Innovate versus Imitate:
Theory and Experimental Evidence

John Duffy and Jason Ralston

September 8, 2018

Abstract

We model and experimentally evaluate the trade-off between innovation and imitation commonly faced by firms and individuals. Innovation involves searching for a high payoff opportunity, but paying a cost to do so. Imitation involves avoiding this 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 agents. We analyze the consequences of complete versus incomplete information about the distribution of payoffs from innovation on agent's decisions. In the complete information case, our predictions are based on expected utility maximizational while in the incomplete information case they are based on expected regret minimization. We then study these predictions in a laboratory experiment where we find evidence in support of our theoretical predictions.

Keywords: Innovation, Imitation, Risk, Ambiguity, Expected utility maximization, regret minimization, experimental economics

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

1 Introduction

The study of decision making under uncertainty is a mainstay of modern economic theory. In many contexts where agents or firms face uncertainty, they do so in an environment that is novel to them. The first firms to enter a market do so when the returns to their research and development are unclear. The first time agents 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 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 leading innovator or to innovate themselves and perhaps push the frontier forward. The standard procedure, as in all of economic decision making, is to weigh the risks and benefits of each action. However, there are many different types of uncertainty and incentive structures that can affect how agents make their decisions. For example, individuals might know the distribution of possible outcomes from innovation, and so behave as expected payoff maximizers in choosing between innovation or imitation. More realistically, they may not know the shape of the distribution from innovation (ambiguity) in which case they might act as regret minimizers. A further complication is whether innovation preclude imitation, or whether imitation is always possible independently of whether an agent chooses to innovate or not. If one can simply abandon failed innovations, those leading to payoffs that are worse than the current best innovation, this incentive structure will also 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 in four different settings. Specifically we analyze the effects of 1) the ability to imitate following unsuccessful innovation, 2) the absence of distributional information on innovation outcomes and 3) the interaction of these two features. We compare and contrast these three environments to the baseline case where the distribution of payoffs is known, but where agents are unable to imitate (or "recall") in the event of unsuccessful innovation. We characterize agents as regret minimizing, which degenerates into utility maximizing when the distribution of innovation opportunity is known.

Our theory predicts that agents will behave differently when incentive and informational structures vary. Specifically, agents are predicted to play probabilistically in situations where they face ambiguity regarding payoff distributions, but deterministically when the distribution from innovation is known, and the expected maximum draws 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). We test the predictions of our model using an experiment featuring a strategy elicitation method in the context of a sequential innovate versus imitate game. Our prediction is that, when playing this game with 10 players and using a triangular distribution with a support spanning 0 to 100 and a mode at 50, that the expected values at which the players will stop innovating, in order from lowest to highest, are 1) known distribution without recall, 2) unknown distribution without recall, 3) unknown distribution with recall, 4) known distribution with recall.

Our experimental results reveal that subjects' that subjects behave close to the predictions of the model in the case where the distribution is unknown, and that they deviate furthest when the distribution is known. Specifically, the threshold for innovation is higher than predicted in cases where the distributions are known than when they are unknown. The possibility of recall increases the threshold in settings where the distribution is unknown. We analyze individual level correlates of willingness to draw from the payoff distribution and find that males are more likely to draw (innovate) than females and that those with higher grade point averages tend to be less willing to draw, though risk preference plays no role, even when we look at treatments where subjects would be believed to exhibit the most risk averse attitudes.

This study is connected to theoretical work in several areas: evolutionary innovation versus imitative behavior, ambiguity aversion, and strategic experimentation. We often seen innovative and imitative behavior featured in the theoretical literature in the context of research and development and changes in productivity (König et al. 2016), developing new products (Ofek & Turut 2008), and in the composition of industrial structure (Iwai 2000). In these models, firms are typically considering a continuous strategy space, instead of a binary strategy space as we do. Firms in these models typically also have beliefs about the payoff distribution, or how the stochastic elements are distributed, allowing for utility maximization. Our model features utility maximization as a special condition of regret minimization, which occurs only when ambiguity over a distribution is resolved.

Our study 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 is 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 (1738). 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.[1] Decision making under ambiguity has been formalized in models of maximin expected utility (Gilboa & Schmeidler 2004), Choquet expected utility (Schmeidler & Gilboa 1994), and models that allow violations of the reduction of compound lotteries axiom (Halevy 2007). We model ambiguity averse individuals as regret minimizers, or agents who apply minimax strategies, similar to the work of Bergemann & Schlag 2011, Renou & Schlag 2010) where agents encounter uncertainty about the distribution of stochastic elements and seek to minimize their regret given that the state of the world that is the *least* favorable. Again, our study focuses on a simplified version of these models, where agents do not need to consider a continuum of strategies. This results in a simple minimax mixed strategy prediction when the distribution over payoff is ambiguous and a unique threshold prediction when that ambiguity is resolved.

Lastly, our model is related to the literature on strategic experimentation. This literature deals with sets of agents facing double-armed bandit problems where agents must decide to divide their time between a safe option and a risky option with a true payoff plus noise (Bolton & 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 the unknown distribution setting have limits on their information sets that are not present in the standard strategic experimentation framework. Namely, our agents are only able to see the current maximum payoff value obtained, and the game is one shot, so no learning is possible. 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.

[1] Uncertainty is now used as an umbrella term to describe both risk and ambiguity.

Our project takes elements from each of these different strains of literature and combines them to ask the question of how an agent or firm might approach the task of taking risky strategies where the risk can be either known or unknown within the context of a single shot innovate or imitate game. Central to our study is attempting to decipher if subjects will behave as ambiguity averse individuals are modeled in the regret minimization literature and how the ability to copy the best outcome of another individual (imitation) might increase the likelihood of an agent obtaining the maximum payoff possible in a game. The paper is organized as follows: the section 2 covers the theory of the game, section 3 covers the exact experimental design of the game, section 4 summarizes the hypotheses our model makes of the behavior in the context of our experiment, section 5 analyzes the data collected and provides evidence in favor and against the hypotheses made, and section 6 concludes with some thoughts on what was learned.

2 Theory: Cardinal Reward

Consider the following sequential move game. Player i must choose between a costly lottery or simply accepting the highest lottery payoff realization achieved by players who have moved earlier in the game. We interpret the former choice as "innovation" and the latter as "imitation". Let us denote the outcome of the lottery faced by player i in period t by $x_{i,t} \sim \text{dist}(\theta)$. Then we can also denote the full prior history of lotteries up to period t by $X_t = \{x_1, \ldots, x_{t-1}\}$ and the maximum of this set as $x_{\max} = \max\{X_t\}$. We further assume that the cost to choosing the lottery 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:

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

We refer to this payoff function as the "recall" version in that if innovation is unsuccessful, player i can always recall (receive) the imitation payoff $x_{\max}$, but pays the cost of innovating c. Below, we will consider the case where such recall is not allowed.

The decision to participate in a lottery in period t depends on the expected value of the lottery, conditional on the draw being larger than the current maximum of lottery outcomes. Specifically, the expected increase in the outcome must be worth rejecting the current maximum of lottery draws. Mathematically,

$$E\left(x_{i,t}|x_{i,t} \geq x_{max}\right) - x_{max} \geq c. \tag{2}$$

The above expression states that the expected increase in the payoff from an additional lottery being accepted must be larger than the cost associated with taking that additional lottery.

Proposition 1. *In the absence of imitation costs, there will be some $k \in I$ such that all $i < k$ will innovate and all $j \geq k$ will imitate.*

Proof. Under the assumption that agent's play rationally, players will continue to choose to pay for lotteries (innovate) until $E\left(x_{i,t}|x_{i,t} \geq max\{x_{i,max}, x_{max}\}\right) - max\{x_{i,max}, x_{max}\} \geq c$ is no longer true. Assuming $x_i \sim N(\mu, \sigma)$, then the expression $E\left(x_{i,t}|x_{i,t} \geq x_{max}\}\right) - x_{max}$ will be decreasing in x_{max} and approaches 0 and x_{max} approaches infinity.

By the intermediate value theorem, there will exist an x_{max} such that $E\left(x_{i,t}|x_{i,t} \geq x_{max}\}\right) - x_{max} = c$, for a c small enough. Thus there will be some $j \in I$ for whom $E\left(x_{i,t}|x_{i,t} \geq x_{max}\}\right) - x_{max} = c$ or $E\left(x_{i,t}|x_{i,t} \geq x_{max}\}\right) - x_{max} < c$. The j^{th} agent will be either indifferent between imitating and innovating or strictly prefer imitating. 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 $i > k$ $E\left(x_{i,t}|x_{i,t} \geq x_{max}\}\right) - x_{max} < c$, and these agents will all imitate. □

As noted, the payoff function (1) allows an innovator to imitate the current best payoff $x_{\max}$ less the innovation cost c in the event that his/her innovation is unsuccessful ($x_{i,t} < x_{\max}$). We can relax this ability to recall the maximum prior payoff and consider payoff functions without such recall:

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

which states that the agent faces the decision of either taking a lottery $x_i \sim \text{dist}(\theta)$ or taking the highest of the previous innovations. With this framing, the decision calculus is somewhat changed. The agents should choose to innovate whenever $E(x_i) \geq x_{max}$, and imitate otherwise. What makes this second setup different from the first is that since the agents no longer are guaranteed the maximum, they will not condition their expectations on x_i being greater than x_{max}. The model makes the following prediction:

Proposition 2. *For some $k \in I$, k will be the first agent to imitate and for all $i > k$ will imitate. Furthermore, the $k-1^{th}$ or k^{th} will set $x_{max} = E(x_i) = \mu$ for all future agents.*

The proof of Proposition 2 follows a logic similar to 1 and is omitted.

In the real world, there are copyrights and fines for copyright violation. We can easily accommodate this feature by requiring the agent to pay a cost, d, if they choose to imitate. This new cost changes the decision criteria in (2) to:

$$E\left(x_{i,t} | x_{i,t} \geq x_{max}\right) - x_{max} \geq c - d. \tag{4}$$

As expected, adding a cost to imitation decreases the number of agents who play imitate and increases k, *ceteris paribus*. Without loss of generality, we will focus on a single cost, c, for innovation rather than d, the cost for imitation, though one can also think of our innovation cost as the *net* cost of the two actions, i.e., $(c - d)$.

2.1 Unknown Distribution

There are many cases in which an agent or firm does not know the distribution of possible payoffs from innovation. In such a setting, we can no longer appeal to expected utility maximization, since the relevant distributional information is not fully available. We do suppose that agents know *something* about innovation prospects. Specifically, we suppose that the support, $[a, b]$, of the unknown distribution from innovation is perfectly (and commonly known), a setting that corresponds to ambiguity as discussed earlier.[2] In such a setting, we conjecture that agents adopt a minimax regret strategy instead of expected utility maximization.

Let $r(x_{max}, y, j)$ be the regret resulting from action j with the payoffs from imitating being x_{max} and innovating being the draw y. The regret is then a piecewise defined function:

[2]That is, we do not consider the case of complete ignorance!

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

Proposition 3. *In the case where the agent is 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.

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

It is assumed that $b - c > x_{max} \geq a$ so that imitation does not dominate. We examine the two degenerate distributions which maximize regret, $F = \delta_b$ and $F = \delta_a$ (when 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 can be found to be

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

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

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

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

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

This function reaches a minimum at p^*.

□

Using a similar method we can find a solution to the regret minimization problem faced by agents with the ability to recall. 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 payoff under $F = \delta_a$ will be $(1 - p)c$

Proposition 4. *In the case where the agent is allowed to recall, agents 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 **??**, noting the change mentioned about the case where $F = \delta_a$.

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

3 Experimental Design

The model makes distinct predictions about stopping rules and how they differ depending on whether the distribution is known as well as upon 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 **??** provides a summary.

Four treatments of the experiment

	No Recall	Recall
Known Distribution	KDNR	KDR
Unknown Distribution	UDNR	UDR

In each of these four treatments, subjects participated in 4 different stages: the main decision consisting of 1) ten, 10-period, 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 stage

In the first and main stage, subjects made a decision about how likely they were to take a draw from a payoff distribution (known or unknown) versus how likely they were to copy the highest payoff of all previous payoffs earned previous subjects. Subjects were presented with some information about the distribution of payoffs as well as information about their payoff function. They were then asked to make a single choice, how likely they were to take a draw from the payoff distribution. This likelihood also determined the probability that they did not take a draw, and instead copied the highest payoff received by previous subjects. Options for the likelihood to take a draw included two buttons which indicated either a 0% or 100% chance to take a draw as well as a field in which the subject could choose a probability between the two extremes.

Subjects were endowed with 10 points and were told that and were informed that taking a draw would cost them their endowment. If the did not take a draw, then the subject would keep their endowment. In sessions where recall was available, when the subject took a draw and it was below the current maximum, their draw was replaced by the current maximum.

Subjects were also given some information about the payoff distribution. In the full information condition, subjects were shown a graph of the distribution they could draw from, importantly featuring the finite range of the support for the triangular distribution used as well as the modal peak and its value. In the incomplete information treatment, only a picture of the support was shown and no other information that the subject could use to characterize the shape of the distribution. In both treatments a line was placed that showed the subject what the current maximum draw was.

After subjects submitted the likelihood at which they would like to take a draw, the computer drew a number from 0 to 100 with uniform probability. If the number was less than or equal to the submitted likelihood, then the computer drew a value from the payoff distribution and calculated payoffs in the manner specified on the decision screen and displayed the results on the page that followed directly after the decision screen.

Subjects took turns in making these decisions. At the beginning of a session, each subject received a number that denoted the order in which they in which they would make their decision. When the program advanced to the round number that matched their order, that subject was asked to make their decision.

Every session consisted of 10 subjects, and each subject made only one decision in the decision stage. Therefore one decision stage lasted 10 rounds. The subjects played 10 decision stages and one was picked at random for payment.

3.2 Risk Elicitation Stage

After 10 decision stages were played, each subject advanced to the risk elicitation stage. In the risk elicitation stage, subjects were presented with 3 gambles. The subject was told to pick one of the gambles and the computer would randomly determine a payoff, conditional on their choice.

3.3 Cognitive Reflection Task and Survey Stages

After completion of the risk elicitation stage, subjects proceeded to the survey stage. In this stage, subjects first answer three cognitive reflection questions and then proceeded to answer demographic questions. These questions covered nationality, ethnicity, age, major, and GPA.

Following the survey stage, subjects were informed of their experimental earnings, risk elicitation earnings, the show-up payment, and the 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 task, we chose one game randomly from all games played and converted subjects point earnings in that game into money earnings at a fixed and known rate 1 point = $0.15 USD. For the second and third tasks subjects earned additional points ....explain details. The total average payment was ˜$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 1.

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 1: Descriptive statistics regarding subject population

4 Hypotheses

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

First, we hypothesize that subjects will behave in accordance with the theory prescribed. That is, we believe that subject's propensity to draw will match the deterministic of probabilistic draw rules specified in the theory above.

Hypothesis 1. *The probabilities of drawing from the distribution will match those predicted by the theory.*

Second, in the case where we can use a threshold analysis, we can find the threshold value where a risk-neutral agent is indifferent between drawing and not. These values can be found analytically, solving for the value of x_{max} where $E(Draw) = E(NotDraw)$. These values are shown in the Table 2 below.

		Recall	No Recall
	Threshold	40.00	82.10
Known Dist.	*Exp. Max*	60.90	79.59
	Exp. # Draws	1.47	7.32
	Threshold	x	x
Unknown Dist.	*Exp. Max*	69.33	77.00
	Exp. # Draws	4.13	6.16

Table 2: Thresholds and expected maxima for each treatment

Hypothesis 2. *The expected stopping value for each of the treatments will follow the predictions made in Table 1.*

We can decompose this hypothesis into two sub-hypotheses. Subjects will have thresholds of indifference predicted by Table 1. Along these lines, we can also make the prediction that the relation of the thresholds will be KDNR $<$ KDR.

Hypothesis 3. *The distribution of maximum draws will follow the predictions made in Table 1.*

Again, we may divide this hypothesis into two sub-hypotheses. We expect the distribution of maximum draws in a game to match those generated by the behavior of agents in the model. From these distributions, we can also make the prediction that the order of the maximum draws in a treatment will follow the ordering KDNR $<$ UDNR $<$ UDR $<$ KDR.

Hypothesis 4. *The distribution of number of draws will follow the predictions made in Table 1.*

This hypothesis follows the decomposition of the preceding two. Subjects should play close to what is predicted by theory and the order of the estimates should follow the pattern found in Table 1.

Hypothesis 5. *Deviations from risk-neutral play are correlated with individual risk preferences or other personal attributes*

We note that our theory, in the case of known distributions, makes use of expected utility maximization. However, we note that empirically many subjects are not risk-neutral, and so we predict that any deviations away from risk-neutral preferences might influence a subject's propensity to draw or innovate. Specifically, those who are risk-averse might stop before our predicted thresholds or exhibit lower probabilities of drawing. Conversely, a risk-loving subject might stop after our predicted thresholds or exhibit high probabilities of drawing. To the extent that our subjects are risk-neutral, deviations from the theory should be minimal. Other attributes that might explain departures from risk neutral predictions include cognitive ability, which we measure using GPA and CRT scores.

5 Experimental Results

The theory makes predictions about a few main measures. Each measure we look at will be related to the theory's predictions on subjects' deterministic or probabilistic propensity to innovate given the current current maximum reached in a game. First, we look at individual level predictions versus behavior across the four different conditions our subjects could find themselves in. These analyses make significant use of deviations from theory. Following this, we begin to look at comparisons between behavior within a game (10 subjects playing for 10 periods) and compare it to numerical results we generate from large simulations of agents all playing the exact strategies dictated by our theory. Here we make use of expected maximum draw within a game, the number of draws expected within a game, and indifference thresholds (in the case where distributional information is present). When possible we compare these measures across the different treatments as well as with the predicted values for each treatment.

5.1 Deviation analysis

We analyze predicted behavior by examining the root mean squared error between the probabilities of drawing generated by our model and actual subject decisions. From the data we have subject decisions regarding their desired probability of taking a draw from the distribution and the theory makes predictions about the probabilities a risk neutral agent would make when confronted with different levels of x_{max}, the current maximum value drawn in a game. We compute the squared deviations 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 and $\hat{p}_{i_t}$ is the predicted probability of risk-neutral agent drawing. This is reported in Figure 1. As was suggested by the preceding threshold analyses, subjects follow the predictions made by the model most closely when recall is not a factor in the decision making process. That is to say that the squared errors of the no recall treatments are significantly larger than those of the recall treatments ($p < 0.01$). Figure 1 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.

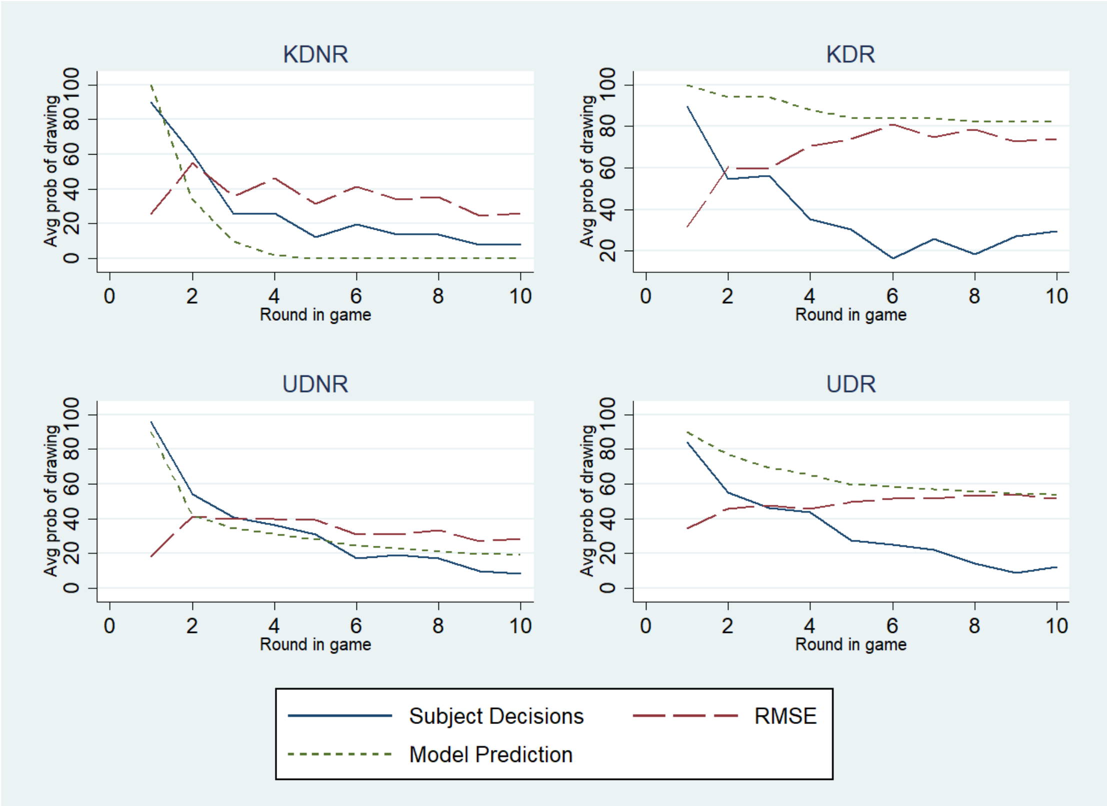

Figure 1: Average deviations from regret minimizing model

We further decompose our sample into early rounds versus later rounds in order to see if subjects are updating their beliefs about the payoff distribution from their initial, regret minimizing beliefs, conjectured to be most prevalent in the early periods before learning can take full force. We define early rounds as the first five rounds, when there are few observations for a subject to condition on to form accurate beliefs about the payoff distribution. Figure 2 demonstrates that, though differences are minor in most cases, that subjects tend to follow the model predictions better over time as more learning can occur. Thus, it seems that it may take many repetitions of the game for subjects to learn the payoff distribution and behave in a manner inconsistent with regret minimization. It also supports the notion that subjects fail to take advantage of the recall condition even after many games have been played.

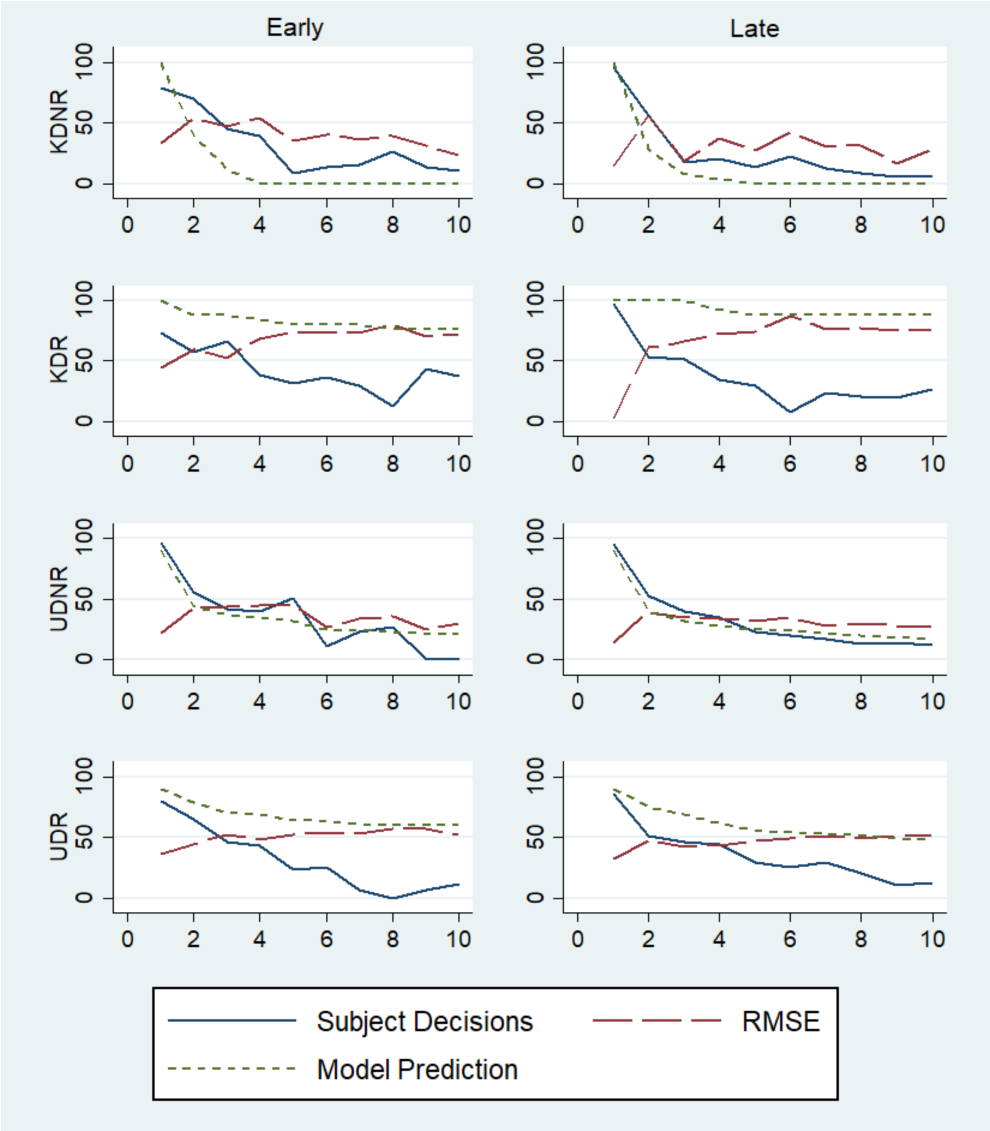

Figure 2: Average deviations from regret minimizing model early and late into sessions

We explore the relationship between deviations and individual characteristics to investigate any correlations that might serve to explain systematic differences in deviations between the four different treatments. Specifically, we examine whether deviations are systematically different between those who are risk neutral and who are not. Using a Kolmogorov-Smirnov test to test for differences in distributions, we find a statistically significant difference ($p < 0.01$).

Next we test for relationships between deviations and individual level characteristics, including age, sex, GPA, cognitive reflection test (CRT) score, and quantitative reasoning (QR) score. We impute a QR score by associating the mean QR score for a major on the GRE test with the major of each subject. For undeclared majors, we used the mean QR score across all test takers, which was 159. The least squares estimates are reported in

Table 3.

	(1) All Treatments	(2) KDNR	(3) KDR	(4) UDNR	(5) UDR
Unknown Dist	-187.5	x	x	x	x
	(194.2)				
Recall	3,430***	x	x	x	x
	(291.4)				
Unknown Dist x Recall	992.2***	x	x	x	x
	(201.0)				
Risk Averse	166.5	468.8	142.4	-265.3	-94.06
	(320.1)	(611.9)	(712.6)	(240.0)	(349.6)
Risk Loving	-42.92	211.8	-144.9	-198.7	-89.79
	(193.3)	(274.0)	(542.0)	(230.2)	(258.1)
Age	-24.99	-14.31	-183.6	84.70	130.6*
	(56.27)	(52.59)	(134.0)	(84.74)	(75.91)
Sex	361.2*	175.1	1,089*	18.97	154.9
	(211.6)	(297.6)	(596.0)	(223.2)	(323.4)
GPA	-186.9**	-457.0**	-230.6	-102.1	-60.85
	(91.04)	(204.4)	(336.8)	(82.88)	(77.65)
CRTscore	15.36	-330.2**	204.8	44.75	-35.35
	(75.65)	(125.1)	(218.4)	(100.8)	(109.2)
QRscore	-2.050	-45.46	5.951	-13.17	-23.48
	(19.75)	(35.39)	(71.82)	(13.35)	(28.51)
Constant	2,618	10,423*	7,561	1,999	3,635
	(2,948)	(5,479)	(12,220)	(2,406)	(4,126)
Observations	2,000	500	500	500	500
R-squared	0.177	0.037	0.021	0.013	0.011

Robust standard errors clustered at the subject level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3: Determinants of deviation from model predictions

Table 3 shows that, though there is no significant correlation between risk preference and deviations after controlling for confounding factors, sex and GPA of the subject make a difference in the degree to which subjects' deviate from the model's predictions. Specifically, men deviate from the model more than women and those with higher GPAs deviate from the model less than those with lower GPAs. These effects are felt most strongly in the KDNR and KDR treatments. Table 6 also shows, again, that KDNR and UDNR have similar deviations, while KDR and UDR create significant differences in deviations. This supports the results of the threshold analysis.

5.2 Threshold Analysis

5.2.1 Known Distribution

Before preceding with threshold estimation for the cases where such analyses are appropriate (i.e. those scenarios where distributional information is known), we present our estimation strategy. For each subject i, we estimate a threshold indifference point using logit regressions of the form $\text{Innovation}_i = \alpha + \beta \text{CurrentMax}_i + \varepsilon_i$, where Innovation_i is an indicator variable for whether a subject attempted to innovate (draw=1) or not (imitate=0). After estimating this equation, we then take the ratio of the estimated values $\frac{-\hat{\alpha}}{\hat{\beta}}$, which indicate the threshold of indifference between drawing and not drawing. Errors are clustered at the subject level. We estimate thresholds at both the game and treatment level.

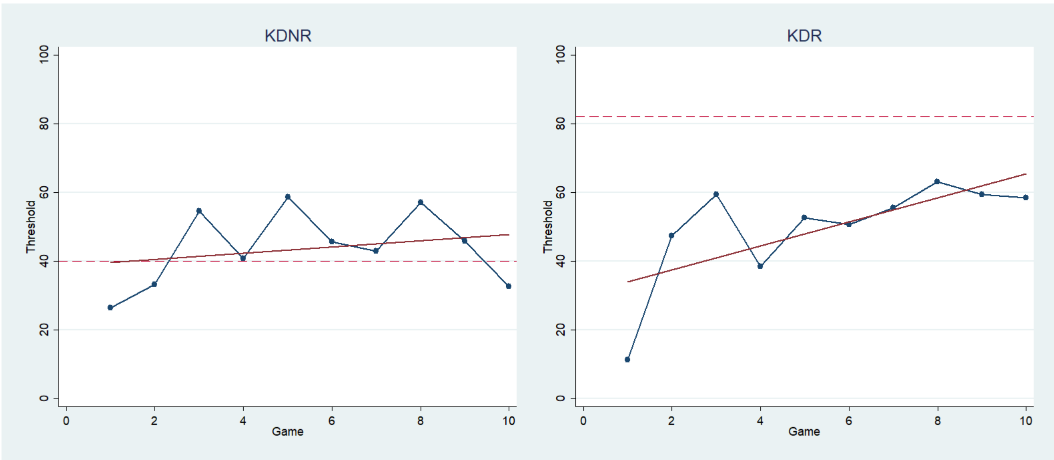

Figure 3: Estimated threshold in games by treatment

Figure 3 shows the average of the estimated thresholds in every game for KDNR and KDR treatments and time trends. Figure 3 supports the notion that subjects behavior is consistent with the model when they are in the no recall condition. In the KDR treatment, where recall is enforced, subjects perform well below how the are predicted to, under-leveraging the benefits of recall. However, it is worth noting that in the KDR treatment, the trend is positive, and that subjects in the KDR treatment had marginally significantly higher estimated indifference points in the last five games than in the first five games ($p = .10$), indicating that the subjects may have been adjusting their behavior.

		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)
	Prediction	40.00	40.00	40.00	82.10	82.10	82.10
Known Dist.	Average	42.54	43.10	42.76	44.13	58.70	51.55
	(s.d)	(3.47)	(3.49)	(2.73)	(4.61)	(1.81)	(2.75)

Table 4: Thresholds by treatment

Table 4 shows the estimated thresholds at which subjects are predicted to be indifferent between drawing and not drawing in each treatment and their associated standard errors. We 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 stated above.

Table 4 shows that KDNR is statistically significantly different from KDR, which is predicted by the model. These findings indicate that our model of innovate versus imitate provides reasonable predictions about the rank *order* of the estimated thresholds. However, we note that the estimates of the thresholds is statistically significantly different from the predicted thresholds in the KDR case.

5.2.2 Unknown Distribution

When analyzing decisions under unknown distributional information, we can no longer use the logit analysis, 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, the analysis we find appropriate is the analysis between x_{max} and the associated probability of taking a draw. This analysis is shown in Figure 4 below. We interpret these probabilities to draw as the average p which minimizes regret subject to each subject's subjective payoff distribution.

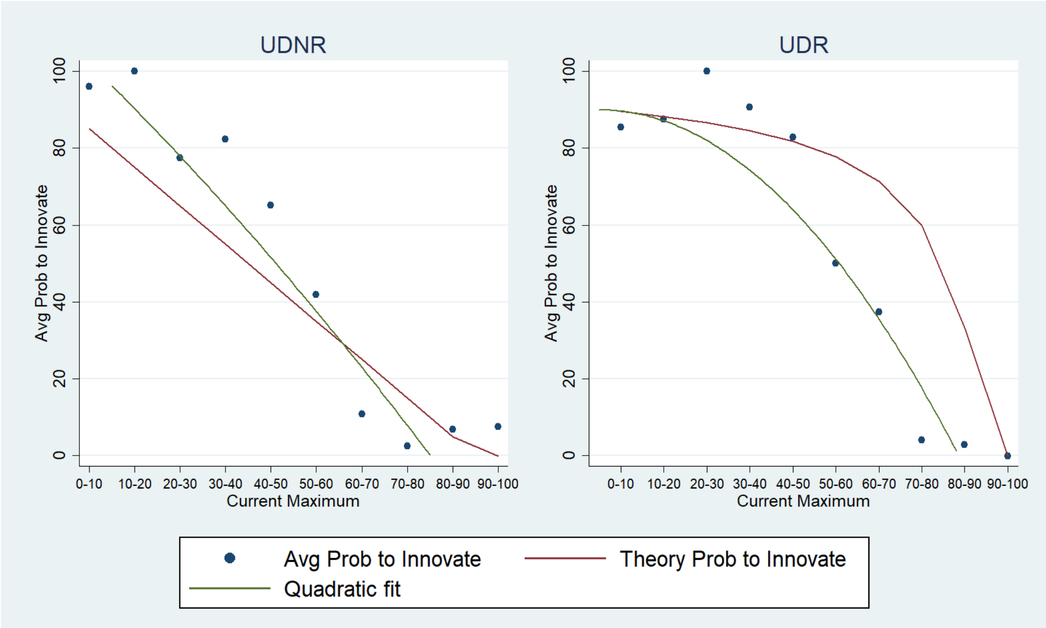

Figure 4: Average maximum draws in games by treatment

This analysis also shows the predicted regret minimizing probabilities as well as a quadratic trend for the data. Using t-tests, we find that there is no significant difference in probabilities to draw predicted by the model and the experimental data ($p = 0.79$) in the no recall condition, but that the difference is large in the treatment with recall ($p < 0.01$). As in the case of the known distributions, we find support for our model when there is no recall condition, but that subjects have difficulty incorporating the incentive structure of the recall condition into their decisions.

5.3 Expected Maxima Analysis

We next investigate subject behavior via the expected maxima, taking the maximum draw within a game in each treatment as our measure of interest. Any predicted value arising from the model is found numerically, using 100,000 simulations of games analogous to those found in the experiment. The 10 agents in these simulated games all play strategies in accordance with the probabilistic innovation strategies predicted by our model.

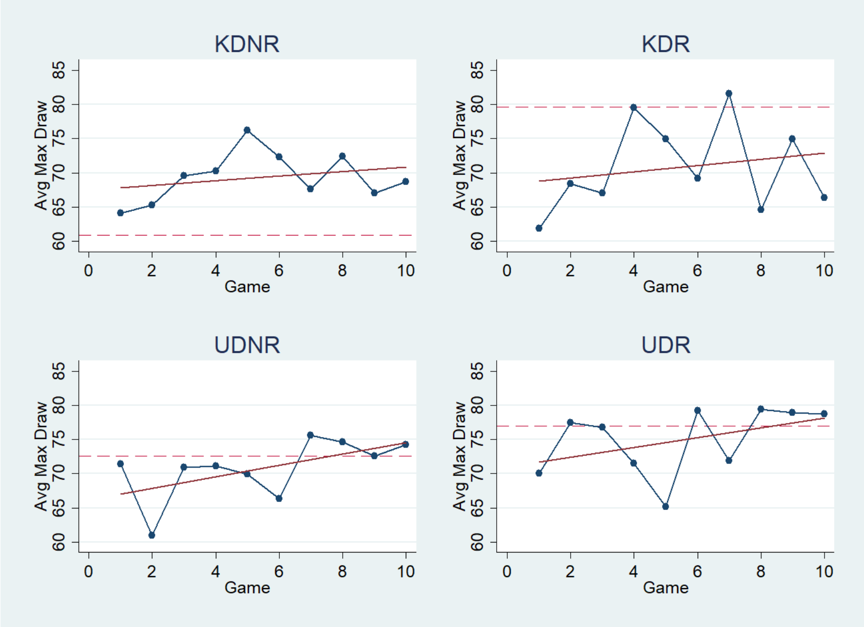

Figure 5: Average maximum draws in games by treatment

Figure 5 shows that, while the maxima within a game increased within treatments with unknown distributions, it remained mostly flat in treatments with known distributions. Table 5 reports the results of pairwise comparisons of maximum draws in a game between treatments using t-tests, Mann-Whitney U-tests, and Kolmogorov-Smirnov tests. We also compare how our close our subjects perform to what is predicted by our models. This is shown in Table 6.

		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	79.59	79.59	79.59
	Average	69.07	69.59	69.33	70.35	71.32	70.83
	(p-value)	(<0.01)	(<0.01)	(<0.01)	(<0.01)	(<0.01)	(<0.01)
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
	(p-value)	(0.12)	(0.974)	(0.282)	(0.02)	(0.767)	(0.164)

Table 6: Predicted versus real maximum values by treatment

Table 6 shows that subjects behave remarkably similar to what is predicted by the model, but only in the cases where the distribution is unknown. Columns (2) and (5) of Table 6 also show that the average maximum draw is increased by the ability to recall, but only in the case where the distribution is unknown. We also note that the average maximum *increases* between the first first five and last five games in a session in all but one of the four treatment conditions, indicating that learning plays some role in behavior.

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

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

5.4 Expected Draws Analysis

We take advantage of the fact that the expected maxima simulation also provides estimates of the expected number of times subjects will draw during a 10 period game. Figure 6 shows the average number of draws within a game by treatment and their respective trends. In no treatment are the trends significant and the first and last five games of a treatment are never significantly different.

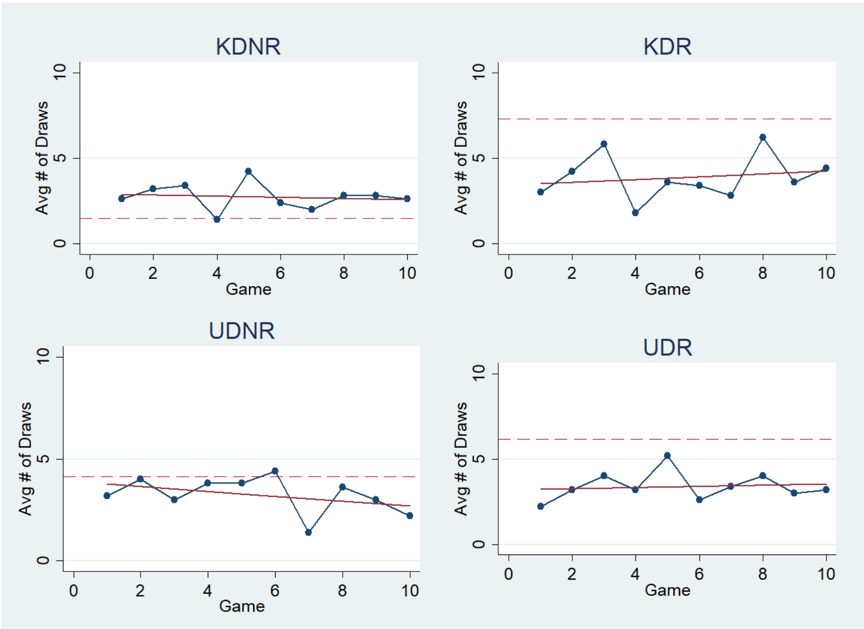

Figure 6: Average number of draws in games by treatment

Table 6 presents the mean number of draws from the simulations along with the average and standard deviation from the experimental data. Table 8 gives the results of all pairwise comparisons between the different treatments using a variety of statistical tests. The only statistically or marginally significant differences exist between KDNR/KDR and KDNR/UDR. In both cases KDNR is lower, which is in live 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	7.32	7.32	7.32
	Average	2.96	2.52	2.74	3.68	4.08	3.88
	(s.d.)	(1.80)	(1.10)	(1.51)	(2.19)	(2.08)	(2.14)
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	3.4
	(s.d.)	(1.77)	(1.63)	(1.73)	(1.63)	(2.14)	(1.91)

Table 7: Predicted versus real number of draws by treatment

However, in no treatment condition does the model accurately predict the number of draws. In every case, using t-test, U-tests, and KS-tests, we can reject the null hypothesis that the number of draws in a simulated game is the same as in the experimental data at the 0.01 level. However, we do find that deviations from the predictions are significantly higher when subjects are in a condition where recall is present. This is in line with the finding from the deviation analysis, suggesting that subjects are under utilizing the benefits of recall.

5.5 Findings

We use the analysis of the preceding section to draw the following findings:

Finding 1. *We find evidence that subjects follow reasonably close to the model in cases when there is no recall condition.*

From our deviation analysis, we find evidence that suggests that subjects deviate much further from the model predictions when they are under the recall condition. We find corroborating evidence in the threshold and expected draws analysis.

Finding 2. *We find support for Hypothesis 2. In general, subjects follow the predicted ordering of thresholds. However, subjects do not exhibit the indifference thresholds implied by the model.*

We find that the order of the thresholds is KDNR < KDR. While there is no order of statistics for the analysis we use for the case of unknown distribution, we can say that the model does a good job at predicting mixing probabilities in the case where there is no recall opportunity. This mirrors the results we find in the known distribution case.

Finding 3. *We find support for Hypothesis 3. We find that the distributions of maximum draws in the UDNR and UDR treatments are statistically insignificantly different from those produced by the theory. We find that subjects do not follow the pattern of predicted expected maxima.*

We find that the subjects' behavior matches very closely to what is predicted for the UDNR and UDR treatments. Interestingly, we can reject that the behavior of our subjects match those predicted by our models in the KDNR and KDR treatments – the

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

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

average is higher than predicted in the KDNR treatment and lower than predicted in the KDR treatment.

What is more, in the UDNR treatment, subjects follow the probabilistic predictions of the regret minimization model that gives rise to the expected maximum draw predictions. Subjects follow these predictions closely even after playing multiple games, indicating that regret minimization can continue to play a role in behavior even after payoffs have been realized several times and learning can occur. We also find that deviations from predicted behavior are statistically significantly different over only the known/unknown distribution treatments: that is to say there is no statistically significant difference in deviations between KDNR and KDR or UDNR and UDR ($p = 0.711$, $p = 0.544$ respectively). The remaining pairwise comparisons all show high levels of statistical significance ($p < 0.01$).

Finding 4. *We find no support for Hypothesis 4. The subjects do not follow the predictions of the model in terms of the order of the predictions or matching the distributions generated by the theory.*

Though we do not find any evidence for Hypothesis 4, our analysis of expected number of draws reveals a similar effect that we found in the deviations analysis: recall increases deviations from the model's predictions.

Finding 5. *We find that risk preferences do not influence deviations from the risk-neutral predictions. There is a positive correlation between GPA and deviations. There are marginal gender effects in the aggregate.*

Table 4 shows that there is no correlation in the subjects' willingness to draw and their risk preference. This results is the same when we consider either the total sample or break down the sample by treatment. Individual characteristics that do seem to influence these decisions are GPA and sex. Higher GPAs negatively correlate to the subjects' willingness to draw and males tend to be more likely to draw than females, all else held equal.

6 Concluding Remarks

Our study 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 agents as regret minimizing. We predict four distinct behaviors which are influenced by the ability to copy the leader in the case of failed innovation (recall) and knowledge of the payoff distribution. These different conditions mirror aspects of the real world, such as the presence of intellectual property rights and the strength of beliefs on the returns to research and development. Our model predicts differences in the probabilities of a regret minimizing individual to innovate conditional on the current maximum draw, which also influences the expected maximum draw to be obtained within a game.

We develop a novel experimental design to test the implications of the model. To our knowledge, this is the first to measure how well subjects behavior corresponds to the predictions of regret minimization where the states of nature form a continuum. We find that our subjects behavior is consistent with that predicted by regret minimization. We find that the biggest driver of differences between the average maximum draw in a game, as well as deviations from expected individual level behavior, is knowledge of the payoff distribution, while recall plays a less central role in explaining differences.

Further, when we compare average propensities to innovate, we find that regret minimization 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. What is more, regret minimization seems to describe 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 its 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, recall does not help in exploring the payoff distribution.

We have plans to modify our design to explore more realistic scenarios. Since one innovation often serves as a complement to new innovations (Romer 1994), it would make sense to explore the results if one were to make the payoff distribution endogenous by allowing successful innovation to shift the parameters of the distribution. Depending on the differences on 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 the changes to the distribution as ambiguous to capture behavior when the changes in the state of nature are not certain as well. Another reasonable modification would be to let the cycle that subjects make decisions repeat, 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. We anticipate that this strand of experimental research will be fruitful.

References

Bergemann, Dirk, and Karl Schlag. “Robust monopoly pricing.” Journal of Economic Theory 146.6 (2011): 2527-2543.

Bolton, Patrick, and Christopher Harris. “Strategic experimentation.” Econometrica 67.2 (1999): 349-374.

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, vol 9: 88-97

Ellsberg, Daniel. “Risk, Ambiguity and the Savage Axioms.” Quarterly Journal of Economics, 75.4 (1961): 643-669.

Iwai, Katsuhito. “Schumpeterian dynamics: An evolutionary model of innovation and imitation.” Journal of Economic Behavior & Organization 5.2 (1984): 159-190.

Iwai, Katsuhito. “A contribution to the evolutionary theory of innovation, imitation and growth.” Journal of Economic Behavior & Organization 43.2 (2000): 167-198.

Halevy, Yoram. “Ellsberg revisited: An experimental study.” Econometrica 75.2 (2007): 503-536.

Hume, David. *A Treatise of Human Nature*, Oxford: Oxford University Press, 1738.

Keller, Godfrey, Sven Rady, and Martin Cripps. “Strategic experimentation with exponential bandits.” Econometrica 73.1 (2005): 39-68.

König, Michael D., Jan Lorenz, and Fabrizio Zilibotti. “Innovation vs. imitation and the evolution of productivity distributions.” Theoretical Economics 11.3 (2016): 1053-1102.

Knight, Frank H. *Risk, Uncertainty and Profit* Boston: Houghton Mifflin Company, 1921.

Ofek, Elie, and Ozge Turut. “To innovate or imitate? Entry strategy and the role of market research.” Journal of Marketing Research 45.5 (2008): 575-592.

Romer, Paul M. “The origins of endogenous growth.” Journal of Economic perspectives 8.1 (1994): 3-22.

Gilboa, Itzhak, and David Schmeidler. “Additive representations of non-additive measures and the Choquet integral.” Annals of Operations Research 52.1 (1994): 43-65.

Renou, Ludovic, and Karl H. Schlag. “Minimax regret and strategic uncertainty.” Journal of Economic Theory 145.1 (2010): 264-286.

Renou, Ludovic, and Karl H. Schlag. “Implementation in minimax regret equilibrium.” Games and Economic Behavior 71.2 (2011): 527-533.

Schmeidler, David, and Itzhak Gilboa. “Maxmin expected utility with non-unique prior.” Uncertainty in Economic Theory. Routledge, 2004. 141-151.

Appendix

7 Simulations of Behavior

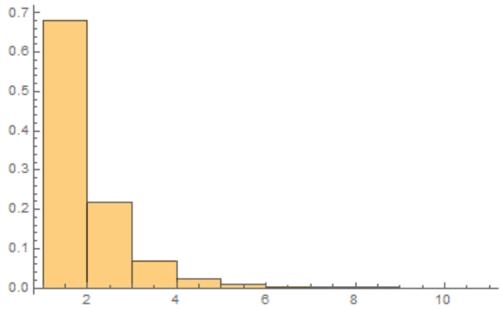

(a) KDNR # of draws

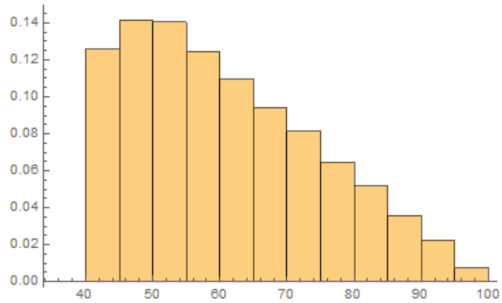

(b) KDNR Max Draw Dist.

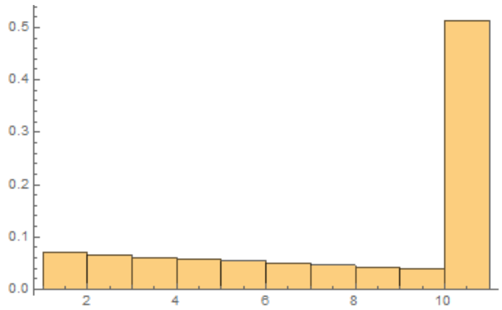

(c) KDR # of draws

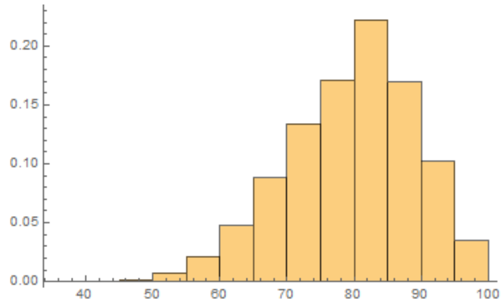

(d) KDR Max Draw Dist.

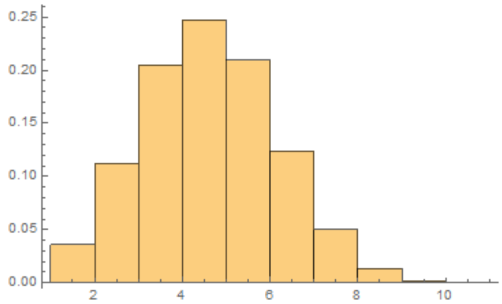

(a) UDNR # of draws

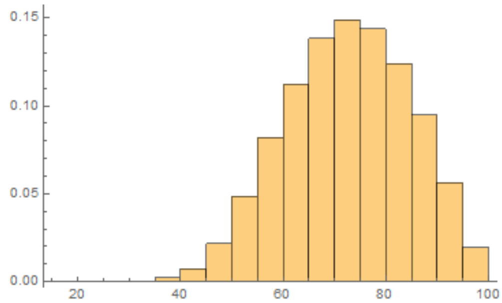

(b) UDNR Max Draw Dist.

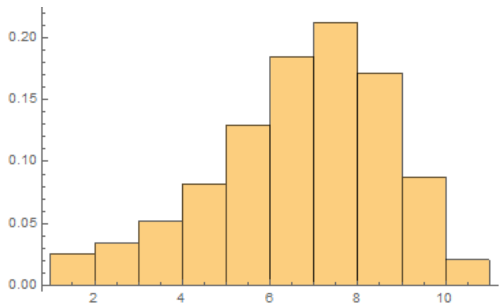

(c) UDR # of draws

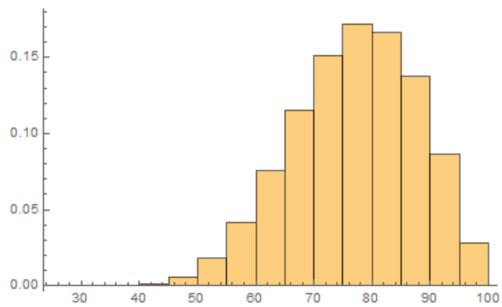

(d) UDR Max Draw Dist.

Figure 7: Simulations of behavior of # of draws and maximum draws in a 10 person game