Innovate versus Imitate: Theory and Experimental Evidence

John Duffy and Jason Ralston

February 13, 2018

Abstract

Agents and firms must exert effort to maximize their utility or profit. Often, they must explore the strategy space to find their optimal outcome. Or, if they find exploration to be too costly, they may copy the most successful strategy from someone else. This is observed in a firm's research and development decisions, an individual's habit setting patterns, and when individuals encounter novel settings. We formulate a novel model of sequential innovation versus imitation decisions made by individuals and implement it in the laboratory. We analyze the effects of complete versus incomplete information of payoff structure and static versus dynamic reward structure on termination of innovation. Among other things, these results help characterize whether subjects make decisions consistent with expected utility maximization or expected regret minimization.

1 Theory: Cardinal Reward

We have a situation where an agent must choose to either engage in a costly lottery or take the highest outcome from previous players. In other words, agents must decide whether to innovate or imitate. Let us denote the outcome of lottery for player i in period t as $x_{i,t} \sim dist(\theta)$. Then we can also denote the full history of lotteries in period t as $X_t = \{x_1, \ldots, x_t\}$ and the maximum of this set as $x_{max} = max\,\{X_t\}$. We also assume that the cost to receiving a lottery in any time period is c. Each agent begins their game with an endowment e, which may be used to purchase lotteries.

The payoff function is then defined to be

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

The decision to participate in a lottery in period t will depend on the expected value of the lottery, conditional on the draw being larger than the current maximum of lottery outcomes.

Specifically, this expected increase in the outcome must be worth rejecting the current maximum of lottery draws. Mathematically,

$$E\left(x_{i,t}|x_{i,t} \geq x_{max}\right) - c + e \geq x_{max} + e$$

which can be reduced to the following expression

$$E\left(x_{i,t}|x_{i,t} \geq x_{max}\right) - x_{max} \geq c. \tag{1}$$

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, for 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, agents will continue to receive lotteries 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 ans 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. □

Above it was assumed that the agent would always get at least the maximum draw from all previous draws. Of course, we could modify the payoff function to be

$$\pi_i = \begin{cases} x_i \text{ if innovate} \\ x_{max} \text{ if imitate} \end{cases}$$

which states that the agent faces the decision of either taking a lottery $x_i \sim dist(\theta)$ or taking the highest of the previous innovations. With this framing, the decision of whether or not taking a lottery is worth rejecting imitation. Mathematically, this can be stated as

$$E(x_i) \geq x_{max}.$$

What makes the 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 above proposition follows from logic similar to 1.

In the real world there are copyrights and fines for copyright violation. Therefore we proceed by introducing a cost to imitation d. The agent will incur this cost if they decide to imitate x_{max}. This induces the following payoff function

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

This new cost changes the decision criteria in (1) to

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

As expected, the cost to imitating will decrease the the number of agents who play imitate and will increase k, *ceteris paribus.*

1.1 Endogenous Distribution Change

Here we explore the possibility where a successful innovation moves the distribution of outcomes in some way. To signify the improvement of technology as a result of innovation, the distribution must be changed in a way such that either the distribution places a larger weight on the more extreme values of the support of the distribution, the support itself is changed to include a higher maximum, or some combination of the two changes. For the moment we will consider shifting the entire distribution to the right by some positive amount δ.

The decision criterion for each agent will not change, of course, but the number of expected draws k will change. In the simple case where agents may *not* copy the best draw in the history so far after a failed innovation, we can see that increasing the expectation of the distribution μ after a successful innovation will lead to a higher $E(x_i)$ for the next agent to consider. In response to to a successful innovation by the previous agent, the following agent now weighs the increased $E(x_i)$ against the new x_{max}.

The amount that k is increased, relative to the case when innovations do not cause a change in the distribution, depends on a race between increases in $E(x_i)$ and what the maximum of

j-many draws is expected to be. In other words, how many draws we expect to take depends on what we expect $E(x_i)$ to be after j draws and what we expect the maximum draw to be after j draws. These expectations are easy to calculate in the case when a distribution does not change as a result of previous draws, but not tractable when changes occur. For this reason, the distribution F of the random variable is specified along with the change to be applied after a successful innovation, and Monte Carlo simulations are run in order to determine an expected k.

As an example, let us consider draws coming from a Gaussian distribution. Initially this distribution is parameterized with $\mu = 5$ and $\sigma = 10$ and the cost of taking a draw is set to 1. In this simulation, every time their is a successful innovation, μ is increased by δ. For each value of δ, 1000 Monte Carlo simulations of the innovate/imitate game are run. The results of this simulation are shown in the appendix. Shown in these tables is that greater values of δ increase the number of draws that are expected to occur during a game. Notice that though the number of expected draws increases, while F is specified to have a support of $(-\infty, \infty)$, the game will end with probability equal to one. This is because as the number of draws n goes to infinity, x_{max} goes to the maximum of the support. Since $E[X]$ will always remain finite, there will at some point be a draw in the sequence where it is no longer beneficial to imitate.

The same is obviously not true of distributions with finite support. We can always find a new parameterization for the distribution such that the new expected draw is above the support's maximum value. This can be easily accomplished by simply changing the support itself. As an example, let us consider a triangular distribution with vertices set to $a = 0$, $b = 7$, and $c = 10$ initially. After every successful innovation, all three parameters are shifted up by δ. The results of these simulations can be found in the appendix. Note here that at $\delta = 4$ and above, innovation is the only action taken because $E[X] > x_{max}$ every time a new innovation is found.

Appendix

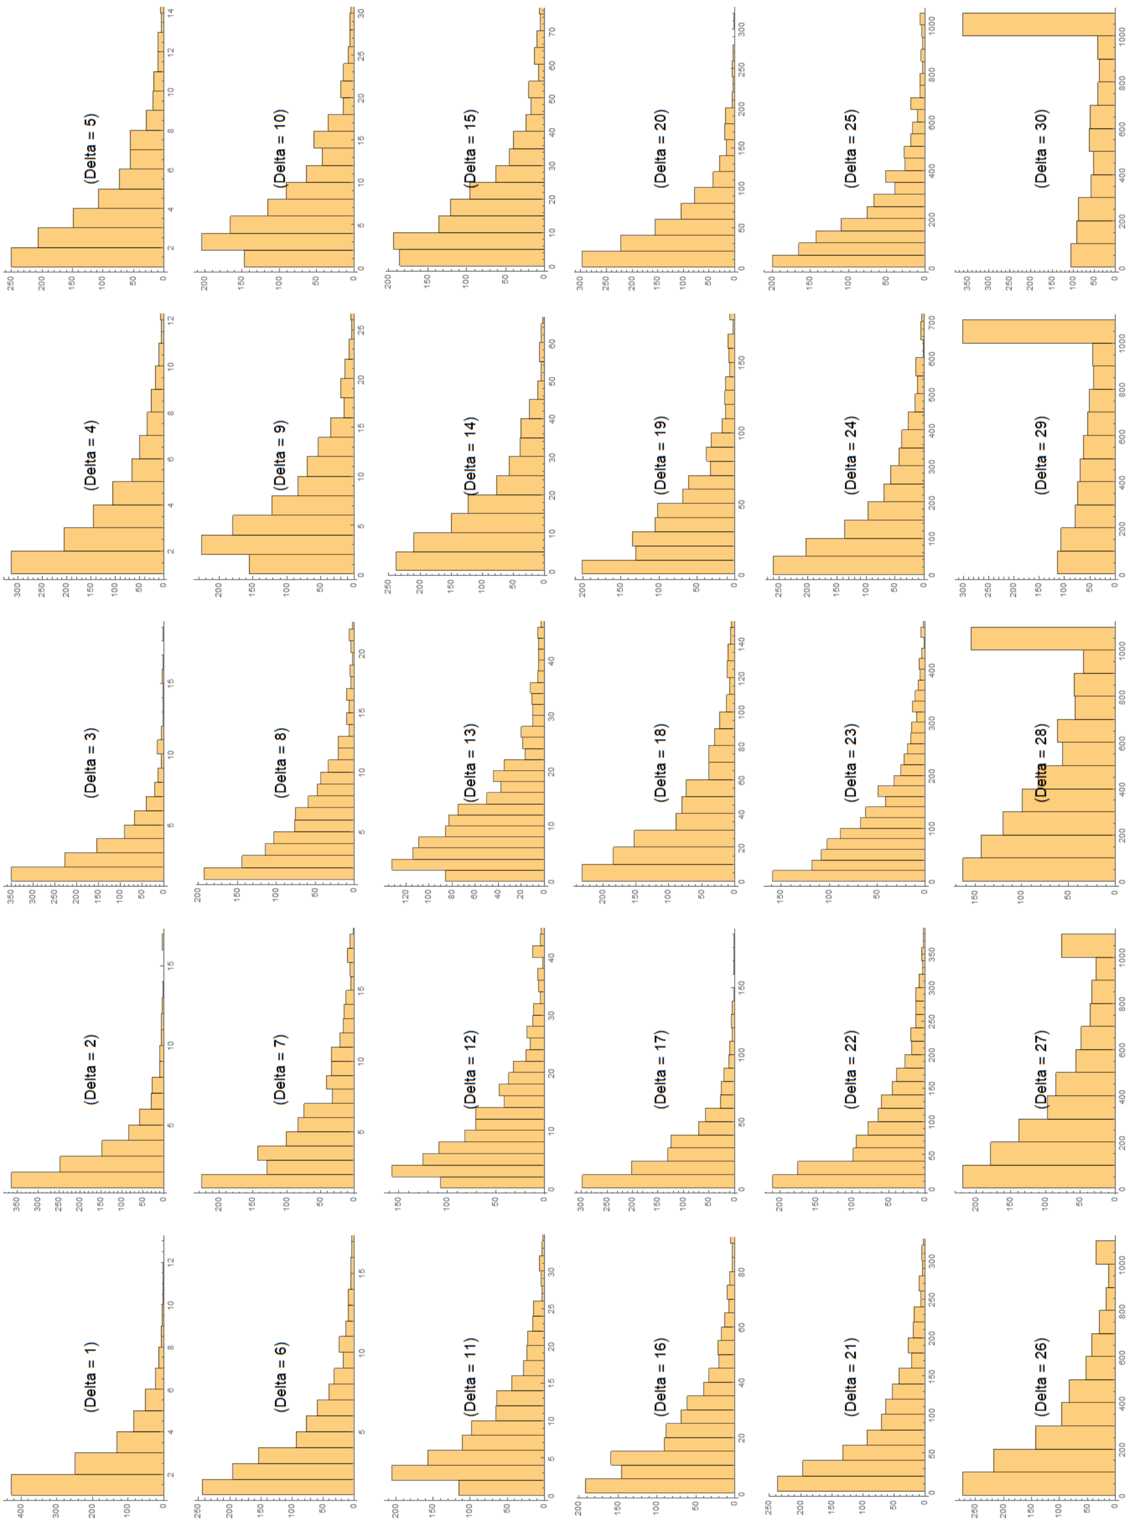
(Delta = 1)
(Delta = 2)
(Delta = 3)
(Delta = 4)
(Delta = 5)
(Delta = 6)
(Delta = 7)
(Delta = 8)
(Delta = 9)
(Delta = 10)
(Delta = 11)
(Delta = 12)
(Delta = 13)
(Delta = 14)
(Delta = 15)
(Delta = 16)
(Delta = 17)
(Delta = 18)
(Delta = 19)
(Delta = 20)
(Delta = 21)
(Delta = 22)
(Delta = 23)
(Delta = 24)
(Delta = 25)
(Delta = 26)
(Delta = 27)
(Delta = 28)
(Delta = 29)
(Delta = 30)

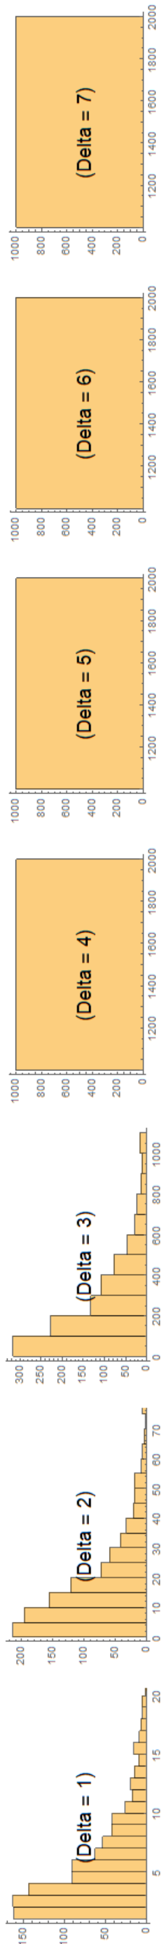
(Delta = 1)
(Delta = 2)
(Delta = 3)
(Delta = 4)
(Delta = 5)
(Delta = 6)
(Delta = 7)