

Dynamics of Strategy Adaptation in a Temporally Extended Monty Hall Dilemma

Stephanie Petrusz (stephanie.petrusz@gmail.com) & Joshua Shields (joshua.n.shields@gmail.com)

Department of Psychology, Franklin & Marshall College
Lancaster, PA, 17603 USA

Theo Rhodes (theo.rhodes@oswego.edu) & Abisha Munroe (amunroe@oswego.edu)

Department of Psychology, SUNY Oswego
Oswego, NY, 13126 USA

Abstract

We present the results of two temporally extended experimental implementations of the Monty Hall dilemma in order to examine the dynamics of belief. In the first experiment, we used the standard three-door version of the dilemma, but biased the probability of the winning door positionally. Participants capitalized on the increased probabilities but did not discover the optimal switch strategy. In the second experiment, we increased the number of doors, in each case removing all but two doors. As the number of doors increased, participants converged on the optimal switch strategy, as well as increasing their confidence in their strategy. This suggests that the information relevant to the MHD is not win frequencies but how the different elements of the dilemma are related.

Keywords: Belief; Monty Hall Dilemma; Dynamical Systems

Introduction

The conditions under which an agent may be said to believe can be cashed out in different ways. Belief is essentially an agent's holding that a proposition or state of affairs is the case, and as such is difficult to externally assess. Two related but distinct ways suggest themselves: an agent may assent to the truth or falsity of a proposition, or an agent may be disposed to act as though a proposition is true—or, if one wishes to avoid speaking as though propositions are the contents of beliefs, to act as though a state of affairs obtains. If we focus on the second type of belief, that is, beliefs as underwriters of actions, two more things make themselves clear. First, the ability of beliefs to guide actions must be relatively robust to uncertainty. Second, beliefs must change in response to new information or evidence, so that actions may be adapted to a dynamic environment.

Formal models of belief therefore often treat beliefs as probabilities, and belief change or belief update as changes in those probabilities and in the connections between the beliefs they represent. Many of these models share certain features: the beliefs in question are framed as propositions and the dynamics are understood as state-to-state changes. Many models further share a Bayesian basis. However, a thread of argument in the epistemological literature claims that a basic species of Bayesian revision is probably inadequate to model belief change (see van Fraassen (1989), cited in Arló-Costa, unpublished). Evaluating the conditionals required by Bayesian update is problematic for

reasons including the complexity of agents' existing commitments (see, for instance, the "web of belief" of Quine and Ullian, 1978) and the difficulty of evaluating the probability of some fact against an unenumerated backdrop of possibilities.

In our view the mathematical considerations particular to formal epistemological accounts of belief update are also appropriate to psychological accounts. For our empirical investigation we required sufficient data to construct a model and a task where we could plausibly discuss what the experimental participants believed about their actions. Because of its simplicity of execution and its popularity, we chose the Monty Hall Dilemma (e.g., vos Savant, 1990). This paradigm satisfied several aims. First, we wanted a task in which outcome probabilities were clear, though not necessarily obvious to the participants. Second, we wanted a task that involved beliefs as the basis for actions, not simply evaluating propositions with degrees of assent. Third, we wanted to offer participants the opportunity to vacillate while still making a desired response—in effect to say, "I'm responding this way but I don't know if I should", as a measure of degree of belief. We hoped that by evaluating participants' approaches to the MHD we would be able to address how people evaluate probabilities and choose between competing strategies for dealing with situations in which probabilities change. We hoped further to investigate whether a separation existed between participants' beliefs as the basis for actions and their propositional beliefs about the task.

The MHD is a good example of the difference between incorporating probabilistic evidence and reasoning probabilistically because it is an example of persistent failure to act in accordance with the rules of probability. In the classic version, a "contestant" is presented with three doors. Behind one of the doors is a car, and behind the other two are goats. The contestant chooses a door, and then one of the two remaining unchosen doors is opened—always revealing a goat. The contestant is then given the chance to stick with their initially chosen door and receive whatever lies behind it, or to switch to the remaining unchosen and unopened door, receiving instead whatever prize lies behind that door. Probabilistically, a contestant always maximizes their chances of winning the car by switching their pick, yet in practice people seldom choose to switch. In fact, many people resist the logic of switching even after it is

demonstrated or explained to them that switching is always the better bet.

The classic formulation can be extended in a variety of ways. Granberg (1999) tested a 4-door version of MHD with unequal probabilities, where the optimal strategy is to choose the lowest probability door initially and then switch to the highest probability door remaining after the reveal stage. Participants in Granberg's study were explicitly informed of the door probabilities on each trial, yet even so no participants reliably used the optimal strategy. Baratgin and Politzer (2010) showed that both the dilemma and the resistance to the winning strategy are rooted in the manner in which information is presented throughout the scenario. They presented several versions of MHD and concluded that when information is presented as being relevant only to the unchosen doors, people do not commit the typical MHD error. They argue that the opening of the "goat" door is taken to be relevant to the overall probabilities, and that players do not effectively isolate the new information to the real space of outcomes, which includes only the unchosen doors.

The MHD has been extensively studied in cognitive science but a full review is beyond the scope of this paper. Krauss and Wang (2003) discuss the psychological mechanisms that might underpin successful MHD strategies and suggest that correctly reasoning through the problem depends on the way in which information is represented throughout the scenario. Their treatment of the "natural frequencies" approach suggests that the solution is more apparent when the frequencies of win and loss events in the aggregate are considered, rather than the probability of winning or losing a single run. Saenen, Van Dooren, and Onghena (2015) used feedback to improve participants' performance on the MHD but noted that the increases in performance did not come with associated improvement in understanding of the dilemma. These results are of a piece with our considerations: the information needed to identify a correct strategy becomes available to be acted upon with practice, but does not become available to conscious knowledge in the same fashion.

We devised several variants of the MHD to explore whether manipulation of some of the dilemma's parameters would affect participants' exploration and adoption of different strategies. If this is the case then we may say participants are able to incorporate information that is probabilistic (that is, not fully reliable) as a basis for action. We also explored whether manipulations yielding more successful performance than in the classic case would translate to higher participant confidence (that is, explicitly held beliefs about the value of the new information or about the utility of an adopted strategy) in their attempts.

Experiment 1

In experiment 1, we use the standard 3-door MHD, biasing the probability of winner based on door position in order to test whether this bias would in fact be apparent to

participants, and whether it would enhance or diminish their discovery of the switching strategy.

Methods

Participants 20 undergraduate students (mean age = 19.9, 13 male, 7 female) from Franklin & Marshall College participated in the experiment for course credit.

Design Experiment 1 used a 3-door MHD with color (shades of red) serving as the property distinguishing the "prize" (brighter shade) from the "goats" (darker shade). We chose color as the target stimulus property because it is both salient to participants and easy to program.

The experiment was divided into 8 blocks: an initial 100-trial block in which each of the three doors had an equal chance of concealing the "prize", followed by 4 100-trial blocks with unequal probabilities of the "prize" being behind each of the doors, and finally 3 50-trial blocks with equal probabilities. In the unequal-probabilities blocks either the leftmost or rightmost door had a higher chance of concealing the "prize". Participants were randomly assigned to either the left-biased (Group L) or right-biased (Group R) condition. Participants were asked both to select a door by clicking on it, and also to register their confidence in their choice, measured in how far from the center of the chosen square *in any direction* they clicked, while still making the response within the confines of the door. This lack of directionality was included specifically to avoid laterality effects, and also to simplify the verbal instructions to the participants. The center represented the highest possible confidence and the inside edge of the door represented the lowest confidence. In all there were 650 trials and 1300 distinct responses for each participant.

Procedure Each individual trial consisted of an instance of MHD, as follows: three "doors" (square boxes) were displayed on screen, along with the text prompt "Select a door". A participant would click on one box, whereupon the color behind a remaining non-prize door was revealed. The participant would then be asked to stick or switch, via the text prompt: "Switch your pick?" They would then either switch (by clicking on the unopened door they did not originally choose) or stick (by once again clicking on the originally-chosen door). The color behind the door they ultimately chose would then be revealed, along with text telling them either "You win" or "You lose" depending on the outcome of their final choice. At the end of the experiment participants were asked to briefly describe any strategy they might have employed in the course of the experiment.

Results

Because of the large number of data points, it is especially important to consider results in terms of both significance and effect size. We report only those results that meet both criteria.

A 2 (Group) \times 8 (Block) ANOVA was conducted on the measures of door chosen, whether the correct door was selected, confidence, and whether a trial was a “switch” or “stick”. Group significantly affected all four measures, but with a meaningful effect size only for door chosen, $F(1, 25984) = 1828.63, p < .001$, partial eta squared = .07. The “door chosen” measure was significantly higher for Group R ($M = 2.19, SD = .83$) than for Group L ($M = 1.67, SD = .77$), with doors coded as 1 (left), 2 (middle), and 3 (right). Block significantly affected all four measures, but with meaningful effect sizes for confidence, $F(7, 25984) = 83.65, p < .001$, partial eta squared = .02 and correct door, $F(7, 25984) = 73.95, p < .001$, partial eta squared = .02. Because lower values of the confidence measure indicate higher confidence (as measured in pixel distance from the center of the chosen door), we report the negative of the values of confidence. Confidence is lowest in Block 1 ($M = -48.57, SD = -31.97$) before increasing steadily until its maximum in Block 5 ($M = -35.5, SD = -24.15$) and then decreasing again until Block 8 ($M = -47.37, SD = -28.17$). The proportion of trials on which participants end up choosing the winning door is lowest in Block 1 ($M = .57, SD = .5$) and steadily increases until Block 5 ($M = .85, SD = .36$) before decreasing sharply in Block 6 ($M = .6, SD = .49$) and remaining relatively steady until the final block ($M = .63, SD = .48$). There was a significant Group \times Block interaction for door chosen, $F(7, 25984) = 296.41, p < .001$, partial eta squared = .07.

Having characterized each trial as either a “switch” or “stick” trial, we conducted a second ANOVA using this strategy variable as an IV. Error degrees of freedom are smaller because the analysis is in terms of trials rather than responses. For the 2 (Group) \times 2 (Strategy) \times 8 (Block) ANOVA, Group significantly affected door chosen, $F(1, 12968) = 2710.61, p < .001$, partial eta squared = .17. Group R chose the right-most door ($M = 2.32, SD = .85$) significantly more than Group L ($M = 1.45, SD = .75$). Strategy significantly affected correct door, $F(1, 12968) = 979.99, p < .001$, partial eta squared = .07. Trials that ended with a switch response had a higher proportion of correctly chosen doors ($M = .82, SD = .39$) than trials that ended with a “stick” ($M = .58, SD = .49$). Strategy also significantly affected confidence, but with an effect size so small as to be considered negligible. Block significantly affected door chosen, $F(7, 12968) = 20.33, p < .001$, partial eta squared = .01, confidence, $F(7, 12969) = 42.08, p < .001$, partial eta squared = .02, and correct door, $F(7, 12968) = 144.6, p < .001$, partial eta squared = .07. There were significant two-way interactions of Group \times Block, Group \times Strategy, and Strategy \times Block. The three-way interaction of Group, Block, and Strategy was also significant.

Finally, we found that correct responses were only weakly correlated with confidence, $r = .31, p < .001$.

Discussion

The winning MHD strategy is switching, but as described above, a preponderance of players not only fail to discover

or employ this strategy, but resist acknowledging that it is better even after having it explained. This fact is often leveraged as part of an argument about irrationality or about the difficulty of reasoning probabilistically. As Granberg (1999) reported, in an unequal-probabilities MHD the best strategy is even less obvious: to optimize the chances of winning a player should initially select a low-probability door and then switch to the highest remaining probability door after a “goat” has been revealed.

In Granberg’s 4-door unequal-probabilities study participants were informed of the probabilities of finding the prize behind each door. In the present experiment participants were not informed of the probabilities at any point, yet were able over the course of the unequal-probabilities blocks to hone in on the location of the winning door (leftmost for Group L, rightmost for Group R). We take this to be evidence of participants’ ability to *perceptually* coordinate with the structure of the experimental task, and believe this ability is conditioned on the availability of very many trials as a basis for learning, and on the fact that the unequal win probabilities made the presence of a winner more salient than in the standard case. We also believe that the availability of a tactic more obvious than switching (which, as shown by the second ANOVA, remains the better strategy) accounts in part for the strength of the correlation between performance and confidence, a subject to which we return in the general discussion.

In Experiment 1 we attempted to approach the Monty Hall paradigm in a way more streamlined than the classic case, but in a way that also allowed us to look at the effects of making the probabilities involved more obvious, and of how repeated exposures to the problem affected participants’ approaches. However, further refinements were required.

Experiment 2

Experiment 2 was designed to not only manipulate the number of doors and the associated win probabilities, but also to further refine and streamline the experimental design.

Methods

Participants Twenty-eight undergraduate students (6 male, 22 female) from SUNY Oswego participated in the experiment for course credit.

Design In Experiment 2 we returned to an equal-probability design for the doors, but manipulated the *number* of doors in the Monty Hall scenario. Each block consisted of 30 MHD trials with between 3 and 11 doors, presented in ascending (from 3 to 11 doors), descending (from 11 to 3 doors) or random increments across 8 blocks. We also changed how we denoted the “prize” and “goat” doors, using a check symbol for a winning door and an x symbol for a losing door.

Procedure The procedure was the same as in Experiment 1, with the important exception that, after a participant made their initial choice, *all but one* of the remaining doors was opened, instead of a single door. This design was based on an experiment suggested in Spivey (2008), and was intended to approach the issue of making the winning *strategy* as salient as the winning door had been in Experiment 1.

Results

We conducted a 3 (Group) × 8 (Block) × 8 (Number of Doors) ANOVA on the dependent measures of correct responses, confidence (calculated for the second response in a trial), and the total number of “switch” trials per participant and block. As in Experiment 1, we report only those results that achieved both significance and meaningful effect size. Group significantly affected number of switches, $F(2, 7485) = 203.66, p < .001$, partial eta squared = .05. On average, the Descending group had the highest number of switches per block ($M = 27.21, SD = 2.74$), followed by the Ascending group ($M = 24.23, SD = 7$), with the random group having the lowest average number of switches ($M = 22.59, SD = 6.8$). Block significantly affected total number of switches (see Figure 1), $F(8, 7485) = 24.15, p < .001$, partial eta squared = .025. Switching was on average highest in Block 6 ($M = 26.57, SD = 3.77$) and lowest in Block 1 ($M = 19.64, SD = 6.99$).




Figure 1: Mean switches made by participants by number of doors and condition.

Number of doors significantly affected correct responses (see Figure 2), $F(8, 7485) = 9.96, p < .001$, partial eta squared = .011 and total number of switches, $F(8, 7485) = 68.06, p < .001$, partial eta squared = .068. On average, the 10-door block had the highest proportion of correct responses ($M = 83.8\%, SD = 36.9\%$), while the 3-door block had the lowest ($M = 56.7\%, SD = 49.5\%$). On average, the

10-door block had the highest number of switches ($M = 27.14, SD = 3.96$) while the 3-door block had the lowest ($M = 18.54, SD = 7.37$).




Figure 2: Mean correct choices made by participants by number of doors and condition.

Having calculated the number of switches, we included this variable in a second analysis. We regressed Condition, Block, Number of Doors, and Total Switches on correct responses. These predictors explained 12% of the variance ($R^2 = .12, F(4, 7559) = 252.76, p < .001$). The number of doors significantly predicted correct responses, $\beta = .09, p < .001$, as did total number of switches, $\beta = .3, p < .001$. We then regressed Condition, Block, Number of Doors, and Total Switches on confidence (registered on the second response per trial). Confidence was normalized so that all values fell between 0 and 1. We did this to account for the fact that participants varied in their range of confidence values, and we were primarily interested in how confidence changed over the course of the experiment (see Figure 3).




Figure 3: Mean normalized confidence estimates for by number of doors and condition.

These predictors explained only 3% of the variance in the normalized confidence ($R^2 = .03$, $F(4, 7559) = 65.45$, $p < .001$). All four predictors significantly affected normalized confidence: Condition ($\beta = -.06$, $p < .001$), Block ($\beta = .05$, $p < .001$), number of doors ($\beta = -.1$, $p < .001$), and total number of switches ($\beta = -.11$, $p < .001$).

Finally, we found that correct responses were only weakly correlated with confidence, $r = .24$, $p < .005$.

Discussion

In Experiment 2 we varied the MHD by manipulating the number of doors in each block. The presence of additional doors alters the win probabilities, though not the fact that switching still has a higher probability of winning than sticking. In theory, the additional doors create a more difficult problem of probabilistic reasoning than the 3-door case, but the removal of all but one door changes those calculations. As Spivey (2008) points out, seeing the subset of possibilities reduced so drastically may actually make switching a more obvious strategy. Our results bear this hypothesis out.

The significant effect of Group on number of switch trials means that the order of presentation of door numbers makes a difference to how participants approach the task—in other words, the task exhibits hysteresis. Here, the Descending group had the highest number of switch trials: the presence and subsequent removal of more doors makes it easier to see that the original door has a lower probability of being a winner than the one remaining out of all the rest, which persists as the number of doors is reduced over the following blocks. This interpretation is also supported by the fact that the 10-door block had the highest numbers of both switches and correct responses, while the 3-door block had the lowest of both. Although a case could be made that

the 11-door block should have been even higher, the 11-door block was also the first block presented for a third of the participants, and therefore performance likely suffered because of the novelty of the task.

We attempted to better isolate the effect that switching has on performance and on confidence using linear regression, with mixed results. In particular, all of our independent variables together accounted for only 3% of the variance in participants' confidence. Taken together with the weak correlation between performance and confidence, it seems clear that we have not yet isolated a manipulation that primarily affects confidence, though this might also be a limitation of how confidence is measured.

General Discussion

The present set of experiments was designed to investigate belief from an empirical standpoint. To this end, a distinction was drawn between two different ways of thinking about belief. There is belief-as-proposition, which is perhaps the better-known way of considering belief, and there is belief as a guide to action. Showing that one *possesses* both species of belief but that they are not equivalent, then, takes a step toward an understanding of belief that is not beholden to its propositional character. Allowing participants to provide an overt measure of behavior (door choices) as well as a more covert measure of their own state (confidence) lets us see belief's different guises separately.

Our other aim was to consider why it is that the MHD presents such difficulty. We considered that participants might assume the initial symmetry between doors held throughout the scenario and attempted to introduce new information via the introduction of the positional bias in Experiment 1. We also considered that participants fail to see that the doors remaining after a reveal constitute a distinct subset of probabilities and attempted to make this fact more salient by manipulating the number of doors throughout Experiment 2.

The results of both Experiments 1 and 2 speak to the distinction between a belief on which action (e.g. choosing a door that might conceal a prize) is based and assent to the propositional description of that belief (e.g. "I believe that Door Number 3 is the winner"). Participants are able to accomplish the former regardless of their degree of belief in the latter. Performance and confidence hold together only weakly even when performance is successful. As acknowledged above, we do not feel that we have as yet hit upon a way of really manipulating confidence, but we do feel that we have demonstrated that successful action need not be conditional on any particular degree of belief. This finding is of a piece with the hypothesis that the aspects of the task about which a participant might have propositionally-construed beliefs are not necessarily the aspects that are used to guide behavior (Anderson & Runeson, 2008; Runeson, Juslin, & Olsson, 2000).

In Experiment 1 we hoped to shed light on the dynamics of beliefs in the face of changing evidence, by altering the

probabilities that a given door would be the winner. The results show that participants are able to perceive that one door is likelier to pay off and are able to adapt their strategy accordingly. This result, coupled with the understanding that the beliefs that guide actions are usually many more in number than those considered here, led us to believe that adaptation to changing information might be considered a dynamical system. Experiment 2 was designed to investigate this question more fully, by incorporating the changing number of doors. This provided not only a candidate control parameter for the MHD but also allowed us to test the hypothesis that failure to separate the probabilities associated with chosen doors versus unchosen doors (or, to put it another way, improperly conditioning beliefs about chosen doors on information about unchosen doors) is responsible for the difficulty of discovering and adopting the switching strategy.

The results of Experiment 2 support both contentions. The different conditions show signs of hysteresis, a hallmark of dynamical systems, and suggest further avenues of exploration in future manipulations within this paradigm. The results also support the hypothesis that information in one part of the dilemma is taken to affect other parts is a matter of framing. Additional manipulations regarding numbers of choices, distributions of probabilities, and physical versus virtual objects, are all natural extensions of the work presented here.

Our paradigm needs refinement, but works at a first past. Participants change their actions based on changing information, and these changes persist in some fashion. The actions are undertaken not wholly in keeping with degree of assent to the correctness of the action. Finally, there are large-scale dynamics with patterns of consistency that provide theoretical and methodological insights into investigating changes in belief.

References

- Andersson, I. E., & Runeson, S. (2008). Realism of confidence, modes of apprehension, and variable-use in visual discrimination of relative mass. *Ecological Psychology*, 20(1), 1-31.
- Arló-Costa, H. (unpublished). *Computational and probabilistic models of abduction*. Department of Philosophy, Carnegie-Mellon University.
- Baratgin, J., & Politzer, G. (2010). Updating: A psychologically basic situation of probability revision. *Thinking & Reasoning*, 16(4), 253-287.
- Granberg, D. (1999). A new version of the Monty Hall Dilemma with unequal probabilities. *Behavioural Processes*, 48(1), 25-34.
- Krauss, S., & Wang, X. T. (2003). The psychology of the Monty Hall problem: discovering psychological mechanisms for solving a tenacious brain teaser. *Journal of Experimental Psychology: General*, 132(1), 3-22.
- Quine, W. V. O., & Ullian, J. S. (1978). *The web of belief* (Vol. 2). R. M. Ohmann (Ed.). New York: Random House.
- Runeson, S., Juslin, P. and Olsson, H. (2000). Visual perception of dynamic properties: cue heuristics versus direct-perceptual competence. *Psychological Review*, 107: 525–555.
- Saenen, L., Van Dooren, W., & Onghena, P. (2015). A randomized Monty Hall experiment: The positive effect of conditional frequency feedback. *Thinking & Reasoning*, 21(2), 176-192.
- Spivey, M. (2008). *The continuity of mind*. Oxford University Press.
- Van Fraassen, B. (1989). *Laws and symmetry*. Oxford University Press.
- Vos Savant, M. (1990). Ask Marilyn. *Parade*, 15.