HOOVER IP²

WORKING PAPER SERIES
NO. 18003

PATENTLY RISKY: FRAMING, INNOVATION AND ENTREPRENEURIAL PREFERENCES

ELIZABETH HOFFMAN
Iowa State University

DAVID SCHWARTZ
Northwestern University Pritzker School of Law

Matthew Spitzer
Northwestern University Pritzker School of Law

ERIC TALLEY
Columbia Law School

JANUARY 8, 2018

Hoover Institution Working Group on
Intellectual Property, Innovation, and Prosperity
Stanford University

www.hooverip2.org

Patently Risky: Framing, Innovation and Entrepreneurial Preferences

Elizabeth Hoffman • David Schwartz • Matthew Spitzer • Eric Talley[1]
December 29, 2017

Abstract Innovation policy balances static monopoly rights against dynamic entrepreneurial incentives. In striking this balance, researchers commonly presume that decision makers in innovative settings react to their economic environments in a manner similar to their counterparts in other contexts. This paper reports on a series of experiments that call this presumption into question. Subjects were offered a choice between a sure thing and a risky choice, where our principal manipulation was to alter the decisional frame. Subjects in the control group confronted an unadorned choice between safe and risky options; subjects in the Invest Frame, in contrast, were told that the risky choice was tantamount to an investment in an innovation-related project. In all other respects, we controlled for both the language of the experimental instrument as well as the economic stakes entailed. In some subgroups, the risky choice also included a potential monetary loss. We administered the experiments across three subgroups of settings/subjects: a brick-and-mortar lab using university students, an Internet protocol again using university students, and an Internet protocol using Mechanical Turkers as subjects. Our main result, which appears quite strong and robust, is that the Invest Frame induced subjects to manifest greater degrees of risk tolerance on average, across all three settings, and across specifications involving positive and negative payoffs. We calibrate our results to an estimate of a downward "shock" that the invest frame introduces to subjects' coefficient of relative risk aversion, benchmarking our results against Holt and Laury's (2002) risk-aversion elicitation scale.

Keywords Risk aversion; framing; patents; intellectual property; entrepreneurship.
JEL Codes C9; C92; L2; L26; K1; O34.

Acknowledgments Many thanks to Stefan Bechtold, Christopher Buccafusco, Paul Edelman, Morgan Ricks, Andres Sawicki, Alexander Stremitzer, seminar participants at PatCon 7 at Northwestern University Pritzker School of Law, Vanderbilt Law, 2017 Economic Science Association World Meeting, the 4th Annual Roundtable on Empirical Methods in Intellectual Property, the 2017 Conference on Empirical Legal Studies, and the 2017 IP Scholars Conference for helpful comments and discussions, and to Hoffmann Kim and Taylor Weidman for excellent research assistance. All errors are ours.

---

[1] Elizabeth Hoffman
Department of Economics, Iowa State University, 518 Farm House Lane, Ames, IA 50011
E-mail: bhoffman@iastate.edu
Phone: 515-294-0039

David Schwartz • Matthew Spitzer
Northwestern University Pritzker School of Law, 375 E. Chicago Ave., Chicago, IL 60611
E-mail: {David.schwartz, Matthew.spitzer}@law.northwestern.edu
Phone: 312-503-{2646, 8434}

Eric Talley
Columbia Law School, 435 West 116th Street, New York, NY 10027
E-mail: etalley@law.columbia.edu
Phone: 212-854-0437

Patently Risky: Framing, Innovation and Entrepreneurial Preferences

Elizabeth Hoffman • David Schwartz • Matthew Spitzer • Eric Talley[1]

December 29, 2017

Abstract Innovation policy balances static monopoly rights against dynamic entrepreneurial incentives. In striking this balance, researchers commonly presume that decision makers in innovative settings react to their economic environments in a manner similar to their counterparts in other contexts. This paper reports on a series of experiments that call this presumption into question. Subjects were offered a choice between a sure thing and a risky choice, where our principal manipulation was to alter the decisional frame. Subjects in the control group confronted an unadorned choice between safe and risky options; subjects in the Invest Frame, in contrast, were told that the risky choice was tantamount to an investment in an innovation-related project. In all other respects, we controlled for both the language of the experimental instrument as well as the economic stakes entailed. In some subgroups, the risky choice also included a potential monetary loss. We administered the experiments across three subgroups of settings/subjects: a brick-and-mortar lab using university students, an Internet protocol again using university students, and an Internet protocol using Mechanical Turkers as subjects. Our main result, which appears quite strong and robust, is that the Invest Frame induced subjects to manifest greater degrees of risk tolerance on average, across all three settings, and across specifications involving positive and negative payoffs. We calibrate our results to an estimate of a downward "shock" that the invest frame introduces to subjects' coefficient of relative risk aversion, benchmarking our results against Holt and Laury's (2002) risk-aversion elicitation scale.

Keywords Risk aversion; framing; patents; intellectual property; entrepreneurship.
JEL Codes C9; C92; L2; L26; K1; O34.

Acknowledgments Many thanks to Stefan Bechtold, Christopher Buccafusco, Paul Edelman, Morgan Ricks, Andres Sawicki, Alexander Stremitzer, seminar participants at PatCon 7 at Northwestern University Pritzker School of Law, Vanderbilt Law, 2017 Economic Science Association World Meeting, the 4th Annual Roundtable on Empirical Methods in Intellectual Property, the 2017 Conference on Empirical Legal Studies, and the 2017 IP Scholars Conference for helpful comments and discussions, and to Hoffmann Kim and Taylor Weidman for excellent research assistance. All errors are ours.

---

[1] Elizabeth Hoffman
Department of Economics, Iowa State University, 518 Farm House Lane, Ames, IA 50011
E-mail: bhoffman@iastate.edu
Phone: 515-294-0039

David Schwartz • Matthew Spitzer
Northwestern University Pritzker School of Law, 375 E. Chicago Ave., Chicago, IL 60611
E-mail: {David.schwartz, Matthew.spitzer}@law.northwestern.edu
Phone: 312-503-{2646, 8434}

Eric Talley
Columbia Law School, 435 West 116th Street, New York, NY 10027
E-mail: etalley@law.columbia.edu
Phone: 212-854-0437

1 Introduction

It has long been known that patent law and policy must navigate a precarious trade-off between static monopoly on the one hand, and dynamic incentives on the other. By offering the prospect of monopoly rights over inventions (a static "bad"), patent law provides incentives for would-be inventors and entrepreneurs to develop new product markets and improve old ones (a dynamic "good"). These competing concerns stake out polls of a delicate balance that can be challenging, in practice, to reconcile. Intellectual property (IP) institutions are charged with the task of providing sufficient incentives to catalyze socially valuable innovation, but they must also take care not to over-incentivize such efforts by promising IP rights that are too capacious, too long-standing, or too preclusive of successors' efforts (Heller & Eisenberg 1998).

Consequently, a critical input into a defensible IP policy is knowledge of how and when incentives work in entrepreneurial contexts. Among economists, it is natural to presume that prospective inventor-entrepreneurs respond to incentives and risks in a manner that is similar to other economic actors, blending risk preferences, marginal utilities, and subjective probabilities in predictable ways. That presumption—while seemingly sensible—is clearly critical: for if the innovation context change decision-making behavior in a manner diverging from other contexts, an efficiency-minded IP policy would need to take such changes into account when striking the balance between monopoly and incentives.

This article reports some simple experiments demonstrating that individual decision-making behavior *does*, in fact, shift in contexts that involve innovative entrepreneurship. Our experimental inquiry produces one striking result, and several others that are less striking but also of interest. In each experimental setting, we confronted subjects with a choice between a sure thing and a risky choice. Our principal manipulation was to vary the frame of this choice. In the *invest frame*, subjects were told they could either keep their sure thing payoff, or invest it in creating a hypothetical invention with risky (but actuarially attractive) payoffs. In the *simple lottery frame* we gave the subjects the same substantive choices (with identical payoff structures), but one that stated an unadorned choice between a safe and a risky option, bereft of other framing. Each of the two frames fully specified and controlled for the risk attributes of the lotteries. Beyond the simple frames, we used exactly the same language to describe the lotteries. We administered the experiments in a brick-and-mortar lab at Iowa State University, over the Internet using Iowa State students as subjects, and over the Internet using Amazon's Mechanical Turkers (M-Turkers) as subjects.

The striking result, which appears quite strong and robust, is that the Invest Frame induces significantly greater risk tolerance among subjects in all three settings. These results hold up regardless of whether we control for the subject's age, gender, ethnicity, and several metrics of stated risk aversion. The results also persist when subjects face the prospect of a potential negative payoff associated with the risky project. We calibrate our results to an estimate of a downward "shock" that the invest frame introduces to subjects' coefficient of relative risk aversion, benchmarking against Holt and Laury's (2002) results. Because the effects of the invest frame continue to hold even in the presence of negative payoffs, our results contrast with (though do not directly contradict) the predictions of Kahneman and Tversky (1979), who find that preferences in the presence of negative payoffs (relative to a reference point) behave fundamentally differently from those with strictly positive payoffs.

A secondary result is that—consistent with other literature—the M-Turkers are not strictly comparable to the student subjects across several dimensions. Most notably, in addition to their demographic differences, M-Turkers manifest greater risk aversion, regardless of frame, than students on the Internet and in the Lab. A related result pertains to the utility of M-Turk subjects more generally. Although there are many papers exploring whether results on M-Turk are different from those in the lab, there have been none (that we can find) that consider the sort of framing that we utilize. Our results preliminarily confirm that—despite their various observable differences from conventional subjects—M-Turkers can be used successfully to test the types of framing manipulations at issue here.

2 Literature Review

2.1 Intellectual Property

Intellectual property (IP) broadly includes patents, copyrights, trademarks, and trade secrets. For our purposes, patents are the most relevant, followed by copyrights which are of some relevance. Patents basically protect functional inventions. Patents are awarded by the government to inventors of new, useful, and non-obvious inventions. Copyrights, which are justified on the same economic theory, protect original works of authorship such as books and music. Because some of the copyright literature is potentially extendable to patents, we discuss both the patent and copyright (and more broadly intellectual property) literature in this section.

There are several places in the intellectual property literature in which incentives and risk preferences matter. Below, we set forth two broad categories.

2.1.1 *Motivation for Inventing and Creating*:

There is a literature in economics, as well as in sociology and psychology, that attempts to explain why individuals and firms generate new creative and innovative works. The classic economic theory is that incentives such as the limited exclusive rights offered by patents – are necessary to properly encourage the generation of new works. Innovation is quite risky, and the need to provide incentives, such as patents, is predicated on the view that individuals will choose to avoid risky enterprises. Arrow (1962) has suggested that risk-aversion may lead to under-investment in invention. *Id*. According to this theory, the exclusive rights provided by the patent system help individuals overcome the risk aversion-induced market failure and innovate in ways that are socially desirable.

The risk of copying provides another economic justification for patents (Lemley 2005). Outside of financial incentives provided by patents, the literature sets forth other motivators of innovation, including reputational effects, career rewards, and other intrinsic motivations (Lach and Schankerman 2008).

2.1.2 *Risk Preferences of Individuals and Firms with Respect to Creating*:

There is little solid empirical or experimental evidence on the risk preferences of individuals and firms in the innovation ecosystem. The majority of the IP literature *assumes* that creators and inventors are risk-averse, although a minority of scholars assert the opposite, namely, that creators and inventors are risk-seeking.

Joseph Stiglitz, when discussing intellectual property, articulates the classic view that "[p]eople and firms are risk averse, and if they have to bear risk, they have to be compensated for doing so"

(Stiglitz 2008). Under this view, potential creators and others in the innovation system suffer from risk-aversion like regular people. Without the financial rewards of the patent and copyright systems, societally sub-optimal levels of creative works will be produced. Steven Horowitz makes a similar claim about copyright, arguing that copyright holders are "risk averse, valuing clear entitlements more than equivalent murky ones" (Horowitz 2012).

Relying on the American mineral system for public lands, in 1977 Edmund Kitch propounded the prospect theory of patents, which claims that patent rights are useful in channeling and coordinating development activities in a new technology. By awarding exclusivity shortly after invention, prospect theory asserts that the patent system provides the first inventor with an incentive to develop the broad field of invention (Kitch 1977). Ghosh notes that prospect theory assumes a risk-averse inventor who needs strong property rights to be incentivized to develop the field (Ghosh 2004).

On the other hand, some scholars assert that inventors and creators are risk-seeking. Economist F.M. Scherer offered the "lottery theory" of patents, arguing that patents are like lottery tickets, with most patents being essentially worthless while a small minority of them have substantial value (Scherer 2001; Crouch 2008). Building upon Joseph Schumpeter's theory that investors overestimate their chances of success when presented with a potentially great reward, Scherer posited that potential inventors are sufficiently incentivized to create new inventions by the tiny chance of a large payoff from a patent. Gideon Parchomovsky and R. Polk Wagner note that "the lottery theory critically depends on the assumption that inventors, like lottery ticket buyers, are risk-seeking—indeed, so risk-seeking that they are willing to engage in an activity with a negative expected value" (Parchomovsky and Wagner 2005, p. 24). They argue that corporations, rather than firms, pursue most patents and assert that "the decisions of corporate managers appear both rational and even risk-averse" (Parchomovsky and Wagner 2005).

There is little reliable data on this issue of risk tolerances relating to intellectual property, and most of it is inconclusive (Sawicki 2016). Perhaps the best study is by Thomas Astebro (2003). Astebro studied a sample of approximately 1,000 Canadian inventions that had been evaluated before commercialization by a non-for-profit organization, the Canadian Innovation Centre (CIC) (Astebro 2003). Astebro surveyed the inventors many years after the CIC evaluation to learn whether they had commercialized after receiving the CIC evaluation, and if so, what the return on investment was. He reported that independent inventors develop and commercialize inventions that have *negative* expected returns. Astebro concludes that "risk-seeking is one of several plausible reasons why so many inventors proceed to develop their inventions while only a small fraction can reasonably expect to earn positive returns on their efforts. Another plausible explanation is that inventors are unrealistic optimists in that they overestimate their abilities to succeed" (Astebro 2003, p. 236).

2.1.3 Prior Experiments on Intellectual Property

Our invest frame directly references an *invention*. There are some relevant works, but none of them preempts our study. There are several prior experimental papers on intellectual property law, many of them by Christopher Buccafusco and Christopher Sprigman and various coauthors, (E.g. Buccafusco and Sprigman 2010, 2011; Buccafusco, Burns, Fromer and Sprigman 2014; Buccafusco and Heald 2013; Bechtold, Buccafusco, and Sprigman 2015; Buccafusco, Bechtold, and Sprigman 2013; Sprigman, Buccafusco, and Burns 2016; Buccafusco, Heald, and Sprigman

2017). These experiments are aimed at figuring out how people respond creatively to various types of incentives, and how they value and trade the IP once it is created.

The closest experiments to our own are probably Buccafusco and Sprigman (2010, 2011), who ran a series of experiments designed to test for the existence and size of the endowment effect in intellectual property rights. They find, in general, that the endowment effect is huge for the rights to a prize for a winning poem or painting. However, the Buccafusco and Sprigman papers, while very valuable in their own right, do not preempt ours. First, they test for bids and offers for a prize in a copyright context, not the decision to *invest in an invention*. Second, their endowment effect frame is fundamentally different from ours. See the discussion, below, in our section on framing. Third, they do not test for the difference between laboratory experiments and M-Turk. There is at least one prior work using M-Turk for an IP experiment by Buccafusco, Heald, and Bu (2016). However, we have found no prior work testing for the difference between a brick-and-mortar laboratory and M-Turk in any IP experiment.

There are a number of other important experimental works in IP. For example, Buccafusco, Burns, Fromer and Sprigman tests the different incentives provided by copyright and patent on creativity.[2] Several prior works have focused on sequential innovation – the problem of needing to get permission to use prior, protected works in creating new works. The first, Torrance and Tomlinson (2009), was an extremely complicated, multiple stage game. Some subsequent experiments have been less complex, Bechtold, Buccafusco & Sprigman (2016), Brueggermann, et al, (2015) and suggest that IP rights in a first invention hinder sequential innovation, although Bechtold, et al, obtain results partially *inconsistent* with inventor rationality. Others, such as Boudreau and Lakhani (2013), suggest that a *lack* of rights in a first invention, as against sequential invention, discourages the initial invention.

In sum, although there are a number of interesting works at the intersection of IP and experimental methods, there is nothing that we have found that addresses the issues covered in our paper.

2.2 Framing

Our experiments rely on a "frame." However, in the literature, *frame* means several different things. In order to situate our paper in the literature, we must briefly review some of the previous papers that synthesize categories of frame.

2.2.1 *Previous Categorizations of Frames*

There are already some categorization schemas in the political science and psychology literatures. For example, Druckman (2001) contrasts *equivalence* framing– "the use of different, but logically equivalent, words or phrases (e.g., 5% unemployment or 95% employment, 97% fat-free or 3% fat) causes individuals to alter their preferences," with *emphasis* framing effects, which "lead the subject to focus on one aspect of a problem, thereby affecting his opinions and preferences." Banerjee and Chakravarty (2012), on the other hand, contrast *label* framing, invoked "if subjects are confronted with alternative wordings, but objectively equivalent material incentives and unchanged reference points (with regard to how the endowment is initially allocated)" with *value*

[2] These legal rules can be quite idiosyncratic. For a superb experimental test of the fairness of the German "Bestseller Paragraph" provision in copyright, and its effect on the market, see Engel and Kurschilgen 2011.

framing, where “subjects are confronted with alternative wordings and objectively equivalent material incentives but changed reference points.” Levin, Schneider and Gaeth (1998) contrast *risky choice* framing (similar to value framing) with *attribute* framing, where “people are more likely to evaluate a gamble favorably when it is described positively in terms of winning rather than when it is described negatively in terms of losing,” and *goal* framing, in which, not surprisingly, “the goal of an action or behavior is” described differently.

Unfortunately, none of these categorizations relates sufficiently precisely our treatment. Thus, we synthesize the literature into three broad categories.

2.2.2 *Light Computation*:

First, there are frames which require light computation by subjects to understand that the choices they have are equivalent. These include the “reference point” frames for which Kahneman and Tversky (1981) are most famous. This category also includes circumstances where frames induce asymmetric errors in understanding games (Fosgaard, Hansen, and Wengström 2016). There are also experiments that use compound lotteries. For example, Abdellaoui, Kilbanoff, and Placido (2015) measured compound risk and found that subjects valued compound risks differently than simple risks and that the risk attitudes displayed “more risk aversion as the reduced probability of the winning event increases.” There is also a fascinating paper by Brooks, Stremitzer and Tontrup (2016) which studies *effort* participants exerted when they entered into a contract and completed monetarily incentivized economic tests. The authors determined that thresholds and framing affect effort, noting particularly that loss framing with “poorly selected thresholds may reduce effort.” (pg. 1) But none of these versions of light computation correspond to the type of frame we used.

2.2.3 *Emphasis and Priming*:

Second, there are frames that emphasize one aspect, or another, of a choice in a negative or positive light. An excellent example comes from Chong and Druckman (2007) at 104:

> What is particularly vexing in public opinion research is a phenomenon known as “framing effects.” These occur when (often small) changes in the presentation of an issue or an event produce (sometimes large) changes of opinion. For example, when asked whether they would favor or oppose allowing a hate group to hold a political rally, 85% of respondents answered in favor if the question was prefaced with the suggestion, “Given the importance of free speech,” whereas only 45% were in favor when the question was prefaced with the phrase, “Given the risk of violence.”

In this sort of frame, there is no real difficulty or mental computation required in understanding the basic choice of allowing a hate group to hold a rally or not. The frame, instead, prompts the subject to concentrate on either a positive aspect (the value of free speech) or a negative aspect (the risk of violence) inherent in the choice. Again, this does not seem to correspond to the frame in our paper.

2.2.4 *Imagine Yourself in a Context*:

Finally, Imagine Yourself in a Context frames are found in experiments that either tell subjects that they are in a particular setting, or ask the subjects to imagine themselves in a particular setting when making choices. In these frames the subjects are prompted to imagine themselves in a casino, or imagine themselves buying insurance, or imagine themselves making an investment. This is, in essence, the natue of the frame we used in our experiment.

3 Description of Experiment

The most fundamental question we consider in our experiments is whether subjects manifest different risk preferences when a risky choice is framed as a simple lottery versus investing in an innovative technology. For convenience here, we refer to these as our "Simple Lottery" frame and our "Invest" frame. The language of our "Invest" frame is:

> *"Before filling out a brief questionnaire, will be given $8 either to **Keep** or to **Invest** in creating a hypothetical invention . . . . If you choose to **Keep**, your earnings will be $8. If you choose to **Invest** there is a 1/3 chance that the creative and commercialization process will be successful and return $30, and a 2/3 chance that it will be unsuccessful in the market and return $3. A role of a die will determine your earnings, either $30 or $3." [Emphasis in original.]*

We are closest to the "Imagine Yourself in a Context" version of framing, albeit with real economic stakes. In the Invest frame, we inform subjects that they have the opportunity to invest in a "hypothetical invention." The payoffs correspond to whether or not the invention succeeds and is a success in the market. Beyond the (accurate) financial rewards, clearly none of this is true. Instead, by being prompted that this is a hypothetical invention, the subjects are being asked to imagine that it is true, and act accordingly. Our frame is clearly not a light computation frame, similar to the reference point frame used by Kahneman and Tversky (1981). Our gambles are stated in absolutely identical terms. And, just as in the other papers that use this frame, we assume that the subjects are imagining in precisely the way that we ask of them.

In the Simple Lottery frame, we tell subjects the following:

> *"Before filling out a brief questionnaire, you will be asked to make a choice between **Option A** and **Option B**. You will have only a single opportunity to choose. After you have made your choice, if you chose **Option A,** your earnings will be $8. If you chose **Option B**, there is a 1/3 chance that your earnings will be $30, and a 2/3 chance that your earnings will be $3. A role of a die will determine your earnings, either $30 or $3."*

Note that the Simple Lottery frame and the Invest frame describe exactly the same percentages and payouts.

In addition to choosing either ***Option A*** or ***B*** or ***Keep*** or ***Invest***, each subject provided answers to both a series of demographic questions (related to age, gender, education, and the like), as well as the Holt and Laury (2002)[3] risk aversion scale, generating a risk aversion parameter interval for

[3] We could have used the simpler Eckel and Grossman (2008) risk aversion test. However, as Eckel and Grossman (2008) said themselves of Holt and Laury (2002), "This mechanism imposes a finer grid on the subjects' decisions, and thus produces a more refined estimate of the relevant utility function parameters.

each subject. Cox and Harrison (2008) summarizes the literature on estimating risk aversion up to 2008.

We first conducted a series of the above experiments in the lab at Iowa State University, using Iowa State students as subjects. The data were collected on a paper form and subjects were paid one at a time after an individual roll of a die to determine the payoff for those who chose the risky option. Students were randomly assigned to either the Simple Lottery or the Invest frame and the order of presentation of the certain and the risky options were randomly presented as either the first or the second option.

We replicated the first set of experiments on the M-Turk platform, using a Qualtrics format to collect the data and roll an electronic die. M-Turk subjects were paid in experimental dollars that converted to ¼ the lab payoffs. Then, we replicated the experiments using a Qualtrics survey emailed to Iowa State students and conducted entirely online. Subjects chose to be paid by Amazon gift card, PayPal, or a check. The payoffs were expressed in experimental dollars that converted to ½ the lab payoffs.

Finally, we replicated the Mechanical Turk (M-Turk) and Online experiments with the possibility of negative payoffs (-$3 and +$42). For the negative payoffs iterations, **Option A** or **Keep** provided earnings of $8. For **Option B** or **Invest**, we informed subjects that "there is a 1/3 chance that your earnings will be $42, and a 2/3 chance that your earnings will be -$3….These earnings or losses will be added to or subtracted from your $5 participation fee."

The number of subjects in each version is listed below.

	Laboratory	Mechanical Turk	Qualtrics Online
Invest Frame—Can't Lose $	51	101	59
Simple Lottery – Can't Lose $	49	92	60
Invest Frame – Can Lose $	0	98	78
Simple Lottery – Can Lose $	0	100	80

Table 1: Number of subjects in each version

4 Results

Our identification strategy hinges on detecting whether the experimental manipulation – *i.e.,* introducing the innovation/Invest frame – shifts subjects' degree of revealed risk aversion in the way posited above, causing them to embrace a risky choice more readily than they would in the absence of the manipulation. In the parlance of the above notation, we are attempting to control for subjects' baseline risk aversion parameter (α_0) and other demographic variables (X_i), and estimate the local average treatment effect of a downward shock (λ) that the experimental condition introduces (i.e., revealed risk aversion goes down in the presence of the manipulation).

We suppose that the relevant population exhibits CRRA preferences scaled by a (type dependent) CRRA risk aversion parameter $\alpha(X_i)$, so that:

However, this comes at a cost of increased complexity, which may lead to errors." (p. 2). Charness, et. al (2013) adds "The prevalent use of the Holt–Laury measure has allowed researchers to compare risk attitudes across a wide array of contexts and environments. In turn, this has facilitated a less fragmented approach to the study of risk preferences that minimizes methodological differences and aims to characterize a more general phenomenon. (p. 46) Since we wanted to estimate a risk aversion parameter, we made the decision to use Holt and Laury (2002), despite the increased complexity.

$$\alpha(X_i) = \alpha_0 + \beta \cdot X_i + \varepsilon_i,$$

The only difference from the above is that we now introduce a statistical noise term ε_i, which we assume to be have zero mean and to be distributed according to the cumulative distribution function for the population, $\Phi(\varepsilon_i)$. A natural assumption given the structure of our data is that ε_i is normally distributed (implying a Probit specification); but it easily confirmed that a variety of other distributional assumptions for $\Phi(.)$ work as well.

In proceeding, it is important to remain mindful of whether our experimental data on risk preferences is comparable to that found in the prior literature more generally. We could deploy this literature in two ways. Under the first (a "bootstrapping") approach, we would use the baseline preference parameter estimates from pre-existing studies to impose similar structural constraints on the risk preference distributions of our own subjects. Under the second, we would use the results of the literature as a rough benchmark of comparison for our own sample of subjects, but then (after satisfying ourselves as to rough comparability) use our subjects' own behaviors to identify the distribution of preferences. The advantage of the first approach is that it facilitates comparability of our results to the existing literature. The advantage of the second approach is that it allows us to control for an assortment of variables (e.g., demographic differences) that might be predictive of risk aversion but not easily observed in summary statistics reported in the existing literature.

We employ the latter approach. In Appendix A, we first confirm that our experimental data appear comparable to what has been found in prior literature, focusing particularly on Holt and Laury (2002) (hereinafter, designated HL) as a benchmark; and second, assuming our experimental control group data are comparable, we proceed to use those data as a baseline for teasing out the effect of our manipulation.

The tables below contain the ordinary least squares results of both (a) our baseline specification where subjects could never lose money from opting for the risky choice (Table 2); and (b) the combined specification where negative payoffs are possible (Table 3). In addition to our control/treatment assignment (which was random), we also control for a variety of demographic variables, including fixed effects for HL-bins in the post-experiment elicitation. The results of these estimations suggest a significant effect of our manipulation consistent with our hypothesis. Treatment group subjects manifest a significant reduction in revealed risk aversion, consisting with an average estimated downward propensity to take the riskless choice of slightly more than ten percentage points across all specifications. The magnitude of this estimated shift appears relatively consistent across specifications, and in each specification it is strongly statistically significant (one tail test) under any conventional measure. The only right-hand-side control variable that appears stronger than the manipulation is whether the subject was an M-Turk subject. Moreover, the estimated effect appears to be economically significant as well, as it represents a shift that is approximately or greater than the width of (on average) *seven* of the interior HL. We would thus expect the average shift to move a subject "down" one HL bin for all but the subjects who are on the extremities of the scale. (In Tables 2A and 3A in the appendix, we illustrate the robustness of our OLS results in Probit and Logit specifications; in all specifications the treatment effect is economically and statistically significant)

	OLS 1	OLS 2	OLS 3	OLS 4	OLS 5	OLS 6
LANGUAGE	-0.111*	-0.134***	-0.134***	-0.121***	-0.131***	-0.131***
	(-2.32)	(-3.00)	(-3.00)	(-2.73)	(-2.99)	(-3.01)
GAMBLED			0.021	-0.014	-0.038	-0.036
			(0.44)	(-0.28)	(-0.76)	(-0.72)
AGE				0.009***	0.002	0.001
				(3.62)	(0.52)	(0.29)
GENDER				-0.018	-0.048	-0.001
				(-0.41)	(-1.06)	(-0.02)
HAND				-0.017	-0.017	-0.017
				(-0.30)	(-0.30)	(-0.29)
ETHNICITY				0.064	0.05	0.045
				(1.15)	(0.90)	(0.82)
TURK					0.224***	0.300***
					(2.89)	(2.96)
GENDER x TURK						-0.11
						(-1.21)
CONSTANT	0.443***	0.580***	0.563***	0.328+	0.439***	0.436***
	(12.61)	(3.74)	(3.53)	(1.95)	(2.83)	(2.81)
R-sqd	0.013	0.183	0.183	0.212	0.23	0.233
P	0.0210	0.0000	0.0000	0.0000	0.0000	0.0000
N	412	412	412	412	412	412
HL Switch FE	No	Yes	Yes	Yes	Yes	Yes

Table 2: Baseline Experiments - Losing Money Not Possible

T-Statistics in Parentheses

+ = Significant at 5% (one tailed test); 10% (two tailed test)

** = Significant at 2.5% (one tailed test); 5% (two tailed test)*

*** = Significant at 1% (one tailed test); 2% (two tailed test)*

**** = Significant at 0.5% (one tailed test); 1% (two tailed test)*

	OLS 1	OLS 2	OLS 3	OLS 4	OLS 5	OLS 6	OLS 7
LANGUAGE	-0.123***	-0.132***	-0.124***	-0.125***	-0.122***	-0.128***	-0.129***
	(-3.62)	(-2.96)	(-3.65)	(-2.84)	(-3.61)	(-2.93)	(-2.95)
LOSEMONEY	0.088**	0.078	0.082**	0.08	0.081*	0.074	0.075
	(2.54)	(1.57)	(2.33)	(1.62)	(2.29)	(1.48)	(1.50)
LOSExLANGUAGE		0.019		0.003		0.013	0.1
		(0.28)		(0.05)		(0.19)	(0.15)
AGE			0.005***	0.005***	0.003	0.002	0.002
			(2.86)	(2.86)	(1.03)	(1.00)	(0.67)
GENDER			-0.02	-0.02	-0.032	-0.032	0.016
			(-0.57)	(-0.57)	(-0.90)	(-0.90)	(0.33)
HAND			-0.012	-0.012	-0.01	-0.01	-0.009
			(-0.22)	(-0.22)	(-0.19)	(-0.18)	(-0.17)
ETHNICITY			0.022	0.022	0.012	0.012	0.008
			(0.49)	(0.49)	(0.27)	(0.27)	(0.18)
GAMBLED			-0.027	-0.027	-0.034	-0.034	-0.174
			(-0.73)	(-0.73)	(-0.91)	(-0.91)	(-0.93)
TURK					0.077	0.078	0.151*
					(1.37)	(1.38)	(1.99)
GENDERxTURK							-0.104
							(-1.47)
CONSTANT	0.605***	0.610***	0.488***	0.489***	0.522***	0.527***	0.523***
	(5.62)	(5.62)	(4.04)	(4.00)	(4.32)	(4.30)	(4.27)
R-sqd	0.128	0.128	0.138	0.138	0.14	0.14	0.143
P	0	0	0	0	0	0	0
N	768	768	768	768	768	768	768
HL Switch FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3: OLS Regressions Including Subjects Who Could Lose Money

T-Statistics in Parentheses

+ = Significant at 5% (one tailed test); 10% (two tailed test)

** = Significant at 2.5% (one tailed test); 5% (two tailed test)*

*** = Significant at 1% (one tailed test); 2% (two tailed test)*

**** = Significant at 0.5% (one tailed test); 1% (two tailed test)*

5 Discussion and Implications

It is important to note that we cannot definitively determine that only one set of revealed preferences – the ones in the Simple Lottery frame or in the Invest frame – is the "true" set of preferences for purposes of welfare analysis. In fact, *both* sets of preferences may be true, but for different settings. Consider the Invest frame, where the experiments used the word "invention." The subjects became much more willing to take the significantly positive expected gamble. This could be because being part of an exciting enterprise, leading to new, useful knowledge, produces great utility. (Note that our subjects were not tasked with actually inventing anything. Rather, they were asked if they wanted to *invest* in an invention.) There could also be an effect from knowing that inventions are prosocial, leading to spillover knowledge that helps society. Both of these are perfectly valid reasons for preferring the gamble in the Invest frame, but not in the stripped-down Simple Lottery frame. Separately, investing in general may be also seen as socially desirable and may induce more risk taking even apart from the inventive component. We recognize that our experiments have crossed the entrepreneurial investing and inventing aspects, and future research to separate these may be fruitful.[4]

Our results are potentially important not just for the individual subjects, but for society, as well. When a large number of gambles are repeated, each having significant positive expected value, and not overly correlated with each other, taking the gambles will almost certainly produce more wealth. Framing the risky choice as an investment in an invention induced more subjects to choose the positive expected gamble. This is good for the individuals and, in the case of inventions, where many of the benefits are external to the particular invention, good for society. Inducing individuals to invest in inventions may make society wealthier. Thus, there may be a normative payoff to our results. Again, we should caution against relying, at this stage, too strongly on these implications. Still, we find the direction of the implications comforting.

Second, most of the interesting questions about innovation policy were impacted into the payoffs. Note that when we switched from the Simple Lottery frame to the Invest frame, the payoffs from the gamble did not decrease. That is because there is something implicitly in the background that prevents the "copying will drive the value of the invention to zero" scenario from happening. What is the implicit social structure that keeps the gamble a positive expected value? It is most likely patent, or, slightly less likely (because it is so much less efficient) trade secret law. The expected value of the gamble could stay positive because of prizes, but those are sensitive to expectations about the extent to which cronyism, among other things, will twist the award of the prizes. We must await further experiments before we can say, with confidence, anything about these policy instruments.

---

[4] We do not believe that the endowment effect explains our results. The endowment effect, if applicable in our experiment, only applies in the Invest Frame. In the Invest Frame (but not the Simple Lottery frame) we arguably gave subjects an entitlement to $8, and then asked if they wanted to give up the $8 to invest. If subjects in the Invest Frame thought they were entitled to the $8 before deciding whether to invest, then they should have been less willing to give up the $8, which would show up as *more* risk aversion in the Invest Frame when compared to the baseline Simple Lottery frame. However, we find the opposite. Thus, the endowment effect does not weaken our result, and may strengthen it.

Third, the subjects from M-Turk were consistently more risk averse than our other subjects. This was true even after controlling for age, sex, and ethnicity. But our M-Turk subjects changed behavior in the same way that the other subjects changed in response to the invest frame; M-Turk subjects became less risk averse. Thus, it appears that M-Turk can be used to test the effect of frames like the one we used. However, there is an underlying difference in risk aversion on M-Turk that must be accounted for in experiments that are looking for that output. This will be the subject of a short paper on methodology that we hope to produce.

6 Conclusion

Our experiments provided two results, the most robust of which was that giving subjects a choice between a sure thing and a gamble in a Simple Lottery frame or in an Invest frame interacted significantly with revealed risk preferences; subjects were more risk-tolerant when situated in the invest frame. Male and female subjects responded in approximately the same way, and did so regardless of whether they were in the brick-and-mortar lab, on the internet, or using Amazon's M-Turk.

These experiments represent just a first step in a series of experiments on patents and their role in economics and law. Many of the most interesting questions, having to do with the responsiveness of investment to the strength of patent protection, were left embedded within our payoff structure. Future experiments will be directed to testing, directly, the questions that were unaddressed in this first set of experiments.

References

Abdellaoui, M., Klibanoff P., & Placido L. (2015). Experiments on Compound Risk in Relation to Simple Risk and to Ambiguity. *Management Science* 61(6), 1306-1322.

Arrow, K. (1962). Economic Welfare and the Allocation of Resources for Invention, reprinted Pp. 609-626 in THE RATE AND DIRECTION OF INVENTIVE ACTIVITIES: ECONOMIC AND SOCIAL FACTORS, (Richard R. Nelson ed.).

Astebro, T. (2003). The Return to Independent Invention: Evidence of Unrealistic Optimism, Risk Seeking or Skewness Loving? *The Economic Journal,* Vol. 113, No. 484 (Jan., 2003), pp. 226-39.

Banerjee, P. & Chakravarty, S. (2012). The Effect of Minimal Group Framing in a Dictator Game Experiment, DOI: 10.2139/ssrn.2071006.

Bechtold, S., Buccafusco, C., & Sprigman, C. (2015). Innovation Heuristics: Experiments on Sequential Creativity in Intellectual Property. *Indiana Law Journal,* 91, 1251-1307.

Boudreau, K., & Lakhani K. (2013). How Disclosure Policies Impact Search in Open Innovation, Harvard Business School Technology & Operations Management Working Paper.

Brooks, R., Stremitzer, A., & Tontrup, S. Stretch It but Don't Break It: The Hidden Cost of Contract Framing. Working Paper (July 2016)(forthcoming in *Journal of Legal Studies*).

Brueggemann, J., Crossetto, P., Meub, L., & Bizer, K. (2015). Intellectual Property Rights Hinder Sequential Innovation -- Experimental Evidence. Available at http://ssrn.com/abstract=2545950.

Buccafusco, C., Bechtold, S., & Sprigman, C. (2017). Innovate or Borrow? A Model for Regulating Sequential Innovation.

Buccafusco, C. & Heald, P. (2013). Do Bad Things Happen when Works Enter the Public Domain?: Empirical Tests of Copyright Term Extension, *Berkeley Technology Law Journal,* 28(1), 1-43.

Buccafusco, C., Heald P., & Bu, W. (2016). Testing Tarnishment in Intellectual Property. *Washington University Law Review*, 94, 341-421.

Buccafusco, C., Burns, Z., Fromer, J., & Sprigman, C. (2014). Experimental Test of Intellectual Property Laws' Creativity Thresholds. *Texas Law Review,* 92, 1921-1979.

Buccafusco, C., & Sprigman, C. (2010). Valuing Intellectual Property: An Experiment. *Cornell Law Review,* 91, 1-45.

Buccafusco, C., & Sprigman, C. (2011). The Creativity Effect. *University of Chicago Law Review,* 78, 31-51.

Camerer, C. & Talley, E. (2007). Experimental Study of Law. *Handbook of Law and Economics*, 2, 1619-1650.

Charness, G., et. al. (2013). Experimental Methods: Eliciting Risk Preference. *Journal of Economic Behavior and Organization,*87, 43-51.

Chong, D. & Druckman, J. (2007). Framing Theory. *Annual Review Political Science*, 10:103–26.

Cox, J. and Harrison, G. (2008). (Eds.) *Risk Aversion in Experiments. Research in Experimental Economics. Vol. 12.* Emerald Group Publishing, Ltd.

Crouch, D. (2008). The Patent Lottery: Exploiting Behavioral Economics for the Common Good. *Geo. Mason L. Rev*iew, 16, 141-72.

Druckman, J. (2001). The Implications of Framing Effects For Citizen Competence. *Political Behavior,* 23(3), 225-256.

Eckel, C. & Grossman, P. (2008). *Forecasting* Risk Attitudes: An Experimental Study Using Actual and Forecast Gamble Choices. *Journal of Economic Behavior & Organization,* 68, 1–17.

Engel, C. & Kurschilgen, M. (2011). Fairness Ex Ante and Ex Post – An Experimental Test of the German 'Bestseller Paragraph.' *Journal of Empirical Legal Studies*, 8, 682-708.

Fosgaard, T., Hansen, L. & Wengström, E. (2016). Framing and Misperception in Public Good Experiments. *Scand. Journal of Economics.* doi:10.1111/sjoe.12165.

Ghosh, S. (2004). Patents and the Regulatory State: Rethinking the Patent Bargain Metaphor after Eldred. *Berkeley Tech. Law Journal*, 19, 1315-88.

Heller, M. & Eisenberg, R. (1998). Can Patents Deter Innovation? The Anticommons in Biomedical Research. *Science*, 280 (5364): 698–701.

Holt, C. & Laury, S. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 95(5):1644-1655.

Horowitz, S. (2012). Copyright's Asymmetric Uncertainty. *University of Chicago Law Re*view, 79:332-385.

Kahneman, D. & Tversky, A. (1981). The Framing Decisions and the Psychology of Choice. *Science,* 211(4481), 453-458.

Kihlstrom, R.E. & Laffont, J.J. (1979). A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion. *Journal of Political Economy,* 87(4), 719-748.

Kitch, E. (1977). The Nature and Function of the Patent System. *Journal of Law & Economics,* 20:265-68.

Lach, S. & Schankerman, M. (2008). Incentives and Invention in Universities. *RAND Journal of Economics*, Vol. 39. Issue 2:403-33.

Lemley, M. (2005). Property, Intellectual Property, and Free Riding. *Texas Law Review,* 83:1031-1075.

Levin, I., Schneider, S., & Gaeth, G. (1998). All Frames Not Created Equal: A Typology and Critical Analysis of Framing Effects. *Organizational Behavior and Human Decision Process,* 76(2), 149-188.

Parchomovsky, G. & Wagner, R.P. (2005). Patent Portfolios. *University of Pennsylvania Law Re*view, 154:1-76.

Sawicki, A. (2016). Risky IP. *Loyola-Chicago Law Review,* 48: 82-143.

Scherer, F. (2001). The Innovation Lottery, in EXPANDING THE BOUNDARIES OF INTELLECTUAL PROPERTY: INNOVATION POLICY FOR THE KNOWLEDGE SOCIETY 3 (Rochelle Cooper Dreyfuss et al. eds.).

Sprigman, C., Buccafusco, C., & Burns, Z.C. (2016). What's a Name Worth? Experimental Tests of the Value of Attribution in Intellectual Property. *Boston College Law Review,* 93, 1389-1435.

Stiglitz, J. (2008). Economic Foundations of Intellectual Property Rights. *Duke Law Journal,* 57:1693-1724.

Torrance, A. & Tomlinson, B. (2009). Patents and the Regress of Useful Arts. *Columbia Science and Technology Law Review*, 10, 130-68.

Appendix A: Theoretical Framework and Identification

As a theoretical matter, we represent subject choices within a generalized expected utility (GEU) choice-theoretic framework (See, e.g., Camerer and Talley 2008). In our framework, our experimental manipulation (the "Invest" frame) represents a controlled shock to subjects' underlying risk preferences, possibly inducing them to think about risk aversion differently than they would otherwise behave were the equivalent economic choice framed as a strict gamble (e.g., Kihlstrom and Laffont 1979).

The discussion below proceeds in two stages: First, we discuss the underlying choice-theoretic framework, and the predicted effect of the manipulation. Second, we consider an empirical calibration and identification strategy, along with giving results from the first set of "baseline" experiments.

Choice Theoretic Framework

Each subject i is presumed to have individual risk preference characteristics summarized by a (potentially type-dependent) risk aversion parameter $\alpha(X_i) \in \mathbb{R}$, where X_i represents a vector of subject characteristics (e.g., demographics). While $\alpha(X_i)$ could take any functional form, we will frequently concentrate on linear relationships, so that:

$$\alpha(X_i) = \alpha_0 + \beta \cdot X_i,$$

where α_0 is a constant representing a "baseline" level of risk aversion and β is a vector of coefficients on subject characteristics X_i.

In both treatment and control groups, the subject faces a choice between a "sure thing" (ST) and a "risky venture" (RV). Project ST pays off $V > 0$ with certainty, while RV pays off $V_H > V$ with probability q and $V_L \in (0, V)$ with probability $(1-q)$, where $q \in (0,1)$. We assume that $qV_H + (1-q)V_L > V$, so that an unbiased, risk-neutral party would always prefer RV to ST. (As noted above, the experimental vignette set forth V = \$8; V_H = \$30; V_L = \$3; and q = 1/3, which clearly satisfies this condition.)

We suppose for concreteness that subjects are heterogeneously risk-averse, exhibiting constant relative risk aversion (CRRA) utility functions. Equivalently, the utility subject i gets from realized income y_i, or $u(y_i; \alpha_i)$, can be represented as follows:

$$u(y_i; \alpha_i) = \frac{{y_i}^{1-\alpha(X_i)}}{1-\alpha(X_i)}$$

(Recall that this function converges to $\ln(y_i)$ as $\alpha(X_i) \longrightarrow 1$.) The special case of $\alpha(X_i) = 0$ corresponds to risk neutrality, while $\alpha(X_i) > 0$ corresponds to risk aversion, and $\alpha(X_i) < 0$ corresponds to a preference for risk.

Given this set of preferences, subject *i* will (weakly) prefer the risky venture (RV) to the sure thing (ST) if and only if:

$$u(RV; \alpha(X_i)) = q \cdot \frac{V_H^{1-\alpha(X_i)}}{1-\alpha(X_i)} + (1-q) \cdot \frac{V_L^{1-\alpha(X_i)}}{1-\alpha(X_i)} \geq \frac{V^{1-\alpha(X_i)}}{1-\alpha(X_i)} = u(ST; \alpha(X_i))$$

or equivalently:

$$q \cdot V_H^{1-\alpha(X_i)} + (1-q) \cdot V_L^{1-\alpha(X_i)} \geq V^{1-\alpha(X_i)}$$

Given our parameterization, there is a unique risk aversion level, $\alpha(X_i) = \alpha^*$, in which the above expression is satisfied at equality, and the subject is indifferent between ST and RV. She thus prefers ST when $\alpha(X_i) > \alpha^*$, and prefers RV when $\alpha(X_i) < \alpha^*$. For the specific numerical values utilized in our experimental setting,[5] it is easily verified that the unique indifference point occurs at $\alpha^* \approx 0.66$.

We represent our experimental manipulation as potentially introducing a "shock" to the baseline level of risk aversion, or α_0 from above, to a new value $\alpha_1 = \alpha_0 + \lambda < \alpha_0$. Note that because our "Invest" frame is designed to *reduce* manifest aversion to risk, we hypothesize the shock to be negative, so that $\lambda < 0$. The shock will not affect all subjects equally: For infra- and extra-marginal subjects (for whom risk aversion $\alpha(X_i)$ was much less or much greater than the critical switch value α^*), the manipulation will not affect preference orderings. However, for near "marginal" subjects where $\alpha(X_i)$ is in the vicinity of α^*, our manipulation can induce a change in behavior from favoring ST to favoring RV. That is, we would expect to find a group of subjects for which:

$$\alpha_0 + \beta \cdot X_i + \lambda < \alpha^* < \alpha_0 + \beta \cdot X_i$$

In other words, if our manipulation has the effect we posit, we would expect a disproportional preference for RV relative to ST in the treatment group compared to the control group. We therefore seek an identification strategy that will allow us to estimate λ, and to test the null hypothesis that $\lambda = 0$ against the (one-sided) alternative that $\lambda < 0$.

(a) Calibration to the Literature

As noted above, one unavoidable limitation of drawing on results from prior literature is that granular information on the subjects' demographics (or the X_is) is rarely if ever reported in usable form. Thus, the best we can do is to benchmark on summary statistics (effectively dropping all of the X_is other than a dummy variable indicating whether the subject was in our experimental control group).

Moreover, in both our experiment and in the prior literature one cannot observe the subjects' true baseline values of α_0. The best one can do is to observe the first "switch point" on the Holt-Laury (2002) scale (hereinafter "HL scale") at which probabilities grow sufficiently favorable that a

[5] I.e., $V = \$8$; $V_H = \$30$; $V_L = \$3$; and $q = 1/3$.

subject first chooses the high-variance (Project B) over the low variance project (Project A). This switching point, in turn, can be converted into a range of risk aversion values (α), as depicted in the final column of the table below[6]:

Option A (Low Variation)	Option B (High Variation)	Switch Point $\Rightarrow \alpha$
10% chance of \$2.00 and 90% chance of \$1.60	10% chance of \$3.85 and 90% chance of \$0.10	$\alpha \leq -1.713$
20% chance of \$2.00 and 80% chance of \$1.60	20% chance of \$3.85 and 80% chance of \$0.10	$-1.713 < \alpha \leq -0.947$
30% chance of \$2.00 and 70% chance of \$1.60	30% chance of \$3.85 and 70% chance of \$0.10	$-0.947 < \alpha \leq -0.487$
40% chance of \$2.00 and 60% chance of \$1.60	40% chance of \$3.85 and 60% chance of \$0.10	$-0.487 < \alpha \leq -0.143$
50% chance of \$2.00 and 50% chance of \$1.60	50% chance of \$3.85 and 50% chance of \$0.10	$-0.143 < \alpha \leq 0.146$
60% chance of \$2.00 and 40% chance of \$1.60	60% chance of \$3.85 and 40% chance of \$0.10	$0.146 < \alpha \leq 0.411$
70% chance of \$2.00 and 30% chance of \$1.60	70% chance of \$3.85 and 30% chance of \$0.10	$0.411 < \alpha \leq 0.676$
80% chance of \$2.00 and 20% chance of \$1.60	80% chance of \$3.85 and 20% chance of \$0.10	$0.676 < \alpha \leq 0.971$
90% chance of \$2.00 and 10% chance of \$1.60	90% chance of \$3.85 and 10% chance of \$0.10	$0.971 < \alpha \leq 1.368$
100% chance of \$2.00 and 0% chance of \$1.60	100% chance of \$3.85 and 0% chance of \$0.10	$\alpha > 1.368$

Table 1A: Holt-Laury (2002) Risk-Aversion Elicitation Bins

In addition, we must further allow for the possibility that a subject would *never* switch within the Holt-Laury experimental protocol, even when the chance of the high payoff reached 100%. This is no doubt inconsistent with any type of rational choice theoretically, but we found that approximately 2.7 percent of our subjects never switched to option B in our Holt-Laury elicitation. We therefore place these subjects into an 11$^{\text{th}}$ bin, which we call A_{11}, and which cannot be rank-ordered against the others.[7] Through the HL elicitation question, we observe a series of dummy variables $z_{i,k}$, which reflect whether bin A_k contains the first bin at which i switches to Option B, for bins $k \in \{1,2,\ldots,10,11\}$.

To assess our experimental data side-by-side against the HL results, we simulated a data set replicating the summary statistics of Holt & Laury (2002). Because the HL data do not include any granular controls, we control (at this stage) only for a single dummy variable: whether the subject was part of our experimental data, and in particular part of the control group. Given the normality assumption on error terms, an ordered probit is the natural choice.

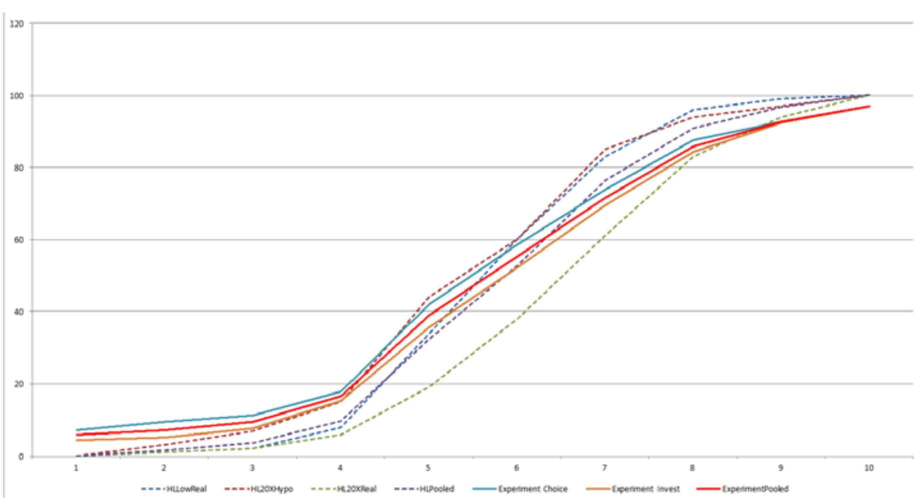

[6] The HL elicitation subdivides the risk aversion domain A into K=10 ordered "bins" coinciding with:

$$\{A_1|A_2 \ldots A_9|A_{10}\} = \{(-\infty,-1.713]|(-1.713,-0.947]|\ldots|(0.971,1.368]|(1.368,\infty)\}$$

[7] Our results change little if the "never switch" subjects are dropped entirely from our data set.

Fig. 1 Subjects' Holt-Laury Switch Bins (Solid Lines; Gains Only & Lose-Money Condition) versus Original Holt-Laury (2002) Switch Distribution (Dotted Lines)

Consider Figure 1, which illustrates the cumulative frequency of switch-point bins, both for the four original Holt-Laury (2002) conditions (dashed lines) and our various experimental baseline subjects (solid lines). As can be seen from the figure, our subjects appear to manifest a somewhat greater degree of risk aversion at the upper end of the HL scale than most of the HL conditions (other than the 20x real stakes condition). That said, our subjects appear to behave consistently in a manner that sits comfortably within the range of responses in Holt & Laury (2002). Moreover, note that our treatment and control subjects manifest nearly identical switch point distributions – a fact that we will utilize in our identification strategy below. Overall, we consider this to be reasonable grounds to believe that our data are highly comparable to Holt & Laury (2002), albeit possibly skewed slightly (but insignificantly) towards greater risk aversion.[8] This comparison provides some comfort that our data are comparable to both prior literature, as well as one another regardless of whether subjects they were assigned to the control or treatment group.

(b) Estimation Results

Having largely satisfied ourselves of the compatibility of our experimental data with prior literature, we now proceed to estimate the effect of the manipulation variable "language" (representing the use of an invest frame) on whether the subject takes a "safe" choice in an experimental setting.

Let $y_i \in \{0,1\}$ denote whether the subject takes the {risky, safe} decision. (Note that we normalize the "safe" decision as $y_i = 1$, so that this fits into the standard framework for limited dependent variables). We use the standard limited dependent variable approach to estimate coefficients underlying the binary choice between projects. Assume that there is some "latent" risk aversion variable $\hat{y}_i$ for each experimental subject, which cannot be observed directly. For subject i the latent variable is defined by:

$$\hat{y}_i = \alpha_0 + \lambda w_i + \beta X_i + \delta z_i + \varepsilon_i$$

The subject's action in is dictated by this latent variable, such that:

$$y_i = \begin{cases} 1 \; if \; \hat{y}_i \geq 0 \\ \quad else \end{cases}$$

In the above setup, α_0 is an estimated constant, representing baseline risk aversion; β is a vector of control-variable coefficients on demographic variables X_i, and δ is a vector of "fixed effect" coefficients for (K-1) of the HL "bins" subjects fall into. Our coefficient of interest in this expression will be λ, which embodies the marginal effect of being placed in the innovation "language" treatment group, (where $w_i = 1$), as opposed to the pure risk frame (where $w_i = 0$). The ε_i denotes an error term on the latent variable. Because we predict that the invest frame will

[8] Beyond eyeballing, we checked whether our subjects appeared comparable to the simulated H-L data based on switching bins in an ordered probit/logit specification. When we compare the pooled HL data to our control group, we found a modest bias in the direction of risk aversion among our experimental controls. However, this bias is not statistically significant under conventional measures (z=1.55 & 1.63, respectively).

make subjects *less* risk averse and *more* risk preferring, we will test a null hypothesis that $\lambda = 0$ against the one-sided alternative that $\lambda < 0$.[9]

Given the framework from above, the risky choice will be taken whenever

$$\varepsilon_i \leq -(\alpha_0 + \beta X_i + \delta z_i + \lambda w_i)$$

which occurs with probability:

$$\Phi\left(\frac{-(\alpha_0+\beta X_i+\delta z_i+\lambda w_i)}{\sigma}\right)$$

And the safe choice will be taken whenever

$$\varepsilon_i > -(\alpha_0 + \beta X_{\square} + \delta z_i + \lambda w_i)$$

which occurs with probability:

$$1 - \Phi\left(\frac{-(\alpha_0+\beta X_i+\delta z_i+\lambda w_i)}{\sigma}\right)$$

Let's suppose that out of our N subjects, we observe *n<N* of them choose the safe choice ($y_i = 1$) and the remaining *N-n* choose the risky choice ($y_i = 0$). The appropriate likelihood function is defined as follows:

$$\Lambda(\alpha_0, \beta, \delta, \lambda) = \prod_{i=1}^{N} \left[\Phi\left(\frac{-(\alpha_0+\beta X_i+\delta z_i+\lambda w_i)}{\sigma}\right)\right]^{1-y_i} \left[1 - \Phi\left(\frac{-(\alpha_0+\beta X_i+\delta z_i+\lambda w_i)}{\sigma}\right)\right]^{y_i}$$

The log likelihood function is:

$$\ln\big(\Lambda(\alpha_0, \beta, \delta, \lambda)\big) = \sum_{i=1}^{N} (1 - y_i) \cdot \ln\left(\Phi\left(\frac{-(\alpha_0+\beta X_i+\delta z_i+\lambda w_i)}{\sigma}\right)\right) + y_i \cdot \ln\left(1 - \Phi\left(\frac{-(\alpha_0+\beta X_i+\delta z_i+\lambda w_i)}{\sigma}\right)\right)$$

[9] One caveat deserves mention here: Because our other control variables (X_i and z_i) are both elicited *after* the experimental manipulation, it is conceivable that the experimental manipulation itself affected post-manipulation responses. This fear is less salient with the demographic variables X_i, such as age, left-handedness, etc. However, the HL risk aversion elicitation, z_i, might well be altered by being assigned to the treatment or control group. Were this to happen, it would likely attenuate any results we find, which is good news for us. That said, this possible treatment effect on a RHS variable is worth keeping in mind in interpreting the regressions below; we will thus consider specifications that both exclude and include fixed effects for HL bins reported by the subjects. (We note, however, that the HL elicitations from our experimental control and treatment subjects appear virtually identical, giving us some confidence that the HL bins are not infected by our experimental manipulation – see Figure 1 above.)

The maximum likelihood approach chooses $\alpha_0, \beta, \delta, \lambda$ -- as well as σ -- to maximize the above function. As before, given our normality assumptions on ε_i, a Probit specification is appropriate.

As noted above, if the invest frame has no effect, then one would predict $\lambda = 0$. If, in contrast, treatment makes subjects *less* risk averse and *more* risk preferring on the margin, then we would predict $\lambda < 0$, we will test the null hypothesis that $\lambda = 0$ against the one-sided alternative that $\lambda < 0$.

(c) Estimation Robustness

The following tables report on alternative probit and logit estimations of Tables 2 and 3 in the text, which used OLS linear probability models.

	Probit 1	Probit 2	Probit 3	Probit 4	Probit 5	Probit 6	Logit 1	Logit 2	Logit 3	Logit 4	Logit 5	Logit 6
LANGUAGE	-0.291*	-0.421***	-0.421***	-0.384***	-0.432***	-0.435***	-0.470*	-0.685***	-0.686***	-0.646***	-0.719***	-0.725***
	(-2.31)	(-3.11)	(-3.10)	(-2.80)	(-3.11)	(-3.14)	(-2.31)	(-3.01)	(-3.01)	(-2.77)	(-3.03)	(-3.06)
GAMBLED			0.064	-0.049	-0.129	-0.126			0.115	-0.097	-0.22	-0.218
			(0.44)	(-0.31)	(-0.80)	(-0.79)			(0.47)	(-0.36)	(-0.79)	(-0.78)
AGE				0.027***	0.004	0.002				0.047***	0.008	0.004
				(3.40)	(0.36)	(0.15)				(3.41)	(0.43)	(0.21)
GENDER				-0.059	-0.16	-0.018				-0.07	-0.252	0.01
				(-0.42)	(-1.08)	(-0.09)				(-0.29)	(-0.99)	(0.03)
HAND				-0.052	-0.056	-0.055				-0.084	-0.098	-0.092
				(-0.28)	(-0.29)	(-0.28)				(-0.26)	(-0.29)	(-0.28)
ETHNICITY				0.181	0.14	0.126				0.316	0.265	0.24
				(1.07)	(0.82)	(0.74)				(1.11)	(0.91)	(0.83)
TURK					0.688***	0.896***					1.129***	1.488***
					(3.04)	(2.96)					(2.96)	(2.93)
GENDER x TURK						-0.307						-0.536
						(-1.05)						(-1.08)
CONSTANT	-0.144	0.254	0.202	-0.471	-0.1	-0.11	-0.23	0.412	0.321	-0.81	-0.198	-0.213
	(-1.62)	(0.64)	(0.49)	(-1.05)	(-0.23)	(-0.25)	(-1.62)	(0.65)	(0.49)	(-1.15)	(-0.29)	(-0.31)
Chi-sqd	5.335	74.016	74.127	76.842	93.296	93.197	5.314	66.627	66.658	68.9	82.057	81.526
p	0.021	0	0	0	0	0	0.021	0	0	0	0	0
N	412	412	412	412	412	412	412	412	412	412	412	412
HL Switch FE	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes

Table 2A: Baseline Experiments - Probit and Logit Specifications when Losing Money Not Possible

T-Statistics in Parentheses

+ = Significant at 5% (one tailed test); 10% (two tailed test)

** = Significant at 2.5% (one tailed test); 5% (two tailed test)*

*** = Significant at 1% (one tailed test); 2% (two tailed test)*

**** = Significant at 0.5% (one tailed test); 1% (two tailed test)*

	Probit 1	Probit 2	Probit 3	Probit 4	Probit 5	Probit 6	Probit 7	Logit 1	Logit 2	Logit 3	Logit 4	Logit 5	Logit 6	Logit 7
LANGUAGE	-0.346***	-0.389***	-0.351***	-0.372***	-0.348***	-0.385***	-0.388***	-0.567***	-0.629***	-0.577***	-0.607***	-0.571***	-0.624***	-0.630***
	(-3.61)	(-3.01)	(-3.64)	(-2.87)	(-3.61)	(-2.98)	(-3.00)	(-3.58)	(-2.95)	(-3.61)	(-2.84)	(-3.57)	(-2.93)	(-2.96)
LOSEMONEY	0.246**	0.201	0.232**	0.21	0.230**	0.193	0.196	0.401**	0.339	0.377**	0.347	0.371*	0.318	0.323
	(2.56)	(1.50)	(2.38)	(1.56)	(2.36)	(1.42)	(1.44)	(2.53)	(1.55)	(2.33)	(1.57)	(2.29)	(1.43)	(1.45)
LOSExLANGUAGE		0.09		0.044		0.075	0.068		0.128		0.063		0.11	0.1
		(0.47)		(0.23)		(0.39)	(0.35)		(0.41)		(0.20)		(0.35)	(0.31)
AGE			0.014***	0.014***	0.007	0.007	0.005			0.023***	0.023***	0.012	0.011	0.008
			(2.85)	(2.83)	(0.99)	(0.94)	(0.62)			(2.86)	(2.85)	(1.03)	(0.99)	(0.65)
GENDER			-0.054	-0.054	-0.09	-0.091	0.047			-0.087	-0.087	-0.145	-0.146	0.092
			(-0.54)	(-0.54)	(-0.88)	(-0.88)	(0.34)			(-0.53)	(-0.53)	(-0.86)	(-0.86)	(0.40)
HAND			-0.036	-0.034	-0.029	-0.026	-0.023			-0.053	-0.05	-0.048	-0.043	-0.041
			(-0.23)	(-0.22)	(-0.19)	(-0.16)	(-0.15)			(-0.20)	(-0.19)	(-0.19)	(-0.17)	(-0.16)
ETHNICITY			0.054	0.054	0.027	0.025	0.011			0.1	0.1	0.057	0.055	0.034
			(0.43)	(0.43)	(0.21)	(0.19)	(0.09)			(0.48)	(0.48)	(0.26)	(0.26)	(0.16)
GAMBLED			-0.079	-0.079	-0.1	-0.101	-0.104			-0.132	-0.132	-0.164	-0.165	-0.174
			(-0.74)	(-0.74)	(-0.92)	(-0.93)	(-0.96)			(-0.74)	(-0.74)	(-0.91)	(-0.92)	(-0.96)
TURK					0.22	0.226	0.430*					0.345	0.354	0.703*
					(1.41)	(1.44)	(2.02)					(1.33)	(1.36)	(1.99)
GENDERxTURK							-0.291							-0.498
							(-1.44)							(-1.48)
CONSTANT	0.272	0.298	-0.05	-0.035	0.052	0.081	0.071	0.445	0.48	-0.086	-0.066	0.075	0.115	0.099
	(0.96)	(1.04)	(-0.16)	(-0.11)	(0.16)	(0.24)	(0.21)	(0.97)	(1.04)	(-0.17)	(-0.13)	(0.14)	(0.22)	(0.18)
R-sqd														
Chi-sqd	93.285	94.437	100.252	100.606	104.513	105.341	105.617	85.351	86.294	91.844	92.123	95.851	96.494	96.37
p	0	0	0	0	0	0	0	0	0	0	0	0	0	0
N	768	768	768	768	768	768	768	768	768	768	768	768	768	768
HL Switch FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3A : Probit and Logit Regressions Including Subjects Who Could Lose Money

T-Statistics in Parentheses

+ = Significant at 5% (one tailed test); 10% (two tailed test)

** = Significant at 2.5% (one tailed test); 5% (two tailed test)*

*** = Significant at 1% (one tailed test); 2% (two tailed test)*

**** = Significant at 0.5% (one tailed test); 1% (two tailed test)*