A NUDGE OR A CRUTCH? TEMPTATION AND LEARNING IN SEQUENTIAL DECISION TASKS

Very Preliminary and Incomplete: Contents may Shift in Flight

William Neilson
University of Tennessee
Knoxville, TN
wneilson@utk.edu

Michael Price
Georgia State University
Atlanta, GA
mprice25@gsu.edu

Mikhael Shor
University of Connecticut
Storrs, CT
mike.shor@uconn.edu

January 2013

Abstract

Research in a variety of settings has shown that providing interim incentives can help people achieve long-term favorable outcomes. For example, researchers have found that paying children to read books can lead to higher performance on reading tests. At the same time, businesses have offered interim incentives, such as teaser rates on cable TV packages, to entice people to make decisions that may not be in their best long term interest. Both of these patterns attest to the ability of short-term incentives to have short-term effects in the desired direction. However, they leave unanswered the question of whether the learning they prompt transfers to new situations in a changing environment. This study uses a laboratory experiment to determine how interim payments, both on and off the optimal decision path, impact learning. Subjects play a simple game against a computer in which winning a prize requires five or six correct moves. Interim payments, smaller than the prize for winning, are inserted in some games, and treatment effects concern whether these interim payments lie on or off the optimal path, and whether they occur early or late in the game. After 30 rounds, the game is changed so that a new path becomes optimal.

We find that interim payments on the optimal path help subjects learn game-specific patterns, but impede them from learning to backward induct. In contrast, interim payments off the optimal path hinder learning, but those who do learn are able to backward induct when the game changes. The most favorable results, however, arise from the untreated game: learning in the absence of teaser payments is highly transferable. These results may have implications for social programs that try to reward good behavior in dynamic environments. They are also consistent with the premise that harder learning is more transferable.

Temptation, Learning, And Backward Induction In Sequential Decision Tasks

I. Introduction

Researchers have identified several settings in which interim payments assist people in achieving long-term goals. For example, Charness and Gneezy (2009) find that paying people to go to the gym can induce higher exercise levels even after the payments end. Fryer (2011) shows that paying students to read improves classroom performance but paying directly for test performance does not. Duflo, Kremer, and Robinson (2010) find that farmers in Kenya under-invest in fertilizer, but significantly increase overall investment in response to small temporary discounts. Ashraf, Karlan, and Yin (2006) show that voluntary acceptance of withdrawal restrictions on a bank account until attainment of a certain balance had a positive and long-lasting effect on savings rates.

All of these studies involve paying people to make a correct intermediate step, and all of these studies find "nudges" helpful both in the short-term and the long-term. Much more common, though, are interim payments designed to entice individuals to make a possibly-incorrect intermediate step. Mortgage lenders, credit card companies, cable TV providers, and mail-order clubs all offer promotional "teaser" rates to lure customers into entering long-term contracts, often reducing their long-run welfare. The use by businesses of interim payments, whether they be beneficial nudges or harmful teasers, testify to the fact that these interim payments stimulate the desired behavior, and the economics literature concurs. The economics literature is silent, however, on how these interim payments impact welfare in a changing environment. This paper seeks to provide a first look at this topic.

A pervasive finding in psychology is that people are surprisingly bad at transferring experiences from one environment into another. For example, schoolchildren who receive substantial practice with algebra problems can later solve identical forms of problems, but not slight variations (Cooper & Sweller 1987). College graduates who had taken economics courses showed no better general economic reasoning skills than graduates without economics courses (Voss et al. 1986). Levitt, List & Sadoff (2011) report that professional chess players often fail to transfer principles of backward induction to laboratory experiments, though the presumption that chess prowess implies an understanding of backward induction in the first place has been questioned by many psychologists. Chess grandmasters do not consider a larger game tree or apply more sophisticated logic than lower-ranked players. Instead, they rely on a much larger number of board positions retained in long-term memory and called upon in game play (de Groot 1966, Chase & Simon 1973). If even chess players learn patterns rather than principles, there is little reason why prowess in one strategic task should translate to non-related tasks. This is consistent with nearly a century of learning research, summarized thusly: "The lack of general transfer is pervasive and surprisingly consistent" (Detterman, 1993, p. 18).

Psychologists have long drawn a distinction between two types of learning (Thorndike & Woodworth 1901, Salomon & Perkins 1989). First, *reflexive* learning, represented by the instinctive learning of patterns, can be applied only to the task at hand or to nearly-identical contexts. Second, *mindful* learning involves a deliberate cognitive search for abstract general principles that can transfer to sufficiently different contexts. Does paying a person to go to the gym, for example, condition gym attendance, or more general principles about healthy behavior? Similarly, does paying children to read lead foster higher scores on *that* reading test, better performance on future *reading* tests, or generally better study skills? Conversely, if one accepts a teaser rate to enter into a welfare-reducing long-term contract, this delays adherence to the optimal savings path, but does any learning come out of this experience? Our goal is to understand not only whether interim payments affect success on the given task, but also whether they affect subjects' ability to learn in a way that transfers to new environments.

We run an experiment to determine how teasers affect learning and learning transfer in a relatively short, simple sequential decision setting. The experiment uses a race game (Dufwenberg, Sundaram, & Butler 2010, Gneezy, Rustichini, & Vostroknutov 2010, Levitt, List, & Sadoff 2010). The game starts with 21 stones. The subject and a computer opponent alternate by removing between 1 and *a* stones per turn. The subject wins a monetary prize if he, and not the computer, removes the last stone. The game has a unique optimal strategy. We examine the effect of teasers on a subject's ability to learn this optimal path. Accordingly, we reward subjects not only for taking the final stone, but also include a smaller "teaser" interim payment that subjects can earn.

Our teasers take two forms, with the interim payment either on or off the optimal path. When the teaser is on the optimal path, a subject who removes the last stone will also necessarily capture the teaser. Here, like paying people to go to the gym, teasers can help guide people to the optimal goal. When the teaser is off the optimal path, capturing the teaser stone and the last stone are mutually exclusive. A subject trying to learn the optimal path may be misguided to capture the teaser instead. Notably, given that the teaser stone pays less than the last stone, the presence of teasers does not alter the optimal path. Following 30 iterations of this game, subjects play a different race game where the only change is the action space, from between one and three stones per turn to between one and four stones per turn, or vice versa. This subtle change alters both the optimal path and whether a teaser is on or off this path. A subject who learned only a pattern of play (e.g., take stone 1, then stone 5, etc) or a general principle about teasers (take the teaser or don't) will not succeed upon a change in the action space. However, subjects who learned the concept underlying the game should be able to apply this principle to the new action space.

We address two primary questions. First, does the presence of teasers alter one's ability to learn the optimal path? Second, if a subject does learn the optimal path, does that knowledge transfer to other contexts? Or, put differently, does the presence of a teaser affect whether learning is reflexive with subjects merely mastering a pattern that coincides with the optimal path, or mindful with subjects gaining a more general understanding of backward induction that transfers to a different, but closely related, race game that follows.

We begin with a theoretical framework based on the notion of learning through epiphanies as proposed by Dufwenberg, Sundaram, and Butler (2010).[1] Epiphanies are random events, and when an epiphany occurs the subject learns the game. We assume that epiphanies come in two types. When a subject has a game epiphany he develops a sufficiently general understanding of backward

[1] Gneezy, Rustichini, and Vostroknutov's (2010) notion of learning through insight can be modeled in essentially the same way.

induction to solve the current game and also to solve the game when the action space changes. A game epiphany, then, corresponds to mindful learning. In contrast, when a subject has a pattern epiphany he only learns the optimal pattern, and this learning does not transfer when the action space changes. Pattern epiphanies correspond to reflexive learning.

The model implies that game epiphanies and pattern epiphanies crowd each other out. Consequently, mindful learning crowds out reflexive learning and vice versa. This makes it possible for interim payments to increase mindful learning and reduce reflexive learning, leading to beneficial outcomes even in a changing environment, but also for interim payments to increase reflexive learning at the expense of mindful learning, causing long-term harm in a changing environment.

The results of the experiment suggest that interim payments, whether they represent nudges on the optimal path or teasers off the optimal path, promote reflexive learning at the expense of mindful learning. Race games without any interim payments have the most transferable learning with 40% of subjects falling into this category, and games with nudges on the optimal path have the least with only 13% of subjects becoming mindful learners. Strikingly, with no interim payments mindful learning almost fully crowds out reflexive learners, with 95% of all learning being transferable. Games with on-path nudges show almost the opposite pattern, with 80% of learning being reflexive.

The data are consistent with the notion that learning in more difficult environments tends to be more mindful. Games with on-path payments prove to be the easiest to learn, while games with off-path or no teasers are more difficult. Nevertheless, in these more difficult games what learning there is tends to be transferable.

The findings raise concerns about the long-term benefits of paternalistic programs that provide payments for making the right step along a path. The data suggest that these programs will, in fact, help people complete the task at hand. The problem arises when the nature of the task changes because the results imply that helping them solve the original tasks actually makes it more difficult for them to solve the new task in the new environment. Moreover, this negative implication is not offset by people learning their lessons when fooled by interim payments like teaser credit card rates. These off-path teasers also make learning less transferable, perhaps because they make learning easier overall despite the bad start.

Section II constructs a model of learning through epiphanies, establishes the crowding-out result, and offers hypotheses. Sections III and IV describe two experiments based on the race game and conveys their results. Section V offers a discussion of the two experiments and relates evidence showing that "harder" learning is "better" learning. Finally, Section VI contains some conclusions.

II. Theory and Hypotheses

Let D denote a (nonstochastic) sequential decision problem in the set of such problems Ω. Individuals have an action space A which is the set of all possible sequences of actions the individual can take. Let $\pi(a_1,\ldots,a_m; D)$ denote the individual's payoff from choosing the sequence $\mathbf{a} = a_1,\ldots,a_m$ when confronted with decision problem D. Each decision problem in Ω is assumed to have a unique solution sequence $\mathbf{a}^* = a_1^*,\ldots,a_n^*$ so that any action sequence $\mathbf{a} \neq \mathbf{a}^*$ results in strictly suboptimal payoffs. Let $f: \Omega \rightarrow A$ be the mapping which assigns the payoff-maximizing action sequence for each problem D.

Because of its assumption of a perfect ability to maximize, standard economic theory predicts that individuals always play the sequence $f(D)$. When actually confronted with a decision problem, though, a boundedly rational individual may or may not figure out the optimal sequence. Ultimately, the outcome depends on whether and what the individual learns. Some individuals will not figure out the optimal sequence $f(D)$ in the time allotted, and we label these *non-learners*. In contrast, *learners* figure out the optimal path. However, there are two ways that an individual can do this, either by learning the problem-specific pattern $f(D)$, which only applies to the precise task at hand, or by learning how the general class of problems Ω works, in which case she learns the entire mapping f. One who learns the mapping can transfer that knowledge to solve any problem in Ω. We call a subject that learns the mapping f a *transferable learner* and one that only learns the specific pattern $f(D)$ a *non-transferable learner*. These notions have counterparts in the psychology literature. Transferable learning corresponds to mindful learning in which agents deliberately search for abstract general principles that apply across a broad range of contexts, and non-transferable learning corresponds to reflexive learning in which agents acquire patterns appropriate for only a narrow range of contexts (Thorndike and Woodworth 1901, Salomon and Perkins 1989).

One can distinguish between transferable and nontransferable learners in the following way. Suppose that an individual confronts the decision problem D several times in a row, and label someone who eventually figures out how to play the optimal sequence $f(D)$ by the last period as a *learner*. Now change the decision problem to D' with a different optimal sequence $f(D') \neq f(D)$. A transferable learner, having learned the mapping f, will be able to follow the optimal sequence immediately. A nontransferable learner, on the other hand, only learned the specific pattern $f(D)$ and so must learn the pattern $f(D')$ anew, which takes time.

We incorporate interim payments into our model. In the literature, *nudges* are interim payments for some element of the optimal sequence $a_1^*,\ldots,a_{n-1}^*$. They serve as an additional payment for taking a specific action along the optimal path before the final action and payment. Because the nudge only adds to the payment from the optimal action sequence, it does not change which sequence is optimal. To allow for interim payments off the equilibrium path, we generalize this notion with the idea of a *teaser payment*, which is a payment for taking a specific action which may or may not be in the optimal action sequence, but that does not change which action sequence maximizes the payoff from that decision problem. For example, suppose that the optimal action sequence is {1,5,9,13,17,21} and the individual wins $1 for completing the sequence all the way to action 21. A nudge would provide an additional payment for taking a specific action like 13, while a teaser might offer a payment for taking an off-the-optimal-path action like 11. To maintain the optimality of the optimal path, the teaser payment must be smaller than the winning payment.

To inform the construction of testable hypotheses, we offer a model of learning through epiphanies consistent with the ideas put forth in Dufwenberg, Sundaram, and Butler (2010) as well as the notion of insight found in Gneezy, Rustichini, and Vostroknutov (2010). An individual confronts the same decision problem D in periods $1,\ldots,N$. The individual learns how to solve the problem by having an epiphany. To allow for a distinction between transferable and non-transferable learning, we assume that there are two types of epiphanies. When an individual has a *game epiphany,* she learns the entire mapping f, and this learning transfers when the decision problem chances from D to the alternative problem D'. In contrast, when she has a *pattern epiphany,* she learns only the specific pattern $f(D)$ and this solution does not transfer when the decision problem changes to D'.

We assume that a subject who learns the optimal path, either through a game epiphany or a pattern epiphany, does not pursue further learning.[2] We treat epiphanies as random, independent events. Let p be the probability that the individual has a game epiphany in period t conditional on not having had an epiphany through period t-1. Similarly, let q be the probability of a pattern epiphany in period t conditional on no prior epiphanies. We assume that $p + q < 1$ and that p and q are independent of t.

Ultimately a policy-maker is interested in three measures: the rate of overall learning, the rate of transferable learning, and the conditional probability that learning is transferable. The overall learning rate is important because if the environment is stable and the optimal path does not change, the overall learning rate reflects overall success. The transferable learning rate, on the other hand, matters for a changing environment in which the agent must periodically adjust to a new optimal path. The conditional probability that transferable learning, then, tells whether a treatment tends to make learners more or less adaptable to changes in the environment. All of these are governed by the probabilities of the two types of epiphanies, as described in the following proposition.

Proposition. *The following statements hold.*

(a) When the probability p of a game epiphany rises, ceteris paribus, the probability of solving the problem by period t increases, the probability of becoming a transferable learner by period t increases, and the conditional probability of learning being transferable increases.

(b) When the probability q of a pattern epiphany rises, ceteris paribus, the probability of solving the problem by period t increases, the probability of becoming a transferable learner by period t decreases, and the conditional probability of learning being transferable decreases.

Proof (to be moved to the appendix). The probability of solving the game in period t is the probability that the player has had an epiphany at some point through period t, given by

$$P_t = 1 - (1 - p - q)^t. \tag{1}$$

For the learning to be transferable the individual must have a game epiphany and not a pattern epiphany. The probability that the individual has a game epiphany in period 1 is simply p. She becomes a transferable learner by period 2 if either (i) she became a transferable learner in period 1, or (ii) she did not learn anything in period 1 but has a game epiphany in period 2. The probability of becoming a transferable learner by period 2, then, is $p + (1 - P)p$. Continuing this logic yields that the probability of becoming a transferable learner by period t is

[2] This assumption could be formalized by a model in which agents do not foresee the change in the decision problem and in which cognitive effort is costly. With no anticipated change in the decision problem, the pecuniary benefit of further deliberation is zero, so agents will not undertake the costly activity.

$$P_t^* = \sum_{s=1}^{t}(1-P)^{s-1}\,p^* = [1-(1-p-q)^t]\,\frac{p}{p+q} = \frac{p}{p+q}P_t. \tag{2}$$

Given these formulations, it becomes simple to characterize the conditional probability that learning is transferable, and it is

$$P_t^*/P_t = p/(p+q). \tag{3}$$

To prove part (a), note that $\partial P_t/\partial p = t(1-p-q)^{t-1} \geq 0$ and that $\partial(P_t^*/P_t)/\partial p = q/(p+q)^2 \geq 0$. It follows that $\partial P^*_t/\partial p \geq 0$ by the product rule.

To prove part (b), note that $\partial P_t/\partial q = t(1-p-q)^{t-1} \geq 0$ and that $\partial(P_t^*/P_t)/\partial q = -p/(p+q)^2 \leq 0$. It remains to show that P^*_t decreases when q increases. We do this by induction. Let P^*_t and R^*_t denote the probabilities of a having a game epiphany by period t when the single-period probabilities of a pattern epiphany are q and r, respectively, where $r > q$. For the first step of the induction process, $P^*_1 = P^{**}_1 = p$ which is independent of q. Now suppose that $R^*_{t-1} \leq P^*_{t-1}$. The new transferable learners in period t are those who have not already learned but have a game epiphany in period t, yielding $P^*_t = P^*_{t-1} + p(1 - P_{t-1})$ when the probability of a pattern epiphany is q and $R^*_t = R^*_{t-1} + p(1 - R_{t-1})$ when the probability of a pattern epiphany is r. Subtracting yields $P^*_t - R^*_t = [P^*_{t-1} - R^*_{t-1}) + p(R_{t-1} - P_{t-1})$ which is nonnegative because $R^*_{t-1} \leq P^*_{t-1}$ by hypothesis and $P_{t-1} \leq R_{t-1}$ follows from the result that overall learning increases with q. □

Part (a) of the proposition examines what happens if game epiphanies become more likely but pattern epiphanies stay the same. The overall learning rate, the transferable learning rate, and the conditional likelihood that learning is transferable all increase. Part (b) looks instead at what happens if game epiphanies stay equally likely but pattern epiphanies become more frequent. The learning rate increases, the transferable learning rate decreases, and what learning does occur is more likely to be non-transferable. This has the important implication that one type of learning crowds out the other. Intuitively, when the probability of a pattern epiphany increases it has no effect on the single-period learning of an agent who has not yet learned. However, because more agents have had pattern epiphanies in the past, the pool of agents who have not yet learned shrinks, and so fewer agents can have game epiphanies in subsequent periods. This crowding-out phenomenon leads to our first hypothesis.

Hypothesis 1. Transferable and non-transferable learning rates are negatively correlated.

If one type of learning crowds out the other, the possibility arises that a treatment could be prove beneficial to overall learning but detrimental to transferable learning. Thus, the impact of interim payments ultimately becomes an empirical question. If a payment on the optimal path increases overall learning, it might do so by increasing transferable learning at the expense of non-transferable learning, but it might instead do the reverse. Because the model does not relate the placement of interim payments directly to the two epiphany probabilities, the model is silent on whether on-path or off-path interim payments will promote transferable or non-transferable learning. It does predict, though, that an increase in one comes at the expense of the other.

In the absence of guidance from the model, we generate hypotheses using insight from the empirical literature. Some of the evidence cited in the introduction, namely that of Shraf, Karlan, and Yin (2006) regarding saving behavior, Charness and Gneezy (2009) regarding exercise, Duflo, Kremer, and Robinson (2010) involving fertilizer use in Kenya, and Fryer (2011) regarding reading and test performance among school children, suggests that interim payments placed on the optimal path have positive long-lasting effects. In the setting described here this would arise from the interim payments promoting transferable learning through an increase in the probability of a game epiphany. This yields the second hypothesis.

Hypothesis 2. Compared to games without teasers, games with on-path teasers should have higher success probabilities in every period, higher rates of transferable learning, and higher conditional rates that learning is transferable.

There is less research about unfavorable, off-path teasers. Still, the fact that firms use them in the marketplace suggests that they should be effective. Short-term effectiveness translates into lower overall values of p and q because then fewer individuals have epiphanies that teach them to avoid the off-path teasers. Long-term effectiveness means that the off-path teasers should work again when the environment changes, which means that p declines by more than q. These considerations lead to our second hypothesis.

Hypothesis 3. Compared to games without teasers, games with off-path teasers should have lower success probabilities in every round, lower rates of transferable learning, and lower conditional rates that learning is transferable.

Hypotheses 2 and 3 together imply that, in terms of overall successful performance by agents, on-path teasers should do the best and off-path teasers should do the worst.

III. EXPERIMENT 1

The first experiment is designed to identify how the presence and location of a teaser payment impact learning in a decision task. The problem itself is based on the race-to-21 game studied by Dufwenberg, Sundaram, and Butler (2010) and Gneezy, Rustichini, and Vostroknutov (2010).

A. DESIGN

An individual decision-maker plays against a computer. The decision-maker moves first choosing an integer number of steps between 1 and s, inclusive, where s is the *step length* and is one of the treatment variables in the experiment. The computer follows, and the players alternate, with the player ending on step 21 winning a prize. We identify a path according to where a turn ends, which allows us to identify a path that is not contingent on where the other player ended her turn. We consider two values of s, either 3 or 4. The game can be solved using backward induction yielding an optimal path that ends turns on steps {1, 5, 9, 13, 17, 21} when the maximum number of steps per turn is $s = 3$ and one that ends turns on steps (1, 6, 11, 15, 21) when $s = 4$.[3]

[3] To see this, consider $s = 3$. If the individual ends a turn on 18, 19, or 20, the computer can end the next turn on 21 and win the final payoff. To avoid this, the individual must end the turn on $21 - (s + 1) = 17$. Similarly, she must end the previous turn on $21 - 2(s + 1) = 13$, and so on to construct the optimal path.

To the standard race game we add an interim teaser payment on some intermediate step $T < 21$. We call the reward for ending a turn on *21* (and winning the game) the *final payoff* and the reward for ending a turn on T the *teaser payoff*, with the teaser payoff half the size of the final payoff. This relationship between the two payoffs implies that the optimal path depends only on s and X and not on T, the placement of the teaser. An individual can only collect the teaser payoff by ending a turn on step T, and passing through step T means that no one collects the teaser payoff. Some games do not have teaser payments, and to allow for this possibility we use the convention $T = 0$ for such a case.

Holding the two payoffs constant, a decision problem D can be described by the pair (s, T). So, for example, the decision problem (3,13) corresponds to the race-to-21 game with a maximum step length of $s = 3$ and a teaser on step $T= 13$, and (3,0) corresponds to the same game but without a teaser payoff. The mapping $f(s, T)$ identified in the preceding section has $f(3, T) = \{1, 5, 9, 13, 17, 21\}$ and $f(4, T) = \{1, 6, 11, 15, 21\}$, both independent of T. The experiment is designed to determine how the presence and placement of T impacts the learning of the rule f versus the problem-specific pattern $f(s, T)$.

Subjects first play the race game (s, T) for 30 rounds and then play the game (s', T) for 15 rounds, where s and s' are either 3 or 4 and $s \neq s'$. With these pairings the game length is the same in both treatments and the teaser placement is the same in both, but the step length differs between treatments which is sufficient to change the optimal path. The different treatments are shown in Table 1, which also shows the number of subjects in each treatment. The treatments are labeled according to the problem faced in the first 30 rounds, and they differ according to whether there is or is not a teaser, whether the teaser is on or off the optimal path, and whether the teaser comes early late in the decision problem. There are a total of 10 treatments in a (2 step length) × (5 teaser location) design, with step lengths $s \in \{3, 4\}$ and teasers $T \in \{0, 5, 6, 11, 13\}$. Payments were held constant across treatments with a final payment of \$1 and a teaser payment of 50¢. Subjects were paid their earnings for all 45 rounds. Subjects participated in only a single 45-round treatment.

The sessions took place in the University of Tennessee Experimental Economics Laboratory. The lab is set up with 27 individual client workstations networked to an intranet server and separated by study carrel walls. Subjects remained anonymous to each other and their decisions remained private throughout the experiment. The race-to-21 game was programmed in Perl6, and a screen shot from the first turn in the on-path late-teaser game (3, 13) is shown in Figure 1. In the interface the steps are referred to as "stones" and the steps that yield payoffs are colored green and marked with the payoffs, in points, for ending a turn on that stone. Points were converted at the rate of 150 points to a dollar.

TABLE 1: Experiment I treatments and number of subjects

Treatment (Game 1)	teaser location	step length Game 1	step length Game 2	N
On-path teaser				
Early	5	3	4	22

	6	4	3	19
Late	11	4	3	23
	13	3	4	20
				84
Off path teaser				
Early	5 (unachievable)	4	3	19
	6	3	4	23
Late	11	3	4	25
	13	4	3	24
				91
No teaser				
		3	4	24
		4	3	21
				45

FIGURE 1: Screen shot of race-to-21 interface

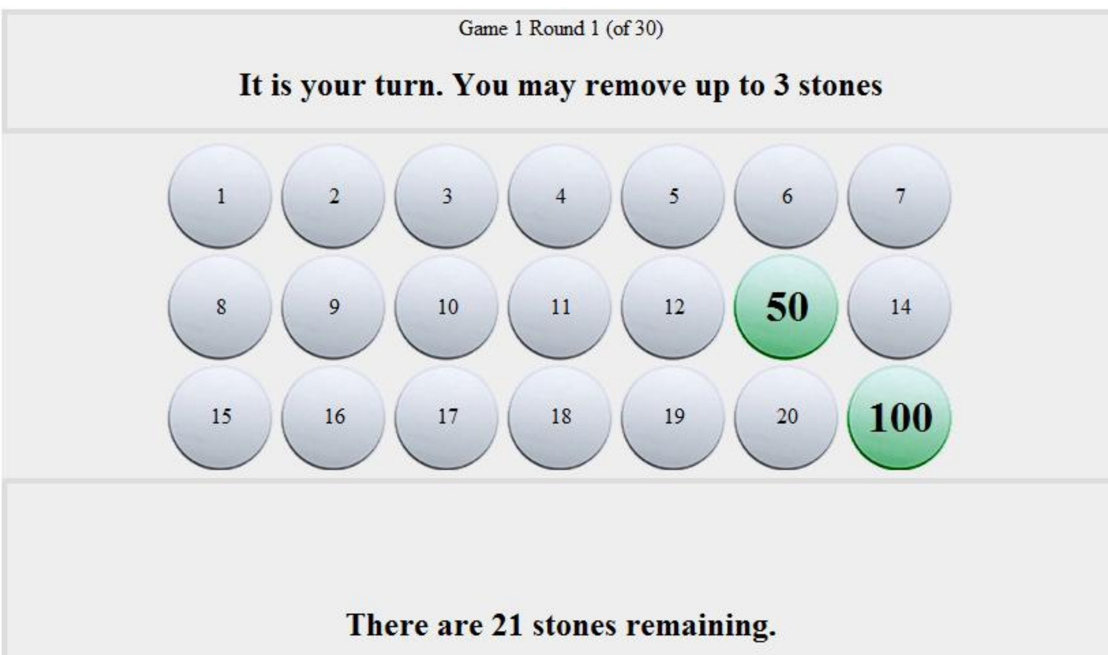

At the start of each lab session the experimenters gave subjects written instructions for the game to complement the same instructions visible on the participants' computer screens. Subjects were asked to follow along as experimenters read the instructions aloud. Subjects were told they earned points when they, and not the computer, removed the green stones. Players removed stones by mousing over the stones they wanted to remove (which highlighted them on the screen), and then clicking on the last stone in the series they wanted to remove. They were given the opportunity to practice with the interface and ask questions before actual play began. All navigation through the pages of the experiment occurred by clicking a button at the bottom of each screen labeled "Proceed." At the completion of round 45 subjects were asked to complete a short questionnaire (the results of which are not used in this paper). Immediately following the experiment subjects were paid their earnings privately in cash. Sessions lasted about an hour and earnings averaged $18.48.

B. RESULTS

We begin by examining the raw data on the rate of success in the first race-to-21 game. Table 2 shows the overall rates of winning over all thirty rounds. When there is no teaser, subjects follow the optimal path and remove the winning stone in 29% of the rounds. A teaser on the optimal path increases this proportion to 52% of the rounds ($p<0.001$) using either a Mann-Whitney or t-test.[4] On the other hand, off-path teasers do not have much effect, reducing the fraction of rounds won by only a statistically insignificant 4 percentage points ($p>0.500$). This suggests that failures to achieve the final payoff cannot be explained solely by subjects pursuing the off-path teaser, as doing so would make subjects win less often than in the no-teaser treatment.

TABLE 2. Summary of first-game activity

	Achieve teaser	Achieve winning stone
On-path	74%	52%
Off-path	58%[†]	25%
No teaser		29%

[†]Excluding game (4, 5) where the off-path teaser is impossible to reach

Nevertheless, teasers do serve as a distraction, with 82% of subjects taking the off-path teaser in the very first round of the attainable off-path teaser treatments, and with *every* subject taking the off-path teaser at least once by round 5. Furthermore, subjects do not appear to abandon off-path teasers over time as roughly half of all subjects take the off-path teaser in each of the last eight rounds.

The first 30 rounds (game 1) allow us to identify players as learners. Game 2 involves a change in the optimal path and thus allows us to classify learners as transferable or non-transferable. We identify a subject as having learned game 1 if she wins game 1 in at least two of the final three rounds. We require two wins to offset the unlikely but non-negligible chance that random play can occasionally win. As a practical matter, however, alternate definitions of learning yield nearly identical qualitative results. Among subjects who learn game 1, we classify a subject as a transferable learner if she goes on to win at least two of the *first* three rounds of game 2, and a non-transferable learner otherwise.

[4] For these significance tests, the unit of observation is the fraction of wins for a single subject.

Table 3 reports the fractions of subjects classified as learners by treatment. The on-path treatments had the highest fraction of subjects who learned game 1 at 70%, while the no-teaser treatments had the lowest fraction at 42%, a significant difference (Fisher exact test p=0.002). As one might expect from Table 2, off-path treatments, with 47% learning game 1, were not significantly different from the no-teaser treatments (p=0.590). The unachievable teaser treatment appears to be different from achievable teaser treatments, though, with the unachievable off-path teaser treatment generating a 74% while the achievable off-path teaser treatments combining for a 40% learning rate.

Our hypotheses in Section II concern the conditional probability that learning is transferable, and those fractions are reported in the final column of the table. Treatments without teasers generate the highest proportion of transferable learners. Specifically, 95% of those who learn to win in a game without teasers transfer that learning to the next game. This does not leave much room for "nudges," i.e. on-path teasers, to improve transferability rates. The table shows that rather than improving transferability rates, on-path teasers dramatically harm transferability rates with only 19% of subjects who learned the first game with the on-path teaser transferring that learning to the second game. Early on-path teasers do somewhat worse in the number of learners but better in the proportion of transferable learners. Off-path teasers lead to higher transferability rates than on-path teasers, with about half of learning being transferable in the off-path teaser games.

TABLE 3. Classification of subjects by learning type

Treatment	Non-learners	Learners		Fraction of learning that is transferable
		Nontransferable	Transferable	
No teaser	58%	2%	40%	95%
Off path	53%	22%	25%	53%
Early	38%	33%	29%	46%
Late	65%	12%	22%	65%
Achievable	60%	19%	21%	52%
Not achievable	26%	32%	42%	57%
On path	30%	57%	13%	19%
Early	39%	41%	20%	32%
Late	20%	72%	7%	9%

The model in Section II predicts that one type of learning crowds out the other, and the numbers in Table 3 are consistent with this prediction. No-teaser treatments have the lowest levels of non-transferable learning and the highest levels of transferable learning, on-path teaser treatments have the most non-transferable and least transferable learning, and the off-path teaser treatments lie in between. Pooling the data into three categories (no teaser, on-path, off-path), the correlation between transferable and non-transferable learning rates is −0.86, providing support for Hypothesis I, that non-transferable learning crowds out transferable learning.

Result 1. Non-transferable learning crowds out transferable learning.

The implications for the ability of on-path teasers to promote transferable learning are stark in the table. Compared to the no-teaser treatments, on-path teasers promote overall learning, increasing it from 42% of subjects to 70%. One might hope that with all of these additional learners, *some* of them would be transferable learners so that the unconditional rates of both types of learning would increase. Table 3 shows, however, that the *unconditional* probabilities of transferable learning are lowest for the on-path teaser treatments at 13% overall and highest for the no-teaser treatments at 40%. Clearly the data do not support Hypothesis 2, that on-path teasers promote transferable learning.

Result 2. O-path teasers lead to more overall learning but less transferable learning than no-teaser treatments.

Off-path teaser treatments generate conditional and unconditional rates of transferable learning that lie between those of on-path treatments and no-teaser treatments. Because the overall learning rates for the off-path and no-teaser treatments are so similar, the primary impact of the off-path teaser is to change the mix of transferable and non-transferable learning by promoting pattern learning, at least when compared to the no-teaser treatment. The data are partially consistent with the patterns in Hypothesis 3, with the exception being the prediction for overall learning rates.

Result 3. Off-path teasers lead to lower conditional and unconditional rates of transferable learning than no-teaser treatments.

The qualitative results in Table 3 are robust to many alternative ways of measuring learning and transferability. As a robustness check, we can ask what each type of learner did in the *last* period of game 2. Over all games, 100% of transferable learners won in the last round of game 2, 69% of non-transferable learners did so, and 23% of non-learners, with all three binary comparisons highly significant (Fisher exact test $p<0.001$). This suggests that behavior in the first three periods of game 2 was not by luck or accident, but really was driven by what subjects learned in game 1.

IV. EXPERIMENT 2

In the first experiment subjects played 30 rounds of a particular race-to-21 game and then played 15 rounds of an otherwise identical game with a different action space. Specifically, those who played a game with a step length of three in the first 30 rounds played a game with a step length of four in the subsequent 15 rounds, and vice versa. The teaser, if there was one, remained in the same position throughout the 45 rounds. Consequently, an on-path teaser game was followed by an off-path teaser game, an off-path teaser game was followed by an on-path teaser game, and a no-teaser game was followed by a no-teaser game. This raises two issues for identifying transferability of learning. First, subjects for whom the first game had an on-path teaser had to learn in the second game to avoid the teaser. In contrast, subjects who first learned to avoid the teaser in an off-path treatment had to learn to take the teaser in the subsequent on-path treatment. These two learning processes might be different, skewing transferability rates in favor of one treatment or the other.

Second, and more compelling, Table 3 shows that the on-path teaser games are easier to learn than the off-path teaser games. Subjects who played the easy, on-path game first were faced with the more difficult, off-path teaser game second, while subjects who played the harder game first moved to the easier game second. Thus, it is possible that our experiment conflates learning with the relative ease of the second game. Experiment 2 addresses this issue by holding the second task constant across treatments.

A. Design

We ran two additional sessions at the University of Tennessee. In one treatment, 21 subjects played 30 rounds of the off-path teaser game (4, 13) followed by 15 rounds of the no-teaser game (3, 0). In the other treatment, 24 subjects played 30 rounds of the on-path teaser game (4, 11) followed by the no-teaser game (3, 0). For comparison purposes, we use the Experiment 1 treatment in which 21 subjects played 30 rounds of the no-teaser game (4, 0) followed by 15 rounds of the no-teaser game (3, 0).

All other aspects of the experiment, including the interface and payoff parameters, were identical to those used in Experiment 1.

B. Results

Figure 2 eloquently tells the story. The horizontal axis shows the round number and the vertical axis measures the fraction of the subjects that won the final payoff in that round. The game changes from one of three teaser conditions to a no-teaser game in round 31. Panel A shows subjects who did not learn game 1. Recalling that a subject is classified as a learner if he or she won at least two of the last three rounds of the first game, performance of the non-learners is poor in rounds 28-30 by definition. Unsurprisingly, it is poor in earlier rounds.[5] More surprisingly, failure to learn in the first game perfectly predicts failure to learn in the second.

Panel B shows only subjects who learned game 1. Subjects who began with the on-path teaser game appear to learn game 1 earlier, as the curve corresponding to the on-path treatment begins higher and rises more quickly than the curves corresponding to the other two treatments. By definition, though, learners in the other two treatments caught up so that they all had perfect performance for the last several rounds of game 1.

[5] The single "blip" reflects the only non-learner in the on-path teaser treatment who appears to have won game 1 once, but not again.

The results of experiment 2 mirror those of experiment 1, both in terms of learning game 1 and in the transferability of learning to game 2. First, on-path teasers lead to significantly higher rates of learning game 1 than either off-path or no teaser treatments (Fisher exact test $p<0.001$). Second, upon beginning of game 2, there is a sharp reduction in performance in both teaser treatments. Once again, the on-path treatment fares the worst. Only about 10% of learners transfer that learning immediately to game 2. In contrast, about 40% of the off-path treatment learners transfer immediately, and about 90% of the no-teaser treatment learners transfer immediately. Beyond immediate transferability, Figure 2 indicates that differences in immediate success of subjects upon encountering game 2 persist through later rounds.

FIGURE 2. Performance of subjects by learning type

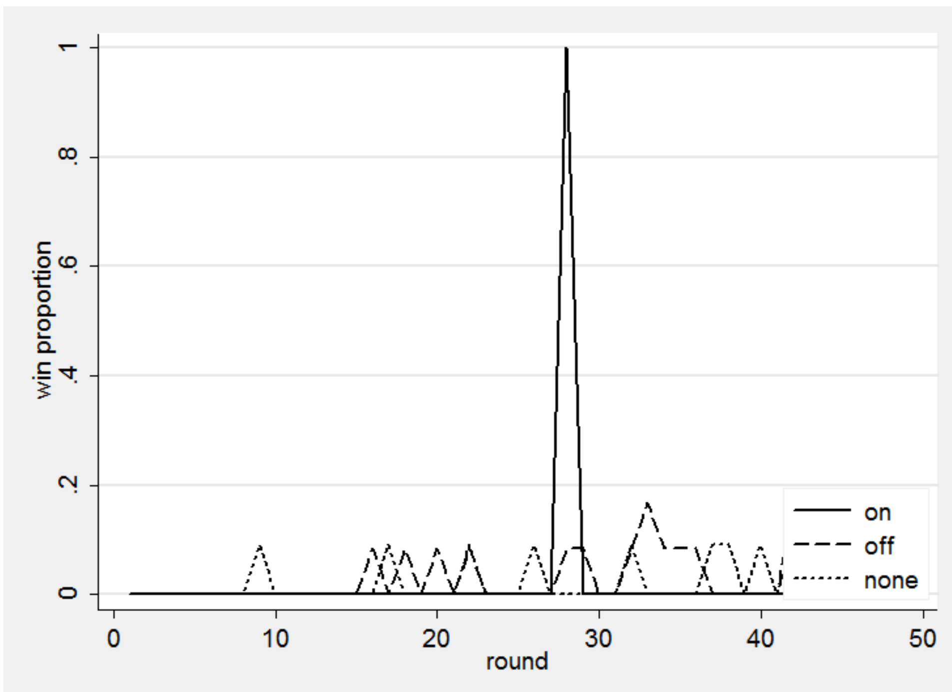

Panel A: Non-learners

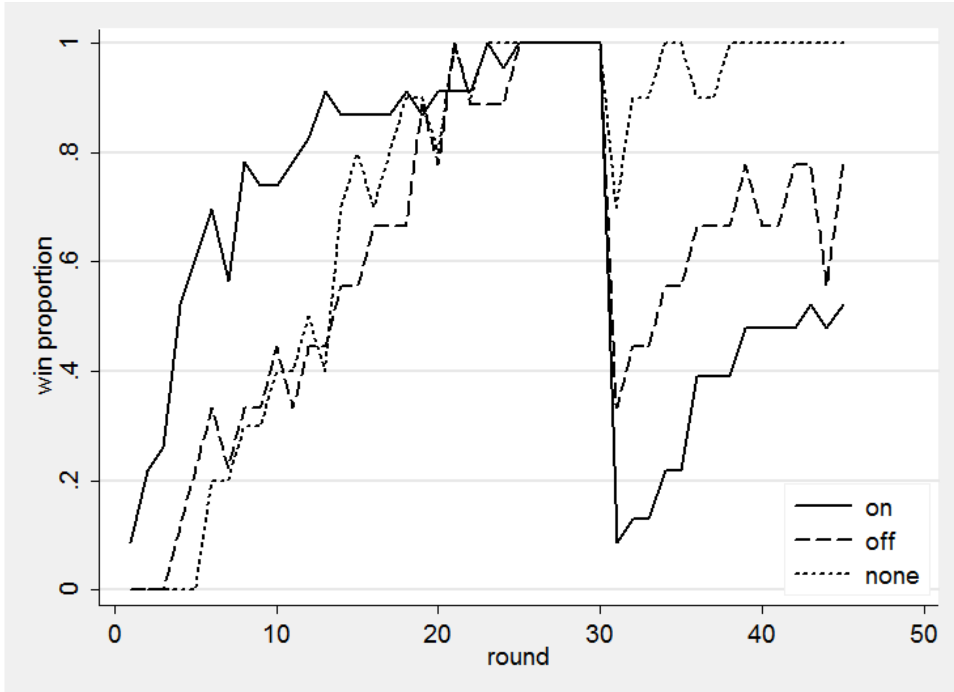

Panel B: Learners

V. Discussion

The data from the two experiments show that interim payments, either on or off the optimal path, deter transferability of learning. There is a clear pattern. Looking only at game 1 data, subjects tended to learn on-path teaser games frequently but off-path or no-teaser games less frequently. The learning from the no-teaser and off-path treatments tended to transfer well to game 2 with the new action space, but the learning from the on-path teaser treatments tended not to. We interpret these results as having implications for the wisdom of using nudges or other interim incentive payments in a changing environment. They have an additional interpretation, though.

The data on game 1 behavior suggest that some games are harder than others in the sense that epiphanies occur less frequently in some games than in others. Furthermore, the easiest games are the ones with the on-path teasers, which are also the ones with the least transferable learning. Hypotheses 2 and 3 in Section II were informed by the empirical literature suggesting long-term benefits from nudges. The experimental results suggest that a different mechanism might be at work here, in that making the optimal path easier to find may lead to shallower learning, and making the task harder makes learning "better."

To examine this possibility, we analyze how the difficulty of game 1 impacts performance in game 2 across both of our experiments. In Table 5, we present regression results on the determinants of a subject's performance in game 2. The dependent variable is a subject's overall win percentage in game 2. Independent variables are a subject's win percentage in game 1 and measures of the ease of both games 1 and 2 that the subject played. Game ease is measured as average win rate over the first thirty rounds for all subjects who played that game, excluding oneself.[6] One would suspect that subjects would have more trouble with harder second games, so that is controlled for by including the fraction of the time that other subjects won that particular game when it was the first game they faced. One might also suspect that a subject who won more often in game 1 would be more adept at learning these games, all else equal, and the subject's own first-game winning percentage controls for these ability differences. The coefficient of interest is the one on game 1 ease, because its sign shows whether easier or harder first tasks promote more transferable learning.

The results show that subjects won the second game more often when the second game was easier and when they, themselves, did better in the first game. However, all else equal, subjects did better in the second game when the first game was harder. The exact meaning of the coefficient on game 1 ease deserves further consideration. Fixing the subject's win percentage in the first game, the subject would have won game 2 about 8.3 percentage points *less* often if other subjects had won game 2 ten percentage points more often. In other words, and reversing the direction of impact, a subject does better in game 2 if he did well on a harder game 1 than if he did well on an easier game 1.[7]

The evidence in Table 2 suggests that placing the teasers on the optimal path makes the game more winnable and subjects more successful, but the regression in Table 5 suggests that this success comes at a price. Making the first game easier proves detrimental to being able to adapt to a changing environment. Furthermore, one of the most striking findings in Table 4 is that the single-period probability of a pattern epiphany in the unmanipulated, no-teaser game is miniscule compared to the rest of the entries. While on-path nudges make winning the game more common, they also turn non-transferable learning from an anomaly into a majority outcome.

TABLE 5. Determinants of Game 2 win percentage

	OLS estimate (std. err.)
Game 2 ease	0.526***
	(0.130)
Game 1 win percentage	0.869***
	(0.053)
Game 1 ease	-0.830***
	(0.116)
Constant	0.208***
	(0.080)

N=264. *** indicates significance at 1%. R^2=0.523, *F*=95.02.

[6] This approach mirrors Guryan, Kroft, and Notowidigdo (2009) to disentangle individual and group effects. They suggest first testing for random assignment by including a variable for average performance across all participants excluding oneself ("experiment difficulty"). Regressing individual game 1 performance on both experiment difficulty and game 1 difficulty, we obtain an insignificant parameter (p=0.284) for game 1 difficulty, evidencing random allocation.

[7] The coefficients on game 1 win percentage and game 1 ease have almost equal magnitudes. This implies that if a subject wins an additional ten rounds in the first game, but so does everyone else, then that subject is no more or less likely to win rounds in the second game because the two coefficients cancel each other out.

VI. Conclusion

Sixty years ago, noted psychologist James Deese noted that "There is no more important topic in the whole psychology of learning than transfer of learning." (Deese 1958, p. 213). With the failure to follow optimal paths to retirement, savings, and educational success, the sentiment has taken on economic import. Many successes have been reported with providing interim incentives to guide people towards their goals. We find in a simple decision setting that such interim incentives have conflicting effects. Guiding people along the optimal path increases their likelihood of success, but at the expense of future success in decisions unaccompanied by similar guides. Conversely, policies that guide people off optimal paths, such as pricing schemes that capitalize on consumers' time-inconsistent preferences (DellaVigna & Malmendier 2004), beget short-term failures but possibly also arm people with a better understanding of how to handle future environments.

In our experiment, teasers effectively provide a point of salience that allows people to focus on shorter horizon problems. Dufwenberg, Sundaram, and Butler (2010) conclude that such shorter horizon games spur an epiphany about backward induction. Gneezy, Rustichini, and Vostroknutov (2010) similarly conclude that focusing on a game within a game can help identify optimal paths. Our results suggest that such epiphanies need to happen without significant guidance, else they fail to achieve the insights required for transferability to new environments.

Of course, people also face many teasers that attempt to discourage them off of the optimal path. Our results with respect to these off-path teasers are more qualified. Overall, off-path teasers reduce short-term success, but encourage substantially more mindful learning than on-path teasers. The most stark findings relate to off-path teasers that are not actually obtainable, perhaps analogous to offers that sound—and indeed are—too good to be true. These are associated with the highest rates of both short-term success and long-term mindful learning.

Rates of mindful learning in the no teaser treatments exceed those with on-path and with off-path teasers. While an unguided environment, free of either helpful or distracting rewards, is not easy to master, its mastery almost always carries with it a mindful understanding that readily translates to new environments. Overall, our results suggest that efforts to guide good behavior, be they rewarding students for reading, adults for exercising, or people for saving, are likely to increase the desired activity in the short term, but may potentially hinder the natural process of learning about good study, lifestyle, and financial habits, in general.

For those of us who teach, the evidence we present is consistent with the premise that we do our students more good by being hard than by being easy. Being easy allows students to perform better on tests and earn higher grades, but being hard ensures that those who do manage to learn are able to apply their knowledge in new settings. Relatedly, teachers often face the dilemma of how to set the target for a heterogeneous class. Do they teach for the top students, the bottom students, or somewhere in between? The evidence from this experiment provides an argument for teaching to the top. Doing so may lead to fewer students learning, but those who do learn are more likely to do so mindfully and more likely to be able to transfer their knowledge to new settings. The real danger of teaching to the bottom, though, is that it promotes reflexive learning not just among those at the bottom, but also among those at the top.

References

Asheim, Geir B., and Martin Dufwenberg (2003), "Deductive reasoning in extensive games," *Economic Journal* 113(487), 305–325.

Ashraf, Nava, Dean Karlan, and Wesley Yin (2006), "Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines," *Quarterly Journal of Economics* 121(2), 635–672.

Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg (2001), "What accounts for the variation in retirement wealth among U.S. households?" *American Economic Review* 91(4), 832–857.

Brown, Alexander L., Zhikang Eric Chua, and Colin F. Camerer (2009), "Learning and visceral temptation in dynamic saving experiments," *Quarterly Journal of Economics* 124(1), 197–231.

Charness, Gary, and Uri Gneezy (2009), "Incentives to exercise," *Econometrica* 77(3), 909–931.

Chase, William G., and Herbert A. Simon (1973), "Perception in chess," *Cognitive Psychology* 4(1), 55–81.

Cooper, Graham, and John Sweller (1987), "Effects of schema acquisition and rule automation on mathematical problem-solving transfer," *Journal of Educational Psychology* 79(4), 347–62.

Cutler, David M., Edward L. Glaeser, and Jesse M. Shapiro (2003), "Why have Americans become more obese?" *Journal of Economic Perspectives* 17(3), 93–118.

Deese, James (1958), *The Psychology of Learning,* New York: McGraw-Hill.

DellaVigna, Stefano, and Ulrike Malmendier (2004), "Contract design and self-control: theory and evidence," *Quarterly Journal of Economics* 119(2), 353–402.

Detterman, Douglas K. (1993). "The case for the prosecution: Transfer as an epiphenomenon," in D. Detterman & R. Sternberg (Eds.), *Transfer on Trial: Intelligence, Cognition, and Instruction,* Norwood, NJ: Ablex, pp. 1–24.

Duflo, Esther, Michael Kremer, and Jonathan Robinson (2010), "Nudging farmers to use fertilizer: theory and experimental evidence from Kenya," NBER working paper no. 15131.

Dufwenberg, Martin, Ramya Sundaram, and David J. Butler (2010), "Epiphany in the game of 21," *Journal of Economic Behavior and Organization* 75(2), 132–143.

Fryer, Roland G., Jr. (2011), "Financial incentives and student achievement: evidence from randomized trials," *Quarterly Journal of Economics* 126(4), 1755–1798.

Gneezy, Uri, Aldo Rustichini, and Alexander Vostroknutov (2010), "Experience and insight in the race game," *Journal of Economic Behavior and Organization* 75(2), 144–155.

Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo (2009), "Peer effects in the workplace: Evidence from random groupings in professional golf tournaments," *American Economic Journal: Applied Economics, American Economic Association*, 1(4), 34–68.

de Groot, Adriaan D. (1966), "Perception and memory versus thought: some old ideas and recent findings," in B. Kleinmuntz (Ed.), *Problem Solving: Research, Method, and Theory,* New York: Wiley & Sons, pp. 19-50.

Houser, Daniel, Michael Keane, and Kevin McCabe (2004), "Behavior in a dynamic decision problem: an analysis of experimental Bayesian type classification algorithm," *Econometrica* 72(3), 781–822.

Levitt, Steven D., John A. List, and Sally Sadoff (2011), "Checkmate: Exploring Backward. Induction among Chess Players," *American Economic Review* 101(2), 975–990.

McKelvey, Richard D., and Thomas R. Palfrey (1992), "An experimental study of the centipede game," *Econometrica* 60(4), 803–836.

Salomon, Gavriel, and David N. Perkins (1989), "Rocky roads to transfer: Rethinking mechanisms of a neglected phenomenon," *Educational Psychologist* 24(2), 113–142.

Thorndike, Edward L., and Robert S. Woodworth (1901), "The influence of improvement in one mental function upon the efficiency of other functions," *Psychological Review* 8, 247–261.

Voss, James F., Jeffrey Blais, Mary L. Means, Terry R. Greene, and Ellen Ahwesh (1986), "Informal reasoning and subject matter knowledge in the solving of economics problems by naive and novice individuals," *Cognition and Instruction* 3(4), 269–302.

APPENDIX: EXPERIMENTAL INSTRUCTIONS

LABORATORY SCRIPT:

Welcome to the UT Experimental Economics Laboratory. My name is ________, and joining me today is ________. We are researchers from the Department of Economics. We understand that many of you have busy schedules and really appreciate your willingness to participate.

Before we begin I need to go over a few lab rules. Once the experiment begins, please refrain from communicating with each other (talking, texting) and please do not open or play any games on the computer (no solitaire or minesweeper).

In this study, you will be asked to make a series of market-like decisions. Your earnings in this experiment are based on the decisions you make. The money you will be paid with comes from a research grant, and this money can only be used to pay experiment participants. You will be paid in cash after the experiment is completed.

The decision-making setting may be unfamiliar to you. This is common. Therefore, in writing the instructions for this experiment, we have done our very best to clearly describe to you all relevant information from which to base your decisions.

There are two important protocols in experimental economics that we would like you to be aware of. First, the instructions contain only true information. There are no hidden tasks, and the experiment works exactly as stated in the instructions. Second, your decisions are confidential. What this means is that you have been randomly assigned an ID number. All decisions you make will be associated with this ID number and not your name. Therefore, when we analyze the data and present results, your name will in no way be affiliated with this study.

We have provided everyone with a pencil, calculator, and paper. Use these items, if you wish, as you make your decisions. But please do not write on the instructions.

Has everyone had a chance to read the informed consent sheet? Is everyone comfortable with the risks involved with participation in this experiment? If you would, please raise your hand to indicate you have read the Informed Consent Sheet and you agree to participate in the experiment.

Today, you will play two experimental games and then answer a short questionnaire. We will proceed by reading the instructions for the first game. I will read the instructions aloud and ask that you follow along on your copy.

Let's begin…<read instructions for Race to 21 game>

After you are finished with the final round of the first game, please DO NOT proceed to the second game. We will go through the instructions to the second game as a group before proceeding.
If any question should arise during the experiment, please raise your hand and one of us will address your question privately. Good luck, and we hope you earn lots of money!

Race to 21 game instructions:

Welcome

Thank you for participating in this experiment. You will be playing games against a computer for money. You will be paid for your participation in cash, immediately at the end of the experiment. How much you earn depends on your decisions. A research foundation has contributed the money for this study.

It is very important that you read all instructions carefully and that you strictly follow the rules of this experiment. If you disobey the rules, you will be asked to leave the experiment. You should never use the browser's forward or back button. All navigation through this experiment should be done by hitting the "proceed" button on each screen.

Your id number for this experiment is ________.

Do not press the PROCEED button until instructed to do so.

Rules of the Game

You will be playing a simple game against a computer opponent. You will see 21 stones, numbered 1 through 21. You will move first. On each of your turns, you may remove between 1 and <3, 4> consecutive stones beginning with the lowest-numbered remaining stone. You will remove stones by mousing over the stones you want to remove and clicking on the last stone you want to remove. After your turn, the computer will remove between 1 and <3, 4> consecutive stones. The stones removed by the computer will flash briefly, and then it will be your turn again. You and the computer alternate moves until all of the stones have been removed.

Profit

Most of the stones in the game are colored gray but one or more may be colored green. You make money when you, and not the computer, click on a green stone. Green stones have boldface numbers equal to the number of points that you earn for removing that stone. At the end of the experiment, points will be converted into dollars at the rate of $1.00 for every 150 points. If you remove a green stone in passing but without clicking on it, you do NOT earn points on that turn.

Practice

On your screen is a panel to help you practice removing stones. In the practice panel, there are only 8 stones. The last stone is a green stone and has a value of 100 points. You may remove between 1 and <3, 4> on each turn. Take this opportunity to try removing different numbers of stones. Notice that moving your mouse over a stone highlights it and the remaining stones with lower values, but that if you try to take more than <3, 4> stones, nothing happens. If you would like more tries, hit the RESET button and continue the practice round.

To sum up, you will play a game against a computer opponent in which you and the computer will alternate removing between 1 and 4 stones until all 8 stones have been removed. If, on your turn, you click on a green stone, you will earn the amount written on that stone.
You will play 30 rounds of this game. Then, you will play 15 rounds of a different game.
Do not press the "Proceed" button until instructed to do so.

Game 2

You have finished the first game. You will now play 15 rounds of game 2.

In game 2, there will be a total of 21 stones. You and the computer will alternate removing between 1 and <3, 4> stones on each turn. You will go first.

Practice

If you wish to practice removing stones, the practice panel on your screen contains 8 stones. You may remove between 1 and <3, 4> on each turn. If you would like more tries, hit the RESET button and continue the practice round.

Do not press the "Proceed" button until instructed to do so.

Survey

Once you have finished the second game, you will complete a brief survey.
These questions will be used for statistical purposes only.

THIS INFORMATION WILL BE KEPT STRICTLY CONFIDENTIAL and WILL BE DESTROYED UPON COMPLETION OF THE STUDY.