Would you pay for transparently useless advice? A test of boundaries of beliefs in the folly of predictions

Nattavudh Powdthavee[*]

Centre for Economic Performance, London School of Economics

The Melbourne Institute of Applied Economics and Social Research, University of Melbourne

Yohanes E. Riyanto

Nanyang Technological University

26th September 2012

[*] *Corresponding author*: Nattavudh Powdthavee
Address: Centre for Economic Performance, London School of Economics, Houghton Street, London, UK, WC2A 2AE.
Email: n.powdthavee@lse.ac.uk.
Tel: +44(0) 7990 815924.
Acknowledgment: We are grateful to Gigi Fosters, Paul Frijters, Daniel Houser, Jack Knetsch, Nuarpear Lekfuangfu, Pongsak Luangaram, Nicolas Motz, Andrew Oswald, Worrawat Sritrakul, Matthew Rabin, Jon Wongswan, Dimitri Vayanos, Marie-Claire Villeval for helpful comments on the draft. We also thank Liu Jia, Weerachart Kilenthong, and students and staffs at University of Thai Chamber of Commerce and Nanyang Technological University for their help in carrying out the experiments. The research was supported by the SUG grant provided by Nanyang Technological University, and was carried out whilst the first author was employed at Nanyang Technological University, Singapore.

Abstract

Standard economic models assume that the demand for expert predictions arise only under the conditions in which (i) individuals are uncertain about the underlying process generating the data, and (ii) there is a strong belief that past performances predict future performances. Thus, when it is transparently clear that the event in question is truly random, there should be no market for expert predictions. We set up the strongest possible test of these assumptions in the laboratory. In contrast to the theoretical suggestions made in the literature, people are willing to pay for predictions of truly random outcomes after witnessing only a short streak of accurate predictions *live* in the lab. Potential explanations and implications of such "irrational learning" are discussed in the contexts of economics and finance.

Keywords: behavioral finance; expertise; hot hand; dynamic-inference model; random streaks
JEL: C91; D03

"[Expert] intuition cannot be trusted in the absence of stable regularities in the environment" – Daniel Kahneman in *Thinking, Fast and Slow* (2011)

1. Introduction

Why do humans pay for expert predictions when most future events are predominantly unpredictable? What explains, for example, the significant amount of money spent in the finance industry on people who appear to be commenting about random walks,[2] payments for services by political and economic forecasters who are often only slightly better at forecasting the future than non-experts,[3] or some other false-expert setting?[4]

Economists typically dismiss such behaviors as random error in decision-making. This is the notion that an average person is disinclined to commit such errors, and that people rationally pay for expert predictions only when they are *a priori* uncertain about the underlying data-generating process, yet maintain the belief that some systematic predictions are potentially possible. What this implies is that economists tend to set a high bar for the way people rationalize their actions *ex ante*, i.e., a significant degree of uncertainty would need to be present in our mind before we could be persuaded to pay for a prediction that is ultimately useless.

By contrast, the psychology literature assumes that human beings are hypersensitive at detecting agency, even when none exists, to help them to explain phenomena that cannot be easily explained. This implies that on average people will be happy to pay for advice that is generally counterintuitive if they believe there is an intelligent agent making the predictions for them. Such a divide between the two social-science disciplines in the beliefs of how people's preferences for expert predictions are formed is scientifically unattractive.

[2] Hung et al. (2008), Chater et al. (2010).
[3] Tetlock (2006).
[4] http://en.wikipedia.org/wiki/Paul_the_Octopus.

While we cannot deny from innumerable observations that individuals can be induced to believe in the "hot hands" of experts (and consequently have been exploited by charlatans, witch-doctors, fortune tellers, casino operators, and mutual-fund managers the world over for centuries), we know much less about the *boundaries* of such seemingly irrational behaviors. For example, we do not know the extent to which people can be induced to believe in the hot hand of an expert if it is *a priori* clear that the event is truly random and if accurate expert predictions can be explained only by luck and not skills.

Our paper attempts to contribute to this research area by investigating whether people are willing to pay for predictions in a situation in which there is true randomness and predictions are blatantly useless. In a laboratory setting of clearly non-existent expertise, can an average person be challenged to switch from having the correct prior belief that "outcomes are determined by chance and predictions are inherently useless" to the false belief that "predictions provide useful information about the future" – thus leading the person to buy subsequent predictions at a fixed price – if she had recently observed an improbable streak of successful predictions being made in front of her *live*? In contrast to the literature, we found that the answer is yes and that the size of the error made systematically by people is large.

This paper is structured as follows: section 2 presents a brief review of the relevant literature, section 3 outlines the stylized dynamic-inference model, and section 4 describes the data and experimental procedures. Empirical strategy is discussed in section 5. Results are then laid out in section 6. Section 7 provides our discussions, the implications of our results, and the concluding remarks.

2. Background

There is little economic theory in this area. A few exceptions are the work of Matthew Rabin and Dimitri Vayanos. In their papers, Rabin (2002) and Rabin and Vayanos (2010) outline a model in which believers of “the law of small numbers” – i.e., those who believe that a small sample of signals represents the parent population from which it is drawn (Tversky & Kahneman, 1971) – will be willing to pay for the services of financial analysts after observing randomly occurring streaks of profitable financial performances provided by mutual-fund managers. Such a belief in the hot hand of a financial expert arises because individuals over-infer the financial manager’s ability following a streak of successful stock performances. In other words, an investor who believes that the performance of a mutual fund is a combination of the manager’s skills and luck will, at first, underestimate the likelihood that a manager of average ability will exhibit a streak of above- or below-average performance. Following good or bad streaks, however, the investor will revert to overestimate the likelihood that the manager is above or below average, and so in turn will over-infer that the streak of unusual performance will continue (see Gilovich et al., 1985). Rabin and Vayanos’s model thus predicts that following a streak of “good” signals in settings where there is an element of skill involved in generating such a streak – e.g., a sequence of successful performances by stockbrokers or managers of actively managed funds – believers of the law of small numbers will be happy to pay for real-time price information provided by these financial “experts” even though it is well-documented that actively managed funds do not outperform their market benchmark on average (see, e.g., Fama, 1991; Cahart, 1997). The key assumption here is that, in order to form a belief in the hot hand of a manager, individuals must be unsure about the data-generating process of stock movements in the first place, but nevertheless hold the belief that past returns help to predict future returns (and that past management performances help to predict future performances).

Econometric evidence on the evolution of beliefs in expert predictions is also scarce. Much of the literature in this field tends to focus on situations in which no experts were present to generate predictions of future independent and identically distributed (i.i.d.)

signals. For example, Croson and Sundali (2005) show that, in a game of roulette, casino gamblers tend to bet against a sufficiently long streak rather than with a streak, which is consistent with the “gambler’s fallacy,” while at the same time gamblers tend to bet on more numbers after winning than after losing, which is a behavior consistent with the hot hand effect. Terrell and Farmer (1996) and Terrell (1998) find evidence of the gambler’s fallacy in horse and dog racing. The authors show that gamblers are less likely to bet on repeat winners by post position. For example, if the animal in post position 3 wins a race, then in the next race the (different) animal in post position 3 is significantly under-bet. Using a computerized roulette game, Ayton and Fischer (2004) show that subjects tend to believe in the gambler’s fallacy with respect to the sequence of outcomes of the roulette wheel. Yet when the subjects’ role is to predict the outcomes of the roulette wheel, they tend to over-predict how well or badly they would do at predicting based on their previous streak of predictions, thus exhibiting the hot hand fallacy.

More recently, Guryan and Kearney (2008) found unique evidence of the potential hot hand effect in stores that sell lottery numbers. The authors show that, in the week following the sale of a large-prize-winning ticket, the winning store experiences a significant increase in relative sales for the winning lottery game. Using a unique panel data of lottery players, Jørgensen et al. (2011) present evidence that, while most lottery players tend to pick the same set of numbers week after week without regard to the numbers drawn in the lottery in previous weeks, those people who do change the set do so in such a way that is consistent with the law of small numbers. On average, these “switchers” move away from numbers that have been recently drawn (the gambler’s fallacy), and move toward numbers that are on a streak (the hot hand fallacy). However, in both scenarios – “lucky” lottery stores and “lucky” lottery numbers – the switches in preferences that are consistent with the hot hand effect are often short-lived; for example, Clotfelter and Cook (1991, 1993) and Teller (1994) show that, shortly after a lottery number wins, individuals are significantly more likely to bet on it. The effect soon diminishes: A few months later the winning number is as popular as the average number.

To the best of our knowledge, only one other study has explicitly tested the implications of expert predictions in a pure i.i.d. setting. In experiments ran by Huber, Kirchler, and Stockl (2010), participants were asked to bet on computer-generated outcomes of coin flips, or rely on randomized "expert" predictions, or choose a risk-free alternative in their investment decisions. The researchers were able to show that people who rely on randomized experts tend to pick those experts who have been successful in the past, which is consistent with the hot hand fallacy. Those who decide the outcome of coin flips on their own tend to behave consistently with the gambler's fallacy, as the frequency of betting on heads (tails) decreases after streaks of heads (tails). However, a potential shortcoming of their study is that, instead of real coins, a "virtual" coin was used in the experiment to generate the signals. This could have induced the erroneous belief that the coin was not fair and that its outcomes were predictable. They also informed participants that the predictions of coin flips were made by "experts" rather than allow the participants to form their own beliefs, which may from the outset have given the participants the confounding impression that skills were involved in predicting the outcomes of these computerized coin flips.

Two main lessons may be learned from the literature. The first is that humans tend to behave in such a way that is consistent with the gambler's fallacy in situations where the underlying mechanism generating the streak of signals is transparently random. The second is that the hot hand fallacy normally arises in situations where human skills are *a priori* – albeit erroneously – perceived as part of the streak-generating process. Or to quote Huber, Kirchler, and Stockl (2010):

The hot hand belief is usually attributed to human skilled performance, whereas the gambler's fallacy is often attributed to inanimate chance mechanism. (p. 446)

The current article builds on this little literature and sets out to test one of the key assumptions in the model by Rabin and Vayanos (2010): In a truly random situation in which experts' predictions are *a priori* known to be transparently useless, people's behaviors will be influenced only by the gambler's fallacy and not by the hot hand of the non-existent expert.[5]

3. The stylized model

Our experimental predictions are motivated by the predictions of the dynamic-inference model by Rabin and Vayanos (2010). We assume that an individual observes two sequences of signals whose probability distributions depend on some underlying states. The first set of signals is a sequence of truly i.i.d. outcomes, e.g., a series of "fair" coin flips, s_t. The second set of signals, a_t, is a sequence of an agent's successful predictions of the i.i.d. outcomes s_t.[6] The signal s_t in periods t = 1, 2, …, is

$$s_t = \mu + \varepsilon_t, \qquad (1)$$

and the signal a_t in periods t = 1, 2, …, is

$$a_t = \varphi_t + \nu_t, \qquad (2)$$

where μ is the long-run mean of the i.i.d. signals, which is fixed at 0.5 for a series of "fair" coin flips; φ_t is the state, which is interpretable as the ability of the agent making the

[5] This assumption is supported by the work of Burns and Corpus (2004). The authors gave participants information about a streak of events but varied the scenarios in such a way that the mechanism generating the events would vary according to the participants' judgment of its randomness. They found that, when participants are presented with less random scenarios, they are more likely to believe that a hypothetical streak will continue. A similar conclusion is also observed in an experiment by Asparaouhova, Hertzel, and Lemmon (2009). They conclude that, with increasing perception of randomness for the outcome-generating process, subjects are more likely to predict systematic reversal of signals (the gambler's fallacy) rather than a continuation of signals (the hot hand fallacy).

[6] It should be noted that the second set of signals is viewed privately and allowed to vary across subjects in our experiment.

prediction of s_t; and ε_t and v_t are i.i.d. normal shocks with means zero and variances $\sigma_\varepsilon^2, \sigma_v^2 > 0$. We can also interpret the shock v_t as the agent's luck at predicting s_t in Period *t*. More importantly, s_t is assumed to be determined independently from a_t, i.e., an outcome of a coin flip is determined by the actual flip rather than by a prediction given by someone.

According to Rabin and Vayanos, we can model the gambler's fallacy as the mistaken belief that the sequences $\{\varepsilon_t, v_t\}_{t\geq 1}$ are not i.i.d. but exhibit systematic reversal. More specifically, we assume that, for the individual,

$$\varepsilon_t = \tilde{\varepsilon}_t - \alpha \sum_{k=0}^{\infty} \delta^k \varepsilon_{t-1-k}, \tag{3}$$

and,

$$v_t = \tilde{v}_t - \alpha \sum_{k=0}^{\infty} \delta^k v_{t-1-k}, \tag{4}$$

where the shocks $\tilde{\varepsilon}_t$ and $\tilde{v}_t$ are i.i.d. and normal with means zero and variances $\tilde{\sigma}_\varepsilon^2, \tilde{\sigma}_v^2$, and $\alpha, \delta \in [0,1)$ are exogenous parameters. We denote the parameter α as the strength of the belief in the gambler's fallacy, and δ as the memory duration of the gambler's fallacy since *k*th period in both scenarios.[7] When α equals zero, the individual is an error-free Bayesian who treats the sequences correctly as i.i.d. When $\alpha > 0$, the individual has the wrong model of how signals are generated. In the spirit of Rabin and Vayanos, we shall call the individual who believes in the law of small numbers ($\alpha > 0$) "Freddy."

To see how this corresponds to the gambler's fallacy, we can take the expectation of both sides of Eqs (3) and (4) conditional on Freddy's information as of Period *t*−1.

$$E_{t-1}(\varepsilon_t) = -\alpha \sum_{k=0}^{\infty} \delta^k E_{t-1}(\varepsilon_{t-1-k}), \tag{5}$$

and,

[7] In a more complex model, both parameters α and δ can vary across different i.i.d. settings for the same individual.

$$E_{t-1}(v_t) = -\alpha \sum_{k=0}^{\infty} \delta^k E_{t-1}(v_{t-1-k}). \tag{6}$$

If Freddy believes that all shocks up to $t-1$ have been positive, then he expects the shocks ε_t and v_t to be negative because $\alpha > 0$ and $\delta \geq 0$. With respect to the predicting agent, Freddy believes that, if the agent has been lucky at predicting the outcomes of a transparently i.i.d. process up to Period $t-1$, then his luck should reverse in Period t.

In principle, the state variable φ_t can be defined as the probability that the agent makes successful predictions of i.i.d. outcomes, a probability which is known and constant. This is because it should be transparent to anyone able to carry out simple probabilistic calculations that there is no systematic method to accurately predict pure i.i.d. outcomes, in our case coin flips. Nevertheless, we will allow for the possibility that the state φ_t evolves according to the auto-regressive process

$$\varphi_t = \psi + \rho(\varphi_{t-1} - \psi) + \eta_t, \tag{7}$$

where $\rho \in [0,1)$ is the reversion rate to the long-run average ψ, and η_t is an i.i.d. normal shock with mean zero and variance σ_η^2 and that is independent of v_t. One of the key assumptions is that, when Freddy is confident that the state is constant ($\sigma_\eta^2 = 0$), he will not be able to develop a belief in the time-varying state, and his predictions will be influenced only by the gambler's fallacy. It is only when Freddy is *uncertain* about the mechanisms generating the data, and where a belief in serially correlated variation is *a priori* plausible, i.e., $\sigma_\eta^2 > 0$ and $\rho > 0$, that a belief in the hot hand of the predicting agent may develop – overtaking the gambler's fallacy – after observing a persistent streak of successful signals.

Alternatively, the model by Rabin and Vayanos (2010) predicts that Freddy will never come to believe in the hot hand of the predicting agent – let alone pay for the agent's predictions – if it is *a priori* clear to Freddy that the event is truly random, and that the agent's predicting ability is time-invariant and will never be better than Freddy's ability. Therefore, it is more likely in cases of transparently i.i.d. events that Freddy will treat a streak of an agent's successful predictions up to Period $t-1$ as luck and will believe that the luck should revert in Period t. Alternatively, a rejection of the hypothesis suggests that Freddy is sufficiently *a priori* uncertain about the state to change his belief from "predictions are useless" to "predictions are useful" after observing a streak of successful signals. In short, the empirical predictions of Rabin and Vayanos's dynamic-inference model are:

1. In transparently i.i.d. events, people will generally behave according to the gambler's fallacy.
2. Only when individuals perceive that human skills may be involved in making accurate predictions of truly random outcomes – i.e., that past performances do indeed predict future performances – can a belief in the hot hand of an expert develop.
3. In a market where predictions are bought and sold, people will never rationally pay for what they believe to be *transparently* useless advice.

4. Experimental framework

4.1. Data

To investigate the conditions under which people would be willing to pay for predictions of truly random events, a series of laboratory experiments was conducted on voluntary participants in Thailand and Singapore. We ran our first set of experiments in Thailand in December 2011, and the randomly selected participants were undergraduate students at the University of the Thai Chamber of Commerce and Chulalongkorn University in Bangkok (N=177). We then ran our second set of experiments in Singapore in March 2012. Here, the volunteer participants originated from randomly drawn undergraduate students at Nanyang

Technological University (*N*=201). Overall, participants were from various schools and faculties, including humanities and social sciences, engineering, sciences, and business and accounting. We ran 12 sessions in total (4 in Thailand and 8 in Singapore) and were able to recruit approximately 45 people per session in Thailand and 30 people per session in Singapore.

Participants were randomly assigned to cubicles as they entered the labs. They were then asked to complete two tasks. The first task involved participants placing bets on the outcomes of five rounds of "fair" coin flips. To ensure the fairness of the coins used in the experiment, from the beginning we made explicitly clear to participants the following:

(i) The coins will be supplied by the participants rather than the experimenters.

(ii) The coins will be changed after the second and fourth flip.

(iii) A volunteer participant in the experiment will be randomly chosen to step out in front of everyone and flip the coin.

(iv) The coin-flipper will be changed in every round.

Participants were each given an initial endowment with which to make their bets in the five rounds of coin flips. There was a minimum bet of 10 tokens per round, and participants were not allowed to go bankrupt before the final round was reached. Participants in Thailand were given an initial endowment of 100 tokens. Since a few participants in Thailand went bankrupt before the final round, when we ran it in Singapore we decided to give each one a higher endowment level at the start of the experiment of 300 tokens. Placing a correct bet was worth double, and an incorrect one was worth zero in return. Each participant was also given at the beginning of the experiment five numbered envelopes, which were taped on each cubicle's table. Contained within was a "prediction" of the coin flip that *had not happened yet* in each of the numbered rounds. In each round, participants were given an opportunity to pay a fixed price of 10 tokens to see the inside of the corresponding numbered envelope *before a bet was* *placed and the coin flipped*. All other participants who decided not to pay were instructed to open the corresponding numbered envelope for *free* immediately after the flip in order to view whether the prediction matched the outcome. In addition, great care was taken not to provide any misleading information, e.g., on who made the predictions, whether the predictions were made by an expert or a group of experts, or how the predictions were generated, which could have potentially primed participants into buying (or not buying) the predictions by evoking the impression that the underlying process generating accurate predictions was humanly possible. Participants were then told that the remaining endowment at the end of the fifth round would be converted to either Singapore dollars (SG$) or Thai baht at the exchange rate of 50 points = SG$1 (25 Thai baht) or approximately US$0.9.

To guarantee a significant number of participants received at least four consecutive correct predictions in five rounds of fair coin flips, predictions were generated and assigned in such a way that approximately ½ of *N* received one correct prediction after round 1, ½×½ of *N* received two correct predictions after round 2, etc. (see Figure 1). This method of **randomization-in-randomization** (R-in-R) – i.e., the process of randomizing people within the same session into control and treatment groups – ensured that at least 1 out of *N* participants per session would randomly receive all-correct predictions *irrespective of the actual outcomes of the coin flips*. The R-in-R design also ensures that we have a random split of participants with divergent beliefs about the predictability of the coin in the control and treatment groups. Of the total of 378 participants from the two countries, 191 received a correct prediction in the first round of coin flips, 92 all-correct predictions after the first two rounds, 48 after the first three rounds, and 23 after the first four rounds.[8]

The second task of the experiment involved participants completing a set of probability tests (which was incentivized with each correct answer given = SG$0.20), as well

[8] The method was first seen on UK television in 2008, demonstrated by the British magician Derren Brown. In his show called *The System*, Brown used the technique to illustrate how he was able to predict, for one particular person, six consecutive wins at a horserace.

as a set of standard control questionnaires. At the end of the experiment, all participants were debriefed on the nature of the experiment either immediately (Thailand) or later via email (Singapore). The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).[9]

4.2.The first-round buyers

Of the 378 participants, 55 (or 14.5%) were first-round prediction buyers, i.e., the prediction was bought by individuals before any sequence of correct or incorrect predictions was observed. Of these 55, 41 (74.55%) were Singaporeans and 14 (25.45%) were Thais. For these subjects, it may be that (a) they did not completely grasp all the facts that we had stated earlier in the experiment about the fairness of the coins, or (b) they confidently believed that coins are predictable, i.e., those with a positive prior on predictions being useful from the start $(\alpha > 0)$. Since we cannot distinguish (a) from (b), our main focus will be on the 323 non-first-round buyers (α is presumably zero, at least in round *j*=1). We briefly return in section 6.1 to analyze the first-round buyers.

We report in Appendix A the descriptive statistics of the buying behaviors of first-round buyers – all and by nationality – and all non-first-round buyers.

5. Econometric specifications

To test whether participants' betting behaviors are influenced by the gambler's fallacy, we estimate the following equation:

$$h_{ij\geq 2} = \alpha_j + \pi_j s\_h_{j-1} + \varphi_j s\_t_{j-1} + \delta_j b_{ij} + \theta_j (s\_h_{j-1} \times b_{ij}) + \varsigma_j (s\_t_{j-1} \times b_{ij}) + \upsilon_{it}, \qquad (8)$$

where *i* indexes the individual; h_{ij} is an indicator variable that takes the value of 1 if the individual bets "head" in round $j \geq 2$, and 0 if the individual bets "tail"; s_h_{j-1} is a dummy

[9] We outline the experimental procedures at the end of this paper in Appendix C.

representing a streak of heads up to round $j-1$; s_t_{j-1} is a dummy representing a streak of heads up to round $j-1$; *b* is an indicator variable representing whether the subject paid 10 tokens to see the prediction in round *j*; and υ_{it} is the error term. The hypotheses to be tested are: (i) If a streak of heads is observed up until round *j*−1, then the individual is more likely to bet tail in round *j*, i.e., π<0 (similarly, if a streak of tails is observed until round *j*−1, then the individual is more likely to bet head in round *j*, i.e., φ>0), and (ii) depending on the prediction in round *j*, those who paid for it will place a bet that corresponds to the hot hand of the expert's prediction rather than be influenced by the gambler's fallacy, i.e., conditioning on prediction in round *j* is "head," the interaction parameter θ should be positive to cancel out the negative π of the gambler's fallacy.

A second equation tests whether a streak of past predictions matters to the subject's purchasing decision in the current round of coin flip, and can be written as

$$b_{ij} = \beta_j p_{i(j-1)} + \gamma b_{ij-1} + \varepsilon_{ij}, \tag{9}$$

where *p* is a set of dummy variables representing a streak of successful or failed predictions up to round *j*−1; b_{ij-1} is the decision to buy in the previous round; and *ε* is the error term. Our key parameter of interest here is *β*, which represents the effect of observing successful (or failed) streaks of past predictions on the subject's buying decision in round *j*. Note that, for *j*=2, the estimated effect of obtaining a correct prediction in round 1 is the effect relative to obtaining an incorrect prediction in round 1. For *j*={3,4,5}, the estimated effect of obtaining all-correct, or all-incorrect, predictions prior to round *j* is thus the effect relative to obtaining some correct and some incorrect signals, which is typically the outcomes that subjects *a priori* expected to see for predictions made on truly random events.

A third, final equation tests whether the paid-for predictions were treated seriously by the subjects who bought them. Here, a natural question is whether the amount of endowment used in each bet is larger on average among buyers than non-buyers. The hypothesis can be tested by estimating the following specification for round j={2, 3, 4, 5}:

$$g_{ij} = \psi_j + \omega_j b_{ij} + v_{ij}, \tag{10}$$

where g denotes the log of the endowment amount used to bet by individual i in round j, and b is a dummy variable indicating whether the prediction was bought by the subject in round j. Assuming that buyers would tend to treat the paid-for predictions seriously and consequently place larger bets than non-buyers, we would expect to see the estimated parameter ω to be positive and statistically significant for j={2, 3, 4, 5}.

6. Results

6.1. Are people's betting behaviors influenced by the gambler's fallacy?

Table 1 indicates whether non-first-round buyers are subject to the gambler's fallacy in their betting behaviors. Here, the dependent variable is an indicator variable that takes a value of 1 if the participant chooses to bet "head" in round j, and 0 if he chooses to bet "tail" in round j. Since the same sequence of coin flips is observed by all participants in a given session and there are 12 sessions in total, we have 118 observations of two heads after the first two rounds and 152 observations of two tails after the first two rounds (i.e., j=1 and j=2). The number of consecutive heads and tails becomes even smaller at the start of round 4; there are only 23 observations of three heads after the first three rounds and 48 observations of three tails after the first three rounds (i.e., j=1, j=2, and j=3). Our focus will be on the subject's propensity to bet "head" following streaks of heads and tails in rounds j=3 and j=4.

Looking at columns 1 (All; Round j=3) and 4 (All; Round j=4) of Table 1, we can see – based on the specification in Eq. (8) – that participants generally behave according to the gambler's fallacy in that:

(i) Participants are approximately 27 to 32 percentage points *less* likely to bet "head" in round j if a streak of *heads* has been observed up to round $j-1$.

(ii) Participants are approximately 12 to 18 percentage points *more* likely to bet "head" in round j if a streak of *tails* has been observed up to round $j-1$.

These estimated effects are also statistically robust at conventional confidence levels.

However, a more interesting pattern emerges when the "bought prediction in j" dummy and its interaction with the streak variables are introduced as independent variables in the "betting 'head' in round j" regression equations. Conditioning on the prediction in round j being "head," there is an offsetting effect to the gambler's fallacy for participants who bought the prediction and recently observed two consecutive heads up to $j-1$: The net effects on betting "head" for these individuals are ($-0.283 + 0.616 = 0.333$) in round j=3 and ($-0.253 + 0.628 = 0.375$) in round j=4. Those who "bought the prediction in j" and observed "a streak of tails up to $j-1$" are statistically indifferent between placing a bet on head or tail in round j if the prediction for round j is "tail": The net effects are ($0.163 + 0.087 = 0.250$) in round j=3 and ($0.0789 + -0.279 = -0.200$) in round j=4, which means that we cannot reject the null of zero at conventional confidence levels.

The results of Table 1 imply that participants who paid for the prediction in round j – presumably because they had (randomly) received accurate predictions in the previous rounds – will likely use them to guide how they place their bets in the same round. This is the case even when the paid-for predictions are predicting outcomes that go against their primitive intuition that, following an unexpected streak of outcomes in an i.i.d. event, future deviations in the opposite direction are then more likely.

6.2. Would people pay for transparently useless information?

We next turn to two questions: Do people who randomly receive correct predictions then perceive in the hot hand of an invisible agent and thus start paying for such transparently useless information? If so, how long is it before they start buying? We found the answers to be: yes, and not long.

Figure 2 illustrates this using the raw data of the pooled Thai and Singaporean sample. Conditioning on non-first-round buyers, the ratio of people who paid for round j's prediction increases from around 9% in round j=2 to 43% in round j=5, providing that they had just previously observed a streak of correct predictions up to round $j-1$. The rise is also monotonic and statistically well-determined throughout. The ratios of buyers to non-buyers are much lower (~1% to 4%) among those who observed a mix of correct and incorrect predictions. There is, however, a puzzling finding that requires further explanation: After observing four consecutive incorrect predictions, there is an increase in tendency for the participants to start paying for the prediction in round j=5, with the ratio of buyers to non-buyers at 17%. A similar picture is observed when we split the sample by countries (see Figures 3A and 3B), although the aforementioned patterns are noticeably more robust for the Thai sample than for the Singaporean sample.

Table 2 presents the results more formally through the use of multivariate regressions. Based on Eq. (9), we can see that the general patterns in our earlier figures are preserved here in our linear probability estimates, with or without additional control variables.[10] For example, in the full specification, the estimated probability of buying round 2's prediction for those who previously received a correct prediction in the first round are 5 percentage points higher than for those who previously received an incorrect prediction (p-value=0.040). The positive gradient in the buying decision as we move through rounds is also noticeable and individually statistically well-determined; holding the decision to buy in round $j-1$ constant,

[10] A non-linear model (marginal probit) estimated in Appendix B produces qualitatively the same results as our linear probability model.

probabilities of buying are 15 percentage points in round j=3 (p-value=0.001), 21 percentage points in round j=4 (p-value=0.004), and 27 percentage points in round j=5 (p-value=0.010). By contrast, the coefficient *All predictions up to j−1 had been incorrect* is statistically significantly different from zero only in the final round of coin flip: For individuals whose previous predictions had been incorrect four times in a row the probability of buying is 16 percentage points higher than those whose previous predictions were a mixture of successes and failures (p-value=0.047). The buying decision is also persistent in itself: The coefficients *Bought prediction in round j−1* are positive and statistically significant at conventional confidence levels in rounds j=4 and j=5.

So what explains the monotonic rise in the ratio of buyers to non-buyers for those whose previous predictions up to $j-1$ had been correct, and a sharp rise in the same ratio in the final round for those whose previous predictions up to $j-1$ had been incorrect? To answer these questions, let us recall the dynamic-inference model of Rabin and Vayanos (2010), which can reasonably explain one result but not the other. When the predictions are *wrong* four times in a row, which is an extraordinary outcome in itself, participants' prior beliefs are further strengthened that no human "skills" – which are typically associated with good performances – could possibly be involved in the underlying process generating the predictions; otherwise the envelopes would not have been performing poorly and the four consecutive wrong predictions must have been due to chance. Moreover, according to Rabin and Vayanos, when participants are confident that the predictions are generated by an i.i.d. process, the only primitive bias here is the gambler's fallacy. Recall that a streak of unsuccessful predictions is not necessarily the same as a streak of coin signals, i.e., unlike a streak of heads or tails of coin flips, it remains equally improbable in any round that participants will predict correctly the contents of the final fifth envelope. Thus, strangely, provided there is a sufficiently strong desire for a systematic reversal after witnessing a streak of "unlucky" envelopes, then there could be an incentive for participants whose envelopes had been unlucky four times in row to pay for the first time for the final round's prediction in the hope that the unlucky streak of predictions will reverse.

The more surprising result, which is intuitively more difficult to explain, remains the monotonically increasing effect of past realizations of correct predictions up to $j-1$ on people's belief in the hot hand of the envelopes (or of the non-existent expert). Economic models such as the dynamic-inference model developed by Rabin and Vayanos predict that, when people are confident that no human skill is involved in making accurate predictions of truly i.i.d. events, a prediction of a future coin flip is worth as much as a blank sheet of paper in an envelope, and people will not pay for such transparently useless advice. Our experiment provides statistically convincing evidence that rejects the idea that when individuals are absolutely rational (recall that these are non-buyers in the first round) they will remain so throughout the repeated random trial.

Put differently: Instead of exhibiting the anticipated gambler's fallacy following an improbable streak of correct predictions up to $j-1$, participants arrive at the belief that the remaining envelopes have the "ability" to accurately predict coin flips that had yet to take place.

Other results in Table 2 are also interesting. There is no difference in the buying behaviors between men and women, Singaporean and Thai subjects, or endowment levels. The coefficients *Proportion of correct answers in statistical test* have the anticipated negative sign, although they are not statistically significantly different from zero in the first three rounds of coin flip. The relationship between observing a streak of heads or tails up to $j-1$ and the propensity to buy a prediction in round j is not statistically robust in either round j=3 or round j=4. Finally, people who bought predictions that later turned out to be "bogus" – i.e., the interaction term between *Bought prediction in j−1* and *Made wrong bet in j−1* – are significantly less likely to buy in future rounds.

One question of interest is whether an individual's tendency to buy is different for different groups of people. Though not reported here, we found no statistically significant slope differences by gender or nationality, or in the test score of buying behavior among subjects who randomly received correct as well as incorrect predictions. This implies that the perceived hot hand in the envelopes predicting the future is not statistically more pronounced for males compared to females, or for the Singaporean sample compared to the Thai sample, or for those who scored better on average in statistics and probabilities. In short, there is no statistical evidence to support the notion that some people are systematically more (or less) susceptible to such an irrational behavior.

What about the betting and subsequent buying behaviors of those with a (potentially) positive prior, i.e., the first-round buyers? Are they significantly different from the non-first-round buyers? Table 3, which re-estimates the specification of Table 1 on the much smaller sub-sample of first-time buyers (N=55), suggests that both first-round buyers and non-buyers exhibit similar propensities to the gambler's fallacy in their betting behaviors: The coefficient *A streak of coins coming up heads up to j−1* in the betting "head" equation is −0.342 [*S.E.*=0.144] in round j=3, which is similar to the −0.324 [*S.E.*=0.061] observed earlier for the non-first-round buyers.

On the other hand, we can see from the buying equation of Table 4 that the estimated hot hand effect in round j=3 is almost twice the size of the same coefficient obtained without additional control variables in Table 2: The coefficients *All predictions up to j−1 had been correct* are 0.295 [*S.E.*=0.142] in round j=3 for the first-round buyers and 0.153 [*S.E.*=0.043] in the same round for the non-first-round buyers. This suggests evidence that the first-round buyers converge more quickly toward the belief that the envelopes have some predictive power, compared to the non-first-round buyers, i.e., those who have either a zero or a negative prior at the beginning of the experiment. This supports the earlier decision to exclude this smaller sample from our initial analysis.[11]

There are several potential objections to our results. The first is that participants only bought the predictions to please the experimenters. However, if this was the case, then we would have expected to see a similar buying pattern between the treated and those in the control group, but we did not. The second is that the "presence" of envelopes themselves may have added an unknown element to how each individual calculate his or her expected utility from placing such bets. For instance, one could imagine each participant thinking at the start of the experiment, "*I know that predictions contained within these envelopes are useless. But if they are really useless, then why would they be here in the first place?*" It is therefore easy to rationalize after seeing a streak of correct predictions being made by these envelopes that the experimenters are "up to something" – perhaps through some type of magic trick – and it would thus be better for the participants to buy the predictions even if they do not really believe in their predictability.[12]

While we cannot completely deny such beliefs in our volunteer participants, we can nevertheless shed further light on this issue. To do this, we asked our participants in Singapore to state their reasons in the post-experimental questionnaire for buying at least one of the predictions (if they did at some point) and for not buying the predictions (if they did not). The results provide interesting insights into people's purchasing decisions.

[11] However, for the sake of completeness, similar conclusions are obtained from the pooled data analysis (combined first-round and non-first-round buyers). See Powdthavee and Riyanto (2012) for the pooled data estimates.

[12] Similarly, one could also imagine investors in the stock market thinking along the same lines about financial experts: "*If stock prices really follow random walk, then why do financial experts exist? And not only do they exist, they also get paid incredible sums of money to comment on something that is i.i.d. in nature*." And then when we see a streak of successful performances by these financial experts, we readily use it as justification for their existence and for the large paychecks they typically receive.

Participants whose predictions were of mixed successes gave reasons for *not buying* that are mostly rational, e.g., "*Coin tosses are random*," "*For this experiment, the outcomes are not predictable*," "*I believe in my own choices*," "*Predictions are as good as my own guess*," "*Predictions are random and a waste of money*." For a larger proportion of participants whose predictions had been correct *four* times in a row, their reasons for *buying* are mainly consistent with the belief in the hot hand of an invisible agent, rather than explicitly state that the experimenters were up to some trickery, e.g., "*Predictions had been more successful than placing own bets*," "*The past predictions were correct*," "*Based on the accuracy of past predictions*." Participants whose predictions had been correct up to round $j-1$ but not in j gave reasons for *buying* that are a mixture of *ex post* rationalizations, excuses, and sometimes regrets, e.g., "*I could not think of an outcome to bet on my own*," "*Out of curiosity*," "*I am stupid to buy it*."

6.3.Do buyers treat paid-for predictions seriously?

One way of inferring whether buyers treated the paid-for predictions seriously is to examine the amount of a bet placed by buyers compared to non-buyers. The hypothesis is that people who bought the predictions will feel – through the process of rationalization – more confident about future outcomes and subsequently place higher bets than non-buyers, i.e., those who still believe that coin flips are i.i.d. and therefore systematically unpredictable. We formally test this hypothesis in Table 5 by running, for each round j, a regression equation in which the dependent variable is the log of the endowment amount used to bet in round j, i.e., Eq. (10).

On average, the bet amount placed by buyers is between 39 and 76 percentage points higher than non-buyers in the final three rounds. The estimated effects are statistically significant at conventional levels and are robust to controlling for the current endowment level and for gender and nationality. These results provide strong evidence – of both statistical and economic significance – that buyers place a significant level of trust on the "hot" envelopes rather than buy them to satisfy their own curiosity or simply for fun.

7. Discussions and concluding remarks

An experimental game in which people guess and bet on the outcomes of "fair" coin tosses is ostensibly simple. Yet behind its simplicity lies its unique strength. Given that the coin is fair (in that it had been proven fair by various explicit processes), it should be universally irrefutable that the exact sequence of future coin tosses is systematically unpredictable.[13] Participants' predictions should therefore be influenced by the gambler's fallacy rather than by the random accuracy of the envelope's predictions. Such a setting provides us with a perfect experimental setting in which we could test the boundaries of people's beliefs in the folly of coin-flip predictions.

However, our findings take us and perhaps many others[14] by surprise. When randomized predictions are introduced into the game as a potential information good, it takes only a few past realizations of correct predictions for individuals to start forming the belief that transparently i.i.d. outcomes are systematically predictable and that the prefixed envelopes contain information worth paying for. The influence of such "bogus" predictions is sufficiently large to offset the gambler's fallacy and induce buyers to place higher bets on average than non-buyers.

[13] This is unlike the belief that past returns do not predict future returns in the stock market, the truth of which many people spend years debating and researching. Prior beliefs about the predictability of randomly selected coin flips should be more absolute and universal.

[14] This comment is made in part to reflect how the earlier draft of this paper has been surprisingly received in the media and by the general public. Since its online publication as an IZA discussion paper in May 2012 (Powdthavee & Riyanto, 2012), our results have been discussed at great length – despite having had no press release – on the Freakonomics blog www.freakonomics.com/blog/ ("Paying for 'transparently useless advice'," 6 June 2012), *New Statesman* ("If you've got lucky, it's easy to convince people you're a sage," 6 June 2012), *The Economist* ("Buttonwood: Not so expert," 9 June 2012), *The Financial Times* ("Heads or tails? Just don't bet on it," 15 June 2012), *The Wall Street Journal* ("People will pay for 'transparently useless advice' about chance events," 18 June 2012), among others. Even we, the experimenters, nearly gave up with the idea of this project before we had even started simply because we had difficulty believing anybody would be willing to pay for predictions of obviously fair coin flips.

Indeed, we believe that our laboratory experiment provides results that are essentially predictable by the dynamic-inference model by Rabin and Vayanos (2010). Our contribution, however, is that these results are obtained even without the key assumption – i.e., that individuals are *a priori* uncertain about the underlying process generating the data – needing to hold. In other words, our experiment establishes a new lower bound of how people's beliefs in transparently useless predictions can be formed, i.e., participants are not entirely protected from "irrational learning" even when, at the beginning of the experiment, these individuals had started with a very strong prior about the ability of the predicting agent. Our results call for the need to look outside the economic discipline for an alternative explanation of such puzzling findings.

One potential explanation is in the work of the psychologist Justin Barrett. In his Hypersensitive Agency Detection Device theory, humans are hardwired to detect patterns in otherwise unrelated events, details that defy straightforward explanations, or consequences that seem out of proportion to the alleged cause. Such sensitiveness to detecting agents even when none exists has several evolutionary advantages. For example, spotting and understanding other agents could have been key to survival for early man and thus continuing to reproduce: It is far better to avoid several imaginary predators than be eaten by a real one. Since there is such a high cost of failing to detect agents and a low cost of wrongly detecting them, evolution will select an inheritable tendency to over-detect agents, even when we do not see them, as a survival strategy (Barrett, 2004; Grey & Wegner, 2010). The way we set up our experiment, which randomly allowed some participants to experience implausible streaks of accurate predictions, may have helped to trigger the hypersensitive agency detection device for these individuals, thus leading them to believe in the hot hand of an invisible agent when there was none.

Our results also open up a new discussion that perhaps people buy expert predictions (which, remember, are transparently useless in our case) not for their predictability but rather for other psychological reasons.[15] For example, our participants may have bought the predictions to:

- **Delegate decision-making**. Participants know that their choice is no better than a random envelope's choice. Nevertheless, if the decision is wrong, at least they can blame the envelope and not themselves.
- **Avoid regret**. Participants know that they will learn the information in the envelope at the end of each coin flip for free. However, they may not want to regret not having that information if it turns out to be correct. Since regret is powerful, they may instead want to pay the small fee to avoid it.
- **Feel in control of the situation**. Participants may buy the predictions simply to feel in control of a situation in which they have no control over the outcomes. Like regrets, the ability to feel in control of an uncontrollable situation is a powerful emotion, and participants may be willing to pay a small fee to be in the possession of it (see, e.g., Langer, 1975).
- **Psychic hedge**. It may be the case that participants formally go with the predictions in the envelope even when they do not formally believe that it has any predictive power. This way, regardless of what happens, they win (at least psychically).

Equally interesting are the potential implications of our findings across social-science disciplines. What we have learned, other than that people are not completely immune to irrational learning, is that it is unimaginably easy and costs almost nothing for "experts" to manufacture fallacious beliefs that even truly i.i.d. events are predictable. This is primarily because, according to Philip Tetlock (2006), experts are not punished sufficiently when they

[15] We are especially grateful to Dan Houser for giving us advice on likely alternative psychological explanations of our results.

are wrong, nor will they admit to being wrong when they are. On the contrary, the benefit of predicting unpredictable events correctly, even if rarely, significantly outweighs the cost of not getting them right. Since experts are usually paid for their services, paid better when they give "good" advice, and are not severely punished when incorrect ones are made, we are essentially rewarding bad judgments over good ones at no cost to the overall demand for expert advice.

Hence, with the existing system of expertise, financial firms could in principle provide their customers with various mutual funds containing randomly chosen stocks, and then build their advertisements around those that outperform the market by chance. Since people also buy transparently useless information for many psychological reasons other than for their predictability, people's demands for financial advice are likely to be inelastic with respect to the prices that financial firms typically charge their clients for their services.[16] This is the case even when it may be clear to everyone involved that market evaluations revert to the mean over the long term and that such advice is not needed very often in reality. Moreover, economists and political forecasters can hand out random predictions about future economic and political crises, and would be tantalizingly rewarded when they are right and yet would not be held accountable for misprediction when they are wrong.

Our experiment highlights these flaws in the current system of expertise and argues that it may not be sufficient to leave it to individuals to judge for themselves whether an expensive expertise is worth paying for. This is the case even in situations where it should be transparent to everyone that outcomes are dominated by pure randomness. Our results also underscore the problem that years of using statistics cannot offset the erroneous intuitions we sometimes have about the hot hand of an expert after streaks of accurate predictions are observed.

[16] A good reference for this comes in the review by Inderst and Ottaviani (2012). According to a recent survey of the market, the majority of US retail investors (approximately 73%) consult a financial adviser before purchasing shares.

This raises an important question: If teaching people about statistics do not help to minimize these mistakes, how can we create a mechanism that will? One potential solution to this problem may come in the form of prediction markets (see, e.g., Wolfer & Zitzewitz, 2006; Arrow et al., 2008). Prediction markets, or information markets, allow participants to generate contracts based on predictions of outcomes of future events and then trade them in the market at a price. A typical example of a contract in the prediction market is a forecast that candidate X will win the next presidential election, which could be traded in the market at, for example, US$1. If the market price of such a contract is currently 40 cents, an interpretation is that the market "believes" candidate X has a 40% chance of winning the election. Provided that the market is efficient, then the price of the contract perfectly aggregates dispersed information about the probability of candidate X being elected.

The promise of prediction markets is that, owing to the law of large numbers, the market prices will often produce aggregated forecasts of event outcomes that have lower prediction errors than predictions generated by a single or a few "experts." Participants in the prediction markets also have a clear financial incentive for truthful revelation of true beliefs as losses will be made if predictions turn out to be wrong. More importantly, prices in the prediction market can also be used as an aggregated indicator of whether the event in question is random or predictable – i.e., aggregate prices should be lower for events that are significantly more random in nature – which may help to improve investment decisions for individuals looking to invest.

We began this paper by noting a divide between the economics and psychology literature. Our experimental results seem to provide novel evidence in favor of psychological explanations for the apparent demands for useless predictions. It appears that economists may need to readjust their prior about where the lower bound – between rational and irrational beliefs – actually lies. Future research may need to return to this to construct a general quantitative model of such “irrational learning” that can be applied across a wide variety of settings, and to determine cost-effective ways (other than the promise of prediction markets) to help minimize these mistakes in the market.

References

Arrow, Kenneth J., and 21 other colleagues. 2008. The promise of prediction markets. *Science*, 877-878.

Asparouhova, Elena, Michael Hertzel and Michael Lemmon. 2009. Inference from streaks in random outcomes: experimental evidence on beliefs in regime shifting and the law of small numbers. *Management Science*, 55(11), 1766-1782.

Ayton, Peter and Ilan Fischer. 2004. The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Memory and Cognition*, 32(8), 1369-1378.

Barrett, Justin L. 2004. *Why would anyone believe in God?* Walnut Creek, CA: AltaMira Press.

Burns, Bruce D. and Bryan Corpus. 2004. Randomness and inductions from streaks: "gambler's fallacy" versus "hot hand". *Psychonomic Bulletin and Review*, 11(1), 179-184.

Cahart, Mark. 1997. On persistence in mutual fund performance. *Journal of Finance*, 52, 57-82.

Chater, Nick, Steffen Huck, and Roman Inderst. 2010. Consumer Decision-Making in Retail Investment Services: A Behavioural Economics Perspective. Report to the European Commission/SANCO.

Clotfelter, Charles and Phil Cook. 1991. Lotteries in the real world. *Journal of Risk and Uncertainty*, 4, 227-232.

Clotfelter, Charles and Phil Cook. 1993. The gambler's fallacy in lottery play. *Management Science*, 39,

Croson, Rachel and James Sundali. 2005. The gambler's fallacy and the hot hand: empirical data from casinos. *Journal of Risk and Uncertainty*, 30(3), 195-209.

Fama, Eugene F. 1991. Efficient capital markets: II. *Journal of Finance*, 46, 1575-1617.

Fischbacher, Urs. 2007. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.

Gilovich, Thomas, Robert Vallone and Amos Tversky. 1985. The hot hand in basketball: on the misperception of random sequences. *Cognitive Psychology*, 17, 295-314.

Gray, Kurt and Daniel M. Wegner. 2010. Blaming God for our pain: human suffering and the divine mind. *Personality and Social Psychology Review*, 14(1), 7-16.

Guryan, Jonathan and Melissa S. Kearney. 2008. Gambling at lucky stores: empirical evidence from state lottery sales. *American Economic Review*, 98(1), 458-473.

Huber, Jurgen, Michael Kirchler, and Thomas Stockl. 2010. The hot hand belief and the gambler's fallacy in investment decisions under risk. *Theory and Decision,* 68(4), 445-462.

Hung, Angela A., Noreen Clancy, Jeff Dominitz, EricTalley, Claude Berrebi, and Farrukh Suvankulov. 2008. Investor and Industry Perspectives on Investment Advisers and Broker-Dealers. RAND Institute for Civil Justice Technical Report.

Inderst, Roman and Marco Ottaviani. 2012. Financial advice. *Journal of Economic Literature*, 50(2), 494-512.

Jørgensen, Claus Bjørn, Sigrid Suetens and Jean-Robert Tyran. 2011. Predicting lotto numbers. CEPR Discussion Paper 8314.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. Allen Lane: Penguin Books.

Langer, Ellen. 1975. The illusion of control. *Journal of Personality and Social Psychology*, 32(2), 311-328.

Powdthavee, Nattavudh, and Yohanes E. Riyanto. 2012. Why do people pay for useless information? Implications of gambler's and hot-hand fallacies in false expert setting. IZA Discussion Paper Series No. 6557.

Rabin, Matthew. 2002. Inferences by believers in the law of small numbers. *Quarterly Journal of Economics*, 117(3), 775-816.

Rabin, Matthew; and Dimitri Vayanos. 2010. The gambler's and hot hand fallacies: theories and applications. *Review of Economic Studies*, 77(2), 730-778.

Teller, Dek. 1994. A test of the gambler's fallacy: evidence from Pari-Mutuel games. *Journal of Risk and Uncertainty*, 17, 151-166.

Tetlock, Philip. 2006. *Expert Political Judgment: How Good is It? How Can We Know?* Princeton University Press.

Tversky, Amos and Daniel Kahneman. 1971. Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105-110.

Wolfers, Justin and Eric Zitzewitz. 2006. Prediction markets in theory and practice. NBER Working Paper 12083.

Figure 1: The coin tree

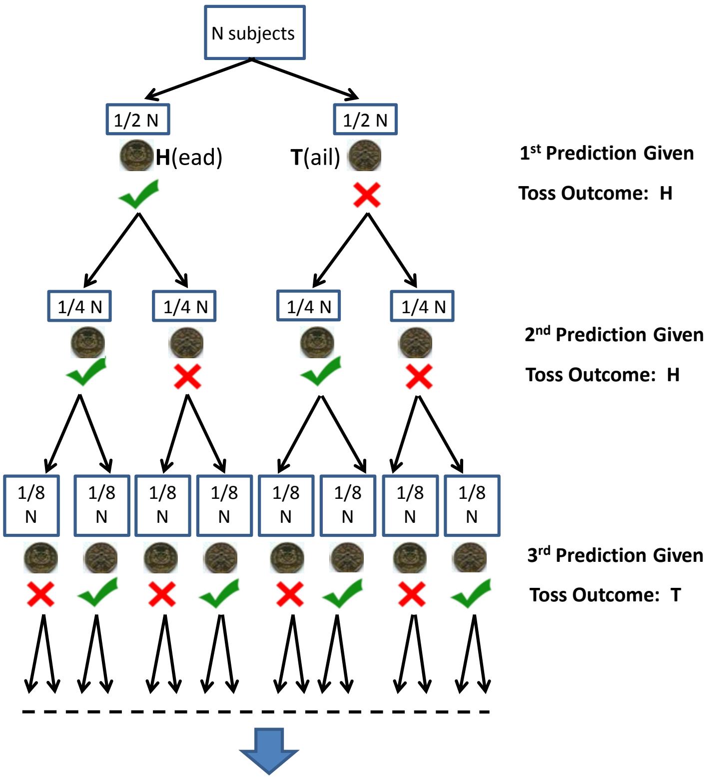

Table 1: Testing the gambler's fallacy in people's betting behaviors

Dependent variable: Betting head in round *j*	**Round *j* = 3** All	**Round *j* = 3** Prediction = H	**Round *j* = 3** Prediction = T	**Round *j* = 4** All	**Round *j* = 4** Prediction = H	**Round *j* = 4** Prediction = T
A streak of coin coming up heads up to *j*-1	-0.324***	-0.283***	-0.425***	-0.265**	-0.253	-0.333**
	[0.0696]	[0.102]	[0.1000]	[0.111]	[0.200]	[0.145]
A streak of coin coming up tails up to *j*-1	0.182***	0.165*	0.163*	0.117	0.187*	0.0789
	[0.0609]	[0.0882]	[0.0900]	[0.0768]	[0.101]	[0.118]
Bought the prediction in round *j*		0.0284	-0.694***		0.0391	-0.405***
		[0.286]	[0.0780]		[0.179]	[0.137]
A streak of heads × bought prediction in *j*		0.616**			0.628**	0.133
		[0.295]			[0.265]	[0.194]
A streak of tails × bought prediction in *j*		0.168	0.0873			-0.279
		[0.291]	[0.238]			[0.175]
Constant	0.640***	0.638***	0.694***	0.581***	0.586***	0.605***
	[0.0511]	[0.0714]	[0.0780]	[0.0307]	[0.0442]	[0.0467]
Implied net effects on betting head in round *j*						
Streak of heads up to j-1 and bought prediction		*0.333*	-		*0.375***	*-0.200*
		[0.277]			*[0.173]*	*[0.128]*
Streak of tails up to j-1 and bought prediction		*0.333*	*0.250*		-	*-0.200*
		[0.277]	*[0.220]*			*[0.128]*
Observations	322	164	158	322	165	157
R-squared	0.192	0.167	0.308	0.025	0.032	0.095

Note: *<10%; **<5%; ***<1%.

All regressions are estimated using a linear probability model and conditioning on non-prediction-buyers in the first round.

Figure 2: Proportion of prediction buyers by types of prediction streaks

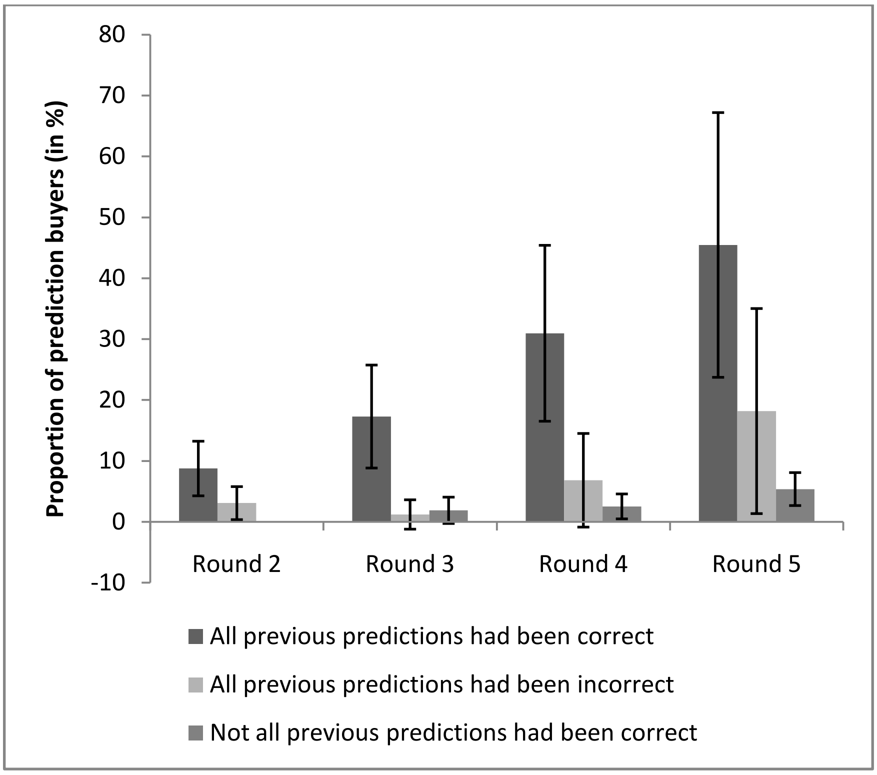

Note: These are raw means and conditioning on non-prediction-buyers in the first round. 2-standard-error bands (95% C.I.) are reported, i.e. 2 above and 2 below se.

Figures 3A-3B: Proportion of prediction buyers by types of prediction streaks and by nationality

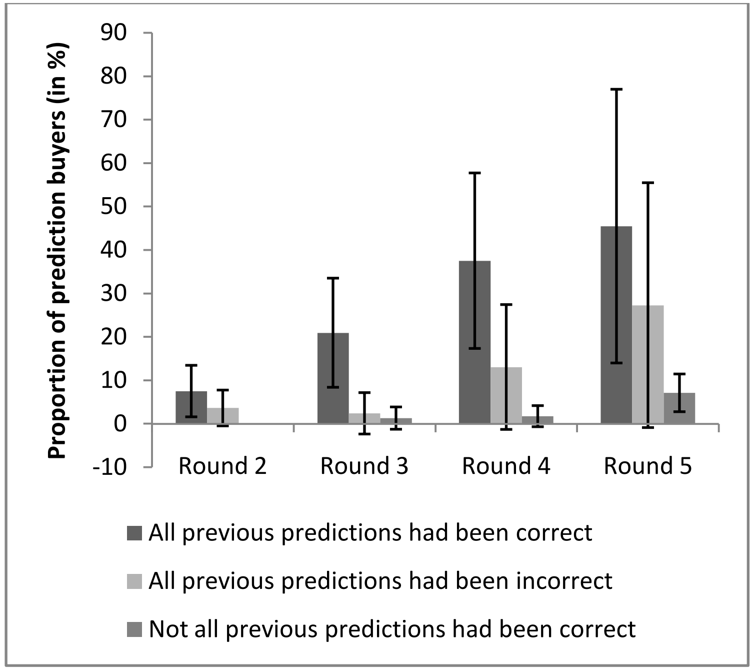

Figure 3A: *Thai sample*

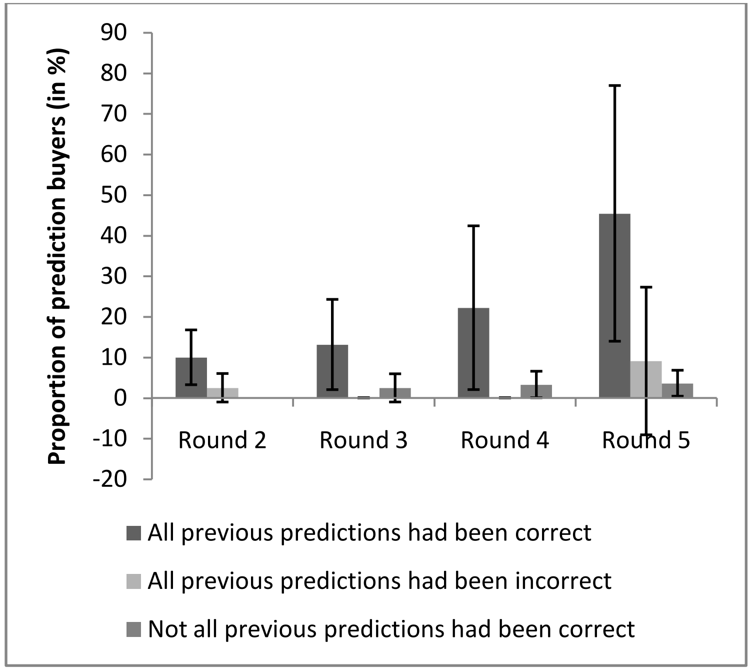

Figure 3B*: Singaporean sample*

Note: These are raw means and conditioning on non-prediction-buyers in the first round. 2-standard-error bands (95% C.I.) are reported, i.e., 2 above and 2 below se.

.

Table 2: Linear probability model estimates of factors determining the decision to buy prediction in each round

Dependent variable: Buy prediction in round *j*	***j*=2**	***j*=3**	***j*=4**	***j*=5**
a) Without controls				
All predictions up to *j*-1 had been correct	0.0568**	0.153***	0.265***	0.341***
	[0.0262]	[0.0434]	[0.0688]	[0.102]
All predictions up to *j*-1 had been incorrect		-0.00242	0.0458	0.144*
		[0.0160]	[0.0400]	[0.0837]
Bought prediction in round *j*-1		0.0976	0.153*	0.278***
		[0.0811]	[0.0839]	[0.0977]
Constant	0.0307**	0.0121	0.0188*	0.0378***
	[0.0135]	[0.0106]	[0.0111]	[0.0126]
Observations	323	323	323	323
R-squared	0.015	0.097	0.159	0.189
b) With controls				
All predictions up to *j*-1 had been correct	0.0539**	0.146***	0.214***	0.265**
	[0.0262]	[0.0425]	[0.0729]	[0.103]
All predictions up to *j*-1 had been incorrect		0.00284	0.0408	0.161**
		[0.0168]	[0.0379]	[0.0807]
Bought prediction in round *j*-1		0.104	0.412**	0.510***
		[0.130]	[0.165]	[0.148]
Male	-0.0154	0.0119	-0.0123	-0.0214
	[0.0257]	[0.0245]	[0.0246]	[0.0265]
Singaporean	0.0313	0.0349	-0.00718	-0.0496
	[0.0625]	[0.0509]	[0.0365]	[0.0390]
Proportion of correct answers in statistical test	-0.0135	-0.0861	-0.0233	-0.107*
	[0.0500]	[0.0638]	[0.0507]	[0.0647]
Endowment in *j*th round	-0.000145	-0.000138	-0.000149	-0.000175*
	[0.000230]	[0.000143]	[0.000149]	[9.03e-05]

A streak of coin coming up heads up to *j*-1		-0.0561	0.0537	
		[0.0482]	[0.0731]	
A streak of coin coming up tails up to *j*-1		0.0108	-0.0428	-0.124***
		[0.0362]	[0.0428]	[0.0385]
Made wrong bet in *j*-1	0.0115	-0.00653	0.0249	-0.0704**
	[0.0331]	[0.0285]	[0.0333]	[0.0315]
Bought prediction in *j*-1 x Made wrong bet in *j*-1		-0.0217	-0.457***	-0.491***
		[0.162]	[0.166]	[0.175]
Constant	0.0523	0.0990	0.0748	0.243***
	[0.0521]	[0.0660]	[0.0521]	[0.0719]
Observations	323	323	323	323
R-squared	0.024	0.122	0.212	0.289

Note: *<10%; **<5%; ***<1%.

All regressions are estimated using a linear probability model and conditioning on non-prediction-buyers in the first round. Dependent variable is a binary variable that takes a value of 1 if the subject paid to see the prediction in the corresponding numbered envelope, and 0 otherwise. Reference groups include: female; Thai students; no streak up to *j*-1; and made correct bet in *j*-1.

Table 3: The gambler's fallacy in first-round buyers, OLS regressions

	Round *j* = 3		
Dependent variable: Betting head in round *j*	**All**	**Prediction = H**	**Prediction = T**
A streak of coin coming up heads up to *j*-1	-0.342**	-0.500**	-0.338
	[0.144]	[0.179]	[0.229]
A streak of coin coming up tails up to *j*-1	-0.0921	-0.167	-0.131
	[0.141]	[0.172]	[0.181]
Bought the prediction in round *j*		-1.03E-15	-0.576*
		[1.18e-08]	[0.317]
A streak of heads x bought prediction in *j*		-1.31E-15	0.00433
		[0.438]	[0.379]
A streak of tails x bought prediction in *j*			-0.202
			[0.352]
Constant	0.842***	1.000***	0.909***
	[0.0860]	[3.73e-09]	[0.0962]
Observations	55	23	32
R-squared	0.104	0.213	0.295

Note: *<10%; **<5%; ***<1%. See Table 1.

Table 4: Buying behaviors of first-round buyers, OLS regressions

Dependent variable: Buy prediction in round *j*	***j*=2**	***j*=3**
All predictions up to *j*-1 had been correct	0.165	0.295**
	[0.108]	[0.142]
All predictions up to *j*-1 had been incorrect		0.0356
		[0.0942]
Bought prediction in round *j*-1		0.260*
		[0.131]
Constant	0.125*	0.0213
	[0.0688]	[0.0489]
Observations	55	55
R-squared	0.039	0.185

Note: *<10%; **<5%. See Table 2.

Table 5: OLS estimates of the log of bet amount placed in each round

Dependent variable: **ln(bet amount in round *j*)**	***j=2***	***j=3***	***j=4***	***j=5***
Bought prediction in round *j*	-0.0262	0.507**	0.391*	0.760***
	[0.206]	[0.242]	[0.224]	[0.150]
Male	0.0233	0.105	-0.0390	0.195**
	[0.0768]	[0.0815]	[0.0833]	[0.0980]
Singaporean	0.476***	0.169	0.314**	0.0582
	[0.172]	[0.145]	[0.125]	[0.135]
Proportion of correct answers in statistical test	0.251	0.114	-0.00115	0.143
	[0.169]	[0.182]	[0.169]	[0.184]
Endowment in round *j*	-0.00104	0.00108**	1.52e-05	0.000257
	[0.000712]	[0.000508]	[0.000426]	[0.000440]
Constant	1.982***	1.260***	1.493***	1.331***
	[0.231]	[0.233]	[0.206]	[0.248]
Observations	321	320	321	319
R-squared	0.178	0.255	0.329	0.283

Note: *<10%; **<5%; ***<1%.

Dependent variable is log of the bet amount placed in round *j*. Robust standard errors are reported in parentheses.

Appendix A: Descriptive statistics

	Non-first-round buyers			First-round buyers
	All	**Thais**	**Singaporeans**	**All**
Round *j* = 2				
Correct predictions up to *j*-1	160	80	80	31
% of all predictions	50.56%	49.08%	50.00%	56.36%
Bought prediction in *j* = 2	19	6	8	9
% of all correct predictions up to *j*-1	**8.75%**	**7.50%**	**10.00%**	**29.03%**
Round *j* = 3				
Correct predictions up to *j*-1	81	43	38	11
% of all predictions	25.08%	26.38%	23.75%	20.00%
Bought prediction in *j* = 3	14	9	5	4
% of all correct predictions up to *j*-1	**17.28%**	**20.93%**	**13.16%**	**36.36%**
Round *j* = 4				
Correct predictions up to *j*-1	42	24	18	6
% of all predictions	13.00%	14.72%	11.25%	10.91%
Bought prediction in *j* = 4	13	9	4	1
% of all correct predictions up to *j*-1	**30.95%**	**37.50%**	**22.22%**	**16.67%**
Round *j* = 5				
Correct predictions up to *j*-1	22	11	11	1
% of all predictions	6.81%	6.75%	6.88%	1.82%
Bought prediction in *j* = 5	10	5	5	0
% of all correct predictions up to *j*-1	**45.45%**	**45.45%**	**45.45%**	**0.00%**
N	**323**	**163**	**160**	**55**

Note: Total N=378 (Thai: N = 177; Singaporean: N=201).

Appendix B: Marginal effects probit model

Dependent variable: Buy prediction in round *j*	***j=2***	***j=3***	***j=4***	***j=5***
a) Without control variables				
All predictions up to j-1 had been correct	0.0568**	0.145***	0.271***	0.335***
	[0.0261]	[0.0502]	[0.0735]	[0.108]
All predictions up to j-1 had been incorrect		-0.00735	0.0605	0.172*
		[0.0284]	[0.0533]	[0.0966]
Bought prediction in round j-1		0.0792	0.100	0.271***
		[0.0719]	[0.0645]	[0.104]
Observations	323	323	323	323
Pseudo R-squared	0.0338	0.1892	0.209	0.2016
b) With control variables				
All predictions up to j-1 had been correct	0.0517**	0.116**	0.218***	0.426***
	[0.0250]	[0.0491]	[0.0776]	[0.138]
All predictions up to j-1 had been incorrect		0.000488	0.0492	0.299**
		[0.0180]	[0.0441]	[0.131]
Bought prediction in round j-1		0.0386	0.258	0.660**
		[0.0528]	[0.162]	[0.278]
Male	-0.0122	0.0127	-0.00105	-0.0299
	[0.0225]	[0.0136]	[0.0192]	[0.0250]
Singaporean	0.0275	0.0192	0.00623	-0.0291
	[0.0573]	[0.0217]	[0.0313]	[0.0319]
Proportion of correct answers in statistical test	-0.0161	-0.0362*	-0.0103	-0.0970**
	[0.0468]	[0.0208]	[0.0361]	[0.0467]
Endowment in jth round	-0.000143	-0.000102	-0.000172	-0.000211**
	[0.000227]	[8.67e-05]	[0.000136]	[9.97e-05]
A streak of heads up to j-1		-0.0310**	0.0569	

		[0.0151]	[0.0780]	
A streak of tails up to j-1		0.000106	-0.0236	
		[0.0122]	[0.0231]	
Made wrong bet in j-1	0.0103	-0.00192	0.0163	-0.0921***
	[0.0318]	[0.0132]	[0.0267]	[0.0336]
Bought prediction in j-1 x Made wrong bet in j-1		0.0325		-0.0521***
		[0.0787]		[0.0143]
Observations	323	323	323	323
Pseudo R-squared	0.0514	0.2744	0.2523	0.3701

Note: *<10%; **<5%; ***<1%. See Table 2.

Appendix C: Experimental instructions and screenshots

REMARKS

Please be reminded that throughout this study:

1. You are not allowed to communicate with other participants at any point in the experiment. Any person caught violating this will be asked to leave the experiment without pay.
2. You are **not allowed to open any of the envelopes in front of you unless instructed to do so by the experimenter.**

General Information

Welcome to all of you! You are now taking part in an interactive study on decision making. **Please pay attention to the information provided here and make your decisions carefully. If at any time you have questions to ask, please raise your hand and we will attend to you in private.**

Your participation in this study is voluntary. You will receive **2 S$** show-up fee for participating in this study. You may decide to leave the study at any time. Unfortunately, if you withdraw before you complete the study, we can only pay you for the decisions that you have made up to the time of withdrawal, which could be substantially less than you will earn if you complete the entire study.

The amount of your earnings from this study depends on the decisions you and others make. At the end of this session, your earnings will be paid to you privately and in cash. It would be contained in an envelope (indicated with your unique user ID). You will need to sign a claim card given to you and exchange your claim card with your payment.

General Instructions

Each of you will be given **a unique user ID** and it **will be clearly stated on your computer screen**. At the end of the study, you will be asked to key in your user ID and other information pertaining to you earnings from this study in the claim card. **Please key in the correct user ID to make sure that you will get the correct amount of payment**.

Rest assured that your **anonymity will be preserved** throughout the study. You will **never be aware of** the personal identities of other participants **during or after** the study. Similarly, other participants will also **never be aware** of your personal identities **during or after** the study. You will only be identified by your user ID in our data collection. All information collected will **strictly be kept confidential** for the sole purpose of this study.

Specific Instructions

The total duration of this study is approximately **1 hour**. You will have to go through two stages.

STAGE 1

In stage 1, you will be asked to play a game of predicting the outcomes of coin tosses. The coin toss game will use a two-sided coin, i.e. head and tail (H-T). Coins will be borrowed from one of the volunteered participants and one will be randomly selected to be used in the coin toss. The game, which will last for 5 (five) rounds, will consist of tossing the coin in total **five (5) times**, one in every round. Participants will be asked **to make a bet on the outcome of each of the subsequent coin tosses**. Below is the screenshot of the welcome page of the experiment.

Screenshot 1

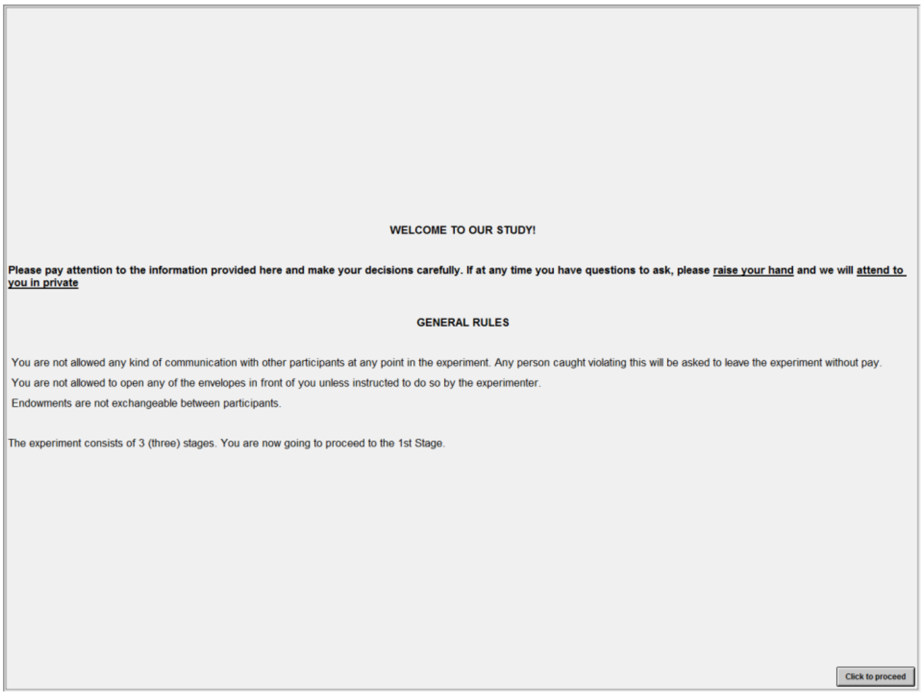

At the beginning of Round 1, you will be given **<u>300 endowment points</u>** which you could use to place a bet on the outcome of each subsequent coin tosses. There is a minimum bet of **<u>10 points per round.</u>**[17]

> **You are not allowed to go bankrupt before the final round of the coin toss, i.e. the 5th round.**

For every **1 point** you use to bet correctly, i.e. **your bet matches with the actual outcome of the coin toss**, your return will be **1 point $\times$ 2 = 2 points**. So, if you have an endowment of 300 points and you used 50 points to make a bet, in a case where you have placed a correct bet, you will earn $50 \times 2 = 100$ points. In terms of endowment that can be used for the next round of betting: 250 points remaining + 100 points from the bet = 350 points. Below is the screenshot.

<u>Screenshot 2</u>

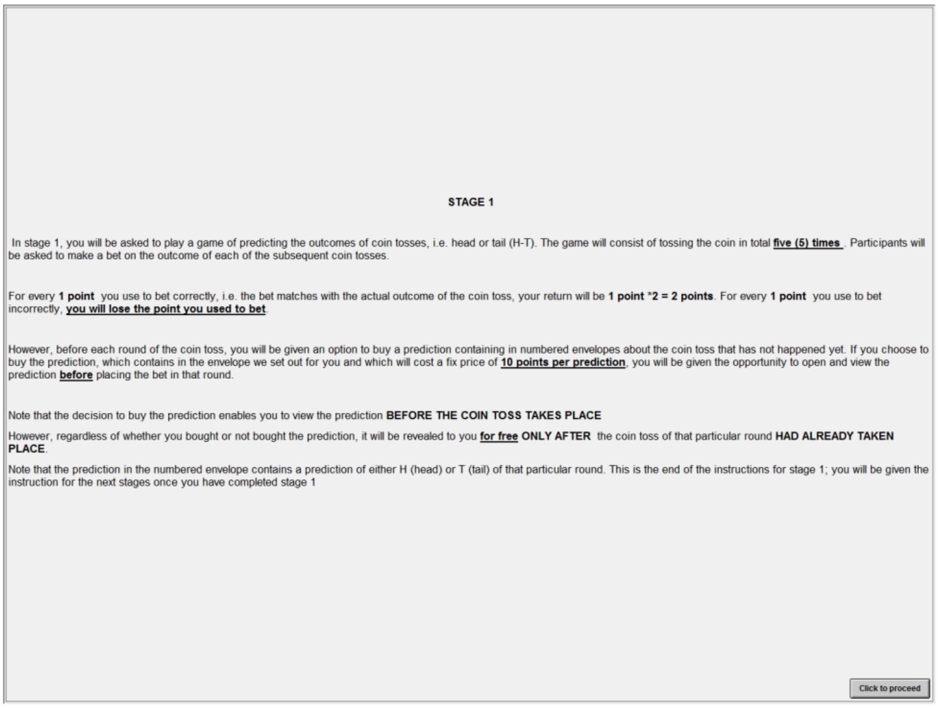

STAGE 1

In stage 1, you will be asked to play a game of predicting the outcomes of coin tosses, i.e. head or tail (H-T). The game will consist of tossing the coin in total **<u>five (5) times</u>** . Participants will be asked to make a bet on the outcome of each of the subsequent coin tosses.

For every **1 point** you use to bet correctly, i.e. the bet matches with the actual outcome of the coin toss, your return will be **1 point *2 = 2 points**. For every **1 point** you use to bet incorrectly, **<u>you will lose the point you used to bet</u>**.

However, before each round of the coin toss, you will be given an option to buy a prediction containing in numbered envelopes about the coin toss that has not happened yet. If you choose to buy the prediction, which contains in the envelope we set out for you and which will cost a fix price of **<u>10 points per prediction</u>**, you will be given the opportunity to open and view the prediction **<u>before</u>** placing the bet in that round.

Note that the decision to buy the prediction enables you to view the prediction **BEFORE THE COIN TOSS TAKES PLACE**

However, regardless of whether you bought or not bought the prediction, it will be revealed to you **<u>for free</u> ONLY AFTER** the coin toss of that particular round **HAD ALREADY TAKEN PLACE**.

Note that the prediction in the numbered envelope contains a prediction of either H (head) or T (tail) of that particular round. This is the end of the instructions for stage 1; you will be given the instruction for the next stages once you have completed stage 1

Click to proceed

For every **1 point** you use to bet incorrectly, **<u>you will lose the point you used to bet</u>**. So, if you have an endowment of 300 points and you used 50 points to make a bet, in a case where you have placed an incorrect bet, your endowment that can be used for the next round of betting will be 250 points.

[17] Note that only the Singaporean sample received 300 endowment points. Participants in the experiment in Thailand, which was ran a few months earlier, received only 100 endowment points each.

A photo of the numbered envelopes

Prediction

However, **before each round of the coin toss**, you will be given an option to buy a prediction containing in numbered envelopes about the coin toss that has not happened yet. The numbered envelope 1 will contain a prediction of the coin toss for round 1. The numbered envelope 2 will contain a prediction of the coin toss for round 2, and so on.

Screenshot 3

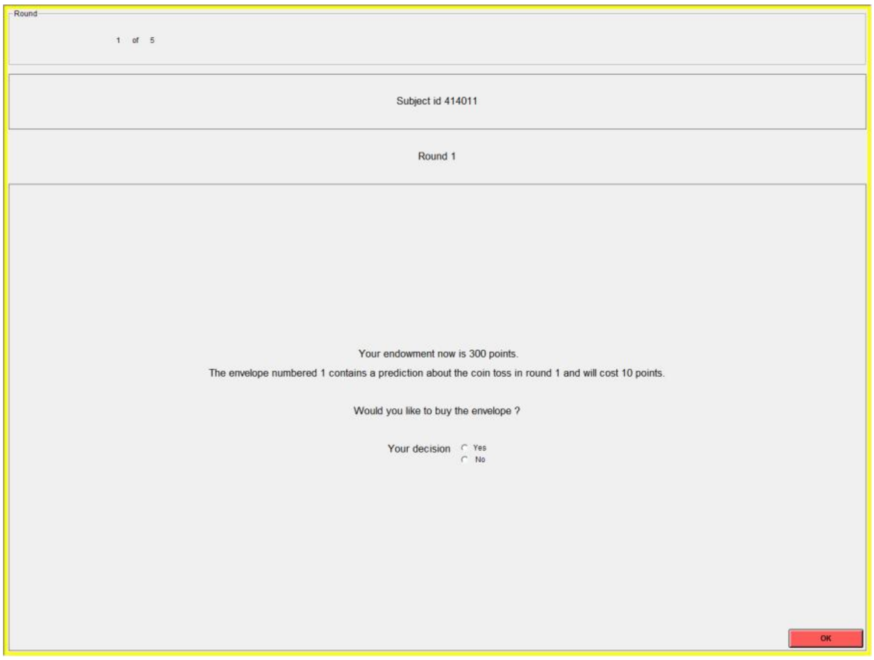

If you choose **to buy** the prediction, which contains in the envelope we set out for you and which will cost a fix price of **10 points per prediction**, you will be given the opportunity to **open** and **view** the expert's prediction **before** placing the bet in that round.

Screenshot 4

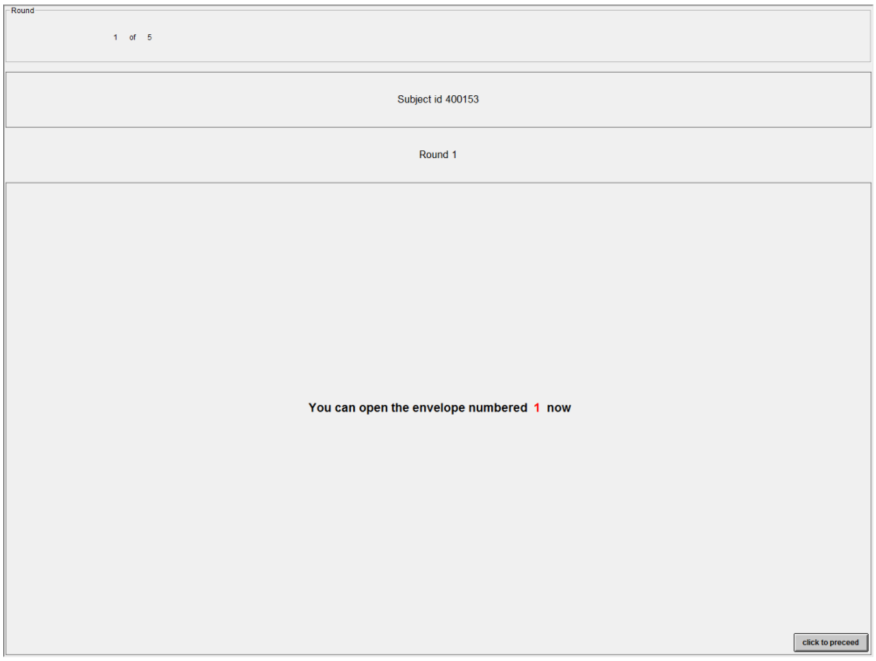

The prediction in the numbered envelope contains a prediction of either H (head) or T (tail) of that particular round.

If you decide **not to buy** the prediction we have provided for you, you will jump directly to the betting (Screenshot 5). Note that you will be given the opportunity to open the envelope **for free** and view the prediction **only after** the coin toss of that particular round had already taken place.

Once you have decided whether or not to buy the prediction, you will have to place your bet on the outcome of the coin toss that is about to be undertaken.

Betting

Below is the screenshot of the betting stage in Round 1 of the experiment.

Screenshot 5

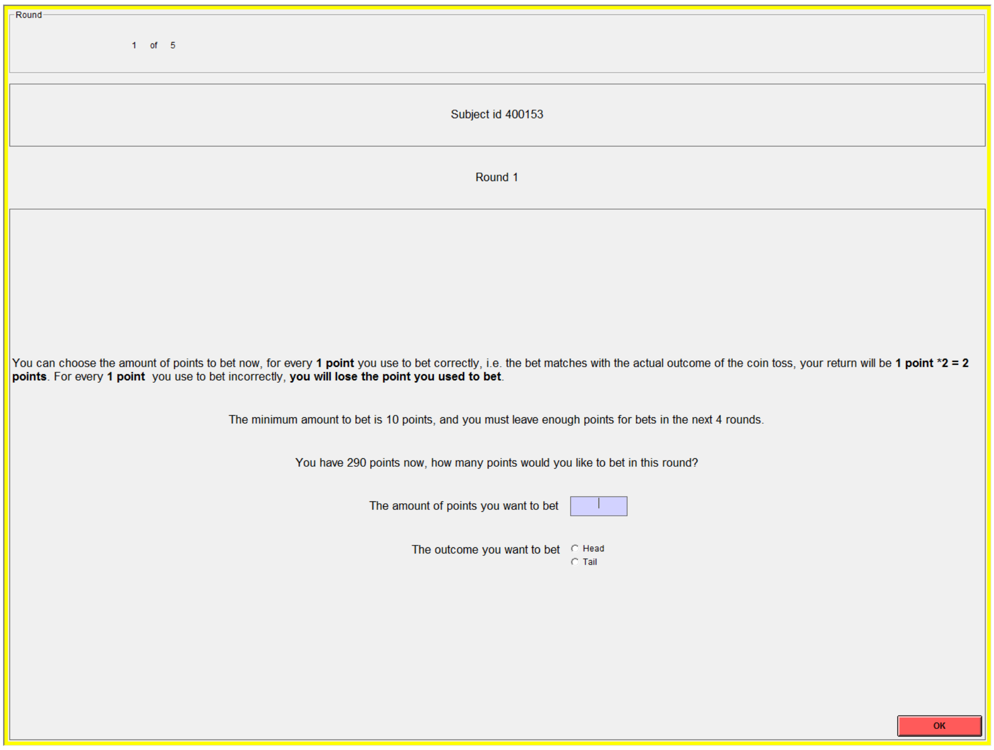

Once the betting is done, the coin is tossed.

Screenshot 6

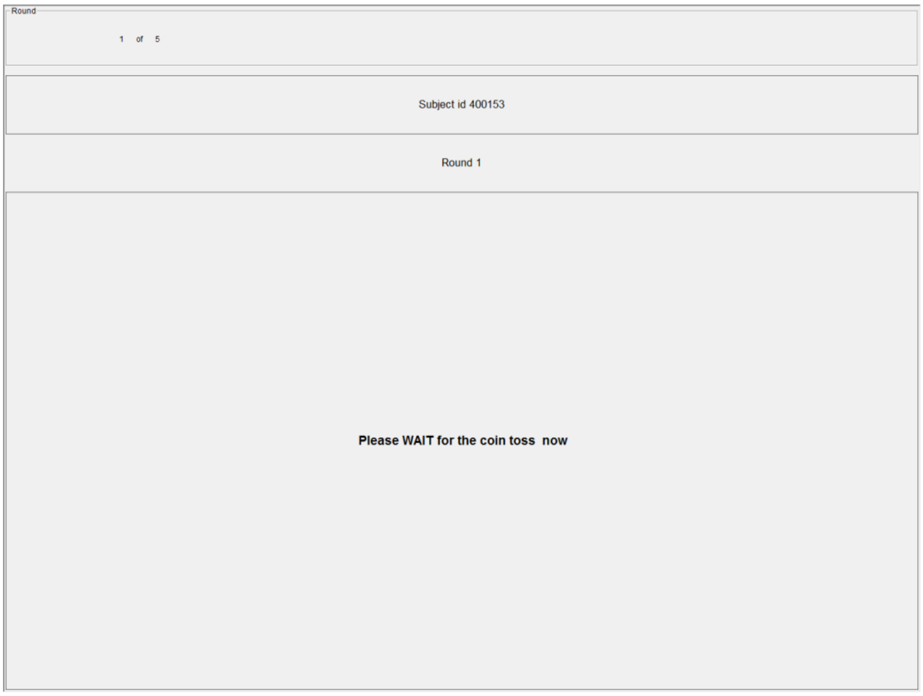

You will have to input the result of the coin toss. Likewise we also input the result of the coin toss, and use it to verify whether or not you have inputted the result of the coin toss correctly.

Screenshot 7

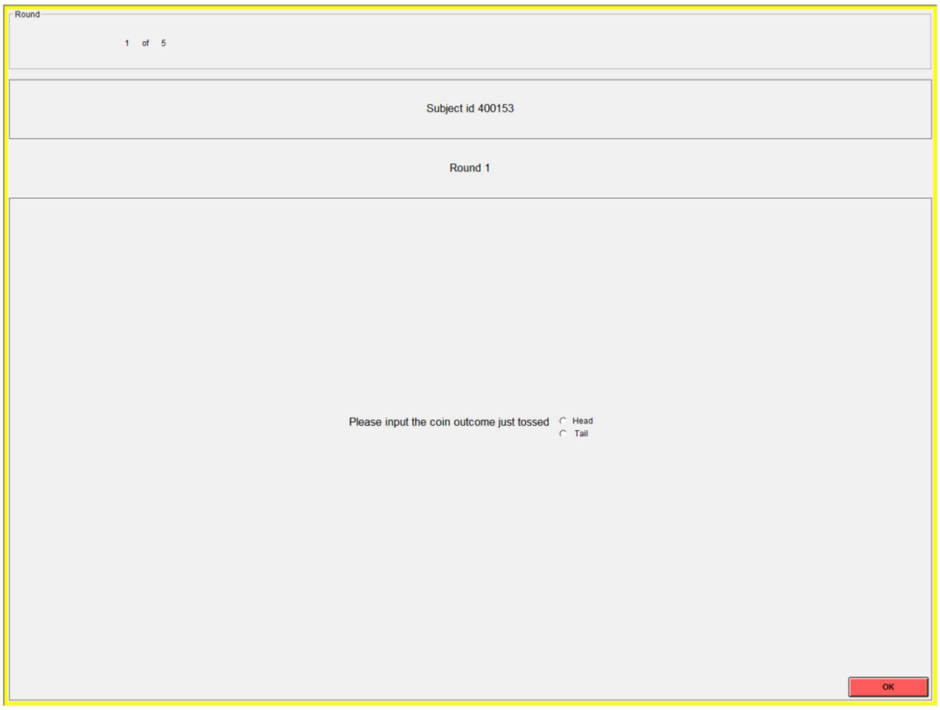

After that, the summary of result from the betting is shown.

Screenshot 8

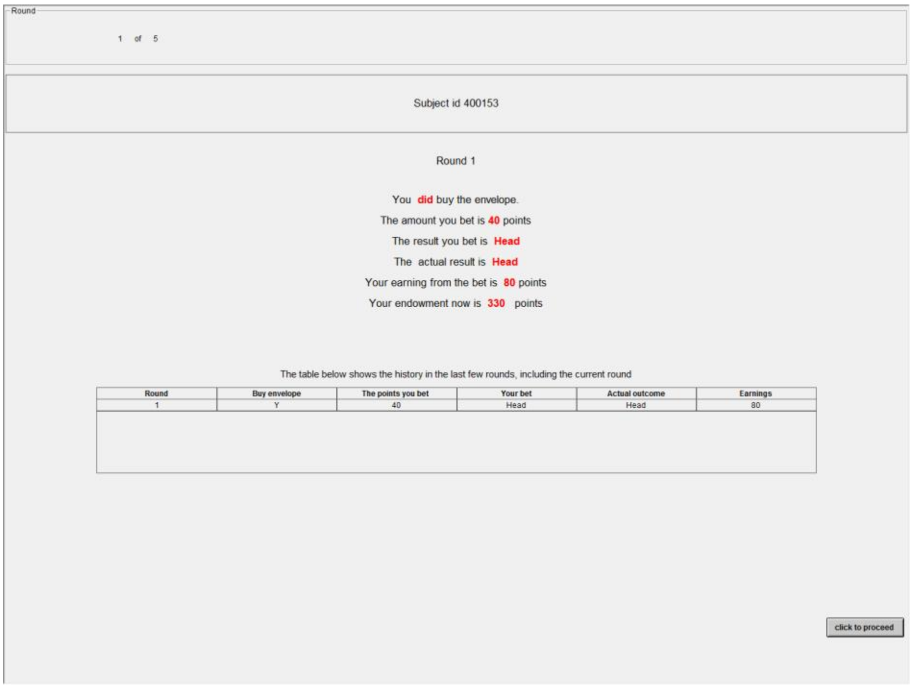

At this stage, those of you who decided not to buy the prediction can open the envelope numbered 1 to find out the prediction of the coin toss. Those of you who decided to buy the prediction can re-open again the envelope numbered 1 if you wish to do so. This is the end of Round 1 of the experiment.

After this, the procedures repeat again until all 5 (five) rounds of betting are completed. Below is the sample of the summary page at the end of Round 5.

<u>Screenshot 9</u>

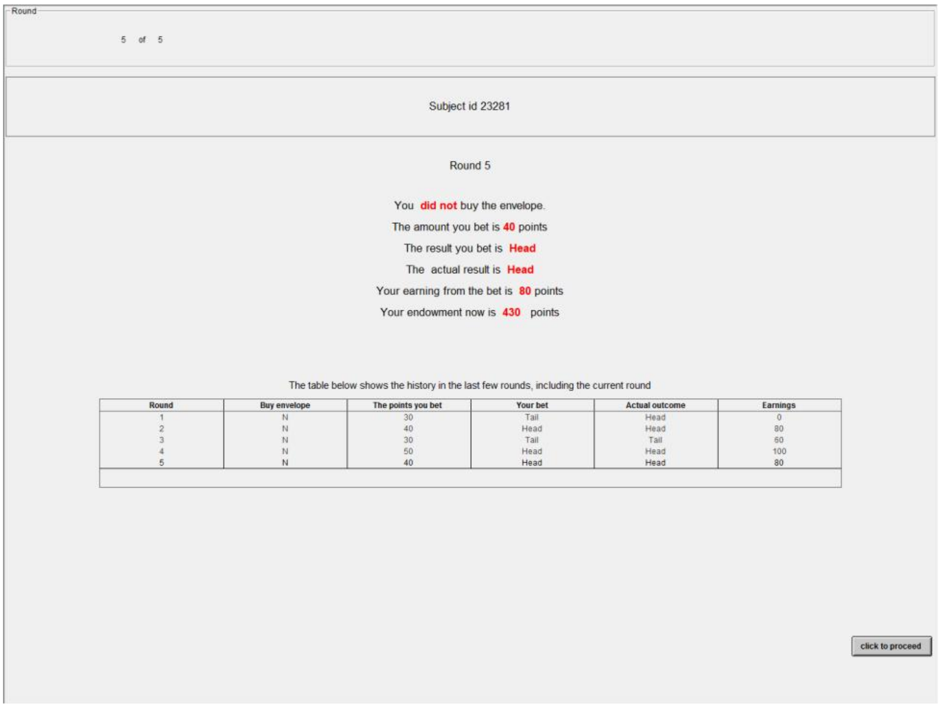

Round

5 of 5

Subject id 23281

Round 5

You **did not** buy the envelope.

The amount you bet is **40** points

The result you bet is **Head**

The actual result is **Head**

Your earning from the bet is **80** points

Your endowment now is **430** points

The table below shows the history in the last few rounds, including the current round

Round	Buy envelope	The points you bet	Your bet	Actual outcome	Earnings
1	N	30	Tail	Head	0
2	N	40	Head	Head	80
3	N	30	Tail	Tail	60
4	N	50	Head	Head	100
5	N	40	Head	Head	80

click to proceed

STAGE 2 (Probability Test)

In stage 2, you will be asked to complete **ten (10) questions on probability**. Please do your best to answer the questions correctly. You will be compensated according to the number of correct answers you made. For every correct answer you will get **+15** points. For every incorrect answer you will get **-15** points. If you leave a question unanswered, you will get **0** point for that particular question. Here it is the screenshot.

<u>Screenshot 9</u>

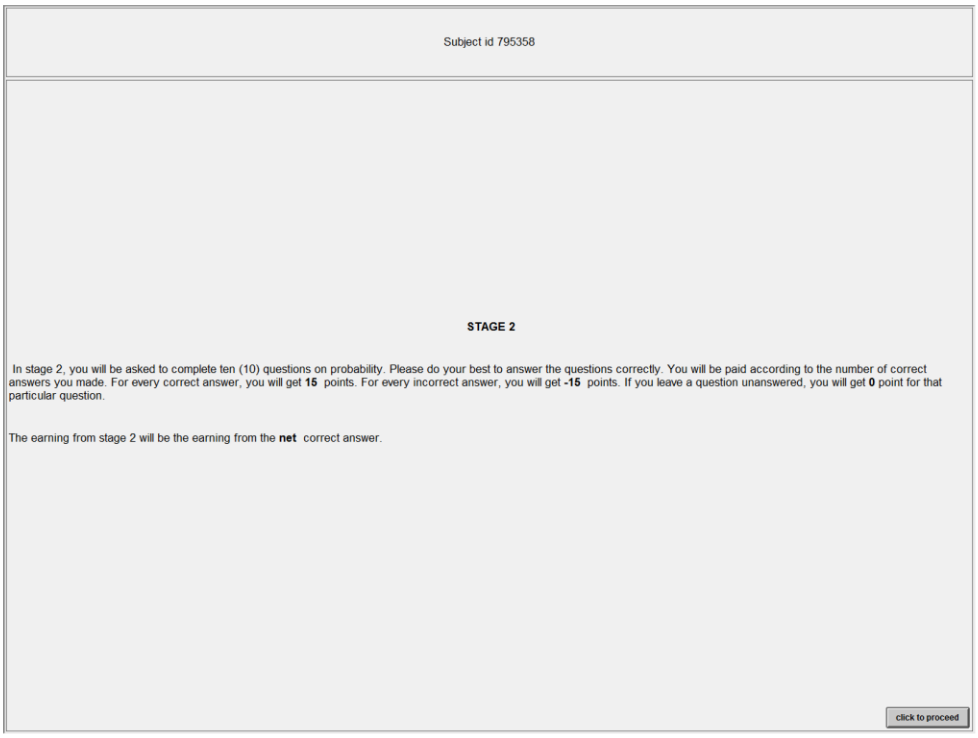

You will be given 8 (eight) minutes to complete all questions. Once the time is out and you have not finished answering all questions, the questions left unanswered will be given 0 point. Below are samples of these questions:

<u>Screenshot 10a</u>

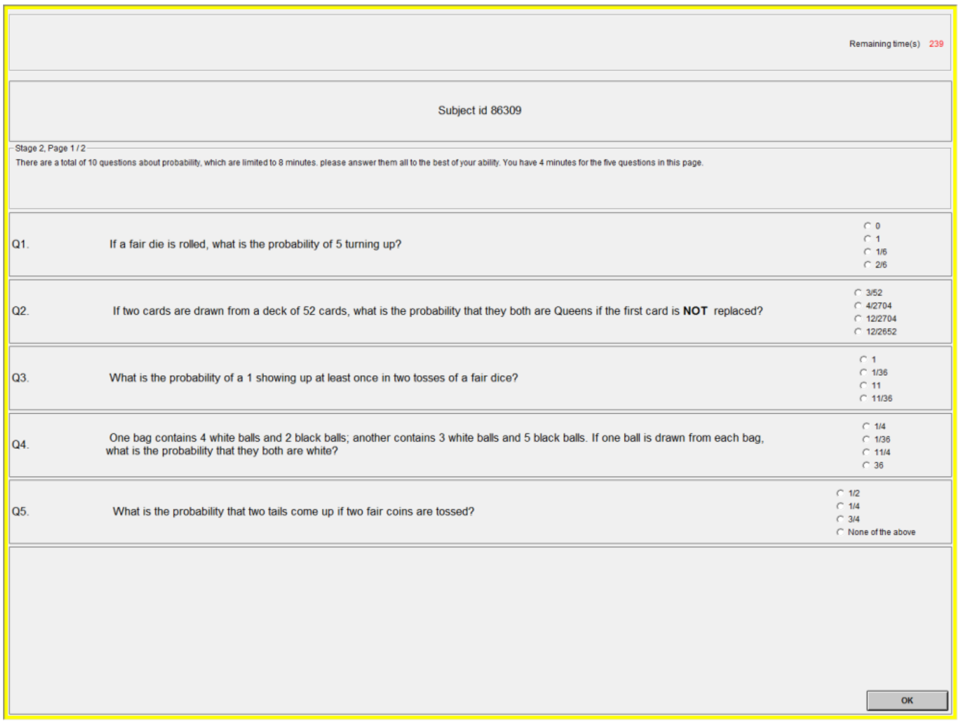

<u>Screenshot 10b</u>

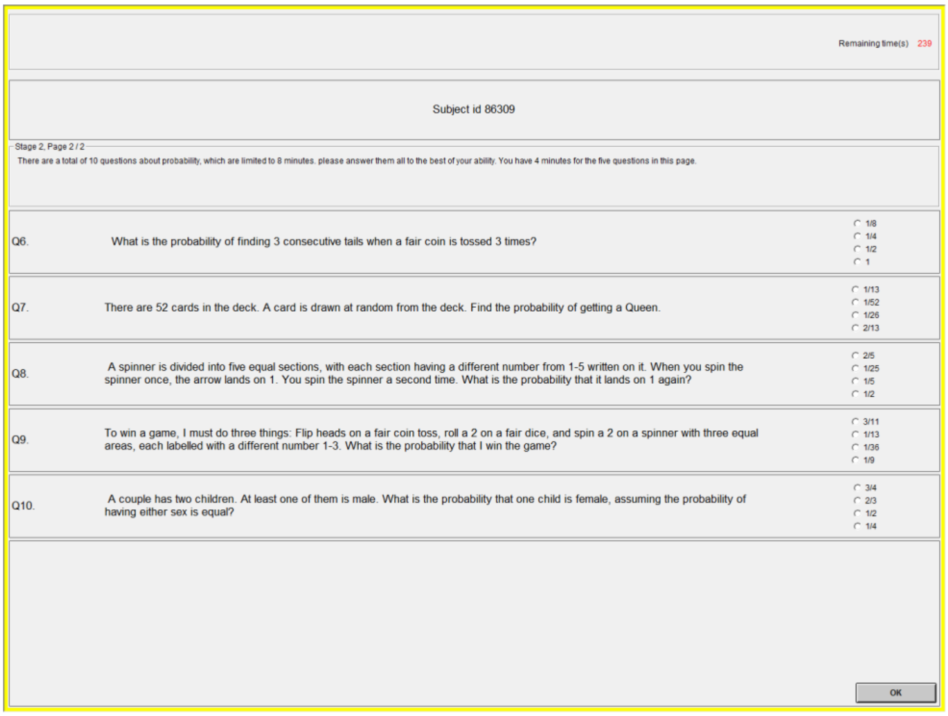

The following screenshot depicts the summary page shown at the end of Stage 2.

Screenshot 11

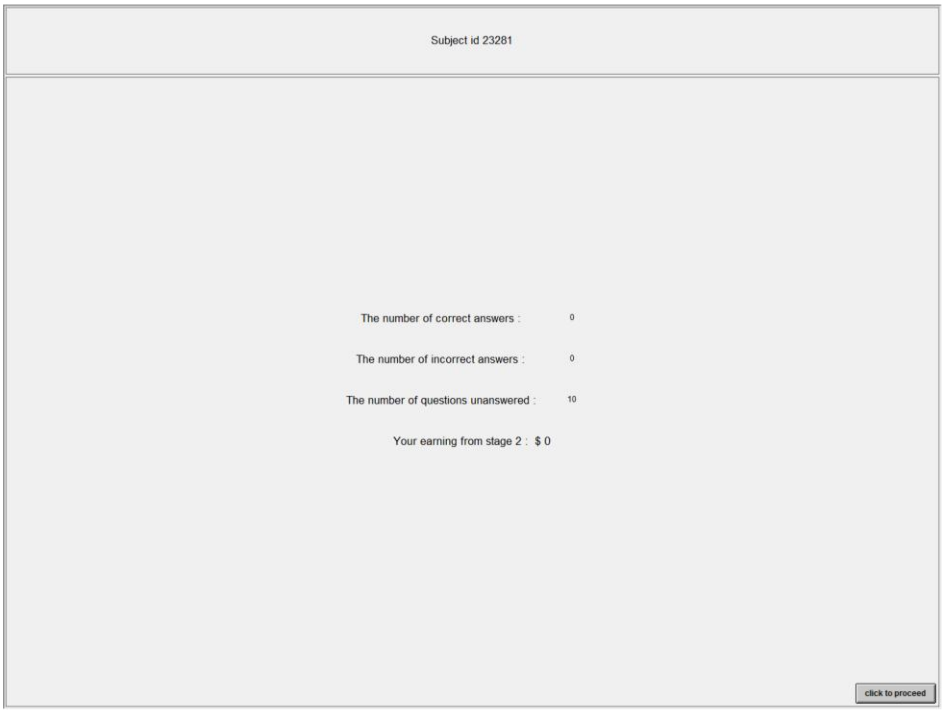

STAGE 3 (Elicitation of Risk Preferences)[18]

In this part of the experiment you will be asked to make a series of choices. How much you receive will depend partly on chance and partly on the choices you make. The decision problems are not designed to test you. What we want to know is what choices you would make in them. The only right answer is what you really would choose.

For each line in the table in the next page, please state whether you prefer option A or option B. Notice that there are a total of 10 lines in the table but just one line will be randomly selected for payment. You do not know which line will be paid when you make your choices. Hence you should pay attention to the choice you make in every line. Below is the screenshot.

Screenshot 12

Line #	Option A	Option B	Please Choose A or B
1	**$1**	**$3** with 0% chance, **$0** with 100% chance	A B
2	**$1**	**$3** with 10% chance, **$0** with 90% chance	A B
3	**$1**	**$3** with 20% chance, **$0** with 80% chance	A B
4	**$1**	**$3** with 30% chance, **$0** with 70% chance	A B
5	**$1**	**$3** with 40% chance, **$0** with 60% chance	A B
6	**$1**	**$3** with 50% chance, **$0** with 50% chance	A B
7	**$1**	**$3** with 60% chance, **$0** with 40% chance	A B
8	**$1**	**$3** with 70% chance, **$0** with 30% chance	A B
9	**$1**	**$3** with 80% chance, **$0** with 20% chance	A B
10	**$1**	**$3** with 90% chance, **$0** with 10% chance	A B

OK

Reward Scheme (Stage 3)

After you have completed all your choices, the computer will randomly generate a number, which determines which line is going to be paid. Your earnings for the selected line depend on which option you chose: If you chose option A in that line, you will receive $1. If you chose option B in that line, you will receive either $3 or $0.

[18] The entire section 3 only applies to the Singaporean sample and not the Thai sample.

To determine your earnings in the case you chose option B, there will be second random draw. The computer will randomly determine if your payoff is 0 or $3, with the chances stated in Option B.

Post-Experiment Questionnaire

After Stage 3 is completed, participants will be prompted with post-experiment questionnaire. Below is the screenshot of the questionnaire.

Screenshot 13

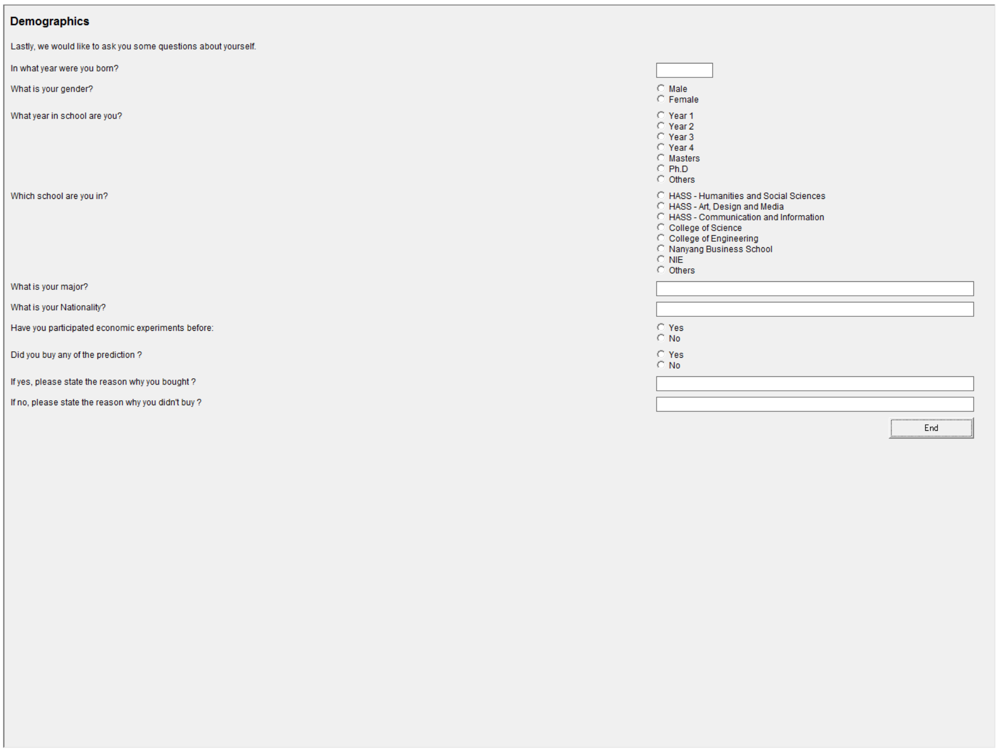

Summary Page of Earnings

Once the questionnaire is completed, the following summary page will be shown to all participants.

Screenshot 14