

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

By GARY CHARNESS AND DAN LEVIN*

An important issue in decision making concerns the manner in which people process new information and update prior beliefs. The Bayesian updating rule, in combination with expected utility theory, is ubiquitous in economic theory, and its application is an important paradigm for examining decision making under risk. A number of experimental studies suggest, however, that people may often ignore prior information when forming beliefs, contrary to Bayes's rule.¹ Another heuristic for processing new information involves some form of *reinforcement*, where one is more likely to pick

choices (actions) associated with successful past outcomes than choices associated with less successful outcomes.²

These separate approaches often prescribe a similar course of action in the face of new information; however, this is not always the case. We construct an individual choice task in which Bayesian updating with expected utility maximization (BEU) sometimes coincides and sometimes opposes a "win-stay lose-shift" heuristic. Here, 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 how one's propensity to use Bayes's rule is affected by whether this rule is aligned with the win-stay lose-shift heuristic or clashes with it. We also consider whether this propensity differs according to whether the information provided by an earlier outcome is accompanied by payment for the outcome and the corresponding feelings of success or failure.

To the best of our knowledge, this is the first study to examine explicitly what happens when these forces work against one another.³ Our constructed case where the reinforcement heuristic leads one astray can be applied more generally to situations where favorable direct

*Charness: Department of Economics, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210 (e-mail: charness@econ.ucsb.edu); Levin: Department of Economics, The Ohio State University, 1945 N. High Street, Columbus, OH 43210-1172 (e-mail: dlevin@econ.ohio-state.edu). We thank Mark Brinkman and Arun Qamra for their help with designing the software and conducting the experimental sessions. We thank Doug Bernheim and an anonymous referee for valuable editorial comments. We have benefited from discussions with, and suggestions from, George Akerlof, Ted Bergstrom, Antonio Cabrales, James Choi, Carl Christ, Stefano DellaVigna, Matthew Ellman, Ido Erev, Guillaume Fréchette, Rod Garratt, Itzhak Gilboa, Robin Hogarth, Edi Karni, Botond Kozzegi, David Laibson, Cade Massey, Jim Peck, Matthew Rabin, David Schmeidler, Bill Zame, Richard Zeckhauser, and seminar participants at the University of California, Santa Barbara, Columbia University, Harvard University, the Econometric Society Summer 2003 meeting, the Stony Brook 2003 Workshop on Experimental Economics and Game Theory, University of California-Berkeley, The Ohio State University, the Fuqua School at Duke University, the 2004 SITE Workshop in Psychology and Economics, Universitat Pompeu Fabra, University of Montpellier, GATE in Lyon, the University of Zürich, and Johns Hopkins 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.

¹ See, for example, Amos Tversky and Daniel Kahneman (1971, 1973), Kahneman and Tversky (1972), and David M. Grether (1980, 1992). More recent studies (e.g., Hans Ouwensloot et al., 1998; Daniel John Zizzo et al., 2000) also provide strong evidence that experimental subjects are often not even close to being "perfect Bayesians."

² The basic reinforcement learning models (e.g., Alvin E. Roth and Ido Erev, 1995; Erev and Roth, 1998) assume an initial propensity for a particular choice and utilize a payoff sensitivity parameter. Colin Camerer and Teck-Hua Ho (1998, 1999) combine reinforcement and belief learning by using experience-weights and updated levels of attraction. Case-based decision theory (Gilboa and Schmeidler, 1995, 2001) formalizes the thrust of reinforcement heuristics in a nonexpected 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.

³ Note that our setup is quite different from the "two-armed bandit" problem, where the separate machines have independent distributions rather than a common state.

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 a rookie with a wage of \$15 per hour, while the second is an experienced veteran with a wage of \$30 per hour. Suppose, further, that the veteran was unavailable in the first week, so that the agency had to send out the rookie. As it happens, the rookie's campaign was quite successful; the agency must now decide whom to send out in the second week. The first reaction might be: "Well the rookie did so much better than expected that we ought to send him/her again." Upon reflection, however, it may occur to the agency that the unexpected successes may also contain information that the relevant business conditions 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.⁴

In our design, each participant chooses a "ball" from either the left or the right "urn," 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. We are interested in two kinds of choices that a participant makes: starting choices refer to the choice of urn from which to draw the first ball, whereas switching choices refer to the choice of urn from which to draw the second ball, after having experienced the first draw.

We found considerable non-BEU behavior in our first treatment when the two heuristics—Bayesian updating and reinforcement—clashed. We then designed two additional treatments to gain further insight. In our second treatment, we increased the informative content

of the information of the first draw from the left that also simplified the calculations for the updating task, while controlling for "emotional reinforcement" (*affect*). We expected that this treatment would reduce the overall error rate. In our third treatment, we maintained the distribution of balls in the second treatment, but eliminated much of the affect associated with this draw by not paying for the realized outcome and not even associating it with success or failure.⁵ 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.⁶

We observe striking patterns: when payoff reinforcement and Bayesian updating are aligned, nearly all people respond as expected; on the other hand, when these forces clash and people are paid for their initial choice, nearly half of all decisions are inconsistent with Bayesian updating. While we find shifts in the predicted direction for both starting choices and switching choices in Treatment 2, the overall error rate was not reduced. In Treatment 3, however, when a draw provides only information, switching errors occur much less frequently, suggesting that the affect induced by payments is a critical factor in deviations from Bayesian updating.

There is a strong correlation between the likelihood of a deviation from BEU optimization and the cost of the deviation, as might well be expected by economists. This is not the only factor influencing the observed behavior, however, as the effect of affect indicates. There is considerable behavioral heterogeneity in our experimental population. We also find a gender effect, in that women in our sample are more likely to deviate from BEU behavior.

In Section I we describe the design of our

⁵ We thank Bill Zame for suggesting the essence of our Treatment 3.

⁶ It is standard to consider reinforcement only in terms of material payoffs, although some psychology studies (e.g., Patrick Suppes and Robert C. Atkinson, 1960) ignore material payoffs, arguing that a "win" is a reinforcing event, in and of itself. Experimental work on learning in games by reinforcement rarely attempts to associate the choice itself with affect, and we are not aware of any reinforcement model that considers affect in the payoffs.

⁴ We thank Richard Zeckhauser for providing the essence of the example above.

Treatment 1			Treatments 2 and 3		
	Left urn	Right urn		Left urn	Right urn
Up ($p = 0.5$)	●●●●○○	●●●●●●	Up ($p = 0.5$)	●●●●○○	●●●●●●
Down ($p = 0.5$)	●●●○○○	○○○○○○	Down ($p = 0.5$)	●●○○○○○	○○○○○○○

FIGURE 1. DISTRIBUTION OF BALLS, BY TREATMENT AND STATE

experiments. In Section II we present our results and consider the underlying determinants of the observed behavior. Section III concludes and points to future research.

I. Experimental Design

We conducted Web-based sessions on the University of California, Santa Barbara, campus, with students recruited by e-mail from the general student population. Average earnings were \$23.14, \$22.06, and \$17.82 for Treatments 1, 2, and 3, respectively; sessions averaged 45 minutes in duration. Participants met in the lab and were given a handout explaining the experiment setup; detailed, hands-on instructions were provided on the Web site, and participants were required to correctly answer questions testing comprehension.⁷

In our design, there are two equally likely states of the world, Up and Down, and two lotteries, Left and Right, consisting of “urns” from which the individual can draw “balls” that are black or white. There are more black balls with Up than with Down in both the Left and Right urns; also, there is always a mix of colors in the Left urn, while the Right urn contains balls of only one color. Figure 1 shows the distribution on balls in our three treatments.

In all cases, the Right urn contains only black balls in Up and only white balls in Down. In Treatment 1, the Left urn has four black balls and two white balls in state Up, compared to three balls of each color in state Down. In Treatments 2 and 3, the Left urn has four black

and two white balls in state Up, compared to two black and four white balls in state Down. Note that a draw from Left provides a more precise signal about the state of the world in Treatment 2.⁸

A decision maker 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 that is clearly explained. Thus, each person has the choice whether to stay with the same-side lottery or switch to the other-side lottery for the second draw.

In Treatments 1 and 2, participants are paid for drawing black balls. In our Treatment 3, we attempted to avoid any form of affect as the result of the outcome of the first draw. One obvious consideration is whether an individual is actually paid on the basis of the outcome. There is a more subtle emotional response, however, that may result from the first draw: one may still receive a feeling of success or failure when observing the first draw if one immediately recognizes that the color that was drawn is good or bad. To avoid this, we did not tell the individual whether black or white would pay on the second draw until *after* the first ball was drawn. This was feasible without loss of generality due to the symmetry of the distribution of black and white balls across the Up and Down states in this treatment.

We summarize the particulars of our treatments in Table 1.

There were 60 periods in each of Treatments 1 and 2. We wished to ensure that participants gained some familiarity with a variety of strat-

⁷ The Web-based instructions can be found at <http://www.econ.ucsb.edu/~gcsurvey/Bayesian Updating/>. The supplemental instructions can be found in Charness and Levin (2003), along with a considerably more detailed presentation of the results.

⁸ A black draw from Left means $p(\text{Up}) = 4/7$ in Treatment 1, compared to $2/3$ in Treatment 2; for white draws, $p(\text{Up}) = 2/5$ in Treatment 1, compared to $1/3$ in Treatment 2.

TABLE 1—TREATMENT SUMMARY

Treatment	First draw restriction	Payment	No. of participants	No. of periods
1	Alternate in first 20 periods	Both draws	59	60
2	Alternate in first 20 periods	Both draws	54	60
3	Left draw only	Second draw only	52	80

egies and outcomes, and also wished to insure that we would have a number of observations on *switching choices* after draws from both urns. Therefore, during the first 20 periods, we required participants to start with the Left (Right) urn in each odd (even) period; in the subsequent periods people chose freely. In Treatment 3, since we paid only for successful second draws and wished to keep the same marginal incentives, we increased the number of periods in each session to 80 to even payoffs across sessions to some extent. In each of these periods, we required the first draw to be made from the Left, but allowed freedom of choice for the second draw.

In the first 50 periods of Treatments 1 and 2 (and the first 70 periods of Treatment 3), each black ball drawn from the Left urn paid one experimental unit (\$0.30), while each black ball drawn from the Right urn paid 7/6 experimental units (\$0.35). This parameterization meant that the expected payoff of an uninformed draw in Treatment 1 was the same from Right or Left.⁹

With these payoff parameters, a risk-neutral BEU participant should start with the Right urn whenever possible. Since any ball drawn from the Right urn fully resolves the uncertainty regarding the state of the world (Up or Down), choosing the urn for the second draw is, in this case, a simple task—switch to the Left urn after failure and stay with the Right urn after success.

The main innovation of our design is the switching choice after a participant draws the first ball from the Left urn. Starting with the Left urn may be required or may reflect a BEU error. After starting with Left, the expected payoff is highest by drawing a second time from

Left after an unsuccessful (or would-be unsuccessful, in the case of Treatment 3) first draw, but to switch to Right after a successful (would-be successful) first draw, thus reversing the win-stay lose-shift heuristic! In Treatment 1, switching to Right after a black draw gives an expected payoff of 28/42, compared to 25/42 from staying with Left; staying with Left after a white draw gives an expected payoff of 17/30, compared to 14/30 from switching to Right.¹⁰ In Treatments 2 and 3, switching to Right after a black/favorable draw gives an expected payoff of 7/9, compared to 5/9 from staying with Left; staying with Left after a white/unfavorable draw gives an expected payoff of 8/18, compared to 7/18 from switching to Right.¹¹

While observing the outcome of the first Left draw yields information about the state of the world, it does not completely resolve uncertainty as a first draw from Right does. The more precise information about the state of the world from a Right initial draw improves the expected payoffs from the second draw. Taking everything into account, in Treatment 1 starting with Right and updating correctly gives an expected payoff of 102/72 for the two draws, while starting with Left and updating correctly gives an expected payoff of 87/72; the corresponding comparison for Treatment 2 is 12/9 versus 10/9.¹²

¹⁰ The calculations after a black initial draw: $pr(\text{Up}|\text{black}) = 4/7$. Drawing from Left yields expected payoffs of $(4/7) \cdot (2/3) + (3/7) \cdot (1/2) = 25/42$. Drawing from Right gives expected payoffs of $(7/6) \cdot (4/7) = 28/42$.

¹¹ We maintained the same payoff structure for successful draws from (Left, Right) in Treatments 2 and 3 to maintain comparability with Treatment 1. Otherwise, we could have paid the same in the later treatments for success from either side, while retaining the feature that the two heuristics clash after a Left initial draw.

¹² Note that the switching decision after a black draw from Left in Treatments 2 and 3 is unaffected by risk preferences, since we have first-order stochastic dominance.

⁹ During the final ten periods, a successful draw from Left paid 7/6 units and a successful draw from Right paid one unit. In the interest of space and clear exposition, we will largely ignore the results from these ten periods with the reversed payoffs in all three treatments.

TABLE 2—SWITCHING-ERROR RATES AFTER RIGHT DRAWS

Treatment	Good outcome, forced draw	Good outcome, voluntary draw	Bad outcome, forced draw	Bad outcome, voluntary draw	Aggregate error %
1	4.2% (12/288)	3.4% (20/596)	5.3% (16/302)	3.8% (26/683)	4.0%
2	12.9% (33/255)	4.8% (31/642)	3.9% (11/285)	2.6% (16/627)	5.0%

TABLE 3—SWITCHING-ERROR RATES AFTER LEFT DRAWS

Treatment	Good outcome, forced draw	Good outcome, voluntary draw	Bad outcome, forced draw	Bad outcome, voluntary draw	Aggregate error %
1	51.5% (171/332)	66.2% (180/272)	38.0% (98/258)	29.2% (64/219)	47.5%
2	36.8% (95/258)	48.4% (75/155)	56.4% (159/282)	51.5% (101/196)	48.3%
3	13.5% (245/1811)	—	42.4% (780/1829)	—	28.2%

II. Results

Our benchmark for analyzing our results is the behavior of a risk-neutral BEU decision maker, and we define “errors” to be deviations from that benchmark. Tables 2 and 3 below present aggregated data for each treatment; note that in Treatment 3 there are no starting errors or switching errors after Right draws, since people were required to draw first from the Left.

RESULT 1: *Switching-error rates are very low when Bayesian updating and a reinforcement heuristic are aligned, but are quite large when these are opposed.*

The reinforcement and Bayesian heuristics are aligned after an initial Right draw, and it is clear that switching errors are rare in this case; the aggregate error rate after Right draws is no more than 5 percent in any treatment. Even this low proportion masks the fact that there are “worst offenders,” with 12 percent of partici-

pants making about 75 percent of all switching errors after a Right draw; apart from those people, the aggregate error rate is between 1 percent and 2 percent in every category.

On the other hand, the two heuristics clash after an initial draw from the Left. We predicted “poorer” performance when one is “pulled” in opposite directions, as Bayesian updating directs one to switch (not switch) after a success (failure). Indeed here we see a very different picture. In Treatments 1 and 2, when people are paid for their initial draws, the switching-error rates after Left draws are substantial, ranging between 29 percent and 66 percent. Comparing each individual’s switching-error rates, we see that the error rate is higher after Left draws than after Right draws for 111 of 113 individuals ($Z = 10.44$, $p = 0.000$, binomial test).

RESULT 2: *Removing affect from the initial draw (by not paying for its outcome and not associating it with success or failure) reduces the error rate, particularly in the case of positive affect.*

The reduction in switching-error rates from Treatment 2 to Treatment 3 is remarkable, since the payoff structure is identical across these treatments. This reduction is particularly dramatic after black/favorable draws from Left, dropping to 13.5 percent from 36.8 percent in Treatment 2, when the first draw is forced. The switching-error rate also drops after white/unfavorable initial Left draws, but this decline is

Furthermore, risk aversion cannot reverse the optimal choice after a white draw from Left, since a risk-averse person should favor the Left side even more. In Treatment 1, if participants have CRRA utility, we must have $\rho > 0.735$ to reverse the optimal choice; a person with this CRRA coefficient would prefer \$10 for certain over a lottery that gives a 50-percent chance of \$137 and a 50-percent chance of \$0. Thus, we feel that while risk aversion may influence some decisions, it is not at the heart of the phenomena we describe.

smaller, from 56.4 percent to 42.4 percent.¹³ A conservative Wilcoxon-Mann Whitney rank-sum test (see Sidney Siegel and N. John Castellan, Jr., 1988) comparing the individual error rates for these decisions indicates a significant difference in behavior across treatments after both black/favorable first draws ($Z = 2.29$, $p = 0.022$, two-tailed test) and white/unfavorable first draws ($Z = 2.13$, $p = 0.033$, two-tailed test).

RESULT 3: *Increasing the informativeness of Left draws does not reduce the overall switching error rate from its Treatment 1 level. On the other hand, there is a clear decrease in the likelihood of drawing from Left when the expected number of black balls in the Left urn is reduced, both for starting and switching decisions.*

We see that the overall switching-error rate after a Left draw is nearly identical in Treatment 1 and Treatment 2, about 48 percent in each case. Our conjecture had been that the greater precision of the information obtained with a Left draw and the simpler calculations would result in fewer switching errors in Treatment 2. Whenever an error rate is close to 50 percent, one is tempted to conclude that play is random. A more careful examination, however, reveals that this 50-percent aggregated switching-error rate is an artifact of the data, as rates are very different after black or white first draws in each treatment, and also vary broadly over the population.¹⁴

¹³ Matthew Ellman points out that, after an unsuccessful draw, negative affect might be a stronger force if people actually suffered a loss, rather than merely not being rewarded.

¹⁴ The switching-error rate after a Left draw in Treatment 1 was 58.1 percent (351/604) when a black ball was drawn and 34.0 percent (162/477) when a white ball was drawn; the error rate after drawing black was higher than the rate after drawing white for 37 people, and vice versa for 18 people. A binomial test rejects the premise that the switching error rates are the same ($Z = 2.56$, $p = 0.010$). In Treatment 2, the switching-error rate after drawing black (staying Left) is reduced to 41.2 percent (170/413) and the rate after drawing white (switching to Right) increases to 54.4 percent (260/478). The error rate after drawing black was lower than the rate after drawing white for 31 people, and vice versa for 16 people; a

Note that the level of each of these switching-error rates shifts by 15 to 20 percentage points toward more Right draws, consistent with the black balls being relatively scarcer on the Left in Treatment 2. This relative scarcity is also reflected in starting choices, as 78.3 percent of all voluntary first draws were (correctly) made from the Right in Treatment 2, compared to 72.3 percent in Treatment 1.¹⁵ The Wilcoxon rank-sum test comparing each individual's overall starting-error rate indicates only marginal statistical significance ($Z = 1.26$, $p = 0.104$, one-tailed test), although a random-effects probit regression that includes every starting choice indicates a significant difference in starting-error rates across treatments ($Z = 4.09$, $p = 0.000$).

RESULT 4: *There is a taste for consistency. A person who has elected to make the first draw from Left is substantially more likely to make the second draw from Left than a person who was required to make the first draw from Left.*

It is interesting to compare switching-error rates after voluntary and required initial draws from Left. We might expect people who voluntarily start with the Left urn, an error, to be more likely also to make switching errors. While this is the case after a black draw, the reverse is true after a white draw. On the other hand, we do see evidence of a "taste for consistency" (Erik Eyster, 2002). Aggregating over the four comparisons available in Table 3, a person who has *elected* to make the first draw from Left is about 20 to 25 percent more likely to make a second draw from Left than a person who was required to start with Left. A random-effects probit regression (with robust standard errors) with a second Left draw as the dependent variable confirms this relationship, as the coefficient for the dummy variable for a voluntary first draw is highly significant ($Z = 2.87$, $p = 0.002$).

binomial test rejects the premise that the rates are the same ($Z = 2.19$, $p = 0.029$).

¹⁵ We observed a similar shift in starting tendencies in the last ten periods of these treatments, with 54.4 percent starting with Right in Treatment 1 and 62.8 percent starting with Right in Treatment 2.




FIGURE 2. COST AND FREQUENCY OF ERRORS

RESULT 5: *The cost of an error has a strong influence on the frequency of the error. However, the presence of affect for the first draw, the transparency of the updating (Left versus Right), and the gender of the participant also appear to play important roles.*

A natural question for economists is whether the frequency of decision errors is inversely related to the cost (reduction in expected payoffs) of such errors, even if complex calculations may be beyond the ability of the general population. In Figure 2, we graph the frequency of each of the possible types of decision error against its cost in expected payoffs. Each point represents the aggregated frequency of one type of error (different across treatments, restrictions, affect, etc.) over the entire relevant population.

While the relationship is not completely smooth, we do find that the higher the cost of an error, the less likely it is to be made. Thus, overall, people are responding to the economic incentives provided. Nevertheless, there are some patterns to the deviations from this cost/frequency relationship. For example, the reduction in one's expected payoff from a second draw from the Left urn after observing black is 0.222 in both Treatment 2 and Treatment 3, yet the aggregate frequencies differ considerably.

We attempt to diagnose the underlying determinants of deviations from the BEU predictions by using a regression in which the frequency of error is the dependent variable. We consider each individual's error frequency for each possible updating error, and use a random effects specification to account for the multiple obser-

vations for each individual. In the regression below, frequency refers to each individual's aggregated error rate for each possible error, ranging continuously from 0 to 1. Standard errors are in parentheses, $N = 1362$, and the Overall $R^2 = 0.266$. Left = 1 if the first draw was from the Left urn and is 0 otherwise. Affect = 1 if the participant was paid for the first draw and is 0 otherwise. And, Female = 1 if the participant was female and is 0 otherwise.

$$\begin{aligned} \text{Frequency} = & 0.317 - 1.49 * \text{Cost} \\ & (0.045) \quad (0.203) \\ & + 1.31 * \text{Cost}^2 + .097 * \text{Left} \\ & (0.334) \quad (0.024) \\ & + .099 * \text{Affect} + .081 * \text{Female} \\ & (0.031) \quad (0.025) \end{aligned}$$

We see that all of the coefficients are highly significant. There is a strong influence from the cost of an error, with the marginal effect declining as the cost increases. Even when the cost effect is taken into account, the updating error rate is substantially and significantly higher when the first draw is from the Left. The regression confirms effects of similar size and direction for affect being present from the first draw, and for female participants.¹⁶ It would appear that the reinforcement heuristic is relatively stronger for women in our study.

RESULT 6: *We generally see considerable behavioral sensitivity to previous outcomes.*

To what extent are people 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? Of the 84 of 113 (74 percent) participants in Treatments 1 and 2 who make initial voluntary draws from both sides, 71 (85 percent) are less likely to change their starting choice after two good outcomes in the previous period than after fewer good outcomes. A binomial test confirms that this differs

¹⁶ The rate of deviation from BEU predictions is higher for women in 20 of the 24 comparisons available. A binomial test rejects the hypothesis of no difference at $p = 0.002$, two-tailed test.




FIGURE 3. LEFT SWITCHING-ERROR RATES OVER TIME

significantly from random behavior ($Z = 6.33$, $p = 0.000$).

We can calculate the difference in the likelihood that a person who drew first from 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.¹⁷ We do so using only unrestricted first draws, obtaining comparisons for 72 individuals. Of the 56 people for whom there was a difference, 37 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 ($Z = 2.41$, $p = 0.016$). The positive difference in rates was at least 25 percentage points for 24 people (33 percent).¹⁸

RESULT 7: *The observed behavior changes little over the course of our sessions.*

An important consideration is whether the behavior that we observe generally persists into later periods. Switching-error rates after Left draws do not appear to diminish over time; we show the switching-error rates for time segments of each treatment in Figure 3. Each point represents the aggregated population switching-error rate in the range of periods shown.

The switching-error rate is either slightly increasing over time or is relatively flat after Left




FIGURE 4. INDIVIDUAL STARTING-ERROR RATES

draws in all three treatments. Thus, overall we see little in the data to indicate that these switching-error rates drop over time. The switching-error rates after Right draws generally show a slightly decreasing trend, but actually increase over time after a white draw in Treatment 1. There is little change in starting choices over time.¹⁹

RESULT 8: *There is considerable heterogeneity in behavior across individuals*

The aggregate error rates in Tables 2 and 3 do not reflect the considerable heterogeneity in behavior across individuals, particularly in starting-error rates and switching-error rates, after Left draws in Treatments 1 and 2. We show this with histograms that display individual aggregate starting-error rates and switching-error rates after mandatory Left draws (Figures 4 and 5).

We see a good number of individuals with overall error rates in nearly every range shown; the diversity is particularly apparent in Figure 5.

III. Conclusion

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 a simple hedonic capacity suffices, no elaborate cognitive process is needed. In contrast,

¹⁷ We thank Stefano DellaVigna for suggesting this idea.
¹⁸ Prior outcomes also affect behavior in Treatment 3. Forty-seven of 52 persons (90 percent) were more likely to change their second draw if it was unsuccessful than if it was successful, significantly different from random behavior ($Z = 5.82$, $p = 0.000$).

¹⁹ For example, in periods 21 to 30 of Treatment 1, 70.0 percent of the starts were from Right, increasing slightly to 72.2 percent in periods 31 to 40, and to 74.4 percent in periods 41 to 50. In periods 21 to 30 of Treatment 2, 78.9 percent of starts were from Right, compared to 77.8 percent in periods 31 to 40 and to 78.3 percent in periods 41 to 50.




FIGURE 5. INDIVIDUAL LEFT SWITCHING-ERROR RATES

the ability to update priors, even if in an imperfect (Bayesian) way, seems to require enormously more. Good solutions to many decision problems in general, and economic ones in particular, require one to update priors when new information arrives.

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. There is a mixture of behavior, however, when these predictions point in opposing directions. Nearly 50 percent of all switching decisions after Left draws violate the Bayes's updating rule in both Treatment 1 and Treatment 2; the similarity across treatments indicates that overall error rates are not reduced when the precision of the information is slightly increased. The error rate drops substantially, however, when the initial draw is largely 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.

Our results, thus, 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 in error rates over time. On the other hand, we see a strong influence from affect, which could be considered a "sensation-based" form of reinforcement. We find that removing positive affect has a stronger effect than removing negative affect—reducing the error rate by 63 percent after a successful Left draw versus 24 percent after an unsuccessful Left draw. Since affect seems to play such a large role, perhaps reinforcement models should consider it to be part of the overall "payoff."

One might question how common it is in real economic environments for the Bayesian updating rule to clash with the "stay after success but switch after failure" heuristic. We suspect that in most situations both heuristics are at least partially aligned, but we also suspect that there tends to be at least a kernel of updating from such inferences in many decisions. Consider, for example, an investor whose recent portfolio has performed very well, beyond investor expectations. James T. Choi et al. (2003) find the puzzling result that people who experience higher appreciation than expected in their 401k retirement accounts do not increase their consumption, in violation of standard neoclassical theory. 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. Based on our data, we would further conjecture that investors who personally select their own 401k portfolios would be particularly prone to this influence.

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 one's age, education level, and sophistication (e.g., math and stat background) play in the weights assigned to the different heuristics when they are opposed. Another conjecture that emerges and can be tested is that lower animals, say rats, would have even higher switching error rates from Left.

We note that overall error rate was not lower in Treatment 2 than in Treatment 1, even though Treatment 2 featured an increase in the informativeness of Left draws and a simplification in

calculations, while errors are rare in the transparent situation after a Right draw. Further research might try to find the “cross-over” threshold between simple updating and difficult updating by varying the distribution of balls in the urns and the payoffs for a successful draw (e.g., what would happen if the Right urn contained nine balls of one color and one ball of the other color?). Another issue is the effect of imposing a cost for switching (or the opposite).

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 may not be the true independent variable, as people are hardly calculating the cost of an error and then choosing how careful to be.

REFERENCES

- Camerer, Colin and Ho, Teck-Hua.** “Experience-Weighted Attraction in Games: Estimates from Weak-Link Games.” In David Budescu, Ido Erev, and Rami Zwick, eds., *Games and human behavior: Essays in honor of Amnon Rapoport*. Hillsdale, NJ: Lawrence Erlbaum Associates (LEA), 1998, pp. 31–53.
- Camerer, Colin and Ho, Teck-Hua.** “Experience-Weighted Attraction Learning in Games: A Unifying Approach.” *Econometrica*, 1999, 67(4), pp. 827–74.
- Charness, Gary and Levin, Dan.** “When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect.” University of California at Santa Barbara, Economics Working Paper Series: No. 1164, 2003.
- Choi, James T.; Laibson, David; Madrian, Brigitte and Metrick, Andrew.** “The Marginal Propensity to Consume Out of Retirement Wealth Shocks Is Weakly Negative.” Unpublished Paper, 2003.
- Erev, Ido and Roth, Alvin E.** “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria.” *American Economic Review*, 1998, 88(4), pp. 848–81.
- Eyster, Erik.** “Rationalizing the Past: A Taste for Consistency.” Unpublished Paper, 2002.
- Gilboa, Itzhak and Schmeidler, David.** “Case-Based Decision Theory.” *Quarterly Journal of Economics*, 1995, 110(3), pp. 605–39.
- Gilboa, Itzhak and Schmeidler, David.** *A theory of case-based decisions*. Cambridge: Cambridge University Press, 2001.
- Grether, David M.** “Bayes Rule as a Descriptive Model: The Representativeness Heuristic.” *Quarterly Journal of Economics*, 1980, 95(3), pp. 537–57.
- Grether, David M.** “Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence.” *Journal of Economic Behavior and Organization*, 1992, 17(1), pp. 31–57.
- Kahneman, Daniel and Tversky, Amos.** “Subjective Probability: A Judgment of Representativeness.” *Cognitive Psychology*, 1972, 3(3), pp. 430–54.
- Ouwensloot, Hans; Nijkamp, Peter and Rietveld, Piet.** “Errors in Probability Updating Behaviour: Measurement and Impact Analysis.” *Journal of Economic Psychology*, 1998, 19(5), pp. 535–63.
- Roth, Alvin E. and Erev, Ido.** “Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term.” *Games and Economic Behavior*, 1995, 8(1), pp. 164–212.
- Siegel, Sidney and Castellan, N. John, Jr.** *Nonparametric statistics for the behavioral sciences*. New York: McGraw Hill, 1988.
- Suppes, Patrick and Atkinson, Robert C.** *Markov learning models for multiperson interaction*. Palo Alto: Stanford University Press, 1960.
- Tversky, Amos and Kahneman, Daniel.** “Belief in the Law of Small Numbers.” *Psychological Bulletin*, 1971, 76(2), pp. 105–10.
- Tversky, Amos and Kahneman, Daniel.** “Availability: A Heuristic for Judging Frequency and Probability.” *Cognitive Psychology*, 1973, 5(2), pp. 207–32.
- Zizzo, Daniel J.; Stolarz-Fantino, Stephanie; Wen, Julie and Fantino, Edmund.** “A Violation of the Monotonicity Axiom: Experimental Evidence on the Conjunction Fallacy.” *Journal of Economic Behavior and Organization*, 2000, 41(3), pp. 263–76.