

Innovate versus Imitate: Theory and Experimental Evidence

John Duffy* Jason Ralston†

June 10, 2020

Abstract

We model and experimentally evaluate the trade-off between innovation and imitation commonly faced by firms. Innovation involves searching for a high payoff opportunity, but paying a cost in order to do so. Imitation involves avoiding that search cost and copying the most successful payoff opportunity uncovered thus far. We formulate a novel model of sequential innovation versus imitation decisions made by a group of n regret minimizing agents. We analyze the consequences of complete versus incomplete information about the distribution of payoffs from innovation on agents' decisions. We then study these predictions in a laboratory experiment where we find evidence in support of our theoretical predictions.

Keywords: Innovation, Imitation, Risk, Ambiguity, Regret Minimization, Experimental Economics

JEL Codes: C91, C92, D21, D81, O31.

*Department of Economics, University of California, Irvine, CA 92697 USA Corresponding author.
*email:*duffy@uci.edu

†Department of Economics, Whitman College, Walla Walla, WA 99362 USA *email:*ralstonj@whitman.edu
Acknowledgments: We thank the editor, two referees, and seminar participants at UC Irvine and Baylor University for valuable comments and suggestions on earlier drafts. We also thank Karl Schlag and Cole Williams for their help with the theory. We gratefully acknowledge funding for this project from the UC Irvine School of Social Sciences. The experimental protocol was approved by UC Irvine's Institutional Review Board, HS #2015-2082.

1 Introduction

In many contexts where consumers or firms face uncertainty, they do so in an environment that is novel to them. The first firm to enter a market does so when the returns to their research and development are unclear. The first time consumers encounter a new item they must decide whether or not to acquire it. On a personal level, we face this scenario when we choose how to order off a menu in a restaurant that we have never been to before.

For those who follow the first mover (the “innovator”), the decision space is doubled: “followers” must decide whether to imitate the innovator or to further innovate themselves and perhaps push the innovation frontier forward. The standard procedure, as in all of economic decision making, is to weigh the benefits and costs of further innovation versus imitation. However, there are many different types of uncertainty and incentive structures that can affect how agents make such decisions. For example, individuals might know the distribution of possible outcomes from further innovation and so behave as expected payoff maximizers in choosing between innovation or imitation. More realistically, they may not know the distribution of possible payoffs from further innovation (the case of ambiguity) in which case they might act as regret minimizers. A further complication is whether innovation precludes imitation, or whether imitation is always possible independently of whether an agent chooses to innovate or not. If one can simply abandon failed innovations, i.e., those leading to payoffs that are worse than the current best innovation, and instead imitate the currently best available innovation, then that incentive structure will change the dynamics of innovation, even when agents are placed in an ambiguous setting.

In this paper, we use theory and experiments to address the innovate versus imitate decision. We have in mind the decisions to innovate or imitate as made by *firms*, in particular.¹ By innovation, we have in mind the choice of a new product or service, production process, marketing strategy or organizational structure resulting in a different payoff to the firm. By imitation, we mean copying these same choices made by another firm, and thereby obtaining the same payoff as that firm. For example, smart phone manufacturers can choose to innovate in the design of a new phone or copy/match feature from other phone manufacturers.

¹Nevertheless, we think the innovation/imitation framework that we introduce is relevant to other actors as well including consumers and governments.

In addressing this decision, we consider cases where the distribution of innovation payoff outcomes is *known* or is *unknown* to firms, who are always informed of the best innovation payoff uncovered in the game thus far. In these two environments, we further consider the case where firms can choose to *imitate* or “recall” the currently best innovation payoff in the event that their innovation payoff is lower (innovation failure) as well as the case where firms cannot recall that best innovation payoff outcome in the event of unsuccessful innovation, that is, the case where innovation *precludes imitation*. We characterize firms as regret minimizing, which, in our setting degenerates to utility maximizing behavior when the distribution of innovation opportunities is made known.

Our theory predicts that firms behave differently depending on the incentive and informational structures that they face. Specifically, firms are predicted to play probabilistically in situations where they face ambiguous innovation payoff distributions, but deterministically when the distribution of payoffs from innovation is known. Further, the expected maximum innovation value (or draw) increases when we allow for recall (imitation following unsuccessful innovation). The four different conditions also give rise to four different values for the expected number of innovation attempts (or draws) and the expected maximum draw. We test all of these predictions of our novel, sequential innovate versus imitate model using an experiment that elicits subjects’ probabilities of innovating.

Our experimental results reveal that behavior is generally well described by our model. Behavior is closest to the theoretical predictions in the case where subjects cannot recall the current best innovation payoff in the event of worse innovation outcomes, and furthest from model predictions when subjects have the ability to recall, especially when recall interacts with ambiguity. Our evidence is based on analysis of subjects’ deviations from model predictions, specifically, their innovation stopping behavior, the expected number of innovation decisions and the expected maximum payoff value from innovation. We speculate that the inability to imitate the current best available innovations (the absence of recall) makes subjects think harder about the choice problem they face, which results in a closer accordance between subjects’ behavior and the predictions of our model. This problem is amplified when subjects are placed in conditions of recall and ambiguous innovation outcomes, a situation where one also feels the effects of ambiguity aversion (Borghans et al., 2009). We also

find that elicited risk preferences play no role in innovation choices nor do any individual subject-level characteristics.

This paper is related to theoretical work in several areas: evolutionary innovation versus imitative behavior, ambiguity aversion, strategic experimentation and search. The trade-off between innovative and imitation behavior has been explored theoretically in the literature on research and development and changes in productivity (König et al. 2016), development of new products (Ofek and Turut 2008), in the composition of industrial structure (Iwai 1984, 2000), and as coexisting behaviors among firms in evolutionary models (Hommes and Zeppini (2014)). In these models, firms typically face a continuous strategy space, instead of a binary strategy space as in this paper. They also possess information about the distribution of payoff-relevant variables, allowing for belief formation and utility or profit maximization.

Our work is influenced by the literature on ambiguity aversion, or aversion to unknown risks, which occurs when some aspect(s) of the distribution of possible states of the world are unknown, so that agents are not able to assign probabilities to states of the world. This notion dates back to the work of David Hume (1739). Knight (1921) made a distinction between risk, where the probabilities of all possible states of nature are known (distributional information is complete), and uncertainty where this is not the case. Ellsberg (1961) added the term ‘ambiguity’ to describe settings in between “complete ignorance” and risk, where decision makers have less than perfect confidence in their estimates of relative likelihoods.² Decision making under ambiguity has been formalized in models of maxmin expected utility (Gilboa and Schmeidler 1989), Choquet expected utility (Gilboa and Schmeidler 1994), and models that allow violations of the reduction of compound lotteries axiom (Halevy 2007). Ambiguity preferences have been studied extensively in the laboratory, with findings that point to high levels of ambiguity aversion, ambiguity prudence, and ambiguity temperance (Abdellaoui et al. 2015, Baillon et al. 2018). We model ambiguity averse individuals as *regret minimizers*, or agents who apply minimax strategies, similar to the work of Bergemann and Schlag (2011) and Renou and Schlag (2010), where agents encounter uncertainty about the distribution of stochastic elements and seek to minimize their regret given the state of the world that is the *least* favorable. Our study focuses on a simplified version of these models,

²Uncertainty is now used as an umbrella term to describe both risk and ambiguity.

where agents do not need to consider a continuum of strategies. This results in a simple minimax mixed strategy prediction when the distribution over payoffs is ambiguous and a unique threshold prediction when that ambiguity is resolved.

We think that the modeling of agents as regret minimizers comes quite naturally in our ambiguous information setting. Regret minimization might also be a quite reasonable assumption in the presence of agency problems. If decisions are made by a manager rather than by the firm owner, regret minimization, rather than expected profit maximization, might be the more relevant objective. Nevertheless, we are not aware of any application of regret minimization to a firm’s choice of whether to innovate or imitate. There *is* evidence that regret minimization explains behavior in a variety of *other* choice settings including consumer choice (Simonson, 1992), recreation choice, (Boeri et al. 2012), transportation choice (Chorus et al. 2008), and health choices (Boeri et al. 2013). Regret minimization has also been used to understand behavior in a variety of laboratory games including the Travelers Dilemma, the Centipede Game, Nash bargaining, and Bertrand competition (Halpern and Pass 2012).

Our model is also related to the literature on strategic experimentation in multi-agent settings. In this literature, agents face double-armed bandit problems and must decide how to divide their decisions between a safe option and a risky option with a true payoff plus noise (Bolton and Harris 1999, Keller et al. 2005). Over time, agents learn the true payoff by witnessing the realized payoffs of other agents who “experiment”, or receive payoffs from the risky strategy. Our game has a similar flavor to it, though agents who play in our unknown distribution setting have limits on their information sets that are not present in the standard strategic experimentation framework. Namely, our agents only see the current maximum payoff value obtained, and the game is one-shot, so learning is quite limited. In a laboratory setting, it is possible that subjects may learn about the distribution of payoffs through their own payoffs resulting from strategic experimentation, though recall complicates this process by censoring the observation of payoffs that fall below the current maximum.

Finally, our model is also reminiscent of certain types of search models in that there is an optimal stopping rule for further innovation. For example, in labor search models (see, e.g. Lippman and McCall (1976)) a worker with a current wage offer in hand has to repeatedly

consider whether to accept that offer or to pay the opportunity cost of waiting to sample again from the distribution of possible wage offers. In that literature, a distinction is usually made between the case where wage offers not immediately accepted are lost, and the case where past wage offers are retained, which are referred to as sampling without or with *recall*. We adopt this same recall /no recall terminology in our innovate-versus-imitate model. The main difference between our model and labor search models is that the search process for the best innovation in our model is a sequential move game played by n different agents (firms), and *not* a model of repeated, individual decision-making. Firms in our model only get a single opportunity to innovate, and they take as given the value of the current best innovation as determined earlier in the game by another firm. Further, we study the case of both known distributions and unknown distributions (ambiguity) for the distribution of rewards to innovation.

This paper takes elements from each of these different strains of the literature and combines them to ask how a firm might approach the risky task of innovating where the payoff to further innovation is either known or unknown within the context of a single shot innovate or imitate game. Our model provides a rational choice explanation as to why firms' innovation efforts may not lead to the best of all possible products and are instead quite homogeneous, for example, why all smart phones seem to be roughly the same; if the distribution of rewards from innovations is known and stationary, then at some point, the costs of further innovation cannot be rationalized and all firms switch to imitating one another. It is less clear what happens in a world where the distribution of rewards from innovation are unknown (ambiguity). Indeed, central to our study is our evaluation of the reasonableness of modeling agents as regret minimizers. Further, we examine whether the ability to copy the best outcome of another individual, an explicit choice of imitation, increases the likelihood that an agent obtains the maximum payoff possible in a game. Indeed, there is much anecdotal evidence that costless imitation, rather than retarding innovation, actually fosters further innovation (Raustiala and Sprigman (2012)).³

The rest of the paper is organized as follows. Section 2 presents the game and our

³See also Engel and Kleine (2015) for a related experiment examining a budget-constrained choice between costly innovation and *costly* imitation.

theoretical predictions; section 3 describes our experimental design; section 4 presents our experimental hypotheses; section 5 reports on the results of the experiment and tests of our hypotheses; finally, section 6 concludes with a summary of the main findings and some suggestions for future research.

2 Theory

2.1 Known Distribution

Consider the following sequential move, n -player game. Agent i must choose between a costly lottery or simply accepting the highest lottery payoff realization achieved by agents who have moved earlier in the game. We interpret the former choice as “innovation” and the latter as “imitation”. The lottery is a random draw from a continuous probability distribution $F(\theta)$. Specifically, let X denote a random variable with distribution $F(\theta)$, and let x_i denote agent i ’s realization of that random variable (lottery draw). We denote the full prior history of lottery realizations up to round t by $H_t = \{x_1, \dots, x_{t-1}\}$ and the maximum of this set by $x_{\max} = \max\{H_t\}$. We assume that an agent’s index, $i \in I$, the set of all n agents, also denotes the time period, or order, in which agents make their decision. We further assume that the cost to choosing the lottery (choosing to innovate) in any time period is fixed and equal to c . Each agent begins the game with an endowment, e , which may be used to purchase lotteries.

The payoff function is given by:

$$\pi_i = \begin{cases} x_i - c + e & \text{if innovate} \\ x_{\max} + e & \text{if imitate} \end{cases} \quad (1)$$

which states that agent i faces the decision of either taking the lottery draw, $x_i \sim F(\theta)$ at cost c or taking the highest of the previous innovations (imitation).

This “no recall” case can be justified as follows. In addition to the cost of innovating, c , choosing to innovate may be so costly in terms of *time and effort* that switching, ex-post, following a failed innovation to a strategy of imitation is simply not possible. Alternatively,

one can imagine that there are legal restrictions, e.g., patents and copywrites, that make it costly to imitate the innovations of others, for example, a smart phone manufacturer has to pay licensing fees to use a technology patented by a rival. With such costs, our no recall condition can be viewed as capturing the case where payment of the cost of innovation, c , precludes payment of the licensing fees or legal costs needed to imitate the innovations of others, so that imitation following a failed innovation is no longer possible.⁴

We characterize our agents as expected regret minimizers. Let $r(x_{max}, x, j)$ be the regret resulting from action j when the payoff from imitating is x_{max} and the innovation payoff draw is x . Regret is then a piecewise-defined function:

$$r(x_{max}, x, j) = \begin{cases} \max\{x - c, x_{max}\} - x_{max} & \text{if imitate} \\ \max\{x - c, x_{max}\} - (x - c) & \text{if innovate} \end{cases} \quad (2)$$

Taking expectations reduces this problem to one of expected utility maximization where agents should choose to innovate whenever $E(X) - x_{max} \geq c$, and choose to imitate otherwise. The model makes the following prediction:

Proposition 1. *In the case where agents are not allowed to recall, for some $k \in I$, agent k will be the first to imitate and all agents, $i > k$, will also choose to imitate.*

Proof. Under the assumption that agents play rationally, they will continue to choose to pay for lottery draws, that is they will choose to innovate, so long as the expected regret from innovating is lower than the expected regret from imitating. Assuming that $F(\theta)$ is a well-behaved continuous probability distribution, the expected regret from innovating will be increasing in x_{max} . Conversely, the expected regret from imitating is decreasing in x_{max} and reaches zero as x_{max} approaches the upper limit of the support.

By the intermediate value theorem, there exists an x_{max} such that the expected regret

⁴To provide one motivating example of our no recall case, in 1982 the gaming company Atari, Inc. contacted with Universal Pictures to produce a video game based on the movie ET. The licensing fees alone cost Atari \$20-25 million, a very high figure for the time. This was one of the first video games based on a movie, but the development was rushed due to the necessity of shipping the game for Christmas. In total, the game was programmed from scratch in 5 and a half weeks and no time was given to audience testing. 3.5 million games were shipped to stores, but only 1 million were ever sold, and Atari experienced huge losses. The licensing fees, unsold game product and decline in the popularity of the Atari 2600 game console led to the collapse and sale of Atari. Somewhat ironically, Steven Spielberg, the director of ET, had urged the developer to imitate Namcos Pacman video game, but Atari felt it was too derivative (Brumfiel 2017).

from innovating and imitating are equivalent, for a c small enough. Thus, there is some $j \in I$ for which the expected regret from the two strategies are exactly equal, or expected regret from imitating will be less than expected regret from innovating. At this point, the j^{th} agent will be either indifferent between imitating and innovating or strictly prefer to imitate. If the j^{th} agent innovates, then the first agent who will imitate will be $k = j + 1$ and if the j^{th} agent imitates, the first agent who will imitate will be $k = j$. For all agents $i > k$, the expected regret from innovating is larger than the expected regret from imitating, and such agents will all choose to imitate. \square

In the “no recall” condition, the point at which agents should switch from innovate to imitate is $x_{max} = E[X] - c$.

There is also the possibility that investing in innovation is not so costly as to preclude imitation in the event of a failed innovation. In such a case, if innovation is unsuccessful, agent i can always “recall” (receive) the imitation payoff x_{max} , but in that case agent i must still pay the cost of innovating, c . We call this condition “recall”, following the convention from labor search models (Mortensen, 1986; Stokey et al. 1989).

$$\pi_i = \begin{cases} \max\{x_i, x_{max}\} - c + e & \text{if innovate} \\ x_{max} + e & \text{if imitate} \end{cases} \quad (3)$$

Under recall, our regret function becomes

$$r(x_{max}, x, j) = \begin{cases} \max\{x - c, x_{max}\} - x_{max} & \text{if imitate} \\ \max\{\min\{x - c, c\}, 0\} & \text{if innovate} \end{cases} \quad (4)$$

A regret minimizer’s objective is to pick the strategy that produces the minimum expected regret, which is

$$\begin{aligned} & E[r(x_{max}, X, j) | X > x_{max}] \\ &= \begin{cases} E[(X - c) - x_{max} | X - c > x_{max}] \times P(X - c > x_{max}) & \text{if imitate} \\ E[x_{max} - (X - c) | x_{max} < X \leq x_{max} + c] \times P(x_{max} < X \leq x_{max} + c) \\ + c \times P(X \leq x_{max}) & \text{if innovate.} \end{cases} \end{aligned}$$

Our agent will choose to innovate when the expected regret from imitating is greater than the expected regret from innovating.

Proposition 2. *In the case where agents are allowed to recall, there will be some agent/period $k \in I$ such that all agents $i < k$ will innovate and all agents $j \geq k$ will imitate.*

The proof for Proposition 2 follows the same form as that of Proposition 1 and is excluded for brevity.

We are assuming that imitation is always a costless possibility in the recall setting, and is costless at the time of deciding between innovation and imitation in the no recall setting. As noted earlier, in many environments there are legal restrictions, e.g., patents and copyrights, on imitation of others' innovations that might prevent imitation or that make imitation costly, e.g. for copying features of smart phones. On the other hand, there are also industries where there is effectively *no cost* to imitation. For example, for historical and legal reasons, imitation is common and effectively costless in the fashion, financial services, and restaurant food industries (Raustiala and Sprigman (2012)). Still, we can easily allow for *costly* imitation by requiring agents to pay a cost, d , if they choose to imitate. In the recall setting, this imitation cost would remain the same if the firm chose to imitate rather than to innovate, or later chose to imitate following a failed innovation (recall). However, in the no recall treatment, the imitation cost would change from d at the innovate/imitate choice stage to being effectively infinite in the case of a failed innovation, capturing the inability to recall.

While we set d to 0 in our experiment, changing d to be positive only changes a constant in our theory. Such a change would decrease the number of agents who play imitate and increase k in expectation, *ceteris paribus*. Without loss of generality, we will focus on a single cost, c , for innovation rather than a cost d , for imitation, though one can also think of our innovation cost as the *net* cost of the two actions, i.e., $(c - d)$.

2.2 Unknown Distribution

More realistically, agents or firms are unlikely to know the distribution of possible payoffs from innovation, and we also consider this case. In such a setting, we suppose that agents do

know *something* about the innovation prospects they face. Specifically while agents do not know the distribution of possible payoffs, we assume that the support of possible innovation outcomes, $[a, b]$, is perfectly and commonly known, a setting that corresponds to the case of ambiguity, as discussed earlier.⁵

As before, let $r(x_{max}, x, j)$ be the regret resulting from action j with the payoffs from imitating being x_{max} and from innovating being the draw x . Regret is then modeled as it was in equation (2), where recall is not allowed.

Proposition 3. *In the case where agents are not allowed to recall, agents will play the mixed strategy $p^* = \frac{x_{max}-a+c}{b-a}$, where p^* is the equilibrium probability that an agent imitates.*

Proof. We start by finding the distribution, F , that maximizes regret in our framework. To this end, we examine the expected regret function $r(p, F, x_{max})$, where p is the probability that an agent chooses to imitate.

$$r(p, F, x_{max}) = \int_a^b [p * r(x_{max}, x, Im) + (1 - p)r(x_{max}, x, In)] dF(x) \quad (5)$$

It is assumed that $b - c > x_{max} \geq a$ so that imitation does not dominate. We examine the two *degenerate* distributions where all mass lies at the boundaries of the support, which maximizes regret, i.e., $F = \delta_b$ and $F = \delta_a$, where payoffs from innovation are at their most extreme.

In the case where $F = \delta_b$, the expected regret is

$$p [\max\{b - c, x_{max}\} - x_{max}] + (1 - p) [\max\{b - c, x_{max}\} - (b - c)]$$

which simplifies to

$$p [b - c - x_{max}] + (1 - p)0 = p [b - c - x_{max}].$$

In the second case, where $F = \delta_a$, expected regret is found to be

$$(1 - p) [x_{max} - a + c].$$

⁵That is, we do not consider the case of complete ignorance!

It follows that the regret from $F = \delta_b$ will be higher than the regret from $F = \delta_a$ when

$$p[b - c - x_{max}] \geq (1 - p)[x_{max} - a + c],$$

which simplifies to

$$\frac{x_{max} - a - c}{b - a} \leq p.$$

Similarly for the case where $F = \delta_a$,

$$\frac{x_{max} - a - c}{b - a} \geq p.$$

Let $p^* = \frac{x_{max} - a - c}{b - a}$. Then we return to our regret minimization problem, where we wish to minimize

$$MR(p, F, x_{max}) \equiv \max_{F \in \Omega} r(p, F, x_{max}) = p(b - c - x_{max})\mathbb{1}(p \geq p^*) + (1 - p)(x_{max} + c)\mathbb{1}(p < p^*). \quad (6)$$

We minimize this function by finding the mixing probabilities for our agents.

$$\frac{\partial(MR)}{\partial p} = \begin{cases} a - c - x_{max} & \text{if } p < p^* \\ b - c - x_{max} & \text{if } p > p^* \end{cases}$$

This function reaches a minimum at p^* . □

Using a similar method we can find a solution to the regret minimization problem faced by agents who have the ability to recall x_{max} in the event of a worse payoff from innovation. Here, the only thing that changes is when we consider the case where $F = \delta_a$. In that case, when there is recall, agents know they cannot do worse than the current maximum draw, thus the expected regret under $F = \delta_a$ will be $(1 - p)c$ instead of $(1 - p)[x_{max} - a + c]$.

Proposition 4. *In the case where agents are allowed to recall, they will play the mixed strategy $p^* = \frac{c}{b - x_{max}}$, where p^* is the equilibrium probability that an agent imitates.*

The proof follows the same form as the proof provided for proposition 3, noting that we revert to our formulation of regret under recall found in (4).

Propositions 1-4 provide sharp testable predictions as to the strategies that agents should play in our innovate versus imitate game. In the next section, we describe our experimental design for evaluating these theoretical predictions.

3 Experimental Design

The model makes distinct predictions about stopping rules, innovation probabilities, and how they differ depending on whether the distribution is known as well as on the ability to recall prior payoffs in the event of an unsuccessful innovation. Thus our experiment employs a 2×2 experimental design where the two treatment variables are: 1) knowledge/lack of knowledge about the distribution of possible payoffs from innovation and 2) the presence or absence of the ability to recall the maximum prior payoff from innovation in the event that an innovation choice leads to a lower payoff. Table 1 provides a summary.

	No Recall	Recall
Known Distribution	KDNR	KDR
Unknown Distribution	UDNR	UDR

Table 1: The four treatments of the experiment

In each of these four treatments, subjects participated in 4 different stages: the main decision consisting of 1) ten, 10-round, 10-player innovate/imitate games, 2) a risk elicitation task, 3) a cognitive reflection task, and 4) a short demographic survey.

3.1 The main task

In the first and main stage, subjects participate in ten, 10-round “games” with a fixed group of $n = 10$ subjects. At the start of each game, subjects were assigned a random number from 1 to 10 representing their position in the order of moves for the current game. A subject was never informed of their position number. Thus, subjects did not know precisely how many other subjects had come before them or how many others would follow them.⁶ For each round of each game, one of the 10 subjects took a turn deciding whether or not to take

⁶The only inference a subject could have made with certainty was that they were not the first mover if their screen showed $x_{max} \neq 0$.

a draw from a payoff distribution (known or unknown) or copy the highest payoff drawn in the current game thus far by subjects who had drawn before them in that game, i.e., x_{\max} . For the first player to draw in round 1 of each game, x_{\max} was set to 0.

The distribution we used in both treatments (known and unknown) was a finely discretized approximation of the symmetric triangular distribution, with support $[0, 100]$ and a modal peak of 50. The theory is agnostic when it comes to the choice of distribution, and the triangular distribution was chosen for its continuous properties, which conforms to the assumptions of the theory. Additionally, we believe this particular continuous distribution calls attention to the unconditional mean, which is important for calculating the risk-neutral stopping rule under known distribution conditions. Also, when judging the probability of receiving a draw above x_{\max} , subjects unfamiliar with integration need not be familiar with concepts from calculus – in principle, they only need to know how to calculate the area of a triangle. Draws from that distribution were truncated at the hundredths place. Though unlikely, it was possible for subjects to draw the same number more than once. In practice this never occurred.

In the known distribution (KD) treatments, subjects were told about the distribution of innovation payoffs, while in the unknown treatment subjects were only informed of the support of the unknown distribution, $[0, 100]$. Specifically, in the KD treatments, subjects were shown a graph of the distribution they were drawing from featuring the finite range of the support as well as the triangular distribution from which innovation draws were made, including the modal peak of that distribution and its value. Figure 1 displays a screen shot of the main decision screen for the known distribution treatment. The dashed line indicates the current value of x_{\max} in the current game, and this value was updated with changes to x_{\max} as they occurred. The value of x_{\max} was reset to 0 at the start of each new 10 round game.

The decision screen for the unknown (UD) treatments is not shown, but it is similar in all respects except that the triangle distribution is *not* shown; rather, subjects just see the support interval for possible innovation draws, $[0, 100]$ and a dashed line again indicates the current value of x_{\max} . Subjects in the unknown treatment were specifically instructed that the distribution of payoffs from choosing to draw (innovate) was unknown to them and could

be any distribution.⁷

In all treatments, subjects knew their payoff function and the value of x_{max} at the time they were asked to make a choice of whether to innovate (draw) or imitate (not draw). In making choices, subjects could choose between two buttons, Do Not Draw (Imitate) or Draw (Innovate), which indicated either a 0% or 100% chance to take a draw. In addition, subjects also had access to a randomization device to make probabilistic decisions consistent with the mixed strategy prediction of our UD treatments –see again Figure 1. Subjects could enter a probability with which they would take a draw (innovate) from the payoff distribution, which also determined the probability that they did not take a draw, and instead copied the highest payoff received by previous subjects. If subjects entered a probability, then the decision to draw (innovate) was made for them with that probability by the computer program and the decision to not draw (imitate) was made for them by the computer program with 1 minus their entered probability.

In each game, subjects were endowed with $e = 10$ points and were informed that taking a draw in that game would cost them this 10 point endowment, i.e., in terms of the theory, we set $c = e = 10$. If they did not take a draw (did not innovate), then they would keep their 10 point endowment, and get the payoff from imitation, x_{max} . In sessions where recall was available, when a subject took a draw and that draw was below the current maximum, x_{max} , their draw was automatically replaced by the current maximum, x_{max} for payoff calculation purposes. In neither case did a draw below the current maximum change the current value of x_{max} . That is, the current value of x_{max} was determined at the 10 player game level and would always be defined as the current maximum draw taken in a game up until that round; this means that the value of x_{max} was non-decreasing over the 10 rounds of a game. Prior to the start of each new game, x_{max} was reset to 0.

Subjects made one innovation/imitation choice per game, according to their position number for that game. For each of the 10 games, position numbers were randomly and anonymously assigned; subjects were never informed of their position number in a game and

⁷We realize that subjects might update their pessimistic priors using observations of the current maxima across multiple games. However, subjects were never told that the distribution did not change from game to game (which was in fact the case), no history of draws (other than x_{max} was shown to them, and density estimation with small sample sizes is unlikely to update priors in a significant manner.

Choice

Pick from the options below. If you would like to draw with some probability, please enter a number between 1 and 99 to represent the percent chance that you will take a draw. When you draw from the distribution, you lose your 10 point endowment.

The largest draw so far is 46.95. If you do not take a draw from the payoff distribution, you will receive 46.95 points plus your 10 point endowment.




Figure 1: The main decision screen for the known distribution treatment.

position could not be inferred from any information displayed on the decision screen. When the program advanced to the round number that matched their position order, the subject was shown the current value of x_{max} and asked to make their decision.

Every session consisted of 10 subjects playing 10 games. Since each game lasted 10 rounds, each subject made only one decision in one round per game. At the end of the experiment, one game was picked at random for payment.

3.2 Risk Elicitation Stage

After the 10 games were played, each subject advanced to the risk elicitation stage. In the risk elicitation stage, subjects were presented with 6 gambles as in Dave et al. (2010). The subject was told to pick the gamble they most preferred and the computer would randomly determine a payoff, conditional on their choice. The risk elicitation stage can be seen in Figure 2.

Please consider the options below and indicate your choice.

The payment from this part of the experiment will be added to your payment for the previous section of the experiment and the \$7 show-up payment for your total take-home amount.

NOTE: The outcome of your choice will be decided by a random number generator when applicable.

Select the option you most prefer:

- 50% chance of \$2.80, 50% chance of \$2.80
- 50% chance of \$2.40, 50% chance of \$3.60
- 50% chance of \$2.00, 50% chance of \$4.40
- 50% chance of \$1.60, 50% chance of \$5.20
- 50% chance of \$1.20, 50% chance of \$6.00
- 50% chance of \$0.20, 50% chance of \$7.00

Next

Figure 2: The risk elicitation screen.

Assuming subjects exhibit CRRA risk preferences, we can find ranges of the coefficient of relative risk aversion, r , by comparing adjacent lotteries in the table of possibilities. This further allows us to classify subjects as risk-averse, risk-neutral, or risk-loving in our analysis. We report the modal range of coefficient of relative risk aversion in Table 2, below.

3.3 Cognitive Reflection Task and Survey Stages

After completion of the risk elicitation stage, subjects proceeded to a survey stage. In this stage, subjects first answered three cognitive reflection test (CRT) questions and then proceeded to answer demographic questions. The latter questions covered nationality, ethnicity, age, major, and GPA. An illustration of the cognitive reflection task question screen is shown in Figure 3.

Survey

Please answer the following questions.

A bat and a ball cost 22 dollars in total. The bat costs 20 dollars more than the ball. How many dollars does the ball cost?

"If it takes 5 machines 5 minutes to make 5 widgets, how many minutes would it take 100 machines to make 100 widgets?" :

In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how many days would it take for the patch to cover half of the lake? :

[Next](#)

Figure 3: A cognitive reflection task screen.

Following the survey stage, subjects were informed of their experimental earnings, risk elicitation earnings, the show-up payment, and their grand total earnings. Subjects were then paid discreetly.

3.4 Subjects and Data Collection

The experiments were conducted at the University of California, Irvine at the Experimental Social Sciences Laboratory (ESSL). Subjects were undergraduate students at UC Irvine with no prior experience with the game. These subjects were recruited using the SONA systems software.

We collected data from 5 groups of 10 players for each of the four treatments (cells) of our experimental design. Thus we have data on the behavior of $5 \times 10 \times 4 = 200$ subjects. For the first stage, main decision task, we chose one game randomly from all 10 games played

and converted subjects’ point earnings from the chosen game into money earnings at a fixed and known rate 1 point = \$0.15 USD. For the second stage, subjects earned money in an incentivized risk elicitation following the design of Dave et al. (2010). Subjects could earn between \$0.20 and \$7.00 in this stage. The CRT and demographic survey questions in the final stage were unincentivized.

The total average payment was \sim \$22, including a \$7 show up payment. On average, subjects spent about an hour in the laboratory, and of that time about 20 minutes were spent reviewing instructions verbally and taking a comprehension quiz. The remaining 40 minutes were devoted to the experiment, which used a web browser and was programmed in Python using the oTree package (Chen et al. 2016).

Some statistics on our subject population, as taken from our demographic survey, are provided in Table 2.

Age	19.94 (1.94)
GPA	3.01-3.50
CRT score	1.13 (1.15)
CRRA coef.	$0.50 < r < 0.71$
% female	70%

Table 2: Descriptive statistics regarding the subject population; primarily means with standard errors in parentheses.

4 Hypotheses

Based on our theory and experimental design, we have several related hypotheses about behavior in our four different treatments – known distribution without recall (KDNR), known distribution with recall (KDR), unknown distribution without recall (UDNR), and unknown distribution with recall (UDR).

Our main hypothesis is that subjects will behave in accordance with our theoretical predictions. That is, a subject’s propensity to draw (innovate) will match the deterministic or probabilistic predictions of the theory. Further, when the distribution is known, we hypothesize that subjects will draw up to the maximum predicted threshold and imitate thereafter

in the manner characterized by either proposition 2 when there is recall or proposition 1 when there is no recall.

Hypothesis 1. *The probabilities of drawing from the distribution will match those predicted by the theory. When the distribution is known, subjects will choose to draw if x_{max} is below a certain threshold and will imitate otherwise. When the distribution is unknown, subjects will employ probabilistic regret minimizing strategies that are functions of x_{max} .*

For the parameterization of the model we implement, the threshold or probability values are shown in the first two rows of Table 3. Note in particular that the threshold for KDNR < KDR. The probability in UDNR is linear in x_{max} and nonlinear in UDR, and, in expectation, for a given x_{max} the probability of innovation is higher in UDR than in UDNR.

		No Recall (NR)	Recall (R)
Known Dist. (KD)	<i>Probability</i>	<i>Deterministic</i>	<i>Deterministic</i>
	<i>Threshold</i>	40.00 [†]	46.89 [†]
	<i>Exp. Max</i>	60.90	64.71
	<i>Exp. # Draws</i>	1.47	1.78
Unknown Dist. (UD)	<i>Probability</i>	$\frac{(b-c)-x_{max}}{b-a}$	$\frac{(b-c)-x_{max}}{b-x_{max}}$
	<i>Threshold</i>	<i>None</i>	<i>None</i>
	<i>Exp. Max</i>	69.33	77.00
	<i>Exp. # Draws</i>	4.13	6.16

Table 3: Thresholds and expected maxima for each treatment. †See the online Appendix for details.

In addition to considering individual subject behavior we also simulated theoretical play by 10 agents, playing according to theoretical predictions in all four treatments of our experiment 100,000 times. We use these simulated distributions to make additional aggregate, game-level distributional hypotheses.

Hypothesis 2. *The mean number of draws and the expected maximum draw will correspond to the simulation results reported on in Table 3.*

From our simulation results, the mean number of draws and the expected maximum draw across our four treatments should follow the order KDNR < KDR < UDNR < UDR.

We note that in the case of known distributions, our theory presumes risk neutral risk preferences. However our subjects may not be risk-neutral with respect to uncertain money

earnings. Departures from the assumption of risk neutral preferences may affect subject's propensity to innovate. Specifically, subjects who are risk-averse might stop innovating *before* our predicted thresholds or exhibit lower probabilities of drawing. Conversely, risk-loving subjects might continue drawing (innovating) *after* our predicted thresholds or exhibit high probabilities of drawing. Other attributes that might explain departures from risk neutral predictions include subjects' cognitive abilities, which we measure using GPA and CRT scores.

Hypothesis 3. *Deviations from risk-neutral play are correlated with individual risk preferences or other personal attributes.*

5 Experimental Results

The theory makes predictions about a few main outcome measures. For each measure, we compare behavior the theory's predictions with subjects' deterministic or probabilistic propensity to innovate (or draw), given the current, realized value for x_{\max} at the time the subject made their innovation/imitation choice. First, we look at individual-level behavior versus predicted behavior across the four different treatments. These analyses make significant use of *deviations* from theoretical predictions. Second, we compare behavior within a game (10 subjects playing for 10 rounds) with numerical results we generated from large simulations of agents playing exactly according to the strategies predicted by our theory. We use the comparison of the experimental data with the simulation analysis to address the expected maximum draw and number of draws and indifference thresholds (in the case where distributional information is known). When possible we compare these measures across the different treatments as well as with the predicted values for each treatment. Our main results are summarized in Findings 1-3 which map directly to evaluation of Hypotheses 1-3.

5.1 Deviation analysis

We first consider the difference between subjects' probability of drawing (innovating) and the predicted probabilities, using the root mean squared error. From our experimental data we have subject's elicited probabilities of taking a draw from the known or unknown in-

novation distribution and our theory makes predictions about the probabilities that a risk neutral agent would choose when confronted with different values for x_{max} , the current maximum value drawn in a game. We compute the squared deviations of actual from predicted probabilities using the metric

$$dev_{i,t} = (p_{i,t} - \hat{p}_{i,t})^2$$

where $p_{i,t}$ is the subject's reported probability of drawing (innovating) and $\hat{p}_{i,t}$ is the predicted probability of a risk-neutral agent drawing. Note that both probabilities are conditional on the value of x_{max} that the subject faced at time t . We take the square root of these squared deviations to create round-averaged probabilities, which are reported in Figure 4, separated according to the four different treatments. Subjects follow the theoretical predictions most closely in the two no recall treatments, KDNR and UDNR. The root mean squared error (RMSE) of deviations from predictions in the no recall treatments are significantly smaller than those of the recall treatments ($p < 0.01$, two-tailed Mann-Whitney U-test) and the RMSE of deviations in the KDR treatment are smaller than those in the UDR treatment ($p < 0.01$, two-tailed Mann-Whitney U-test). Figure 4 also supports the notion that subjects' decisions are consistent with the predictions of our regret minimization model in the UDNR treatment, but less consistent in the UDR treatment.




Figure 4: Average deviations from the regret minimizing model

We further decompose our sample into early games versus late games to look for evidence of subject learning. We define the early rounds as the first five rounds and the late rounds as the last five rounds. Figure 5 demonstrates that, though the differences are minor in most cases, subjects tend to follow the model predictions a little better with experience. This same figure also supports the notion that subjects in the KDR and UDR treatments fail to take advantage of the recall opportunity, especially in the first five games.

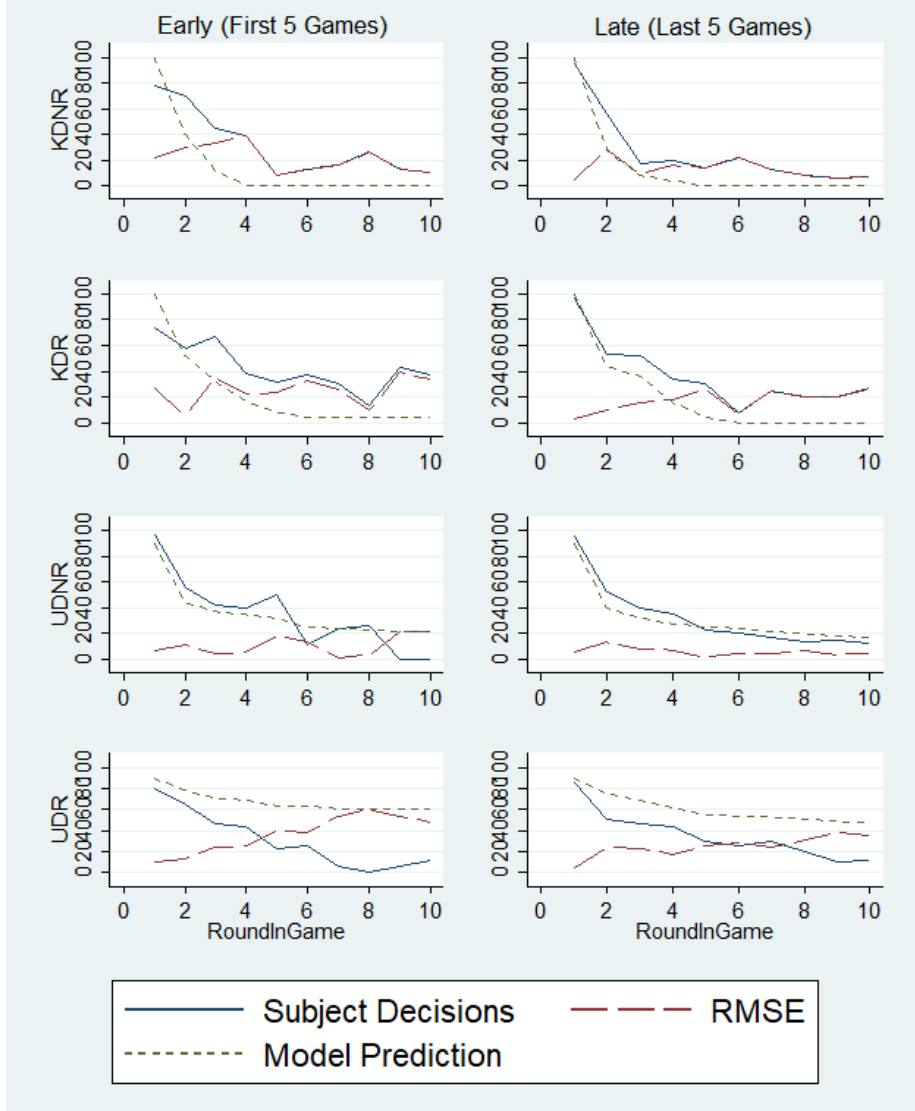


Figure 5: Average deviations from the regret minimizing model, early and late rounds of sessions. Theoretical predictions are based upon the actual current x_{max} values in an experimental game.

5.1.1 Bifurcated drawing sequence

We define a “bifurcated drawing sequence” as a 10-round game in which there is a one-time-only switch-over from innovation to imitation. Our theory predicts that when the distribution is known, as in our KD treatments, all games should involve such bifurcated drawing sequences. In neither the no recall nor the recall condition of the KD treatments do subjects bifurcate perfectly ($p < 0.01$, two-tailed Mann-Whitney U-test). However, when subjects do bifurcate, the maximum draw achieved within a game is significantly higher

than when they do not bifurcate ($p < 0.01$, two-tailed Mann-Whitney U-test). This finding indicates that when the theory is followed more closely, payoff outcomes are better. Indeed, when subjects adhere to the predicted bifurcating behavior, their payoffs are, on average, $\sim \$2.09$ higher than when they do not, a significant difference. Moreover, there is a significant negative correlation between earnings and the root mean squared error ($p < 0.01$, two-tailed Mann-Whitney U-test). These two facts indicate that when subjects play according to the theoretical predictions, they stand to earn considerably more than if they depart from these predictions.

5.2 Threshold Analysis

5.2.1 Known Distribution

We next consider whether subjects in the known distribution treatments (KDR and KDNR) were playing according to the thresholds predicted by the theory as reported in Table 3. We estimate each subject i 's threshold using logit regressions of the form $Pr[\text{Innovation}_i | x_{max,i}] = \alpha + \beta x_{max,i} + \varepsilon_i$, where $Pr[\text{Innovation}_i | x_{max,i}]$ is an indicator variable for whether a subject attempted to innovate (draw=1) or not (imitate=0), conditional on the current value of $x_{max,i}$ that he/she faced. After estimating this equation, we take the ratio of the estimated values $\frac{-\hat{\alpha}}{\hat{\beta}}$, as an indicator of the threshold of indifference between drawing and not drawing.⁸ Errors are clustered at the subject level. We estimate thresholds at both the game and treatment levels.

⁸In a logistic regression, we estimate $Pr[\text{Innovation} | x_{max}] = [1 + \exp(-\alpha - \beta x_{max})]^{-1}$. The indifference threshold is the value of x_{max}^* for which $Pr[\text{Innovation} | x_{max}] = 0.5$, which is equal to the ratio of the estimates, $\frac{-\hat{\alpha}}{\hat{\beta}}$.




Figure 6: Estimated threshold in games by treatment

Figure 6 shows the average of the estimated thresholds in every game for the KDNR and KDR treatments along with time trends and the model threshold values for x_{max} for which subjects are predicted to be indifferent between innovating and imitation arising from Propositions 3 and 4. Figure 6 supports the notion that subjects' behavior is consistent with the model when they are in the no recall condition, KDNR. In the KDR treatment, the trend in estimated thresholds is more strongly positive, and subjects in the KDR treatment had marginally significantly higher estimated indifference points in the last five games as compared with the first five games ($p = .10$, two-tailed Mann-Whitney U-test), indicating that subjects are adjusting their behavior with experience.

	KDNR			KDR		
	Games 1-5 (1)	Games 6-10 (2)	All Games (3)	Games 1-5 (4)	Games 6-10 (5)	All Games (6)
Prediction	40.00	40.00	40.00	46.86	46.86	46.86
Average (s.e.)	42.54 (3.47)	43.10 (3.49)	43.76 (2.73)	44.13 (4.61)	58.70 (1.81)	51.55 (2.75)
p-value	0.47	0.38	0.17	0.56	<0.01	0.09
Obs	500	500	500	500	500	500

Table 4: Thresholds by treatment

Table 4 reports the average thresholds (Average) for which subjects are estimated to be indifferent between drawing and not drawing in each treatment and the associated standard errors using all data of each known distribution treatment (500 observations).⁹ The p-values report on tests of the null hypothesis of no difference between these average estimated

⁹A corresponding session-level threshold analysis is available in the online Appendix A.5 in Table 11.

thresholds and the predicted thresholds of the theory.¹⁰ We further average over the first five and last five games to show the effects of learning. We acquire standard errors using 1000 repetitions of nonparametric bootstrapping of the logit regression specification given earlier. We see that in both treatments, the data align closely with model predictions. We do note, however, that subjects in games 6-10 of KDR stray significantly above the predicted threshold ($p < 0.01$, two-tailed Mann-Whitney U-test), a result of a more strongly positive trend in estimated thresholds in this treatment. We further observe that over all games, the average threshold in KDR is significantly greater than in KDNR, which is also predicted by the model. These findings indicate that our model provides reasonable predictions about both the *level* and the rank *order* of the innovate/imitate thresholds used by subjects.

5.2.2 Unknown Distribution

When analyzing decisions under unknown distributional information, we can no longer use our logit estimation strategy since each p represents the probability that minimizes regret, given a certain x_{max} . Thus, when a regret minimizing agent plays $p = 0.5$, it does not represent indifference between innovating and imitating. Instead, it is the mixing probability that minimizes regret given “worst-case” priors on the distribution of x . Therefore, for the unknown distribution case, we compare the mixed strategy prediction of the regret minimization theory which conditions on the value of x_{max} (see Table 3 with subjects’ own probability of taking a draw (innovation)). This analysis is shown in Figure 7. We interpret these probabilities to draw as the average p which minimizes regret subject to each subject’s subjective payoff distribution.

¹⁰Unless otherwise stated, each table which compares experimental data to theoretical predictions, as in Table 4, uses Wilcoxon signed-rank tests. The results of these tests is summarized in the “p-value” row of each table.

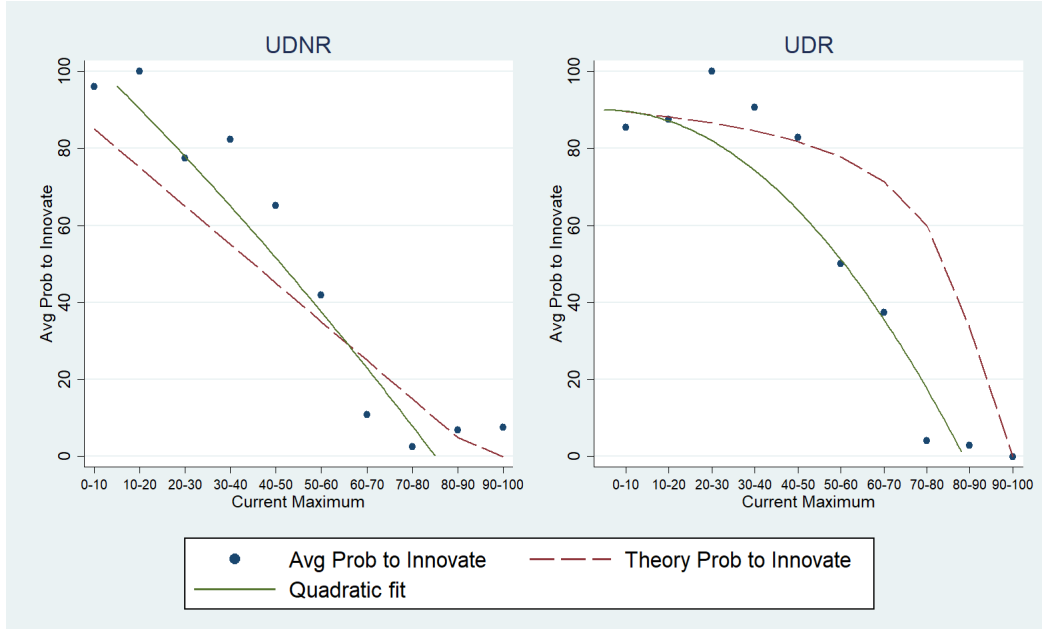


Figure 7: Average maximum draws in games by treatment

This analysis also shows the predicted regret minimizing probabilities of the theory as well as a quadratic trend line fit to the experimental data. Using t-tests, we find that there is no significant difference in the probabilities to draw as predicted by the model and found in the experimental data ($p = 0.79$) in the no recall condition (UDNR), but that there is a significant departure between predicted and actual probabilities to draw in the treatment with recall (UDR, $p < 0.01$). As in the case of the known distributions, subjects appear to have greater difficulty incorporating the incentive structure of the recall condition into their decision-making. Turning to the quadratic fit of the data, our theory is supported in both treatments by the result that decision probabilities are only influenced linearly by the current maximum draw ($p = 0.06$, two-tailed t-test) in UDNR and only by the quadratic term in UDR ($p = 0.08$, two-tailed t-test). This finding is important because our theory of regret minimization shows that decisions should be linear in the current maximum in UDNR and non-linear in the current maximum in UDR.

5.2.3 Subjective Expected Utility

It is possible that instead of minimizing regret, subjects in our unknown distribution treatments instead chose to maximize subjective expected utility. Though we gave the subjects

no information at all about the nature of the unknown distribution in the UDR and UDNR treatments, subjects might have (subjectively) applied an uniformed prior and assumed that the distribution was *uniformly* distributed over $[0, 100]$. Under such a distributional assumption, subjects would be predicted to all stop innovating at $x_{max} = 40$ in the UDNR treatment and at $x_{max} = 80$ in the UDR treatment¹¹. Instead, using the same estimation strategy outlined in Section 5.2.1, we find that the empirical stopping thresholds in our experimental data from the UDNR and UDR treatments are 48.14 and 52.31, respectively. Using a two-tailed z-test, we find that these values are both significantly different from the uniform prior predicted values at the $p = 0.001$ level. This finding suggests that subjects did not act consistent with a belief that the distribution being drawn from was uniform and that our instructions were successful in creating some ambiguity regarding the actual distribution they faced in the UDNR and UDR treatments.

Summarizing our results thus far, we have:

Finding 1. *There is strong support for Hypothesis 1. The probabilities of drawing from the distribution match those predicted by the theory, especially in the KDNR, KDR, and UDNR treatments. When considering the case of known distributions, subjects follow the predicted threshold stopping rule (KDNR and KDR). When the distribution is unknown, subjects employ probabilistic strategies which resemble closely the regret minimizing theoretical predictions.*

5.3 Expected Maxima Analysis

In the next two sections, we use numerical methods to determine theoretical predictions regarding expected maxima and the number of draws (innovation decisions) that would arise if subjects were playing according to the theory. Specifically we compare our experimental data with 100,000 independent simulations of our 10 round innovation/imitation game in which the simulated agents play strategies in accordance with the deterministic or probabilistic strategies predicted by our theory. For each simulated game we collect the value of the maximum draw and the total number of draws made and we use the distribution of these 100,000 simulations¹² for comparisons with our experimentally generated data.

¹¹Proofs for these thresholds are provided in the online Appendix.

¹²These distributions can be found in the online Appendix.

We first investigate subject behavior via the expected maxima, taking the maximum draw within a game in each treatment as our measure of interest.

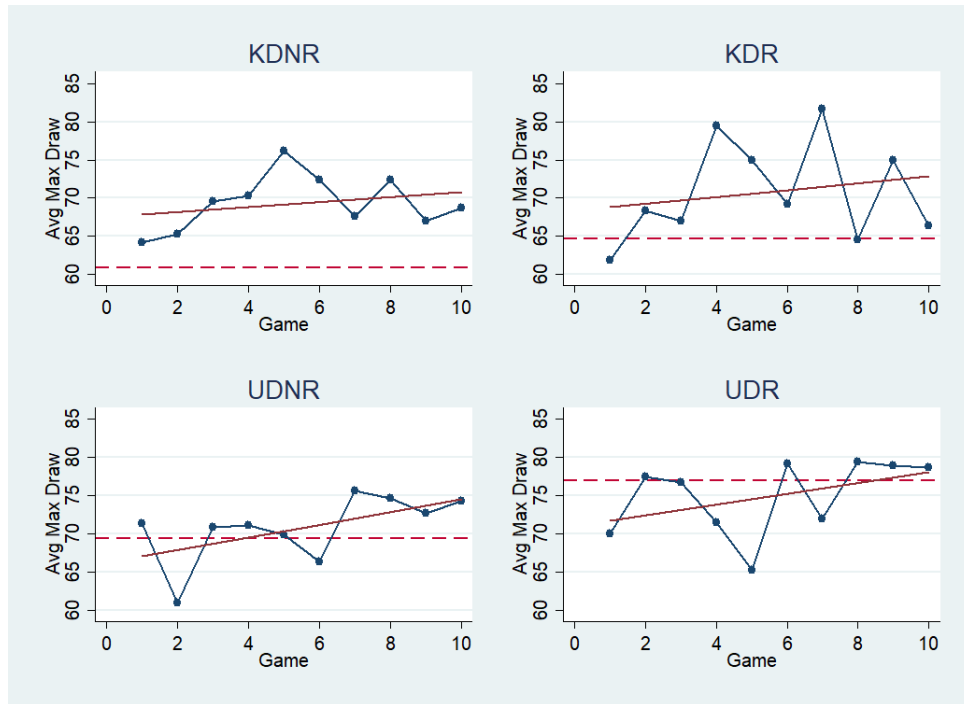


Figure 8: Average maximum draws in games by treatment

Figure 8 shows the average maximum draw in the experimental data, along with a linear trend line, and a horizontal line indicating the mean maximum draw from our 100,000 simulations. The figure reveals that, while the average maximum increased over the 10 games in the treatments with unknown distributions, it remained mostly flat in the treatments with known distributions. Table 5 reports the results of pairwise comparisons of the average maximum draws in games across treatments using t-tests, Mann-Whitney U-tests, and Kolmogorov-Smirnov tests.¹³

¹³A corresponding session-level analysis is available in the online Appendix A.5 in Table 13.

<i>test</i>	KDNr			KDR			UDNR			UDR		
	Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games
KDNr												
<i>t-test</i>												
<i>U-test</i>	x			0.696	0.570	0.495	0.943	0.274	0.497	0.284	0.007	0.007
<i>KS-test</i>				0.720	0.698	0.647	0.892	0.313	0.549	0.273	0.013	0.012
KDR				0.994	0.906	0.864	0.994	0.468	0.864	0.468	0.040	0.040
<i>t-test</i>							0.644	0.602	0.974	0.538	0.021	0.045
<i>U-test</i>	x				x		0.554	0.347	0.901	0.587	0.015	0.030
<i>KS-test</i>							0.699	0.468	0.864	0.468	0.037	0.068
UDNR												
<i>t-test</i>												
<i>U-test</i>	x				x			x		0.247	0.045	0.032
<i>KS-test</i>										0.204	0.056	0.028
										0.281	0.281	0.112

Note: Tests always compare 50 observations from one treatment to 50 observations from another.

Table 5: Differences in maximum draws in a game between treatments

		No Recall			Recall		
		Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games
		(1)	(2)	(3)	(4)	(5)	(6)
Known Dist.	Prediction	60.90	60.90	60.90	64.71	64.71	64.71
	Average	69.07	69.59	69.33	70.35	71.32	70.83
	(s.e.)	(1.59)	(1.60)	(1.58)	(1.66)	(1.42)	(1.53)
	p-value	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
	Obs	50	50	50	50	50	50
Unknown Dist.	Prediction	72.61	72.61	72.61	77.00	77.00	77.00
	Average	68.85	72.69	70.77	72.17	77.64	74.90
	(s.e.)	(1.56)	(1.17)	(1.40)	(1.26)	(1.23)	(1.29)
	p-value	0.12	0.97	0.28	0.02	0.77	0.16
	Obs	50	50	50	50	50	50

Table 6: Predicted versus actual average maximum draws by treatment

We also compare how close our subjects perform in terms of the average maximum draw to what is predicted by our model.¹⁴ The results are reported in Table 6¹⁵, which compares average maxima in the experimental data relative to the theoretical predictions across the four treatments, for the first five, last five and all 10 games. The results there suggest that subjects behave remarkably similar, in terms of their average maximum draw, with the predictions of the model in the cases where the distribution is unknown, and achieve somewhat greater maxima on average than predicted when the distribution is known. We further observe that, consistent with the theory, the average maximum draw is increased somewhat by the ability to recall, relative to the no recall case, but only significantly so in the unknown distribution case (Table 5).

5.4 Expected Number of Draws

We take advantage of the fact that the expected maxima simulation also provides estimates of the expected number of times subjects will draw (innovate) during a 10 round game. In Figure 9 we compare the predicted number of draws from the simulations (shown as horizontal lines) with the average actual number of draws made in each of the 10 games, along with a linear time trend fit to the experimental data. As can be seen, this trend is flat

¹⁴While we do not report the standard errors from the 100,000 simulations, our t-tests make use of these simulated standard errors. Further, we compare the simulated and empirical distributions of maximum draws in a game using Kolmogorov-Smirnoff and Mann-Whitney U-tests. In the following section we use these same tests to examine the expected number of draws.

¹⁵A corresponding session-level analysis is available in the online Appendix A.5 in Table 12.

and, indeed, the difference between the number of draws in the first and last five games of a treatment is never significantly different.

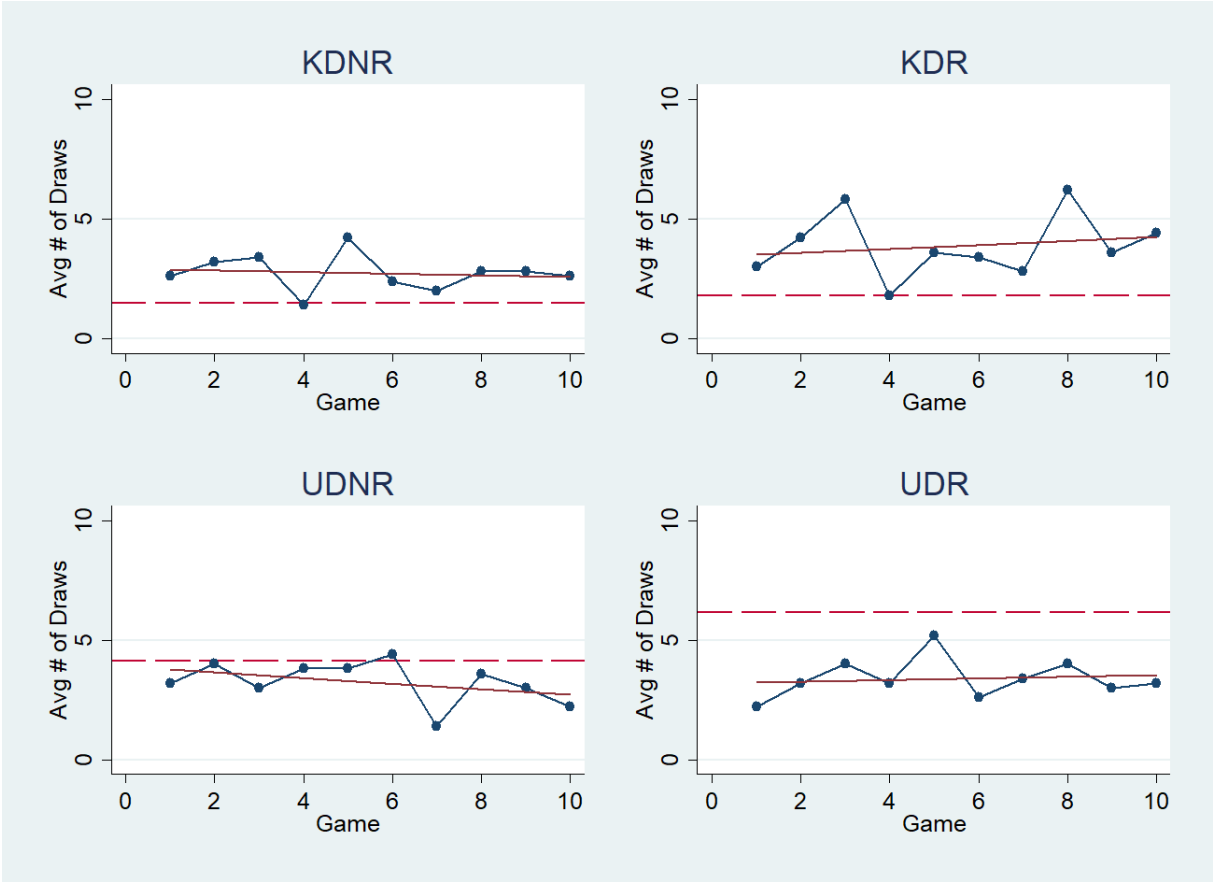


Figure 9: Average number of draws in games by treatment

Table 7 presents the mean number of draws from the simulations along with the average and standard errors from the experimental data over the first 5, last 5 and all 10 games of the four treatments.¹⁶

Figure 9 and Table 7 reveal that the model predictions for the expected number of draws are closest to the experimental data in the two no recall treatments (KDNR and UDNR) and furthest from the data in the two no recall treatments (KDR and UDR). This is in line with the finding from the deviation analysis, suggesting that subjects are under-utilizing the benefits of recall. Still, for all 4 treatments, using t-tests, U-tests, and KS-tests, we can reject the null hypothesis that the average number of draws in a simulated game is the same

¹⁶A corresponding session-level threshold analysis is available in the online Appendix A.5 in Table 14.

as in the experimental data at the 0.01 level. We further note that in our known distribution treatments subjects tend to over innovate, while in our unknown distribution treatments, subjects tend to under innovate, which may simply reflect the ambiguity that they face in the unknown distribution treatments.

Table 8 reports on the results of all pairwise comparisons of the average number of draws between the different treatments using a variety of statistical tests.¹⁷ The only statistically or marginally significant differences exist between KDNR and KDR and between KDNR and UDR. In both comparisons, KDNR has significantly fewer draws (innovation decisions), which is in line with the predictions of the model.

		No Recall			Recall		
		Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games
		(1)	(2)	(3)	(4)	(5)	(6)
Known Dist.	Prediction	1.47	1.47	1.47	1.78	1.78	1.78
	Average	2.96	2.52	2.74	3.68	4.08	3.88
	(s.e.)	(0.25)	(0.16)	(0.21)	(0.31)	(0.29)	(0.30)
	p-value	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
	Obs	50	50	50	50	50	50
Unknown Dist.	Prediction	4.13	4.13	4.13	6.16	6.16	6.16
	Average	3.56	2.92	3.24	3.56	3.24	3.40
	(s.e.)	(0.25)	(0.23)	(0.46)	(0.23)	(0.31)	(0.27)
	p-value	0.07	<0.01	<0.01	<0.01	<0.01	<0.01
	Obs	50	50	50	50	50	50

Table 7: Predicted versus actual number of draws by treatment

We summarize these findings as follows:

Finding 2. *We find mixed support for Hypothesis 2. Regarding expected maxima, subjects match closely with the simulated distributions of maximum draws in the UDNR and UDR treatments, but are significantly different from the predictions of the KDNR and KDR treatments. Regarding the expected number of draws, in all treatments but KDNR, subjects significantly underdraw from the distributions, and only in the KDNR treatment do they overdraw.*

5.5 Risk Aversion and Other Individual Characteristics

We next consider if there is any relationship between individual characteristics and the propensity to draw (innovate) across our four experimental treatments. We were particularly

¹⁷A corresponding session-level analysis is available in the online Appendix A.5 in Table 15.

<i>test</i>	KDNr			KDR			UDNR			UDR		
	Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games
KDNr												
<i>t-test</i>				0.219	0.002	0.003	0.250	0.323	0.130	0.232	0.149	0.060
<i>U-test</i>	x			0.220	0.006	0.005	0.187	0.505	0.126	0.123	0.428	0.087
<i>KS-test</i>				0.281	0.155	0.068	0.155	0.994	0.178	0.281	0.468	0.112
KDR							0.835	0.036	0.107	0.830	0.174	0.244
<i>t-test</i>					x		0.922	0.050	0.200	0.898	0.131	0.330
<i>U-test</i>							0.906	0.699	0.711	0.994	0.281	0.964
<i>KS-test</i>										1.00	0.562	0.664
UDNR								x		1.00	0.790	0.742
<i>t-test</i>	x				x					1.00	0.994	1.00
<i>KS-test</i>												

Note: Tests always compare 50 observations from one treatment to 50 observations from another.

Table 8: Differences in expected number of draws in a game between treatments

concerned about the role that might be played by differences in individual risk preferences. As we use the elicitation procedure of Dave et al. (2010), we followed their prescriptive advice of using the risk elicitation choices to make a binary classification of subjects as either risk-averse or risk-neutral.¹⁸ If there is some curvature in the regret function, then the threshold stopping value of a subject may differ from our theory, which assumes risk-neutral preferences. A risk-averse subject would be predicted to stop innovating earlier than a risk-neutral agent and a risk-loving subject would be predicted to stop innovating later than a risk-neutral agent. In addition to risk preferences, we also consider the role of age, sex, GPA, cognitive reflection test (CRT) score, and a quantitative reasoning (QR) score for innovation decisions.¹⁹ Finally, we test whether a subject’s individual position order in the sequence of moves mattered for their decision to innovate, even after taking into account the value of x_{max} that they faced. While the position order was not known, subjects might infer they were the first or a later mover by the value of x_{max} , which is initially 0.²⁰ Even accounting for the value x_{max} , early movers (particularly those with social preferences) might be more predisposed to innovate, as compared with late movers out of concerns for the welfare of other group members.

The least squares estimates from our regression of the propensity to innovate on individual characteristics are reported in Table 9.

¹⁸That is, we classify a subject as risk-averse if they chose the first, second, third, or fourth option in the elicitation found in Figure 2 and risk-neutral if they chose option five or six.

¹⁹We impute a QR score by associating the mean GRE quantitative score associated with the major of each subject (such scores can range from 130-170). We had data on the major of each subject participant, and we used the mean GRE quantitative score for each major as reported by the Educational Testing Service, which administers the GRE. For undeclared majors, we used the mean GRE quantitative score across all test takers, which was 159.

²⁰This concern motivates our additional session-level analyses, found in Section A.5, since individual decisions may be contaminated by the effects of learning from others both within games and across games.

Treatment	(1) All	(2) KDNR	(3) KDR	(4) UDNR	(5) UDR
x_{max}	-1.142*** (0.0383)	-1.109*** (0.0786)	-1.133*** (0.0787)	-1.270*** (0.0785)	-1.106*** (0.0671)
Risk Averse	-2.237 (2.025)	-1.730 (3.110)	-6.430 (4.819)	1.752 (3.997)	1.329 (3.627)
Age	0.554 (0.503)	-0.0824 (0.620)	1.681 (1.049)	1.622 (1.665)	-0.468 (0.966)
Female	-3.522 (2.199)	3.698 (3.811)	-6.623 (4.929)	-5.367 (5.079)	-5.792 (4.396)
GPA	-0.787 (0.869)	-3.596** (1.626)	-1.359 (2.007)	-1.482 (2.044)	0.979 (1.122)
CRTscore	-0.276 (0.929)	-1.717 (1.619)	0.697 (2.125)	2.404 (2.216)	-0.982 (1.615)
QRscore	0.0358 (0.211)	0.0684 (0.465)	0.200 (0.451)	-0.294 (0.395)	0.452 (0.461)
Position Order	-0.702** (0.330)	-1.561** (0.672)	0.512 (0.701)	-0.985 (0.661)	-0.627 (0.549)
Constant	87.99** (35.58)	106.9 (78.05)	59.66 (82.86)	141.4** (63.51)	55.66 (74.57)
Observations	2,000	500	500	500	500
R-squared	0.456	0.501	0.412	0.503	0.443

Robust standard errors clustered at the subject level in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 9: Individual characteristics and the propensity to draw (innovate)

Table 9 reveals that subjects draw less frequently the higher is the value of x_{max} that they saw when it was their turn to draw, which is always consistent with the theory. Controlling for the value of x_{max} , we find little evidence that any of subjects' individual level characteristics significantly affected their propensity to draw across the four treatments, with one exception: the order in the game negatively impacts the propensity to draw ($p = 0.034$). This finding suggests that, even controlling for the value of x_{max} , one's position order may matter even though the order of moves was *not* provided and was randomly assigned for each game. However, disaggregating the regression results by treatment, we see that position order only matters in the KDNR treatment, as in the other three treatments, the effect of a subjects' position order is insignificant.

We summarize these findings as follows:

Finding 3. *Hypothesis 3 is rejected. Personal characteristics are observed to have insignificant correlations with subject deviations from predicted behavior.*

6 Concluding Remarks

We have provided a model of the decision by firms as to whether to innovate new products/services or production processes or to imitate those of other firms. We further consider case where innovation precludes imitation (no recall) and where it does not (recall), reflecting different costs of innovation and imitation. Our model makes several key contributions to the regret minimization and strategic experimentation literatures. First, we merge the two into a model of experimentation in a single shot game, where we model firms as regret minimizers. We predict four distinct behaviors depending on knowledge of the innovation distribution and the ability to recall in the event of failed innovations. These different conditions mirror aspects of the real world, such as the presence or absence of intellectual property rights affecting imitation, and the strength of beliefs on the returns to research and development of new innovations. Our model predicts differences in the probabilities with which a regret minimizing firm innovates conditional on the current maximum draw, which also influences the expected maximum draw to be obtained within a game. Our model can explain why there can be substantial initial innovation in an industry that later tapers off, as firms switch from a strategy of innovating to imitating, resulting in a homogeneous selection of products or services across firms, as seems to have occurred in the smart phone industry.

We develop a novel experimental design to test the implications of our model. To our knowledge, this is the first experiment to measure how well subjects' behavior corresponds to the predictions of regret minimization where the states of nature form a continuum. We find that subject behavior is largely consistent with the regret minimization model predictions. We find that the biggest driver of differences between model predictions and behavior in most of our subject- and game- level measures whether recall (imitation following failed innovation) is possible, while knowledge of the distribution plays a less central role in explaining differences.

Further, when we compare average propensities to innovate, we find that regret minimiza-

tion describes patterns in subject behavior well. Again, regret minimization best describes subject behavior when there is no recall, which is again mirrored in results pertaining to the average maximum reached within a game. Regret minimization describes behavior just as well in early rounds as it does in later rounds, which indicates that the beliefs of our subjects are not updated enough to move them away from their regret minimizing behavior. This, in turn, points to the fact that our environment is not similar enough to the classical strategic experimentation paradigm to generate results consistent with that theory. It is likely that there are too few observations for subjects to condition on and estimate the payoff distribution in a meaningful way.

Since recall does not substantially change the propensity to innovate and serves to increase deviations between our model and the data, we conclude that recall does not help in exploring the payoff distribution. This may be due to the fact that while recall lowers innovation costs, it also lowers the difference in payoffs between innovating and imitating, especially after a sufficiently high maximum has been reached within a game. This smaller differential, compared to the case where there is no recall, may lead to subjects thinking less critically about what their best decision is when attempting to minimize regret under the recall condition. In essence, the reduced salience of decision making brought about by the recall condition leads to less critical thinking about an agent's optimal strategy. This effect is only exacerbated in the UDR treatment, where subjects are presented with the additional complication of decision making under ambiguity, allowing for ambiguity aversion (Borghans et al., 2009) to add to the increased deviations brought about by recall.

In future research, it would be of interest to modify our design to explore some more realistic scenarios that might alter our main findings. For instance, innovation may not perfectly preclude imitation as in our no recall treatment. It may be that both innovation and imitation are always possible as in our recall treatment, but that there is always some cost to imitation that would more naturally limit its use in the recall setting. It may be that innovation and imitation interact in complementary ways that our model does not yet capture as emphasized in the industry studies of Raustiala and Sprigman (2012). Further, since one innovation often serves as a complement to new innovations (Romer 1994), a useful change to our model would be to make the payoff distribution *endogenous* by allowing successful

innovation to shift the parameters of the distribution. Depending on the differences in the changes to the parameters, innovation may continue indefinitely or stop earlier than it might have if the distribution was static. One could also model changes to the distribution as ambiguous to capture behavior when the changes in the state of nature are not known. Another reasonable modification would be to let subjects play the game for a longer period of time, e.g., cycling through the n-round game multiple times, allowing for a more realistic depiction of the research and development process and transforming the game into one of strategic interaction and repeated play, allowing for more learning to develop. Finally, we believe that additional attention should be given to regret minimization, especially where it can be contrasted with the predictions of expected utility maximization. We leave all of these extensions and considerations to future research.

References

- Abdellaoui, M., Klibanoff, P., Placido, L., 2015. Experiments on compound risk in relation to simple risk and to ambiguity. *Management Science* 61, 1306-1322.
- Baillon, A., Schlesinger, H., Van de Kuilen, G., 2018. Measuring higher order ambiguity preferences. *Experimental Economics* 21, 233-256.
- Bergemann, D., Schlag, K., 2011. Robust monopoly pricing. *Journal of Economic Theory* 146, 2527-2543.
- Boeri, M., Longo, A., Doherty, E., Hynes, S., 2012. Site choices in recreational demand: a matter of utility maximization or regret minimization? *Journal of Environmental Economics and Policy* 1, 32-47.
- Boeri, M., Longo, A., Grisolia, J.M., Hutchinson, W.G., Kee, F., 2013. The role of regret minimization in lifestyle choices affecting the risk of coronary heart disease. *Journal of Health Economics* 32, 253-260.
- Bolton, P., Harris, C., 1999. Strategic experimentation. *Econometrica* 67, 349-374.

- Borghans, L., Golsteyn, B.H.H., Heckman, J.J., Meijers, H., 2009. Gender differences in risk aversion and ambiguity aversion. *Journal of the European Economic Association* 7, 649-658.
- Brumfiel, G., 2017. Total failure: The world's worst video game, National Public Radio, May 31, <https://www.npr.org/2017/05/31/530235165/total-failure-the-worlds-worst-video-game>
- Chen, D.L., Schonger, M., Wickens, C., 2016. oTree - An open-source platform for laboratory, online and field experiments. *Journal of Behavioral and Experimental Finance* 9, 88-97.
- Chorus, C.G., Arentze, T.A., Timmermans, H.J.P., 2008. A random regret-minimization model of travel choice, *Transportation Research Part B: Methodological* 42, 1-18.
- Dave, C., Eckel, C.C., Johnson, C.A., Rojas, C., 2010. Eliciting risk preferences: When is simple better? *Journal of Risk and Uncertainty* 41, 219-243.
- Ellsberg, D., 1961. Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics* 75, 643-669.
- Engel, C., Kleine, M., 2015. Who is afraid of pirates? An experiment on the deterrence of innovation by imitation. *Research Policy* 44, 20-33.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 141-153.
- Gilboa, I., Schmeidler, D., 1994. Additive representations of non-additive measures and the Choquet integral. *Annals of Operations Research* 52, 4365.
- Halevy, Y., 2007. Ellsberg revisited: An experimental study. *Econometrica* 75, 503-536.
- Halpern, J.Y., Pass, R., 2012. Iterated regret minimization: A new solution concept. *Games and Economic Behavior* 74, 184-207.
- Hommes, C., Zeppini, P., 2014. Innovate or imitate? Behavioural technological change. *Journal of Economic Dynamics and Control* 48, 308-324.

- Hume, D., 1739. *A Treatise of Human Nature*. London: John Noon.
- Iwai, K., 1984. Schumpeterian dynamics: An evolutionary model of innovation and imitation. *Journal of Economic Behavior & Organization* 5, 159-190.
- Iwai, K., 2000. A contribution to the evolutionary theory of innovation, imitation and growth. *Journal of Economic Behavior & Organization* 43, 167-198.
- Keller, G., Rady, S., Cripps, M., 2005. Strategic experimentation with exponential bandits. *Econometrica* 73, 39-68.
- Knight, F.H., 1921. *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- König, M.D., Lorenz, J., Zilibotti, F., 2016. Innovation vs. imitation and the evolution of productivity distributions. *Theoretical Economics* 11, 1053-1102.
- Lippman, S.A., McCall, J.J., 1976. The economics of job search: A survey. *Economic Inquiry* 14, 155-189.
- Mortensen, D.T., 1986. Job search and labor market analysis. In: Ashenfelter, O.C., Layard, R., (Eds.). *Handbook of Labor Economics*, Vol. 2. Amsterdam: North-Holland, 849-919.
- Ofek, E., Turut, O., 2008. To innovate or imitate? Entry strategy and the role of market research. *Journal of Marketing Research* 45, 575-592.
- Raustiala, K., Sprigman, C., 2012. *The Knockoff Economy: How Imitation Sparks Innovation*. Oxford: Oxford University Press.
- Renou, L., Schlag, K.H., 2010. Minimax regret and strategic uncertainty. *Journal of Economic Theory* 145, 264-286.
- Romer, P.M., 1994. The origins of endogenous growth. *Journal of Economic Perspectives* 8, 3-22.

Simonson, I., 1992. The influence of anticipating regret and responsibility on purchasing decisions. *Journal of Consumer Research* 19, 105-119.

Stokey, N.L., Lucas, R.E. Jr., Prescott, E.C., 1989. *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.

Appendices to “Innovate versus Imitate: Theory and Experimental Evidence”

by John Duffy and Jason Ralston

A Threshold Proofs

The proof below will make extensive use of the generalized form of the triangular distribution, with lower bound a , upper bound b , and midpoint m , as defined by:

$$f(x) = \begin{cases} \frac{2(x-a)}{(b-a)(m-a)} & \text{if } a \leq x < m, \\ \frac{2}{b-a} & \text{if } x = m, \\ \frac{2(b-x)}{(b-a)(b-m)} & \text{if } m < x \leq b, \\ 0 & \text{elsewhere.} \end{cases} \quad (7)$$

When we parameterize the distribution with $a = 0$, $b = 100$, and $m = 50$, the above definition reduces to

$$f(x) = \begin{cases} \frac{2x}{5000} & \text{if } 0 \leq x < 50, \\ \frac{2}{100} & \text{if } x = 50, \\ \frac{2(100-x)}{5000} & \text{if } 50 < x \leq 100, \\ 0 & \text{elsewhere.} \end{cases} \quad (8)$$

A.1 Known Distribution, No Recall

We first consider the no recall case. Thus, the decision whether to draw or not draw for a risk neutral individual is defined as

$$E(X) - x_{\max} \geq 10$$

Since $E(X) = \mu = 50$, the threshold will be set as $x_{\max} = 40$.

A.2 Known Distribution, Recall

Here the agent is allowed to condition on the current x_{\max} , which, in effect, raises the lower support of X . This leads to a dynamic reduction in cost that decreases as x_{\max} increases. Recall that the decision to draw or not draw for a risk neutral agent is based upon

$$E[r(x_{\max}, X, j)|X > x_{\max}] = \begin{cases} E[(X - c) - x_{\max}|X - c > x_{\max}] \times P(X - c > x_{\max}) & \text{if imitate} \\ E[x_{\max} - (X - c)|x_{\max} < X \leq x_{\max} + c] \times P(x_{\max} < X \leq x_{\max} + c) \\ \quad + c \times P(X \leq x_{\max}) & \text{if innovate.} \end{cases}$$

Due to the distribution's piece-wise nature, and the decision criterion's dependence on x_{\max} , we derive the threshold stopping value using numerical methods. For simplicity, we show the graphs of expected regret from innovation (yellow) and from imitation (blue) in Figure 10.




Figure 10: Expected regret from innovation (yellow) and imitation (blue) as a function of x_{\max}

The intersection of the two curves occurs at $x_{\max} \approx 46.89$.

A.3 Subjective expected utility threshold proofs

For the following proofs, we assume that subjects believe the distribution they are drawing from is a uniform distribution with support spanning from $a = 0$ to $b = 100$, which entails the pdf defined below.

$$f(x) = \begin{cases} \frac{1}{100} & \text{if } 0 \leq x \leq 100 \\ 0 & \text{elsewhere} \end{cases} \quad (9)$$

A.3.1 No recall

In the case of no recall, the payoff function is defined by

$$\pi_i = \begin{cases} x_i - c + e & \text{if innovate} \\ x_{\max} + e & \text{if imitate} \end{cases}$$

Thus, when a risk-neutral agent must decide whether to innovate or imitate, they only consider the expected payoff of innovation versus imitation. They will be indifferent between the two alternatives when their expectations are equal. Specifically this happens when

$$E[X] - c + e = x_{\max} + e.$$

Keeping in mind that X is distributed according to the uniform distribution and that $c = 10$ in our experiment, we arrive at the conclusion that our agent will be indifferent between innovation and imitation when $x_{\max} = 40$.

A.3.2 Recall

Under recall, the payoff function is given by

$$\pi_i = \begin{cases} \max\{x_i, x_{\max}\} - c + e & \text{if innovate} \\ x_{\max} + e & \text{if imitate} \end{cases}$$

Thus, when a risk-neutral agent must decide whether to innovate or imitate, they will again only consider the expected payoff of innovation versus imitation. They will be indif-

ferent between the two alternatives when their expectations are equal. This happens when

$$E[X|X > x_{max}] - c + e = x_{max} + e,$$

which is equivalent to

$$\frac{\int_{x_{max}}^{100} x \frac{1}{100} dx}{\int_{x_{max}}^{100} \frac{1}{100} dx} - c + e = x_{max} + e.$$

This can be simplified to

$$100 + x_{max} - 2(c - e) = 2(x_{max} + e)$$

Again, since $c = e = 10$, we can solve for x_{max} . Specifically, $x_{max} = 80$.

A.4 Use of Randomization Device

Our experimental design features a randomization device that subjects may use if they are uncertain about their innovation/imitation choice. We included this feature because our regret minimization framework predicts that agents will randomize when the distribution of outcomes is unknown, but not when it is known. Thus, if a subject were to prefer resolving randomization using a computer instead of internally, the randomization device feature could be useful.

	Known Dist	Unknown Dist
No Recall	15	21
Recall	22	17

Table 10: Randomization device usage

We report usage of the randomization device in Table 10. Table 10 indicates that the number of subjects using the randomization device does not vary by distributional knowledge (Mann-Whitney, $p = 0.884$). While this process measure does not support the hypothesis that regret minimizers are more likely to randomize under conditions of distributional uncertainty, we argue that our subjects may randomize in a wide variety of ways that do not involve the use of our computer-aided randomization device before selecting whether to in-

novate or imitate. We think the analysis of section 5.2.2 more clearly addresses the relevance of the mixed strategy in the environment without distributional information (ambiguity).

A.5 Session-Level Analyses

Since the same group of subjects repeatedly interact with one another over the 10 games of each session, there may be learning within a session, even though feedback about the behavior of other was limited; subjects only observed the value of x_{max} when making their decision. Still, it is of interest to ask whether our results continue to hold using more conservative tests on average session-level data only.²¹ Since each treatment consisted of just five sessions, our statistical power is limited. Nevertheless, we continue to find significant differences between our treatments using these session-level observations.

We begin by examining threshold stopping behavior in our treatments with known distributions. The theoretical stopping rule, estimated sample averages, and the standard error of the estimates are provided in Table 11. This table 11 shows that, in general, the session-level data is statistically indistinguishable from model predictions. The only exception is for KDR in games 6 - 10 ($p < 0.01$), which is also what we found using the individual-level data in Section 5.1. Although thresholds across all games and in games 1 - 5 were not significantly different between the two known distribution treatments, we do note that in games 6 - 10 we observe higher thresholds in KDR than in KDNR (t-test, $p = 0.022$), which is consistent with our previous individual-level data analysis (compare with Table 4).

	KDNR			KDR		
	Games 1-5 (1)	Games 6-10 (2)	All Games (3)	Games 1-5 (4)	Games 6-10 (5)	All Games (6)
Prediction	40.00	40.00	40.00	46.86	46.86	46.86
Average	44.52	46.27	45.02	44.04	58.16	51.02
(s.e.)	(4.44)	(3.85)	(3.82)	(7.38)	(1.68)	(5.03)
p-value	0.35	0.14	0.35	0.89	0.04	0.35
Obs	5	5	5	5	5	5

Table 11: Session level thresholds by treatment

²¹In section 5 of our paper, we use either individual-level data (10 individual decisions/game \times 10 games/session \times 5 sessions/treatment = 500 observations per treatment) or game-level data (10 games/session \times 5 sessions/treatment = 50 observations per treatment) to analyze the predictions of our model.

We next consider the average maximum draw in a game using session-level data. Table 12 shows that our regret minimization model makes accurate predictions for the average maximum draw in the cases where the distribution is unknown, but that participant behavior deviates significantly in the known distribution setting. These findings are again in line with our results using the game-level analysis (compare with Table 6). Table 13 gives some evidence that, even using session-level averages, the maximum draw reached within a game is significantly higher in UDR than in KDNR and KDR. This finding is also in line with our prior analyses (compare with Table 5).

		No Recall			Recall		
		Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games
		(1)	(2)	(3)	(4)	(5)	(6)
Known Dist.	Prediction	60.90	60.90	60.90	64.71	64.71	64.71
	Average	69.07	69.59	69.33	70.35	71.32	70.83
	(s.e.)	(1.91)	(3.81)	(2.83)	(2.65)	(1.22)	(1.16)
	p-value	0.04	0.14	0.08	0.08	0.04	0.04
	Obs	5	5	5	5	5	5
Unknown Dist.	Prediction	72.61	72.61	72.61	77.00	77.00	77.00
	Average	68.85	72.69	70.77	72.17	77.64	74.90
	(s.e.)	(2.88)	(2.07)	(1.40)	(1.90)	(2.45)	(1.75)
	p-value	0.35	0.89	0.35	0.04	0.89	0.22
	Obs	5	5	5	5	5	5

Table 12: Session-level average maximum draws by treatment

We also compare the average number of draws taken within a game at the session level. These comparisons are displayed in Table 14. As in the game-level analysis (see Table 7), we see that subjects systematically draw too much in the known distribution settings and too little in unknown distribution settings. Table 15 shows that treatment differences found in Table 8 are weakened by the diminished statistical power. Nevertheless, we still find evidence that the number of draws taken during a game is substantially lower in KDNR than in KDR and UDR, suggesting that recall may have a slight positive effect on a subject’s willingness to draw, as per the predictions of the theory.

We also use session-level data to examine differences in deviations from the theory, as in Section 5.1. We find that session-level data provide similar results to individual-level data: the RMSE between behavior and theory is higher when subjects are not afforded the ability to recall (Mann-Whitney, $p = 0.049$). The RMSE in the KDR treatment is also significantly lower than in the UDR treatment (Mann-Whitney, $p = 0.009$). These findings

	KDNr			KDR			UDNr			UDR		
	Games 1 - 5	Games 6 - 10	All Games	Games 1 - 5	Games 6 - 10	All Games	Games 1 - 5	Games 6 - 10	All Games	Games 1 - 5	Games 6 - 10	All Games
KDNr	t-test			0.707	0.677	0.636	0.950	0.495	0.661	0.284	0.114	0.132
	U-test	x		0.602	0.602	0.917	0.754	0.465	0.917	0.347	0.175	0.047
	KS-test			0.819	0.329	0.819	0.819	0.819	0.819	0.819	0.819	0.082
KDR	t-test	x					0.711	0.585	0.972	0.591	0.050	0.089
	U-test				x		0.465	0.754	0.754	0.917	0.047	0.076
	KS-test						0.819	0.819	0.819	0.819	0.082	0.082
UDNr	t-test	x								0.363	0.161	0.102
	U-test				x					0.465	0.251	0.175
	KS-test									0.819	0.819	0.329

Note: Tests always compare 5 observations from one treatment with 5 observations from another.

Table 13: Differences in session-average maximum draws between treatments

confirm our initial assertion that recall increases deviations from theory, and that deviations are especially high in UDR.

The findings of this section suggest that, even using *session level-data* to avoid possible correlations between games within a session and learning between subjects within a game, that the treatment effects and differences we found using game-level or individual level data continue to obtain. Differences from the theory are especially pronounced for the UDR treatment, a result that is mirrored in our main analyses.

		No Recall			Recall		
		Games 1-5	Games 6-10	All Games	Games 1-5	Games 6-10	All Games
		(1)	(2)	(3)	(4)	(5)	(6)
Known Dist.	Prediction	1.47	1.47	1.47	1.78	1.78	1.78
	Average	2.96	2.52	2.74	3.68	4.08	3.88
	(s.e.)	(0.80)	(0.49)	(0.68)	(0.98)	(0.93)	(0.96)
	p-value	0.04	0.04	0.04	0.04	0.04	0.04
	Obs	5	5	5	5	5	5
Unknown Dist.	Prediction	4.13	4.13	4.13	6.16	6.16	6.16
	Average	3.56	2.92	3.24	3.56	3.24	3.40
	(s.e.)	(0.79)	(0.73)	(1.45)	(0.73)	(0.49)	(0.85)
	p-value	0.04	0.04	0.04	0.04	0.04	0.04
	Obs	5	5	5	5	5	5

Table 14: Session level number of draws by treatment

	KDNr			KDR			UDNR			UDR		
	Games 1 - 5	Games 6 - 10	All Games	Games 1 - 5	Games 6 - 10	All Games	Games 1 - 5	Games 6 - 10	All Games	Games 1 - 5	Games 6 - 10	All Games
KDNr												
t-test				0.222	0.028	0.051	0.093	0.389	0.106	0.110	0.193	0.072
U-test	x			0.207	0.01	0.093	0.141	0.435	0.094	0.169	0.316	0.059
KS-test				0.329	0.082	0.329	0.329	0.819	0.329	0.329	0.329	0.329
KDR												
t-test					x		0.810	0.141	0.235	0.814	0.297	0.386
U-test	x						0.399	0.143	0.249	0.454	0.465	0.346
KS-test							0.329	0.329	0.329	0.819	0.819	0.329
UDNR												
t-test					x					1.000	0.635	0.633
U-test	x									0.830	0.673	0.530
KS-test										1.000	1.000	0.819

Note: Tests always compare 5 observations from one treatment to 5 observations from another.

Table 15: Differences in session's average number of draws in a game between treatments

B Laboratory Instructions

We provide our laboratory instructions for the KDNR and UDR treatments, in that order. The KDR instructions are identical to the KDNR instructions, except that they included a full description of recall and how it worked, as is detailed in the UDR instructions. Similarly, the UDNR instructions are changed from the UDR instructions to remove all mention of how recall would work.

Instructions

Overview

Welcome to this experiment in the economics of decision-making. Funding for this experiment has been provided by the UC Irvine School of Social Sciences. We ask that you not talk with one another and that you silence your mobile devices for the duration of today's session.

For your participation in today's session you will be paid in cash at the end of the experiment. The amount you earn depends on the choices you make and on the choices made by others. Thus, it is important that you listen carefully and fully understand these instructions before we begin. There will be a short comprehension quiz following the reading of these instructions which you will all need to complete before we can begin the experimental session.

The experiment will make use of the computer workstations and all interactions by you and others will take place through these networked computers. You will interact anonymously with one another and your data records will be stored only by your ID number; your name or the names of other participants will not be revealed at any time during today's session or in any write-up of the findings from this experiment.

Today's session consists of two parts. You will receive instructions for part two at the end of part one.

In the first part of the session, you will participate in number of "games." Each game consists of a number of "rounds." In each round you will view some information and make a choice. Your choice, and possibly the choices of others determines the amount of points that you earn each round. At the end of the session, we will randomly select one game from all of the games played in today's session. Your point earnings from the chosen game will be converted into dollars at a conversion rate of \$0.15 per point earned. Your earnings from the chosen game and your \$7 show-up payment will be paid to you in cash and in private. You will also have the opportunity to earn additional earnings in the second part of the experiment.

Specific Details

There are N individuals in today's session. At the start of each new game, each individual will be randomly assigned a position number for the game. This position number indicates the round in the game, $1, 2, \dots, N$, at which you will be called upon to make a choice.

When it is your turn to make a choice you will see the Choice screen (you will see a waiting screen until that time). On the Choice screen you will be asked to make a choice. Specifically, you can decide whether or not to draw a number from a distribution having a mean of 50.00 and a standard deviation of 26.36. The distribution is shown in the computer screen and depicted below in Figure 1. The horizontal axis shows the numbers (in points) that you could draw, from (0.00 to 100.00) non-inclusive. The vertical axis reveals the probability or likelihood of drawing each possible number. As the distribution reveals,




Figure 1: The triangular distribution used in this experiment

the most likely outcome is 50.00, with the likelihood of numbers away from 50 declining equally in both directions.

Drawing a number from the distribution is costly. Specifically, a draw costs you 10 points. However, every individual is given an endowment of 10 points at the start of each new game, so the choice you face is whether to spend your endowment of 10 points drawing a number from the distribution.

Prior to making this choice, you are informed of the highest number that has been drawn by another participant in the current game. The choice you face is whether or not you want to try to draw a new number, at a cost to you of your 10 point endowment for the game.

If you choose not to draw a number, then your points for the game will equal the highest number chosen in the game so far plus your 10 point endowment.

If you choose to draw a number, then your points for the game we equal the number you drew for the game minus your 10 point endowment (the cost of drawing a number).

Please note the following:

- First, draws from the distribution are *with replacement* which means that the same number can be drawn more than once. The likelihood of drawing any number, as illustrated in Figure 1 and shown on your computer screen, does NOT change across all games played in today's session.
- Second, if the highest number drawn in the current game is 0.0, then EITHER you are the first person to make a choice in the current game OR no prior participant has chosen to draw a number in the current game.
- Third, if you choose to draw a number, it is possible that the number you draw is higher or lower than the current highest number drawn by another participant in the current game. The distribution shown in Figure 1 reveals the likelihood of each of the possible numbers you could possibly draw.

- Fourth, your earnings from drawing a number can be *higher* or *lower* than your earnings from NOT drawing a number. If you choose to draw a number and the number you draw is higher than the current highest number drawn, then your earnings from drawing that number will be higher than your earnings from not drawing a number only if the number you drew is at least 10 more than the current high number, since by drawing a number you lose your endowment of 10 points. On the other hand, if you choose to draw a number, your earnings will be lower than if you had not drawn a number whenever the number you drew is less than the current high number plus 10.

In making a choice of whether or not to draw a number for the current game, you face three options:

1. Do NOT draw
2. Enter a probability to draw
3. Draw

If you choose Do NOT draw, then you definitely DO NOT draw a number for the current game. If you choose Draw, then you definitely DO draw a number for the current game. If you choose Enter a probability to draw a number, then you must also enter a number (an integer) between 1 and 99 representing the probability, in percentage terms, that you will draw a number for the current game. After entering your percent chance of drawing a number, the computer program randomly draws a number (an integer) between 1 and 99 inclusive. If this randomly drawn number is less than or equal to your entered percent chance of drawing a number, then you will draw a number from the distribution for the current game: it will be automatically drawn for you; otherwise, you will not draw a number for the current game. Thus, the higher (lower) is the percent chance you enter for the probability of drawing a number, the more (less) likely it is that you draw a number in the current game.

Payment

The first part of today's sessions consists of 10 games. At the end of each game you will learn your points earned for the game, which, as explained above, depend on whether or not you draw a number. After all 10 games have been completed, one of the 10 games will be randomly chosen for payment. Each game has an equal chance of being chosen and so you will want to do your best in each game. Your points from the chosen game will be converted into dollars at the rate of 1 point = \$0.15 (15 cents). Thus, the more points you earn, the greater are your monetary earnings. In addition, you are guaranteed \$7 for showing up to today's experiment. You will be paid your show-up payment, together with your earnings from the first part of the experiment and your earnings from the second part of the experiment at the end of today's session. All payments will be made in cash and in private. At the end of this first part of the experiment, you will receive further instructions for how to complete part 2 of the experiment.

Questions?

Now is the time for questions. If you have a question about any aspect of these instructions, please raise your hand and an experimenter will answer your question.

Quiz

Before we start today's experiment we ask you to answer the following quiz questions that are intended to check your comprehension of the instructions. The numbers in these quiz questions are illustrative; the actual numbers in the experiment may be quite different. Before starting the experiment we will review each participant's answers. If there are any incorrect answers we will go over the relevant part of the instructions again.

1. Each game consists of _____ rounds. Your position in the game circle one: changes remains the same in each game.
2. Before deciding whether to draw a number you can see the highest number drawn so far in the current game. Circle one: True False
3. If you choose NOT to draw a number, then your point earnings will be equal to: Circle one:
 - a. the highest number drawn earlier in the current game, (or 0.00 if no number has been drawn yet)
 - b. the highest number drawn earlier in the current game (or 0.00 if no number has been drawn yet) plus your 10 point endowment.
 - c. your 10 point endowment.
4. If you choose to draw a number, you lose your 10 point endowment for the current game. Circle one: True False Do you get a new endowment of 10 points for each new game? Circle one: Yes No.
5. Consider the following scenario. The current highest number is 65.56. You choose to draw a number which turns out to be 73.21 What is your payoff in points for the game in this case? _____ What would have been your payoff in points if you did not choose to draw a number? _____
6. Consider the following scenario. The current highest number is 31.03. You choose to draw a number, which turns out to be 56.71. What is your payoff in points this case? _____ What would have been your payoff in points if you did not choose to draw a number? _____
7. At the end of the experiment, one game will be randomly chosen. Your point earnings from that game will be converted into dollars at the conversion rate of 1 point = \$_____.

Instructions

Overview

Welcome to this experiment in the economics of decision-making. Funding for this experiment has been provided by the UC Irvine School of Social Sciences. We ask that you not talk with one another and that you silence your mobile devices for the duration of today's session.

For your participation in today's session you will be paid in cash at the end of the experiment. The amount you earn depends on the choices you make and on the choices made by others. Thus, it is important that you listen carefully and fully understand these instructions before we begin. There will be a short comprehension quiz following the reading of these instructions which you will all need to complete before we can begin the experimental session.

The experiment will make use of the computer workstations and all interactions by you and others will take place through these networked computers. You will interact anonymously with one another and your data records will be stored only by your ID number; your name or the names of other participants will not be revealed at any time during today's session or in any write-up of the findings from this experiment.

Today's session consists of two parts. You will receive instructions for part two at the end of part one.

In the first part of the session, you will participate in number of "games." Each game consists of a number of "rounds." In each round you will view some information and make a choice. Your choice, and possibly the choices of others determines the amount of points that you earn each round. At the end of the session, we will randomly select one game from all of the games played in today's session. Your point earnings from the chosen game will be converted into dollars at a conversion rate of \$0.15 per point earned. Your earnings from the chosen game and your \$7 show-up payment will be paid to you in cash and in private. You will also have the opportunity to earn additional earnings in the second part of the experiment.

Specific Details

There are N individuals in today's session. At the start of each new game, each individual will be randomly assigned a position number for the game. This position number indicates the round in the game, $1, 2, \dots, N$, at which you will be called upon to make a choice.

When it is your turn to make a choice you will see the Choice screen (you will see a waiting screen until that time). On the Choice screen you will be asked to make a choice. Specifically, you can decide whether or not to draw a number from from 0 to 100, non-inclusive. The likelihood that you draw any particular number is fixed, but unknown to you. That is, the distribution of numbers you are drawing from could take the shape of ANY valid probability distribution, defined over the interval between 0 and 100, non-exclusive. All numbers drawn within this interval are limited to two decimal places. Thus, the smallest

possible number you could draw is 0.01 and the largest possible number you could draw is: 99.99.

Drawing a number from the distribution is costly. Specifically, a draw costs you 10 points. However, every individual is given an endowment of 10 points at the start of each new game, so the choice you face is whether to spend your endowment of 10 points drawing a number from the distribution.

Prior to making this choice, you are informed of the highest number that has been drawn by another participant in the current game. The choice you face is whether or not you want to try to draw a new number, at a cost to you of your 10 point endowment for the game.

If you choose not to draw a number, then your points for the game will be equal the highest number chosen in the game so far plus your 10 point endowment.

If you choose to draw a number, then your points for the game will depend on the number you draw. Specifically:

- If the number you draw is greater than the highest number drawn in the game so far, then your points for the game will equal the number that you draw minus your 10 point endowment (the cost of drawing a number).
- If the number you draw is less than or equal to the highest number drawn in the game so far, then your points for the game will equal that highest number drawn in the game so far minus your 10 point endowment (the cost of drawing a number).

Please note the following:

- First, draws from the unknown distribution are *with replacement* which means that the same number can be drawn more than once. The likelihood of drawing any number does NOT change across all games played in today's session.
- Second, if the highest number drawn in the current game is 0.00, then EITHER you are the first person to make a choice in the current game OR no prior participant has chosen to draw a number in the current game.
- Third, if you choose to draw a number, it is possible that the number you draw is higher or lower than the current highest number drawn by another participant in the current game. Remember, you don't know the likelihood of each of the possible numbers you could possibly draw; you only know that the unknown distribution of numbers is constant over time.
- Fourth, your earnings from drawing a number can be *higher* or *lower* than your earnings from NOT drawing a number. If you choose to draw a number and the number you draw is higher than the current highest number drawn, then your earnings from drawing that number will be higher than your earnings from not drawing a number only if the number you drew is at least 10 more than the current high number, since by drawing a number you lose your endowment of 10 points. On the other hand, if you choose to draw a number, your earnings will be lower than if you had not drawn a number whenever the number you drew is less than the current high number plus 10.

In making a choice of whether or not to draw a number for the current game, you face three options:

1. Do NOT draw
2. Enter a probability to draw
3. Draw

If you choose Do NOT draw, then you definitely DO NOT draw a number for the current game. If you choose Draw, then you definitely DO draw a number for the current game. If you choose Enter a probability to draw a number, then you must also enter a number (an integer) between 1 and 99 representing the probability, in percentage terms, that you will draw a number for the current game. After entering your percent chance of drawing a number, the computer program randomly draws a number (an integer) between 1 and 99 inclusive. If this randomly drawn number is less than or equal to your entered percent chance of drawing a number, then you will draw a number from the distribution for the current game: it will be automatically drawn for you; otherwise, you will not draw a number for the current game. Thus, the higher (lower) is the percent chance you enter for the probability of drawing a number, the more (less) likely it is that you draw a number in the current game.

Payment

The first part of today's sessions consists of 10 games. At the end of each game you will learn your points earned for the game, which, as explained above, depend on whether or not you draw a number. After all 10 games have been completed, one of the 10 games will be randomly chosen for payment. Each game has an equal chance of being chosen and so you will want to do your best in each game. Your points from the chosen game will be converted into dollars at the rate of 1 point = \$0.15 (15 cents). Thus, the more points you earn, the greater are your monetary earnings. In addition, you are guaranteed \$7 for showing up to today's experiment. You will be paid your show-up payment, together with your earnings from the first part of the experiment and your earnings from the second part of the experiment at the end of today's session. All payments will be made in cash and in private. At the end of this first part of the experiment, you will receive further instructions for how to complete part 2 of the experiment.

Questions?

Now is the time for questions. If you have a question about any aspect of these instructions, please raise your hand and an experimenter will answer your question.

Quiz

Before we start today's experiment we ask you to answer the following quiz questions that are intended to check your comprehension of the instructions. The numbers in these quiz questions are illustrative; the actual numbers in the experiment may be quite different. Before starting the experiment we will review each participant's answers. If there are any incorrect answers we will go over the relevant part of the instructions again.

1. Each game consists of _____ rounds. Your position in the game circle one: changes remains the same in each game.
2. The distribution of numbers you are drawing from: Circle one:
 - a. Changes over time.
 - b. Is unknown.
 - c. Is unknown, but all possible numbers lie between 0 and 100, non-inclusive, and the distribution does not change over time.
3. Before deciding whether to draw a number you can see the highest number drawn so far in the current game. Circle one: True False
4. If you choose NOT to draw a number, then your point earnings will be equal to: Circle one:
 - a. the highest number drawn earlier in the current game, (or 0.00 if no number has been drawn yet)
 - b. the highest number drawn earlier in the current game (or 0.00 if no number has been drawn yet) plus your 10 point endowment.
 - c. your 10 point endowment.
5. If you choose to draw a number, you lose your 10 point endowment for the current game. Circle one: True False Do you get a new endowment of 10 points for each new game? Circle one: Yes No.
6. Consider the following scenario. The current highest number is 65.56. You choose to draw a number which turns out to be 73.21 What is your payoff in points for the game in this case? _____ What would have been your payoff in points if you did not choose to draw a number? _____
7. Consider the following scenario. The current highest number is 56.71. You choose to draw a number, which turns out to be 31.03. What is your payoff in points this case? _____ What would have been your payoff in points if you did not choose to draw a number? _____
8. At the end of the experiment, one game will be randomly chosen. Your point earnings from that game will be converted into dollars at the conversion rate of 1 point = \$_____.

C Simulations of Behavior

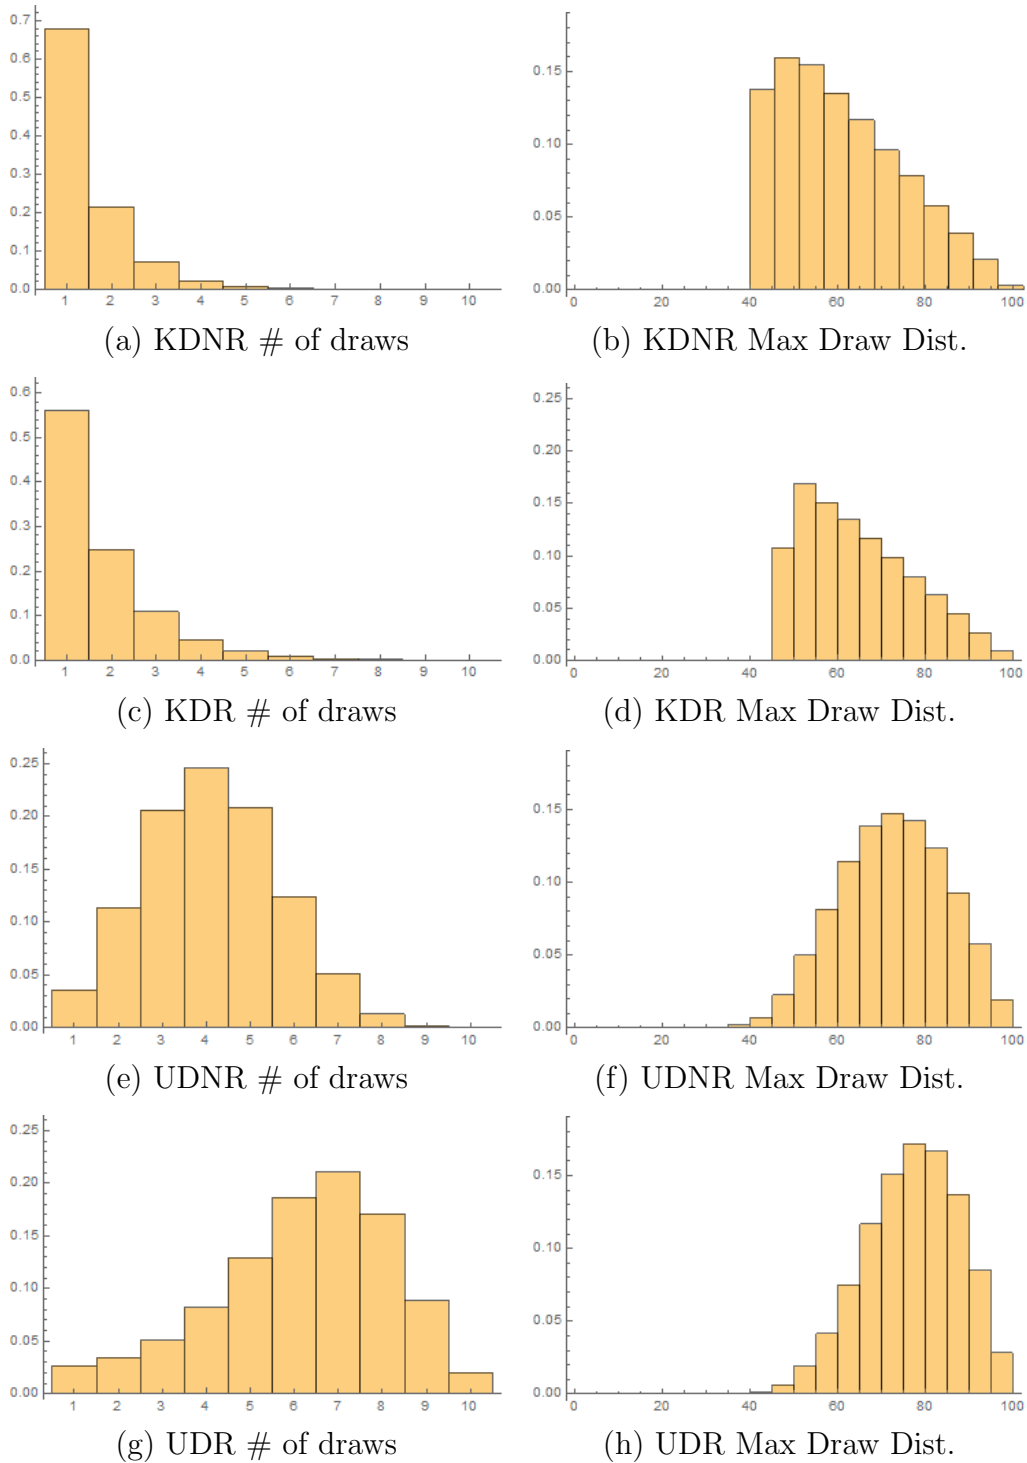


Figure 11: Simulations of behavior of # of draws and maximum draws in a 10 person game.