- - -

acknowledged. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

However, there is mounting empirical and experimental evidence suggesting that humans do not always behave in the way traditional economic models predict, but often make seemingly random and suboptimal decisions [10]. These behavioral anomalies and psychological traits are especially pronounced when elements of risk and probability are involved. Examples include loss aversion [11±14], overconfidence [15, 16], overreaction [17], herding [18] psychological accounting [19], miscalibration of probabilities [20], the uncertainty effect [21], and confirmation bias [22]. The spectacular rise of US stock market prices during the tech bubble in the early 2000s, and the even more spectacular crash following the 2007±2008 financial crisis, has intensified the controversy surrounding the rationality of investors.

One particularly interesting behavioral anomaly is probability matching, also known as the ^amatching law,^o or Herrnstein's Law [23±30] Dthe tendency of the relative frequency of predictions of outcomes of an independent randomized event to match its underlying probability distribution. The best-known example of probability matching is the human tendency to choose randomly between heads and tails when asked to guess the outcomes of a series of biased coin tosses. When individuals are asked to guess the repeated outcomes of a biased coin, say with a bias of 70% heads, and rewarded based on whether they guessed correctly, most subjects seem to randomize their guesses at around 70% heads, instead of engaging in the economically optimal behavior of always guessing heads.

Probability matching has long puzzled economists and psychologists because of its apparent inconsistency with basic self-interest. The idea of randomizing behavioris especially difficult to reconcile with the standard economic paradigm of expected utility theory, in which individual behavior is non-stochastic and completely determined by the individual's utility function, budget constraints, and the probability laws governing the environment. For example, Kogler and Kuhberger [31] report that, ^aExperimental research in simple repeated risky choices shows a striking violation of rational choice theory: the tendency to match probabilities by allocating the frequency of response in proportion to their relative probabilities^o.

Nevertheless, probability matching has been observed in thousands of geographically diverse human subjects over several decades, as well as in other animal species, including ants $[32\pm35]$, bees $[36\pm38]$, fish [39, 40], pigeons [41, 42], and primates [43]. In virtually any setting where an animal is able to make a choice between A versus B in a randomized experiment, we observe probability matching.

The source of these irrational behaviors is often attributed to psychological factors, such as fear, greed, and other emotional responses. However, the fact that some of these behaviors are observed so consistently across species suggests that they may have a more fundamental and common origin, one with an evolutionary role that belies their apparent shortcomings. For example, the neurological basis of probability matching has been investigated extensively [44± 49]. In the context of a binary choice model, Brennan and Lo [50] show that probability matching behavior is perfectly consistent with evolution, arising purely from the forces of natural selection and population growth. Moreover, under generalized environmental conditions, i.e., broad assumptions about the conditions required for reproductive success, they derive more general types of behavior that involve randomization, but not necessarily strict probability matching.

In this paper, we present the first experimental test of the evolutionary model of Brennan and Lo [50]. We design an experiment in real-world decision making with monetary payoffs to measure the degree of probability matching among individuals, its determining factors, and its level of variation. Here by probability matching we mean the ^amatching law,^o or Herrnstein's Law discussed above D the tendency to choose randomly between heads and tails when asked to guess the outcomes of a series of biased-coin tosses, where the randomization frequency matches the probability of the biased coin.

We recruited a sample of 82 volunteers from the MIT Behavioral Research Laboratory to participate in our experiment. Each participant played a computer game consisting of 200 trials of a binary choice decision. In each trial, an image of either Angelina Jolie or Brad Pitt was displayed with a certain probability, and subjects were paid according to the number of trials in which they correctly guessed which image appeared.

By varying the payoff structure of the game, we were able to test whether subjects showed probability matching behavior, and whether deviations occurred as predicted by the model in Brennan and Lo [50]. Specifically, we designed several payoffs where the evolutionarily dominant behavior was either to maximize, i.e., always to choose one option, or to randomize, i.e., to choose randomly between two options. We found strong evidence for a behavioral difference between theoretical maximizers and theoretical randomizers, as predicted by Brennan and Lo [50]. After controlling for a wide range of demographic and socioeconomic variables, theoretical randomizers still engaged in randomizing behavior more often than theoretical maximizers. When facing different environments (i.e., payoffs in the experiment), our subjects responded differently by adapting to the new conditions and showing different stable behaviors.

We were also able to study individual differences in the tendency to maximize or randomize by collecting basic demographic and socioeconomic information from the anonymous participants. We found that subjects with a higher level of financial assets tended to randomize less often, while subjects with children tend to randomize more often. Moreover, subjects who had taken probability and statistics classes and those who self-reported finding a pattern in the game (none existed) also tended to randomize more often, contrary to our prior expectation that those participants with a better understanding of probability might be more likely to adopt the economically maximizing behavior. In fact, we found that those subjects engage in the exact opposite behavior. This may be due to an attempt to ^abeat the game,^o based on the qualitative answers to our post-trial survey by participants.

From the evolutionary perspective, the key to understanding these behavioral predictions lies in the assumption of systematic reproductive risk [50, 51]. The experiment we describe in this article involves a binary choice in which the risks to the population are idiosyncratic, that is, the outcomes of one individual's choice are independent of those of another. However, when individuals with preferences formed in response to systematic risks are placed in the different environment, there is the potential for probability matching to occur, creating what appears to be irrational behaviors for those environments.

Our results contribute to the growing literature on rationalizing the existence of probability matching. As far back as the 1950s, researchers [52±54] developed statistical models that attempted to explain and predict matching behavior. Since then, several behavioral reasons have been offered, including its emergence as a consequence of pattern searching [55], through the greater utility gained from guessing the rarer event correctly [56], and by the role of diversification to avoid boredom [57]. More recently, explanations of probability matching have been proposed from an evolutionary point of view. Wolford, Miller, and Gazzaniga [45] argue that early human beings look for explanatory causal relationships as a survival strategy. Wozny, Beierholm, and Shams [58] have shown that humans match probabilities not only in cognitive tasks, but also in perceptual tasks. This implies that the human nervous system has a built-in function that samples from a distribution of hypotheses, and updates its belief after each observation.

Our results provide experimental validation for the predictions of Brennan and Lo [50], as well as additional evidence that individuals engage in randomized behavior and probability matching, even those with prior experience in probability and investing. More importantly, our results may provide an explanation for several notable departures from exact probability

matching [31, 59, 60]. Randomizing behavior that matches environmental probabilities depends on the relative reproductive success of the outcomes, and the evolutionary framework proposed in Brennan and Lo [50] offers a simple and specific set of conditions for understanding and predicting such behavior.

2 Materials and methods

2.1 Evolutionary origins of probability matching

Brennan and Lo [50] proposed an evolutionary framework for the origin of several behaviors that are considered anomalous^o in economic theories based on the assumption of rational behavior. In particular, probability matching D the tendency of the relative frequency of guesses of the outcomes of a sequence of independent random events to match the underlying probability distribution of events D can be explained when the uncertainty in environment is systematic across all individuals, an example demonstrating that natural selection is able to yield behaviors that may be individually sub-optimal but are optimal for the population. For expositional convenience, we present a brief review of this framework here, and then turn to our experimental design.

We begin with a population of individuals that live for one period, produce a random number of offspring asexually, and then die. During their lives, individuals make only one decision: they choose one of two possible courses of action, denoted a and b, and this choice results in one of two corresponding random numbers of offspring, x_a and x_b , given by:

$$Prob(x_{a} = c_{a1}; x_{b} = c_{b1}) = p \in [0; 1]$$

$$Prob(x_{a} = c_{a2}; x_{b} = c_{b2}) = 1 - p \equiv q$$
(1)

where *p* is some probability between 0 and 1.

We further assume that x_a and x_b are independently and identically distributed over time, and identical for all individuals in a given generation. In other words, if two individuals choose the same action a, both will produce the same number of random offspring x_a . This implies that the variation in offspring due to behavior is wholly systematic, i.e., the link between action and reproductive success is the same throughout the population.

A ^amindless^o individual's behavior in this world is fully specified by the probability of choosing action a. Following the notation in Brennan and Lo [50], we denote this probability as f. Each individual dies after one period, and we assume its behavior f is heritable: offspring will behave in a manner identical to their parents, i.e., they choose between the two actions according to the same probability f.

From the individual's perspective, always choosing the action with a higher expected reproductive success (f = 0 or 1) will lead to more offspring on average. However, Brennan and Lo [50] showed that from the perspective of the population, this individually optimal behavior cannot survive. In fact, the evolutionarily dominant behavior will depend on the relationship between the probability p and the relative fecundity ratios $r_j = c_{aj}/c_{bj}$ for each of the two possible states of the world, j = 1, 2, where f can be anywhere between 0 and 1 in general, implying randomized behavior. See Proposition 3 of Brennan and Lo [50] for more detail.

Fig 1 illustrates the evolutionarily dominant behavior f^* as a function of r_1 and r_2 . If r_1 and r_2 are not too different in value \ominus i.e., the ratio of fecundity between choices a and b is not very different between the two states of the world \ominus then random behavior yields no evolutionary advantage over deterministic choice. In this case, the individually optimal behavior ($f^* = 0$ or 1) will prevail in the population.



Fig 1. Regions of the $(\mathbf{r}_1, \mathbf{r}_2)$ -plane that imply deterministic $(\mathbf{f}^* = 0 \text{ or } 1)$ or randomizing $(0 < \mathbf{f}^* < 1)$ behavior, where $\mathbf{r}_j = \mathbf{c}_{aj}/\mathbf{c}_{bj}$ measures the relative fecundities of action \mathbf{a} to action \mathbf{b} in the two states $\mathbf{j} = 1, 2$. The asymptotes of the curved boundary line occur at $\mathbf{r}_1 = \mathbf{p}$ and $\mathbf{r}_2 = \mathbf{q}$. Values of \mathbf{r}_1 and \mathbf{r}_2 for which exact probability matching is optimal are given by the solid black curve. Source: Brennan and Lo [50, Fig 1].

https://doi.org/10.1371/journal.pone.0252540.g001

However, if one of the *r* variables is large while the other is small, then random behavior will be more advantageous for the population than a deterministic one. In such cases, there are times in which each choice performs substantially better than the other. Under those conditions, it is evolutionarily optimal for a population to diversify between the two choices, rather than always choosing the outcome with the highest probability of progeny in a single generation.

A simple numerical example from Brennan and Lo [50] will illustrate the basic mechanism of this model. Consider a population of individuals, each facing a binary choice between one of two possible actions, **a** and **b**. 70% of the time, environmental conditions are positive, and action **a** leads to reproductive success, generating 3 offspring for the individual. 30% of the time, environmental conditions are negative, and action **a** leads to 0 offspring. This corresponds to p = 70%, $c_{a1} = 3$, $c_{b1} = 0$ in the notation of (1). Suppose action **b** has exactly the opposite outcomesDwhenever **a** yields 3 offspring, **b** yields 0, and whenever **a** yields 0, **b** yields 3. This corresponds to $c_{a2} = 0$, $c_{b2} = 3$ in the notation of (1). From the individual's perspective, always choosing **a**, which has the higher probability of reproductive success, will lead to more offspring on average. However, if all individuals in the population behaved in this ^arational^o manner, the first time that a negative environmental condition occurs, the entire population will become extinct. Assuming that offspring behave identically to their parents, the behavior ^aalways choose **a**^o cannot survive over time. For the same reason, ^aalways choose **b**^o is also unsustainable. In fact, one can show that in this special case, the behavior with the highest reproductive success over time is for each individual to choose **a** 70% of the time and **b** 30% of the time, matching the probabilities of reproductive success and failure. Eventually, this particular randomizing behavior will dominate the entire population.

The key to understanding these behavioral predictions lies in the assumption of *systematic* reproductive risk. This dependence on risk has implications that go far beyond the current setting. For example, Zhang, Brennan, and Lo [51] show that environments with a mix of systematic and idiosyncratic reproductive risks cause different risk preferences to emerge. While our risk preferences may be determined by the nature of the risks to which we and our evolutionary ancestors have been exposed, we do not necessarily have the ability to distinguish between systematic and idiosyncratic risks in our day-to-day decision making.

2.2 The binary choice game

Turning to our experimental design, we presented live human subjects with a binary choice game in which the risks to the population are idiosyncratic, that is, the outcomes of one individual's game are independent of those of another. However, when individuals apply preferences formed in response to systematic risks to the wrong environment, there is the potential for probability matching to occur, creating what appears to be irrational behaviors for those environments.

In our experiment consisted of four particular payoff structures by varying the parameters in (1) (equivalently, four particular points in Fig 1), and observe whether participants show behaviors predicted by Brennan and Lo [50] (equivalently, behaviors indicated by different colors in Fig 1).

We recruited a sample of 82 volunteers and conducted our binary choice experiment at the MIT Behavioral Research Laboratory. Our subjects were varied in their personal and socioeconomic characteristics. We provide a summary of their statistics in Section 3.1.

The full experimental session typically lasted 45 to 60 minutes for a given participant. Each participant used a computer program that completed 200 iterations of a binary choice trial, in essence playing a lottery. On each iteration of the trial, subjects were shown an image of one of two popular film starsĐAngelina Jolie or Brad PittĐwith specific fixed probabilities that were unknown to the subjects. Each participant was randomly assigned to one of four experimental designs, as shown in <u>Table 1</u>. In designs 1 and 2, Angelina Jolie appeared 70% of the time and Brad Pitt 30% of the time. Designs 3 and 4 used the opposite probabilities. The participant guessed which image would appear before it was revealed, and the participant would receive a certain amount of virtual dollars if their guess was correct. Fig 2 shows a screenshot of the computer interface used in the experiment.

In designs 1 and 4, subjects received two virtual dollars when they guessed correctly, and zero virtual dollars when they guessed incorrectly. In designs 2 and 3, subjects received two virtual dollars when they guessed correctly and one virtual dollar when they guessed

Design	Image Probability	Payoff	Utility Maximizing Behavior	Evolutionarily Dominant Behavior [50]
1	$\mathbb{P}(\text{Angelina}) = 0.7$ $\mathbb{P}(\text{Brad}) = 0.3$	Correct: v\$2 Incorrect: v\$0	Always Guess Angelina	$f^* = \mathbb{P}(\text{Guess Angelina}) = 0.7$
2	$\mathbb{P}(\text{Angelina}) = 0:7$ $\mathbb{P}(\text{Brad}) = 0:3$	Correct: v\$2 Incorrect: v\$1	Always Guess Angelina	$f^* = \mathbb{P}(\text{Guess Angelina}) = 1:0$
3	$\mathbb{P}(\text{Brad}) = 0.7$ $\mathbb{P}(\text{Angelina}) = 0.3$	Correct: v\$2 Incorrect: v\$1	Always Guess Brad	$f^* = \mathbb{P}(\text{Guess Brad}) = 1:0$
4	$\mathbb{P}(\text{Brad}) = 0.7$ $\mathbb{P}(\text{Angelina}) = 0.3$	Correct: v\$2 Incorrect: v\$0	Always Guess Brad	$f^* = \mathbb{P}(\text{Guess Brad}) = 0.7$

Table 1. Experimental design.

https://doi.org/10.1371/journal.pone.0252540.t001

incorrectly. These designs correspond to four different evolutionarily dominant behaviors in Fig 1 (see also Brennan and Lo [50]), as shown in the last column of Table 1. Designs 1 and 4 are meant to yield randomized behavior according to theory, while designs 2 and 3 are meant to yield deterministic behavior. In terms of parameters in Fig 1 using p = 0.7, Design 1 corresponds to $r_1 = \infty$ and $r_2 = 0$, which yields the dominant behavior $f^* = 0.7$; Design 2 corresponds to $r_1 = 2$ and $r_2 = \frac{1}{2}$, which yields the dominant behavior $f^* = 1$; Design 3 corresponds to $r_1 = \frac{1}{2}$ and $r_2 = 2$, which yields the dominant behavior $f^* = 0$; Design 4 corresponds to $r_1 = 0$ and $r_2 = \infty$, which yields the dominant behavior $f^* = 0.3$.

Fig 3 shows the trial-by-trial outcome of two representative subjects. The subject in Fig 3a



Fig 3. Experimental outcomes for a representative randomizer (3a) and a representative maximizer (3b). The highest row of triangles displays the randomly generated appearances of Angelina Jolie for each trial. The second row of triangles displays the instances when the subject's response was Angelina Jolie. The bottom two rows of triangles represent the same information for Brad Pitt appearances and Brad Pitt responses. The middle row of red ticks represents trials that the subject guessed correctly. The diagonal line shows the cumulative payout to the subject over time.

https://doi.org/10.1371/journal.pone.0252540.g003

addition, 64% of our subjects have reported taking some probability and statistics classes, an unsurprising finding, given that the experiment took place at MIT.

Subjects each received \$5 in base pay for showing up, and \$0.05 for each virtual dollar they earned. Total dollar earnings ranged from \$14.80 to \$22.20. Table 2 also reports the total

Table 2. Participant	demographics and	summary statistics
----------------------	------------------	--------------------

Variable	Distribution ($n = 75$ subjects)					
		Personal Characteristics				
Gender	Male 53.3%	Female 45.3%				
Has Children	Yes 12.0%	No 86.7%				
Age	≤23 33.3%	[24, 46] 34.7%	>46 30.7%			
Marital Status	Single 77.3%	Partnered 9.3%	Married 8.0%	Other 4.0%		
	Soc	ioeconomic Characteristics				
Taken Probability & Statistics Class	Yes 64.0%	No 34.7%				
Gambling Experience	Yes 33.3%	No 65.3%				
Housing Status	Rent 88.0%	Own 10.7%				
Working Status	Student 40.0%	Currently Working 38.7%	Unemp			
		·	·		·	

number of correct guesses for all subjects. The best performer guessed 154 (77%) trials correctly, while the worst performer only guessed 98 (49%) trials correctly. The median subject guessed 129 (64.5%) out of the 200 trials correctly, slightly less than the expected number of correct guesses for a perfect maximizer, who would always guess the dominant image.

The post-game survey also asked participants about their perceptions of the binary choice game. 44% of our subjects reported that they found a pattern in the game. It is clear that many participants were looking for patterns throughout the game, despite its completely random nature. This is consistent with the ^arepresentativeness heuristic^o first documented by Tversky and Kahneman [61, 62]. We include quotes from two representative subjects.

^aI kept losing count, but clearly the ratio of appearance of Jolie's picture to Pitts's kept going up until it was something 7:1, then it went down (not always in increments of one, I think) until it was 1:1, and then it went back up again.^o

^a70% Angelina. If we picked her too many times, Brad was introduced as a counter-pick.^o

In addition, 74.7% of the subjects reported that they had a specific strategy in the game. By reading the post-study surveys, we realized that our subjects exhibited a wide range of heterogeneous strategies for the game. Here we show a few representative quotes from the two extremes of these strategies, where some subjects indicated clearly that they were always choosing one image:

^aAlways pick Angie.^o

^aChoosing Brad Pitt all the time. His image appeared more frequently and even if the probability was 50% it would not have mattered who I choose, so why not choose him all the time. Also minimizes thinking effort and time to click.^o

Other subjects seemed to engage in more complicated strategies:

^aChose Brad Pitt the majority of the timeÐif Brad Pitt appeared at least 6 times in a row, chose Angelina Jolie.^o

^a`...I was switching between one and another until I noticed some sort of pattern and then I favored Angelina Jolie's picture for the higher number and Brad Pitt for the lower number in the pattern of 5-1-3-1-2.°

These self-reported strategies are also reflected in the wide heterogeneity in behavior when we analyze participant choices.

3.2 A model for individual behavior

Brennan and Lo [50] predict that subjects assigned to designs 1 and 4 of our binary choice game will randomize their behavior. We refer to them as ^atheoretical randomizers.^o On the other hand, subjects assigned to designs 2 and 3 are predicted to choose the dominant image deterministically, and we refer to them as ^atheoretical maximizers^o (see <u>Table 1</u>). In this section, we study whether theoretical randomizers indeed randomize more often than theoretical maximizers.

We first describe a simple model of individual behavior. Define *D* to be the dominant option in the game. In our experiment, *D* represents Angelina Jolie in designs 1 and 2, and Brad Pitt in designs 3 and 4.

Each individual *i* chooses the dominant option D with probability *f*, where *f* represents the individual's (unobserved) behavior. In other words, the individual's decision in each trial is generated by a Bernoulli random variable:

$$y_t = \begin{cases} 1; \text{ with probability } f; \\ 0; \text{ with probability } 1 - f; \end{cases}$$
(2)

where $y_t = 1$ represents choosing the dominant option D, and t = 1, ..., 200. Suppose in T trials, an individual chooses the dominant option

$$N = \sum_{t=1}^{T} y_t \tag{3}$$

times. From observed data T and N, our goal is to estimate and understand the factors which

determine individual behavior f in different payoff structures. The sample average proportion

$$\hat{f} = N = T$$
 (4)

is the obvious choice as the point estimate of behavior *f*.

If an individual's decisions are independent over time, it follows from (2) and (3) that $N \sim Binomial(T, f)$, and \hat{f} is approximately normally distributed with mean f and variance f(1 - f)/T. More generally, if an individuals' decisions are not independent over time, \hat{f} still has mean f, but its variance may be different. In Section 3.4, we estimate whether individual decisions are independent, and in Section 3.5 we discuss its implications for the variance of \hat{f} and the hypothesis tests we carry out.

3.3 Initial learning

During the experiment, subjects required a number of trials to estimate the frequency of each image. This means that their first few guesses tended to show unstable behavior. To account for this, we divided each individual's total number of trials into eight consecutive segments of 25 trials each, and estimated the aggregate behavior *f* for each segment across individuals within the same trial design. Individual behavior was too noisy for successful functional estimates over the initial trials, so we used the aggregate pattern across individuals to better understand the speed of participant learning.

Fig.4 shows the estimated aggregate behavior for theoretical maximizers (designs 2 and 3, $f^* = 0.7$) and theoretical randomizers (designs 1 and 4, $f^* = 1.0$), segmented into eight consecutive batches. We used the sample average proportions \hat{f} in (4) as the point estimate of behavior f, and the normal approximation for binomial distributions to estimate its confidence interval: $\hat{f} \sim N(f; f(1 - f) = T)$. More specifically, for a given confidence level $1 - \alpha$ (e.g., $\alpha = 0.05$, or 95% confidence), the $(1 - \alpha)$ -confidence interval is given by:

$$\left(\hat{f} - z\sqrt{\frac{\hat{f}(1-\hat{f})}{T}}; \hat{f} + z\sqrt{\frac{\hat{f}(1-\hat{f})}{T}}\right)$$

where *z* is the $1 - \frac{a}{2}$ quantile of a standard normal distribution corresponding to the target error rate α . For a 95% confidence level, the error $\alpha = 1 - 0.95 = 0.05$, so $1 - \frac{a}{2} = 0.975$ and *z* = 1.96. There are other approximations for confidence intervals of binomial random



Fig 4. Estimated aggregate behavior for batches of 25-trial segments.

https://doi.org/10.1371/journal.pone.0252540.g004

reproductive success is the payoff from the game. This is obviously an extreme simplification of reality. Nonetheless, the model still provides the important insight that the presence of probability matching, and the degree at which individuals engage in probability matching, are determined by the environment, which we can specifically test through our experiment.

In particular, we are able to test the hypothesis that:

$$H_0: f_{maximizer} \le f_{randomizer}; \quad H_a: f_{maximizer} > f_{randomizer}$$
(6)

as predicted by Brennan and Lo [50], where $f_{maximizer}$ is the behavior of individuals in designs 2 and 3, and $f_{randomizer}$ is the behavior of individuals in designs 1 and 4. As specified in (2), we observe repeated individual decisions that are, in our model, determined by the unobserved behavior $f_{maximizer}$ and $f_{randomizer}$. We can pool together data from all theoretical maximizers and compare with data from all theoretical randomizers. This is a standard two sample proportion test, except that decisions for the same individual might be correlated, as shown in Section 3.4.

Given a particular individual, we use vector $y (y_1, \dots, y_7)'$ to denote her sequence of T random Bernoulli trials. For simplicity and analytical tractability, we assume the sequence has equicorrelation of (estimated as 10.8% in our dataset). In other words, the covariance matrix of y is given by

$$\operatorname{Cov}(\mathbf{y}) = \operatorname{Var}(y_1) \cdot \operatorname{Corr}(\mathbf{y}) = f(1-f) \cdot \begin{pmatrix} 1 & r & \cdots & r \\ r & 1 & \cdots & r \\ \cdots & & & \\ r & r & \cdots & 1 \end{pmatrix};$$
(7)

where the first equation follows from the fact that y_1, \dots, y_T are identically distributed as specified in (2).

Therefore, the variance of the estimated behavior \hat{f} for individual *i*, as defined in (4), is given by

$$\operatorname{Var}(\hat{f}^{(l)}) = \operatorname{Var}\left(\frac{\sum_{t=1}^{T} y_t}{T}\right) = \frac{\operatorname{Var}(\mathbf{1}'\mathbf{y})}{T^2} = \frac{\mathbf{1}'\operatorname{Cov}(\mathbf{y})\mathbf{1}}{T^2}$$
$$= \frac{f(1-f)T(1+(T-1)r)}{T^2}$$
$$= \frac{f(1-f)}{T} \cdot (1+(T-1)r)$$
(8)

where 1 is the unit vector of all 1's, and we have omitted the superscript (*i*) in our derivation for notational simplicity.

Note that the first term in (8), $\frac{f(1-f)}{T}$, is simply the variance of $\hat{f}^{(l)}$ if individual decisions are independent Bernoulli random variables. Therefore, the second term in (8), (1 + (T - 1)), can be treated as an adjustment factor of $\hat{f}^{(l)}$'s variance when individual decisions are correlated.

For a set of n independent individuals with T trials each, the overall estimated behavior for them is simply the average estimated behavior of each individual. Therefore the variance for their overall behavior is:

$$\operatorname{Var}\left(\hat{f}\right) = \operatorname{Var}\left(\frac{\sum_{i=1}^{n} \hat{f}^{(i)}}{n}\right) = \frac{\operatorname{Var}(\hat{f}^{(1)})}{n} = \frac{f(1-f)}{nT} \cdot (1 + (T-1)r):$$

Foatd28ctabilitv260 Td

As a result, for an (unpaired) two sample proportion test between two groups of subjects with n_1 and n_2 individuals respectively, we have the test statistic:

$$Z = \frac{(\hat{f}_1 - \hat{f}_2) - 0}{\sqrt{f^*(1 - f^*)\left(\frac{1}{n_1 T} + \frac{1}{n_2 T}\right)}} \cdot \frac{1}{\sqrt{1 + (T - 1)r}}$$
(9)

where \hat{f}_1 , \hat{f}_2 , and f are the average behavior for individuals from group 1, group 2, and all pooled together, and they do not depend on y's covariance structure. The first term in (9) is simply the standard *z*-score for the two sample proportion test, and the second term in (9) can be treated as the adjustment factor for correlation, which in our case is:

$$\frac{1}{\sqrt{1 + (T - 1)r}} = \frac{1}{\sqrt{1 + (150 - 1) \cdot 10.8\%}} \approx 0.24$$

for stable trials.

With this correlation adjustment, the null hypothesis in (6) is rejected with a *z*-statistic of 2.014 (or a *p*-value of 0.022), providing evidence for a difference in behavior between theoretical maximizers and theoretical randomizers. As predicted, when facing different environments (i.e., different payoffs in the experiment), theoretical randomizers indeed randomize more often than theoretical maximizers. Subjects responded differently by adapting to the environment and showing different stable behaviors.

After the game, we asked subjects in a survey whether they employed a specific strategy, and 74.7% of the subjects reported ^aYes^o. We perform the same hypothesis test (6) for subjects who reported ^aYes^o and those who reported ^aNo^o separately. We find that the effect holds strongly for individuals who reported that they used a specific strategy (adjusted *z*-statistic of 2.489, adjusted *p*-value of 0.006), but not for those who did not (adjusted *z*-statistic of –0.058, adjusted *p*-value of 0.523). This serves as another robustness check that the effect is driven by intentional behavior on the part of the subjects, and is not purely noise. This also provides empirical evidence for theories that attempt to explain probability matching through pattern seeking [55] and searching for causal relationships [45].

In principle, one can also perform the same test for different slices of the subjects across demographic, socioeconomic, and game-specific dimensions shown in Table 2. This helps to build intuition on whether the same effect holds true universally, and what types of individuals have stronger effects. However we acknowledge that the power of our study is limited due to the sample of 75 subjects, particularly after multiple-testing adjustment, and we leave this to a future study.

3.6 Individual differences

To jointly study individual differences in decision-making with the variables considered in Table 2, we consider a logistic regression model at the level of each guess by the individual subject. Specifically, for individual i, at

probability of choosing the dominant option D is modeled by:

$$f^{(i)} = \mathbb{P}(y_t^{(i)} = 1)$$

$$= Logistic(IsTheoryMaximizer_i + IsMale_i + HasChild_i + AgeBucket_i$$

$$+ HasProbClassExp_i + HasGamblingExp_i + HasInvestExperience_i \qquad (10)$$

$$+ IsStudent_i + IsOwn_i + IncomeBucket_i + TotalAssetBucket_i$$

$$+ HasPattern_i)$$

where *Logistic*(x) = $(1 + \exp(-x))^{-1}$.

We have seen in Section 3.4 that individual decisions are correlated over time. Therefore, the errors for regression (10) may be autocorrelated. We group trials from the same individual together, and order their decisions chronologically. In particular, the response variable is organized as:

$$\underbrace{(\underbrace{y_1^{(1)}}; \underbrace{y_2^{(1)}}; \cdots; \underbrace{y_T^{(1)}}_T; \underbrace{y_1^{(2)}}; \underbrace{y_2^{(2)}}; \cdots; \underbrace{y_T^{(2)}}_T; \cdots; \underbrace{y_1^{(n)}; \underbrace{y_2^{(n)}}; \cdots; \underbrace{y_T^{(n)}}_T)';}_{n-\text{th subjects } T \text{ trials}})';$$

and we apply the Newey-West heteroskedasticity and autocorrelation consistent (HAC) estimator for the variance of the coefficients in the following results. We adopt Newey and West's [65] suggestion to choose the truncation parameter to be the integer part of $4(nT/100)^{2/9}$, which is 11 in our case. This indeed increases the variance estimate of our coefficients compared with the case of independent errors, and our results are not materially different with respect to the choice of the truncation parameter. In fact, we have tried an estimation with truncation parameter to be 150, the number of total valid decisions for one individual. The main variable *IsTheoryMaximizer* remains statistically significant at 5%.

<u>Table 3</u> summarizes the independent variables in Eq. (10). These variables correspond to the collected personal information of the subjects (see <u>Table 2</u>), categorized to make them proper binary or ordinal variables. We have dropped several variables that are highly collinear with the covariates in (10). The *p*-value of the log-likelihood ratio test of the full model is 4×10^{-54} , implying a high degree of significance.

The first variable, IsTheoryMaximizer, encodes whether the

demographic groups, to other financial and non-financial contexts, and at different magnitudes of payoffs.

In addition to testing the evolutionary model of Brennan and Lo [50], our experimental results suggest that it is valuable to derive behavioral predictions and implications through an evolutionary lens. Traditional utility-based theories would yield the same maximizing behavior for all four designs in our experiment. Yet we find evidence for differences in reality, and the evolutionary framework offers a potential explanation and prediction for such behaviors: the environment matters.

More generally, financial marketsĐa collection of individual decision makersĐcan also be studied using the same principles, leading to the Adaptive Markets Hypothesis [66, 67] and its many empirical implications [68±70]. In the same way that micro-level individual decision making can be better understood through an evolutionary lens, markets and societies at the system-wide and macroscopic level can also benefit an adaptive perspective.

Acknowledgments

The authors thank the MIT Behavioral Research Laboratory for hosting the experiment, and Jennifer L. Walker, Jayna Cummings, and Allison McDonough for their assistance in the experiments. Insightful comments and discussions from Jason Anthony Aimone (editor), Kim Fairley (reviewer), Taibo Li and Alex Huang are also greatly appreciated.

Author Contributions

Conceptualization: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Data curation: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Formal analysis: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Investigation: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Methodology: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Supervision: Andrew W. Lo.

Validation: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Writing ± original draft: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

Writing ± review & editing: Andrew W. Lo, Katherine P. Marlowe, Ruixun Zhang.

References

- 1. von Neumann J, Morgenstern O. Theory of Games and Economic Behavior. Princeton, NJ: Princeton University Press; 1944.
- Nash JF. Equilibrium points in n-person games. Proceedings of the national academy of sciences. 1950; 36(1):48±49. https://doi.org/10.1073/pnas.36.1.48 PMID: 16588946
- Lucas RE Jr. Expectations and the neutrality of money. Journal of Economic Theory. 1972; 4(2):103± 124. https://doi.org/10.1016/0022-0531(72)90142-1
- Samuelson PA. Proof that properly anticipated prices fluctuate randomly. Industrial Management Review. 1965; 6(2):41±49.
- 5. Fama EF. Efficient Capital Markets: A Review Of Theory And Empirical Work. The Journal of Finance. 1970; 25(2):383±417. https://doi.org/10.1111/j.1540-6261.1970.tb00518.x
- Black F, Scholes M. The Pricing of Options and Corporate Liabilities. Journal of Political Economy. 1973; 81(3):637±654. https://doi.org/10.1086/260062
- Merton RC. Theory of rational option pricing. The Bell Journal of economics and management science. 1973; 4(1):141±183. https://doi.org/10.2307/3003143

- 8. Kocherlakota NR. Modern macroeconomic models as tools for economic policy. The Region (Federal Reserve Bank of Minneapolis). 2010; p. 5±21.
- 9. Hu HT. Efficient markets and the law: A predictable past and an uncertain future. Annual Review of Financial Economics. 2012; 4(1):179±214. https://doi.org/10.1146/annurev-financial-110311-101733
- Kahneman D, Slovic P, Tversky A. Judgment under Uncertainty: Heuristics and Biases. Cambridge, UK: Cambridge University Press; 1982.
- 11. Tversky A, Kahneman D. Judgment under uncertainty: Heuristics and biases. Science. 1974; 185 (4157):1124±1131. https://doi.org/10.1126/science.185.4157.1124 PMID: 17835457
- 12. Kahneman D, Tversky A. Prospect theory: An analysis of decision under risk. Econometrica. 1979; 47 (2):263±291. https://doi.org/10.2307/1914185
- Shefrin H, Statman M. The disposition to sell winners too early and ride losers too long: Theory and evidence. The Journal of finance. 1985; 40(3):777±790. https://doi.org/10.1111/j.1540-6261.1985. tb05002.x
- 14. Odean T. Are investors reluctant to realize their losses? The Journal of finance. 1998; 53(5):1775± 1798.
- Barber BM, Odean T. Boys will be boys: Gender, overconfidence, and common stock investment. The quarterly journal of economics. 2001; 116(1):261±292. https://doi.org/10.1162/003355301556400
- 16. Gervais S, Odean T. Learning to be overconfident. The Review of Financial Studies. 2001; 14(1):1±27. https://doi.org/10.1093/rfs/14.1.1
- De Bondt WF, Thaler R. Does the stock market overreact? The Journal of finance. 1985; 40(3):793± 805.
- Huberman G, Regev T. Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. The Journal of Finance. 2001; 56(1):387±396. https://doi.org/10.1111/0022-1082. 00330
- Tversky A, Kahneman D. The Framing of Decisions and the Psychology of Choice. Science. 1981; 211:453±458. https://doi.org/10.1126/science.7455683 PMID: 7455683
- Lichtenstein S, Fischhoff B, Phillips LD. Calibration of probabilities: The state of the art to 1980. In: Kahneman D, Slovic P, Tversky A, editors. Judgment under uncertainty: Heuristics and biases. Cambridge, UK: Cambridge University Press; 1982. p. 306±334.
- Gneezy U, List JA, Wu G. The uncertainty effect: When a risky prospect is valued less than its worst possible outcome. The Quarterly Journal of Economics. 2006; p. 1283±1309. https://doi.org/10.1093/ qje/121.4.1283
- Mahoney MJ. Publication prejudices: An experimental study of confirmatory bias in the peer review system. Cognitive therapy and research. 1977; 1(2):161±175. https://doi.org/10.1007/BF01173636
- Grant DA, Hake HW, Hornseth JP. Acquisition and extinction of verbal conditioned responses with differing percentages of reinforcement. Journal of experimental psychology. 1951; 42(1):1±5. https://doi.org/10.1037/h0054051 PMID: 14880648
- Hake HW, Hyman R. Perceptions of the Statistical Structure of a Random Series of Binary Symbols. J Exp Psychol. 1953; 45:64±74. https://doi.org/10.1037/h0060873 PMID: 13035003
- Herrnstein RJ. Relative and absolute strength of responses as a function of frequency of reinforcement. Journal of the Experimental Analysis of Behaviour. 1961; 4(3):267±272. https://doi.org/10.1901/jeab. 1961.4-267
- Herrnstein RJ. On the law of effect. Journal of the Experimental Analysis of Behavior. 1970; 13:243± 266. https://doi.org/10.1901/jeab.1970.13-243 PMID: 16811440
- 27. Herrnstein RJ. The Matching Law. Cambridge, Massachusetts: Harvard University Press; 1997.
- Bradshaw CM, Szabadi E, Bevan P. Behavior of humans in variable-interval schedules of reinforcement. Journal of the Experimental Analysis of Behavior. 1976; 26:135±141. https://doi.org/10.1901/ jeab.1976.26-135 PMID: 16811935
- Davison M, McCarthy D. The matching law: A research review. Hillsdale, NJ: Lawrence Erlbaum; 1988.
- Vulcan N. An Economist's Perspective on Probability Matching. Journal of Economic Surveys. 2000; 14:101±118. https://doi.org/10.1111/1467-6419.00106
- Kogler C, Kichberger A. Dual Process Theories: A Key for Understanding the Diversification Bias? J Risk Uncertainty. 2007; 34:145±154.
- Deneubourg JL, Aron S, Goss S, Pasteels JM. Error, communication and learning in ant societies. European Journal of Operational Research. 1987; 30(2):168±172. https://doi.org/10.1016/0377-2217(87) 90093-2

- Wozny DR, Beierholm UR, Shams L. Probability Matching as a Computational Strategy Used in Perception. PLoS Comput Biol. 2010; 6(8):e1000871. <u>https://doi.org/10.1371/journal.pcbi.1000871</u> PMID: 20700493
- 59. Baum WM. On two types of deviation from the matching law: Bias and undermatching. Journal of the Experimental Analysis of Behavior. 1974; 22:231±242. <u>https://doi.org/10.1901/jeab.1974.22-231</u> PMID: 16811782
- Horne PJ, Lowe CF. Determinants of human performance on concurrent schedules. Journal of the Experimental Analysis of Behavior. 1993; 59:29±60. https://doi.org/10.1901/jeab.1993.59-29 PMID: 8433066
- Tversky A, Kahneman D. Belief in the law of small numbers. Psychological bulletin. 1971; 76(2):105. https://doi.org/10.1037/h0031322
- Kahneman D, Tversky A. Subjective probability: A judgment of representativeness. Cognitive psychology. 1972; 3(3):430±454. https://doi.org/10.1016/0010-0285(72)90016-3
- Newcombe RG. Two-sided confidence intervals for the single proportion: comparison of seven methods. Statistics in medicine. 1998; 17(8):857±872. https://doi.org/10.1002/(SICI)1097-0258(19980430) 17:8%3C857::AID-SIM777%3E3.0.CO;2-E PMID: 9595616
- Brown LD, Cai TT, DasGupta A. Interval estimation for a binomial proportion. Statistical science. 2001; p. 101±117.
- Newey WK, West KD. A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix. Econometrica. 1987; 55(3):703±708. https://doi.org/10.2307/1913610
- Lo AW. The adaptive markets hypothesis. Journal of Portfolio Management. 2004; 30(5):15±29. https:// doi.org/10.3905/jpm.2004.442611
- Lo AW. Adaptive Markets: Financial Evolution at the Speed of Thought. Princeton, NJ: Princeton University Press; 2017.
- Neely CJ, Weller PA, Ulrich JM. The adaptive markets hypothesis: evidence from the foreign exchange market. Journal of Financial and Quantitative Analysis. 2009; 44(2):467±488. <u>https://doi.org/10.1017/</u> S0022109009090103
- Kim JH, Shamsuddin A, Lim KP. Stock return predictability and the adaptive markets hypothesis: Evidence from century-long US data. Journal of Empirical Finance. 2011; 18(5):868±879. https://doi.org/ 10.1016/j.jempfin.2011.08.002
- Khuntia S, Pattanayak J. Adaptive market hypothesis and evolving predictability of bitcoin. Economics Letters. 2018; 167:26±28. https://doi.org/10.1016/j.econlet.2018.03.005