3439. https://doi.org/10.1287/mnsc.2019.3439.

Keupp MM, Palmié M, Gassmann O (2012) The Strategic Management of Innovation: A Systematic Review and Paths for Future Research. *International Journal of Management Reviews.* 14(4): 367–90. https://doi.org/10.1111/j.1468-2370.2011.00321.x.

Krieger J, Li D, Papanikolaou D (2019) *Missing Novelty in Drug Development*. (Cambridge, MA). https://doi.org/10.3386/w24595.

Li D, Raymond LR, Bergman P (2020) Hiring as Exploration. *National Bureau of Economic Research*. 27736.

Lotka AJ (1926) The Frequency Distribution of Scientific Productivity. *J. Washington Academy of Sci.* 16(12): 317–23.

Louviere JJ, Woodworth G (1983) Design and Analysis of Simulated Consumer Choice or Allocation Experiments: An Approach Based on Aggregate Data. *Journal of Marketing Research.* 20(4): 350-367.

Louviere JJ, Flynn TN, Marley AAJ (2015) *Best-Worst Scaling: Theory, Methods and Applications.* (Cambridge University Press).

Luce RD (1959) On the Possible Psychophysical Laws. *Psychology Review.* 66(2).

Lyons EE, Zhang L (2018) Who Does (Not) Benefit from Entrepreneurship Programs? *Strategic Management Journal.* 39(1): 85-112.

March JG (1991) Exploration and Exploitation in Organizational Learning. *Organization Science.* 2(1): 71-87.

Marks AR (2011) Repaving the Road to Biomedical Innovation Through Academia. *Science Translational Medicine.* 3(89): 89cm15.

Manso, G (2011) Motivating Innovation. *Journal of Finance.* 66(5): 1823-1860.

McFadden, D (1974) The Measurement of Urban Travel Demand. *Journal of Public Economics.* 3(4): 303-328.

McKinsey (2020) *The State of AI in 2020*. https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/global-survey-the-state-of-ai-in-

Merton RK (1973) *The Sociology of Science: Theoretical and Empirical Investigations*.

Munos BH, Chin WW (2011) How to Revive Breakthrough Innovation in the Pharmaceutical Industry. *Science Translational Medicine.* 3: 89cm16.

Porter ME (1985) *Competitive Advantage: Creating and Sustaining Superior Performance*. https://doi.org/10.1182/blood-2005-11-4354.

Scherer FM, Harhoff D (2000) Technology Policy for a World of Skew-distributed Outcomes. *Research Policy*. 29(4-5): 559-566.

Silverberg G, Verspagen, B (2007) The Size Distribution of Innovations Revisited: An Application of Extreme Value Statistics to Citation and Value Measures of Patent Significance. *Journal of Econometrics*. *139*(2): 318-339.

Stephan PE (2010) The Economics of Science. In *Handbook of the Economics of Innovation*. 1:217–73. (Elsevier B.V.). https://doi.org/10.1016/S0169-7218(10)01005-1.

Teece DJ (2010) Business Models, Business Strategy, and Innovation. *Long Range Planning.* 43(2–3): 172–94. https://doi.org/10.1016/j.lrp.2009.07.003.

Thurstone LL (1927) A Law of Comparative Judgment. *Psychology Review.* 34(4): 273.

Tian X, Wang TY (2014) Tolerance for Failure and Corporate Innovation. *The Review of Financial Studies.* 27(1): 211-255.

Tishler A (2008) How Risky Should an R&D Program Be? *Economics Letters.* 99(2): 268–71. https://doi.org/10.1016/j.econlet.2007.07.006.

Figures and Tables

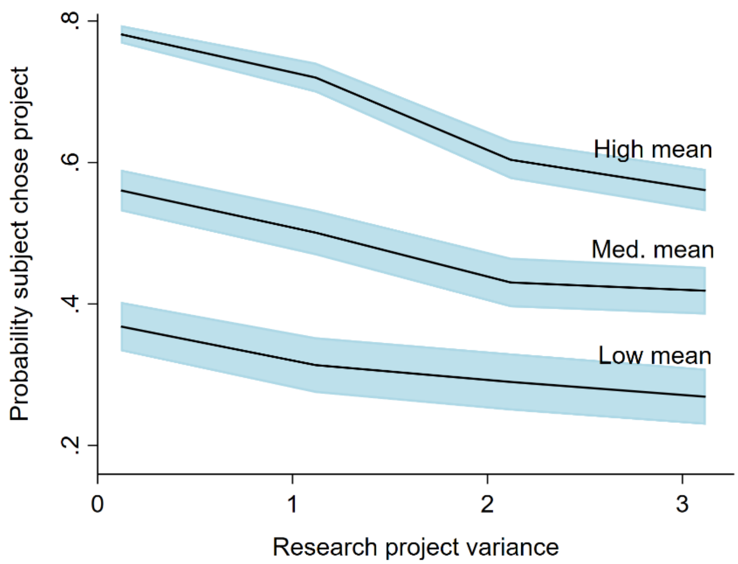

Figure 1. Likelihood of Choosing a Project with a Given Mean and Variance
The figure shows the average probability of a subject choosing a project with a given mean and variance. Each solid line shows how the choice probability changes as variance increases for three different mean project scores: a high mean score of 4, a medium mean score of 3.5, and a low mean score of 3. The light blue areas are 95% confidence intervals. The estimates are based on a GMNL model fit to data from the first set of choice scenarios from the phase 1 experiment. The underlying GMNL model results are presented in Table 2.

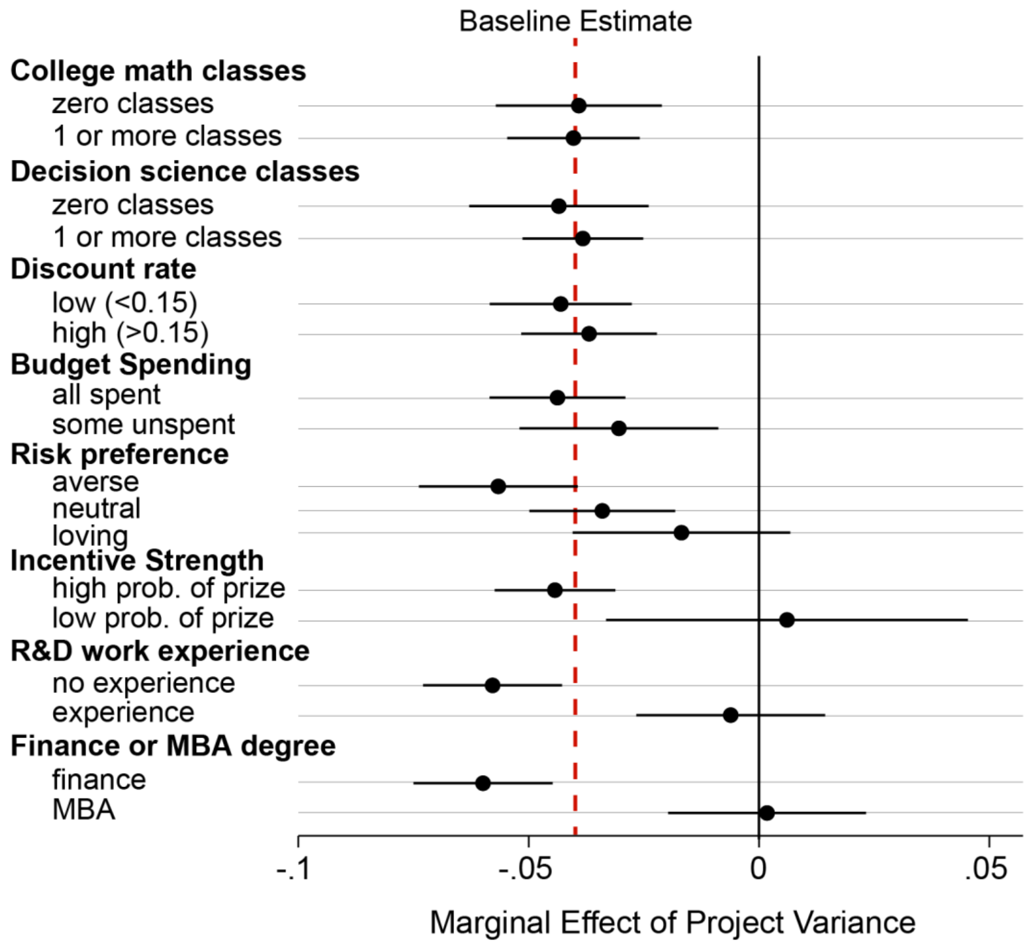

Figure 2. Heterogeneity in Marginal Effect of Variance on Project Choice
The figure shows marginal effects of variance on the probability of choosing a project in the first set of choice scenarios from the phase 1 experiment, broken down by subject demographics and characteristics. The points are based on coefficient estimates from the eight G-MNL regressions reported in Table 3. The black, solid, horizontal lines show the 95% confidence intervals based on standard errors clustered at the subject level. The vertical, red, dashed line shows the marginal effect based on the G-MNL estimates in Table 2. For continuous heterogeneity measures (coursework, discount rate, and incentive strength) the measure is discretized by dividing into above and below mean values.

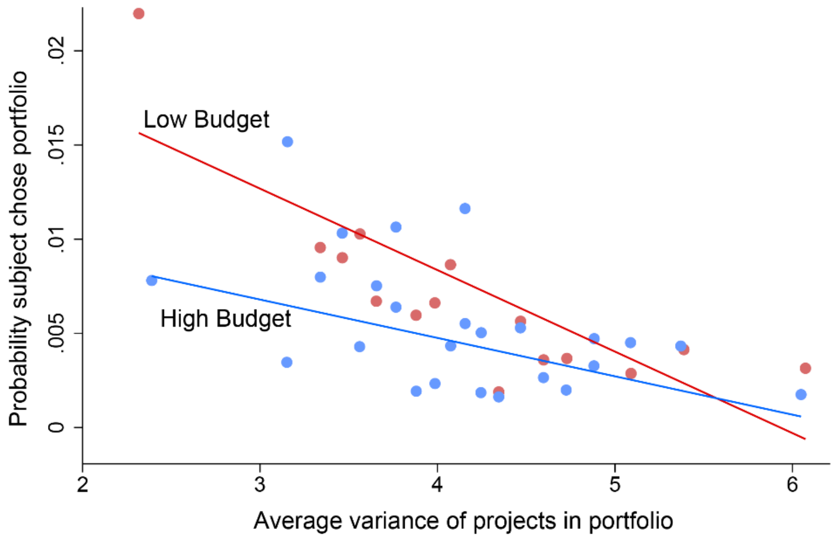

Figure 3. Effect of Budget on Preference for Portfolio Variance
The figure shows the average effect of variance on portfolio choice (choice scenario set 3) in the phase 1 experiment, broken down by budget. The red circles show the effect of variance on portfolio choice for low budgets (less than $15 million). The red line is a linear fit through the low budget points. The blue circles show the same relationship for higher budgets (between $16 and $19 million). The blue line shows a linear fit through the high budget points. All values are conditional on average project mean, average project cost, the interaction between average project mean and cost, the interaction between average project variance and cost, and choice scenario indicator variables.

Panel A: Phase 1 Experiment

Variable	Mean	Std. Dev.
Age	26.68	5.29
Years of work experience	2.62	4.35
Has worked in R&D	0.30	0.46
Coef. of rel. risk aversion	1.26	0.27
Discount rate	0.25	0.21
Math classes	4.64	1.46
Decision science classes	4.03	1.56
Observations	150	

Panel B: Phase 2 Experiment

	Arm 1		Arm 2		Arm 3	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Age	31.41	6.09	31.70	5.44	32.04	7.47
Years of work experience	7.34	6.49	6.41	5.76	7.95	6.96
Has worked in R&D	0.54	0.50	0.53	0.50	0.53	0.50
Coef. of rel. risk aversion	1.26	0.27	1.14	0.29	1.15	0.33
Discount rate	0.29	0.27	0.27	0.27	0.24	0.21
Loss averse	0.67	0.47	0.77	0.43	0.74	0.44
Math classes	4.15	1.56	4.11	1.75	3.79	1.74
Decision science classes	2.74	1.73	2.77	1.81	2.81	1.50
Observations (total=140)	46		47		47	

Table 1. Summary Statistics
The table shows summary statistics for the participants in the experiments. Panel A shows statistics for the 150 participants in the phase 1 experiment, and Panel B shows statistics for the 140 participants in the phase 2 experiment.

Dependent variable:	(1) Phase 1 Choice Scenario Set 1 Project choice	(2) Phase 1 Choice Scenario Set 1 Project choice
Average Utility Weight		
Average Project Score	5.09***	2.48***
	(0.51)	(0.18)
Project Score Variance	-0.63***	-0.45***
	(0.075)	(0.051)
Utility Weight Heterogeneity		
Average Project Score	2.35***	
	(0.30)	
Project Score Variance	0.75***	
	(0.076)	
Tau	0.029	
	(0.090)	
Model	G-MNL	C-Logit
Observations	13,500	13,500
Subjects	150	150

Table 2. Project choice as a function of mean and variance

The table shows results estimated using choice-scenario-level data from choice scenario set 1 in the phase 1 experiment. Column 1 is estimated using a G-MNL model (Equation 1). Column 2 is estimated using a conditional logit model. The outcome variable is an indicator for whether the project was chosen. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Dependent variable:	(1) Project Choice	(2) Project Choice	(3) Project Choice	(4) Project Choice	(5) Project Choice	(6) Project Choice	(7) Project Choice	(8) Project Choice
Dimension of heterogeneity:	Coef. RRA	Budget unspent	RD Work Exp.	Discount rate	Math classes	Decision Science	Finance of MBA	Prize Prob.
Average Utility Weight								
Average Score	5.46***	5.73***	5.10***	4.90***	6.26***	4.71***	4.55***	3.73***
	(1.11)	(0.63)	(0.41)	(0.93)	(1.35)	(0.43)	(0.47)	(0.50)
Score Variance	0.80**	-0.52***	-0.74***	-0.74***	-0.46**	-0.47***	-0.23***	-0.25*
	(0.37)	(0.11)	(0.076)	(0.13)	(0.19)	(0.14)	(0.083)	(0.14)
Average x Het.	-0.24	-0.051	0.33	0.49	-1.71	1.00**	1.11***	2.04***
	(0.59)	(0.69)	(1.07)	(4.68)	(1.44)	(0.47)	(0.42)	(0.28)
Variance x Het.	-1.11***	-0.13	0.38	0.49**	-0.21	-0.29	-0.67***	-0.50***
	(0.30)	(0.15)	(0.32)	(0.23)	(0.23)	(0.18)	(0.13)	(0.18)
Utility Weight Heterogeneity								
Average Score	0.57***	1.44***	2.38***	1.79***	0.018	2.31***	2.24***	1.92***
	(0.11)	(0.23)	(0.23)	(0.22)	(0.15)	(0.24)	(0.25)	(0.24)
Score Variance	0.72***	0.79***	0.64***	0.74***	0.68***	0.69***	0.58***	0.56***
	(0.068)	(0.076)	(0.078)	(0.068)	(0.085)	(0.085)	(0.084)	(0.11)
Average x Het.	2.16***	1.37**	1.91***	4.99	1.69***	0.99***	1.42***	2.02***
	(0.23)	(0.56)	(0.27)	(4.05)	(0.28)	(0.13)	(0.28)	(0.31)
Variance x Het.	0.014	0.16	0.70***	0.23	0.44***	0.44**	0.56***	0.57***
	(0.049)	(0.11)	(0.27)	(0.38)	(0.095)	(0.20)	(0.12)	(0.13)
Tau	0.37***	0.58***	0.075	0.024	0.39**	0.14***	0.11***	0.058*
	(0.058)	(0.080)	(0.11)	(0.098)	(0.16)	(0.015)	(0.026)	(0.031)
Subjects	150	150	150	150	150	150	150	150
Observations	13,500	13,500	13,500	13,500	13,500	13,500	13,500	13,500

Table 3. Project Choice Heterogeneity, Choice Scenario Set 1, Phase 1

The table shows results from estimating Equation (1) using choice-scenario data from choice scenario set 1 in the phase 1 experiment. The outcome variable is an indicator for whether the project was chosen. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. Each column also shows the effect of project the interaction between those variables and a dimension of heterogeneity. The dimension of heterogeneity is given at the top of the column. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Choice Scenario 2 Data Only	(2) Both Choice Scenarios 1 and 2	(3) Choice Scenario 2 Data Only	(4) Both Choice Scenarios 1 and 2
Dependent variable	Project choice	Project choice	Project choice	Project choice
Average Utility Weight				
Average project score	5.01***	4.82***	2.07***	2.48***
	(0.46)	(0.29)	(0.11)	(0.18)
Project score variance	-1.59***	-0.62***	-0.85***	-0.45***
	(0.16)	(0.085)	(0.073)	(0.051)
Average x Choice scenario 2		0.12		-0.41***
		(0.28)		(0.15)
Variance x Choice scenario 2		-0.92***		-0.40***
		(0.11)		(0.064)
Utility Weight Heterogeneity				
Average project value	2.68***	1.76***		
	(0.31)	(0.18)		
Project variance	1.57***	0.81***		
	(0.13)	(0.086)		
Average x Choice scenario 2		1.14***		
		(0.23)		
Variance x Choice scenario 2		1.19***		
		(0.086)		
Tau	0.046	0.27***		
	(0.042)	(0.045)		
Model	G-MNL	G-MNL	C-Logit	C-Logit
Subjects	150	150	150	150
Observations	13,500	27,000	13,500	27,000

Table 4. Comparison of Choice Scenarios 1 and 2, Phase 1 Experiment

The table shows results using choice-scenario-level data from choice scenario sets 1 and 2 in the phase 1 experiment. The sample restrictions are indicated at the top of each column. Columns 1 and 2 are estimated using G-MNL models (Equation 1). Columns 3 and 4 are estimated using conditional logit models. The outcome variable is an indicator for whether the project was chosen. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. "Choice scenario 2" is an indicator equal to 1 if the data come from choice scenario 2 and 0 otherwise. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Dependent variable:	(1) Project choice	(2) Project choice
Average Utility Weight		
Average project value	6.43***	2.80***
	(0.65)	(0.28)
Project variance	-0.36***	-0.31***
	(0.100)	(0.085)
Average x Arm 2	-1.25***	-0.69
	(0.20)	(0.43)
Average x Arm 3	-0.29	-0.52
	(0.18)	(0.35)
Variance x Arm 2	-0.23	-0.025
	(0.14)	(0.12)
Variance x Arm 3	-0.34***	-0.21
	(0.13)	(0.13)
Utility Weight Heterogeneity		
Average project value	3.75***	
	(0.53)	
Project variance	0.70***	
	(0.064)	
Average x Arm 2	0.16*	
	(0.086)	
Average x Arm 3	0.52**	
	(0.20)	
Variance x Arm 2	0.11	
	(0.071)	
Variance x Arm 3	0.11***	
	(0.038)	
Tau	0.16***	
	(0.014)	
Model	G-MNL	C-Logit
Subjects	140	140
Observations	12,600	12,600

Table 5. Choice as a Function of Project Attributes and Treatment Arm: Phase 2 Experiment

The table shows results from estimating Equation (1) using choice-scenario-level data from choice scenario set 1 in the phase 2 experiment. The outcome variable is an indicator for the chosen project. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. The interactions are indicators for experimental arm. The base category is the replication arm. "Arm 2" is the no-loss-framing treatment, and "Arm 3" is the objective costs treatment. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix for
The Risk of Caution: Evidence from an R&D Experiment

A. Experiment Details

A.1. Instructions for Baseline Experiment and Replication of Baseline

The exact text of the instructions given to research subjects at the start of the experiment, framing the choices in the context of R&D and describing the incentives is

> Welcome to the survey and thank you for agreeing to participate in our study. Over the course of the next hour you will be asked to assume the role of the manager of a research division of an organization in the biomedical/health sector. You will be presented data on a series of potential research projects that you could fund. Some questions will ask you to rank individual projects to fund; other questions will ask you to construct a portfolio of projects from a selected list. We will also ask you some additional demographic and preference questions to better understand your decision-making processes.
>
> Research project outcomes always involve some uncertainty, which is partly reflected in the diversity of evaluation scores that can be assigned to any given proposal. After you complete the survey, the computer will draw a number from a random number generator that is consistent with the characteristics of the research projects you selected for each of the R&D investment questions you completed. Better ranked proposals will tend to have better outcomes and proposals where there is more disagreement in the ranking will tend to have more variable, both good and bad, outcomes. When proposals have different costs, expected payoffs are proportionate to proposal cost.
>
> The random numbers generated for each question will then be added to provide an aggregate score for each participant in the survey. While all participants will receive $15 for participation, bonus payments will be offered for top performers. Those that score within the top 25 percent of this survey round will receive an additional $25 bonus, with that bonus increasing to $100 for those within the top 10 percent.

Exact text of the instructions introducing the first experiment:

> For the next ten questions, assume that you are the head of the research division of an organization and are considering funding four research project proposals (A, B, C, D).

- Each project proposal has received a rating on a scale from 1 to 5 (with 5 being the top rating) by seven scientific experts unaffiliated with the projects under consideration (1, 2, 3, 4, 5, 6, 7) on your advisory board.
- All the proposals have the same cost.
- The matrix below displays how the four proposals (columns) you should use for this question were rated by the seven reviewers (rows).
- The average of the reviewers' scores for each proposal is shown at the bottom of the proposal's column.
- The order in which proposals appear is randomized.
- Proposal rankings should be treated as an indication of potential financial return. Negative returns (financial losses) are possible. Remember that your final compensation for participating in this study will depend on the choices you make here.

Prior to the second experiment, where we explicitly showed the variance, we added the following note to the above instructions:

> In addition to the average of the reviewers' scores for each proposal, we also report the **variance** of scores, a measure of the variability of the reviewers' assessments.

Exact text of the instructions introducing the portfolio selection questions (experiment 3):

> The next ten questions are similar in spirit to the to the earlier ones concerning the project that you would most/least like to fund. In this case, however, you will be asked to put together a portfolio of research project proposals to fund. As before, assume that you are the head of the research division of an organization conducting R&D.
>
> - There are proposals for 10 possible research projects (denoted A through J).
> - Each proposal has received a rating on a scale from 1 to 5 (with 5 being the top rating) by seven scientific experts unaffiliated with the projects under consideration (reviewer 1 through 7) on your advisory board.
> - The average of the reviewers' scores for each proposal and the variance are also displayed in each proposal's column.
> - The order in which proposals appear is randomized.
> - Proposal rankings should be treated as an indication of potential financial return. Negative returns (financial losses) are possible.
>
> In contrast to our earlier questions, each proposal now has a **different cost**, which is displayed in the last row of the table. The cost of specific proposals

will now influence what research projects you are able to fund. The portfolio you choose must cost the same or less than your budget constraint. Any leftover funds from each question will be returned to your organization's headquarters and will be unavailable for future R&D investments by your Division. Remember that your final compensation for participating in this study will depend on the choices you make here.

Please repeat the exercise from the previous question for this **new** set of research projects with a **different** set of scores from your science advisory panel.

Round (2 of 10)

Reviewer	Proposal A	Proposal B	Proposal C	Proposal D
1	2	4	5	2
2	5	4	2	2
3	3	4	3	5
4	2	4	4	2
5	5	1	1	4
6	5	3	1	5
7	2	4	2	5
Average score	3.43	3.43	2.57	3.57
Variance	1.96	1.10	1.96	1.96
Which proposal you would **most** like to fund.	☐	☐	☐	☐
Which proposal you would **least** like to fund.	☐	☐	☐	☐

Fig. S1. Example of Project Selection Question from Experiment 2

The screen shows a real project selection question shown to a subject in the experiment. The screen comes from experiment 2, which showed both the mean score for each proposed project as well as the variance of scores. The screens for experiment 1 were identical except that they did not explicitly show the variance.

Your R&D budget to fund this portfolio of proposals is $14 million dollars.

Check the boxes of the projects you would like to fund located below each proposal. As you enter proposals, the costs of each one will be subtracted from your R&D budget and your remaining funds will be displayed. Depending on the cost of the individual proposals you fund you may not be able to fill in all of the boxes and you may not fully exhaust your R&D budget. You can swap proposals in and out by selecting and unselecting boxes until you are satisfied with your choices. Then click on the arrow to advance to the next question.

Remaining Budget: $14 million

Reviewer	Proposal A	Proposal B	Proposal C	Proposal D	Proposal E	Proposal F	Proposal G	Proposal H	Proposal I	Proposal J
1	4	3	5	5	5	1	4	5	2	4
2	1	2	5	5	3	4	5	2	3	1
3	1	1	5	5	5	1	5	4	1	5
4	1	1	5	4	5	1	3	2	3	4
5	1	1	2	3	3	2	5	2	4	4
6	2	1	5	5	5	1	5	5	2	4
7	1	2	5	5	5	1	5	5	3	3
Average Score	1.57	1.57	4.57	4.57	4.43	1.57	4.57	3.57	2.57	3.57
Variance	1.29	0.62	1.29	0.62	0.95	1.29	0.62	2.29	0.95	1.62
Cost	$10 M	$10 M	$10 M	$10 M	$7 M	$1 M	$1 M	$4 M	$4 M	$1 M
	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

Fig. S2. Example Portfolio Selection Question from Experiment 3

The screen shows a real portfolio selection question (experiment 3) given to one of the subjects as part of the experiment. The subject could choose projects for the portfolio using the checkboxes below the project proposals. The subject could add and remove projects, seeing their effect on "Remaining Budget", until they were satisfied.

A.2. Instructions for Non-loss Framed Version of the Experiment

The instructions for the non-loss framed experiment removed any mention of potential losses.

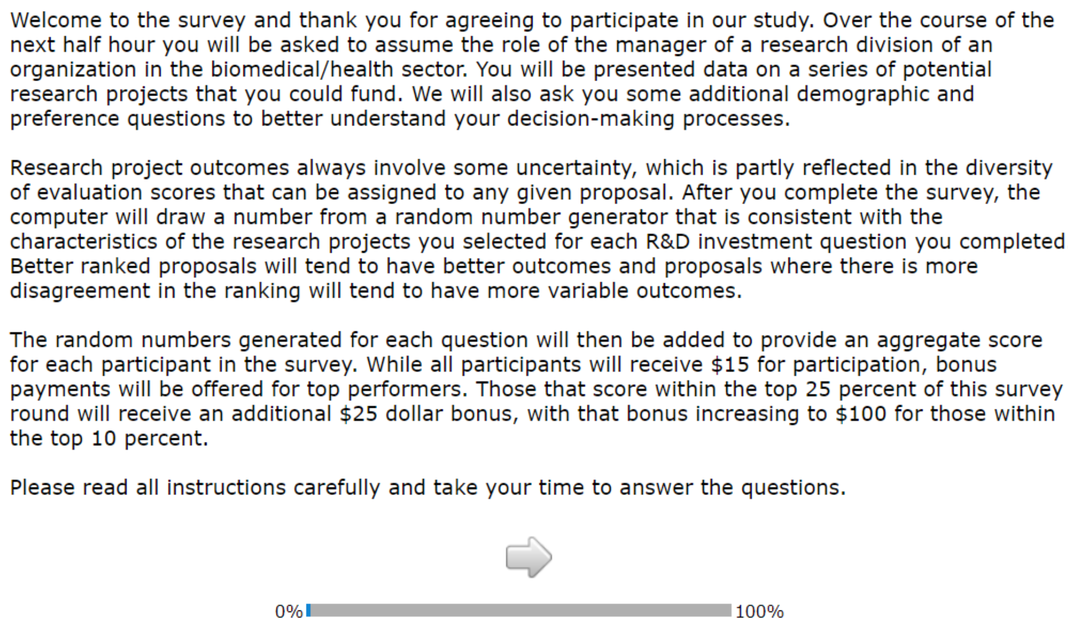

Welcome to the survey and thank you for agreeing to participate in our study. Over the course of the next half hour you will be asked to assume the role of the manager of a research division of an organization in the biomedical/health sector. You will be presented data on a series of potential research projects that you could fund. We will also ask you some additional demographic and preference questions to better understand your decision-making processes.

Research project outcomes always involve some uncertainty, which is partly reflected in the diversity of evaluation scores that can be assigned to any given proposal. After you complete the survey, the computer will draw a number from a random number generator that is consistent with the characteristics of the research projects you selected for each R&D investment question you completed. Better ranked proposals will tend to have better outcomes and proposals where there is more disagreement in the ranking will tend to have more variable outcomes.

The random numbers generated for each question will then be added to provide an aggregate score for each participant in the survey. While all participants will receive $15 for participation, bonus payments will be offered for top performers. Those that score within the top 25 percent of this survey round will receive an additional $25 dollar bonus, with that bonus increasing to $100 for those within the top 10 percent.

Please read all instructions carefully and take your time to answer the questions.

Fig. S3. Introductory Instructions for Non-loss Framed Experiment

The image shows a screenshot of the instruction provided to subjects at the beginning of the non-loss framed experiment. Relative to the baseline experiment, the instructions differ only in the removal of language stating that losses or bad outcomes might occur with high variance or low scoring projects.

For the next ten questions, assume that you are the head of the research division of an organization and are considering funding four research project proposals (A, B, C, D).

- Each project proposal has received a rating on a scale from 1 to 5 (with 5 being the top rating) by seven scientific experts unaffiliated with the projects under consideration (1, 2, 3, 4, 5, 6, 7) on your advisory board.
- All the proposals have the same cost.
- The matrix below displays how the four proposals (columns) you should use for this question were rated by the seven reviewers (rows).
- The average of the reviewers' scores for each proposal is shown at the bottom of the proposal's column.
- The order in which proposals appear is randomized.
- Proposal rankings should be treated as an indication of potential financial return. Remember that your final compensation for participating in this study will depend on the choices you make here.

Round (1 of 10)

Reviewer	Proposal A	Proposal B	Proposal C	Proposal D
1	2	3	4	2
2	2	5	4	5
3	2	4	4	3
4	3	5	3	2
5	3	5	3	5
6	4	4	3	5
7	2	5	3	2
Average score	2.57	4.43	3.43	3.43
Which proposal you would **most** like to fund.	☐	☐	☐	☐
Which proposal you would **least** like to fund.	☐	☐	☐	☐

0% 100%

Fig. S4. Example of Project Choice Question from Non-loss Framed Experiment

The image shows a screenshot from an example project choice shown to the subjects in the non-loss framed experiment. Note that compared to the baseline experiment, the phrase "Negative returns (financial losses) are possible" has been removed from the final bullet point in the instructions.

A.3. Instructions for Objective Costs Version of the Experiment

The instructions for the non-loss framed experiment removed any mention of potential losses.

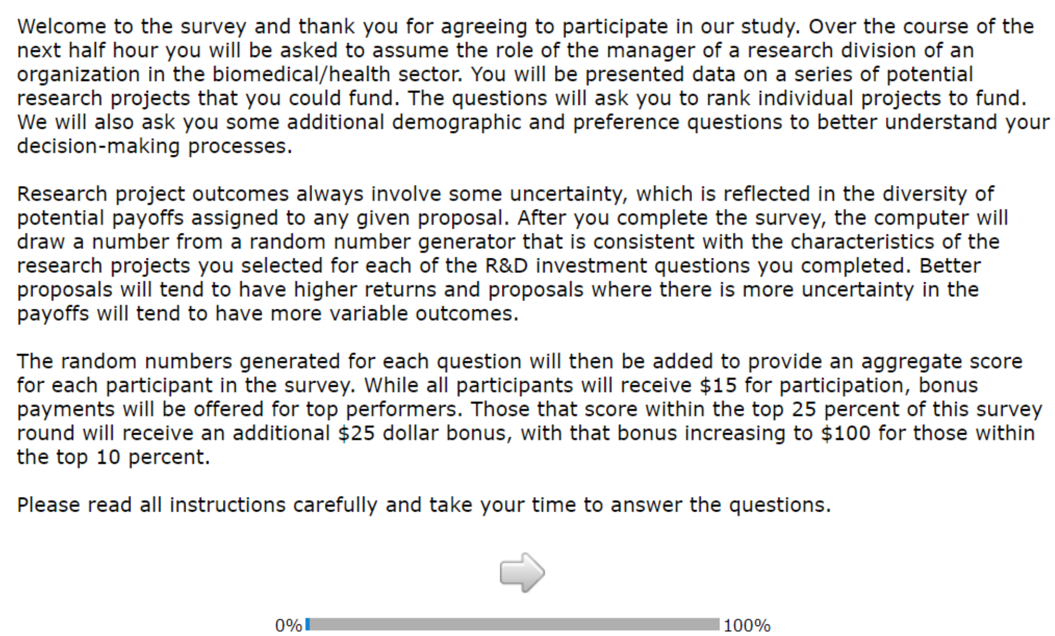

Welcome to the survey and thank you for agreeing to participate in our study. Over the course of the next half hour you will be asked to assume the role of the manager of a research division of an organization in the biomedical/health sector. You will be presented data on a series of potential research projects that you could fund. The questions will ask you to rank individual projects to fund. We will also ask you some additional demographic and preference questions to better understand your decision-making processes.

Research project outcomes always involve some uncertainty, which is reflected in the diversity of potential payoffs assigned to any given proposal. After you complete the survey, the computer will draw a number from a random number generator that is consistent with the characteristics of the research projects you selected for each of the R&D investment questions you completed. Better proposals will tend to have higher returns and proposals where there is more uncertainty in the payoffs will tend to have more variable outcomes.

The random numbers generated for each question will then be added to provide an aggregate score for each participant in the survey. While all participants will receive $15 for participation, bonus payments will be offered for top performers. Those that score within the top 25 percent of this survey round will receive an additional $25 dollar bonus, with that bonus increasing to $100 for those within the top 10 percent.

Please read all instructions carefully and take your time to answer the questions.

0% 100%

Fig. S5. Introductory Instructions for Objective Costs Experiment

The image shows a screenshot of the instruction provided to subjects at the beginning of the objective cost version of the experiment. In the baseline experiment, subjects were instructed that their choices would be based on project rankings. In the objective cost version of the experiment, subjects were instructed that the values shown when selecting projects were the set of possible objective, monetary-valued returns for each project.

For the next ten questions, assume that you are the head of the research division of an organization and are considering funding four research project proposals (A, B, C, D).

- Each project proposal has potential payoffs reported by seven scientific experts unaffiliated with the projects under consideration (1, 2, 3, 4, 5, 6, 7) on your advisory board.
- The payoff estimates range from $6 to $11 million. You should consider these payoff estimates to be unbiased, accurate measures of the true potential returns for each project.
- All the proposals have the same cost of $5 million.
- The matrix below displays the possible net returns for the four proposals (columns) you should use for this question. The realized net return will be drawn randomly from the set of values shown.
- The average net return for each proposal is shown at the bottom of the proposal's column.
- The order in which proposals appear is randomized.
- Remember that your final compensation for participating in this study will depend on the choices you make here.

Round (1 of 10)

Reviewer	Proposal A	Proposal B	Proposal C	Proposal D
1	$2M	$3M	$4M	$2M
2	$3M	$5M	$4M	$1M
3	$3M	$4M	$4M	$4M
4	$3M	$5M	$3M	$3M
5	$1M	$5M	$3M	$4M
6	$3M	$4M	$3M	$5M
7	$3M	$5M	$3M	$5M
Avg. net return	$2.57M	$4.43M	$3.43M	$3.43M
Which proposal you would **most** like to fund.	☐	☐	☐	☐
Which proposal you would **least** like to fund.	☐	☐	☐	☐

0% 100%

Fig. S6. Example of Project Choice Question from Objective Returns Experiment

The image shows a screenshot from an example project choice shown to the subjects in the non-loss framed experiment. In this version of the experiment, the subjects were told that the set of values associated with each project were accurate measures of the set of possible dollar-valued returns. To reinforce this point, each selection screen presented the values in terms of millions of dollars.

A.4. Illustration of Incentives to Choose Riskier Projects

Consider a simplified version of the experiment where 10 participants each choose one project, and the subject with the highest score wins a prize. Assume that all projects have a mean of zero and that nine of the participants choose projects $Y_1, \dots, Y_9 \sim i.i.d.\, N(0,1)$. Should the tenth participant choose a project, X, with a variance higher than 1? If the participant chooses a project with a variance of 1, all subjects will be symmetric, so the probability of any one of them receiving the max score will be 1/10. The participant can do better by choosing a higher variance project. If the participant chooses a project with variance approaching infinity, the probability of winning will approach 1/2 because the probability of a draw from a normal distribution with arbitrarily high variance exceeding any given positive value goes to 1/2.

Intermediate values can be approximated using order statistics. Consider the case with n participants and let $Y_{(n)} = \max\{Y_1, \dots, Y_n\}$ be the n^{th} order statistic of the choices from all other participants. By Blom (1961), the expected value of this order statistic, denoted $E(n{:}\, n)$, can be well approximated by

$$E(n{:}\, n) \approx \Phi^{-1}\left(\frac{n-\alpha}{n-2\alpha+1}\right)$$

where α*=0.375.* For *n=9,* this expected value evaluates to 1.494, so to derive the probability of the tenth participant winning given a particular choice of variance, one can evaluate the probability that a normal random variable with mean zero and that variance exceeds 1.494. For instance, the probability of X exceeding this value if $X \sim N(0,2)$ is 0.23. If the participant instead chose a project with a variance of 4, then the probability would rise to 0.35. The probability of exceeding $E(9{:}\, 9)$ for a range of possible values for the variance of X is given in Figure S3. Winning the prize in this case is a Bernoulli random variable, so an increasing probability of winning, as depicted in the figure, indicates that a higher variance choice first-order stochastically dominates a lower variance choice, and a participant should therefore choose higher variance projects regardless of their risk preferences.

The simple payment structure shown above encourages higher variance choices no matter the risk preferences of the subjects. The incentives are similar to those in the R&D model of Cabral (2003). Cabral noted that because of the winner-take-all nature of research (for instance, due to patents), a firm that is lagging behind the research frontier should engage in higher risk R&D activity. If the firm's investments pay off, they will capture the market, but if the investments do not pay off, they will not lose relative to the status quo (conditional on equal funding requirements for higher and lower risk R&D investments).

The actual incentive structure in the experiment has two payment thresholds—one for a score in the top 25% and one for a score in the top 10% of participants. With a generic two-threshold payment system, individual risk preferences could matter for optimal behavior. If a risk averse subject believed they were likely to score above the first threshold, a marginal increase in variance would raise their expected monetary payment by increasing the probability of a score above the second threshold, but it could also increase the likelihood that the subject scores in the bottom group. With sufficient risk aversion, the subject might prefer not to take this trade-off. By setting the first threshold above the 50th percentile of scores, we avoid this concern. If all participants made the same choices, then there would be a 75% chance of getting the low prize, a 15% chance of winning the middle prize, and a 10% chance of winning the high prize. If one agent deviates to reduce the variance of their choice, then they simply increase the probability of receiving the lowest prize while reducing the probability of receiving either of the higher prizes. The opposite holds for an increase in variance. Again, the distribution of earnings from higher variance choices first-order stochastically dominates the distribution from lower variance choices. All other things equal, subjects should choose higher variance projects regardless of risk preferences.

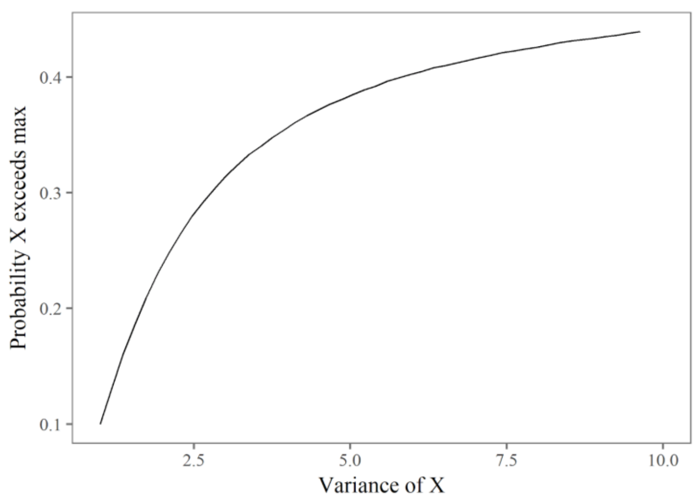

Fig. S3. Probability of Winning Prize Rises with Variance

Approximations of $E(9{:}9)$ for variance from 1 to 10 illustrating the incentive for choosing projects with higher variance.

A.5. Calculation of the Coefficient of Relative Risk Aversion

Subjects were shown the list of choices in Figure S4 between guaranteed payments and gambles. We calculated a coefficient of relative risk aversion from the choices by solving for x in the equation

$$20{,}000^{x} = 0.5 * 10{,}000^{x} + 0.5 * G^{x}$$

Where G is the simple average of high value of the two gambles from the first time that the subject switched from choosing Option B to choosing Option A. For instance, if the subject chose Option B until the offered gamble was for \$10,000 or \$30,000, then we set G equal to \$35,000. For this subject, the elicited degree of risk aversion would be 1.21. We chose the average value because we ideally want to find the point of indifference between the guaranteed payment and the gamble.

We classified subjects as risk averse if they switched to the guaranteed payment prior to the \$10,000:\$30,000 choice. Subjects who switched at the \$10,000:\$30,000 choice were classified as risk neutral, and subjects who switched at the \$10,000:\$20,000 choice were classified as risk loving. Four subjects always chose Option B and could not be classified. In the estimation sample, 52% of subjects were risk averse, 36% were risk neutral, and 12% were risk loving.

Suppose you have inherited investment property and you are given a choice between two different types of investments. One will provide a one-time guaranteed payout and the other will pay a one-time uncertain one.

Option A	Option B
\$20,000 guaranteed	50:50 chance at \$10,000:\$90,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$80,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$70,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$60,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$50,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$40,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$30,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$20,000
\$20,000 guaranteed	50:50 chance at \$10,000:\$10,000

Fig. S4. Risk Preference Elicitation Method
The screen shows the question used to elicit risk preferences from study participants.

A.6. Calculation of Loss Aversion Preferences

We elicited loss aversion parameters using a similar elicitation to Imas, Sadoff, and Samek (2017), based on the design of Abdellaoui et al. (2008). The elicitation was done through a set of three choice menus. The menus were similar to those used to elicit risk aversion, as shown in the previous section. Importantly, however, the choices in the menus were over hypothetical gambles that either did or did not include losses. Using the choices, we estimated risk aversion over gains and losses as well as a measure of loss aversion, λ, the difference in slope in the utility function between the gain and loss domain after accounting for risk-aversion-based curvature. Formally, λ is the parameter in the prospect theory value function given by

$$v(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x)^{\beta} & \text{if } x < 0 \end{cases}$$

Overall, most subjects were loss averse in our sample. The 25th percentile λ of 0.96 and the median λ was 1.74 (loss aversion is indicated by λ > 1). For comparison, the median λ in Imas, Sadoff, and Samek (2017) was 1.59.

B. Additional Tables

	(1) I.AB	(2) I.BA	(3) II.AB	(4) II.BA
Discount rate	-0.25	0.079	0.39**	-0.22
	(0.16)	(0.19)	(0.18)	(0.15)
Coef. RRA	-0.13	0.078	0.21*	-0.16
	(0.13)	(0.17)	(0.11)	(0.12)
College math	0.020	-0.024	0.048*	-0.044
	(0.029)	(0.030)	(0.027)	(0.027)
Dec. science class	-0.0071	0.012	-0.023	0.018
	(0.027)	(0.028)	(0.025)	(0.023)
Finance degree	-0.090	0.10	-0.076	0.064
	(0.085)	(0.11)	(0.096)	(0.095)
Work in RD	0.034	-0.092	-0.11	0.17*
	(0.090)	(0.091)	(0.088)	(0.089)
Age	-0.0032	0.028*	-0.018	-0.0063
	(0.014)	(0.014)	(0.014)	(0.011)
Years worked	-0.012	-0.011	0.028	-0.0041
	(0.018)	(0.017)	(0.019)	(0.010)
Observations	150	150	150	150
Omnibus *F*-test	0.76	1.04	1.72	1.35
p-value	0.64	0.41	0.099	0.22

Table B1. Test of Randomization of DCE Treatment in Phase 1 Experiment, Choice Scenario Sets 1 and 2

The table shows omnibus balance tests for the DCE treatment arm randomization for the first two sets of choice scenarios in the phase 1 experiment. Each column is a separate linear regression where the left-hand side variable is the DCE treatment shown to the subject. The right-hand side variables are subject demographics. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The *F*-test at the bottom of the table jointly tests whether all coefficients are equal to zero. Imbalance in the randomization is indicated by low *p*-values for this *F*-test. None of the treatments were subject to significant imbalance.

	(1) budget1	(2) budget2	(3) budget3	(4) budget4	(5) budget5	(6) budget6	(7) budget7	(8) budget8
Discount rate	-0.88	2.16***	-0.21	-0.83	-0.17	0.62	-0.99	0.30
	(0.80)	(0.82)	(0.92)	(0.94)	(0.88)	(0.87)	(0.81)	(0.95)
Coef. RRA	-1.14	0.22	0.012	0.55	0.45	1.05	-1.13*	-0.015
	(0.72)	(0.60)	(0.75)	(0.69)	(0.62)	(0.80)	(0.63)	(0.60)
College math	-0.13	0.23*	0.019	0.11	0.11	0.15	-0.25*	-0.23
	(0.14)	(0.13)	(0.15)	(0.14)	(0.15)	(0.15)	(0.13)	(0.15)
Dec. science class	0.11	0.032	-0.019	-0.33**	0.043	0.096	-0.024	0.094
	(0.14)	(0.13)	(0.14)	(0.14)	(0.14)	(0.14)	(0.13)	(0.14)
Finance degree	-0.15	0.18	0.11	-0.17	-0.16	-0.48	0.35	0.32
	(0.43)	(0.43)	(0.44)	(0.45)	(0.42)	(0.47)	(0.42)	(0.39)
Work in RD	-0.021	-0.59	0.23	-0.28	0.40	0.45	-0.24	0.044
	(0.42)	(0.44)	(0.45)	(0.45)	(0.45)	(0.47)	(0.45)	(0.43)
Observations	150	150	150	150	150	150	150	150
Omnibus *F*-test	0.79	2.06	0.055	1.53	0.34	1.23	1.47	0.49
p-value	0.58	0.062	1.00	0.17	0.92	0.30	0.19	0.81

Table B2. Test of Randomization of Budget in Phase 1 Experiment, Choice Scenario Set 3

The table shows omnibus balance tests for the budget randomization for the third set of choice scenarios (portfolio choices) in the phase 1 experiment. Each column is a separate linear regression where the left-hand side variable is the budget offered to the subject in the choice scenario. The right-hand side variables are subject demographics. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The *F*-test at the bottom of the table jointly tests whether all coefficients are equal to zero. Imbalance in the randomization is indicated by low *p*-values for this *F*-test. Budget 2 subjects had a significantly higher discount rate and exhibit differences overall that were significant at the 10% level, but none of the budgets were subject to imbalance significant at the 5% level.

	(1) Arm 1	(2) Arm 2	(3) Arm 3
Discount rate	-0.16	0.056	-0.11
	(0.14)	(0.14)	(0.16)
Coef. RRA	0.16*	-0.034	-0.0037
	(0.098)	(0.12)	(0.11)
College math	-0.020	0.00088	0.016
	(0.023)	(0.024)	(0.027)
Dec. science class	-0.0019	0.050*	-0.056**
	(0.023)	(0.026)	(0.026)
Work in RD	0.016	0.15*	-0.015
	(0.081)	(0.080)	(0.080)
Age	0.0086	0.0070	-0.011
	(0.0098)	(0.0100)	(0.0079)
Years worked	-0.0098	-0.0042	0.0045
	(0.0092)	(0.010)	(0.0087)
Lambda (loss aversion)	-0.0018	-0.00064	0.0019
	(0.0011)	(0.0015)	(0.0017)
Observations	140	140	140
Omnibus *F*-test	1.30	1.02	0.89
p-value	0.25	0.42	0.52

Table B3. Test of Randomization of Treatment Arm in Phase 2 Experiment

The table shows omnibus balance tests for the treatment arm randomization for the phase 2 experiment. Each column is a separate linear regression where the left-hand side variable is an indicator for the treatment arm of the subject. The right-hand side variables are subject demographics. Robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The F-test at the bottom of the table jointly tests whether all coefficients are equal to zero. Imbalance in the randomization is indicated by low p-values for this F-test. None of the arms were subject to significant imbalance.

	(1) Phase 1 Choice Scenario Set 1 Project choice	(2) Phase 1 Choice Scenario Set 1 Project choice
Average Utility Weight		
Average Project Score	5.86***	2.42***
	(0.51)	(0.16)
Project Score Variance	-0.53***	-0.44***
	(0.070)	(0.043)
Utility Weight Heterogeneity		
Average Project Score	2.08***	
	(0.18)	
Project Score Variance	0.81***	
	(0.077)	
Tau	0.55***	
	(0.078)	
Model	G-MNL	C-Logit
Observations	17,550	17,550
Subjects	195	195

Table B4. Phase 1, Choice Scenario Set 1 Results Including All Subjects

The table shows results estimated using choice-scenario-level data from choice scenario set 1 in the phase 1 experiment using the sample of all subjects (including session 1 and preference elicitation multiple switchers; compare to Table 2). Column 1 is estimated using a G-MNL model (Equation 1). Column 2 is estimated using a conditional logit model. The outcome variable is an indicator for whether the project was chosen. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Phase 1 vs. Phase 2 Choice Scenario Set 1 Project choice	(2) Phase 1 vs. Phase 2 Choice Scenario Set 1 Project choice
Average Utility Weight		
Average Project Score	4.91***	2.48***
	(0.38)	(0.18)
Project Score Variance	-0.62***	-0.45***
	(0.071)	(0.051)
Average x **1**{Phase 2}	2.28***	0.33
	(0.48)	(0.33)
Variance x **1**{Phase 2}	0.23	0.14
	(0.16)	(0.099)
Utility Weight Heterogeneity		
Average Project Score	2.46***	
	(0.21)	
Project Score Variance	0.73***	
	(0.065)	
Average x **1**{Phase 2}	4.85***	
	(0.65)	
Variance x **1**{Phase 2}	0.19	
	(0.19)	
Tau	-0.12***	
	(0.027)	
Model	G-MNL	C-Logit
Observations	17,640	17,640
Subjects	196	196

Table B5. Comparison of Choice Scenario Set 1 Results for Phase 1 and Phase 2 (replication)

The table shows results estimated using choice-scenario-level data from choice scenario set 1 in the phase 1 and 2 experiments. Column 1 is estimated using a G-MNL model (Equation 1). Column 2 is estimated using a conditional logit model. The outcome variable is an indicator for whether the project was chosen. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. The indicator **1**{Phase 2} is equal to 1 if the subject was part of Phase 2 and is zero otherwise. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Standard errors, clustered at the subject level, are in parentheses, * p < 0.10, ** p < 0.05, *** p < 0.01.

	(1) Phase 1 Project rank	(2) Phase 2: Replication Project rank
Rank 1		
Average Project Score	4.29***	4.46***
	(0.17)	(0.19)
Project Score Variance	-0.78***	-0.56***
	(0.087)	(0.16)
Rank 2		
Average Project Score	3.06***	3.17***
	(0.13)	(0.15)
Project Score Variance	-0.74***	-0.50***
	(0.068)	(0.14)
Rank 3		
Average Project Score	2.41***	2.44***
	(0.098)	(0.13)
Project Score Variance	-0.38***	-0.25**
	(0.049)	(0.10)
Observations	7640	9520
Pseudo R^2	0.289	0.292

Table B6. M-Logit Version of the Choice Scenario Set 1 Results for Phase 1 and Phase 2 (replication)

The table shows results from estimating a multiple logit version of Equation (1) using choice-scenario data from choice scenario set 1 in the phase 1 experiment (Column 1) and for the phase 2 experiment (Column 2). The outcome variable is the rank of each of the 4 projects in a given choice scenario (the base outcome if the bottom ranked project). "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. All variables in the model aside from a constant term are shown in the table. Standard errors, clustered at the subject level, are in parentheses, * p < 0.10, ** p < 0.05, *** p < 0.01.

	(1) Portfolio choice	(2) Portfolio choice	(3) Portfolio choice	(4) Portfolio choice
Avg. Variance	-0.54***	-0.0042***	-0.52***	-0.0039***
	(0.068)	(0.00034)	(0.068)	(0.00033)
Avg. Mean	4.43***	0.010***	4.44***	0.010***
	(0.33)	(0.00039)	(0.33)	(0.00038)
Budget	-0.20***		-0.18***	
	(0.013)		(0.025)	
Avg. Cost	-4.62***	-0.0045***	-4.62***	-0.0047***
	(0.41)	(0.00036)	(0.41)	(0.00036)
Avg. Variance x Budget			0.066*	0.0015***
			(0.035)	(0.00014)
Avg. Mean x Budget			0.0015	-0.0019***
			(0.015)	(0.00023)
Model	C-Logit	Linear Reg.	C-Logit	Linear Reg.
Observations	210,302	210,302	210,302	210,302
Subjects	150	150	150	150

Table B7. Effect of Budget in Choice Scenario Set 3 (Portfolio Choice)

The table shows conditional logit and linear regression estimates of portfolio choices, estimated on the choice set data from the third set of choice scenarios in the phase 1 experiment. The choice set is all portfolios that had total cost less than or equal to the subject's budget. The outcome variable is an indicator for the chosen portfolio. *Budget* is an indicator equal to 1 if the budget is higher than median. *Avg. Variance* is the average variance of the projects in the portfolio. *Avg. Mean* is the average mean of the projects in the portfolio. *Avg. Cost* is the average cost of projects in the portfolio. All explanatory variables are standardized. In addition to the shown variables, Column 1 contains subject fixed effects. Column 2 further contains fixed effects for the interaction of choice scenario and budget (which is why the "Budget" coefficient is excluded). Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)
	Phase 1 Experiment	Pooled Phase 2	Arm 1: Replication	Arm 2: No Loss Framing	Arm 3: Objective Returns
Dependent variable:	Project choice	Project choice	Project choice	Project choice	Project choice
Average project value	2.19^{***}	2.16^{***}	2.51^{***}	2.08^{***}	2.01^{***}
	(0.12)	(0.13)	(0.17)	(0.28)	(0.19)
Project variance	-0.64^{***}	-0.57^{***}	-0.55^{***}	-0.48^{***}	-0.69^{***}
	(0.053)	(0.055)	(0.086)	(0.089)	(0.11)
Observations	27000	25200	8280	8460	8460
Subjects	150	140	46	47	47

Table B8. Choice as a Function of Mean and Variance in Choice Scenario Set 1: Comparison of All Treatment Arms using Conditional Logit

All models are estimated using conditional logit and contain subject-by-choice scenario fixed effects. The dependent variable is equal to 1 if the subject chose the project and 0 otherwise. The average project score and project score variance are both standardized. Standard errors clustered at the subject level are in parentheses. Significance indicated by * p < 0.10, ** p < 0.05, *** p < 0.01.

Dependent variable:	(1) Project Choice
Average Utility Weight	
Average Score	5.41***
	(1.32)
Score Variance	0.74
	(0.56)
Average x Coef. RRA	-1.48
	(1.31)
Average x Unspent budget	-0.60
	(0.37)
Average x Work in R&D	0.29
	(0.32)
Average x Discount rate	0.86
	(0.63)
Average x Math classes	-3.41***
	(0.88)
Average x Decision science classes	2.52***
	(0.66)
Average x Mfin program	-0.23
	(0.41)
Average x Pr(higher prize)	2.57***
	(0.70)
Variance x Coef. RRA	-1.25*
	(0.66)
Variance x Unspent budget	-0.20*
	(0.11)
Variance x Work in R&D	0.23*
	(0.13)
Variance x Discount rate	0.12
	(0.24)
Variance x Math classes	-0.019
	(0.18)
Variance x Decision science classes	-0.32*
	(0.19)
Variance x Mfin program	-0.39**
	(0.16)
Variance x Pr(higher prize)	-0.073
	(0.16)
Utility Weight Heterogeneity	
Average Score	0.35***
	(0.083)
Score Variance	0.20*
	(0.10)
Average x Coef. RRA	0.17
	(0.29)
Average x Unspent budget	1.31***
	(0.50)
Average x Work in R&D	-0.74***
	(0.15)
Average x Discount rate	-0.96***
	(0.23)
Average x Math classes	1.30***

	(0.18)
Average x Decision science classes	1.24***
	(0.16)
Average x Mfin program	0.50***
	(0.12)
Average x Pr(higher prize)	-1.04***
	(0.24)
Variance x Coef. RRA	-0.030
	(0.086)
Variance x Unspent budget	0.077
	(0.26)
Variance x Work in R&D	0.30**
	(0.12)
Variance x Discount rate	-0.26**
	(0.11)
Variance x Math classes	0.29**
	(0.11)
Variance x Decision science classes	-0.56***
	(0.13)
Variance x Mfin program	0.56***
	(0.082)
Variance x Pr(higher prize)	-0.19
	(0.12)
Tau	0.27***
	(0.058)
Model	G-MNL
Observations	13,500
Subjects	150

Table B9. Simultaneous Estimates of Observable Dimensions of Heterogeneity, Phase 1, Choice Scenario Set 1

The table shows results from estimating Equation (1) using choice-scenario data from choice scenario 1 of the phase 1 experiment. The outcome variable is an indicator for the chosen project. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. The interactions are with dimensions of heterogeneity summarized in Section 5.2.1. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Overall, the models show that elicited loss aversion did not strongly mediate the effect of project variance. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Dependent variable:	(1) Project Choice	(2) Project Choice
Average Utility Weight		
Average Score	7.43***	3.21***
	(1.20)	(0.36)
Score Variance	-0.31**	-0.32***
	(0.14)	(0.095)
Average x α	-0.61*	-0.17
	(0.32)	(0.18)
Average x β	-0.16***	-0.050***
	(0.039)	(0.017)
Average x λ	0.11	0.032
	(0.080)	(0.046)
Variance x α	-0.022	-0.0024
	(0.015)	(0.019)
Variance x β	-0.042*	-0.035**
	(0.023)	(0.016)
Variance x λ	0.015***	0.0089*
	(0.0038)	(0.0048)
Utility Weight Heterogeneity		
Average Score	0.20**	
	(0.093)	
Score Variance	0.63***	
	(0.12)	
Average x α	0.0052	
	(0.0065)	
Average x β	0.0048	
	(0.0072)	
Average x λ	-0.00070	
	(0.00081)	
Variance x α	-0.0020	
	(0.0073)	
Variance x β	-0.012	
	(0.014)	
Variance x λ	-0.00066	
	(0.0043)	
Tau	0.76***	
	(0.077)	
Model	G-MNL	C-Logit
Observations	4,140	4,140
Subjects	46	46

Table B10. Heterogeneity by Elicited Loss Aversion (Phase 2)

The table shows results from estimating Equation (1) using choice-scenario data from choice scenario 1, arm 1 of the phase 2 experiment. The outcome variable is an indicator for the chosen project. "Average project score" is the average of the five scores for the project. "Project score variance" is the variance of the scores. Both explanatory variables are standardized. The interactions are with measures of risk and loss aversion. The parameters correspond with those given in the description of the loss aversion preference elicitation in Section A.6. α is risk aversion over the gain domain, β is risk aversion over the loss domain, and λ is the coefficient of loss aversion. All models contain subject and choice scenario random effects in addition to the variables shown in the table. Overall, the models show that elicited loss aversion did not strongly mediate the effect of project variance. Standard errors, clustered at the subject level, are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Statement summary	Frac. of particip-ants	Frac. cond. on var. maximizing	Example statements
Mean-variance tradeoff	0.71	0.0	"I looked for low price projects with high scores and low variance."
Variance loving	0.21	1.0	"I went with the highest number and then often the highest number with the highest variance. Some people can have a difference of opinion."
Loss aversion	0.12	0.05	"I tried to only fund projects with a rating of over 3 even when considering the variance. I tried to only risk a loss if the R&D project was $1 million."
Ambiguity aversion	0.02	0.0	"There's a lot of gray area, do some better investments return 10000x or more like 100x in this scenario?"
Made statement contrary to instruction	0.01	0.0	"I left $1 million in R&D budget after each round because I was instructed to not fully spend it."
Looked at individual scores	0.26	0.52	"I used a '5' as indication that a Project Hero could be identified within the company to carry the project through to completion"

Table B11. Summary of Participant Debriefing Following Phase 1 Experiment

The table shows summary statistics for the number of participants who gave responses of different types during the debriefing after the Phase 1 experiment. The column "Frac. cond. on var. maximizing" is the fraction of participants who made each statement conditional on that participant indicating that they maximized variance across their choices.

Dependent variable:	(1) Var. maximizing	(2) Var. maximizing
Project mean pref.	-1.38*	
	(0.74)	
Project var. pref.	2.01***	
	(0.62)	
Risk neutral		0.015
		(0.048)
Risk loving		0.17*
		(0.10)
R&D work experience		0.20***
		(0.068)
	(0.080)	(0.046)
Model	Linear regression	Linear regression
Observations	150	150

Table B12. Correlation Between Stated Variance Loving Behavior and Project Choices or Participant Covariates

The table shows results from estimating linear regressions using the variables shown in the table and participant-level data from the phase 1 experiment. The outcome variable is an indicator for the participant saying that they tried to maximize variance in the debriefing (see Table B11). "Project mean pref" is the participant-specific estimate of the utility weight placed on average project scores based on the G-MNL model reported in Table 2. "Project var. pref." is the same but for project score variance. "Risk neutral" is an indicator for the elicited coefficient of relative risk aversion being near 1, and "risk loving" is an indicator for a coefficient greater than 1. "R&D work experience" is an indicator for prior work in the R&D sector. Heteroskedasticity robust standard errors are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C. Additional Figures

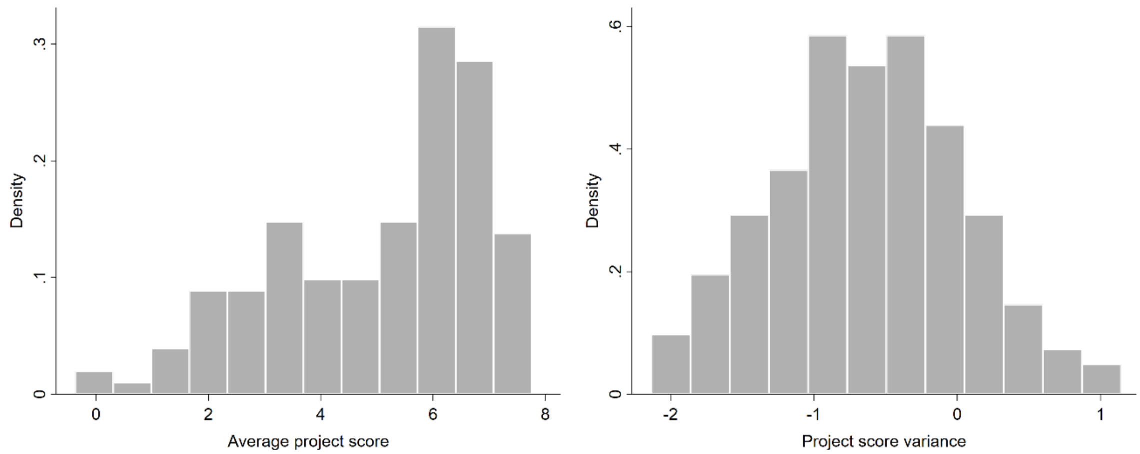

Figure C1. Heterogeneity of Estimated Preference Parameters, Phase 1, Choice Scenario Set 1

The figures show estimates of the individual-level preference parameters based on the estimates from the G-MNL model in Table 2. The left panel shows estimates of the preference for higher project mean score and the right panel shows the same for higher project score variance.

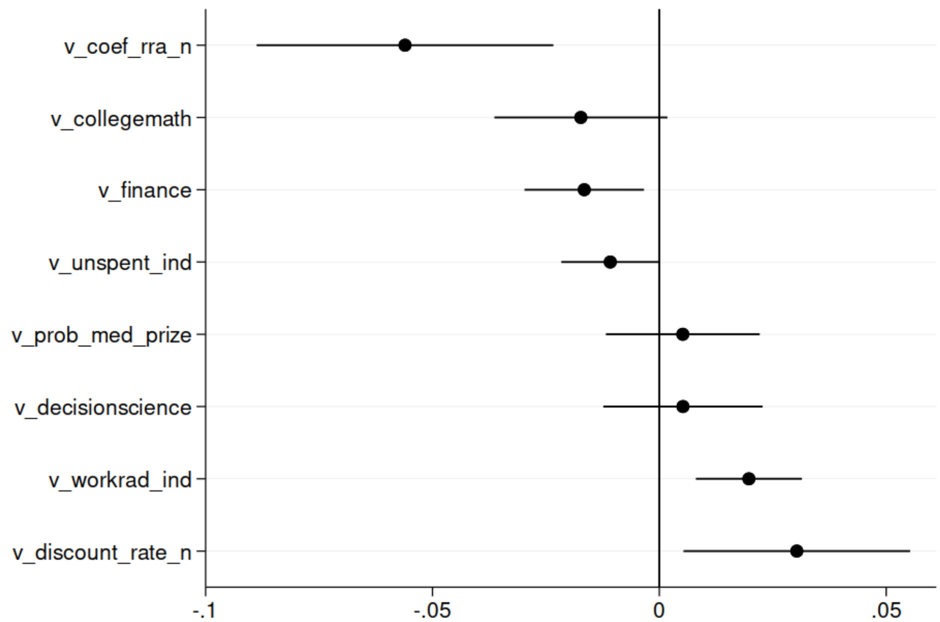

Figure C2. Marginal Effects of Variance Interacted with Dimensions of Heterogeneity, Simultaneously Estimated

The figures show marginal mediation effects of project score variance interacted with different subject demographics, elicited preferences, and characteristics. The estimates are from a single G-MNL model estimated on data from the phase 1 experiment, first set of choice scenarios. The data and estimation are identical to Table 3 from the body of the paper except that all dimensions of heterogeneity are included in the regression at the same time.