UC Santa Cruz
UC Santa Cruz Previously Published Works

Title
DISSECTING THE MONTY HALL ANOMALY

Permalink
https://escholarship.org/uc/item/7tp7c0rk

Journal
Economic Inquiry, 56(3)

ISSN
0095-2583

Authors
James, Duncan
Friedman, Daniel
Louie, Christina
et al.

Publication Date
2018-07-01

DOI
10.1111/ecin.12533

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Economic Inquiry

DISSECTING THE MONTY HALL ANOMALY

DUNCAN JAMES, DANIEL FRIEDMAN, CHRISTINA LOUIE and TAYLOR O'MEARA*

We assess competing explanations of irrational behavior in the Monty Hall problem by creating new variants of the problem. Some variants employ a feature that automates the merging of probabilities, thus rendering transparent the probabilistic advantage of the rational choice. That feature also enables systematic variation in informational asymmetry, and in ordering of actions. Data from 77 subjects, each of whom makes 30 binary decisions, indicate that automated merging raises the fraction of rational choices from around 40% to over 80%. Other features examined have much less impact, indicating the importance of a Bayesian updating failure. (*JEL* C91, D02, D81, D83)

I. INTRODUCTION

Nature randomly chooses, with equal probability, one of three doors behind which to put a valuable prize; there is nothing behind the other two doors. A player initially chooses one of the doors, say door i. Nature then opens a door j subject to two constraints: that $j \neq i$ and that j is not the prize door. Due to those constraints, the probability that i is the prize door remains at 1/3. The player now has the option to open either door i or to open the remaining door, $k \neq i,j$. Switching to door k wins with probability $1-1/3=2/3$, and so the rational choice is to switch.

This "Monty Hall problem" (Nalebuff 1987; Selvin 1975; vos Savant 1990/1991) has attracted continuing interest because in most empirical tests the majority of humans do not switch, but rather stick with their original choice (e.g., Friedman 1998; Granberg 1999; Krauss and Wang 2003; Palacios-Huerta 2003). The reasons for this irrational behavior remain unclear. Leading candidates include the illusion of control, escalation of commitment, confirmation bias, reacting to an informational asymmetry, improper use of Bayes rule, and probability matching (e.g., Bazerman 1990; Camerer 1995; Camerer and Weber 1999; Estes 1954; Friedman 1998; Wason 1960).

Much is at stake in evaluating these competing explanations. If the main cause is improper Bayesian updating following restricted choice, then the Monty Hall problem is a decision-theoretic manifestation of longstanding statistical puzzles in education, psychology, and sports science (Miller and Sanjurjo 2015) and in medical decision-making (Cox et al. 2016). If it is a more general failure of hypothetical thinking, then we get new insight into economically important phenomena such as the "winner's curse" in common value auctions and a variety of other "cursed equilibria" (Esponda and Vespa 2014; Eyster and Rabin 2005). On the other hand, if the main cause is escalation of commitment, then we may gain new insight into how to avoid disputable investments like the Space Shuttle program, or the Concorde supersonic passenger plane (Arkes and Blumer 1985; Heath 1995).

Our paper assesses the competing explanations by offering human subjects ten different variants of the Monty Hall problem, each with round-by-round feedback, and comparing the switch (i.e., rational choice) rates. The next

ABBREVIATION

LPM: Linear Probability Model

*D.F., C.L., and T.O'M. are grateful to the National Science Foundation for financial support under grant SES-1357867.

James: Associate Professor, Department of Economics, Fordham University, Bronx, NY 10458. Phone (718) 817-4053, Fax (718) 817-3518, E-mail dujames@fordham.edu

Friedman: Distinguished Professor, Department of Economics, University of California, Santa Cruz, Santa Cruz, CA 95060. Phone (831) 295-1733, Fax (831) 459-5077, E-mail dan@ucsc.edu

Louie: LEEPS Lab Research Assistant, Department of Economics, University of California, Santa Cruz, Santa Cruz, CA 95060. Phone (626) 297-0062, Fax (831) 459-5077, E-mail christina.louie18@gmail.com

O'Meara: LEEPS Lab Research Assistant, Department of Economics, University of California, Santa Cruz, Santa Cruz, CA 95060. Phone (916) 380-2080, Fax (831) 459-5077, E-mail tomeara89@gmail.com

Economic Inquiry
(ISSN 0095-2583)
Vol. 56, No. 3, July 2018, 1817–1826

doi:10.1111/ecin.12533
Online Early publication December 4, 2017
© 2017 Western Economic Association International

section shows how the variants unbundle and rearrange key features of the usual Monty Hall problem. The following section poses the competing explanations as hypotheses that can be tested on our data, and presents the results. The penultimate section shows how the merging-of-probability-masses operation tested in our experiments might be useful in the field, and would thus be a candidate for inclusion in some decision-making expert systems. A concluding discussion summarizes implications and suggests future work.

II. DESIGN

The original Monty Hall task has the following timeline:

1. The true state is determined from among three equally likely states (say Blue, Red, and Yellow).
2. The subject initially nominates one of these states (say $i =$ Blue).
3. The conductor privately learns the true state (this may precede step 2).
4. The conductor discloses a remaining state k (say, Red) subject to two constraints:
   1. $k \neq i$ that is, it cannot be the state initially nominated, and
   2. k is not the true state.
5. The subject is given the opportunity to Switch her nomination from i (the state initially nominated, here Blue) to the remaining state j (Yellow in the current example), or to Stay with the initial nomination.
6. The true state is revealed, and the subject wins a valuable prize if it is her final nomination.
   (a) She wins on j (here, Yellow) if and only if she switched in step 5
   (b) She wins on i (here, Blue) if and only if she did not switch in step 5.

Subjects may misconstrue step 4 as reallocating the 1/3 prior probability of the disclosed state k equally to the remaining state j and to the initially nominated state i, and so believe that the posterior probabilities are .50:.50. But actually the probability of the initially nominated state (1/3) does not change, since constraint 4a *sequesters that state from the updating process*. In effect, then, step 4 just shifts the prior 1/3 probability from the disclosed state k to the remaining state j, boosting its posterior probability to 2/3.

There is a more transparent way to present the apportionment of probability mass implied by step 4. Instead of disclosing k, the conductor can tell the subject that if she chooses to Switch from i, then she wins the prize if the true state is *either* of the two states not nominated initially. In the example where the initial nomination is Blue, a subject who chooses Switch would win if the true state is *either* Red or Yellow, while a subject who chooses Stay would win only if the true state is Blue. We refer to this alternative to step 4 as "Merge probabilities of states not initially nominated," or Merge for short. It tells subjects quite directly that the step 5 decision Stay will win the prize if i was the true state all along, but otherwise Switch will win.

Merge frees up the ordering of actions within the timeline: resolution of the true state in step 1 can now take place as late as between steps 5 and 6, and the conductor no longer has to know the true state before the subject does. Exploiting that freedom, Table 1 below defines ten treatments, which systematically switch on (indicated by a ✓) or off (−) key features of the original Monty Hall problem.

Treatment T7 in Table 1 is the original Monty Hall task, preserving the timeline above. Adjacent treatments in the Table differ in only one feature; for example, T6 uses Merge but otherwise is as in the original task. Not all feature combinations are implementable; for example, without Merge, the conductor must know the true state before offering the standard switch opportunity, and thus that state must be determined before the switch opportunity is offered. Hence, the last three rows in T7–T10 are necessarily identical, while the first two features are varied independently. Merge is used in T1–T6, enabling many combinations of the features in the first four rows of Table 1.

This set of treatments isolates the effect of each listed feature of the Monty Hall problem. For instance, comparing T4 versus T5 switch rates isolates the effect of determining the true state before the subject nominates an initial state. In like manner, competing explanations of behavior in the Monty Hall problem can be mapped to testable implications on how switch rates will differ across treatments. A manifest of our hypotheses is as follows.

A. *Illusion of Control and Escalation of Commitment*

As noted in the introduction, psychologists offer two different explanations of why subjects might feel attached to their initial nomination: either because they feel that they have

TABLE 1
Mapping of Features to Treatments

	Treatments									
Features	**T1**	**T2**	**T3**	**T4**	**T5**	**T6**	**T7**	**T8**	**T9**	**T10**
Subject makes initial nomination (vs. Random)	✓	-	-	-	-	✓	✓	-	-	✓
True state determined before initial nomination	-	-	-	-	✓	✓	✓	✓	-	-
Conductor knows true state before switch opportunity	-	-	-	✓	✓	✓	✓	✓	✓	✓
True state determined before switch opportunity	-	-	✓	✓	✓	✓	✓	✓	✓	✓
Merge probabilities of states not initially nominated	✓	✓	✓	✓	✓	✓	-	-	-	-

some ability to guess (or influence) the true state, or because they feel a sense of ownership and are reluctant to abandon their initial nomination. In either case (we cannot distinguish between them), attachment would be weaker when the initial nomination is random. Hence, we obtain the hypothesis,

H1: The switch rate is higher with random assignment than when subjects make the initial nomination.

Our data support this hypothesis if we can reject the null hypothesis of no difference in switch rates in the direction H1 specifies (e.g., there is a significantly higher switch rate associated with random assignment of the initial nomination). The next several hypotheses will be tested in a similar manner.

B. Asymmetric Information

Often it is rational to refuse to trade with a better informed individual, and subjects might inappropriately apply that rule of thumb when the conductor offers the option to Switch (Friedman 1998). Hence, we obtain the hypothesis,

H2: The switch rate is higher when the conductor does not know the true state before offering the opportunity to switch.

C. Erroneous Updating

As noted earlier, it may be that subjects incorrectly assess .50:.50 posterior probabilities over the remaining states, rather than 2/3:1/3 win probabilities for Switch versus Stay, and that our Merge treatment will help overcome that updating error. Thus, we obtain the hypothesis,

H3: The switch rate is higher with the Merge feature in place than with the information delivery used in the original Monty Hall procedure.

D. Confirmation Bias

It may be that subjects mistakenly regard the disclosure that state k is not the true state as confirming the initial nomination i as the correct identification of an already resolved true state. That effect might weaken when the true state is not determined prior to the initial nomination.

H4: The switch rate is higher when the true state is not determined before the initial nomination, in those treatments where the conductor discloses a nonwinning, non-nominated state.

E. Timing the Resolution of Uncertainty

The confirmation bias hypothesis above is not the only possible way that the time that uncertainty is resolved could affect the switch rate. Subjects might regard even an *undisclosed* resolution of the true state as an updating opportunity, or at least a cue that they *should* update. If the roll resolving the true state takes place after the initial nomination, the subjects may consider switching their nomination, realizing that something has changed. By contrast, if the resolution takes place before the initial nomination, this may fool the subjects into thinking—despite the mass reapportionment that takes place after the initial nomination, by different means, in *all* treatments—that nothing changes between their initial nomination and the switch opportunity, and thus that no updating is called for. Lumping together these (and perhaps other) conjectures, we obtain our next hypothesis.

H5: The switch rate is higher when the true state is not determined before the initial nomination.

Of course, resolving the true state *after* the final switch opportunity should completely eliminate these sorts of biases, hence, we obtain the hypothesis,

H6: The switch rate is higher when the true state is not determined until after the subject's final decision to switch or stay.

F. Probability Matching

Probability matching is a benchmark choice model from psychology (Estes 1954) in which

subjects' choice rates match the objective probabilities. With correct (or erroneous) assessments of the win probabilities, we have the two-part hypothesis:

H7: Up to sampling error, the switch rate is 66.7% (consistent with correct subjective probabilities but suboptimal choice) or 50% (consistent with incorrect subjective probabilities and suboptimal choice).

G. Experiential and Transfer Learning

Do subjects learn from experience within a given treatment? Are they able to transfer lessons learned in one treatment to another?

We address the first question by constructing and including in our econometric analysis the explanatory variable Switchbonus, introduced in Friedman (1998). This variable—defined as the cumulative, across rounds, tally of "Switch payoff minus Stay payoff, in a round"—captures the subject's exposure to accumulating evidence that Switch offers a better chance than Stay of winning the prize. Depending on the sequence of observed dice rolls, that evidence can be strong, weak, or even misleading.

H8: The switch rate is increasing in Switchbonus.

Transfer learning might be captured by sequences in which Original Monty (T7) is used in Segment 2 following a Merge treatment (T1–T6) in Segment 1, and by sequences in which T7 is used in Segment 3 following a Merge treatment in Segment 2. The hypothesis is

H9: Switch rates are higher in Original Monty (T7) segments that follow a segment with treatment (T1–T6) using the Merge feature.

H. Asymptotic Rationality

Do subjects eventually get it right? Our last hypothesis is

H10: Switch rates are asymptotically 100%.

We implement our design using a physical randomization device. The true state is determined by rolling (under an opaque cup) a six-sided die with face pairs colored Blue, Red, and Yellow, and, in treatments with the first feature set to Random (viz., T2–T5, T8, and T9), the initial nomination is similarly determined using another similarly colored die. All rolls are made in the presence of the subject, and disclosed to the subject by the end of the round in all treatments; thus, feedback on all realizations, payoffs, and opportunity costs is made each round in a credible and salient manner. Subjects keep track of the feedback events using a record sheet, as in Friedman (1998).

Each subject in our experiment completes three segments of ten rounds, for a total of 30 rounds. The first and last segment always use the same treatment, while the middle segment uses a different treatment, in ABA fashion (Treatment A—Treatment B—Treatment A, a standard treatment sequencing when there is a possibility of learning or other nonstationarity in behavior). The bottom part of Table 1 shows the number of observations of each treatment in each segment collected from our 77 subjects recruited from the Learning and Experimental Economcs Projects lab subject pool (University of California at Santa Cruz undergraduates from a wide set of majors). See Appendix S2, Supporting Information, for the full set of treatment sequences. Appendix S1 is a complete copy of instructions to subjects.

Subjects earned one point in every round in which their final choice was the true state. On average, they earned 17.08 points out of the maximum 30. (Optimal play, always Switch, would earn 20 points on average, while always Stay would average 10 points.) At the end of the experiment, subjects were paid $0.70 per point plus a show-up payment, typically $7.

III. RESULTS

Table 2 collects summary statistics, of interest in their own right as well as providing preliminary evidence regarding hypotheses.

While the lower portion of Table 2 permits direct inspection of the data by treatment and by segment, our formal statistical results are presented in Table 3. Table 3 reports results from a specification regressing switching (yes (1) or no (0) as a limited dependent variable) on: binary treatment variables, each indicating the presence or absence of each of the five design features (as listed in Table 1 or the upper portion of Table 2); a time trend; interactions between the five design feature binary variables and the time trend; either a constant (when the model is estimated as logit) or subject fixed effects (when the model is estimated as a linear probability model [LPM]); variables indicating whether implementation of the original Monty Hall problem followed implementation of other treatments using the Merge feature ("Segment 2 is Original

TABLE 2
Treatments and Mean Switch Rates

Features	Treatments									
	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10
Subject makes initial nom.	✓	-	-	-	-	✓	✓	-	-	✓
True state determined first	-	-	-	-	✓	✓	✓	✓	-	-
Conductor knows true state	-	-	-	✓	✓	✓	✓	✓	✓	✓
True state before switch opp.	-	-	✓	✓	✓	✓	✓	✓	✓	✓
Merge probabilities	✓	✓	✓	✓	✓	✓	-	-	-	-
	Mean Switch Rate (%)									
N = Number of observations										
Segment 1 $N = 770$	84.0 $N = 50$	70.0 $N = 70$	75.5 $N = 110$	82.1 $N = 140$	78.6 $N = 70$	90.0 $N = 40$	38.6 $N = 140$	26.0 $N = 50$	56.0 $N = 50$	44.0 $N = 50$
Segment 2 $N = 760$	90.0 $N = 40$	91.7 $N = 60$	82.7 $N = 110$	86.0 $N = 150$	86.2 $N = 80$	95.0 $N = 40$	58.5 $N = 130$	52.0 $N = 50$	56.0 $N = 50$	66.0 $N = 50$
Segment 3 $N = 770$	96.0 $N = 50$	80.0 $N = 70$	74.5 $N = 110$	91.4 $N = 140$	80.0 $N = 70$	90.0 $N = 40$	37.9 $N = 140$	36.0 $N = 50$	70.0 $N = 50$	58.0 $N = 50$
Overall $N = 2{,}300$	90.0 $N = 140$	80.0 $N = 200$	77.6 $N = 330$	86.5 $N = 430$	81.8 $N = 220$	91.7 $N = 120$	44.6 $N = 410$	38.0 $N = 150$	60.7 $N = 150$	56.0 $N = 150$

Monty" and "Segment 3 is Original Monty" binary variables allow capture of transfer of learning); and variables measuring feedback information useful in reinforcement learning (Switchwon and Switchbonus).

A. Illusion of Control and Escalation of Commitment

H1: The switch rate is higher with Random assignment than when subjects make the initial nomination.

H1 predicts positive switch rate differences T2−T1 and T5−T6 (with Merge present), as well as positive differences T8−T7 and T9−T10 (with Merge absent). The switch rates reported in the bottom of Table 1 are consistent with this prediction in Segment 2 for T2−T1, and for Segments 1 and 3 for T9−T10, but otherwise the differences go the wrong way.

In Table 3, the effect of the feature "Subject makes initial nomination" is not significant in either a stable or a time varying way, as evidenced by its insignificant (intercept or level) dummy variable and the insignificant interaction term between the preceding and time, respectively; this is true using either estimator (LPM or logit). We conclude that, despite their a priori appeal, we find no evidence that Illusion of Control or Escalation of Commitment provoke irrational choices in our Monty Hall data.

B. Asymmetric Information

H2: The switch rate is higher when the conductor does not know the true state before offering the opportunity to switch.

H2 implies that switch rate difference T3−T4 should be positive in Table 2. This prediction fails: the rates are essentially the same for the middle segment, but in the other segments (and overall) the switch rate is noticeably *higher* in T4. The regression estimates in Table 3 for level and slope for the relevant feature dummy are insignificant, with the possible exception ($p = .09$) of a positive slope in the logit regression. Again, a plausible explanation for irrationality finds no substantial support in our data.

C. Erroneous Updating

H3: The switch rate is higher with the Merge feature than with the original information procedure.

TABLE 3

Dependent Variable: SW(=1 if Switch Chosen)	Linear Probability Model	Logit Model	
Explanatory Variables:	**Coefficient (Standard Error)**		**Marginal Effect**
Constant	—	−1.0923**	—
		(0.4893)	
Time	0.0033	0.0096	0.0021
	(0.0056)	(0.0339)	
Subject makes initial nomination	0.0537	0.4371	0.0960
	(0.0506)	(0.2714)	
True state determined before initial nomination	0.0178	−0.0389	−0.0085
	(0.0526)	(0.2699)	
Conductor knows true state before switch	−0.0993	0.2479	0.0545
	(0.0676)	(0.3833)	
True state determined before switch opportunity	0.1394*	0.3156	0.0693
	(0.0814)	(0.4268)	
Merge probabilities	0.4403***	2.0208***	0.4439
	(0.0587)	(0.3293)	
Segment 2 is Original Monty	−0.0462	−0.6142	−0.1349
	(0.1323)	(0.7064)	
Segment 3 is Original Monty	−0.0692	−0.2714	−0.0596
	(0.2433)	(1.3319)	
Time × Subjects make initial nomination	−0.0010	−0.0008	−0.0002
	(0.0024)	(0.0154)	
Time × State determined before initial nomination	−0.0050**	−0.0307**	−0.0068
	(0.0024)	(0.0153)	
Time × Conductor knows state before subject	0.0042	0.0374*	0.0082
	(0.0032)	(0.0221)	
Time × State determined before switch opportunity	−0.0056	−0.04417*	−0.0097
	(0.0036)	(0.0252)	
Time × Merge probabilities	0.0009	0.0197	0.0043
	(0.0045)	(0.0254)	
Time × Segment 2 is Original Monty	0.0196**	0.0826*	0.0182
	(0.0087)	(0.0478)	
Time × Segment 3 is Original Monty	0.0052	0.0209	0.0046
	(0.0100)	(0.0543)	
Switchwon[a]	−0.0245	−0.1142	−0.0251
	(0.0177)	(0.1125)	
Switchbonus[b]	0.0069	0.0589***	0.0130
	(0.0044)	(0.0165)	
Adjusted R^2	0.12714	—	
N obs[c]	2,223	2,223	
F-statistic	19.1705	—	
Mean fixed effect	0.3022	—	

[a]Switchwon = 1 if Switch won in previous period.

[b]Switchbonus = cumulative earnings for Switch minus Stay.

[c]With 30 rounds and 77 subjects, we have 77 × (30−1)− 10 = 2,223 observations; the −10 reflects observations lost due to a recording error (for one subject in Segment 2), and the 77 × (−1) is due to a lagged variable.

*$p < .10$; **$p < .05$; ***$p < .01$.

The evidence for this hypothesis is very strong. Switch rates in all segments of all treatments with the Merge feature are at least as high as any switch rate in any segment without Merge. The direct comparison T6−T7 is an impressive 91.7−44.6 = 47.1%. Table 3 corroborates, showing a huge level effect in the LPM (or marginal effect in the logit model) of 44 percentage points, with miniscule p values. Slope effects are not significant. We conclude that erroneous probability updating accounts for the lion's share of observed irrationality, and that Merge enables subjects to understand right away (with no gradual adjustment) the advantages of Switch.

D. *Confirmation Bias*

H4: The switch rate is higher when the true state is not determined before the initial nomination, in

those treatments where the conductor discloses a nonwinning, non-nominated state.

T7 and T10 differ only in whether or not resolution of the true state takes place before or after the initial nomination; the same is true of T8 and T9. Thus one possible test of H4 takes each switch decision as a unit of observation, and compares switch rates across treatments. The chi-squared test rejects the null hypothesis of equal switch rates for T7 and T10 ($p = .02$), and for T8 and T9 ($p = .0002$), in the direction predicted by H4. A conservative approach is to take individual subject switch rate in each segment as the unit of observation, and test for differences across treatments using a rank-sum test. Pooling T7 and T10 on one hand, and T8 and T9 on the other, one fails to reject the null of no difference in central tendency ($p = .12$).

E. Timing the Resolution of Uncertainty

H5: The switch rate is higher when the true state is not determined before the initial nomination (over all treatments).

As reported in Table 2, the difference in mean switch rates for T10−T7, T9−T8, and T4−T5 is positive for all three comparisons, consistent with the prediction. More formally, we see that while the level coefficient for "true state determined before initial nomination" is insignificant in both LPM and logit, the slope (time interaction) coefficient is significant ($p < .05$) in both regressions. Thus, there is some evidence that earlier resolution of the true state does lower switching, via suppression of adjustment toward the optimal strategy.

H6: The switch rate is higher when the true state is not determined until after the subject's final decision to switch or stay.

Table 2 shows that the relevant treatments are T2 and T3, and that the prediction is simply that the difference T2−T3 in switch rates is positive. This prediction is correct in two out of three segments. (It would be desirable to make a similar comparison with the Merge feature turned off, but that is not logically possible.) In Table 3, however, none of the level or slope coefficients for "true state determined before switch" are significant at usual levels. There does not appear to be a clear role for behaviors relating to resolution of the true state beyond that already captured by relative positioning of initial nomination and state resolution (as in H5). Put another way, rendering the problem as a hypothetical statistics problem—with no resolving roll until after all reasoning and responses are completed—does not significantly boost or suppress the switch rate.

F. Probability Matching

H7: Up to sampling error, the switch rate is 66.7% (consistent with correct subjective probabilities but suboptimal choice) or 50% (consistent with incorrect subjective probabilities and suboptimal choice).

Applying an exact binomial test to each subject's data, we are able for each subject to make an assessment of whether or not their behavior over the 30 rounds could be consistent with probability matching either at 66.7% or 50%. For 28 of 77 subjects, we can reject probability matching at either benchmark. These rejections are due to switch rates well in excess of 66.7%. The picture is less distinct for other subjects—including 28 for whom we cannot reject either matching at 50% or matching at 66.7%. However, subject choice in this range could be generated during dynamic adjustment toward higher switch rates (see below). We conclude that probability matching does little to explain our data.

G. Experiential and Transfer Learning

H8: The switch rate is increasing in Switchbonus.

Table 2 has some favorable preliminary evidence. With the exception of T7, the Segment 3 switch rates are always at least as high as in Segment 1 for every treatment. Of course, this may be partly due to transfer learning (see below), so we look to Table 3 for sharper evidence. Both Switchbonus coefficient estimates have the predicted (positive) sign; the estimate is not significant in the LPM ($p = .11$), but it is highly significant ($p < .001$) in the logit model.

In passing, we note that Switchwon, also introduced by Friedman (1998), is the indicator for whether switching would have won the prize in the immediately preceding period. The coefficient estimates reported in Table 3 are not significant (and negative in sign) in contrast to the significant, positively signed estimates in Friedman (1998).

H9: Switch rates are higher in T7 segments that follow a segment using the Merge feature.

Table 2 contains evidence that supports that prediction. T7 switch rates are below 40% in Segments 1 and 3, but are nearly 60% in Segment 2, which is always sandwiched by Merge treatments (see line 7 in the online Appendix Table S1). Again, Table 3 yields sharper tests; the relevance of variable "Segment 2 is Original Monty" was just noted, and "Segment 3 is Original Monty" is also very relevant because (see column 7 of Table S1) T7 only sandwiches Merge treatments. Both variables have insignificant level effects, but positive slope estimates. The Segment 2 slope estimate is significant at $p < .0256$ in the LPM and is borderline significant ($p = .084$) in the logit model. Thus, between H8 and H9, we have evidence consistent with experiential learning and transfer learning, respectively.

H. Asymptotic Rationality

H10: Switch rates are asymptotically 100% for a typical subject.

The evidence from Table 2 (higher switch rates in Segment 3 than in Segment 1 for almost all treatments) is encouraging but inconclusive. Table 3 considers linear time trends (or linear effects in Switchbonus), which are not especially well suited for the question at hand. As an alternative, we adapt a technique introduced in Noussair, Plott, and Riezman (1995) and estimate the coefficients B_{Init} and B_{Asymp} in the equation

$$\text{Absolute Deviation from Switch}_{it} = B_{Init,i}(1/t) + B_{Asymp,i}\,((t-1)/t) + \epsilon_{it},$$

using data from the final segment for each individual subject i. Here, $t = 1, 2, \ldots, 10$ is the period (round) number in that segment. The coefficients B_{Init} and $B_{Asymp,}$ respectively, represent the initial value and constant asymptote of the hyperbola that best fits the ten observations. Of course, there is considerable sampling error and heterogeneity across individual subjects, but the medians by treatment are of interest. For treatments T1 through T6 (all treatments with the Merge feature), the median B_{Asymp} estimate is 0, implying a 100% switch rate, as H10 predicts. For treatments T7–T10 (all treatments without the Merge feature), the median estimated B_{Asymp} is 0.1703, implying an 83% switch rate. Among subjects in T7 (the original Monty Hall), the median B_{Asymp} estimate is 0.4988, implying a 50% switch rate. Under the model, subjects with a positive asymptote, such as the median subjects in the non-Merge treatments, would never attain 100% switching.

IV. APPLICATION TO FIELD SITUATIONS

Abstract statistical reasoning can be difficult in the calmest of circumstances, and can be impractical when it is most needed, in fast-moving, stressful environments such as armed combat (a new blip appears on the screen—is it friend or foe?) or medical diagnosis. For such applications, simplification (e.g., Gigerenzer and Todd 1999) and automated Bayesian tools (e.g., Cox et al. 2016) may help prevent fallible humans from making serious errors.

To see the connection to our experiment, consider the following special but not unrealistic medical diagnosis situation. There are three mutually exclusive possible causes of a patient's aliment, and a priori all three are equally likely. The insurance company (or perhaps the current state of knowledge) will permit tests for only two of the possible causes. A negative result for either test would definitively and quickly reject that associated cause. Thus in the class of scenarios we consider here, the first test received back will necessarily be a negative one. Both possible tests are performed.

Next, the first test result comes in—it is, of course, negative for one of the two testable causes. What now is the updated probability that the true cause is the one for which a test is not permitted? We confidently conjecture that most people (and likely most medical professionals) would say 50%. The correct answer, of course, is that the prior probability of 1/3 remains valid, because that cause was "sequestered" in exactly the same way as the initial choice in the original Monty Hall problem. Knowing the correct answer—or having it enforced by means of an expert system that not only requests and stores information on testing of individual states but also automatically reallocates probability mass across states in real time in the manner tested by our Merge feature—could help prevent bad decisions when immediate action is required. Note: Eventually the second test result will arrive, and in general it will change the Bayesian posterior probability of the sequestered cause. The point is that action may be required in the meantime.

The example generalizes. Suppose there are $n > 3$ possible causes, $k < n$ of them testable, with negative results received before positive results. Contrary to most people's intuition, the prior on a nontestable cause would be unaffected after

receipt of $j < k$ negative results. Our experiment thus underlines the need for automated Bayesian tools in medical diagnosis (as in Cox et al. 2016) and in a variety of other field situations ranging from angel investment (in startup companies) to warfare. Expert systems can easily be augmented to allow a branching that establishes whether a particular state is (not) sequestered from updating; this step shapes the kind of updating that takes place across the states not sequestered, but also not contemporaneously revealed.

V. DISCUSSION

Our laboratory dissection of the Monty Hall puzzle suggests that the root problem is a failure in Bayesian updating—specifically, a failure to merge correctly the probability masses of the states not initially nominated. Most subjects not exposed to the transparent merging-of-probability-mass design feature (Merge) do not learn to switch consistently by the end of the experiment. On the other hand, most subjects do switch consistently (i.e., behave rationally) while exposed to Merge. Switching is also elevated to some degree among subjects in Original Monty Hall who have previously experienced a treatment employing Merge.

Other conjectured explanations for failure to switch have much smaller impact in the data. Of those explanations, only confirmation bias (or some other forms of sensitivity to when uncertainty is resolved) might receive any support, with some indication of operation via suppression of learning. Our data overall suggest that experiential learning and transfer learning do increase switch rates, though the effect sizes are dwarfed by those of Merge.

Taken together with the cross-species results of Herbranson and Schroeder (2010), one might conclude from our results that, compared to pigeons, humans are relatively bad at learning from experience. A more generous interpretation is that human choice is more responsive to insight. We are more likely to persist in suboptimal behavior driven by an incorrect belief (misconstrued win probabilities for switching), but we are faster to respond to new insight.

There are important practical implications of our results. It is an unfortunate reality in medicine that failure to correctly exploit Bayesian updating opportunities on the part of medical professionals can result in sickness or death on the part of patients (Cox et al. 2016). Anything that can be done to ease the correct implementation of Bayesian reasoning—such as the engineering of information displays and information provision so as to promote correct Bayesian reasoning (Gigerenzer and Hoffrage 1995), or providing econometric decision aids (Cox et al. 2016)—is potentially of the greatest value. Our findings, particularly as related to the Merge design feature, suggest a specific improvement that (medical diagnostic and other) decision aids might incorporate. That is the identification and sequestration of states that cannot be the object of updating. Explicitly introducing and accounting for this step early on in updating procedures, which facilitates the calculation of conditional probabilities, would constitute a simple and practical transfer of our findings to real-world decision-making.

REFERENCES

Arkes, H. R., and C. Blumer. "The Psychology of Sunk Cost." *Organizational Behavior and Human Decision Processes*, 35(1), 1985, 124–40.

Bazerman, M. H. *Judgment in Managerial Decision Making*. New York: Wiley, 1990.

Camerer, C. "Individual Decision Making," in *The Handbook of Experimental Economics*, edited by J. H. Kagel and A. E. Roth. Princeton, NJ: Princeton University Press, 1995, 587–703.

Camerer, C. F., and R. A. Weber. "The Econometrics and Behavioral Economics of Escalation of Commitment: A Re-Examination of Staw and Hoang's NBA Data." *Journal of Economic Behavior & Organization*, 39(1), 1999, 59–82.

Cox, J. C., V. Sadiraj, K. Schnier, and J. F. Sweeney. "Higher Quality and Lower Cost from Improving Hospital Discharge Decision Making." *Journal of Economic Behavior & Organization*, 131(B), 2016, 1–16.

Esponda, I., and E. Vespa. "Hypothetical Thinking and Information Extraction in the Laboratory." *American Economic Journal: Microeconomics*, 6(4), 2014, 180–202.

Estes, W. K. "Individual Behavior in Uncertain Situations: An Interpretation in Terms of Statistical Association Theory," in *Decision Processes*, edited by R. M. Thrall, C. H. Coombs, and R. L. Davis. New York: Wiley, 1954, 127–37.

Eyster, E., and M. Rabin. "Cursed Equilibrium." *Econometrica*, 73(5), 2005, 1623–72.

Friedman, D. "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly." *American Economic Review*, 88(4), 1998, 933–46.

Gigerenzer, G., and U. Hoffrage. "How to Improve Bayesian Reasoning without Instruction: Frequency Formats." *Psychological Review*, 102(4), 1995, 684–704.

Gigerenzer, G., and P. M. Todd. "Fast and Frugal Heuristics: The Adaptive Toolbox," in *Simple Heuristics that Make Us Smart*. New York: Oxford University Press, 1999, 3–34.

Granberg, D. "Cross-Cultural Comparison of Responses to the Monty Hall Dilemma." *Social Behavior and Personality*, 27(4), 1999, 431–38.

Heath, C. "Escalation and De-escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting." *Organizational Behavior and Human Decision Processes*, 62(1), 1995, 38–54.

Herbranson, W., and J. Schroeder. "Are Birds Smarter Than Humans? Pigeons (*Columba livia)* Perform Optimally on a Version of the Monty Hall Dilemma." *Journal of Comparative Psychology*, 124(1), 2010, 1–13.

Krauss, S., and X. Wang. "The Psychology of the Monty Hall Problem: Discovering Psychological Mechanisms for Solving a Tenacious Brain Teaser." *Journal of Experimental Psychology: General*, 132(1), 2003, 3–22.

Miller, J. B., and A. Sanjurjo. "A Bridge from Monty Hall to the (Anti-)Hot Hand: Restricted Choice, Selection Bias, and Empirical Practice." SSRN Working Paper No. 2709837, 2015.

Nalebuff, B. "Puzzles: Choose a Curtain, Duel-ity, Two Point Conversions, and More." *Journal of Economic Perspectives*, 1(2), 1987, 157–63.

Noussair, C. N., C. R. Plott, and R. G. Riezman. "An Experimental Investigation of the Patterns of International Trade." *American Economic Review*, 85(3), 1995, 462–91.

Palacios-Huerta, I. "Learning to Open Monty Hall's Doors." *Experimental Economics*, 6(3), 2003, 235–51.

Selvin, S. "A Problem in Probability (Letter to the Editor)." *The American Statistician*, 29, 1975, 67.

vos Savant, M. "Ask Marilyn." *Parade*, September 8, 1990; December 2, 1990; February 17, 1991.

Wason, P. C. "On the Failure to Eliminate Hypotheses in a Conceptual Task." *Quarterly Journal of Experimental Psychology*, 12(3), 1960, 129–40.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix S1. Instructions to subjects

Appendix S2. Supplementary analysis