

When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect

Gary Charness and Dan Levin*

October 24, 2003

Abstract: We examine decision-making under risk and uncertainty in a laboratory experiment. The heart of our design examines how one's propensity to use Bayes' rule is affected by whether this rule is aligned with reinforcement or clashes with it. In some cases, we create environments where Bayesian updating after a *successful* outcome should lead a decision-maker to make a change, while no change should be made after observing an *unsuccessful* outcome.

We observe striking patterns: When payoff reinforcement and Bayesian updating are aligned, nearly all people respond as expected. However, when these forces clash, around 50% of all decisions are inconsistent with Bayesian updating. While people tend to make costly initial choices that eliminate complexity in a subsequent decision, we find that complexity alone cannot explain our results. Finally, when a draw provides only information (and no payment), switching errors occur much less frequently, suggesting that the 'emotional reinforcement' (*affect*) induced by payments is a critical factor in deviations from Bayesian updating. There is considerable behavioral heterogeneity; we identify different types in the population and find that people who make 'switching errors' are more likely to have cross-period 'reinforcement' tendencies.

Key Words: Bayesian updating, Reinforcement; Affect; Experimental economics

JEL Classification: B49, C91, D81, D89

Acknowledgements: We thank Mark Brinkman and Arun Qamra for their help with designing the software and conducting the experimental sessions. We have benefited from discussions with, and suggestions from, George Akerlof, Ted Bergstrom, James Choi, Stefano dellaVigna, Ido Erev, Guillaume Frechette, Rod Garratt, Itzhak Gilboa, Botond Kozsegi, David Laibson, Jim Peck, Matthew Rabin, David Schmeidler, Bill Zame, Richard Zeckhauser, and seminar participants at UCSB, Columbia University, Harvard University, the Econometric Society Summer 2003 meeting, the Stony Brook 2003 Workshop on Experimental Economics and Game Theory, UC Berkeley, and The Ohio State University. All errors are our own. This version of the paper was completed while the second author was a visiting scholar in Harvard Business School and he wishes to thank HBS for its hospitality.

-
- Contact: Gary Charness, Department of Economics, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210, charness@econ.ucsb.edu, <http://www.econ.ucsb.edu/~charness>. Dan Levin, Department of Economics, The Ohio State University, 1945 N. High Street, Columbus, OH 43210-1172, dlevin@econ.ohio-state.edu, <http://www.econ.ohio-state.edu/levin/>

I. Introduction

Economists and psychologists have long been interested in how people make decisions under uncertainty. One particular issue concerns the manner in which people process new information and update prior beliefs. The Bayesian updating rule is ubiquitous in economic theory and its application. Bayesian updating is also the way we come to think about applying Expected Utility (EU) theory; together these concepts form the linchpin of the standard approach to decision-making under risk. Loosely speaking, Bayes' rule is a formula that instructs us how to update our *prior* beliefs about distribution functions once we observe realizations of random variables from the true distribution. The *posterior* beliefs formed represent a precise way of weighting the base belief with the 'recent evidence'.

There have been a number of experimental investigations in economics and psychology that study whether and to what extent people update information according to Bayes' rule. In three early papers,¹ Daniel Kahneman and Amos Tversky find evidence that people instead evaluate "the probability of an uncertain event ... by the degree to which it is similar in essential properties to its parent population." Grether (1980, 1992) tests this *representativeness* heuristic with the methodology of experimental economics, finding support for the view that individuals often ignore prior information when forming beliefs, contrary to Bayes' rule. More recent studies (e.g., Ouwensloot, Nijkamp & Rietveld, 1998; Zizzo, Stolarz-Fantino, Wen & Fantino, 2000) also provide strong evidence that at least experimental subjects are not 'perfect Bayesians' and often are not even very close.

Although we realize that *reinforcement* has specific meanings in the learning literature, we use the term in a very general sense, to simply reflect the idea that choices (actions)

¹ Tversky and Kahneman (1971, 1973) and Kahneman and Tversky (1972).

associated with successful outcomes are likely to be picked more often than choices associated with less successful outcomes.² There are several models of reinforcement learning in both the economics and psychology literature; these models successfully capture much observed behavior in laboratory experiments, using only a modest number of parameters. The basic reinforcement models (e.g., Roth and Erev 1995; Erev and Roth 1998) assume an initial propensity for a particular choice and utilize a payoff sensitivity parameter. Camerer and Ho (1998, 1999) combine reinforcement and belief learning (e.g., Cheung and Friedman 1997; Fudenberg and Levine 1998) by using experience-weights and updated levels of attraction.³

Expected utility theory and reinforcement models often prescribe a similar course of action, and both can potentially explain much observed behavior in the field and in the laboratory. We devise an experiment that successfully discriminates between these two approaches: At some nodes, predictions for a risk-neutral DM who uses Bayesian updating and maximizes expected utility (from here on we denote such DM by BEU) leads to one continuation, while reinforcement and case-based decision theory lead to another continuation. To the best of our knowledge, this is the first study to explicitly examine what happens when these forces clash.⁴ Our constructed case where the reinforcement heuristic leads one astray can be applied more generally to situations where favorable direct information about one choice may be indirectly even more favorable for an alternative choice.

As an example, imagine an agency having to choose between two salespersons. The first is young and inexperienced, with a wage of \$15 per hour. The second is a veteran, experienced

² Reinforcement is by nature a descriptive theory, while Bayesian updating is normative (of course, to the extent that people are ‘believers’ in Bayesian updating, it may also become descriptive).

³ Case-based decision theory (Gilboa and Schmeidler 1995, 2001) formalizes the thrust of reinforcement heuristics in a non-expected utility framework wherein people follow a decision rule that chooses an act with the highest relative score, based on performance in past cases and the similarity of those cases to the current decision case.

⁴ Note that our set-up is quite different from the “two-armed bandit” problem, where the separate machines have independent distributions rather than a common state.

salesperson with a wage of \$30 per hour. Suppose further that the agency decided to send the rookie, and that after a week this person's campaign has been quite successful. The agency must now decide whom to send out in the second week. The first reaction might be: "Well the young person did so much better than expected that we ought to send him/her again." However, upon reflection, it may occur to the agency that the unexpected successes may also contain information that the relevant business conditions out there are very favorable, so that the stakes are much higher than anticipated. Under these circumstances, a switch to the more experienced salesperson may well be the better course of action.⁵

In such situations there may be a clash between two fundamental 'heuristics' and it is not clear how a decision maker (DM) will resolve this tension.⁶ Thus, our main objective in designing, conducting and reporting our experiments is not merely to provide an additional evidence on whether or not Bayes' rule is also a good descriptive model. Rather, we wish to highlight the roles of the informational context and behavioral environment in explaining the pattern of departures from Bayes' rule and/or reinforcement.

In our design, each participant chooses a 'ball' from one of two 'urns', where the same (but undisclosed) state of the world determines the composition of valuable balls in each of the urns. After observing the outcome and replacing the ball, the participant then chooses the urn from which to draw a second ball, knowing that *the state of the world is the same across these two draws*. One urn contains either only valuable balls or only valueless balls (depending on the state), while the other contains a mix that is more favorable in the good state. A critical element of our environment is the different character of the information provided by the first draw made,

⁵ We thank Richard Zeckhauser for providing the essence of the above example.

⁶ We use the term 'heuristic' loosely, as Bayes' rule is indeed quite precise.

depending on which urn is chosen for this draw. We shall see that this plays a major role in the behavior that we observe.

We observe striking patterns of behavior: When both heuristics, reinforcement and Bayesian updating, are aligned, nearly all people respond as expected. However, results are dramatically different when these two heuristics clash. In our first treatment, nearly 50% of all second draws are made from the wrong urn, from the perspective of BEU. We also find a pronounced tendency toward ‘complexity aversion’, as people tend to make costly initial choices that mitigate or eliminate uncertainty/complexity in a subsequent decision. This behavior changes little over time, even though it is costly. In this sense, reinforcement from overall lower payoffs does not seem to lead to improved choices, although there is a strong relationship between the cost of an error and its frequency.

Although it is natural to be concerned that risk aversion plays a part in our results, our conclusions are essentially unchanged if we assume that participants have constant-relative-risk-aversion (CRRA) utility of the form: $u(x) = x^{1-\rho}$, where $\rho \in [0,1)$ is the CRRA coefficient. Risk aversion may well influence the decisions made, but it is not at the heart of the phenomena we describe, as we find that a very high degree of CRRA is typically needed to drive a wedge between optimal decisions based on expected payoffs and optimal decisions based on expected utility. In addition, given the stakes and the fact that each DM makes the same choice quite a few times, there is much less scope for risk aversion to affect our predictions.⁷

We develop hypotheses regarding the source of this behavior and test them in two follow-up treatments. In our second treatment, we maintain the same level of *affect* (a general term in psychology for feelings, emotions, or moods attached to reinforcement) as in the previous study

⁷ Calculations for the degree of risk aversion needed to reverse the BEU predictions are provided in Appendix C.

but significantly increase the informativeness of the first draw and simplify the updating task. We find modest shifts in behavior, in the directions predicted by BEU. However, even though the Bayesian calculations in this case are substantially less complex, the overall error rate is not reduced at all.

It is standard to only consider reinforcement in terms of material payoffs. In addition, experimental work on learning in games by reinforcement rarely attempts to associate the choice itself with affect.⁸ In contrast, our design explicitly enables us to examine separately the effects on behavior of overall payoff reinforcement and the immediate reinforcement from the success or failure of the choice. Often in life, choices we make and/or the actions we take result in simultaneously experiencing sensations (e.g., success or failure, warm glow) and receiving additional information about the state of our environment.

In our third treatment, we maintain the information provided by the first draw from the urn; however, we eliminate the affect associated with this draw by neither paying for the realized outcome nor even associating it with success or failure. Here there is a dramatic decline in the (switching) error rate after what would have been a successful outcome, and a more modest improvement after what would have been an unsuccessful outcome. This suggests that, to a large extent, errors are driven by the affect experienced from the initial draw, with positive affect being a particularly strong force.

In the next section we describe the design of our experiments, presenting the Bayesian predictions and the conjectures/hypotheses that motivated our follow-up treatments. In Section III we present our results and evaluate our hypotheses. In Section IV, we address five questions

⁸ However, some psychology studies (e.g., Suppes and Atkinson 1960) ignore material payoffs, arguing that a ‘win’ is a reinforcing event, in and of itself). Yet we are not aware of any reinforcement model that considers affect in the payoffs.

concerning the underlying determinants of the observed behavior; in Section V we discuss implications. Our last section provides a summary and points to possible future research.

II. Experimental Design and Predictions

Treatment 1

We conducted a Web-based experiment on the UCSB campus. Participants met in the lab and were given a handout (read aloud to the group) explaining the experimental set-up; these supplemental instructions can be found in Appendix A. After answering questions, we had participants direct their browsers to: <http://www.econ.ucsb.edu/~gcsurvey>. Instructions on the series of pages provided more detailed, hands-on instructions regarding the mechanics of making choices in the experiment.⁹ There were also questions testing comprehension (calculating payoffs in scenarios) that had to be answered correctly before a participant could move on.

In our design, there are two equally likely states of the world, the ‘good state’ (Up) and the ‘bad state’ (Down) and two lotteries (Left and Right), consisting of ‘urns’ from which the individual can draw ‘balls’. The urns contain some combination of black balls and white balls, where only black balls have value. In state Up, the Left urn has four black balls and two white balls, while the Right urn has six black balls. In state Down, the Left urn has three black balls and three white balls, while the Right urn has six white balls.

	Left Urn	Right Urn
Up ($p = .5$)	●●●●oo	●●●●●●
Down ($p = .5$)	●●●ooo	oooooo

⁹ The instructions on the Web initially were designed to be self-sufficient, but we discovered some serious comprehension problems in a pilot session when no additional instructions were given.

A decision-maker (DM) who does not know the state of the world makes two draws with replacement. The state of the world in the first draw remains the same for the second draw in that period, a fact clearly explained to the DM. Each DM made choices in 60 periods. We wished to ensure that participants gained some familiarity with a variety of strategies and outcomes, and also wished to insure that we would have a number of observations on *switching decisions* after draws from both urns. Thus, we imposed restrictions on the first draw during the first 20 periods (Phase I): In each odd (even) period, participants were required to start with the Left (Right) urn. During the next 30 periods (21-50, Phase II) we retained the same parameters, but people chose whichever side they wished for all draws.

In these first 50 periods, every time the DM draws a black ball from the Left urn, the DM receives 1 experimental unit; the pay is $7/6$ experimental units for each black ball drawn from the Right urn. White balls have no payoff value. After the first ball is drawn (and observed), the DM is asked to choose the lottery (Left/Right) from which to draw the second ball. In other words, the DM is asked whether s/he would stay with the same lottery (side) or switch to the other lottery (side). In the final 10 periods (51-60, Phase III), we reversed the payoffs, with a black draw from Left paying $7/6$ units and a black draw from Right paying 1 unit; in all other respects Phase III was identical to Phase II.

The results for each participant were tallied and participants were paid individually, at the rate of \$0.30 per experimental unit. Average earnings were \$23.14, with sessions averaging 40 minutes in duration. There were 59 participants in our first treatment, and we have a complete record of behavior for each person.

One characteristic of this decision task is that both urns offer the same expected payoff if there is only one draw, but one should optimally start with Right on the first draw if there are two draws. The reason for such a change is quite clear: Observing the outcome of the first Left draw does not resolve uncertainty regarding the state of the world as the first Right draw does. The more precise information about the state of the world from a Right initial draw improves the expected payoffs from the second draw. If there is only one draw, the expected payoff for both Left and Right is the same:

$$EU(Left) = \frac{1}{2} \times \frac{2}{3} + \frac{1}{2} \times \frac{1}{2} = \frac{7}{12} \qquad EU(Right) = \frac{1}{2} \times \frac{7}{6} + \frac{1}{2} \times 0 = \frac{7}{12}$$

With two draws, first suppose that Right is chosen on the first draw. Then the expected payoff (utility, assuming risk neutrality) of both draws is 17/12 (calculations for this and subsequent cases can be found in Appendix B).

On the other hand, suppose Left is chosen on the first draw. The expected utility of the first draw is 7/12, as before. Suppose the first ball drawn is black; then the expected utility of making the second draw from Left is 25/42 and the expected utility of making the second draw from Right is 28/42. Thus, if the DM draws a black ball from the Left urn on the first draw (a good outcome), it is optimal to switch (!) to the Right urn for your second draw. On the other hand, suppose that the first ball drawn is white; then the expected utility of making the second draw from Left is 17/30 and the expected utility of making the second draw from Right is 14/30. Thus, if a white ball is drawn from the Left urn on the first draw (a bad outcome) it is optimal to stay (!) with the Left urn for the second draw.

Thus, starting with Left gives an expected utility for both draws of $87/72$. The expected utility of both draws after a Right start is $102/72$, about 17% greater than that expected by starting with Left. It is therefore BEU-optimal to start from Right.

It is particularly interesting to notice the updating after a Left initial draw. After having a *successful* result, an initial draw of a black ball, a BEU decision-maker should switch to the Right urn. But such a DM should prefer to stay with the Left urn and not switch after a bad outcome, an initial draw of a white ball. Quite clearly, such predictions are not aligned with reinforcement learning models (and do not sit well with case-based decision theory).

We also consider the same environment with a slight modification: Instead of paying $7/6$ for black balls drawn from the Right urn and 1 unit for black balls drawn from the Left urn, we reverse these proportions. With these new prizes a BEU decision-maker should draw first from Left and stay with the Left urn on the second draw **no matter what the observed outcome** (calculations omitted).

Hypotheses and Follow-up Treatments

After starting from Right the switching decisions are simple, as the outcome (the color of the first drawn ball) perfectly reveals the state of Nature. Thus, the decision to stay with Right, after drawing successfully (a black ball), or the decision to switch to Left, after drawing unsuccessfully (a white ball), both ‘feel right’ and are ‘rational’. After starting from Left the decisions how to continue are much more complex. Again, drawing a black (white) ball here means success (failure) and is likely to reinforce continuing drawing from the Left (switching to drawing from the Right). However, even without precise (Bayes) calculations one may sense that here as well, as in Right, observing a black ball raises the likelihood that the state of nature is Up (Down), thus creating an incentive to switch to the Right (stay with Left). These two

opposite forces are not so easily reconciled.¹⁰ Thus, we expected a higher error rate after starting from Left.

H1 – Switching-error rates after Left draws will be higher than after Right draws.

Two natural primary, and not mutually exclusive, explanations come to mind after observing the dramatically higher switching-error rates on Left starts: 1) The clash between the two heuristics, and 2) The relative complexity of the switching decisions after a Left start. Accordingly, we designed two additional treatments to gain further insight into the behavioral influences at play.

In Treatment 2 we significantly simplified the Bayesian inference after a first draw from Left. We had 54 participants in Treatment 2, who earned an average of \$22.06. Treatment 2 is obtained from Treatment 1 by a single change: We replaced one black ball with a white ball in the Down state of the Left urn, resulting with two black balls and four white balls, and so reversing the composition of the Left urn in the Up state:

	Left Urn	Right Urn
Up ($p = .5$)	●●●●○○	●●●●●●
Down ($p = .5$)	●●○○○○	○○○○○○

¹⁰ The careful reader may wish to keep in mind that the reinforcement forces after starting from Left may have a different impact in Phase I as compared to Phase II, due to the restriction on the initial draw. It seems plausible that the emotional reaction to the results of the first draw might be stronger when one has chosen Left voluntarily. We shall revisit this point in section V.

Left's Up and Down states are now symmetric, with $\Pr[\text{Up} \mid \text{black}] = \Pr[\text{Down} \mid \text{white}] = 2/3$. This change increases the informative sharpness of a first draw from the Left. This change ought to 'move' behavior toward the BEU benchmark, leading to the following hypothesis:

H2 – Switching-error rates after Left draws will be lower in Treatment 2 than in Treatment 1.

In addition, as a consequence of the manner in which we increase the informational precision, Left draws are less attractive (fewer black balls in the Left urn in the Down state). In the absence of countervailing forces, this gives the following hypothesis:

H3 – There will be fewer Left start choices in Treatment 2 than in Treatment 1, in phases where a DM makes the starting decision.

Does *affect* play an important role in the observed behavior? In Treatment 3, we significantly reduce the affect received from the first draw, while maintaining the precision of the information obtained. Treatment 3 has the same color composition and payoffs as in Treatment 2; however, now we require people to start from the Left side throughout the session. The result of this first draw doesn't count towards payoffs and the 'color of success' is not known at that time. In this way, we remove the emotional reaction of 'winning' or 'losing' that otherwise goes along with observing the outcome and deciding what to do next. This first draw *does* provide information about the state of the world, and the DM is told the (random) color of success after observing the outcome of the first draw. We had 52 participants in Treatment 3, who earned an average of \$17.82.

Given the lower affect (emotional response) of the outcome of a Left draw in Treatment 3 compared to Treatment 2 (identical informational precision), we would expect fewer Updating errors in Treatment 3 after a Left draw. To the extent that affect is involved in the observed switching errors after Left, we have the hypothesis:

H4 – The error rates will be lower in Treatment 3 than in Treatment 2 for switching after initial Left draws.

III. Results

Our benchmark for analyzing our results is the behavior of a BEU decision-maker, and we define ‘errors’ to be deviations from that benchmark. We first present aggregated data for each treatment and then evaluate our hypotheses of the previous section. We also illustrate patterns of individual behavior with histograms for starting-error and switching-error rates.¹¹

Aggregate Behavior

Tables 1-3 display BEU switching-error rates in the corresponding treatments.

Table 1 – Treatment 1 Switching-error Rates

After a Right draw				
Draw	Phase			Aggregated
	I	II	III	
Black	12/286 (4.2%)	20/596 (3.4%)	20/161 (12.4%)	52/1043 (5.0%)
White	15/304 (4.9%)	23/683 (3.4%)	7/160 (4.4%)	45/1147 (3.9%)
Combined	27/590 (4.6%)	43/1279 (3.4%)	27/321 (8.4%)	97/2190 (4.4%)
After a Left draw				
Draw	Phase		Aggregated	Phase III
	I	II	I and II	
Black	179/332 (53.6%)	180/272 (66.2%)	359/604 (59.4%)	31/157 (19.7%)
White	98/258 (38.0%)	63/219 (28.8%)	161/477 (33.8%)	17/112 (15.2%)
Combined ¹²	277/590 (46.9%)	243/491 (49.4%)	520/1081 (48.1%)	48/269 (17.8%)

¹¹ Error rates for each decision made by each individual are given in Appendix D.

¹² Note that combining Black and White with the weights based on the number of draws creates a bias in favor of the Black rate, since (on the Left) Black should be drawn 7/12 of the time. If we give the Black and White rates equal weight, the corresponding percentages are 45.8%, 47.5%, 46.6%, and 17.5%.

It is immediately apparent that switching errors are relatively rare after a Right draw, but are quite common after a Left draw. The error rate appears to be much higher after drawing a black ball from the Left than after drawing a white ball from the Left.

Starting errors apply only to Phase II and Phase III. There were 491 Left initial draws of the total of 1770 in Phase II, or 27.7%. Recall that starting from Left is optimal in Phase III; we see that 321 of 590 (54.4%) initial draws in Phase III were nevertheless from the Right urn.

Table 2 – Treatment 2 Switching-error Rates

After a Right draw				
Draw	Phase			Aggregated
	I	II	III	
Black	32/255 (12.5%)	31/642 (4.8%)	20/178 (11.2%)	83/1075 (7.7%)
White	12/285 (4.2%)	16/627 (2.6%)	3/161 (1.9%)	31/1073 (2.9%)
Combined	44/540 (8.1%)	47/1269 (3.7%)	23/339 (6.8%)	114/2148 (5.3%)
After a Left draw				
Draw	Phase		Aggregated I and II	Phase III
	I	II		
Black	95/258 (36.8%)	75/155 (48.4%)	170/413 (41.2%)	69/99 (69.7%)*
White	157/282 (55.7%)	102/196 (52.0%)	259/478 (54.2%)	30/102 (29.4%)
Combined	252/540 (46.7%)	177/351 (50.4%)	429/891 (48.1%)	99/201 (49.3%)*

*In Phase III of Treatment 2, it is BEU-optimal to switch to Right after a Black draw from Left, in contrast to the BEU-optimal strategy in Study 1.

Once again we see that switching errors are relatively rare after a Right draw and quite common after a Left draw. The overall switching-error rates in Phases I and II are quite close to the corresponding rates in Treatment 1. However, now the error rate appears to be lower after drawing a black ball from the Left than after drawing a white ball from the Left.

Concerning starting errors, there were 351 Left initial draws of the total of 1620 in Phase II, or 21.7%.¹³ Recall that starting from Right is (barely) best in Phase III; we see that 201 of 540 (38.7%) initial draws in Phase III were made from the Left urn.

Table 3 – Treatment 3 Switching-error Rates

Draw	Phase		Combined
	I	II	
Favorable	245/1811 (13.5%)	78/259 (30.1%)*	323/2070 (15.6%)
Unfavorable	776/1829 (42.4%)	96/261 (36.8%)	872/2090 (41.7%)
Combined	1021/3640 (28.0%)	174/520 (33.5%)*	1195/4160 (28.7%)

*In Phase II of Study 3, it is BEU-optimal to switch to Right after a Left, Favorable draw, in contrast to the BEU-optimal strategy in Phase III of Study 1.

Since it is not possible to start with Right in Treatment 3, there are no starting errors or decisions after Right draws here. By ‘Favorable’ (‘Unfavorable’) we just mean that our randomization of the prized color in the second draw picked the same (other) color as the DM’s initial draw.

Let us next examine whether DMs make initial draws in accordance with the BEU model. Recall that the correct BEU choice in Treatment 1 is to start from Right in Phase II and to start from Left in Phase III; in Treatment 2, it is best to start with Right at all times.

Table 4 shows voluntary initial draws in Phase II and III of Treatments 1 and 2:

Table 4 – Starting-error Rates

Treatment	Phase	BEU-start	Error Rate
1	II	R	491/1770 (27.7%)
2	II	R	351/1620 (21.7%)
1	III	L	321/590 (54.4%)
2	III	R	201/540 (37.2%)

¹³ There were 70 periods in Phase I of Treatment 3, so we could also choose to compare only the first 50 of these to their equivalent in Treatment 2. However, we shall see that error rates are quite similar in periods 21-50 and 51-70 of Treatment 3.

In Phase II of both Treatments 1 and 2, the strong majority of people choose Right, in accordance with the BEU prediction, although Left is chosen more than 20% of the time.¹⁴ People are somewhat more likely to start with Right in Treatment 2 than in Treatment 1, in accordance with the fact that black balls are relatively less prevalent in Treatment 2. Wilcoxon tests (on individual tendencies) find the difference to be only marginally significant for Phase II ($Z = 1.26, p = 0.104$), but quite significant for Phase III ($Z = 2.32, p = 0.010$).

It is interesting to note that the error rates for Phase II initial draws are much lower than for the switching decisions. This is a bit puzzling, since calculating the correct initial side from which to draw involves solving the switching problem (from both sides). In our later discussion, we provide some possible explanations regarding how DMs who did so poorly on the switching decisions do so much better on this more difficult problem.

Evaluation of Hypotheses

Recall that after an initial draw from the Left the two heuristics clash and we hypothesized that a DM will do worse when he or she is ‘pulled’ in opposite directions. People who are impacted by both heuristics may attach different weights to each and end up either going with one heuristic or experimenting by alternating their choices when it’s not obvious what to do. In fact, we observe a very close correspondence between BEU-predictions and actual decisions when DMs start with Right. The aggregate error rate after Right draws is 4.4% in Treatment 1 and 5.3% in Treatment 2. Even this low proportion masks the fact that there are ‘worst

¹⁴ We can also examine whether people are more likely to start with Right over time, by considering Phase II behavior. In periods 21-30 of Treatment 1, 413 of 590 starts (70.0%) were from Right. This increased slightly to 426 of 590 starts (72.2%) in periods 31-40, and to 439 of 590 (74.4%) in periods 41-50. In periods 21-30 of

offenders.’ We exclude seven DMs as worst offenders in each of Treatments 1 and 2. These people, who constitute roughly 12% of the population (7/59 and 7/54), are responsible for about 75% of all errors.¹⁵ Apart from those seven people, the aggregate error rate is between 1% and 2% in every category.¹⁶

After a Left initial draw, we see a very different picture. There is no particular correspondence between BEU-predictions and actual decisions, as switching-error rates after Left draws in Phases I and II are substantial (between 34% and 59%) for both Treatments 1 and 2. It is clear that these data strongly support H1. One simple test compares each individual’s switching-error rates after all Left initial draws and all initial Right draws; we see that the error rate is higher after Left draws than after Right draws for 111 of 113 of the individuals in Treatments 1 and 2 ($Z = 10.44$, $p = 0.000$, binomial test). Here a BEU must switch (not switch) after a success (failure), the opposite of the prediction of a reinforcement or case-based model. The error rates in Phase III of Treatment 1, where a BEU should never switch from Left regardless of the first draw, are much lower than in the earlier phases, but are still substantial; in Treatment 2 these error rates are definitely higher, particularly after a black ball is drawn.

We can evaluate Hypothesis 2 and Hypothesis 3 on the basis of Tables 1 and 2. Regarding Hypothesis 2, we see that the switching-error rate at LB is indeed lower in Treatment 2 than in Treatment 1, but this is not the case after LW. Comparing the aggregated switching-error rate after Left in Treatments 1 and 2, we find no support for Hypothesis 2. The rates are

Treatment 2, 426/540 starts (78.9%) were from Right. This decreased very slightly to 420 of 540 starts (77.8%) in periods 31-40, and to 423 of 540 (78.3%) in periods 41-50.

¹⁵ For example, the high error rate for Phase III after a Black draw in Treatment 1 reflects three DMs who made 16 of the 20 errors in this case.

¹⁶ This is about as good as it gets in laboratory tasks. A colleague points out that this is lower than the percentage of students who forget to write a name on the exam book after being reminded three times.

almost identical in Phase I and Phase II.¹⁷ It does not appear that the increased precision of the information is effective in reducing this error rate.

Hypothesis 3 finds widespread support, as Left draws are less likely in Treatment 2 for every comparison. In Phase I-II, people switch to Right after LB 40.4% of the time in Treatment 1 vs. 58.8% of the time in Treatment 2; after LW in Phase I-II of Treatment 1 (Treatment 2), people switch to Right 33.8% (54.2%) of the time. In Phase III, people switch to Right after LB 19.7% of the time in Treatment 1 vs. 30.3% of the time in Treatment 2; after LW, the respective rates of switching to Right are 15.2% and 29.4%. With respect to starting rates, 27.7% (21.7%) of starts are from Left in Phase II of Treatment 1 (Treatment 2). Left starts are also more frequent in Treatment 1 during Phase III, 54.4% to 62.8%.

The striking difference between Treatment 3 and Treatment 2 is that the error rate after the equivalent of LB draws in the initial phase are much rarer, dropping to 13.5% (from 36.8% in the first 20 periods, when Left first draws were are involuntary. The switching-error rate also drops after the equivalent of LW draws, but this decline is smaller (42.4% compared to 55.7% in the first 20 periods). Thus, Hypothesis 4 is supported both after LB draws (switching-error rates of 36.8% vs. 13.5%) and after LW draws (55.7% vs. 42.4%) in Phase I, the proper comparison. In Phase III, we see error rates lower in Treatment 3 after LB draws (69.7% to 30.3%), but not after LW draws (29.4% to 36.8%).

Individual Behavior

While aggregating the data is quite informative, doing so obscures individual variations. In this subsection, we document the considerable heterogeneity using simple histograms.

Figures 1-5 display starting-error rates and switching-error rates across individuals:

¹⁷ The rate is actually higher in Treatment 2 than in Treatment 1 during Phase III, but the BEU optimal switching








Figure 1 shows a great deal of variation in starting choices. Figure 2 shows that only a handful of participants make many (over 10%) switching errors after an initial RIGHT draw; however, there is high variance across individuals in switching-error rates after a Left draw. In Phase I-II, there are many individuals in each of the categories displayed, for both Black and White initial observations. In Phase III, there are only a few people who switch to Right more than half the time, but there is still considerable variance across the first three categories.

changes after LB, making comparisons a bit problematic.



The patterns in Treatment 2 are very similar to those in Treatment 1, with starting errors a bit less common. In Treatment 3, we again see substantial variance for Left switching-error rates, with error rates less than 10% for a substantially greater proportion of the participants.

IV. Patterns and Questions

In this section, we consider patterns in individual behavior, as well as some questions that shed light on the observed ‘error structure’.

Given the considerable heterogeneity observed for individual error rates, it may be possible to identify ‘types’. For example, to what extent are DMs sensitive to information learned in the previous period? Further, do people who make switching errors within a period also make more changes across periods, based on prior outcomes?

Q1. Are there identifiable ‘reinforcement types’ in the population? Is individual cross-period behavior correlated with individual within-period behavior?

To answer this question, we create individual measures of susceptibility to reinforcement by examining the degree to which the individual’s choice behavior is sensitive to the previous outcome from a comparable choice.

One such measure can be determined by considering the likelihood that an individual changes the location (Right or Left) of the initial draw, depending on the number of successful outcomes in the previous period. In Treatments 1 and 2, each period has two outcomes, so there can be zero, one, or two successful (black) outcomes. We calculate the rate that each individual switches the location of the initial draw in Phase II after no black draws in a period and also after two black draws in a period; since periods with zero black draws are less common, we also compare the rate after two black draws with the combined rate for zero or one black draws.

This measure has some meaning for both Right and Left initial draws. While switching after a previous Left initial draw is always more common than switching after a previous Right initial draw, the likelihood of switching variously inversely with the number of black draws in the previous period. The aggregate switch rates after voluntary Phase II Left draws are 59.6%, 53.1%, and 37.8% after zero, one, or two black draws, respectively; after Right draws, these rates are 22.1%, 19.4%, and 11.1%, respectively.

We see that about 29 of 113 (26%) participants in Treatments 1 and 2 never change their initial draws in the 30 periods of Phase II. For the remaining participants, reinforcement typically plays a role, as the rate of switching starting choices is higher after zero or one last-period successes than after two last-period successes for 71 of these 84 individuals (85%). A simple binomial test indicates that this is quite significantly different from random behavior ($Z = 6.33$, $p = 0.000$). This rate also declines (weakly) monotonically for 36 people (43%) as the number of successes increases from zero to one to two. The positive difference in individual rates after zero or one last-period successes and two last-period successes is statistically significant, despite the small sample size, for 24 participants (29%).

Another reinforcement measure is the difference in the likelihood that a person who drew first from the Left in the previous period again makes the first draw from Left in the period under consideration, depending on whether the outcome of the previous Left initial draw was black or white.¹⁸ Excluding all cases where the location of the initial draw is mandated, we calculate this measure using choices in Phase II of Treatments 1 and 2; since there were no observations in one or more of these categories for 41 people in these treatments, we obtain differences in rates for 72 individuals. Thirty-seven of these people were more likely to stay with left after drawing black in the previous period than after drawing white, compared to 19 people for whom this

relationship was reversed and 16 people for whom this made no difference. The positive difference in rates was at least 25 percentage points for 24 people (33%).

There are two reinforcement measures for Treatment 3 that parallel those used in Treatments 1 and 2. First, we can check the degree to which an individual's first second draw depends on the number of successful outcomes (zero or one) the last time that individual had a similar outcome (favorable or unfavorable). Second, there is the difference in the likelihood that a person who drew second from the Left in the previous period again makes the second draw from Left in the period under consideration, depending on whether the outcome of the previous second draw from Left was successful or unsuccessful.

Prior outcomes also affect behavior in Treatment 3. Forty-seven of 52 people (90%) were more likely to change their second draw depending on its previous success, significantly different from random behavior ($Z = 5.82, p = 0.000$). The positive difference in individual cross-period switching rates depending on last-period success is statistically significant for 19 people (37%).¹⁹ The differences for the second measure are positive for 40 of the 50 (80%) people with calculable measures ($Z = 4.24, p = 0.000$) and positively greater than 25 percentage points in 10 cases (20%).

It is clear that some people's choices are quite malleable, while these remain fixed for others. We now consider correlations between an individual's cross-period reinforcement susceptibility and his or her tendency to make errors in within-period switching decisions. Table 5 shows Spearman rank-correlation coefficients for all pairwise combinations of cross-period and within-period measures for Treatments 1 and 2 combined:

¹⁸ We thank Stefano dellaVigna for suggesting this idea.

¹⁹ Note that we have more observations in Treatment 3, making statistical significance easier to achieve with the same proportionate differences.

Table 5: Spearman's r for Measures in Treatments 1 and 2

	<i>L diff</i>	<i>Diff 20</i>	<i>LBL</i>	<i>LWR</i>	<i>Diff 210</i>
<i>Diff 20</i>	.417***				
<i>LBL</i>	.152	.072			
<i>LWR</i>	.229***	.254**	-.156		
<i>Diff 210</i>	.393***	.822***	-.054	.294***	
<i>R rate</i>	.222*	.294***	-.154	.517***	.287**

*** indicates significance at $p = 0.01$, ** indicates significance at $p = 0.05$, and * indicates significance at $p = 0.10$ (all two-tailed tests).

In this table, *Diff 20*, *Diff 210*, and *L diff* are the cross-period reinforcement measures discussed in turn above. All of these measures are very highly correlated with each other. *LBL* and *LWR* are the within-period switching-error rates in Phase II after a black draw from Left and a white draw from Left, respectively. Recall that reinforcement should tend to increase these error rates. *R rate* considers individual error rates after initial voluntary draws from the Right.

While none of the cross-period measures is significantly correlated with switching-error rates after a black draw from Left, these measures are all highly correlated with switching-error rates after a white draw from Left. The switching-error rate after voluntary draws from Right is significantly correlated with the cross-period measures and *LWR*, but not *LBL*. Interestingly, there is a negative (but not significant, $p = .189$) correlation between *LWR* and *LBL*.

Thus, there generally is a real correlation between individual cross-period and within-period reinforcement measures, with the exception being *LBL* errors. In this latter case, perhaps the positive affect from the successful first draw overwhelms individual cross-period tendencies.

Table 6 shows Spearman rank-correlation coefficients for all pairwise combinations of cross-period and within-period measures for Treatment 3:

Table 6: Spearman's r for Measures in Treatment 3

	<i>L2 diff</i>	<i>Diff 10</i>	<i>LFL</i>
<i>Diff 10</i>	.362***		
<i>LFL</i>	.172	.118	
<i>LUR</i>	.013	-.050	.531***

*** indicates significance at $p = 0.01$

Here we see essentially no correlation between within-period measures and the cross-period measure. Apparently removing affect eliminates the correlation found between LWR and the cross-period measures seen in Table 5.

Q2. Are the switching-error rates the same after Black and White draws?

After initial Right draws in Treatment 1, the low aggregate switching-error rate is largely independent of whether the first draw was black or white. In contrast, when the first draw is from the Left in Phases I-II of Treatment 1, the error rate after drawing a black ball is nearly double the (still-substantial) error rate after drawing a white ball. For statistical purposes, we cannot treat each observation as being independent, so we must examine individual behavior. A Wilcoxon signed-ranks test (see Siegel and Castellan 1988) comparing each individual's respective switching-error rates rejects the hypothesis that there is no difference in these switching-error rates, with $Z = 3.08$, $p = 0.002$ (two-tailed test). This large difference is also present when we exclude either worst offenders or all offenders. There is little difference in switching-error rates in Phase III of Treatment 1.

There is a slight difference in switching-error rates after Right draws in Treatment 2; however, this vanishes when the worst offenders are excluded. In Phases I-II of Treatment 2, the direction of the difference in switching-error rates after Left draws is reversed, with errors more common after observing a white ball than after observing a black ball. The Wilcoxon

signed-ranks test again finds a significant difference between these switching-error rates, with $Z = 2.01$, $p = 0.044$ (two-tailed test). However, this difference is eliminated when we exclude the worst offenders, and seems to reverse slightly when we exclude all offenders. In Phase III, we see a very large difference in switching-error rates, which appears to actually *increase* when we remove subsets of offenders.

The aggregate Phase I-II switching-error rate from Left in both Treatments 1 and 2 is close to 50%. With a 50% error rate, it is tempting to conclude that play is random after drawing from the Left. However, a more careful examination reveals that this 50% aggregated switching-error rate is an artifact of the data, as rates are very different for black or white draws.

In Treatment 3, the switching-error rate after ‘unfavorable draws’ (different color than what is desired for the second draw) in Phase I triple the corresponding rate after ‘favorable draws’ (same color as what is desired for the second draw). On the other hand, there is only a modest difference between these rates in Phase II, where the payoffs have been reversed.

Q3. Is there any evidence that switching-error rates diminish over time?

A natural consideration is whether the errors that we observe represent initial confusion or whether they generally persist through later periods. At a first rough cut, switching-error rates do not appear to be diminishing over time. One comparison is between the first and second Phases of each Treatment. In Treatment 1, switching-error rates after Right draws drop slightly from 4.9% in Phase I to 3.4% in Phase 2. However, the switching-error rate after LB actually *increases* from 53.6% in Phase I to 66.2% in Phase II. The rate does drop slightly over time after LW. The overall error rate after starting with Left is slightly higher in Phase II than in Phase I (49.7% vs. 46.9%). However, since individuals make their initial choices voluntarily in

Phase II and are forced to do so in Phase I, there is a self-selection issue that may confound comparisons across these phases.

An alternative test of changes in behavior over time is to consider switching-error rates for time segments of the penultimate phase in each treatment, since the environment is constant for a relatively long interval. Table 7 shows these rates:

Table 7 – Switching Error Rates over Time Segments

<i>Treatment 1 Switching-error Rates</i>			
Initial Draw	Periods 21-30	Periods 31-40	Periods 41-50
RB	5.0%	2.5%	2.6%
RW	4.4%	3.3%	2.3%
LB	60.4%	61.4%	77.8%
LW	27.8%	26.7%	33.8%
<i>Treatment 2 Switching-error Rates</i>			
Initial Draw	Periods 21-30	Periods 31-40	Periods 41-50
RB	3.9%	5.4%	5.4%
RW	0.5%	2.3%	4.6%
LB	49.1%	52.0%	58.0%
LW	49.2%	54.3%	52.2%
<i>Treatment 3 Switching-error Rates</i>			
Initial Draw	Periods 1-20	Periods 21-50	Periods 51-70
L‘B’	13.4%	13.9%	13.1%
L‘W’	46.2%	41.0%	41.5%

The switching-error rate after RB decreases over time in both Treatments 1 and 2; however, while the rate after RW also increases over time in Treatment 1, it decreases over time in Treatment 2. In any case, all of these differences are rather small. The switching-error rate is


either slightly increasing over time or is relatively flat after Left draws in all three treatments.²⁰ Thus, overall we see little in the data to indicate that switching-error rates drop over time.

Cost and Frequency of Errors

Q4. Does the cost of an error affect the frequency of the error?

A natural question for economists is whether the frequency of decision errors is inversely related to the cost of such errors, even if complex calculations may be beyond the ability of the general population. A complete Table showing the expected cost of each of the 24 possible errors and the corresponding frequency of the error is provided in Appendix D. Figure 6 uses these data and illustrates the inverse relationship between the expected cost of an error and its frequency:

Figure 6 - Cost and Frequency of errors



While the relationship is not completely smooth, it is clear that, in general, the higher the cost of an error, the less likely it is to be made. Every error that costs more than .30 has a frequency of less than 5%, while every error that costs less than .30 has a frequency greater than

²⁰ Since there are more Black draws than White draws from Left, and since the initial level of the Left switching-error rates is higher for Black, we might expect the Black error rate to decrease more quickly than the White error rate. However, our results are not in accord with this expectation.

11.2%. Thus, while the difference in payoffs between the optimal play and the alternative choice may be small in some cases, overall, people are responding to the economic incentives provided.²¹

Nevertheless, there are some patterns to the deviations from this cost/frequency relationship. For example, the expected cost of drawing from the Left urn after first drawing from the Right urn and observing Black is .222 in Phase III of Treatments 1 and 2, and the frequency of this error is 12.4% and 11.2%, respectively. By comparison, the expected cost of drawing from the Left urn after first drawing from the Left urn and observing Black in Phases I and II of Treatment 2 is also .222, yet the frequency of this error is 41.2%. On the other hand, consider the same error (Left after LB) in Treatment 3. Here the expected cost of the error is still .222, but the error rate is only 13.4%. Apparently the transparency of the updating problem and the presence of affect are also important considerations.

Gender

Q5. Are there differences in error rates across gender?

We have seen that there is considerable heterogeneity in the population. Can we find any determinants for this – specifically, are there any differences in behavior across gender? The motivation for such an inquiry is the possibility that females and male are responding with systematic differences to different triggers, the reinforcement response and the forward and rational reaction to the (pure) informational signal upon observing the color of the first draw.²²

²¹ In Treatment 1, the difference between the expected payoffs from the best strategy and the worst strategy is \$7.96, or 37.8% of the expected payoff from making random choices. In Treatment 2, the corresponding figures are \$6.34 and 32.5%; in Treatment 3, the corresponding figures are \$3.22 and 24.8%.

²² After we ran our experiments, collected the data and wrote the first draft, readers/listener suggested that it would be interesting to control for subpopulation characteristics such as gender, statistical sophistication, major, etc. While it was too late to collect most of these data, we were able to determine each participant's gender from his or her name.

Appendix F shows male and female error rates for each of the 24 possible errors. We use the Wilcoxon-Mann-Whitney ranksum test on individual error rates, with around 25-30 individuals of each gender for each comparison.²³ While only a few of the comparisons are significant on an individual basis, there does seem to be a pattern here. The female error rate is higher in 20 of the 24 comparisons.²⁴ A binomial test thus rejects equality at $p = 0.002$, two-tailed test. In addition, the difference seems to be sharper when errors are more costly. For the 12 costliest errors of the 24 listed (cost at least .222), male error rates were lower in 11 cases, with the single exception being at the lowest cost in the range. Five of these 12 differences were at least marginally significant. It would appear that the reinforcement heuristic is relatively stronger for the female participants in our study.

Regression Results

We supplement our nonparametric tests for the questions posed above with regressions in which the frequency of error is the dependent variable. While the relevant assumptions concerning the underlying error distribution may not be satisfied, regression analysis typically has more power. Table 8 shows the results for a number of specifications including the explanatory variables we have discussed:

²³ Ranksum tests have the advantage of implicitly putting less weight on outliers, since a very high error rate for an individual translates only into a poor ranking. Thus, some of the apparently large differences that are due to outliers are not significant.

²⁴ Twenty-three of the 24 signs for the nonparametric tests are negative, an even stronger result. Apparently three of the cases where the difference was positive were so outlier-driven that the rank-sum test gives them a negative (but not significant) test statistic.

Table 8: Regressions for Frequency of Error in Decisions

Independent variables	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) RE	(6) RE
Constant	.469*** (.042)	.547*** (.066)	.061 (.049)	.305** (.129)	.100** (.051)	.175*** (.052)
Cost	-.928*** (.153)	-1.88*** (.498)	-	-1.84*** (.566)	-.483*** (.090)	-1.48*** (.215)
Cost_squared	-	1.76** (.861)	-	2.14** (.889)	-	1.75*** (.344)
Left	-	-	.310*** (.063)	.172** (.085)	.226*** (.033)	.232*** (.032)
Affect	-	-	-	.114* (.066)	.118*** (.031)	.127*** (.031)
Female	-	-	-	-	.079*** (.025)	.079*** (.019)
N	24	24	20	20	999	999
Adjusted R ²	.609	.656	.547	.723	.252	.268

Standard errors are in parentheses; *** indicates significance at $p = 0.01$, ** indicates significance at $p = 0.05$, and * indicates significance at $p = 0.10$ (all two-tailed tests). Left = 1 if a decision involves updating after a draw from Left and is 0 otherwise. Affect = 1 if a decision involves updating and the first draw counts for payoffs, but is otherwise 0.

Specifications (1) and (2) are simple OLS regressions on the data displayed in Figure 6, with each observation representing the population value for each decision error. Specifications (3) and (4) considers only the 20 switching errors and consider dummy variables for updating transparency and affect. In specifications (5) and (6), each individual's error rate for each possible error is treated as one observation. Since each individual has either eight (Treatments 1 and 2) or four (Treatment 3) possible updating errors, there is the potential problem that observations are not independent. We use a random-effects model to account for the multiple observations for each individual, and include gender in the regressions.^{25,26}

²⁵ Given the potential censoring issue for the frequency of error, we also ran Tobit regressions. As these produced very similar results, we report the OLS regressions.

²⁶ The much higher Adjusted R² for specifications (1)-(3) reflects the fact that the population rates used have already been 'smoothed' and the individual observations in specifications (4) and (5) have not.

All specifications confirm the strong effect of the cost of an error on its frequency. In addition, specifications (4), (5), and (6) confirm the effects of transparency, gender, and affect. The updating error rate is substantially and significantly higher when (in declining order of magnitude) the first draw is from the Left, when there is affect involved with respect to the first draw, and when the chooser is female.

V. Discussion

On a basic level, some form of reinforcement seems a simple and natural mechanism for even simple creatures to use for guidance and learning, since no elaborate cognitive process is needed. In contrast, the ability to update priors, even if in an imperfect (Bayesian) way seems to require enormously more: brain, reasoning, cognition, etc. However, good solutions to many decision problems in general, and economic ones in particular, require one to update priors when new information arrives.

A sophisticated DM when confronted with a decision problem framed in terms of priors and additional information (as in our experiment) may recognize the environment and apply some updating heuristics (perhaps Bayesian). But how likely is it that in a real-life (and unframed) situation a less-than-fully-rational DM will indeed apply something close to Bayes' rule? Another relevant question is whether this distinction is all that important. One might question whether decisions where these two heuristics clash are very common. We suspect that in most situations both heuristics are at least partially aligned, so that relying on reinforcement alone may be effective. Nevertheless, our results indicate that both heuristics are at work, and that when they clash we can expect confusion and divergence from optimality. We feel that this observation provides an important insight, and also suspect that there tends to be at least a kernel of updating from negative inferences in many decisions.

In any case, even if reinforcement and Bayesian updating often give the same advice, we argue that observing the results of this clash and the relative strengths of the conflicting forces may be relevant even when they do not clash. In some cases, DMs experience both forces, while others are experiencing only one. Consider, for example, an investor whose recent portfolio has performed very well, beyond investor expectations. Choi, Laibson, Madrian & Metrick (2003) find the puzzling result that people who experience orthogonal positive wealth shocks (higher appreciation than expected) in their 401k retirement accounts do not increase their consumption, in violation of standard theory. As differences in capital gains in one's own portfolio vs. capital gains in another investor's portfolio should be irrelevant for forecasting future returns, in principle the informational content of the recent idiosyncratic success can be ignored.

However, in practice this does not appear to be the case. Since future returns on any given portfolio are equally available to all investors and information about past returns is public, one might wonder why these two investors appear to process information differently, as seems to be implied by their results. In terms of our work, an explanation might be that the affect of a successful choice gives an additional 'push' for successful investors to invest in the 401k. We would further conjecture that investors who personally select their own 401k portfolios would be particularly prone to this influence. In general, the affect received from one's own experiences is likely to increase the strength of 'learning' relative to that produced by the experiences of others.

Some evidence relevant to this conjecture emerges from our study, as we can test this point by comparing switching decisions after voluntary and involuntary initial draws. A successful portfolio (black, in Treatments 1 and 2) can be the outcome of 'investing' in either Right or Left. Successful investors in Right will continue to invest in Right, as would any observer learning the outcome of the draw. However, a successful investor in Left who has

elected to invest initially in Left (Phase II) is more likely to invest in Left in the next draw than is the investor who has also had a good outcome, but did not choose where to invest (Phase I). LB switching rates are about 12 percentage points higher in Phase II than in Phase I in both Treatment 1 and Treatment 2. This difference corresponds to about a 20-25% increase.

On the other hand, switching errors after LW draws in fact decrease slightly from Phase I to Phase II, in both treatments. Thus, DMs who observe a white ball after an initial Left draw are more likely to persist with Left for the second draw when the first Left draw was voluntary. In terms of the portfolio story, this means that an investor who chooses a portfolio tends to invest more in it after a bad outcome than does an investor who has observed the outcome without personal involvement.

We do not find that increasing the precision of the information received after a Left draw reduces the switching-error rate (Hypothesis 2). On the other hand, we find evidence that removing the affect from ‘winning’ on the first draw (Treatment 3) reduces errors, as average error rates in Treatment 3 are indeed considerably lower (Hypothesis 4). Finally, when Left is made less attractive (by changing the composition of the balls in the Left urn in the Up state), the Left choice is always less frequent in starting and in switching decisions.

We do find a strong relationship between the cost and frequency of errors. However, if the cost of error is the primary consideration, one might expect that error rates diminish over time, as people tried to avoid these costs. Yet we find little or no evidence that error rates decrease over time, as one might expect with some form of reinforcement learning. In fact, our results go to the question of the nature of reinforcement. On the one hand, the reinforcement from payoffs does not appear to be a major factor, given the lack of change over time. On the other hand, we see a strong influence from affect, which could be considered a ‘sensation-based’

form of reinforcement. Since affect seems to play such a large role, perhaps reinforcement models could be improved by considering this to be part of the overall ‘payoff’.

Regarding initial draws, there seems to be a pronounced tendency to prefer to make the first draw from the Right, unrelated to BEU calculations. In Treatment 1, this is selected 72% of the time in Phase II (when it is correct under BEU) and 54% of the time in Phase III (when it is *not* correct under BEU). In Treatment 2, Right is selected first 78% of the time in Phase II (when it is correct under BEU) and 63% of the time in Phase III (when the expected payoff is nearly the same for Right and Left starts).²⁷ Perhaps one explanation for this behavior can be traced to the fact that making the first draw from the Right leads to easy subsequent switching decisions. In contrast, starting from Left leads to a difficult decision node after the first draw; this is true both conceptually and as evidenced by the data.

We see two plausible explanations for this phenomenon. First, there may be a ‘curiosity effect’; people express interest in learning whether the actual state in the period was Up or Down. Making at least one choice from Right ensures that the DM learns this information, while Left choices leave this question unanswered. Second, a DM who selects Left initially and then gets a bad outcome on the next draw might at that point be haunted by the suspicion that he or she updated incorrectly and was therefore instrumental with respect to the bad outcome. Loomes and Sugden (1982) introduce the notion of *regret* in choice under uncertainty. If an individual chooses between two actions and the outcome resulting from her choice is less desirable than the outcome that would have resulted (in the *ex post* state of the world) from the other action, this may serve to diminish the pleasure that the individual might otherwise receive from the less

²⁷ This 63% figure is perhaps the clearest evidence of the start-from-Right bias, since the expected (optimal) payoffs from starting with either urn are nearly identical.

desirable outcome.²⁸ In our context, this could lead to regret over the initial choice. To the extent that DMs anticipate or experience this, first draws from Right become more likely.

Regarding switching decisions, we find that there seems to be a reluctance to switch when the first draw is from Left, since the error rate is much higher when one is supposed to update and switch (black draw) than when one is supposed to update and stay (white draw). People who voluntarily start with Left tend to switch less frequently than when they are constrained to start with Left. This suggests that there is less updating in this more complex cognitive task than in the simpler one after starting with Right, where the switching patterns closely match the predictions. Note that this contrasts with Grether (1980), who finds that “individuals tend to give too much weight to the ‘evidence’ and thus too little weight to their prior beliefs, though priors are *not* ignored.” Perhaps the fact that the first ‘draw’ contained information, but no affective element, can help reconcile this difference. In any case, this behavior seems entirely consistent with the status quo bias (Samuelson and Zeckhauser 1988), which predicts that people are reluctant to make changes when risky outcomes are involved.

Another possible explanation for such different error rates (switching behavior after success vs. after failure) beyond differential cost of errors that we discuss below would be that “carrots are more effective than sticks” – positive reinforcement affects behavior more than negative reinforcement does. However, it is far from clear that this is the case. In fact, Baumeister, Bratslavsky, Finkenauer & Vohs (2001) discuss the pervasive evidence that “bad is stronger than good.” Nevertheless, our results in Treatment 3 strongly suggest that psychological affect plays a major role in switching decisions, and that positive affect is a stronger motivational force than negative affect; while removing affect from initial Left draws

²⁸ The authors motivate the concept by pointing to the difference between the sensation of losing a sum of money at a horse race and the sensation of losing the same sum as the result in income tax rates.

reduces the error rate after ‘black’ (favorable) draws by 63% (from 36.8% in Phase I of Treatment 2 to 13.5% in Treatment 3), the corresponding reduction in the error rate after ‘white’ (unfavorable) draws is only 24% (from 55.7% to 42.4%).

VI. Conclusion

We conduct an experiment that permits us to compare the motivating force of BEU to that of reinforcement (or case-based) learning. The heart of our design is to construct situations where these two motivations are aligned and situations where they conflict. Our results indicate that both heuristics are at work, and that when they clash we can expect divergence from BEU. We believe that examining the relative strengths of the conflicting forces is useful for explaining differential intensities of behavior even when the forces do not clash.

When BEU predictions and reinforcement-based predictions agree, nearly all people respond as expected. However, there is a mixture of behavior when these predictions point in opposing directions. Nearly 50% of all switching decisions after Left draws violate the Bayes updating rule in Treatments 1 and 2; this drops substantially in Treatment 3, where the Left initial draw is stripped of its affect. It appears that much of the power of reinforcement comes from the psychological affect induced by an outcome, rather than from a more cortical consideration of one’s received payoffs.

We find a great deal of variation across individual participants. Some people never change their decisions on the basis of realized outcomes. On the other hand, we find strong cross-period updating behavior for many DMs and the tendency to update across periods is correlated with intra-period switching-error rates (particularly after drawing a white ball from the Left urn) in Treatments 1 and 2, but not in Treatment 3. We also find a correlation for individual

tendencies to make within-period switching errors after Right draws and after Left draws. Finally, people are much less likely to change their initial voluntary draw if it was a Right start.

We find evidence of a preference for transparency, as exemplified by a start-from-Right bias. Choosing to choose a ball from the clearer environment helps the decision-maker avoid potential regret and cognitive dissonance. We also observe that people seem to be reluctant to switch urns for their second draw, particularly when there is affect involved or when the location of the first draw was voluntary. This behavior is consistent with the Samuelson and Zeckhauser (1988) status quo bias. We have argued that it is implausible to rationalize the very high switching-error rates using risk aversion. In addition, while error rates are much lower when the applicable Bayesian updating is obvious, we do not find that these rates are reduced when the complexity of the task is reduced rather than eliminated, as in our Treatment 2.

Our initial foray into this area leaves much more work to be done, and we plan to pursue this rich vein. One area concerns differences in individual behavior. With better controls on individual background (beyond gender), one could assess the roles that age, education level and sophistication (e.g., math and stat background) play in the weights assign to the different heuristics when they clash. Another conjecture that emerges and can be tested is that lower animals, say rats, would have even higher switching error rates from Left.

Finally, in our design, the simple errors (i.e., incorrect updating after a Right draw) are the most costly. It should be interesting to see what happens when the simpler decision errors are not so costly (in expectation), while decision errors in more complex environments are relatively expensive. We suspect that the cost of the error is not the true independent variable, as people are hardly calculating the cost of an error and then choosing how careful to be.

References

- Baumeister, R., E. Bratslavsky, C. Finkenauer & K. Vohs (2001), "Bad is Stronger than Good," *Review of General Psychology*, **5**, 323-370.
- Camerer, C. and T. Ho (1998), "Experience-Weighted Attraction in Games: Estimates from Weak-link Games," in D. Budesu and R. Zwick, editors, *Games and Human Behavior, Essays in Honor of Amnon Rapoport*, 31-53, Hillsdale: NJ.
- Camerer, C. and T. Ho (1999), "Experience-Weighted Attraction Learning in Games: A Unifying Approach," *Econometrica*, **67**, 827-874.
- Camerer, C. and M. Weber (1992), "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty*, **5**, 325-370.
- Choi, J., D. Laibson, B. Madrian & A. Metrick (2003), "The Marginal Propensity to Consume Out of Retirement Wealth Shocks is Weakly Negative," mimeo.
- Cheung, Y. and D. Friedman (1997), "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, **19**, 46-76.
- Erev, I. and A. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, **88**, 848-881.
- Fudenberg, D. and D. Levine (1998), *The Theory of Learning in Games*, MIT Press: Cambridge, MA.
- Gilboa, I and D. Schmeidler (1995), "Case-Based Decision Theory," *Quarterly Journal of Economics*, **110**, 605-639.
- Gilboa, I and D. Schmeidler (2001), *A Theory of Case-Based Decisions*, Cambridge: Cambridge University Press.
- Grether, D. (1980), "Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence," *Journal of Economic Behavior and Organization*, **17**, 31-57.
- Grether, D. (1992), "Testing Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, **95**, 537-557.
- Kahneman, D. and A. Tversky (1972), "Subjective Probability: A Judgment of Representativeness," *Cognitive Psychology*, **3**, 430-454.
- Loomes, G. and R. Sugden (1982), "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *Economic Journal*, **92**, 805-824.
- Ouwensloot, H., P. Nijkamp & P. Rietveld (1998), "Errors in Probability Updating Behaviour: Measurement and Impact Analysis," *Journal of Economic Psychology*, **19**, 535-563.

- Roth, A. and I. Erev (1995), "Learning in Extensive-form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, **8**, 164-212.
- Samuelson, W. and R. Zeckhauser (1988), "Status Quo Bias in Decision-Making," *Journal of Risk and Uncertainty*, **1**, 7-59.
- Suppes, P. and R. Atkinson (1960), *Markov Learning Models for Multiperson Interaction*, Stanford University Press, Palo Alto, CA.
- Tversky, A. and D. Kahneman (1971), "Belief in the Law of Small Numbers," *Psychological Bulletin*, **76**, 105-110.
- Tversky, A. and D. Kahneman (1973), "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology*, **5**, 207-232.
- Zizzo, D, S. Stolarz-Fantino, J. Wen & E. Fantino (2000), "A Violation of the Monotonicity Axiom: Experimental Evidence on the Conjunction Fallacy," *Journal of Economic Behavior and Organization*, **41**, 263-276.