

A Bridge from Monty Hall to the Hot Hand: Restricted Choice, Selection Bias, and Empirical Practice

Joshua B. Miller and Adam Sanjurjo^{*‡}

This Version: October 27, 2017

Abstract

We show how classic conditional probability puzzles, such as the Monty Hall problem, are intimately related to the hot hand selection bias (Miller and Sanjurjo 2015). We explain the connection by way of the *principle of restricted choice*, an intuitive inferential rule from the card game bridge, which we show is naturally quantified as the updating factor in the odds form of Bayes Rule. We illustrate how, just as the experimental subject fails to use available information to update correctly when choosing a door in the Monty Hall problem, we as researchers can neglect analogous information when designing experiments, analyzing data, and interpreting results.

JEL Classification Numbers: C11; C13; C18; C52; C9; D83.

Keywords: Hot Hand; Hot Hand Fallacy; Monty Hall Problem; Selection Bias; Principle of Restricted Choice; Bayes Rule; Contract Bridge; Berkson's Paradox; Boy-or-Girl Paradox; Bertrand's Box Paradox; Alternation Bias; Gambler's Fallacy; Streak Reversal Paradox; Google Problem; ESP; Psi; Precognition.

^{*}Sanjurjo: Fundamentos del Análisis Económico, Universidad de Alicante. Financial support from the Spanish Ministries of Education and Science and Economics and Competitiveness (ECO2015-65820-P) and Generalitat Valenciana (Research Projects Gruposo3/086 and PROMETEO/2013/037) is gratefully acknowledged.

[†]Both authors contributed equally, with names listed in alphabetical order.

[‡]This draft has benefitted from conversations with Maya Bar-Hillel, Ben Cohen, Lola Collado, Vincent Crawford, Carlos Cueva, Martin Dufwenberg, Jordan Ellenberg, Maria Glymour, Íñigo Iturbe, Steve Kass, Mike Lipkin, Massimo Marinacci, Judea Pearl, Lones Smith, and Dan Stone, as well as seminar participants at U. de Alicante, Radboud U. Nijmegen, and IMEBESS 2016 Roma.

The data cannot speak for themselves; and they never have, in any real problem of inference. (Jaynes 2003)

1 Introduction

The Monty Hall problem is a simple probability puzzle that is well-known for its ability to confound the intuitions of everyone from the layperson to the mathematically sophisticated. The most famous version of the problem is as follows:

Suppose you're on a game show, and you're given the choice of three doors. Behind one door is a car, behind the others, goats. You pick a door, say #1, and the host, who knows what's behind the doors, opens another door, say #3, which has a goat. He says to you, "Do you want to pick door #2?" Is it to your advantage to switch your choice of doors?¹ (Vos Savant 1990)

While the intuitively appealing answer is that either choice leads to the same chances of winning the car, this is wrong (under natural assumptions), as the contestant's chances of winning the car actually increase if she switches doors.² People typically get this problem wrong. In particular, a robust finding in laboratory experiments is that roughly 80-90 percent of subjects incorrectly stay with the same door, rather than switch (see, e.g. Friedman [1998]).³ In addition, it is well documented that myriad mathematically-inclined academics (including Paul Erdős) responded incredulously to the original Monty Hall puzzle (Rosenhouse 2009; Vazsonyi 1999; Vos Savant 1990).⁴

The hot hand fallacy is a well-known cognitive illusion that refers to peoples' tendency to believe that success breeds success, even when it does not. In particular, in the seminal study of Gilovich,

¹This version of the problem is written by a reader of Vos Savant's columns, Craig F. Whitaker. The problem (without the goats) can be found previously in Selvin (1975) and Nalebuff (1987a), and is identical to the earlier Three Prisoners problem, which was popularized by Gardner (1959).

²The necessary assumptions are that: (1) Monty always reveals a goat, and (2) if the car is behind the contestant's door, then each goat has a non-zero probability of being revealed by Monty.

³Simple math puzzles of this sort have been noted for their importance in stimulating research ideas, and illustrating principles from microeconomic theory (Fehr and Tyran 2005; Friedman 1998; Kluger and Wyatt 2004; Nalebuff 1987b). The Monty Hall problem, in particular, has been studied extensively: see e.g. Bar-Hillel and Falk (1982); Chen and Wang (2010); Fox and Levav (2004); Kluger and Friedman (2010); Page (1998); Palacios-Huerta (2003); Selvin (1975); Slembeck and Tyran (2004); Tor and Bazerman (2003); Tubau, Aguilar-Lleyda, and Johnson (2015).

⁴The prominent 20th century mathematician Paul Erdős resisted the correct solution until he reluctantly accepted when it was demonstrated to him via computer simulation (Vazsonyi 1999).

Vallone, and Tversky (1985) the authors found that basketball players shoot no better after having just made several shots in a row, despite a near unanimous belief reported by players, coaches, and fans that players shoot better in these situations. Further, when confronted with the scientific evidence against their beliefs, practitioners (players and coaches) were left unpersuaded, leading the hot hand to become known as a “massive and widespread cognitive illusion” (Kahneman 2011).⁵

In order to estimate a player’s probability of making a shot conditional on having made several in a row Gilovich et al. (1985) and subsequent studies: (1) selected only the shot attempts that immediately follow several (e.g. three) made shots, and then (2) calculated the player’s shooting percentage on these shots. While, intuition suggests that this would yield an unbiased estimate of the player’s probability of success on these shots, it has recently been shown that this intuition is wrong (Miller and Sanjurjo 2015). The reason why is that the procedure biases the researcher towards over-selecting missed shots, which leads to an underestimate of the player’s probability of success on these shots. In fact, this bias is large enough to invalidate, and reverse, the results of the original hot hand fallacy study, as well as its replications (Miller and Sanjurjo 2015).

While on the surface it does not appear that there is any connection between why people have difficulty understanding the Monty Hall problem and why researchers have overlooked the bias in common measures of the hot hand for the past three three decades, we show that the two are in fact intimately related. The first step to draw out the connection is to recognize that both environments involve a procedure that selects an observation for analysis based on information that relates to the outcomes of other observations in the same dataset. In particular, just as Monty offers the contestant to switch to door #2, knowing that a goat is behind the door he just opened (#3), the hot hand researcher selects shot 42 from a sequence of basketball shots, knowing that “make” was the outcome of shots 39 through 41. The final step to bridge the two environments is to use an intuitive tool that determines exactly what information is revealed by each selection procedure.

The tool that we use is called the *principle of restricted choice*. It is an intuitive inferential rule

⁵The hot hand fallacy has been given considerable weight as a candidate explanation for various puzzles and behavioral anomalies identified in the domains of financial markets, sports wagering, casino gambling, and lotteries. See, e.g. Arkes (2011); Avery and Chevalier (1999); Barberis and Thaler (2003); Brown and Sauer (1993); Camerer (1989); Croson and Sundali (2005); De Bondt (1993); De Long, Shleifer, Summers, and Waldmann (1991); Durham, Hertz, and Martin (2005); Galbo-Jørgensen, Suetens, and Tyran (2016); Guryan and Kearney (2008); Kahneman and Riepe (1998); Lee and Smith (2002); Loh and Warachka (2012); Malkiel (2011); Narayanan and Manchanda (2012); Paul and Weinbach (2005); Rabin and Vayanos (2010); Sinkey and Logan (2013); Smith, Levere, and Kurtzman (2009); Sundali and Croson (2006); Xu and Harvey (2014); Yuan, Sun, and Siu (2014)

from the card game of contract bridge that makes clear the information revealed by the optimizing behavior of a constrained opponent (Reese 1960, chap. 3 p. 26). In Section 2 we first illustrate the principle with a simple example of (stochastic) choice from a menu, and then show that it is naturally quantified as the updating factor from the odds formulation of Bayes' rule.

In order to illustrate the simplicity, utility, and generality of thinking in terms of restricted choice, in Section 3 we apply it to solve a number of classic probability paradoxes, including the Monty Hall problem.⁶ This exercise makes clear that restricted choice renders intuitive the typically difficult counterfactual (and hypothetical) reasoning that is inherent in Bayesian updating. By contrast, we show that some intuitive heuristic approaches used by others, which provide an arguably simpler path to the correct solution in one problem or another, can lead to mistakes because they do not readily generalize across problems.

In Section 4 we introduce a progression of novel coin flip probability puzzles that culminate in the *Streak Reversal Paradox*, which is a recasting of the hot hand selection bias. We first show that the simplest of the coin flip puzzles is essentially equivalent to the Monty Hall problem; Monty selecting a door with a goat and asking the contestant to think about the odds of a car being behind another door is just like the researcher selecting a heads from the record of a sequence of coin flips and asking for the odds of a tails on another flip. On the other hand, in the Streak Reversal Paradox the researcher now selects k consecutive heads and asks for the odds of a tails on the next (right adjacent) flip.

The explicit connection between the hot hand selection bias, when posed as a coin flip paradox, and the Monty Hall problem, suggests that familiarity with restricted choice thinking could have helped researchers realize (and avoid) the types of counter-intuitive mistakes that can arise when particular observations are selected for analysis on the basis of the outcomes of other observations within the same dataset. As such, in Section 5, aside from the hot hand fallacy application, we consider a few other empirical examples in which restricted choice thinking can help researchers avoid making similar mistakes.⁷ The first is in Schelling (1971), a well-known study of segregation. The second pertains to studies that are subject to the well-known Berkson's bias. Lastly, we provide

⁶Reese (1960, p. 29) illustrates the principle of restricted choice with a problem that is nearly identical to the Monty Hall problem. Gillman (1992) appears to be the first to use the principle to explain the intuition behind the Monty Hall problem.

⁷For a prominent example of the empirical implications of the Monty Hall problem, see Chen (2008), which we discuss in Appendix B.

a hypothetical example of ESP research gone wrong, which illustrates some further pitfalls that we as researchers can avoid by using restricted choice thinking.

2 The Principle of Restricted Choice

“Since East could have played either card indifferently from K-Q, the fact that he has played one affords an indication that he does not hold the other” (Reese 1960)

The *principle of restricted choice* was introduced in the context of the card game contract bridge, in order to illustrate the information that is revealed by the actions of an agent with a known decision rule (Reese 1960, chap. 3 p. 26).⁸ However, because the rules of contract bridge are complex, and not necessary to understand the principle, we illustrate it in the following simpler setting.⁹

Suppose that we work in a restaurant, and know that Ann and Bob are equally likely to come in and order something from the menu. In addition, we know that Ann is indifferent between all 100 alternatives on the menu, so orders one of them at random (uniformly). On the other hand, Bob strictly prefers the hamburger, so will order it for sure. Now, suppose that we receive an order for a hamburger. Who is more likely to be the customer: Ann or Bob?

One incorrect intuition is to conclude that we have learned nothing from the hamburger order because it does not rule out either Ann or Bob. The glaring omission in this reasoning, of course, is that it does not take into account how Ann and Bob choose items from the menu.

By contrast, once we take into account Ann and Bob’s respective choice rules, our intuition immediately tells us the correct answer: because ordering a hamburger is more consistent with Bob than Ann, the customer is now more likely to be Bob. This is the principle of *restricted choice*; Bob is more restricted to choose hamburger than Ann because while Ann may order the hamburger, Bob must. As a result, once we find out that a customer ordered hamburger the odds that the customer is Bob, rather than Ann, increase.

⁸Reese (1960, chap. 3 p. 26) credits Alan Truscott, who wrote the daily bridge column for the New York Times from 1964 to 2005, for introducing restricted choice to the bridge community in the 1950s. Prior to that Borel and Chéron (1940) use the concept, at least implicitly, by applying Bayes’ rule to calculate probabilities in bridge problems.

⁹The professional bridge player Chris Willenken is quoted: “It takes at least 12 hours of study before you should even sit down at a table,” (<http://www.newsweek.com/bridge-world-rocked-top-players-busted-377135>).

2.1 Formal Definition

The principle of restricted choice is an informal *intuition* for why beliefs should shift in a particular direction upon the arrival of new information. The intuition illustrates the essential qualitative feature of Bayesian updating, i.e. the re-scaling of the prior odds by the likelihood ratio.

To see the connection, let \mathcal{A} (e.g. “Ann”) and \mathcal{B} (e.g. “Bob”) represent two hypothetical models that could have determined (i.e. “chosen”) the observed outcome c . Further, let the prior odds in favor of model \mathcal{B} having determined the observed outcome (relative to \mathcal{A}), be defined as $\mathbb{O}(\mathcal{B}, \mathcal{A}) := \mathbb{P}(\mathcal{B})/\mathbb{P}(\mathcal{A})$. The task is, upon observing an outcome c , to infer the (posterior) odds in favor of model \mathcal{B} (relative to \mathcal{A}), which are given by $\mathbb{O}(\mathcal{B}, \mathcal{A}|c) := \mathbb{P}(\mathcal{B}|c)/\mathbb{P}(\mathcal{A}|c)$. Bayes’ Rule then defines the relationship between prior and posterior odds:¹⁰

$$\mathbb{O}(\mathcal{B}, \mathcal{A}|c) = \mathcal{R}_{\mathcal{B}:\mathcal{A}}(c) \times \mathbb{O}(\mathcal{B}, \mathcal{A})$$

where $\mathcal{R}_{\mathcal{B}:\mathcal{A}}(c) := \mathbb{P}(c|\mathcal{B})/\mathbb{P}(c|\mathcal{A})$ is the likelihood ratio, or Bayes factor.

We interpret the likelihood ratio $\mathcal{R}_{\mathcal{B}:\mathcal{A}}(c) \in [0, \infty)$ as the relative degree to which model \mathcal{B} (as compared to \mathcal{A}) is restricted to “choose” c . In this sense, we have quantified the principle of restricted choice. For example, if $\mathcal{R}_{\mathcal{B}:\mathcal{A}}(c) > 1$ then model \mathcal{B} is relatively more restricted to choose c . Thus, in accordance with Bayes’ rule, the principle of restricted choice tells us that the posterior odds in favor of \mathcal{B} must be greater than the prior odds, and $\mathcal{R}_{\mathcal{B}:\mathcal{A}}(c)$ tells us precisely by how much.

To illustrate using the example given above, while either Ann or Bob could have ordered the hamburger, the knowledge that Bob is 100 times more restricted to choose the hamburger than Ann—i.e. $\mathcal{R}_{\mathcal{B}:\mathcal{A}}(\text{hamburger}) = 100 > 1$ —allows us to conclude that the odds in favor of the customer being Bob must increase by a factor of 100. Therefore, because we know that the prior odds in favor of Bob (relative to Ann) are 1:1, i.e. $\mathbb{O}(\mathcal{B}, \mathcal{A}) = 1$, then the posterior odds must be 100:1. Thus, because for every 100 chances in favor of Bob there is 1 chance in favor of Ann, the probability of Bob being the customer is $\mathbb{P}(\mathcal{B}|\text{hamburger}) = 100/101$.

Finally, the principle of restricted choice highlights the advantages of representing uncertainty with odds, rather than probabilities: (1) it is simpler because it involves the binary comparison of just two candidate models at a time, (2) it is easier for computation because it involves a

¹⁰This holds when $\mathbb{P}(c|\mathcal{A}) > 0$.

less complex updating factor, and (3) it makes the restricted choice intuition more salient. See Appendix A for a more complete discussion. In the following section we illustrate how easy it is to apply restricted choice reasoning when uncertainty is represented by odds.

3 Restricted choice in classic conditional probability “paradoxes”

We show how the simplicity and intuition of restricted choice reasoning extends to three related classic conditional probability puzzles that tend to confound the intuitions of laypersons and experts alike: the Monty Hall problem, two versions of the Boy-or-Girl Paradox, and their progenitor: Bertrand’s Box Paradox (Gorroochurn 2012; Rosenhouse 2009).

Problem 1 (Bertrand’s Box Paradox) *Three boxes are identical in external appearance. The first box contains two gold coins, the second two silver coins, and the third one gold coin and one silver coin. You choose a box at random, and draw a coin. Suppose that you draw a gold coin. What is the probability that the other coin is also a gold? (Bertrand [1889]; adapted from Rosenhouse [2009])*

Given that a gold coin was drawn, it is impossible that the all silver box was chosen. Thus, two possible boxes remain: the all gold, and the mixed. Given this, it becomes intuitively appealing to conclude that the probability that the other coin is gold is $1/2$. However, what this reasoning misses is that when a draw occurs from the all gold box one is more restricted to “choose” (draw) a gold coin, because with the gold box one *must* draw a gold coin, whereas with the mixed box one can draw either a gold or silver coin. From this it immediately follows that one is twice as restricted to choose the gold coin from the all gold box relative to the mixed box; i.e. $\mathcal{R} = \mathbb{P}(\text{gold coin}|\text{gold box})/\mathbb{P}(\text{gold coin}|\text{mixed box}) = 1/(1/2) = 2$. Therefore, by the principle of restricted choice, the updated odds in favor of the draw having come from the all gold box are 2:1, double the prior odds of 1:1; i.e.

$$\mathbb{O}(\text{gold box, mixed box} \mid \text{gold coin}) = 2 \times \mathbb{O}(\text{gold box, mixed box}) = 2$$

Thus, given that there are only two boxes from which the draw could have been made, the probability that it was made from the all gold box is $2/3$. As a result, the probability that the other

coin is also a gold is not equal to $1/2$, but $2/3$.

Next we consider the Boy-or-Girl paradox (as presented in Problem 1 of Bar-Hillel and Falk [1982]).¹¹

Problem 2 (Boy-or-Girl Paradox) *Mr. Smith is a father of two. We meet him walking along the street with a young boy whom he proudly introduces as his son. What is the probability that Mr. Smith's other child is also a boy? (Gardner [1961]; Bar-Hillel and Falk [1982])*

The intuitive answer to this problem is $1/2$, and under usual assumptions, this is correct. To illustrate, let us assume that Mr. Smith simply chooses his walking companion at random from among his two children (without discriminating). With this, the problem becomes close to the Bertrand's Box Paradox: Mr. Smith's children are drawn, without-replacement, from either an all-boy "box," an all-girl box, or a mixed-gender box. The key difference, however, is that now the two types of boxes are not all equally likely. In particular, the equivalent of the mixed box—one boy and one girl—has 2:1 prior odds in its favor, relative to any single gender box. This is due to there being two birth order possibilities in the mixed-gender box (Boy-girl and Girl-boy).¹² Indeed, as in Bertrand's Box Paradox, learning that the randomly chosen walking companion is a boy makes the posterior odds in favor of both children being boys (relative to mixed gender) double the prior odds. The reason why is that the choice of a boy is twice as restricted in the all boys case. Thus, because the prior odds were 1:2 "in favor" of all boys (relative to mixed gender), i.e. $\mathbb{O}(\text{all boys, mixed}) = 1/2$, the posterior odds are 1:1 in favor of two boys. Finally, because there remain only two possible compositions of children—all boys, or mixed gender—the probability that Mr. Smith has all boys is $1/2$. As a result, the probability that his other child is a boy is $1/2$.¹³

¹¹The original version appears in Gardner (1961), see Footnote 13 for further discussion.

¹²The use of birth order as a unique identifying label is not strictly necessary (and is not possible if the children are twins). Instead one may use any label that is independent of Mr. Smith's choice rule.

¹³The most common version of the Boy-or-Girl Paradox is as follows: "Mr. Smith says: 'I have two children and at least one of them is a boy.' Given this information, what is the probability that the other child is a boy?" (Fox and Levav 2004). If one assumes that Mr. Smith would say nothing (or its equivalent) in the case that he were to have two girls, then in this version of the problem Mr. Smith is equally restricted to report boy in the cases of Boy-girl, Girl-boy, and Boy-boy, so prior and posterior odds are identical. As a result, the correct probabilities can be computed simply by a careful enumeration of the sample space, and elimination of the impossible Girl-girl combination. Therefore, failure to see the correct answer in this version of the Boy-or-Girl Paradox can arise not because of a failure to incorporate the subtleties of Bayesian reasoning, but simply by a failure to appreciate the subtleties of the sample space. The classic example of this type of mistake is Leibniz's Error, i.e. believing that 11 and 12 are equally probable when rolling a pair of fair dice because there is just one way for each sum to be partitioned into two numbers less than (or equal to) 6 (Gorroochurn 2012).

We now consider another version of the Boy-or-Girl Paradox (presented as Problem 2 in Bar-Hillel and Falk [1982]).¹⁴

Problem 3 (Younger Boy-or-Girl Paradox) *Mr. Smith is a father of two. We meet him walking along the street with a young boy whom he proudly introduces as his eldest child. What is the probability that Mr. Smith's younger child is also a boy? (Bar-Hillel and Falk 1982).*

Because the younger child must be either a boy or a girl, the intuitively appealing response to this question is, again, $1/2$. This response happens to be correct if we assume that Mr. Smith chooses his walking companion at random, from among his two children, regardless of gender, as in Problem 2.

Nevertheless, while gender-neutrality may be natural to assume in one society, it may not be universal. Thus, another possibility is that Mr. Smith has the unfortunate attitude of only being willing to walk with a son. Assume that this is so, but that his discriminatory practices end there; that is, if he has two boys, then he is indifferent between walking companions, so chooses one of them at random. Under these assumptions, observing the gender of Mr. Smith's walking companion yields redundant information. That is, if we had merely observed Mr. Smith walking with a child, without any further information, we would already know that the child must be a boy, and that the possible birth order combinations are thus Boy-girl, Girl-boy, and Boy-boy.¹⁵

However, in the current problem we additionally discover that Mr. Smith's walking companion is his eldest child, which eliminates the possibility of Boy-girl, reducing the possibilities to Girl-boy and Boy-boy. With this, the intuitive response is again $1/2$, but now this response is wrong. This is because the associated reasoning fails to take into account that the degree of restrictedness in Mr. Smith's choice now varies across each of these hypothetical birth orders. In particular, if the younger child is a girl (Girl-boy), then Mr. Smith's choice of walking partner will be the older boy, for sure. On the other hand, when the younger child is also a boy (Boy-boy), then Mr. Smith is equally likely to choose each boy. This means that when the younger child is a girl, Mr. Smith's choice is twice as restricted. Therefore, the posterior (relative) odds in favor of Girl-boy are double the prior odds of 1:1. Thus, the probability of Girl-boy is $2/3$. As a result, the probability of

¹⁴We adapt the language slightly, but the problem is equivalent.

¹⁵See Footnote 13 for a discussion of this case, in which we learn that "at least one is a boy."

Boy-boy is $1/3$, i.e. there is a one-third chance that the younger child is a boy.¹⁶

We now consider the well-known Monty Hall problem, which is essentially identical to the version of the Younger Boy-or-Girl Paradox just discussed, in which Mr. Smith is only willing to walk with sons. With respect to the original statement of the problem in Vos Savant (1990), we change the door numbers (without loss of generality) in order to facilitate comparison with the coin flip problems presented in Section 4.

Problem 4 (Monty Hall problem) *Suppose you're on a game show, and you're given the choice of three doors: Behind one door is a car; behind the others, goats. You pick a door, say number 3, and the host (Monty Hall), who knows what's behind the doors, opens another door, say number 1, which has a goat. He then says to you, "Do you want to pick door number 2?" Is it to your advantage to switch your choice? (Selvin [1975]; adapted from Vos Savant [1990])*

As with the Boy-or-Girl Paradox, the correct answer depends on conditions that have not yet been specified. One possibility is that Monty, the host, follows a rule by which he must always reveal a goat from behind one of the two doors that the contestant does not choose. Further, in the case that Monty has two goats to choose from, he chooses a door (uniformly) at random.

Under these conditions, because one goat and one car will always remain covered once Monty reveals a goat, an intuitively appealing conclusion is that the odds in favor of the car being behind door number 2 are 1:1 (relative to door number 3), meaning that the contestant should be indifferent about switching.

Nevertheless, as in the previous problems, there is an error in this simple reasoning. To see why, notice first that before Monty opens door number 1, the contents behind doors number 1 and 2, respectively, are one of the following, each with equal probability: Car-goat, Goat-car, or Goat-goat. However, once Monty reveals a goat behind door number 1, the remaining possible arrangements become Goat-car and Goat-goat. Now, because Monty must open the first door in the case of Goat-car, whereas he only opens it half of the time in the case of Goat-goat, he is twice as restricted to open it in the case of Goat-car. Therefore, given that the prior (relative) odds in favor of Goat-car were 1:1, the posterior odds must double, i.e. the odds in favor of the car being

¹⁶More formally, let B_e be the event that the elder child is a boy, and B_y be the event that the younger child is a boy. Further, let $G_y B_e$ be the conjunction event $G_y \cap B_e$. Finally, the restrictedness of $G_y B_e$, for the choice of B_e , is given by $\mathcal{R} = \mathbb{P}(B_e | G_y B_e) / \mathbb{P}(B_e | B_y B_e) = 1 / (1/2) = 2$. Therefore, by the principle of restricted choice, the posterior odds are $\mathbb{O}(G_y B_e, B_y B_e | B_e) = 2 \times \mathbb{O}(G_y B_e, B_y B_e) = 2$.

behind door number 2 are now 2:1 (relative to door number 3). As a result, it is in the contestant’s interests to switch doors, as the probability of winning the car by doing so is $2/3$.¹⁷

3.1 Restricted choice as a Simple, Intuitive, and Unified Approach to Solving the Classic Conditional Probability “Paradoxes”

In addition to being simple and effective for solving particular problems, the fact that restricted choice uses Bayesian reasoning to account for each problem’s selection procedure makes it completely general. By contrast, intuitive heuristic approaches to solving these problems are ad-hoc, so do not readily generalize across problems, as we will now discuss.

A common explanation for the Monty Hall problem—which is the highest voted answer on the question-and-answer website math.stackexchange.com (Becker 2017)—involves framing the problem from the *ex ante* perspective, i.e. before Monty opens the door (see e.g. Krauss and Wang [2003]). In this case, it is clear that if the contestant commits to switching, then he will get the best of the other 2 doors, so will win the car $2/3$ of the time, i.e. if the contents of the first two doors are Goat-car or Car-goat then switching (from door 3) yields the car, and if the contents are Goat-goat then switching yields a goat. Indeed, people in experiments have greater success in versions of the problem that encourage this type of *ex ante* reasoning (Krauss and Wang 2003).¹⁸ While intuitively appealing, this solution is ad-hoc, so does not readily generalize to other problems. For one, it is not straightforward how to apply it to, e.g., Problems 1-3 presented above. More substantively, because the *ex ante* reasoning ignores: (i) which door Monty actually opens, and (ii) Monty’s rule for choosing which door to open when the contents of the first two doors is Goat-goat, it cannot possibly provide (unless by coincidence) the conditional probability of interest, i.e. the probability of winning the car by switching, *conditional* on Monty having opened the first door (according to

¹⁷Alternatively, one can slightly change Monty’s selection rule to make the problem essentially identical to the Boy or Girl problem (problem 2): Monty selects one of the two doors that the contestant did not choose, at random, and opens it, regardless of whether it covers the goat or car, and he happens to reveal a goat. This version is known as “Ignorant Monty” (Granberg and Brown 1995; Rosenthal 2008).

¹⁸In Krauss and Wang (2003) the “less-is-more” manipulation (“1D”), which underlies all further manipulations that increase correct responses, involves an *ex ante* framing of the problem, which facilitates reasoning in terms of natural frequencies.

his rule).^{19,20,21}

In a well-known article on conditional probability brain teasers Bar-Hillel and Falk (1982, pp. 116–117) provide an intuitive sample space explanation of the Boy-or-Girl Paradox (Problem 2) as an alternative to their Bayes’ rule calculations. In particular, they argue that the birth order label given to each child is irrelevant, and that sibling pairs in the sample space can be labeled by the gender of the child that we meet (m) and the child that is at home (h), i.e. $\{B_m B_h, B_m G_h, G_m B_h, G_m G_h\}$. While this approach happens to give the correct answer to Problem 2 because Mr. Smith chooses his walking companion at random (uniformly) from among the two children—i.e. independent of any exogenous identifying label—it does not extend to settings such as that of Problem 3, in which Mr. Smith’s choice is not equally restricted across possible sibling pairs. Instead, his choice there among the two children depends on their gender labels, which consequently renders the meet/home labels endogenous.²²

Lastly, there is a well-known way of modifying the Monty Hall problem in order to make more salient that the odds change when Monty opens a door to reveal a goat. In particular, in the Parade magazine article in which she first discussed the Monty Hall problem, Vos Savant (1990) states:

Here’s a good way to visualize what happened. Suppose there are a million doors, and you pick door #1. Then the host, who knows what’s behind the doors and will always avoid the one with the prize, opens them all except door #777,777. You’d switch to that door pretty fast, wouldn’t you?

Vos Savant’s modification effectively conveys the restricted choice intuition in a way that helps make

¹⁹The ex ante approach yields the incorrect probability of winning the car by switching in any version of the Monty Hall problem in which Monty does not happen to deal with the case of Goat-Goat by choosing each of the doors with precisely the same probability. To illustrate, if the probability with which Monty opens the first door (rather than the second) is instead $p \neq .5$, then Monty is $1/p \neq 2$ times as restricted to open door 1 in the case of Goat-car as in the case of Goat-goat, which results in posterior odds of $1/p : 1$ in favor of Goat-car. Further, in contrast with what the ex ante reasoning suggests, in the case of $p = 1$ it is not even advantageous for the contestant to switch.

²⁰For another example in which the requested conditional probability can easily change based on minor alterations in the decision rule of the host, which the ex ante approach cannot take into account, see the coin-flip Monty Hall problem that we introduce in Problem 5 of Section 4, along with the ensuing discussion.

²¹While the ex ante approach leads to incorrect conditional probabilities, it is possible to use a natural frequency intuition to arrive at the correct conditional probabilities for the Monty Hall problem. This can be done by adapting the natural sampling approach that Gigerenzer and Hoffrage (1995) use for Bayesian updating.

²²While Bar-Hillel and Falk (1982) do consider alternative background assumptions for Problem 2, in which Mr. Smith selects a male walking companion whenever possible, they do not mention that the argument for the birth order label being irrelevant is not valid in settings such as Problem 3, in which the birth order is revealed.

the correct answer—to switch doors—more transparent.^{23,24} Indeed, when experimental subjects face a modified Monty Hall problem of this type (a 100 door version) they correctly decide to switch doors at a much higher rate (approx. 85%) than when facing the standard version of the problem (approx. 15%) (Page 1998). However, while distorting the odds gives subjects a better sense of the direction in which the odds have shifted, which is sufficient to improve the rate of correct responses in the Monty Hall problem, it is incomplete in the sense that it does not indicate how to compute the posterior odds.²⁵ This is important because in real-world decisions an informed agent will typically not only wish to decide whether to switch or not, but also ascertain the value of switching. The many-doors intuition provides no guidance for how to do this. Further, as a general approach to solving conditional probability problems, not only does the many-doors approach increase the complexity without providing the requested probability, but in order for it to deliver the correct intuition it requires one to already understand how restricted choice operates in a particular problem.²⁶

4 Restricted choice in novel coin flip puzzles

“The player’s performance, then, can be compared to a sequence of hits and misses generated by tossing a coin.” (Gilovich et al. 1985)

First we break the simplest version of the hot hand selection bias of Miller and Sanjurjo (2015) into two probability puzzles (*Next Flip Paradox*, *Alternation Paradox*), which allows us to illustrate its tight relationship with the classic conditional probability puzzles discussed in Section 3. Then, we recast the more general version of the selection bias as a statistical puzzle (*Streak Reversal*

²³Note that in order for #777,777 to be strictly preferred among the two remaining doors we must assume that when the car is behind door #1, Monty places a non-zero probability on opening a door other than door #777,777.

²⁴In particular, because Monty must leave door #777,777 closed when the car is behind it, whereas he has a 1/999,999 probability of leaving it closed when the car is behind door #1, he is 999,999 times more restricted to leave door #777,777 closed when the car is behind door #777,777 (vs. door #1). Therefore because the prior odds between the two doors are 1:1, the posterior odds become 999,999:1 in favor of door #777,777.

²⁵Further, facing the 100 door version of the problem seems to cause only a marginal improvement in peoples’ understanding of the problem, as when experimental subjects then go back to the standard version they proceed to make the wrong choice at rates similar to subjects that never faced the 100 door version (Page 1998).

²⁶For example, it may not be straightforward to individuals how to apply the one-million-doors modification to the Boy-or-girl problems (Problems 2 and 3) or Bertrand’s box paradox (Problem 1). In particular, in the Boy-or-girl problems the analogous modification is for Mr. Smith to walk with all but one of his 999,999 children, and to meet him walking with only boys. In Bertrand’s box paradox there would be 999,999 coins in each box, 999,998 coins would be drawn, and they would all need to be gold.

Paradox), which we use in Section 5 to illustrate some empirical applications.

The following problem, which we call the Next Flip Paradox, is the basic building block for its successor: the Alternation Paradox, which is the simplest version of the hot hand selection bias (Miller and Sanjurjo 2015).

Problem 5 (Next Flip Paradox) *Jack flips a coin three times, then tells you that the first flip is a heads. What is the probability that the second flip is also a heads?*

As in the problems presented in Section 2, the answer depends on conditions that have not yet been specified. In particular, if Jack had decided to select which of the first two flip outcomes to reveal at random, or had simply planned on always revealing the outcome of the first flip, then the correct answer is $1/2$. This is precisely as in the Boy-or-Girl Paradox discussed in Section 2, in which Mr. Smith chooses a child at random, regardless of the gender (Problem 2). On the other hand, if Jack had instead only been interested in the respondent's beliefs about the probability that a heads follows a heads, so decided to select one of the first two flip outcomes at random based on the criteria that it be a heads, then the answer changes.²⁷

The Next Flip Paradox is nearly identical to the standard Monty Hall problem. In particular, just as Monty looks behind each door before opening one, which in turn reveals information regarding the location of the car, Jack looks at the outcome of each flip before selecting one, which in turn reveals information about the location of heads. To see this, let Jack now be the game show host instead of Monty, with the game proceeding as follows: (1) Jacks flips three coins, leaving each behind a separate door, (2) he asks the contestant to choose one of the three doors, informing her that she will receive a prize if the door she chooses conceals a tails flip, (3) once the contestant has chosen one of the doors (flips) Jack opens one of the other two doors at random, according to the criteria that he must reveal a heads flip (in the case of two tails flips he cannot open either door), and finally (4) he offers the contestant the opportunity to switch. As in Problem 4, we assume (without loss of generality) that the contestant's initial choice is the third door, and that Jack opens the first door, revealing the first flip to be a heads. In this case, the first two flip outcomes must be either Heads-tails or Heads-heads. Then, by the same restricted choice reasoning as in the Monty Hall problem, Jack is twice as restricted to open the first door in the case of Heads-tails, as

²⁷In this case, Jack of course would not have asked the question if the first two flips had come up tails.

in the case of Heads-heads. As a result, the probability that the second flip is a heads is $1/3$.²⁸

While it is true that the contestant can extract information about the second flip from the knowledge that the first flip is a heads, this does not imply that coins have memory. Instead, the contestant exploits the fact that Jack has inspected the outcome of the first two flips before choosing, which means that Jack's choice (probabilistically) reflects his knowledge. More subtly, this also implies that time's arrow is irrelevant, i.e. if Jack were to instead reveal that the second flip were a heads, then the probability of heads on the previous (first) flip would similarly be $1/3$.^{29,30}

This leads to the second version of the problem, The Alternation Paradox, which brings us one step closer to illustrating the hot hand selection bias, and happens to be the exact probabilistic representation of the simple 3-flip example of the bias given in Table 1 of Miller and Sanjurjo (2015).³¹

Problem 6 (Alternation Paradox) *Jack will flip a coin three times, then select a flip that is immediately preceded by a heads, at random. Assuming that Jack has a flip to select, what is the probability that the selected flip is a heads?*

In order for Jack to select a flip, he must inspect the outcomes of the first two flips. Given that at least one of the two has come up heads, it is clearly impossible that the sequence could have started with TT. Let $H_{..}$ be the event that Jack selects the second flip, which is preceded by a heads on the first flip, and $_{.}H_{.}$ be the event that Jack selects the third flip, which is preceded by a heads on the second flip. Conditional on Jack having chosen a flip, these events are equally likely.³² Now, in the case that Jack selects the second flip the probability of heads is simply the

²⁸A slight modification makes the Next Flip Paradox identical to the Younger Boy-or-Girl Paradox from Section 2: Jack flips the coin twice, then chooses one of the heads flips at random (final flip included), and tells you that it is the first flip.

²⁹Problem 5 is similar to Galton (1894)'s Three Coin Puzzle, which, while traditionally a sampling problem (Gor-roochurn 2012), can be re-cast as a conditional probability puzzle. In particular, suppose that a researcher tosses 3 coins, then picks out two of the coins that match, and reports that the first two coins are heads. This raises the question: what is the probability that the third coin is a tails? The answer is, under natural assumptions, that the researcher is three times as restricted to choose the first two flips in the case of HHT than in the case of HHH. As a result, the probability of tails is $3/4$.

³⁰The probability that the third flip is heads does not depend on which flip Jack reveals because its outcome cannot restrict his choice among the first two flips.

³¹Table 1 of Miller and Sanjurjo (2015) demonstrates that the expected value of the natural (maximum likelihood) estimator for the probability of heads given heads—the proportion of heads on flips immediately preceded by heads—is $5/12$.

³²If the sequence begins with HT then Jack will choose the first flip, if the sequence begins with TH then Jack will choose the second flip, and if the sequence begins with HH then Jack will choose each flip with equal chance.

solution to the Next Flip Paradox, i.e. $\mathbb{P}(\text{HH}_-|\text{H}_-) = 1/3$. On the other hand, if Jack selects the third flip then $\mathbb{P}(\text{_-HH}|\text{H}_-) = \mathbb{P}(\text{_-H}) = 1/2$, as the outcome of the last flip cannot restrict Jack's choice of which immediate heads successor to select. Then it immediately follows that:

$$\begin{aligned} \mathbb{P}(\text{heads} \mid \text{flip preceded by a heads}) &= \mathbb{P}(\text{HH}_-|\text{H}_-) \times \mathbb{P}(\text{H}_-) + \mathbb{P}(\text{_-HH}|\text{H}_-) \times \mathbb{P}(\text{H}_-) \\ &= \frac{1}{3} \times \frac{1}{2} + \frac{1}{2} \times \frac{1}{2} \\ &= \frac{5}{12} \end{aligned}$$

The next problem extends the Alternation Paradox to 100 flips, and streaks of length 3. It illustrates the hot hand selection bias found in Miller and Sanjurjo (2015).

Problem 7 (Streak Reversal Paradox) *Jack, now a researcher, observes the outcome of 100 flips of a fair coin. He collects all of the flips that are immediately preceded by 3 consecutive heads, and calculates the proportion of heads, $\hat{p}(H|3H)$, on these flips. He expects this proportion to be 0.5. Is he correct?*

While Jack's expectation is intuitively appealing, it turns out to be incorrect. To see why, notice that a flip is representative of those collected from a particular sequence if it is chosen at random from all of the flips in the sequence that are immediately preceded by 3 consecutive heads. The probability that one of these randomly chosen flips is a heads equals the proportion of heads among the collected flips, i.e. the *empirical probability*. It follows by the law of total probability that, prior to flipping, the probability of heads on a flip that will be chosen at random from all of the flips in a sequence that are immediately preceded by three consecutive heads is equal to the expected value of the proportion across all possible sequences.³³ This expected value happens to be equal to .46 (see Miller and Sanjurjo [2015] for the formula). As a result, the corresponding probability of heads for the randomly selected flip is not .50, but .46.

Just as the principle of restricted choice provided intuition for the answer to the Next Flip Paradox, it can do the same for the Streak Reversal Paradox. To see how, begin by imagining that Jack selects a flip at random, from among all of those that are immediately preceded by three consecutive heads. Now, if he were to select, say, flip number 42, then intuition would lead us to

³³Possible sequences are those that allow the proportion to be well-defined, i.e. those that contain at least one flip that is immediately preceded by three (or more) consecutive heads.

expect the odds of heads to be 1:1. In fact, this intuition would be correct if Jack were to have selected the flip before having examined the sequence. However, because Jack instead examined the sequence first, then selected flip 42 by virtue of it being among those flips that are immediately preceded by three consecutive heads, this intuition is incorrect.

Given how flip 42 was selected, we can now use restricted choice to explain why the selection procedure makes flip 42 more likely to be a tails. In particular, first observe that if flip 42 were a heads, then flips 39-42 would be HHHH, meaning that flip 43 would also immediately follow (at least) three consecutive heads. Thus, in this case Jack could have selected flip 43 instead of flip 42. On the other hand, if flip 42 were instead a tails, then flips 39-42 would be HHHT, meaning that Jack would be restricted from selecting flip 43 (or 44, or 45). This implies that if flip 42 were a tails rather than a heads then Jack would be relatively more likely to select flip 42, as there would be fewer eligible flips (on average) in the sequence from which to select. Thus, the fact that Jack is more restricted to select flip 42 in the case that it is a tails makes the likelihood that the flip he selected was a tails greater than the unconditional probability of flipping a tails.³⁴

5 Empirical applications

We provide four empirical examples of how understanding restricted choice thinking can sometimes help us as researchers to avoid making critical mistakes in our design of experiments, analysis of data, and interpretation of results.³⁵

5.1 Hot hand fallacy

“The hot-hand and streak shooting concern the probability (P) that a player will hit the next shot given his previous hits and misses” (Tversky and Gilovich 1989)

Given that we have shown in Section 4 that the principle of restricted choice can be used to illustrate the selection bias present in the coin flip puzzles, it is now straightforward to demonstrate its connection with the critical mistake made in the seminal study of the hand fallacy by Gilovich,

³⁴This reasoning holds for any flip the Jack may select, unless it happens to be the final flip of the sequence, in which case the posterior odds of a heads vs. tails on the selected flip are the same as the prior odds. The reason why is that the outcome of the final flip does not influence the number of flips that are eligible for selection.

³⁵For a prominent example of the empirical implications of the Monty Hall problem, see Appendix B

Vallone, and Tversky (1985; “GVT”), and others that followed. GVT conduct a controlled shooting experiment in which collegiate basketball players attempt 100 shots, from locations on the court at which they are expected to make half of them. Under this design target basketball shot outcomes are like those of a coin (i.e. generated by an i.i.d. Bernoulli process), so the authors make precisely the same intuitively appealing assumption as the researcher in the Streak Reversal Paradox. In particular, GVT conduct a two-sample test of proportions (paired t-test) for each player, in which the first sample consists of the player’s shots that are immediately preceded by three consecutive makes (i.e. a streak of successes), and the second sample consists of the player’s shots that are immediately preceded by three consecutive misses (i.e. a streak of failures). GVT’s null hypothesis is that the expected proportion of makes should be the same in each sample of shots.

Given the solutions to the coin flip problems discussed above, it immediately follows that GVT’s intuitive assumption is incorrect, and leads to a biased test. In particular, there is a severe downward bias in their estimate of the probability that a “make” will immediately follow three consecutive makes, and an equal and opposite bias in their estimate of the probability that a make will immediately follow three consecutive misses. Specifically, under GVT’s design target and null hypothesis, the true expected proportion of makes when a 50% shooter is on a hit streak is not .50, but .46. By symmetry, the expected proportion of makes when on a miss streak is .54. Thus, when taking the difference, one induces a bias of 8 percentage points.³⁶ Further, when one applies the bias correction to each of the shooters in GVT’s data, who tended to deviate slightly from the design target, the positive 3 percentage point average hot hand effect reported by the authors (not statistically significant), becomes a statistically significant 13 percentage point effect (Miller and Sanjurjo 2015).³⁷

Similarly biased measures were also used in the replications of GVT’s original study: a close replication of GVT’s controlled shooting study with olympic basketball players (Avugos, Bar-Eli, Ritov, and Sher 2013), and an analysis of the data from the NBA three point shootout (Koehler and Conley 2003), which has been referred to as “an ideal situation in which to study the hot hand”

³⁶In fact, the bias is actually a bit more severe than this, due to an additional selection effect that is driven by the exclusion of sequences that do not have both: (1) at least one shot that immediately follows a streak of made shots, and (2) at least one shot that immediately follows a streak of missed shots.

³⁷To put this number into perspective, the difference between the median three point shooter and the top three point shooter in the 2015-2016 NBA season was 12 percentage points. Taken from ESPN, “NBA Player 3-Point Shooting Statistics - 2015-16.” <http://www.espn.com/nba/statistics/player/-/stat/3-points> [accessed September 24, 2016].

(Thaler and Sunstein 2008). As in the case of the re-analysis of the Gilovich et al. (1985) data, in contrast with the original results, a bias-corrected analysis reveals significant and substantial hot hand shooting in both datasets (Miller and Sanjurjo 2015).

5.2 Clustering and segregation

“Patterns, though, can be deceptive, and it is useful to have some measures of segregation or concentration or clustering or sorting. One possible measure is the average proportion of neighbors of like or opposite color.” (Schelling 1971)

While the hot hand selection bias and the coin flip puzzles pertain to time series data, in Section 4 we have noted that the time component is not central to the bias. Instead, the key is that the selection criteria is based on the outcomes of consecutive, i.e. spatially adjacent, flips. This in turn suggests the possibility of a more general spatial selection bias if we were to extend the setting to multiple dimensions.

Consider an $n \times n$ grid of cells, each populated with the outcome of a fair coin flip. Now, suppose that we are interested in the probability that a cell is a heads, given that all of its neighbors are heads. An intuitive way to estimate this probability would be to select the subset of all cells that are surrounded by heads cells, and then calculate the proportion of heads among these. However, this estimate will be biased downward due to a mechanism that is essentially identical to the spatial bias that emerges in the one-dimensional setting of the Streak Reversal Paradox (Problem 7).

That this measure of the similarity of a cell to its neighbors is biased suggests the possibility of such a bias appearing in studies of homophily in location preferences, e.g. racial segregation. Indeed, although it has no bearing on the validity of the main results, Schelling (1971, p. 156) briefly considers a similar biased measure of segregation. Specifically, Schelling writes:

If we count neighbors of like color and opposite color for each of the 138 randomly distributed stars and zeros in Figure 7, we find that zeros on the average have 53% of their neighbors of the same color, stars 46%. (The percentages can differ because stars and zeros can have different numbers of blank neighboring spaces.)

To illustrate the issue with this measure in a simpler way, that is similar to our discussion of the Streak Reversal Paradox (Problem 7), let the zeros and stars be arranged linearly (1-D) in

consecutive cells, with no blank cells in between them. Further, let each cell be a zero, or star, with equal chances. Now, consider the proportion of zeros for which both neighbors are also zeros. Here it is easy to see that in choosing any zero cell at random there are, on average, comparatively fewer cells that could have been chosen in the case that at least one neighbor is a star, i.e. the choice is more restricted, which means that the given cell is relatively more likely to have been chosen in this case. This means that the posterior odds make it relatively less likely that both neighbors are zeros, i.e. the proportion of zeros that are flanked on both sides by zeros is expected to be less than .25. This in turn implies that because the zeros and stars are distributed at random, and a zero is selected at random, the average proportion of neighbors that are also zero is expected to be less than 50%. More generally, the extent to which the magnitude of the bias is empirically relevant will depend on the definition of a cluster and the size of the grid under consideration.

5.3 Berkson’s Paradox

“... the outcomes of two independent coins are rendered dependent by the testimony that at least one of them is a tail.” (Pearl 2009a)

Berkson’s bias (aka Berkson’s paradox) is a classic form of selection bias that was introduced with a hypothetical example of two diseases that, while not associated in the general population, surprisingly become negatively associated in the population of hospitalized patients (Berkson 1946).³⁸ Berkson’s bias can be illustrated with the following simple example:

Problem 8 (Berkson’s Paradox) *Suppose that a randomly selected high school student has a 50 percent chance of having good SAT scores, along with a 50 percent chance of having good grades, and that the attributes are independent. Further, suppose that every student with at least one good attribute applies to university, and highlights his/her single best attribute in the application. If an applicant highlights good grades, then what is the probability that the applicant has good SAT scores? (Adapted from Pearl [2009b]).*

Assuming that an applicant highlights an attribute at random (uniformly) in the case that both attributes are good, this problem is identical to the Younger Boy-or-Girl Paradox (Problem 3).³⁹

³⁸Berkson’s bias is also known as the admission rate bias (Sackett 1979)

³⁹It is also identical to a two coin version of the Next Flip Paradox (Problem 5).

The reason why is that each attribute is good, or not good, with a 50-50 chance, just as each child is a boy, or not a boy, with a 50-50 chance. As a result, an applicant who has good grades and poor SAT scores is twice as restricted to highlight good grades, as compared to an applicant with both good grades and good SAT scores. Thus, as prior odds are even, the principle of restricted choice leads to posterior odds of 2:1 in favor of the applicant having good grades and poor SAT scores, i.e. given good grades, the probability that the applicant has poor SAT scores is $2/3$.

While it remains true that, among the applicants with good grades, half of them also have good SAT scores, this subgroup constitutes just $1/3$ of the applicant pool. The remaining $2/3$ of the applicants, on the other hand, have just one good attribute. This implies that there will be a negative correlation between attributes in the applicant pool, despite the correlation in the general population being zero.⁴⁰

It is not difficult to imagine how this phenomenon can lead the casual observer to hold fallacious beliefs. For example, a student (or professor) who spends enough time in a university environment may come to incorrectly believe that those attributes that are associated with good grades (e.g. diligence), are in general inversely related to those attributes associated with good SAT scores (e.g. brilliance). This mistake is analogous to a gambler holding the belief that streaks are more likely to end rather continue, because in his personal experience this is, in fact, representative of a typical night at the casino (see Streak Reversal Paradox [Problem 7] and Gambler's Verity [Problem 9]).⁴¹

Casual encounters with biased experiential samples is not the only way that people can develop, or maintain, fallacious beliefs. To illustrate, it was recently proposed that empirical work documenting an (internally valid) negative correlation between conscientiousness and cognitive ability may instead merely be reporting a statistical artifact that is driven by a selection bias identical to Berkson's Paradox (Murray, Johnson, McGue, and Iacono 2014). Further, it is not difficult to imagine that a similar bias may be present in experiments in which performance on behavioral tasks that involve cognitive ability is correlated with a personality measure such as conscientious-

⁴⁰The sign of the correlation in the admitted population is negative due to the fact that $\mathbb{P}(\text{good SATs}|\text{good grades, admitted}) - \mathbb{P}(\text{good SATs}|\text{admitted}) = 1/2 - 2/3 < 0$, which is directly proportional to the correlation. On the other hand, if one were to consider an analogous problem in which there are more than two attributes, and admission requires multiple positive (coin flip) attributes, although $\mathbb{P}(\text{good SATs}|\text{good grades, admitted}) > 1/2$ the correlation between attributes in the admitted population will remain negative. This can be shown with a restricted choice argument analogous to the one presented here (see Miller and Sanjurjo [2017].)

⁴¹See Miller and Sanjurjo (2016) for a discussion on the connection between the Streak Reversal Paradox and the Gambler's Fallacy.

ness. Finally, because experimental subjects may be further selected on attributes such as budget constraints, and/or intellectual curiosity, one can similarly imagine the discovery of appealing new correlations, which are nevertheless spurious—e.g. a negative correlation between measures of intellectual curiosity and greedy/selfish behavior in experimental tasks.⁴²

5.4 A Hypothetical Case: Gambler’s Verity and Psi research

The probability of winning is inversely proportional to the amount of the wager.

(Las Vegas proverb)

We show how the same bias that underlies the Alternation Paradox from Section 4 can be used to generate a novel puzzle in which a rule for predicting i.i.d. outcomes can appear to outperform chance rates. In particular, this can happen if a researcher is unaware of the implicit selection bias that the rule generates.

Problem 9 (Gambler’s Verity) *Imagine a roulette wheel in which half of the slots are red, and half are black (for simplicity). Jill will observe exactly 3 spins of the wheel, and has committed to the following betting strategy: whenever observing a red (R) bet black (B) on the next spin; otherwise, do not bet. Do you expect Jill to win as many bets as she loses?*

Because Jill’s betting strategy will restrict her to betting on the second flip, the third, or both, there are three possible outcomes: she will win on none of her bets, half of them, or all of them. While intuition suggests that she is expected to win on half of her bets, this is incorrect, as it overlooks the fact that the three outcomes are not equally likely. To see this, we can simply enumerate the sample space, as follows: if the sequence is BBB or BBR Jill will not bet; otherwise, for the remaining six equally likely sequences, she will bet. Now, given that Jill bets, she has a $3/6 = 1/2$ probability of winning all of her bets (RBR,RBB,BRB), a $1/6$ probability of winning half of them (RRB), and a $2/6 = 1/3$ probability of losing all of them (BRR,RRR). As a result, Jill is expected to win more bets than she loses, with an expected win rate of $(\frac{1}{2} \times 1) + (\frac{1}{6} \times \frac{1}{2}) + (\frac{1}{3} \times 0) = 0.58$. In fact, her high success rate immediately follows from the solution to the Alternation Paradox, which we solved using the same restricted choice thinking as in our solution to the Monty Hall problem (Problem 4). That is, Jill’s expected win rate is equivalent to the statement that

⁴²See Miller and Sanjurjo (2017) for more on Berkson’s Paradox.

for a randomly selected flip that is immediately preceded by a heads, the probability of a tails (an alternation) is $1 - 5/12 = 0.58$.

Finally, given that Jill's probability of walking away ahead is higher than her probability of walking away behind, it would appear that she has discovered a strategy with which she can expect to win money. However, this is not true. The reason why accords with the Las Vegas proverb above: her probability of walking out ahead is inversely proportional to the total amount wagered (the number of bets). In particular, relative to the high probability sequences in which she walks away ahead, in the low probability sequences in which she walks away behind she wagers 50 percent more and her absolute (negative) profit is 50 percent greater.⁴³ As a result, given fair odds, Jill is expected to break even.

A misguided (hypothetical) ESP experiment

It is straightforward to illustrate how The Gambler's Verity problem could have implications for social science research. In particular, consider the hypothetical case of the Amazing Geller, an ESP master who claims to have a scientifically validated method to train people in precognition. In order to validate his method, he devises a test to prove that his students can successfully predict the outcomes of coin flips.

While Geller admits that he cannot train his students to predict every coin flip perfectly, he asserts that his method renders some coin flips predictable. For the test, Geller submits his students to a team of objective third party researchers. The experiment is simple: for each student, a researcher will flip a coin 100 times, and the student will predict only on flips in which the student senses the ensuing outcome. According to Geller, not all of his trainees have the ability to learn how to predict, so he requests that the researchers merely count how many of his students predict at better than chance rates.

The researchers follow his instructions, and find that of the 1000 students tested, 490 students predict at a rate better than chance, 395 at a rate worse than chance, and 115 students predict at the rate of chance. Thus, the odds are found to be substantially in favor of a student predicting

⁴³Specifically, conditional on walking away ahead, the sequences RBR, RBB and BRB are equally likely, and in each sequence Jill wagers once and wins once. On the other hand, conditional on walking away behind, while the sequences BRR and RRR are also equally likely, for the sequence BRR Jill wagers once and loses, but for the sequence RRR she wagers twice and loses twice. As a result, when Jill walks away behind (1/3 probability) she is expected to wager 1.5 times with a net payoff of -1.5 , whereas when she walks away ahead (1/2 probability) she is expected to wager 1 time with a net payoff of 1.

at better-than-chance rates, relative to worse-than-chance rates. Further, the average student is observed to have a 54 percent success rate on his/her predictions.

Upon finding these results to be highly significant, the researchers are mystified, and left with no choice but to conclude that Geller must indeed have amazing abilities.

What the researchers fail to realize is that these results can easily occur in the absence of precognition. In fact, this outcome is close to what would be *expected* if Geller were to have instructed his students to simply predict a tails whenever the previous three flips are heads. That this is true immediately follows from the fact that Geller's students are predicting a streak reversal (tails) in a setting that is quantitatively identical to the Streak Reversal Paradox.

6 Conclusion

The principle of restricted choice is far more than an ad-hoc intuition, only useful for understanding inference in contract bridge and the Monty Hall problem. Because it is naturally quantified as the updating factor in the odds form of Bayes' rule, it provides a unified, intuitive, and simple approach to solving all of the classic conditional probability puzzles. Moreover, restricted choice thinking can be used to identify novel biases in important empirical settings, so can potentially help future researchers avoid making certain intuitively appealing, but critical, errors when designing experiments, analyzing data, and interpreting results.

References

- ARKES, J. (2011): “Do Gamblers Correctly Price Momentum in NBA Betting Markets?” *Journal of Prediction Markets*, 5, 31–50.
- AVERY, C. AND J. CHEVALIER (1999): “Identifying Investor Sentiment from Price Paths: The Case of Football Betting,” *Journal of Business*, 72, 493–521.
- AVUGOS, S., M. BAR-ELI, I. RITOV, AND E. SHER (2013): “The elusive reality of efficacy performance cycles in basketball shooting: analysis of players’ performance under invariant conditions,” *International Journal of Sport and Exercise Psychology*, 11, 184–202.
- BAR-HILLEL, M. AND R. FALK (1982): “Some teasers concerning conditional probability,” *Cognition*, 11, 109–122.
- BARBERIS, N. AND R. THALER (2003): “A survey of behavioral finance,” *Handbook of the Economics of Finance*, 1, 1053–1128.
- BECKER, A. (2017): “The Monty Hall problem,” *Mathematics Stack Exchange*, Retrieved on 2017-09-14 from: <https://math.stackexchange.com/q/96832>(version: 2012-01-06).
- BERKSON, J. (1946): “Limitations of the application of fourfold table analysis to hospital data,” *Biometrics Bulletin*, 47–53.
- BERTRAND, J. (1889): *Calcul des probabilités [Calculus of probabilities]*, Paris: Gautier-Villars et Fils.
- BOREL, E. AND A. CHÉRON (1940): *Théorie mathématique du bridge à la portée de tous*, Paris: Gauthier-Villars.
- BROWN, W. A. AND R. D. SAUER (1993): “Does the Basketball Market Believe in the Hot Hand? Comment,” *American Economic Review*, 83, 1377–1386.
- CAMERER, C. F. (1989): “Does the Basketball Market Believe in the ‘Hot Hand,’?” *American Economic Review*, 79, 1257–1261.
- CHAMLEY, C. P. (2004): *Rational Herds: Economic Models of Social Learning*, Cambridge: Cambridge University Press.
- CHEN, K. (2008): “Rationalization and Cognitive Dissonance: Do Choices Affect or Reflect Preferences?” Cowles Foundation Discussion Paper No. 1669.
- CHEN, M. K. AND J. L. RISEN (2010): “How Choice Affects and Reflects Preferences: Revisiting the Free-Choice Paradigm,” *Journal of Personality and Social Psychology*, 99, 573–594.
- CHEN, W. AND J. WANG (2010): “Epiphany Learning for Bayesian Updating: Overcoming the Generalized Monty Hall Problem,” *Working Paper*, December 7th.
- CROSON, R. AND J. SUNDALI (2005): “The Gamblers Fallacy and the Hot Hand: Empirical Data from Casinos,” *Journal of Risk and Uncertainty*, 30, 195–209.

- DE BONDT, W. P. (1993): “Betting on trends: Intuitive forecasts of financial risk and return,” *International Journal of Forecasting*, 9, 355–371.
- DE LONG, J. B., A. SHLEIFER, L. H. SUMMERS, AND R. J. WALDMANN (1991): “The Survival of Noise Traders In Financial-markets,” *Journal of Business*, 64, 1–19.
- DURHAM, G. R., M. G. HERTZEL, AND J. S. MARTIN (2005): “The Market Impact of Trends and Sequences in Performance: New Evidence,” *Journal of Finance*, 60, 2551–2569.
- EGAN, L. C., L. R. SANTOS, AND P. BLOOM (2007): “The origins of cognitive dissonance evidence from children and monkeys,” *Psychological Science*, 18, 978–983.
- FEHR, E. AND J. TYRAN (2005): “Individual Irrationality and Aggregate Outcomes,” *Journal of Economic Perspectives*, 19, 43–66.
- FOX, C. R. AND J. LEVAV (2004): “Partition-edit-count: naive extensional reasoning in judgment of conditional probability,” *Journal of Experimental Psychology: General*, 133, 626–642.
- FRIEDMAN, D. (1998): “Monty Hall’s Three Doors: Construction and Deconstruction of a Choice Anomaly,” *American Economic Review*, 88, 933–946.
- GALBO-JØRGENSEN, C. B., S. SUETENS, AND J.-R. TYRAN (2016): “Predicting Lotto Numbers A natural experiment on the gamblers fallacy and the hot hand fallacy,” *Journal of the European Economic Association*, 14, working Paper.
- GALTON, F. (1894): “A plausible paradox in chances,” *Nature Neuroscience*, 49, 365–366.
- GARDNER, M. (1959): “Mathematical Games: Problems involving questions of probability and ambiguity,” *Scientific American*, 201, pp. 174–182.
- (1961): *The 2nd Scientific American book of mathematical puzzles and diversions.*, New York: Simon & Schuster.
- GIGERENZER, G. AND U. HOFFRAGE (1995): “How to improve Bayesian reasoning without instruction: Frequency formats,” *Psychological Review*, 102, 684–704.
- GILLMAN, L. (1992): “The Car and the Goats,” *The American Mathematical Monthly*, 99, 3–7.
- GILOVICH, T., R. VALLONE, AND A. TVERSKY (1985): “The Hot Hand in Basketball: On the Misperception of Random Sequences,” *Cognitive Psychology*, 17, 295–314.
- GORROOCHURN, P. (2012): *Classic problems of probability*, New Jersey: John Wiley & Sons.
- GRANBERG, D. AND T. A. BROWN (1995): “The Monty Hall dilemma,” *Personality and Social Psychology Bulletin*, –.
- GURYAN, J. AND M. S. KEARNEY (2008): “Gambling at Lucky Stores: Empirical Evidence from State Lottery Sales,” *American Economic Review*, 98, 458–473.
- JAYNES, E. T. (2003): *Probability theory: the logic of science*, Cambridge: Cambridge University Press.

- KAHNEMAN, D. (2011): *Thinking, Fast and Slow*, Farrar, Straus and Giroux.
- KAHNEMAN, D. AND M. W. RIEPE (1998): “Aspects of Investor Psychology: Beliefs, preferences, and biases investment advisors should know about,” *Journal of Portfolio Management*, 24, 1–21.
- KAHNEMAN, D. AND A. TVERSKY (1972): “On Prediction and Judgement,” *ORI Research Monograph*, 12, (10)(a).
- KLUGER, B. AND D. FRIEDMAN (2010): “Financial engineering and rationality: Experimental evidence based on the Monty Hall problem,” *Journal of Behavioral Finance*, 11, 31–49.
- KLUGER, B. D. AND S. B. WYATT (2004): “Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem,” *Journal of Finance*, 59, 969–998.
- KOEHLER, J. J. AND C. A. CONLEY (2003): “The “hot hand” myth in professional basketball,” *Journal of Sport and Exercise Psychology*, 25, 253–259.
- KRAUSS, S. AND X.-T. WANG (2003): “The psychology of the Monty Hall problem: discovering psychological mechanisms for solving a tenacious brain teaser.” *Journal of Experimental Psychology: General*, 132, 3–.
- LEE, M. AND G. SMITH (2002): “Regression to the mean and football wagers,” *Journal of Behavioral Decision Making*, 15, 329–342.
- LOH, R. K. AND M. WARACHKA (2012): “Streaks in Earnings Surprises and the Cross-Section of Stock Returns,” *Management Science*, 58, 1305–1321.
- MALKIEL, B. G. (2011): *A random walk down Wall Street: the time-tested strategy for successful investing*, New York: W. W. Norton & Company.
- MILLER, J. B. AND A. SANJURJO (2015): “Surprised by the Gambler’s and Hot Hand Fallacies? A Truth in the Law of Small Numbers,” Working Paper, September 24.
- (2016): “Selection Bias, Sample-size Neglect, and the Persistence of Gambler’s Fallacy Beliefs,” *Working Paper*.
- (2017): “Berkon’s Bias and the Hot Hand bias: A Quantitative and Conceptual Comparison,” *Working Paper*.
- MURRAY, A. L., W. JOHNSON, M. MCGUE, AND W. G. IACONO (2014): “How are conscientiousness and cognitive ability related to one another? A re-examination of the intelligence compensation hypothesis,” *Personality and Individual Differences*, 70, 17–22.
- NALEBUFF, B. (1987a): “Puzzles: Choose a curtain, duel-ity, two point conversions, and more,” *Journal of Economic Perspectives*, 157–163.
- (1987b): “Puzzles: Noisy Prisoners, Manhattan Locations, and More,” *Journal of Economic Perspectives*, 1, 185–191.

- NARAYANAN, S. AND P. MANCHANDA (2012): “An empirical analysis of individual level casino gambling behavior,” *Quantitative Marketing and Economics*, 10, 27–62.
- O’HAGAN, A., C. E. BUCK, A. DANESHKHAH, J. R. EISER, P. H. GARTHWAITE, D. J. JENKINSON, J. E. OAKLEY, AND T. RAKOW (2006): *Uncertain judgements: eliciting experts’ probabilities*, West Sussex, England: John Wiley & Sons.
- PAGE, S. E. (1998): “Let’s make a deal,” *Economics Letters*, 61, 175–180.
- PALACIOS-HUERTA, I. (2003): “Learning to open Monty Hall’s doors,” *Experimental Economics*, 6, 235–251.
- PAUL, R. J. AND A. P. WEINBACH (2005): “Bettor Misperceptions in the NBA: The Overbetting of Large Favorites and the ‘Hot Hand,’” *Journal of Sports Economics*, 6, 390–400.
- PEARL, J. (2009a): “Causal inference in statistics: An overview,” *Statistical Surveys*, 3, 96–146.
- (2009b): *Causality: Models, Reasoning and Inference*, New York, NY: Cambridge University Press.
- RABIN, M. AND D. VAYANOS (2010): “The Gamblers and Hot-Hand Fallacies: Theory and Applications,” *Review of Economic Studies*, 77, 730–778.
- REESE, T. (1960): *Master Play in Contract Bridge*, New York: Dover Publications, Inc.
- ROSENHOUSE, J. (2009): *The Monty Hall Problem: The Remarkable Story of Math’s Most Contentious Brainteaser*, New York: Oxford University Press.
- ROSENTHAL, J. S. (2008): “Monty Hall, Monty Fall, Monty Crawl,” *Math Horizons*, 16, 5–7.
- SACKETT, D. L. (1979): “Bias in analytic research,” *Journal of Chronic Diseases*, 32, 51–63.
- SCHELLING, T. C. (1971): “Dynamic models of segregation,” *Journal of Mathematical Sociology*, 1, 143–186.
- SELVIN, S. (1975): “A Problem in Probability (letter to the editor),” *The American Statistician*, 29, 67.
- SINKEY, M. AND T. LOGAN (2013): “Does the Hot Hand Drive the Market?” *Eastern Economic Journal*, Advance online publication, doi:10.1057/ej.2013.33.
- SLEMBECK, T. AND J.-R. TYRAN (2004): “Do institutions promote rationality?: An experimental study of the three-door anomaly,” *Journal of Economic Behavior & Organization*, 54, 337–350.
- SMITH, G., M. LEVERE, AND R. KURTZMAN (2009): “Poker Player Behavior After Big Wins and Big Losses,” *Management Science*, 55, 1547–1555.
- SUNDALI, J. AND R. CROSON (2006): “Biases in casino betting: The hot and the gamblers fallacy,” *Judgement and Decision Making*, 1, 1–12.
- THALER, R. H. AND C. R. SUNSTEIN (2008): *Nudge: Improving Decisions About Health, Wealth, and Happiness*, Yale University Press.

- TOR, A. AND M. H. BAZERMAN (2003): “Focusing failures in competitive environments: Explaining decision errors in the Monty Hall game, the acquiring a company problem, and multiparty ultimatums,” *Journal of Behavioral Decision Making*, 16, 353–374.
- TUBAU, E., D. AGUILAR-LLEYDA, AND E. D. JOHNSON (2015): “Reasoning and choice in the Monty Hall Dilemma (MHD): implications for improving Bayesian reasoning,” *Frontiers in Psychology*, 6.
- TVERSKY, A. AND T. GILOVICH (1989): “The “Hot Hand”: Statistical Reality or Cognitive Illusion?” *Chance*, 2, 31–34.
- VAZSONYI, A. (1999): “Which Door Has the Cadillac?” *Decision Line*, 17–19.
- VOS SAVANT, M. (1990): “Ask Marilyn,” *Parade Magazine*, 15.
- XU, J. AND N. HARVEY (2014): “Carry on winning: The gambler’s fallacy creates hot hand effects in online gambling,” *Cognition*, 131, 173 – 180.
- YUAN, J., G.-Z. SUN, AND R. SIU (2014): “The Lure of Illusory Luck: How Much Are People Willing to Pay for Random Shocks,” *Journal of Economic Behavior & Organization*, forthcoming.

A Appendix: Odds vs. Probability

To see the advantages of representing uncertainty in terms of odds rather than probabilities, first imagine comparing the two posterior probabilities $\mathbb{P}(M_1|c)$ and $\mathbb{P}(M_2|c)$, where the former is given by $\frac{\mathbb{P}(c|M_1)}{\mathbb{P}(c)}\mathbb{P}(M_1)$ and the latter by $\frac{\mathbb{P}(c|M_2)}{\mathbb{P}(c)}\mathbb{P}(M_2)$. Notice that in both cases, to compute the usual updating factor $\frac{\mathbb{P}(c|M_i)}{\mathbb{P}(c)}$, one must either somehow know $\mathbb{P}(c)$, or be able to compute it using the law of total probability over all possible models $i = 1, 2, 3, \dots, n$, i.e. $\mathbb{P}(c) = \sum_{i=1}^n \mathbb{P}(c|M_i)\mathbb{P}(M_i)$.⁴⁴ Of course, in the odds form, the $\mathbb{P}(c)$ term vanishes.⁴⁵

Next, for a concrete illustration of how the odds representation of uncertainty makes Bayesian updating less complex, and reduces avoidable calculation errors, consider the sequential witness version of Kahneman and Tversky (1972)’s taxi problem introduced by Chamley (2004, pp.19–21). In the problem, the prior probability of an encountered cab being red is .1, and the probability of correctly identifying the cab color is .8. First, a witness reports the cab being red, which makes the odds in favor of the red cab rise from 1:9 to 4:9, as the witness is $\mathcal{R} = .8/.2 = 4$ times more restricted to report red when the cab is in fact red (vs. when it is yellow). Thus, this information makes the probability that the cab is red rise from 1/10 to 4/13 $< 1/2$. Then, when a second witness also reports the cab being red the odds in favor of the red cab rise again, by a factor of four, from 4:9 to 16:9. Thus, the posterior probability becomes 16/25, rather than the incorrect value of 16/29 that is reported in Chamley (2004, p.21, footnote 3) after a more extended series of calculations.

Finally, when the number of candidate models is greater than two, but not too numerous, a decision maker who wants to choose the most probable model can achieve this by simply performing multiple binary comparisons, each of which is naturally operationalized by inspecting the odds. A psychological advantage of binary comparisons is that probability often shares features in common with other psychophysical quantities in that the relative, rather than absolute, judgement of numerical probability is a more reliable (and natural) measure of the subjective belief in

⁴⁴This expansion suggests an alternative representation of the updating factor for $\mathbb{P}(M_1)$. In particular $\mathbb{P}(M_1|c) = \frac{1}{E_M[\mathcal{R}_{M:M_1}(c)]}\mathbb{P}(M_1)$ where $E_M[\mathcal{R}_{M:M_1}(c)] = \sum_{i=1}^n \frac{\mathbb{P}(c|M_i)}{\mathbb{P}(c|M_1)}\mathbb{P}(M_i)$, i.e. the prior probability of M_1 , re-scaled by how restricted the average model is to “choose” c (relative to model M_1).

⁴⁵When $\mathbb{P}(c)$ is known, the updating factor for probabilities is intuitive. For example, in the Monty Hall Problem (Problem 4) with the contestants initial pick of door 3, Monty’s ex ante probability of choosing door 1 to open satisfies $\mathbb{P}(c = 1) = 1/2$, whereas conditional on door 2 concealing a car M_2 , Monty’s probability satisfies $\mathbb{P}(c = 1|M_2) = 1$. This implies that the posterior probability $\mathbb{P}(M_2|c = 1)$ must be twice the prior, i.e. the probability doubles from 1/3 to 2/3.

the probability ranking of two (or more) models (O’Hagan, Buck, Daneshkhah, Eiser, Garthwaite, Jenkinson, Oakley, and Rakow 2006).⁴⁶

B Appendix: Modified Monty Hall and the Free Choice Paradigm

We provide a further illustration of the empirical relevance of the Monty Hall problem by using a modified version of it to explain an important error discovered by Chen (2008) in a common experimental design used to test for cognitive dissonance. While restricted choice does not provide additional insight to this error, the error’s similarity in structure to the Monty Hall problem is worthy of discussion.⁴⁷

Problem 10 (Modified Monty Hall: The Car, the Goat, and the Bicycle) *Suppose you’re on a game show, and you’re given the choice of three doors: Behind one door is a car; behind another is a goat; behind another is a bicycle. The host asks which prizes you prefer; you say that you prefer the car over the bicycle, and the bicycle over the goat (and, of course, the car over the goat). You pick a door, say number 3, and the host, who is 100 percent honest, and knows what’s behind the doors, tells you that you prefer door 2 over door 1. He then says to you, “Do you want to pick door number 2?” Is it to your advantage to switch your choice?*

In this version of the Monty Hall problem, the contestant learns from Monty that door number 2 is the better of the two unchosen doors, whereas nothing has been learned about door number 3. Given this, the intuitive answer is to switch—which happens to also be correct. To see why, if door number 1 is worse than 2, then this means that there remain three possible arrangements for the first and second doors, respectively: Goat-car, Bike-car, and Goat-bike. Thus, the host’s choice is equally restricted across the *remaining* arrangements, so the relative odds do not change. As a result, the key to solving the problem is the enumeration of the sample space, and the elimination of impossible arrangements. That is, for the case in which Monty tells the contestant that the

⁴⁶One potential limitation of representing uncertainty with odds is that while 2:1 odds are substantially larger than 1:1 odds, this is arguably not salient in the representation. On the other hand, the difference can be made more salient by scaling up the number of chances for each event, e.g. 20:10 vs. 10:10.

⁴⁷The error that Chen (2008) discovered instead arises from incorrect assumptions about the sample space. In particular, as in the discussion of the classic Boy-or-Girl Paradox, and Leibniz’s error (Footnote 13), the most difficult part of solving the associated Monty Hall problem (Problem 10) are the subtleties involved with enumerating the sample space, as there is no updating of the odds for the remaining possibilities.

inferior prize behind doors 1 and 2 is behind door 1, this immediately eliminates the arrangements: Car-goat, Car-bike, and Bike-goat. As a result, the remaining arrangements (in which Monty must choose door 1) are infinitely more restricted with regards to the choice of door 1, than in the eliminated arrangements (in which Monty will never choose door 1). Finally, given the elimination of these three possibilities, and the unchanged 1:1 relative odds on the remaining three possibilities, there is a $2/3$ chance that the car is behind door number 2.

The Modified Monty Hall problem turns out to be related to a method for measuring cognitive dissonance, called the free choice paradigm. In one version of the paradigm, the researcher starts with three items that s/he has reason to believe the subject is indifferent between; say A, B, and C. Then the researcher selects two of these items at random, and offers them to the subject, who must choose one. Now, suppose that A and B are selected by the researcher, and that the subject chooses A over B. Next, the researcher pairs the *unchosen* alternative B with alternative C. Finally, the subject must choose between B and C. In this type of experiment, what researchers typically find is that C is chosen significantly more often than B. Under the assumption that the experiment was designed perfectly, the result is perhaps surprising. That is, given that subjects began the experiment indifferent between A, B, and C, it would appear that the rejection of B, *at random*, in the first round choice then made B less desirable in the second round choice.

On the other hand, if subjects are not indifferent, but instead have preferences between the items (however slight), and these preferences are revealed by their choices, then the result is not surprising. To see why, when subjects are presented with any two items, then they are restricted to choose the item that they prefer most. This is analogous to Monty revealing that door 2 has the better prize than door 1, above. Now, if Monty were to have instead offered the door with the less preferred prize (door 1), then there would have been a $2/3$ chance that it yielded the goat—the worst option. In this case, staying with door 3 would have been the correct choice. Now, in the case of the experiment, upon observing that the subject prefers A to B, there is similarly a $2/3$ chance that B is the worst option. To see this, notice that for A to be chosen over B, the rank of (A, B), respectively, must be (1, 2), (2, 3), or (1, 3), so there is a $2/3$ chance that the subject ranks B as the least desirable item. As a result, the subject should be expected to choose C over B not $1/2$, but $2/3$ of the time. In sum, by designing an experiment in which the subject's unchosen round one item is offered along with new alternative C in the second round choice, a bias is created in which

C is made more likely to be chosen than B.⁴⁸

⁴⁸This is not the only variant of the free choice paradigm. The one presented here is related to Egan, Santos, and Bloom (2007), as discussed in Chen (2008). Also see Chen and Risen (2010).