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.

Contact: rcarson@ucsd.edu (RTC); jgraffzivin@ucsd.edu (JGZ); Jordan.Louviere@unisa.edu.au (JJL); ssadoff@ucsd.edu (SS); jgs2103@columbia.edu (JGS)

Keywords: Innovation, Research & Development, Risk, Incentives, Hiring

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 three 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. We also elicited participants' personal risk preferences and found that decision-makers who are more tolerant of risk were more likely to fund breakthrough projects. The results suggest that the risk preferences of managers in charge of research investments may have an oversized effect on the rate of breakthrough innovation and the profitability of firms.

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 (Bloom et al. 2017). 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 also 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). As such, the key to a successful R&D program is its ability to encourage appropriate risk taking—one that tolerates failure in pursuit of reward (March 1991, Manso 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 multiple 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."

In this paper, we describe the results of three experiments designed to understand the role of uncertainty in shaping a manager's decision to fund R&D projects. The experiments were conducted with 191 Master of Business Administration (MBA) and Master of Finance students in a program focused on the intersection of business and technology. Many of these students will go on to work at investment firms or serve as managers making R&D decisions at companies in the health and technology sectors.

To mirror 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. They were asked to choose their preferred research projects from a series of hypothetical proposals that had been judged and scored by an objective, outside 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] Participant compensation was determined by the performance of the R&D projects that the participants chose to fund and was designed to mirror the real-world incentives for risk taking in R&D decisions (Dasgupta and Maskin 1987). Payments were based on a competitive, "tournament" structure, with the highest scoring participants receiving a substantially larger monetary reward than their peers. This incentive structure disproportionately rewarded participants for choosing higher-variance (i.e. riskier) projects.

Each participant took part in three choice experiments. The first experiment assessed whether the incentives to choose high-variance projects led to such choices among participants. In the experiment, each participant was presented with ten scenarios where they chose from four potential projects. 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 projects first-order stochastically dominated choosing a lower variance project. Despite this, participants routinely chose dominated projects. Holding average score constant, participants were, on average, significantly *less* likely to choose a project as variance

[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.

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—75% 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.

The second experiment assessed whether a simple informational intervention could lead participants to choose higher-variance projects. Participants were shown the same information as in the first experiment, but in addition to the average scores we reported the variance of project scores. The variance was straightforward to infer from the individual scores in the first experiment, so the second experiment studied the effect of increasing the salience of score variance. Increasing the salience backfired, leading subjects to engage in even more variance averse behavior than in the first experiment. The effects were strong enough that the majority of subjects were willing to choose a project with lower average score just to avoid choosing a higher variance project, a result that is clearly at odds with the incentives to fund breakthrough projects.

The third experiment asked participants to construct eight portfolios of research projects to assess whether the ability to construct a research portfolio would encourage more risk taking, and to determine the effect of funding scarcity on decision-making. As in the second experiment, the projects were presented with their individual scores, average score, and variance. In addition, each project had a cost ranging from $1 to $10 million. Participants were randomly assigned a budget, and they chose projects to fit within that budget. We find that, consistent with the first two experiments, individuals continued to make variance-averse choices. And, that aversion to variance was most pronounced under smaller budgets.

The lack of appropriate risk taking by most subjects in our experiment points to potential inefficiencies in the research investment process, and the results could help explain low rates of breakthrough innovation. To be more concrete, we present an example based on our empirical results to highlight the effect this behavior could have on breakthrough advances. In the experiments, subjects were shown hypothetical projects with ratings from an outside scientific advisory committee. Each rating was on a 1 to 5 scale, with 5 being the highest rating. For tractability and illustrative purposes, assume that the project ratings are an accurate representation of the expected quality of a proposal and that ratings of 5 indicate the potential for a breakthrough. Now 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. In this example, 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 16 percentage points less likely to choose the second project, despite the fact that the first project has no chance of producing a breakthrough (i.e., an outcome of the highest possible quality) and the second project has a 43% chance of doing so.

Our results also show that explicit risk-taking incentives might not be enough to encourage optimal R&D within a firm. Our examination of participant preferences points to an alternative solution. The degree of variance aversion in participants' investment choices was significantly related to a participant's baseline risk preferences. Participants who were generally risk loving chose projects with higher variance, on average, than risk neutral or risk averse subjects. As a result, risk-loving participants performed better, on average, on the experimental tasks and chose projects more in line with optimal theory. 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.

2. Literature

Theoretical models of optimal R&D argue that firms should invest in high variance research projects. An important early contribution to this literature is Dasgupta and Maskin (1987). They point out that both from the private perspective of a firm/scientist, and from the perspective of society, the spoils from R&D discoveries are disproportionately skewed toward novel, high-quality discoveries. For firms, initial discoveries are rewarded both through institutional arrangements (patents, for example) and first-mover advantages in the market. For academic scientists, the goal is to produce novel ideas, and being the first to discover and publish a finding confers large rewards (Merton 1973). From society's perspective, technology is a public good, so Dasgupta and Maskin argue that social surplus is maximized when only the best available technology is employed. For instance, if two researchers develop alternative manufacturing techniques for producing the same good, surplus is greater if only the better of the two techniques is employed. 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]

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 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.

Aside from the incentives provided by highly skewed compensation schemes, choosing high variance research projects can be valuable for exploring the potential space of investments. Allen (1991) shows theoretically that if R&D helps firms gain

[2] Later work has reinforced the insight of Dasgupta and Maskin (1987) while elaborating on cases where private incentives might lead to too much or too little risk taking in R&D relative to the social optimum. For example, Cabral (1994) argues that private risk-taking in R&D might be too low if researchers pay insufficient attention to duplicative efforts by competing researchers. One caveat is that Dasgupta and Maskin (as well as follow-up paper by Cabral) measure variance by the diversity of research projects undertaken by different scientists. These authors primarily focus on the question of duplication of research effort. In our setting, we abstract from duplication of effort and instead focus on the effect of convex payoffs.

information in addition to generating returns, then higher variance projects provide more opportunity for a firm to learn how to distinguish itself from rivals.

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 whether explicit incentives can overcome suboptimal risk taking in R&D decisions within a firm.

Some research has demonstrated that preferences play an important role in determining firm innovation. 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).

The results in this paper speak to the need to think about individual preferences and the behavior of managers and workers below the CEO level who might oversee innovation investments for the firm. 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 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. Indeed, recent evidence suggests that less risk averse individuals generate more novel inventions by pursuing riskier innovation strategies (Graff Zivin and Lyons, 2020).

3. Experimental Design

3.1. Experimental Setup

The experiments were implemented among master's degree students enrolled at the Rady School of Management of the University of California, San Diego—a program focused on the intersection of business and technology. The typical student has five 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.

Study participants were asked to assume the role of the head of a research division of an organization considering whether to fund project proposals based on ratings from a scientific advisory panel (see Appendix for the instructions). Each participant took part in three experiments. In the first experiment, 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 sets of research projects, with each set characterized by different reviewer score profiles.

In the second experiment, the same procedure was repeated for ten more sets 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 experiment did not provide more information than the first one. It did emphasize the variance more than the first experiment. An example of the initial project choice screen is shown in the Appendix.

The third experiment 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 experiment, 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.

By design, participants 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.[3] 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 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 in experiment 3 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 were summed to create the total score for the participant.

We then publicly awarded prizes to the top performers in each session: the top 25% of scores received $25 and the top 10% of subjects received $100 at the end of the session. All subjects received a $15 participation fee. Because we offered large

[3] Specifically, participants were told: "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."

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.

At the end of the experiment—after participants made their decisions but before learning of their performance—subjects completed 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 subjects' coefficient of relative risk aversion (details on this calculation can be found in the Appendix).

We implemented the experiments during regularly scheduled class sessions of the MBA and Master of Finance programs. 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 and received payment at the end of the session, which lasted about an hour.

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. The final sample therefore contained 191 subjects. Each subject faced 10 choice scenarios in experiments 1 and 2. Each scenario involved choosing between 9 potential options, yielding 17,190 total observations for each experiment. The options in experiment 3 varied by budget, which was randomized across subject. The average subject had 1,399 options, leading to a total sample size for experiment 3 of 267,210 observations.[4]

3.2. Design of the Discrete Choice Experiments

The design for the experiments presented to the subjects builds on models of random utility theory to estimate discrete choice models using decisions from discrete choice experiments. 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 and McFadden's models. 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 work. 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

[4] More specifically, in both experiments 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 experiment 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.

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.

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 array the 16 combinations into 20 sets of four "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. In order 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. We then randomly blocked each of the two versions of the DCE—Version I and Version II—into two subsets of 10 choice sets each, Subset A and Subset B. Within each version, we randomized the order of the subsets, "AB" or "BA". This produced four treatment groups: Version I.AB, Version I.BA, Version II.AB and Version II.BA.

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 treatment groups discussed above (i.e., Version I.AB.1, Version I.BA.2, Version II.AB.3, Version II.BA.4).

Participants were randomly assigned to one of the four treatment groups, stratified on session. Approximately 48 subjects participated in each treatment. The order in which projects were presented within each treatment group was also randomized across sessions.[5] 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

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

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

[5] Due to a coding error, the project order was not explicitly randomized in the first session, causing higher mean value projects to be presented further to the left on each question. This ordering could have led subjects to choose higher mean projects if the order in which projects appeared influenced choices. We present robustness to the exclusion of this session in the appendix and show that the results are unchanged.

$$\beta_i = \sigma_i \beta + \eta_i \quad (2)$$

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 $\overline{\sigma} + \theta z_i$ 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]

In the analysis of the final discrete choice experiment which involved choices over budget, we estimate 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. Budget was randomly varied within subject, across choice scenario. Therefore, we rely on between-subject comparisons that preclude the use of individual and choice scenario-specific heterogeneity parameters.

[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.

We verify that the randomization was balanced on observable characteristics in the appendix.

5. Results

5.1. Project Choice

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. The repeated, within-subject sampling of the experimental design allows us to estimate generalized multinomial logit (G-MNL) models that further control for individual characteristics by including indicator variables for each subject.

Table 1 shows results for project choice as a function of the mean and variance of project ratings. In all columns, the explanatory variables are standardized to have an average value of zero and standard deviation of one by subtracting the sample average and dividing by the sample standard deviation. The dependent variable is an indicator equal to one if the subject chose the project.[7] Column 1 reports results from participants' first decision task in which they answered ten questions, each asking them to rank four possible projects. Recall that for each project, the participant was shown individual ratings from seven outside reviewers as well as the average rating across reviewers. Column 2 shows results from the second set of ten questions that

[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 1, 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.

was identical to the first set except that each question also included a row displaying the variance of the projects (see Figure S1).

The coefficients in the top portion of the table (labelled "Mean") can be interpreted as measures of the average utility weight that subjects placed on different attributes. The positive coefficient on average project score indicates that respondents generally preferred projects with higher average ratings, conditional on project variance. The association of project variance and selection is *negative*. That is, on average, subjects preferred a project less if the variance of the scores was higher. In the second experiment, where subjects were explicitly shown the variance, this association was even more strongly negative than in the first experiment. Rather than helping individuals make the correct choice, highlighting the variance pushed participants further away.

The second section of the table (labelled "Standard Deviation") reports estimates of the diagonal elements from the variance-covariance matrix of i. Larger estimates indicate more idiosyncratic preference heterogeneity across subjects. The estimates indicate substantial, statistically significant heterogeneity in how subjects responded to the variance of project ratings. We discuss interpretation of relative magnitudes below using implied choice probabilities given project attributes, as shown in Figure 1.

The third section of the table reports the estimate of the standard deviation of individual-level scale heterogeneity. The significance of this coefficient shows that there was important scale heterogeneity in both sets of responses, supporting the use of the G-MNL model for analysis of these data.

Figure 1 summarizes the results for the first two experiments by showing average choice probabilities implied by the results in Table 1. Panel (a) shows the results for the first experiment and Panel (b) shows the results for the second experiment. The figure plots the average probability that a subject would choose a project with a given mean and variance.[8] For illustrative purposes, the figure shows the probability of choosing a project with three different mean scores: low, medium, and high. The low mean score is equal to 3 out of 5 and corresponds to the 25th percentile of scores shown to subjects in the experiment. The medium mean score is equal to 3.5 out of 5 and is the median score in the experiment. Finally, a high mean score is equal to 4 out of 5 and is the 75th percentile. Therefore, comparing the bottom line in Figure 1 to the top line shows the difference in the probability that a subject would choose a project given a one-point increase in the average score, moving that project from the 25th to 75th percentile of scores. Similarly, the figure shows choice probabilities for a range of project variances. A project with low variance (the 25th percentile of projects shown to subjects) has a variance of 0.45 while a project with high variance (the 75th percentile) has a variance of 1.79. The average project shown to subjects had a variance of 0.95.

The results indicate that in both the first and second experiments, subjects strongly preferred to choose projects with higher mean scores. This is shown by the size of the gap between the solid lines in the figure. Moving from a project with a moderate score (3 out of 5) to a project with a high mean score (4 out of 5) increases the probability that a project was chosen from 0.28 to 0.64, more than doubling the likelihood a subject chose that project. Overall, subjects chose the project with the highest available mean score 70% of the time in both of the first two experiments, regardless of the variance. These results suggest that our subjects were responsive to the incentives in the experiment to choose higher quality projects. However, as discussed below, they were not responsive to the incentives to choose higher variance projects.

The figure also shows that subjects were consistently less likely to choose projects if the projects had higher variance, a pattern of behavior that is contrary to the incentives they faced. In the first experiment, Figure 1 Panel (a) shows that an

---

[8] Because subjects choose their first, then second, then third ranked project, the choice probability is measured for all instances in which a project remained in the choice set.

increase in the variance of a project consistently and significantly decreased the probability that the average subject would choose that project. If a subject saw a low variance project (variance of 0.45), they selected it 36% of the time, on average. In contrast, if the subject saw a high variance project (variance of 1.79) but the exact same mean score, they selected that project 31% of the time. This drop of 5 percentage points is large—a 14% decrease in the probability of selecting the project—and statistically significant (95% confidence interval of 3.7 to 6.3 percentage points). The effect of variance was stronger for projects with a higher mean. For a project with a high mean score, going from a low to high variance reduced the probability that a typical subject chose that project by 10.2 percentage points (95% confidence interval of 8.0 to 12.4).

In the second experiment that explicitly displayed project variance, the subjects exhibited even stronger aversion to choosing projects with higher variance. Comparing Panels (a) and (b), the overall effect of an increase in variance on the probability of choice was significantly stronger in the second experiment—the effect of high versus low variance on probability of choice was -5 percentage points in experiment 1 versus -9 percentage points in experiment 2—and the difference is statistically significant (*p=0.0002*).[9] Emphasizing variance made subjects more likely to pick low variance projects with low mean scores; and less likely to pick high variance projects with high mean scores. The effects are strong enough that subjects were more likely to choose a project with a moderate mean score and low variance than to choose a project with a high mean score and high variance—a decision that runs counter to the incentives in our experiment. That the simple act of reporting variance, which should have made it easier for subjects to respond to the incentives

[9] Based on estimating a model that included interactions between an indicator for the second experiment and project variance. The reported *p-value* is from the two-sided test that the effect of variance was equal across experiments 1 and 2 (n=34,380, z=3.73). Standard errors clustered at the subject level.

of the contest, led them to make more risk averse choices is quite surprising. As shown below, this pattern of results holds across various subsamples of the data.

5.2. Portfolio Selection

Finally, we explore the role of financial scarcity in driving conservatism in R&D strategies. The third experiment asked participants to choose projects for a set of eight different portfolios with eight distinct budgets. The goal of the portfolio selection experiment was to understand how subjects' aversion to variance responded to explicit variation in the amount of money they had available to spend on research projects.

Like the first and second experiment, the average subject strongly preferred portfolios with high mean scores and low variance. 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 value but different variance. For instance, in the portfolio choice scenario 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). Based on the incentives, the subject choosing between two portfolios with the same mean score would have had a higher chance of winning larger prizes if he or she chose the portfolio with the higher variance. However, at all budget levels, most subjects (75%) chose the *lower* variance project.

We analyze the relationship between budget and variance more formally in Table 2, which reports results from estimating conditional logit models on data from the third discrete choice experiment. The dependent variable is equal to 1 if a subject chose a given portfolio and 0 otherwise. The choice set is all portfolios that were feasible given the offered budget (so all portfolios with total cost equal to or less than the budget amount). The model includes choice-scenario fixed effects, so the results are estimated off of differences across individuals who faced the same choice scenario.

The "Avg. Variance" and "Avg. Mean" coefficients show that, on average, subjects preferred portfolios comprised of projects with higher mean scores and lower score variance, consistent with the choices in the project selection questions.[10] Given that these estimates come from conditional logit models, interpretation of the coefficients is more straightforward than for the G-MNL models reported above. The "Avg. Mean" coefficient indicates that a portfolio, where the average project had a mean value one standard deviation higher than a portfolio comprised of projects with the sample average mean score, was $\exp(0.7) \approx 2$ times more likely to be selected. This effect is large. In practice, subjects almost always chose projects with the highest or nearly the highest mean scores.

The effect of variance is less extreme, but the results indicate that increasing the project-average variance of a portfolio by 1 standard deviation, relative to a portfolio with the sample average variance, decreased the odds of selection by 32% (the odds ratio for choice of a portfolio that has projects with average variance one standard deviation higher than average is $\exp(-0.38)=0.68$).

The main effects of interest are the interactions between budget and both average variance and average mean. Budget was randomized across individuals, and the models include choice scenario fixed effects, so the results can be interpreted as the effect of a 1 million dollar increase in budget, holding all available projects fixed. The negative coefficient on budget interacted with mean indicates that subjects chose projects with lower means as their budget is expanded, but this effect is not statistically significant.

The positive coefficient on the interaction between budget and variance shows that as the budget constraint was relaxed, individuals chose higher variance projects on average. In this case, the interaction is statistically significant at the 10 percent

[10] By *average* variance and *average* mean, we are referring to the average of the projects in the portfolio.

level.[11] This effect reflects the same preferences exhibited in the previous experiments. With small budgets, subjects' portfolios were almost entirely composed of low-variance projects. With larger budgets subjects chose portfolios with a higher proportion of high-variance projects. This effect could be partly mechanical—subjects with high budgets might not have had additional low-variance projects available in their choice sets—or it could be due to changes in preference for high versus low-variance projects at different budget levels. Whatever the case, these choice scenarios mimic real-world choices of how to allocate a budget over a fixed set of R&D options, and the results suggest that one way to increase the variance of choices is to offer a larger budget.

Figure 2 summarizes the relationship between the probability a subject chose a portfolio and the variance of the average project in that portfolio, stratified by 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 (two-sided test of equivalence between slopes clustered by question, $F(1, 31)=7.56$, $p=0.0098$) 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. These results suggest

[11] The "Budget" coefficient is a pure control. Because the choice set includes all portfolios with total cost less than or equal to the budget, larger budgets increased the size of the choice set. Therefore, a larger budget makes selection of any one portfolio less likely. The negative coefficient on budget results from this mechanical relationship. The final coefficient, "Avg. Cost", is identified from the sub-sample of subjects who did not exhaust their budgets. For subjects who did exhaust their budget, budget and cost are collinear. We include the cost variables as controls to ensure that the budget variation can be interpreted for both the subjects who did and did not exhaust their budgets.

that while decision-makers are risk averse at all budget levels, they are relatively more willing to take on risk when budgets are less constrained.

5.3. Explaining Project Choice

We now investigate the factors that might explain why subjects preferred projects with lower variance. We saw from the "Standard Deviation" results in Table 1 that there was substantial heterogeneity in individual choices, and the results here explore some observable dimensions of heterogeneity. We first assess the role of inexperience with the setting, inattention, or misunderstanding of the incentives. We then examine the role of individual risk preferences. The results are shown in Table 3 and are summarized in Figure 3.

The results on experience are shown in Columns 1.A and 2.A of Table 3, where "Worked in R&D" is an indicator variable equal to 1 for all subjects who reported that they had work experience in R&D (37% of the sample reported some experience). If inexperience is driving the aversion to risk in the overall sample, we would expect subjects who have worked in R&D to respond positively to higher variance projects. We find some evidence that this is the case. The coefficients in Column 1.A, for example, imply that for subjects with R&D experience, a one standard deviation increase in project score variance led to a 0.005 point increase in the probability of that project being chosen. This is in contrast to the 0.02 point decrease for subjects who did not have R&D experience. R&D experience does not have a significant effect on the standard deviation of choice, as the second panel of Table 3 indicates.

Figure 3 shows the effect of changing score variance from low to high (from 0.45 to 1.79) on the probability of choice for a project with a high mean score (4 out of 5). In each section of the figure, we estimate this effect in experiment 1 and experiment 2 for the different subsamples of the participants corresponding to the different columns of Table 3.

The top set of results show the role of prior experience in R&D estimated in Table 3 Columns 1.A and 2.A. The figure shows that while both groups prefer projects with lower variance, subjects without R&D experience are particularly sensitive to the emphasis on variance in experiment 2, where subjects without R&D experience were about 50% more variance averse than subjects who did have experience. Further results in the Appendix examine the role of academic training. The results show that subjects who had taken more finance classes were more variance averse, on average, than subjects who had taken fewer finance classes. The number of mathematics courses taken was not strongly related to behavior in the experiments.

We next examine the role of inattention or confusion about the structure of the incentives. We proxy for inattention using the 32% of subjects who did not exhaust their budgets in the portfolio construction experiment (experiment 3). We do so because subjects had a strong incentive to spend their entire budget and failure to do so is consistent with the notion that subjects were either confused about the incentive structure or unmotivated to devote the attention needed to succeed in the experiment. The results show that budget use was not a strong predictor of how subjects responded to variance. As can be seen in Column 1.B, subjects who spent all of their budget behaved in a manner largely consistent with the rest of the sample on the initial project choice tasks.

We do find evidence that when project variance was emphasized in the project choice questions (column 2.B), subjects who exhausted their budgets chose projects with higher variance, on average, than subjects who did not exhaust their budgets. However, subjects from both groups exhibited statistically significant variance aversion in both decision experiments. Figure 3 reinforces this point. In both experiments, the response to an increase in variance between the two groups was similar and not statistically distinguishable. These results suggest that the aversion to risky projects in our main results cannot be explained by possible confusion about or inattention to the experiment.

Finally, we turn to the effect of individual risk preferences. Of the subjects, 52% were risk averse, 36% were risk neutral, and 12% were risk loving. The results in Columns 1.C and 2.C of Table 3, as well as the third section of Figure 3, show that more risk loving subjects were consistently more likely to choose higher variance projects than their risk averse peers. In the table, compared to risk averse subjects (the omitted category), risk neutral and risk loving subjects consistently chose higher variance projects, on average. In the first experiment, when variance was not emphasized, risk loving individuals were nearly indifferent to higher risk, on average. In the second set of choice experiments, with the variance emphasized, an increase in score variance was associated with a slight decrease in choice probability for risk loving subjects, but the effect was less than half the size of that for risk averse subjects (the difference, however, is not statistically significant, $p=0.55$).

Figure 3 shows that for the example project, risk averse subjects were about 50% less likely to choose a high variance project than risk loving subjects, but the difference was only statistically significant at the 10% level ($p=0.07$).[12] In the second experiment, the differences are even stronger: risk loving subjects were substantially and significantly less variance averse than the other subjects ($F(1, 190)=19.93$, two-sided p-value<0.0001 on the test of equivalence between risk averse and risk loving subjects and $F(1, 190)=3.57$, two-sided p-value=0.06 on the test of equivalence between risk neutral and risk loving subjects). These results suggest that individuals' risk attitudes are an important and consistent driver of their response to variance when choosing projects.

6. Discussion

Anemic research pipelines and the apparent slowdown of paradigm-shifting discoveries over the past quarter century has 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 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).

[12] Calculated by regressing predicted probability of choice on the three risk categories and using a two-sided *F-test* of equality between coefficients for the risk averse and risk loving groups, with standard errors clustered at the subject level ($F(1, 190)=3.23$).

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 subjects 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.[13] 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. Taking preferences into account could be a way to increase the productivity of scientific research. Whether risk preferences and the decisions they engender are, in fact, malleable is an open question that may have important implications for R&D choices, the advancement of science, and the fate of research-intensive firms.

Acknowledgments

The authors gratefully acknowledge the financial support of the National Science Foundation through its SciSIP Program (Award SBE-1561257); The protocol was approved by the UC San Diego IRB, protocol number 160212.

[13] 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.

References

Aghion, Philippe, and Peter Howitt. "A Model of Growth Through Creative Destruction." *Econometrica* 60, no. 2 (1992): 323–51. https://doi.org/10.2307/2951599.

Allen, Beth. "Choosing R&D Projects: An Informational Approach." *The American Economic Review Papers & Proceedings* 81, no. 2 (1991): 257–61.

Amit, Raphael, and Christoph Zott. "Value Creation in E-Business." *Strategic Management Journal* 22, no. 6–7 (2001): 493–520. https://doi.org/10.1002/smj.187.

Azoulay, P., J. S. Graff Zivin, and G. Manso. "Incentives and creativity: evidence from the academic life sciences." *The RAND Journal of Economics* 42, no. 3 (2011): 527-554.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb. "Are Ideas Getting Harder to Find?" *National Bureau of Economic Research* No. 23782 (2017): 53. http://www.nber.org/papers/w23782.

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

Boudreau, Kevin J., Eva C. Guinan, Karim R. Lakhani, and Christoph Riedl. "Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science." *Management Science* 62, no. 10 (2016): 2765-2783.

Cabral, Luís. "Bias in Market R&D Portfolios." *International Journal of Industrial Organization* 12, no. 4 (1994): 533–47. https://doi.org/10.1016/0167-7187(94)90006-X.

Cabral, Luís M.B. "R&D Competition When Firms Choose Variance." *Journal of Economics and Management Strategy* 12, no. 1 (2003): 139–50. https://doi.org/10.1162/105864003321220760.

Charness, Gary, Uri Gneezy, and Alex Imas. "Experimental Methods: Eliciting Risk Preferences." *Journal of Economic Behavior and Organization* 87, 43–51 (2013).

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

David, H. A., *The Method of Paired Comparisons.* (Griffin, London, 1988), vol. 12.

Fiebig, D. G., M. P. Keane, J. J. Louviere, N. Wasi, "The generalized multinomial logit model: Accounting for scale and coefficient heterogeneity." *Marketing Science* 29, no. 3 (2010): 393-421.

Goel, Anand M., and Anjan V. Thankor. “Overconfidence, CEO Selection, and Corporate Governance.” *The Journal of Finance* 63, no. 6 (December 2008): 2737–84. https://doi.org/10.1111/j.1540-6261.2008.01412.x.

Graff Zivin, Joshua and Elizabeth Lyons. "The Effects of Prize Structures on Innovative Performance." NBER Working Paper, no. 26737 (2020).

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

Hirshleifer, David, Angie Low, and Siew Hong Teoh. “Are Overconfident CEOs Better Innovators?” *The Journal of Finance* 67, no. 4 (2012): 1457–98. https://doi.org/10.1111/j.1540-6261.2012.01753.x.

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

Keupp, Marcus Matthias, Maximilian Palmié, and Oliver Gassmann. “The Strategic Management of Innovation: A Systematic Review and Paths for Future Research.” *International Journal of Management Reviews* 14, no. 4 (2012): 367–90. https://doi.org/10.1111/j.1468-2370.2011.00321.x.

Krieger, Joshua, Danielle Li, and Dimitris Papanikolaou. “Missing Novelty in Drug Development.” Cambridge, MA, (2019). https://doi.org/10.3386/w24595.

Lotka, A. J. “The Frequency Distribution of Scientific Productivity.” *J. Washington Academy of Sci.* 16, no. 12 (1926): 317–23

Louviere, J. J. and G. Woodworth, “Design and analysis of simulated consumer choice or allocation experiments: an approach based on aggregate data.” *Journal of Marketing Research* 20, no. 4 (1983): 350-367.

Louviere, J. J., T. N. Flynn, A. A. J. Marley, *Best-Worst Scaling: Theory, Methods and Applications.* (Cambridge University Press, 2015).

Luce, R. D. “On the possible psychophysical laws.” *Psychology Review* 66, no. 2 (1959).

March, James G. “Exploration and exploitation in organizational learning.” *Organization Science* 2, no. 1, (1991): 71-87.

Marks, Andrew R. "Repaving the road to biomedical innovation through academia." *Science translational medicine* 3, no. 89 (2011): 89cm15.

Manso, G. “Motivating innovation.” *Journal of Finance* 66, no. 5 (2011): 1823-1860.

McFadden, D. “The measurement of urban travel demand.” *Journal of Public Economics* 3, no. 4 (1974): 303-328.

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

Munos, B. H., W. W. Chin, “How to revive breakthrough innovation in the pharmaceutical industry.” *Science Translational Medicine* 3 89cm16 (2011).

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

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

Teece, David J. “Business Models, Business Strategy and Innovation.” *Long Range Planning* 43, no. 2–3 (April 2010): 172–94. https://doi.org/10.1016/j.lrp.2009.07.003.

Thurstone, L. L. “A law of comparative judgment.” *Psychology Review* 34, no. 4 (1927): 273.

Tishler, Asher. “How Risky Should an R&D Program Be?” *Economics Letters* 99, no. 2 (2008): 268–71. https://doi.org/10.1016/j.econlet.2007.07.006.

Figures and Tables

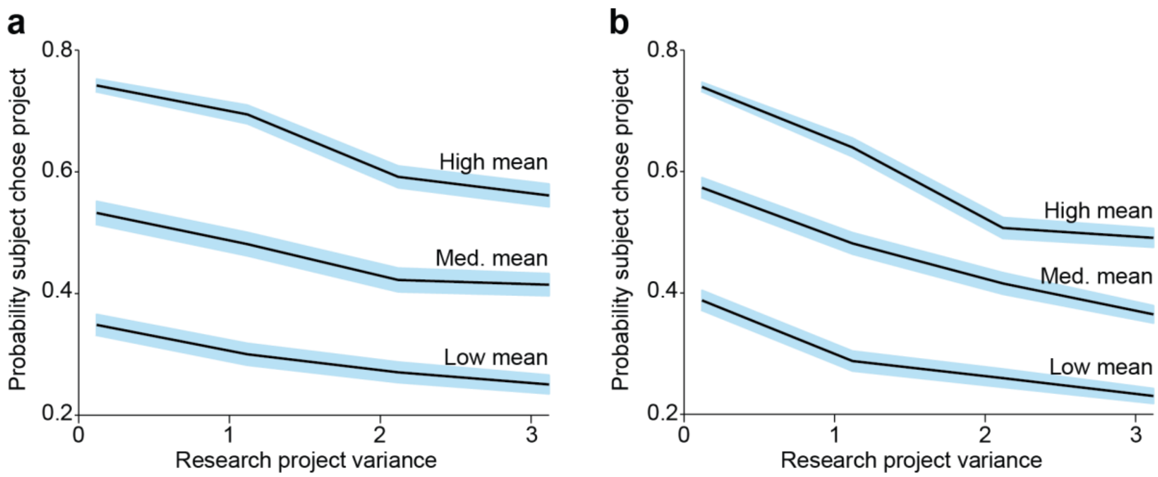

Figure 1. Likelihood of Choosing a Project with a Given Mean and Variance
The figure shows the average probability of a subject choosing a project with the given mean and variance. Each solid line shows how the choice probability changes as variance increases for three different mean project scores: 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. (**a**) Shows estimates based on data from the first experiment that did not explicitly show the variance of scores. (**b**) Shows estimates based on data from the second experiment that emphasized project variance.

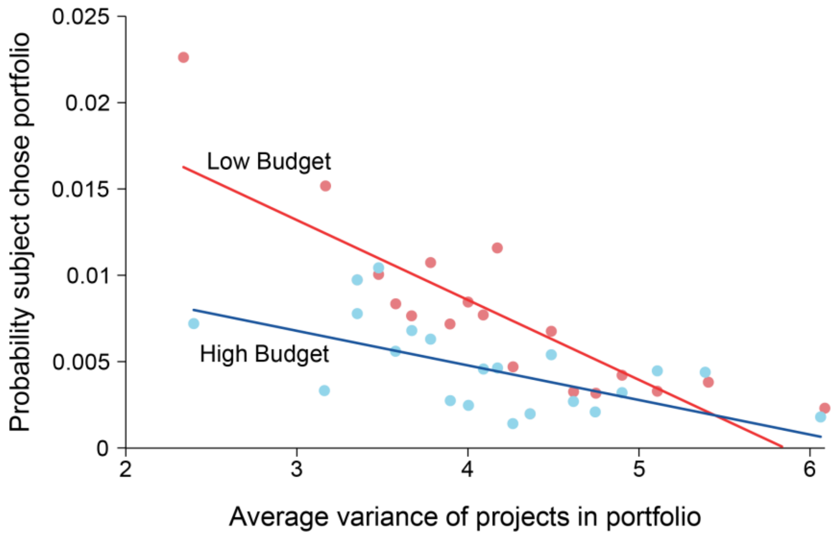

Figure 2. Effect of Budget on Preference for Portfolio Variance
The figure shows the average effect of variance on portfolio choice, broken down by budget. The red circles show the effect of variance on portfolio choice for low budgets (less than $15 million). The red line is a linear fit through the 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.

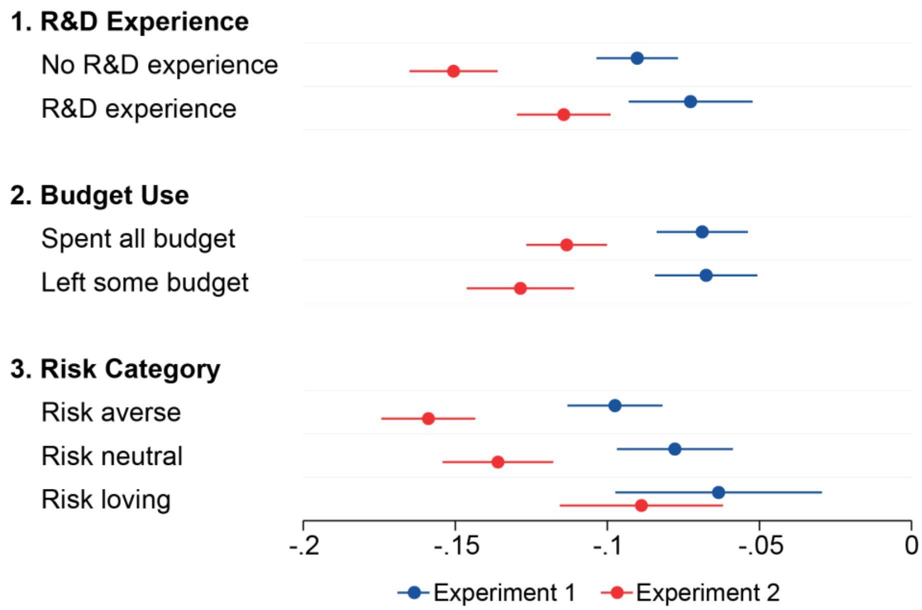

Figure 3. Heterogeneity in Effect of Variance on Likelihood of Choosing Project
The figure shows the effect of an increase in project score variance from 0.45 to 1.79 on the probability a subject chooses that project, broken down by subject characteristics. The estimated choice probabilities are derived from six G-MNL models reported in the Appendix. The blue circles show the effect of variance in experiment 1 (project scores and mean score shown). The red circles show the effect of variance in experiment 2 (variance also shown). The lines show the 95% confidence intervals based on standard errors clustered at the subject level. For section 1, the effect is broken down by whether the subject had prior experience with a R&D firm or organization. For section 2, the effect is broken down by whether the subject exhausted their budget in the portfolio choice experiment. For section 3, the effect of variance is broken down by the subject's baseline risk aversion.

	(1) Experiment 1 Project choice	(2) Experiment 2 Project choice, Variance emphasized
Mean		
Average Project Score	5.20***	4.54***
	(0.38)	(0.39)
Project Score Variance	-0.61***	-1.16***
	(0.056)	(0.10)
Standard Deviation		
Average Project Score	3.19***	0.092
	(0.27)	(0.13)
Project Score Variance	0.74***	1.94***
	(0.059)	(0.19)
Tau	0.14***	0.89***
	(0.035)	(0.060)
Observations	17,190	17,190

Table 1. Project choice as a function of mean and variance
The table shows results from estimating Equation (1) using data from experiments 1 and 2. Standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All models contain subject and choice scenario random effects.

	(1) Portfolio choice
Avg. Variance	-0.38***
	(0.11)
Budget	-0.10***
	(0.023)
Avg. Variance x Budget	0.014*
	(0.0075)
Avg. Mean	0.72***
	(0.064)
Avg. Mean x Budget	-0.0013
	(0.0011)
Avg. Cost	-2.94***
	(0.29)
Model	c-logit
Choice Scenario Fixed Effects	Yes
Demographic Covariates	No
Observations	267,210

Table 2. Experiment 3 Results for Portfolio Choice

Standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the choice scenario level. All models include fixed effects for the choice scenario as well as indicator variables for the experimental session. The table is the conditional logit estimate of the relationship shown in Figure 5. *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. The choice set is all portfolios that had total cost less than or equal to the subject's budget.

	Exp 1: Mean Score Shown			Exp 2: Mean and Variance Shown		
	(1.A)	(1.B)	(1.C)	(2.A)	(2.B)	(2.C)
Mean						
Average Project Score	5.73***	6.33***	5.82***	5.16***	5.05***	5.18***
	(0.37)	(0.69)	(0.49)	(0.43)	(0.34)	(0.55)
Project Score Variance	-0.71***	-0.59***	-0.77***	-1.60***	-1.22***	-1.54***
	(0.074)	(0.075)	(0.081)	(0.26)	(0.11)	(0.20)
Variance×Worked in R&D	0.38***			0.62*		
	(0.13)			(0.35)		
Variance×Exhaust Budget		0.080			0.60***	
		(0.12)			(0.23)	
Variance×Risk Neutral			0.27**			0.12
			(0.12)			(0.26)
Variance×Risk Loving			0.72***			0.61***
			(0.16)			(0.22)
Standard Deviation						
Average Project Score	3.83***	2.64***	2.77***	2.49***	2.04***	2.17***
	(0.28)	(0.30)	(0.24)	(0.30)	(0.16)	(0.26)
Project Score Variance	0.68***	0.72***	0.68***	1.52***	1.52***	1.52***
	(0.053)	(0.058)	(0.057)	(0.13)	(0.12)	(0.13)
Variance×Worked in R&D	0.17			0.17		
	(0.16)			(0.24)		
Variance×Exhaust Budget		0.25**			0.012	
		(0.099)			(0.13)	
Variance×Risk Neutral			0.046			0.11
			(0.092)			(0.14)
Variance×Risk Loving			0.20			0.98***
			(0.15)			(0.15)
Tau	0.34***	0.41***	0.015	0.18***	0.20***	0.046
	(0.019)	(0.061)	(0.036)	(0.025)	(0.042)	(0.039)
Observations	17,190	17,190	17,190	17,190	17,190	17,190

Table 3. Choice as a Function of Variance, Mean, and Subject Characteristics

Standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All models contain subject and choice scenario random effects. The estimating equation in given in Equation (1).