

The Price of Luck: Paying for the Hot Hand of Others

Silvia Bou
Department of Business
Universitat Autònoma de Barcelona

Jordi Brandts
Instituto de Análisis Económico (CSIC)
& Barcelona GSE

Magda Cayón
Department of Business
Universitat Autònoma de Barcelona

Pablo Guillén
The University of Sydney

January 8, 2016

Abstract

We report on the results of an experiment with a statistical choice task involving the toss of a fair coin. In our experiment participants had to decide whether they were willing to pay a price to switch from betting on the future performance of one player to betting on that of another player. The switch was from a player who had been previously less successful in betting on five coin flips to another one who had been more successful in the same task. We conducted a treatment with the Becker-DeGroot-Marschak mechanism and one in which participants were faced with a fixed price. In both cases, participants exhibit a strong bias towards placing their bets on players with a good guessing history in the coin toss task. Participants' behavior is compatible with prescriptive luck beliefs, that is, the idea that luck is a somehow deterministic and personal attribute.

Keywords: decision heuristics, cognitive bias, economic experiments

JEL numbers: C91

The authors gratefully acknowledge financial support from the Antoni Serra Ramoneda (UAB – Catalunya Caixa) Research Chair, the Spanish Ministry of Economics and Competitiveness (Grant: ECO2014-59302-P), the *Generalitat de Catalunya* (Grant: 2014 SGR 510) and the Barcelona GSE.

1. Introduction

We use an experiment to test the ability of advanced undergraduate finance students to understand iid phenomena. Our aim in designing this experiment was to create the simplest possible environment in which to study whether people are willing to pay for what, given the iid nature of the coin tosses, is pure luck. In our set-up there is no basis whatsoever for the ability to influence success in guessing the sequence of coin tosses. We believe that this complete decoupling of luck and elements of ability is an important strength of our design.

Experimental sessions were divided into three phases and participants were randomly assigned to one of two groups, A or B. In Phase I participants in Group A were each asked to place bets on a sequence of five coin tosses in an incentive-based manner, while participants in Group B observed the process without making any decisions. In Phase II participants in Group B had to allocate themselves to one of two members of Group A, one of which had been more successful than the other in betting in Phase I. In Phase III the two members of Group A had to place bets on another sequence of five coin tosses. Each B's earnings in Phase III were the same as the ones of the A they had allocated themselves to in Phase II. That is, each B's earnings dependent on the luck of the A they allocated themselves to in Phase II.

The focus of our study is the Bs' allocation decision in Phase II. By default all B players were initially allocated to the member of Group A who had earned relatively little in Phase I, due to poor guessing of the first sequence of coin tosses. At this point the B participants had to make the crucial decision. Each B had to decide whether to switch from the less successful to the more successful one. Switching came at a price.

To determine the price participants were willing to pay to switch from the less successful to the more successful one, we used two different treatments, the BDM and the Fixed Price treatments. In the BDM treatment students were asked to quote prices at which they would be willing to switch to another Group A participant. Specifically, Group B participants needed to provide their willingness-to-pay to switch from any level of performance to each of all possible higher levels of performance. This task was incentive-based according to the well-known Becker-DeGroot-Marschak mechanism (for short, BDM). The decision that was implemented was the one pertaining to the effective levels of success in the coin toss of the two A participants.

The BDM is a widely used mechanism in experimental economics. It has the advantage of yielding rich data on the issue we are interested in. It is, however, not free of criticism: from game form recognition to weak incentives.¹ To control for this, we ran the Fixed Price treatment (for short, FP) in which we informed Group B participants of the actual lowest and highest performing Group A participants and put a fixed price on switching from the lowest to the highest.

Studies in experimental economics and psychology have found a variety of biases when participants face statistical tasks in the laboratory. Two prominent biases are the so-called Gambler's (Tversky and Kahneman, 1971) and Hot-hand fallacies (Gilovich et al, 1985). In a nutshell, the Hot-hand fallacy is the belief that people who have experienced repeated success with an event involving randomness have a greater chance of further success in additional attempts than others. The Gambler's fallacy refers to the belief that if something happens more frequently than expected during some period, then it will happen less frequently in the future. Rabin (2000) or Rabin and Vayanos (2010) explain both fallacies within the same model by

¹ See Section 4 for a detailed discussion.

assuming boundedly rational agents incapable of performing proper statistical inference. In many of the older experiments participants are not fully informed about the distribution of the data generation process, like basketball players allegedly affected by the Hot-hand fallacy, and they need to perform some kind of statistical inference or Bayesian updating based on a more or less long stream of observations (see Offerman and Sonnemans, 2004, for an experiment). Recent papers (i.e., Asparouhova et al, 2009, Yuan, Sun and Siu, 2014; Suetens et al, 2015; Powdthavee and Riyanto, 2015) rely on a very simple, unequivocal and clear data generation process such that statistical inference or Bayesian updating is not necessary at all.

Our paper addresses the question whether people are willing to pay to bet on an investor whose performance is completely determined by luck, and our experimental design attempts to represent this in the simplest possible way.² Our results show that a sizeable proportion of B participants are easily influenced by past performance of As and are willing to pay to switch from being tied to a previously less successful to a more successful one. We think that our simple set-up makes it easy for other researchers to check whether our results will replicate in other participant pools and with suitable design variations.

The rest of the article is structured as follows. Section 2 describes the experimental design in detail. Section 3 contains the statistical analysis of the experimental results and Section 4 discusses.

² Powdthavee and Riyanto (2015) study the closely related problem whether people are willing to pay for predictions of truly random outcomes after witnessing only a short streak of accurate predictions live in the lab. In their environment, experimental participants could, before betting on coin flip, pay for viewing a prediction of the coin flip. This prediction was provided by the experimenters and was random, since there was no way to know beforehand what the result of the coin flip would be. The results show that participants are indeed willing to pay for seeing such predictions. We do not think that our design is better than theirs, or the reverse. Perhaps, one can say that our design is slightly simpler, so that it may be a little easier to replicate

2. Experimental design

The experiment consists of two treatments, BDM and Fixed Price. We ran a BDM session with a total of 20 participants in Group A and 28 in Group B. We ran two sessions in the FP treatment with 40 participants in Group A and 47 in Group B.

2.1 Procedures common to all treatments

Participants were all students in an undergraduate finance course at the Universitat Autònoma de Barcelona. Participation was voluntary and remunerated and took place in a large classroom, but not during class. The experiment was conducted with paper and pencil. The lecturer of the course was present during all sessions. Upon arrival students were randomly assigned to one of three groups: A, B or X. There was only one participant in Group X, whose task was to toss a coin. Participants were assigned a participant number and then individually led to their seats in the classroom. They were seated apart so that they could not observe others' decisions and were told not to communicate with each other.

Once all participants were seated, one of the experimenters read the instructions.³ Then Phase I of the experiment began. Participant X tossed a coin five times. Before each coin toss by participant X, each Group A participant had to place a bet on either heads or tails. Each Group A participants earned 2 EUR for each hit, but nothing for a miss. Bets had to be made sequentially and this was checked by the experimenters, that is, Group A participants could not make bets on

³ An English translation of the original instructions for the BDM treatment can be found in Appendix 1.

a coin toss before seeing the outcome of the previous coin toss. Group B participants just observed the coin tosses during Phase I and did not make any decisions in this phase.

After all five coin tosses had taken place we asked six participants of Group A to leave the room with one of the experimenters. Among those six participants there were the two Group A members with the highest and the lowest number of effective hits. Note that six is the number of possible different numbers of hits of Group A participants (0 to 5). We used this procedure to convey to students that all outcomes were possible and to make it in the BDM treatment more natural for participants to consider them all. For parallelism we used the same procedure in the Fixed Price treatment sessions.

Then Phase II of the experiment began, in which the Bs had to make their decisions. Group B participants were told that in Phase III:

- The two Group A participants with the highest and lowest numbers of hits would be guessing sequentially on five subsequent throws of the coin.
- All the Group B participants were initially assigned to the Group A participant with the lowest number of hits.
- Each Group B participant could now, in Phase II, either stick to this allocation or switch to the one with the highest number of hits.
- Group B participants were told that they would need to pay a price if they wanted to switch to the participant with the highest number of hits. The way this was done differed between the two treatments.

After each Group B participant had made the decision whether to switch A players, Phase III of the experiment began. The two Group A participants with the highest and lowest number of hits placed bets on five subsequent coin tosses by participant X. As before, bets had to be placed sequentially. The five coin throws were done separately for the two A players. Group B participants who did not change their default Group A participant went with the Group A participant with the lowest number of hits in Phase I. The values of the bets will be explained below. Once Phase III was finished, we paid participants individually and the experiment ended.

2.2 The BDM treatment

In the BDM treatment, participants were asked to fill out a table indicating the conditional prices for each possible combination of highest and lowest numbers of hits. See table 1. They were allowed to give prices from zero to ten EUR in increases of 0.5 EUR.

PLAYER N°	GROUP B					
	Player to switch to					
Assigned player	0 right guesses	1 right guess	2 right guesses	3 right guesses	4 right guesses	5 right guesses
0 right guesses						
1 right guess						
2 right guesses						
3 right guesses						
4 right guesses						
5 right guesses						

Table 1. Decision sheet in BDM treatment.

Once every Group B participant had made these decisions, the two Group A participants with the effectively highest and lowest numbers of hits were asked to come back into the room. At this moment Group B participants were informed what the highest and lowest numbers of hits had effectively been. This determined for each Group B participant the “personal change price,” that is, the price a Group B participant had given to switch from the *actual* lowest number of hits to the *actual* highest number of hits. Following the BDM price elicitation a price between 0 EUR and 10 EUR was then randomly drawn. The draw was done by one of participants by taking one of twenty-one pieces of paper (with prices in 50 cents steps) out of a bag.

If the randomly chosen price was above a Group B participant’s personal change price, she was not allowed to change and, hence, did not have to pay the price. If the randomly chosen price was smaller or equal than the personal change price then the Group B participant had to change, and paid the randomly chosen price. The amount of the personal change price was subtracted from an endowment of 10 EUR.

The Group B participants who turned out not to change went with the Group A participant with the lowest number of hits in Phase I. They bet 2 EUR on each of the Group A participant’s coin bets. Those Group B participants who did turn out to change went with the Group A participant with the highest number of hits in Phase I. They bet $(10 - \text{randomly determined price})/5$ EUR on each of the corresponding Group A participant’s guess of the coin flip. If the Group A participant obtained a hit the Group B participant earned twice his bet, $2 * ((10 - \text{randomly determined price})/5)$ EUR, if the Group A participant obtained a miss, the Group B participant earned nothing.

2.3 The FP treatment

After Phase I concluded, Group B participants were informed about the actual highest and lowest numbers of hits achieved by Group A participants. Group B participants in the FP treatment were also endowed with 10 EUR. They were informed about the actual lowest and highest number of hits and they were initially assigned to the Group A participant with the lowest number of hits in Phase I. This means that initially their luck in Phase III was tied to that of the least successful member of Group A in Phase I. Then Group B members had to make just one decision. That is, to determine whether a Group B participant was allowed to switch she had to answer the following question: Would you like to switch your bet to the participant with the highest number of hits for Y EUR? Where Y was the median value of the prices elicited in the BDM treatment according to the actual distance in hits between the Group A participant with the highest number of hits and the Group A participant with the minimum number of hits. Y was 2 EUR and 1.5 EUR for FP-1 and FP-2 respectively.

Then, Phase III of the experiment started. The two Group A participants with the highest and lowest number of hits placed bets on five subsequent coin tosses by participant X. Group B participants who did not switch their default Group A participant went with the Group A participant with the lowest number of hits in Phase I. The value of the bets in the FP treatment followed the same logic as in the BDM treatment, as explained at the end of section 2.2. For each of the five coin throws of Phase III, the value of the bet was a fifth of what was left from the endowment of 10 euro after subtracting how much a participant had paid to switch from one A player to the other one. That is, if a B player had not switched A players then the bet on each coin throw was 2 euro. If a B player had paid the price to switch A players then the bet on each coin throw was $(10 - \text{switching price})/5$. The betting rule was double or nothing for each of the

five throws. Once Phase III had finished, we paid participants individually and the experiment ended.

3. Results

3.1 BDM treatment

Participant ID	0-1	0-2	0-3	0-4	0-5	1-2	1-3	1-4	1-5	2-3	2-4	2-5	3-4	3-5	4-5
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0.5	1	2.5	3.5	4	0.5	1.5	2.5	3.5	1	2	3	1	2	0.5
4	0	0	0.5	0.5	1	0	0	0.5	0.5	0	0	0.5	0	0	0
5	0	0	0	1	1	0	0	1	1	0	0.5	0.5	0	0	0
6	1.5	2	2.5	2.5	2.5	2	2	2.5	2.5	2	2.5	2.5	2.5	2.5	2.5
7	0.5	1	1.5	2	2.5	0.5	1	1.5	2	0.5	1	1.5	0.5	1	0.5
8	3	3.5	4	4.5	5	2.5	3	3.5	4	2	2.5	3	2	2.5	2
9	1	3	4	8	10	1.5	3	8	10	3.5	7	9.5	7	8	5
10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	1	0	0	0	1	0	0	0.5	0	0.5	0
13	0.5	1	1.5	2	3	0.5	1	1.5	2	0.5	1	1.5	0.5	1	0.5
14	0.5	1	1.5	2	2.5	0.5	1	1.5	2	0.5	1	1.5	0.5	1	0.5
15	1	3	5	7	9	1	3	5	7	1	3	5	1	3	1
16	0	0	0	2.5	2	0	0	2.5	3	0	0	0	0	0	0
17	0	0	0.5	1	1	0	0	0.5	1	0	0	0.5	0	0	0
18	0	0.5	0.5	1	2	0	0.5	0.5	1	0	0.5	0.5	0	0.5	0
19	0	0	0.5	1	1.5	0	0	0.5	1	0	0	0.5	0	0	0
20	0.5	1	1.5	2	2.5	0.5	1	1.5	2	0.5	1	1.5	0.5	1	0.5
21	1	2	2.5	3	3.5	3	4	5	6	6	7	8	7.5	8	10
22	2	2.5	3.5	4	5.5	2.5	3.5	4	4.5	4	4.5	5.5	2.5	4.5	4
23	0	2.5	5	7.5	10	0	2.5	5	7.5	0	2.5	5	0	2.5	0
24	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
25	1	1.5	3	3.5	4	1	1.5	3	3.5	1	1.5	3	1	1.5	0.5
26	0.5	0.1	1.5	2.5	3	1.5	2	3	3.5	2.5	3.5	4	4	4.5	5
27	0	1	1.5	2	2.5	0	1	1.5	2	0	1	1.5	0	1	0
28	0.5	1	3	4	4.5	0.5	1	4	4	0.5	1	3	0.5	1	0.5
Median	0.25	1	1.5	2	2.5	0.25	1	1.5	2	0.25	1	1.5	0.25	1	0.25

Table 2. Willingness-to-pay to switch, BDM treatment

The Group B participants who choose to switch went with the Group A participant with the highest number of hits in phase in Phase I. They bet EUR $(10 - Y)/5$ EUR on each of the corresponding Group A participant's guess of the coin flip. If the Group A participant got a hit

the Group B participant earned twice his bet $2*((10 - Y)/5)$ EUR, if the Group A participant obtained a miss, the Group B participant earned nothing.

Table 2 shows all the prices chosen by all participants and Figure 1 shows the change-prices elicited from the 28 Group B participants for each of the possible differences in previous hits, distributed in the four quartiles (the * denote outliers).⁴ Observe that the larger the increase in hits the larger the price participants are willing to pay tends to be.




Figure 1: Box-plot of distances

The impression one gets by inspecting the information shown in Figure 1 can be backed up by regression analysis. A GLS regression with participant fixed effects explains the variable willingness-to-pay as a linear function of the difference-in-hits. The difference-in-hits coefficient is positive, significant ($p = 0.000$) and equal to 0.55. That is, when the difference in hits increases by 1 the willingness-to-pay to switch increases on average by 0.55 cents of a Euro.

⁴ The outliers reflect the inconsistent behaviour of two subjects, numbers 9 and 21. See Table 2. Information about bets and realizations can be found in Appendix 2.

3.2 Fixed Price treatment

We ran two sessions of the Fixed Price treatment: FP-1, with 19 Group B participants, and FP-2, with 28 Group B participants. In FP-1 the minimum number of hits in Phase I was 1 and the maximum 5. In FP-2 the minimum and maximum number of hits in Phase I were 1 and 4 respectively. The fixed switching prices (Y) offered to Group B participants correspond to the median switching prices in the BDM treatment for the 1-5 and the 1-4 distances. That is, $Y = 2$ EUR for FP-1 and $Y = 1.5$ EUR for FP-2.

With these prices, about 21% of the participants decided to switch both in FP-1 and in FP-2.⁵ These two percentages can be firstly compared to the 50% in the BDM, given that the median divides the distribution in two halves. The 21% found in the Fixed Price correspond to a 42% of the initial effect in the FP treatments. There are two important caveats to this interpretation of the results.

First of all, note that in the BDM treatment we obtained quite a few values exactly at the median (1.5 EUR). This suggests that it is important to check whether the result changes substantially if the 21% switching rate in FP was to be compared to the BDM switching rate strictly above the median (just 43% rather than 50%). In that case we would still obtain a 49% of the initial effect.

Second, and most important, if the BDM overestimates the willingness-to-pay (see below for possible reasons) the median would be biased in the same direction. If in the Fixed Price treatment we had used price levels lower than the observed medians under the BDM, then the

⁵ Exactly 21.05% and 21.43 for FP-1 and FP-2 respectively.

effect in the Fixed Price treatment would be even stronger. Since it is hard to correct for this in a meaningful way we chose to use the original median for comparison.

4. Discussion

Our data suggest that many of our subjects, students in an advanced finance course, both fail to understand the nature of an iid variable and behave in a manner compatible with beliefs in prescriptive luck. We obtain this result using two different procedures.

We find that our participants, who have some training in statistics, make what at first sight appears to be a rather elementary decision error. In our set-up, ability of any kind cannot play any role. That is, the success of A players in Phase I is purely random. Given the transparency of the environment we use, it may come as a surprise that so many participants are willing to pay for luck. The general implication of our results is that it is quite easy to fall into the trap of attributing (prescriptive luck) to others. This tendency may be exploitable by those who know that many people have this tendency.

However, note that our experimental design only allows for detecting the participants who fall prey to the Hot-hand fallacy jointly with those who make a random mistake in decision-making, since mistakes can only be made in a direction consistent with the Hot-hand fallacy. It may well be possible that a number of seemingly rational participants would follow the Gambler's fallacy and that others would make random mistakes in that same direction if the design allowed them to do so. This issue could be investigated within our design with a new treatment in which the members of Group B would be allocated to the best (instead of the worst) performer in Phase I. With this treatment it would be possible to detect belief in the Gambler's

fallacy jointly with mistakes in that direction. Since random mistakes should a priori not be more prevalent in this case than in the one we study in the paper, the difference in the switching rates would reveal whether people's belief in the Hot-hand fallacy is stronger than the one in the Gambler's fallacy. We think that studying this issue in more depth is a very interesting avenue for future research.

We don't think that either of two procedures we use provides a perfect point estimate. For instance, in the Fixed Price treatment an individual makes only one, albeit very simple, decision and so does not have an opportunity to learn by repetition not to choose a dominated option. The BDM provides sixteen opportunities to each participant to choose a dominated option. In this case, one can object that it should not be too surprising to see how many participants behave in such a way, at least once. Nevertheless, we believe that overall our data clearly suggest that a sizeable fraction of people are willing to pay for luck.

Our results can be interpreted in terms of the distinction between two modes of thought, "System 1" and "System 2," as discussed by Kahneman (2011). Although, upon deliberation, it should be clear to anybody that having been successful in betting on a coin flip in five consecutive trials should have no bearing on future betting success, our first reaction to seeing that a person has been successful in betting may be driven by System 1 and, hence, be fast, automatic and emotional.

Our results can also be interpreted in terms of our participants believing in *prescriptive luck*.⁶ According to Darke and Freedman (1997) prescriptive luck is the perception that good luck is a stable characteristic that consistently favours some people but not others. Thompson and Prendergast (2013) reports the results of survey work in which a substantial proportion,

⁶ As opposed to *Descriptive Luck*, that is, the use of the word "luck" as a way to describe fortunate or unfortunate events.

about 40%, of a student population are self-reported believers in luck. Our results are compatible with the findings reported in Suetens et al (2015), that is, participants in the Danish Lotto believe they can predict future draws. Probably the work most related to ours is Powdthavee and Riyanto (2015): many participants in an experiment are willing to pay for a coin toss ‘prediction’ contained in a sealed envelope. Powdthavee and Riyanto (2015) cite some work done by prestigious psychologist, see Barrett (2004) and Gray and Wegner (2010), to find a plausible explanation for their result. That is, the human brain may well have evolved to find patterns from very limited observations: it is far better to avoid several imaginary predators than being eaten by a real one. Thus, ours and their result would be the result of hard-wired, ancestral brain functions.

Our work can also be linked to two very recent articles, Miller and Sanjurjo (2014, 2015) They prove that in a finite sequence generated by iid Bernoulli trials with probability of success p , the relative frequency of success on those trials that immediately follow a streak of one, or more, consecutive successes is expected to be strictly less than p , i.e. the empirical probability of success on such trials is a biased estimator of the true conditional probability of success (p). They also find that under such conditions related to previous successes, many basketball players succeed with a probability higher than p . Therefore such players get “real” hot hands. Note that such “real” hot hands can only be attributed to a temporary increase of player ability that could be explained, for instance, by increased confidence. Also, note that our experimental design *rules out* any temporary increase in ability as success can only be explained by descriptive luck. That is, our Group B participants do seem to believe in the hot hands of some Group A participants, but real hand-hotness is not possible in our experiment.

Our results are also related to those of Offerman and Schotter (2009). They experimentally study behaviour in a problem of optimal production and in a takeover game and find that subjects the tendency to imitate the best performer in the past is so strong that people consistently fail to distinguish between situations where imitation of the best is a reasonable decision-principle and those where it is not. These results suggest that what we observe may be part of a more general phenomenon of imitating successful others. The environments studied in Offerman and Schotter (2009) are interesting in their own right, since they are examples of some economically relevant situations. One advantage of our environment is, in our opinion, its extreme simplicity.

Several considerations need to be taken into account to better qualify our results. The use of the BDM mechanism has been criticized on a number of grounds. First of all, it has been reported that the BDM procedure may not be behaviourally strategy-proof and thus lead to under or overestimating the willingness-to-pay. For instance, Cason and Plott (2014) argue that participants fail to recognize the game form; Horowitz (2006) claims that willingness-to-pay depends on the distribution of potential prices; Noussair et al (2003) conjecture that differences in the shape of the payoff function drive the results in an experiment in which the Vickrey auction outperforms the BDM in terms of finding the true willingness-to-pay.

Bartling et al. (2015) conduct a new experiment to dig deeper into the issue of misunderstanding of the incentive properties of the BDM mechanism. They report that 70% of the experimental participants do understand the incentive properties while 30% do not. They interpret this as support for Cason and Plott's claim that caution is warranted when relying on valuations with this mechanism. However, Bartling et al. (2015) also find that the data obtained with the BDM do contain valuable information.

Another possible critique to our use of the BDM mechanism is the “payoff dominance” problem reported in Harrison (1992), which could explain both under-reporting and over-reporting of the willingness-to-pay. The argument is about the negligible cost of a seemingly sizable deviation from the true willingness-to-pay. The example calculations in Harrison (1992) result in costs as low as \$0.001. In our case the cost of the average deviation when asking participants about their willingness-to-pay to switch from 0 to 5 hits is 3.6 EUR resulting in an expected cost of 1.3 EU (3.6 times the probability of 36%), several orders of magnitude higher than the calculation in Harrison (1992).

In summary, although we are aware of the fact that there may be some difficulties associated with the use of the BDM, we think that, given its advantages, it is a valid procedure and we, therefore, use it. In fact, it is widely used by experimentalists.

Second, the experiments were conducted in the classroom. The lecturer of the fourth year, Advanced Finance course was present. Note that the principles on which our experiment is based are essential for finance. The presence of the lecturer was meant to reinforce students’ rational, deliberate, costly and thoughtful behavior. That is Kahneman’s (2011) System 2. That is, if there was an experimenter demand effect in our experiment it should have worked in the direction of inducing rational behaviour.

We are aware of the fact that our results challenge very basic elements of rational behaviour and we do understand that many economists may be sceptical about them. However, we think that we have not forced the results and invite other researchers to check whether our results replicate.

The question arises whether the tendency towards attributing prescriptive luck can be remedied. Training in basic statistics seems like a natural way to make people aware of the problem. Indeed, given the simplicity of our experimental set-up, it may be used as a demonstration of how easily people fall into the trap and be used in teaching of how easily people can fall prey to a very basic decision error. More generally, perhaps advice from others can, similarly to what occurs in other environments, be helpful.

6. References

- Asparouhova, E., Hertz, M., & Lemmon, M. (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.
- Bartling, B., Engl, F. & Weber, R. A. (2015). Game form misconceptions are not necessary for a willingness-to-pay vs. willingness-to-accept gap. *Journal of the Economic Science Association*, 1(1), 72-85.
- Barrett, J. L. (2004). Why would anyone believe in God? AltaMira Press.
- Cason, T. N. & Plott, C. R. (2014). Misconceptions and game form recognition: challenges to theories of revealed preference and framing. *Journal of Political Economy*, 122(6), 1235-1270.
- Darke, P. R. & Freedman, J. L. (1997). The belief in good luck scale. *Journal of Research in Personality*, 31(4), 486-511.
- Gilovich, T., Tversky & A., Vallone, R. (1985). The Hot-hand in basketball: on the misperception of random sequences. *Cognitive Psychology*, 3, 295-314.
- Gray, K., & Wegner, D. M. (2010). Blaming God for our pain: Human suffering and the divine mind. *Personality and Social Psychology Review*, 14(1), 7-16.
- Harrison, G. W. (1992). Theory and misbehavior of first-price auctions: Reply. *The American Economic Review*, 82(5), 1426-1443.
- Horowitz, J. K. (2006). The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods. *Economics Letters*, 93(1), 6-11.
- Kahneman, D. (2011). Thinking fast and slow. Penguin books, Melbourne.
- Miller, J. B., & Sanjurjo, A. (2014). A cold shower for the hot hand fallacy. *SSRN working paper*.
- Miller, J. B., & Sanjurjo, A. (2015). Is it a fallacy to believe in the hot hand in the NBA three-point contest? *Bocconi University working paper*.
- Noussair, C., Robin, S. & Ruffieux, B. (2004). Revealing consumers' willingness-to-pay: A comparison of the BDM mechanism and the Vickrey auction. *Journal of Economic Psychology*, 25(6), 725-741.
- Offerman, T. & A. Schotter (2009). Imitation and luck: An experimental study on social sampling. *Games and Economic Behavior*, 65, 461-502.
- Offerman, T. & Sonnemans, J. (2004). What's causing overreaction? An experimental investigation of recency and the hot-hand effect. *The Scandinavian Journal of Economics*, 106(3), 533-554.
- Powdthavee, N., & Riyanto, Y. E. (2015). Would you pay for transparently useless advice? A test of boundaries of beliefs in the folly of predictions. *Review of Economics and Statistics*, 97(2), 257-272.

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

Rabin, M. & Vayanos, D. (2010). The gambler's and hot-hand fallacies: Theory and applications. *The Review of Economic Studies*, 77(2), 730-778.

Suetens, S., Galbo-Jørgensen, C. B., & Tyran, J. R. (2015). Predicting lotto numbers: a natural experiment on the gambler's fallacy and the hot-hand fallacy. *Journal of the European Economic Association*, forthcoming.

Thompson, E. & Prendergast, G. (2013). Belief in luck and luckiness: Conceptual clarification and measure validation. *Personality and Individual Differences*, 54(4), 501-506.

Tversky, A. & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105-110.

Yuan, J., Sun, G. & Siu, R. (2014). The lure of illusory luck: How much are people willing to pay for random shocks. *Journal of Economic Behavior and Organization*, 106, 269-280.

Appendix 1

GENERAL INSTRUCTIONS

Welcome to this experiment,

IT IS VERY IMPORTANT TO REMAIN SILENT DURING THE WHOLE EXPERIMENT!!!

If you have any doubts about the instructions please raise your hand and wait until one of the experimentalists comes to your place to solve it.

You will receive 4 € as a show-up fee.

This experiment has three phases that called PHASE ONE (only Group A plays), PHASE TWO (only Group B plays), and PHASE THREE (Group B and part of Group A plays).

There are three types of players:

Players from group A

Players from group B

Player X: (THE INNOCENT HAND)

PHASE ONE

In this phase players in group B do not play, they only observe, so only the group A players act. Group A players must bet heads or tails when player X throws a coin. The players that guess right will obtain 2 € and the ones that don't get 0 €. Player X earns 10€ for her participation.

The process will be as follows:

- 1- Every player A will make her bet in an individual way: if she bets that heads will come out, she writes down a C in the cell "Bet" and if she that tails will come out she writes a + in the cell "Bet" in the folder that has been given to her.
- 2- Player X will throw the coin
- 3- Every player A will check the result (heads or tails) and will fill the cell "result" by writing a ✓ if she had the right answer and an X if she didn't. All these can be found in the documents that have been given to players A.

There will be five rounds of this process. The bets will be made before each one of the throws. We will check that before each throw all players in group A have made their bets. Once the five rounds have been completed the documents will be collected.

After that a group of 6 players A will be selected and we will invite them to go out of the room. The selection will be made according to the different possible results that might come out in the five rounds.

The selected players A will go out of the room and wait for instructions.

The rest of players A will remain seated in their places during the rest of the experiment until they are called to be paid what they earned in the experiment.

PHASE TWO

This is the phase where players from group B participate in an active way in the experiment.

Each player B has an initial endowment of 10 €.

We assign by default to every player B one of the players that have been selected to leave the room, specifically the one that had the lowest number of right answers in PHASE ONE.

In phase three we will repeat the five rounds of coin throws as in the first phase. Players from group B will not bet. Players B's earnings will be determined by the player A that has been assigned to them (minimum number of right answers). This earnings will be of 2 € per right answer. But players B have the opportunity to change from the assigned player A to the one that had the maximum number of right answers in phase one.

To switch from the assigned player to the one with the highest number of right answers B players must pay a price.

The way to determine if a player B will switch from the assigned player A to the one of maximum number of right answers or will remain with the assigned player A works as follows:

- 1- First, every player B will determine the price that she is willing to pay to switch from the A player with the minimum number of right answers to another A player with higher number of right answers. This will be done by filling in the table that has been given to them in the documentation.

PLAYER N°		GRUP B				
	Player to switch to					
Assigned player	0 right answers	1 right answers	2 right answers	3 right answers	4 right answers	5 right answers
0 right answers						
1 right answers						
2 right answers						
3 right answers						
4 right answers						
5 right answers						

The prices for changing can be expressed in fractions of 0.5 €. The maximum price that can be paid is the initial endowment of 10 €.

- 2- Second, the organizers will reveal the results of the selected players A so players B will know the number of right answers that the player they have been assigned to had, and will also know the score of the maximum number of right answers. This will allow players B to know which cell in the table they are playing with.
- 3- Third, we will determine the random price of change by a lottery. If the price of change that player B has set is lower than the lottery price then player B doesn't switch. If the price of change that player B has set is equal or higher than the lottery price then player B switches to the player with the maximum number of right answers by paying that lottery determined random price. Each player B must write down the player that finally is assigned to her in the cell on the first page of the documentation she has been given.
- 4- If player B does not change, she bets 2 € per round
If player B does change, her bet will be:

$$bet = \frac{(10€ - random\ price\ of\ change)}{5}$$

The earnings will be double the bet.

*As an example: If the random price of change is 2.5 € Player B bets $(10-2.5)/5=1.5$ per round, so if she get the right answer then the earnings will be double = $1.5*2=3$. Wrong answers have a cost of 0 €*

After that we will let into the room the 6 players A and identify the player with the minimum number of right answers and the one with the maximum number of right answers and they will proceed to play the five rounds of coin throws.

While players A come into the room we will collect the documentation from B players.

PHASE THREE

In this phase we will proceed as follows:

- 1- Player A will make her bet out loud: Heads or Tails
- 2- Player X throws the coin
- 3- The result will be written on the board.

There will be five rounds for each of the two players (minimum and maximum number of right answers). Bets will be made before each throw.

Players A do not earn money in this phase.

Players B's earnings depend on the amount of right answers that the A player they are assigned gets.

Once the experiment has finished it is important that all players remain seated until they are called by their number to be paid what they have earned in the experiment.

Appendix 2

BETTING SEQUENCES AND RESULTS OF THE COIN THROWS

The BDM session:

The sequence of coin tosses in Phase I was THTHH. The worst Group A participant obtained 2 hits with a betting sequence of TTTTT. The best A player obtained 4 hits in Phase I, with the betting sequence HHTHH. Fortune reversed in Phase III, where the two A players placed bets on separate sequences of coin throws. Now the two participants obtained respectively 4 and 1 hits. For the first player the sequence of bets was HHHHH, while the coin throws yielded HHHTH. For the second player the sequence of bets was HHTHH while the coin throw yielded HTHTT.

The FP-1 session:

The sequence of coin tosses in Phase I was TTHTH. The worst Group A participant obtained 1 hit with a betting sequence of HHHHT. The best A player obtained 5 hits in Phase I. Fortune reversed again in Phase III. Now the two participants obtained respectively 4 and 2 hits. For the first of these players the sequence of bets was THHTH, while the coin throws yielded HHHTH. For the second player the sequence of bets was HHHHH while the coin throw yielded HTHTT.

The FP-2 session:

The sequence of coin tosses in Phase I was THHTT. The worst Group A participant obtained 1 hit with a betting sequence of HTTTH. The best A player obtained 4 hits with a betting sequence of THHTH. In this session fortune did not reverse. Now the two participants obtained respectively 0 and 2 hits. For the first of these players the sequence of bets was HHHHT, while the coin throws yielded TTTTH. For the second player the sequence of bets was TTHHT while the coin throw yielded HTTTT.