

University of Connecticut

Department of Economics Working Paper Series

Nudging Backward Induction

by

William Neilson
University of Tennessee, Knoxville

Michael Price
Georgia State University

Mikhael Shor
University of Connecticut

Working Paper 2016-31
October 2016

365 Fairfield Way, Unit 1063
Storrs, CT 06269-1063
Phone: (860) 486-3022
Fax: (860) 486-4463
http://www.econ.uconn.edu/

This working paper is indexed in RePEc, http://repec.org

Nudging Backward Induction

By WILLIAM NEILSON, MICHAEL PRICE, AND MIKHAEL SHOR[*]

October 2016

A growing literature shows that interim incentives can help people achieve favorable long-term outcomes. We design a laboratory experiment to explore how interim incentives impact learning backward induction using a race game: Subjects play a simple game against a computer in which winning requires a sequence of correct moves. Our data highlight a perverse effect of incentives. Interim rewards inserted along the optimal path—nudges—help subjects learn patterns but crowd-out the skills required to solve a related game. Interim payments off the optimal path—teasers —can either help or hurt learning.

JEL Codes: C91, D03, D60

[*] Neilson: Department of Economics, University of Tennessee, Knoxville, TN 27996-0550 (e-mail: wneilson@utk.edu). Price: Department of Economics, Georgia State University, P.O. Box 3992, Atlanta, GA 30302-3992 (e-mail: mprice25@gsu.edu). Shor: Department of Economics, University of Connecticut, 365 Fairfield Way, U-1063, Storrs, CT 06269-1063 (e-mail: mike.shor@uconn.edu). We thank audiences at the European Economic Association conference in Oslo and the Game Theory Society conference in Istanbul, as well as seminar participants at George Mason University, Indiana University, the University of Arkansas, the University of Connecticut, and the University of Tennessee. This research was supported in part by the United States Air Force under contract number FA7014-06-D-00190.

1. Introduction

This paper reports a laboratory experiment designed to determine whether nudges help subjects learn a skill important to both individual and strategic decision-making: backward induction.[1] We use the race game (Dufwenberg, Sundaram, and Butler 2010; Gneezy, Rustichini, and Vostroknutov 2010; Levitt, List, and Sadoff 2010) played against a computer. In this game, the subject can only win by playing a specific, unique sequence of actions which can be found using backward induction. Typically in the race game, the only payment goes to the winner at the end of the game, but our treatments explore whether subjects respond to smaller interim payments that are on, or off, the winning path. Interim payments on the winning path can be thought of as nudges to promote learning, while payments off the winning path can be thought of as teasers to deter learning.

The rules of the game have two players alternate moving along a path with 21 steps. During her turn, a player must move at least one step forward, but can move no more than three steps at a time. The player who ends the turn on step 21 wins a prize, and her opponent earns nothing. Player 1 can win this game by ending her turns on steps 1, 5, 9, 13, 17, and 21, a path that can be found through a process akin to backward induction. For example, if she ends a turn on 17, player 2 must end a turn on either 18, 19, of 20. Player 1 can reach 21 from any of those three places, and so ending a turn on 17 guarantees a win on her next turn. To guarantee being able to end a turn on 17, she should end a turn on 13, and so on.

There are a number of reasons why the race game against a computer is ideally suited for exploring incentives for and against learning to backward induct.[2] First,

[1] The classic studies of backward induction involve the centipede game (e.g. McKelvey and Palfrey, 1992; Fey, McKelvey, and Palfrey, 1996), a two-player game that should unravel to the beginning.

[2] Other studies have used the race game to explore other topics related to backward induction. Cardella (2012) uses a race game to compare learning-by-doing with learning-by-observing; Carpenter, Graham, and Wolf (2013) show that higher cognitive ability improves performance in the race game; Brosig-Koch, Heinrich, and Helbach (2014) find that teams outperform individuals in the backward induction task; and Brosig-Koch, Heinrich, and Helbach (2015) find that the ability to backward induct increases with children's age;

Dufwenberg, Sundaram, and Butler (2010) and Gneezy, Rustichini, and Vostroknutov (2010) show that solving the game involves insight and a "eureka" moment in which the method of solving the problem suddenly appears to the subject, and he or she can continue to solve the problem successfully once this moment arises.[3] Because of this, learning is rather simple to detect. Second, play against a computer programmed to win means that subjects are highly unlikely to win by pure luck and cannot win through the suboptimal play of an opponent. As such, a player's first win can most likely be attributed to learning as opposed to a myriad of other factors. Third, the uniqueness of the optimal path makes it straightforward to designate interim payments as either positive nudges or negative teasers based upon their placement. That is, a teaser involves playing off the optimal path and thus sacrifices a win against our computer opponent. Conversely, against a human opponent, one could pursue an off-path teaser and still win if the opponent subsequently makes a mistake.

Psychologists have long drawn a distinction between two types of learning (Thorndike and Woodworth 1901; Salomon and Perkins 1989). First, reflexive learning, or the instinctive learning of patterns, can only be applied to the task at hand or nearly identical contexts and does not transfer to more distantly-related contexts. Second, mindful learning, or a deliberate cognitive search for abstract general principles, readily transfers learning to a broader array of contexts. Correspondingly, we allow for epiphanies to take on two distinct forms: non-transferable epiphanies, which involve learning the pattern that wins a game but not the general principle, and transferable epiphanies, through which subjects develop a sufficiently general understanding of backward induction to solve a class of related games.

[3] Eureka moments are not isolated to race games. McKinney and Van Huyck (2013) find them in the game of Nim. That game, though, has no straightforward way to add interim payments.

To distinguish between these types of learning, we employ a within subject experimental design whereby subjects play two distinct versions of the race game. After playing the same race game 30 times against a computer, subjects move on to play a second race game, with a different optimal path, for another 15 periods. In particular, we consider a second version of the race game in which players can move as many as four steps in a turn. This variant of the game has a different optimal path: 1, 6, 11, 16, 21, which corresponds to the three-step optimal path only on the first and last turns.

If a subject truly learns to backward induct in the three-step game, she should be able to figure out the four-step game fairly quickly and vice versa. On the other hand, if she only learns the winning pattern and not backward induction, it should take time before she learns the other version of the game. We classify a subject as a *learner* if she eventually learns to win game 1 and as a *non-learner* otherwise. Among *learners*, we classify a subject as a *transferable learner* if she can readily win in the early rounds of game 2, and a *non-transferable learner* otherwise.

We provide a theoretical framework to guide the empirical analysis and provide a lens for interpreting the data. A fundamental prediction of our framework is that non-transferable learning crowds out transferable learning. We use this prediction to determine whether and how incentives influence learning. Specifically, treatments that induce non-transferable learning should yield less transferable learning and vice versa. This suggests a simple way to test the predictions of our framework—examine the correlation between types of learning across treatments.

Our experiment is designed to identify how the presence and location of interim payments impact learning. In our baseline game, subjects and the computerized opponent alternate choosing integers between 1 and s, which are summed until the sum reaches a target number, $n = 21$. The player who raises this sum to n wins a fixed monetary payment. Treatments augment the standard race game by adding a smaller interim payment for a player who, on her turn, brings the sum to $T < n$.

Treatments further vary both the location of the interim payment and the feasible action space.

The design allows us to nudge subjects onto the optimal path of the first race game by placing a small interim payment somewhere along the optimal path. Likewise, it allows us to tease them away from the optimal path by offering a small payment if a turn ends at a step that precludes them from winning. We look to see if nudges help and teasers hurt learning the first game, as measured by the number of subjects who win at least two of the final three rounds of the first game. Among those learners, we then determine whether learning is transferable or non-transferable by seeing if they are able to win the second game, with a different optimal path, quickly.

We find that nudges help subjects learn, but that learning is largely non-transferable. On the other hand, teasers make it less likely that subjects learn the first game, but what learning there is tends to be transferable. In all, the treatments show the negative correlation between types of learning predicted by the theory.

Our paper adds to the literature on nudges in three ways. First, most experimental studies of nudges involve field experiments promoting positive behavioral changes, often in health, education, environmental, or development economics. For example, Fryer (2011) shows that paying students to read improves subsequent classroom performance, though paying directly for test performance has no discernible impact. Similarly, Ashraf, Karlan, and Yin (2006) show that voluntary acceptance of withdrawal restrictions until attainment of a certain balance has a positive and long-lasting effect on individual rates of saving. And, in the context of healthcare, Charness and Gneezy (2009) find that paying people to go to the gym induces higher exercise levels that persist after the payments end. In contrast, our paper studies the ability of nudges to promote a beneficial behavior, learning backward induction, in a *laboratory* experiment.

Second, our paper provides a potential explanation for why the effects of many nudges tend to wane over time (see, e.g., Volpp et al., 2009; Charness and Gneezy, 2009; Ferraro and Price, 2013; Allcott and Rogers, 2014). To the extent that such nudges promote non-transferrable learning, we would expect long-run impacts to be less pronounced than those observed during the intervention period.

Finally, papers exploring nudges tend to look at ways to influence beneficial behavioral changes, and we introduce the idea of negative nudges intended to move subjects toward detrimental behavior. Such negative nudges are prevalent in the real world, as evidenced by the fact that mortgage lenders, credit card companies, satellite TV providers, and health clubs offer promotional "teaser" rates designed to lure customers into contracts that may reduce welfare in the long run. While such "teasers" have been shown to adversely impact consumer well-being in the short run (Shui and Ausubel, 2005; DellaVigna and Malmendier, 2006), our results suggest that the impacts in the long run may be compounded as those who learn to avoid "teasers" are less likely to transfer this learning and make better decisions in related environments.

Our paper also adds to the literature on transferability of play across games. Grimm and Mengel (2012), Haruvy and Stahl (2012), and Mengel and Sciubba (2014) all find evidence of learning across games, which would be considered transferable learning in our parlance. Closer to ours, Rick and Weber (2010) look at *p*-beauty contests and find that learning is more transferable in sessions where players do not receive feedback than in sessions where they do.[4] Both their no-feedback game and our off-path teaser games make subjects think more deliberately about how to win than their feedback game and our on-path nudge game.

[4] Huck et al. (2011) explore a related question–how the accessibility and coarseness of feedback impacts learning across games. Results from this study highlight that spillovers are more pronounced when feedback is not easily accessible.

Combining this evidence suggests that deliberate thinking about games involving inductive reasoning increases transferability.[5] Further, we add to this general literature by using nudges to manipulate the transferability of play across games.

We report two experiments based on two different design choices for the transition from the first game played 30 times by subjects to the second game played 15 times. Section 2 concerns the first experiment in which the first game played by subjects has either no interim payment, an interim payment on the optimal path, or an interim payment off the optimal path, and one of two action spaces, either three spaces per turn or four spaces per turn. The second game leaves everything the same, including the interim payment placement, but switches to the other action space. This moves on-path payments off of the optimal path, and vice versa. While this design choice follows the usual practice of making only one change to the design at a time, it entails detecting the transferability of learning with different games for different subjects, and some of those games may be easier to learn than others. Experiment 2, reported in section 3, addresses this concern by having all subjects play the same second game. Section 4 offers a discussion and some brief concluding remarks.

2. Experiment 1

A. Design

Our first experiment is designed to identify how the presence and location of interim payments impact learning in a relatively simple task. The experiment is based on variants of the race game. In the race game, two players alternate choosing integers between 1 and *s,* which are summed until the sum reaches a target number, *n*. The player who raises this sum to *n* wins a monetary payment.

[5] Ding and Nicklisch (2013) find that subjects learn more from negative feedback than positive feedback in a simultaneous-move game.

To this standard race game we add a smaller interim payment for a player who, on her turn, brings the sum to $T < n$. We fix the length of the game, n, at 21, and compare performance across games with different interim payment locations and different action spaces. Specifically, we consider four locations for the interim payment, T = 5, 6, 11, or 13, and two action spaces, s = 3 or 4, reflecting the maximum integer that can be chosen each move.

Within this set-up, a decision problem, D, can be described by the pair (s,T). For example, the decision problem (3,13) corresponds to the race-to-21 game with an action space of {1,2,3} and an interim payment on T = 13. The decision problem (3,•) corresponds to the same game but without an interim payment.

The race game without interim payments is dominance solvable and the optimal path can be determined using logic akin to backward induction. For any game with action space s, a player should end the k^{th}-to-last turn by selecting a number given by $21 - k(s + 1)$, when available. The introduction of interim payments makes this no longer true. An off-path interim payment may be sensible to pursue if I believe that the opponent will make a subsequent mistake, allowing me both to win and to receive the interim payment. Thus, moves off the optimal path can be due either to error or to hyper-rational play given uncharitable beliefs about the opponent's rationality.

To eliminate this indeterminacy, subjects play against computer opponents. Unbeknownst to subjects, the computerized opponent was programmed to punish any deviation from the optimal path. In any such instance, the computer's next move placed it on the optimal path where it remained until removing the last stone and winning. Specifically, the computer follows the optimal strategy when possible (following an error by the subject), takes the interim payment when this does not conflict with following the optimal path, and otherwise moves randomly.

We implement the race game on a computer by presenting subjects with 21 numbered stones (Figure 1). The subject moves first by clicking on one of the first

s stones, removing it and all lower-numbered stones. The computer follows, removing between 1 and *s* of the remaining lowest-numbered stones. Players alternate turns in this manner. The last stone is colored green and labeled as being worth 100 "points" (the experimental currency) for the player who ends her turn by removing this stone. Interim payments are presented as an additional green stone worth 50 points.[6] Figure 1 represents the (3,13) game, with an action space of 3 and the interim payment on stone 13.

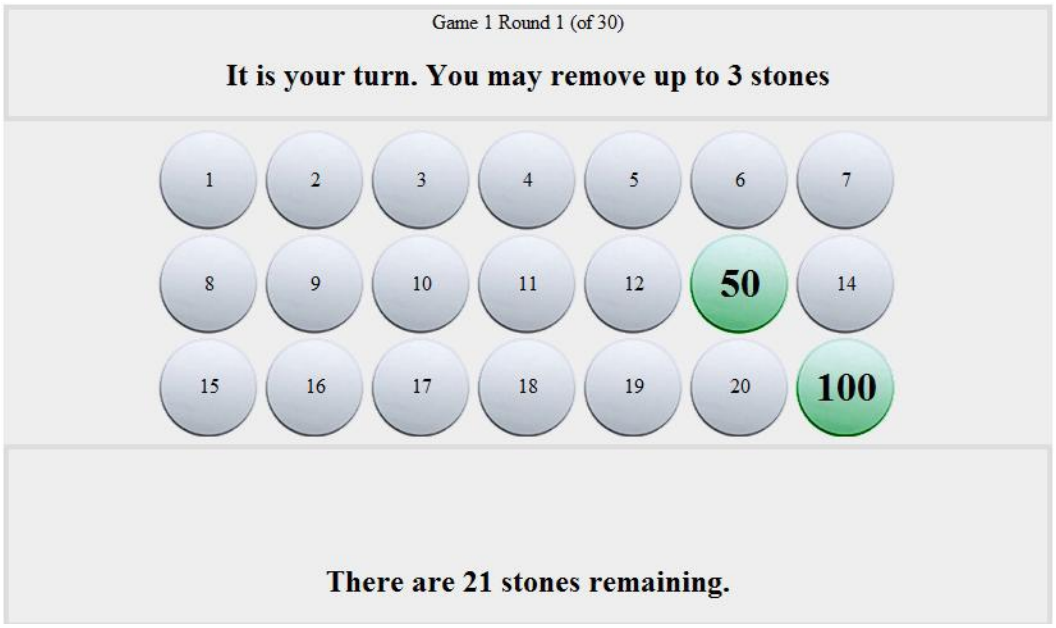


FIGURE 1. SCREEN SHOT OF RACE GAME INTERFACE

Subjects first play the race game (*s*,*T*) for 30 rounds (game 1). Next, subjects play 15 rounds of a race game (*s'*,*T*) with the same interim payment location but a different action space (game 2). As the two games have different optimal paths, this

[6] The player must stop on the interim payment stone to collect the interim payment. For example, if the interim payment is on stone 11 and stones 1-9 have already been removed, then removing stones 10 and 11 on the next turn yields the interim payment but removing stones 10, 11, and 12 (thereby passing through stone 11) does not.

allows us to identify whether learning transfers from the first game to the second game. Further, any interim payment that is on the optimal path in game 1 is off the optimal path in game 2 and vice versa.

We consider 10 distinct treatments that cross two action spaces, $s \in \{3, 4\}$, with interim payments located at five different locations $T \in \{\bullet, 5, 6, 11, 13\}$. The different treatments are shown in Table 1. Treatments are labeled according to the problem faced in game 1, or the first 30 rounds, and differ along three dimensions: (i) the presence or absence of an interim payment, (ii) whether the interim payment is located on or off the optimal path, and (iii) whether the interim payment is located early or late on the path. Earning an on-path interim payment is a necessary condition to winning the game. Earning an off-path interim payment precludes

TABLE 1. EXPERIMENT 1 TREATMENTS AND NUMBER OF SUBJECTS

Treatment (Game 1)	Interim Payment Location	Action Space		N
		Game 1	Game 2	
On-path interim payment				
Early	5	3	4	22
	6	4	3	19
Late	11	4	3	23
	13	3	4	20
				84
Off-path interim payment				
Early	5 (unattainable)	4	3	19
	6	3	4	23
Late	11	3	4	25
	13	4	3	24
				91
No interim payment (baseline)				
		3	4	24
		4	3	21
				45

winning the game as it cedes the optimal path to the computer. Treatments without interim payments correspond to a standard race game and serve as our baseline.

One treatment deserves special mention. In game (4,5), the early off-path treatment with action space 4, the interim payment is unattainable for the subject. If a subject were to play optimally and remove a single stone on her first move, the computer would take the interim payment. If a subject were to remove more than one stone on her initial move, the computer would seize the optimal path, skipping over the interim payment to remove through the sixth stone. We label this treatment as the unattainable interim payment and all other off-path treatments as attainable.

All sessions were conducted 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 experiment was programmed in Perl 5.

At the start of each session, subjects were provided written instructions for the game and asked to follow along as the experimenter read them aloud (see Appendix). The same instructions were also visible on subjects' computer screens. Instructions described the game and detailed how to remove stones using the mouse. Importantly, the instructions outlined how subjects earned points and highlighted that points were *only* earned by ending a turn on a green stone.

Subjects 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 questionnaire and were paid their earnings for all 45 rounds privately in cash. Points were converted at the rate of 150 points to a dollar. A total of 220 subjects participated in Experiment 1 and earned an average of $18.48 for sessions that lasted about an hour.

B. Theory & Hypotheses

Dufwenberg, Sundaram, and Butler (2010) provide evidence from a laboratory experiment that experience from playing a short race game (of length 6) can transfer to a subsequent more complex, longer race game. They reason that the simpler race game makes it clear that the game has an analytical solution which subsequently facilitates a search for that solution. Gneezy, Rustichini, and Vostroknutov (2010) argue that such insights can be gained in longer games directly by focusing attention on simpler sub-games embedded within them. Our on-path interim payment treatments fit well with both explanations. A subject who focuses first on the interim-payment effectively faces a shorter game. Thus, we hypothesize that on-path interim payments will increase the likelihood of winning a race game. Conversely, off-path interim payments teach precisely the wrong patterns and, in life, are generally designed to discourage people from following optimal paths. This leads to a first hypothesis:

Hypothesis 1. On-path interim payments will increase learning and off-path interim payments will decrease learning.

Hypothesis 1 focuses on the how interim payments impact learning within a given decision task. Yet, a primary goal of our experiment is to deduce not only on *whether* subjects learn but also *what* they learn and if such learning is transferable to a different environment.

For a given race game, (s,T), denote by $f(s,T)$ its unique, optimal path. We denote as *learners* those who are able to figure out this optimal path. We consider two ways that subjects can learn and distinguish these forms based on whether they learn a game-specific pattern $f(s,T)$ or the more general mapping f that provides a solution to all such games. We classify those who learn only the specific pattern $f(s, T)$ as *non-transferable learners*. Faced with a different game $f(s',T)$, they will

have to learn the correct pattern anew. We classify subjects who learn the mapping f and can apply this knowledge to solve related games as *transferable learners*. Of course, some subjects may fail to figure out the optimal path $f(s,T)$ in the time allotted. We label such subjects as *non-learners*.

To inform the construction of a testable hypothesis, we offer a model of learning consistent with the notion of epiphanies put forth in Dufwenberg, Sundaram, and Butler (2010) and that of insight as found in Gneezy, Rustichini, and Vostroknutov (2010). A subject confronts the same game in a series of rounds $1,\ldots,R$. The subject learns how to solve the problem by having an epiphany. We treat epiphanies as random events each round. To allow for a distinction between non-transferable and transferable learning, we allow for two types of epiphanies. A *non-transferable epiphany* involves learning the specific pattern $f(s,T)$ which may not transfer when the game changes. A *transferable epiphany* involves learning the entire mapping f and therefore empowers the subject to transfer what is learned to a new game.

Let p^t be the probability that the subject has a transferable epiphany in round r conditional on not having had an epiphany of either type through round r-1. Similarly, let p^n be the probability of a non-transferable epiphany in round r conditional on no prior epiphanies of either type. We assume that $p^n + p^t \leq 1$ and that p^n and p^t are nonnegative and independent of r. Our main assumption is that a subject who has who has one type of epiphany does not later have the other. In particular, a subject who has a non-transferable epiphany about a game-specific pattern does not pursue a deeper understanding of the game. This may follow from a rational person who, already having maximized payoffs by learning a pattern that leads to winning, sees no pecuniary benefit from further deliberation. Alternatively, this may follow from cognitive biases by which success leads one to overestimate the appropriateness of the approach (e.g., Baron and Hershey 1988) and therefore confuse a non-transferable epiphany for a transferable one.

With these preliminaries in mind, consider a policy-maker who is interested in three measures: the rate of overall learning, the rate of transferable learning, and the conditional probability that learning is transferable. The rate of overall learning (brought about through either type of epiphany) foretells success in the current environment as long as the optimal path does not change. The rate of transferable learning matters particularly when the environment changes over time requiring the mastery of new optimal paths. The conditional probability that learning is transferable is relevant for evaluating environments where people face interim payments. It provides a way of measuring if a given incentive makes learners more or less adaptable to changes in the environment. It also provides a way of measuring whether short-term success is predictive of future success in different decision tasks. 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^t of a transferable epiphany rises, ceteris paribus, the probability of learning the game by round r increases, the probability of becoming a transferable learner by round r increases, and the conditional probability that learning is transferable increases.

(b) When the probability p^n of a non-transferable epiphany rises, ceteris paribus, the probability of learning the game by round r increases, the probability of becoming a transferable learner by round r decreases, and the conditional probability that learning is transferable decreases.

Proof. In appendix.

Part (a) of the proposition examines the effect of an increase in the probability of transferable epiphanies. The overall learning rate, the transferable learning rate, and the conditional likelihood that learning is transferable all increase. An increase in

the probability of a transferable epiphany increases overall learning and reduces the proportion of subjects who have non-transferable epiphanies. Part (b) looks instead at the effect of an increase in nontransferable epiphanies. While 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 can crowd out the other.

Intuitively, when the probability of a non-transferable epiphany increases, it has no effect on the single-round probability that a subject has a transferable epiphany. However, because it is now more likely that subjects have had non-transferable epiphanies in the past, the pool of subjects who have not yet learned shrinks. Hence, fewer subjects can have transferable epiphanies in subsequent rounds. By making the winning pattern easier to identify, fewer subjects pursue deeper insights, and thus the learning of patterns crowds out transferable epiphanies. This crowding-out phenomenon leads to a second hypothesis:

Hypothesis 2. Transferable and non-transferable learning rates are negatively correlated.

We do not conjecture, a priori, whether on-path interim payments are more likely to increase transferable or non-transferable learning and therefore correspond to part (a) or part (b) of the proposition. Certainly, both possibilities are reasonable. The reasoning in Dufwenberg, Sundaram, and Butler (2010) and Gneezy, Rustichini, and Vostroknutov (2010) might suggest that the shorter sub-game involved in pursuing the interim payment can help a subject better understand the notion of backward induction and thus increase transferable learning. Conversely, interim payments may simply lead a person down the optimal path without realizing why. If one were to imagine a game with an interim payment for each correct move, this would likely serve to light the path and teach patterns rather than backward induction. Ultimately, this is an empirical question that depends on whether a given

incentive promotes transferable or non-transferable learning. However, our model predicts that one type of learning can crowd out the other. Therefore, the possibility arises that a particular incentive could prove beneficial to overall learning but detrimental to transferable learning.

C. Results

Game 1: Learning to Win

We begin by examining results from the first game to determine whether interim payments help or hinder winning. Table 2 shows the overall proportion of wins by treatment pooled across all subjects and all thirty rounds. When there is no interim payment, subjects follow the optimal path and win in 29% of the rounds. An on-path interim payment increases the win proportion to 52%, a highly significant difference using either a Mann-Whitney or t-test ($p < 0.001$).[7] This is consonant with prior work suggesting that paying people to go to the gym increases gym attendance (Charness and Gneezy, 2009) or that paying students to read improves subsequent class performance (Fryer, 2011). In contrast, off-path interim payments have little impact on performance. Such interim payments reduce the proportion of rounds won by a statistically insignificant 4 percentage points ($p>0.500$).

TABLE 2. SUMMARY OF OUTCOMES IN GAME 1

Interim payment	Win proportion	Interim payment proportion
On-path	52%	74%
Off-path	25%	58%[†]
None	29%	

[†]Excluding the unattainable interim payment treatment, game (4,5)

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

There is significant evidence of learning over time (Figure 2). In our baseline treatments without interim payments, subjects win by taking the last stone 11% of the time during the first 10 rounds but this proportion increases to 43% in the last 10 rounds. The proportion of subjects who win in the first 10 and last 10 rounds increases from 33% to 68% for the on-path treatments and from 9% to 42% for the off-path treatments. Although the proportion of subjects who win increases over time, Figure 2 highlights that the rate of learning is largely independent of treatment. The win proportion is significantly higher in the on-path treatments than either other treatment even when considering only the first, middle, or last ten rounds ($p < 0.010$) while the off-path treatments do not differ from the baseline treatments over any of these ten-round blocks ($p > 0.300$).

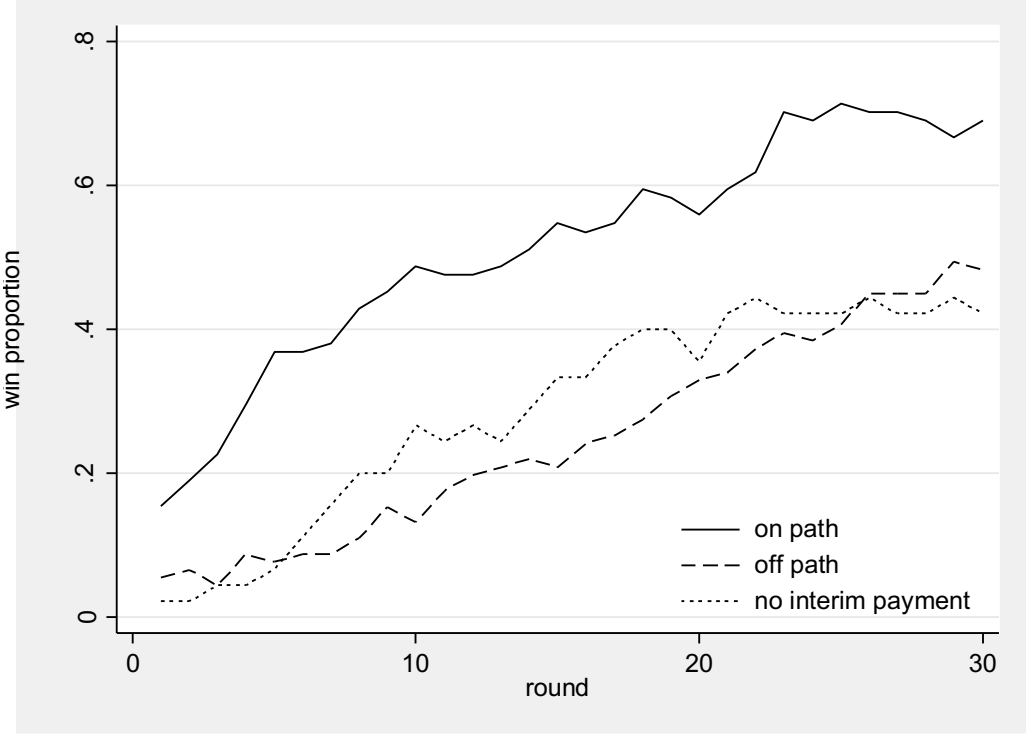


FIGURE 2. PROPORTION OF WINS BY TREATMENT FOR GAME 1 (FIRST 30 ROUNDS)

Exploring the data further, we find some evidence that off-path interim payments initially serve as a distraction. When it is attainable, 82% of subjects take the off-path interim payment in the very first round. Every subject in these treatments takes the off-path interim payment at least once by round 5. Furthermore, a substantial proportion of subjects do not appear to abandon off-path interim payments over time. In each of the last ten rounds, roughly half of all subjects take the off-path interim payment. However, off-path interim payments do not appear to preclude learning as they do not significantly reduce win rates relative to treatments with no interim payment. Rather, off-path interim payments may simply provide a consolation prize for subjects who are unlikely to have learned the optimal path in their absence.

To complement the unconditional results, we estimate a series of regression models. Specifically, we regress whether a subject wins in a particular round on the round of play, the presence and location of interim payments, and the action space (Table 3). We cluster standard errors by subject.

Empirical estimates highlight patterns consistent with the unconditional results. As noted in column (1), subjects in the on-path interim payment treatments are significantly more likely to win than are counterparts in other treatments. For example, in the final round of play, the estimated probability that a subject with an on-path interim payment removes the last stone is 23 percentage points higher than for counterparts in our baseline treatments. Further, win proportion is significantly increasing with time. The average subject is estimated to be eight times more likely to win in the thirtieth round of play than in the first.

TABLE 3. PROBIT ANALYSIS FOR DETERMINANTS OF WINNING GAME 1

	(1)	(2)	(3)
Interim payment is:			
on path	0.647***		
	(0.194)		
on path early		0.429	0.429
		(0.223)	(0.221)
on path late		0.866***	0.877***
		(0.219)	(0.219)
off path	-0.130		
	(0.197)		
off path early		0.120	-0.323
		(0.220)	(0.256)
off path late		-0.374	-0.370
		(0.234)	(0.233)
unattainable			0.887**
			(0.325)
Action space = 4	0.181	0.172	0.005
	(0.137)	(0.138)	(0.155)
Round #	0.051***	0.052***	0.052***
	(0.004)	(0.004)	(0.004)
Constant	-1.482***	-1.494***	-1.431***
	(0.187)	(0.188)	(0.189)
χ^2 p-value	<0.001	<0.001	<0.001

N = 6600. Probit estimates (Robust standard errors clustered by subject).
, * indicate significance at 1% and 0.1%.

Overall, we find partial support for Hypothesis 1. We summarize these data in our first result:

Result 1. Relative to the baseline treatment, (i) on-path interim payments increase win proportion, whereas (ii) off-path interim have no impact on win proportion.

That on-path payments lead to increased rates of winning shares similarity with Gneezy, Rustichini, and Vostroknutov (2010) who show that learning can be enhanced by focusing attention on sub-games embedded within a longer game. In our context, on-path payments serve to split the Race Game into two smaller subgames (i) a game to obtain the interim payment and (ii) a subsequent game from that location to the 21st stone.

Our next set of comparisons explores whether the location of interim payments (early or late on the decision path) and the action space (up to three or up to four stones per turn) impact win proportion. Table 4 summarizes win proportions broken out along these dimensions.

For on-path interim payments, locating them later on the decision path significantly increases the win proportion over locating them earlier (Mann-Whitney $p = 0.028$). Depending on the action space, subjects facing late on-path interim payments are 20% ($s = 3$) to 58% ($s = 4$) more likely to win than counterparts facing early on-path interim payments.

TABLE 4. SUMMARY OF OUTCOMES IN GAME 1 BY TREATMENT

	Win proportion
On-path	
Late $s = 3$	54%
Early $s = 3$	45%
Late $s = 4$	65%
Early $s = 4$	41%
Off-path	
Late $s = 3$	22%
Early $s = 3$	20%
Late $s = 4$	15%
Early $s = 4$	48%
None	
$s = 3$	28%
$s = 4$	30%

For off-path interim payments, we observe a more surprising result. The only interim payment location to significantly affect win proportion is the unattainable interim payment in game (4,5). Win proportion in this treatment is 60% higher than that observed in the corresponding baseline treatment ($p = 0.002$).

Taken jointly, these data suggest a second result:

Result 2. (i) On-path interim payments lead to a higher win proportion when located late on the decision path rather than early. (ii) Off-path interim payments lead to a higher win proportion when unobtainable than when obtainable.

Support for the first part of Result 2 is provided in the probit results shown in the second and third columns of Table 3. To illustrate the estimated parameter values, in the thirtieth round, subjects in the on-path late treatment are 18% more likely to win than their counterparts in the on-path early treatment ($p = 0.032$). Support for the second part of Result 2 is also shown in Table 3. While there is no significant difference in the parameter estimates for the off-path early and off-path late indicators, both are significantly different from the indicator for the unattainable treatment. For the final round of play, subjects facing an unattainable interim payment are 77% to 85% more likely to win than their counterparts in the off-path early and off-path late treatments, respectively.

Conversely, action space has little effect on win proportion (Tables 3 and 4). For example, without interim payments, there is a (statistically insignificant) 2 percentage point difference in win proportion between treatments with an action space of four stones and those with an action space of three stones. These data lead to a third result:

Result 3. Win proportions are independent of the action space.

Across all three model specifications in Table 3, the dummy for an action space of four is positive but statistically insignificant. This suggests that games requiring six correct moves to win are not significantly more difficult than games that require five correct moves. In this regard, though our games exhibit very small differences in length, our data are at odds with findings from McKelvey and Palfrey who show that the frequency of subgame perfect equilibrium play increases nearly 10-fold when moving from a six-move to a four-move version of the centipede game and with Dufwenberg, Sundaram, and Butler (2010) who show that shorter games are easier to learn than longer games.

Game 2: Transferability of Learning

Our results so far have focused on the first 30 rounds (game 1) and whether interim payments affect if subjects learn to win the game. We next consider how interim payments affect the nature of learning. We classify a subject as a *learner* if she wins game 1 in at least two of the final three rounds and as a *non-learner* otherwise. 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. Game 2 involves a change in the optimal path and thus allows us to further classify learners as transferable or non-transferable. Among *learners*, 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.[8]

Table 5 reports the proportion of subjects classified as learners and non-learners by treatment (pooling across the two action spaces). The on-path treatments have the highest proportion of subjects classified as learners (70%). In contrast, the

[8] This is a fairly strict definition of transferability of learning as it requires that one is able to deduce the solution to the new game immediately. In Experiment 2 below, we consider longer-term implications of learning.

TABLE 5. CLASSIFICATION OF SUBJECTS BY LEARNING TYPE

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

baseline treatments have the lowest proportion of subjects classified as learners (42%). Using a Fisher exact test, we find that the difference in the proportion of learners across these treatments is significant ($p = 0.003$). Among subjects in the off-path treatments, 47% are classified as learners, a proportion that is statistically indistinguishable from the baseline treatments ($p = 0.590$). Yet, among the off-path treatments, whether the payoff is attainable or unattainable results in a marked difference in the proportion of learners ($p < 0.011$).

Our baseline treatments without interim payments have the lowest proportion of non-transferable learning and the highest proportion of transferable learning. On-path treatments, in contrast, have the highest proportion of non-transferable learning and the lowest proportion of transferable learning. Off-path treatments lie between the two on both measures. Pooling the data into five categories (on-path early/late, off-path early/late, and no interim payment), the correlation between transferable and non-transferable learning is -0.88. Across all ten treatments, the correlation is -0.72 ($p = 0.019$). This suggests a fourth result:

Result 4. Transferable and non-transferable learning are negatively correlated.

Result 4 is consistent with the framework in Section 2B and provides support for Hypothesis 2. Our data suggest that non-transferable learning crowds out transferable learning. On-path interim payments increase overall learning, but do so by significantly reducing the proportion of subjects whose learning is transferable to new environments. This opens the possibility that incentives that are effective in the short run and enhance learning could prove detrimental in the long run. As noted in the first row of Table 5, 95% of learners in our baseline treatments without interim payments are able to transfer that knowledge immediately to the second game. This suggests that almost all learning in these treatments follows from transferable epiphanies about the nature of these games rather than non-transferable learning of patterns. Conversely, among the learners in the on-path treatments, only 19% are transferable learners. Transfer rates for games with off-path interim payments lie between these two extremes; just over half of the learners in such games are classified as transferable.

Exploring these differences a level deeper, the data in Table 5 highlight a correlation between the proportion of subjects denoted as learners and the proportion of transferable learners. For example, late on-path interim payments have the highest proportion of learning at 80% but the lowest proportion of transferable learning at only 7%. In contrast, early on-path interim payments have twice as many non-learners but, conditional on learning, three times the proportion of transferable learners as late on-path interim payments. We observe a similar pattern of correlation among off-path treatments. Games with higher rates of learning have lower rates of transfer and vice versa. Taken in its totality, our data suggest that when incentives lead to higher overall rates of learning, they appear also to lead to lower rates of learning transfer.

Compared to our baseline treatments, on-path interim payments increase the proportion of learners by nearly two thirds. Yet, they also decrease the proportion of subjects classified as transferable learners by more than two thirds. Off-path interim payments, in contrast, have little impact on the overall proportion of learners. However, such teasers do influence the type of learning. Our off-path treatments have a lower proportion of subjects classified as transferable learners. Taken jointly, these data suggest a fifth result:

Result 5: Interim payments promote non-transferable learning at the expense of transferable learning.

The observed data patterns suggest that on-path interim payments serve to both promote non-transferable learning among a subset of subjects who would otherwise prove unable to learn the game but also crowd out transferable learning among some subjects who would have become transferable learners in the absence of such payments. Thus, on-path interim payments present a tradeoff between better performance on the existing game and preparedness for future decision problems. Off-path interim payments, in contrast, provide no discernible benefit. In fact, such incentives do little more than shift a subset of subjects from being transferable to non-transferable learners without any offsetting increase in the overall rate of learning.

The qualitative results in Table 5 are robust to many alternative ways of measuring learning and transferability. As one robustness check, we can ask what each type of learner did in the *last* round 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 rounds of game 2 was not by luck or accident, but was driven by what subjects learned in game 1.

3. Experiment 2

In Experiment 1, subjects played 30 rounds of a particular race-to-21 game followed by 15 rounds of an identical game except with a change in action space from 3 to 4 or from 4 to 3. However, the location of the interim payment (if any) was held constant throughout all 45 rounds of play. Consequently, games with on-path interim payments were followed by off-path interim payments and vice versa.

This raises two issues for identifying transferability of learning. First, subjects for whom the first game had an on-path interim payment had to learn in the second game to avoid the interim payment. In contrast, subjects who first learned to avoid the interim payment in an off-path treatment had to learn to take the interim payment in the subsequent on-path treatment. These two learning processes might be different, skewing transferability rates in favor of one treatment or the other.[9]

Second, data from Table 5 suggest that on-path interim payment games are easier to learn than off-path games. Subjects who initially faced an on-path game were confronted with a more difficult second game than those who initially faced an off-path game. It is thus possible that performance in game 2 conflates learning from game 1 with the relative difficulty or ease of game 2. Experiment 2 addresses this issue by holding the second task constant across treatments.

A. Design

Experiment 2 consists of two additional treatments that were run at the University of Tennessee Experimental Economics Laboratory. In one treatment, 21 subjects first played 30 rounds of the off-path interim payment game (4,13), corresponding to an action space of 4 and an interim payment on the thirteenth

[9] For example, using animal experiments, Rescorla (2002) finds that the rate of acquiring an association between a stimulus and an outcome is greater than the rate of un-acquiring that association. In our context, this would serve to work against transfer in games with on-path payments and enhance rates of transfer in games with off-path incentives.

stone. This was followed by 15 rounds of the game (3,·) without an interim payment. In the other treatment, 24 subjects played 30 rounds of the on-path interim payment game (4,11) followed by the game (3,·). For comparison purposes, we draw upon data from a baseline treatment from Experiment 1 in which 21 subjects played 30 rounds of the game without interim payments (4,·) followed by 15 rounds of the game (3,·). Thus, in Experiment 2, all subjects face the same game 2. All other aspects of the experiment, including the interface and payoff parameters, were identical to those used in Experiment 1.

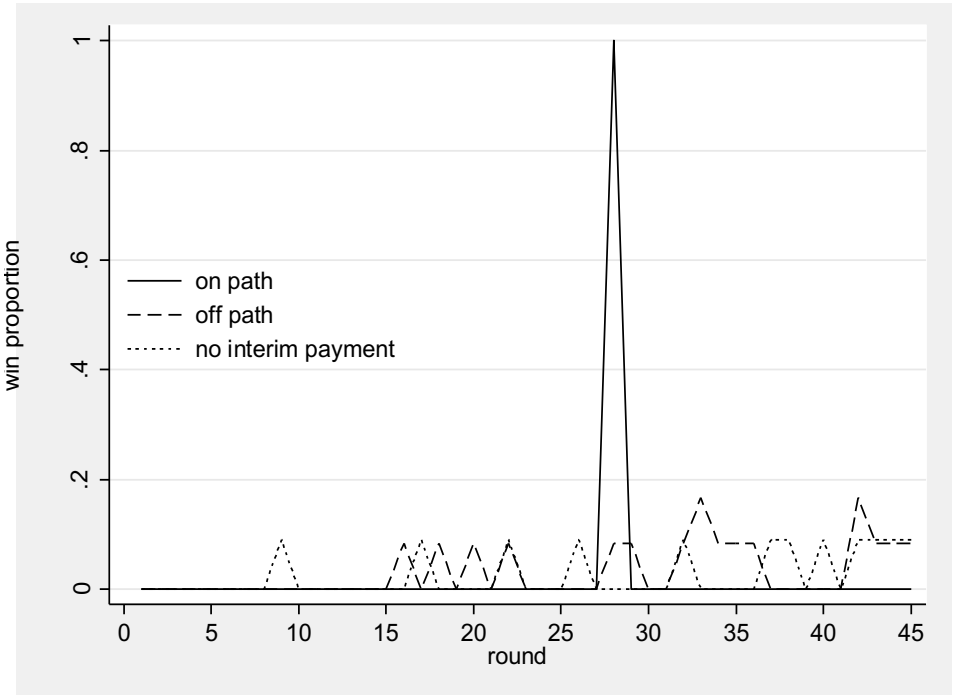
B. Results

Consistent with Experiment 1, all but one subject (96%) learn the on-path game over the first 30 rounds, while only 43% learn the off-path treatment. Figure 3 depicts our main result: without interim payments, learning is almost entirely transferable; with interim payments, much of the observed learning is not transferable. The horizontal axis shows the round number and the vertical axis measures the proportion of subjects who won by taking the last stone in that round for on-path, off-path, and no interim payment treatments.

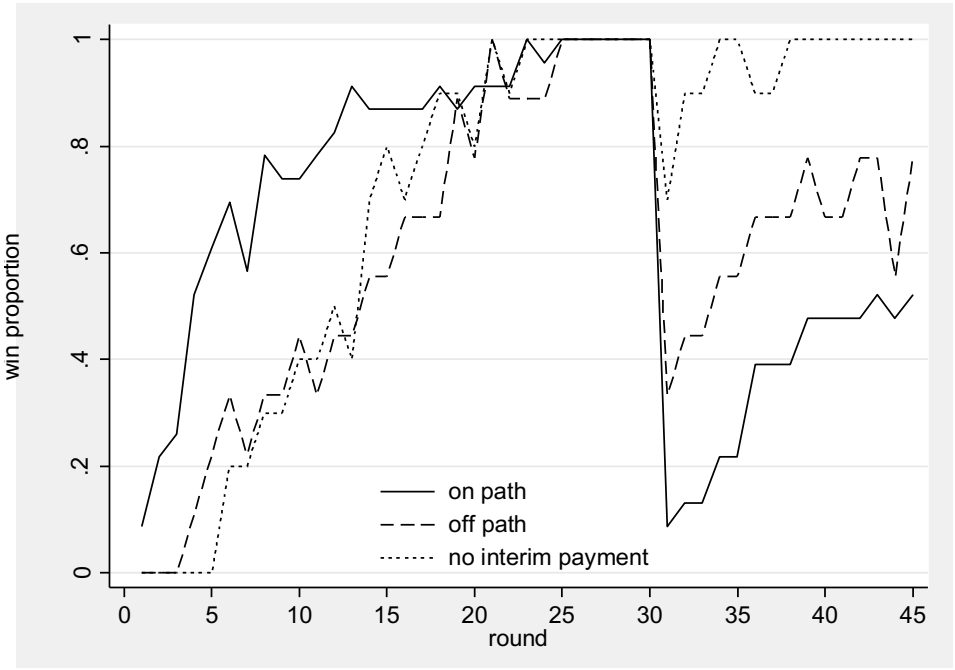
Panel A displays only those who did not learn game 1. Unsurprisingly, such subjects continue to perform poorly in the (3,·) game starting in round 31. Fewer than 10% of these subjects manage to win in any of the first three rounds. More striking is that failure to learn in the first game is a near-perfect predictor of failure to learn in the second; win proportions remain well below 20% over all 15 rounds.[10]

Panel B restricts attention to the subset of subjects who were identified as having learned game 1. Subjects in the on-path treatment appear to learn game 1 earlier. The curve depicting win proportions for the on-path treatment begins higher and

[10] The "blip" in Panel A of the figure represents the only subject who did not learn in the on-path treatment but nevertheless managed to win in one round, but not in subsequent rounds.



Panel A: Non-learners



Panel B: Learners

FIGURE 3. PERFORMANCE OF SUBJECTS BY LEARNING CLASSIFICATION

rises more quickly than those for the other two treatments. By definition, these differences are mitigated over the course of play and disappear by the 25th round at which point performance among learners is perfect.

As illustrated in Panel B, the results of Experiment 2 mirror those found in Experiment 1. Only 10% of learners in the on-path treatment are able to transfer this knowledge and win in at least two of the first three rounds of game 2. In contrast, approximately 40% of the learners from the off-path treatment and nearly 90% of those from the no interim payment treatment prove transferable.

Most notably, Figure 3 highlights that the treatment specific differences in learning game 2 persist through later rounds. Whereas *every* learner in the baseline treatment wins game 2 in the 45th round of play, less than 80% of learners from the off-path treatment and approximately half of those from the on-path treatment do the same. Thus, even after 15 rounds of game 2, subjects who previously experienced on-path interim payments are significantly less likely to win than subjects in the baseline treatment (Fisher exact test $p = 0.007$). Taken jointly, these data highlight that interim payments significantly change the nature of learning.

We next combine the results from our two experiments to analyze how the difficulty of game 1 impacts subsequent game 2 performance. Table 6 presents regression results examining the determinants of a subject's performance in game 2. The dependent variable is a subject's proportion of rounds won in the second game. Independent variables in model 1 include the subject's win proportion in game 1 and measures designed to capture the relative ease of the games played. We measure the ease of each of our ten games (five locations of interim payments and two action spaces) by the average win proportion over the first thirty rounds for all subjects who played that particular game, excluding oneself.[11] As one would

[11] 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 ($p = 0.233$) for game 1 difficulty, evidencing random allocation.

TABLE 6. DETERMINANTS OF GAME 2 WIN PROPORTION

	(1)	(2)
Game 2 ease	0.525**	0.490**
	(0.130)	(0.126)
Game 1 win proportion	0.870**	
	(0.053)	
Game 1 win rank		0.826***
		(0.047)
Game 1 ease	-0.833***	-0.234*
	(0.116)	(0.104)
Constant	0.210***	-0.050
	(0.080)	(0.080)
R^2	0.523	0.523
F	95.53	95.02

$N = 264$. OLS estimates (std. errors).
***, * indicates significance at 0.1% and 5%.

suspect harder second games would prove more difficult, we include a control for the difficulty of the second game: the proportion of the time that other subjects won the game when it was the first game played.

The coefficient of interest is the one on game 1 ease. Its sign indicates how the difficulty of the first game influences performance in the second game. Subjects won the second game more frequently when it was easy and if they were more successful in the first game. However, all else equal, subjects performed better in the second game if they faced a more difficult first game. Fixing the subject's own first-game performance, a 10 percentage point increase in the proportion of times others won that game leads to an 8.3 percentage point *reduction* in the proportion of wins in the second game. In other words, all else equal, a subject does better in game 2 if she faced a harder decision problem over the first thirty rounds.[12]

[12] The coefficients on game 1 win proportion 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.

To control for correlation between performance in game 1 and game 1's ease, we include a second specification. We use a subject's *relative* performance within an experimental treatment. Specifically, the variable *game 1 win rank,* taking on values between zero and one, is a subject's relative position when all subjects in the same experiment and treatment are ranked by number of wins in game 1 from lowest to highest. All subjects who never won in game 1 receive a rank of 0 and all other ties are assigned the average rank. We again confirm that making game 1 easier reduces performance in game 2 ($p = 0.026$).

The evidence in Table 2 suggested that locating interim payments on the optimal path makes the game more winnable and subjects more successful. However, the regression results in Table 6 suggest that this success comes at a price. Making the first game easier makes it less likely the subject will adapt to a changing environment. Hence, while on-path interim payments make winning more common, they also turn non-transferable learning from an anomaly into a majority outcome. These results suggest that making the optimal path easier to find may lead to shallower learning whereas making the task harder promotes a deeper understanding of the decision problem and how best to solve it.

4. Discussion and Conclusions

Sixty years ago, noted psychologist James Deese wrote that "There is no more important topic in the whole psychology of learning than transfer of learning." (Deese 1958, p. 213). With the marked failure of many to follow optimal paths to retirement, savings, and educational success, the sentiment has taken on significant economic import. Although there is a growing body of work suggesting that providing interim incentives (nudges) is an effective way to guide people towards their goals, we find that such incentives can provide an unnecessary crutch and introduce conflicting effects over different time horizons. Hence, the success of a nudge comes at a cost. Incentives that promote desired behavior in the short run may crowd out the insights required to adapt to and succeed in new environments.

Of course, people also face many "teasers" such as pricing schemes that capitalize on time-inconsistent preferences and attempt to discourage them off the optimal path. We find that off-path interim payments lead to unexpectedly few additional short-term failures over games without interim payments. However, teasers do serve to decrease the proportion of transferable learners. Yet, off-path interim payments nevertheless encourage substantially more mindful learning and arm people with a better understanding of the underlying decision problem than on-path nudges. A surprising finding relates to off-path interim payments that are not attainable, e.g., offers that sound—and indeed are—too good to be true. Games with such interim payments are associated with the highest rates of both short-term success and transferable learning.

In understanding the results of our experiment, we posit that interim payments effectively provide a point of salience that allows people to focus on shorter horizon problems. Consonant with prior work, our results highlight that focusing on games within a game help subjects identify optimal paths (Gneezy, Rustichini, and Vostroknutov, 2010) as shorter games spur epiphanies about backward induction

(Dufwenberg, Sundaram, and Butler, 2010). Yet, our results provide a caveat on the prior work. Epiphanies need to happen without significant guidance. Otherwise they fail to promote the insights required for transferability to new environments. While an unguided environment is not easy to master, its mastery carries with it a mindful understanding that readily translates to new environments.

Overall, our results suggest that efforts to guide good behavior, such as rewarding students for reading, adults for exercising, or people for saving, are likely to increase the desired outcome in the short term. However, such efforts may backfire and hinder the natural process of learning the more general principles underlying good study, lifestyle, and financial habits. Future work should explore the effect of such payments in a dynamically-changing field setting and examine in more detail why unachievable temptations create an environment conducive to learning.

References

Allcott, Hunt, and Todd Rogers (2014), "The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation," *American Economic Review* 104(10), 3003–3037.

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.

Baron, Jonathan, and John C. Hershey (1988), "Heuristics and biases in diagnostic reasoning: I. priors, error costs and test accuracy," *Organizational Behavior and Human Decision Processes* 41(2), 259–279.

Brosig-Koch, Jeannette, Timo Heinrich, and Christoph Helbach (2014), "Does truth win when teams reason strategically?" *Economics Letters* 123, 86–89.

Brosig-Koch, Jeannette, Timo Heinrich, and Christoph Helbach (2015), "Exploring the capability to reason backwards: An experimental study with children, adolescents, and young adults," *European Economic Review* 74, 286–302.

Cardella, Eric (2012), "Learning to make better strategic decisions," *Journal of Economic Behavior & Organization* 84, 382–392.

Carpenter, Jeffrey, Michael Graham, and Jesse Wolf (2013), "Cognitive ability and strategic sophistication," *Games and Economic Behavior* 80, 115-130.

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

DellaVigna, Stefano, and Ulrike Malmendier (2006), "Paying not to go to the gym," *American Economic Review* 96(3), 694–719.

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

Ding, Jieyao, and Andreas Nicklisch (2013), "On the impulse in impulse learning," *Economics Letters* 121(2), 294–297.

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.

Ferraro, Paul J., and Michael K. Price (2013), "Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment," *Review of Economics and Statistics* 95(1), 64–73.

Fey, Mark, Richard D. McKelvey, and Thomas R. Palfrey (1996), "An experimental study of constant-sum centipede games," *International Journal of Game Theory* 25, 269–287.

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.

Grimm, Veronika, and Friederike Mengel (2012), "An experiment on learning in a multiple games environment," *Journal of Economic Theory* 147, 2220–2259.

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* 1(4), 34–68.

Haruvy, Ernan, and Dale O. Stahl (2012), "Between-game rule learning in dissimilar symmetric normal-form games," *Games and Economic Behavior* 74, 208–221.

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.

McKinney, C. Nicholas, Jr., and John B. Van Huyck (2013), "Eureka learning: Heuristics and response time in perfect information games," *Games and Economic Behavior* 79, 223–232.

Mengel, Friederike, and Emanuela Sciubba (2014), "Extrapolation and structural similarity in games," *Economics Letters* 125, 381–385.

Rescorla, Robert A. (2002), "Comparison of the rates of associative change during acquisition and extinction," *Journal of Experimental Psychology: Animal Behavior Processes* 28(4), 406–415.

Rick, Scott, and Roberto A. Weber (2010), "Meaningful learning and transfer of learning in games played repeatedly without feedback," *Games and Economic Behavior* 68, 716–730.

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

Shui, Haiyan, and Lawrence M. Ausubel (2005), "Time inconsistency in the credit card market," University of Maryland working paper.

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.

Appendix: Proof

Proof of Proposition. The probability of learning the game by round r is the probability that the player has had either type of epiphany at some point through round r, given by

(1) $$P_r = 1 - (1 - p^n - p^t)^r.$$

A subject has a transferable epiphany in round r with probability $p^t\,[(1 - p^n - p^t)^{r-1}]$ where the part in brackets is the probability of no prior epiphany of either type. Therefore, the probability of becoming a transferable learner by round r is given by

(2) $$P_r^t = \sum_{j=1}^{r}(1 - p^n - p^t)^{j-1}\,p^t = [1 - (1 - p^n - p^t)^r]\frac{p^t}{p^t+p^n} = \frac{p^t}{p^t+p^n}P_r.$$

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

(3) $$P_r^t/P_r = \frac{p^t}{p^t+p^n}.$$

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

To prove part (b), note that $\partial P_r/\,\partial p^n = r(1 - p^n - p^t)^{r-1} \geq 0$ and that $\partial(P_r^t/P_r)/\,\partial p^n = -p^t/(p^n + p^t)^2 \leq 0$. It remains to show that P_r^t decreases when p^n increases. We do this by induction. Let P_r^t and P_r^{t*} denote the probabilities of a having a transferable epiphany by round r when the single-round probabilities of a non-transferable epiphany are p^n and p^{n*}, respectively, where $p^n > p^{n*}$. Similarly, let P_r and P_r^* denote the probabilities of having either type of epiphany by round r. For the first step of the induction process, $P_1^t = P_1^{t*} = p^t$ which is independent of p^n and p^{n*}. Now suppose that $P_{r-1}^t \leq P_{r-1}^{t*}$. The new transferable learners in round r are those who have not already learned and have a transferable epiphany in round r, yielding

$$P_r^{t*} = P_{r-1}^{t*} + p^t(1 - P_{r-1}^*) \text{ and } P_r^t = P_{r-1}^t + p^t(1 - P_{r-1}).$$

Subtracting yields $P_r^{t*} - P_r^t = (P_{r-1}^{t*} - P_{r-1}^t) + p^t(P_{r-1} - P_{r-1}^*)$ which is nonnegative because $P_{r-1}^t \leq P_{r-1}^{t*}$ by hypothesis and $P_{r-1}^* \leq P_{r-1}$ follows from the result that overall learning increases with p^n. □

For online publication

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.