Validating the predictions of
case-based decision theory

Benjamin J E Radoc

Thesis submitted for the degree of
Doctor of Philosophy

University of East Anglia

School of Economics

October 2015

Contents

Appendices   1.1: Experiment instructions
                 1.2: Decision form
                 1.3: Selected experiment photos
                 2.1: Positive treatment instructions
                 2.2: Independent treatment instructions
                 2.3: Negative treatment instructions
                 2.4: Quiz screenshots
                 3.1: Distribution of gains and losses
                 3.2: Broker recommendation
                 3.3: Philippine Stock Exchange Index (2007-2010)
                 3.4: Client interface
                 3.5: PSG/PSL and PDG/PDL subsample comparison

List of Figures

List of Tables

Acknowledgments

The completion of this thesis would not have been possible without the help of several groups of people. I am indebted to the following:

To SSF and ECO for my studentship; their generosity allowed me to pursue full-time studies. I also thank CBESS for the research grants; without these, there was no way I could have afforded any of my experiments.

To my supervisors, Professor Robert Sugden and Professor Theodore L. Turocy; I cannot thank them enough for their guidance and support. No one can argue that I've received top-notch supervision from them.

To my examiners, Arvid Hoffmann and Alasdair Brown; their comments and suggestions helped improve the quality of this thesis.

To my PGR/ECO/couchsurfing friends (and their friends) who have made postgraduate student life more pleasant than I expected. Thank you for the good times at several (home-made/take-away/restaurant) meals, at the poker table, the badminton court, at impromptu parties, and elsewhere.

To my family and friends back home in the Philippines for their unending support throughout my life choices.

I dedicate this body of work to the memory of my father.

Abstract

Real-life decision-makers typically do not know all possible outcomes arising from alternative courses of action. Instead, when people face a problem, they may rely on the recollection of their past personal experience: the situation, the action taken, and the accompanying consequence. In addition, the applicability of a past experience in decision-making may depend on how similar the current problem is to situations encountered previously. Case-based decision theory (CBDT), proposed by Itzhak Gilboa and David Schmeidler (1995), formalises this type of analogical reasoning. While CBDT is intuitively appealing, only a few experimental and empirical studies have attempted to validate its predictions. This thesis reports two laboratory experiments and an empirical study that attempt to confirm the predictive power of CBDT vis-à-vis Bayesian reasoning.

Introduction

How do investors decide whether to buy a stock or not? Investing in financial markets can be daunting, especially to inexperienced individuals. Some investors may have very limited market information, while those with access to multiple data sources may find it difficult to process all available information. For example, Agnew and Szykman (2005) showed that investors without basic knowledge of financial markets remain confused even after financial information is presented in a simplified format, or after investment options have been reduced. Because they are unlikely to possess actual knowledge on all states of the world, decision-makers may instead opt to rely on information gathered from similar experiences in the past.

The importance of a similarity notion in decision-making has long been recognised (for examples, see Rubinstein 1988; Leland 1994; Buschena 2003). Given features of a problem and applying analogical reasoning (Gregan-Paxton and Cote 2000), a decision-maker may be able to match the surface and structural features[1] between a base situation and the target problem. If the target event is perceived as sufficiently similar to the base event, the individual will likely adopt the same successful act taken in the past.

Gilboa and Schmeidler (1995) formalised case-based decision theory (CBDT) using the concept of similarity and utility from past experience to explain behaviour. CBDT predicts that given a new problem, a decision-maker will act based on the memory of actions and associated outcomes in past similar situations. Before deciding, an agent assesses the similarity of the current situation with past problems encountered, and then recalls the actions taken in those similar situations. The theory predicts that a decision-maker will choose a past action in similar situations with the highest similarity-weighted sum of outcomes (Gilboa, Liebermann and Schmeidler 2006).

---

[1] Surface features pertain to the description of the individual elements of a representation, while structural features capture the relationship among the elements of the representation (Gregan-Paxton and Cote 2000). As a decision-maker gains experience, knowledge transfer between decision problems becomes more largely driven by structural features, rather than surface features (Zizzo 2003).

Under CBDT, an experience is encoded into memory as a case with triple elements: problem p, act a and result r. When an agent faces a new problem q, she scans memory M for cases encountered in the past and evaluates similarity vis-à-vis p, conditional on similarity function s. At each similar case in M, the decision-maker recalls the act a chosen and the corresponding outcome r.[2] Given problem q, memory M, similarity function s, and utility function $u(r)$, available acts $a' \in A$ are ranked based on the similarity-weighted sum of utilities from each act:

$$U(a') = U_{q,M}(a') = \sum_{(p, a=a', r) \in M} s(p,q)u(r)$$

Each act is evaluated over a different set of cases so that a decision-maker's memory of cases on one act is disjoint from her memory of cases on another act. This assumption of act separability proposes that decision-makers maintain separate memories of alternative actions undertaken in the past. Since an act is evaluated over past outcomes on that act, experiences from other acts are not taken into account during decision-making.

In decision-making, all that the theory requires is the agent's ability to recall past cases, and the capacity to evaluate similarity between new and past problems encountered. Since only experienced cases are in the agent's memory, CBDT accommodates the possibility for ignorance (when neither outcomes nor probabilities are known). When a new problem is entirely novel, a decision-maker is assumed to randomly choose among possible acts. If a similar problem is repeatedly encountered, available acts are evaluated based on the average similarity-weighted utility of each act. In this instance, the case-based prediction converges with the expected utility from the act (Gilboa and Schmeidler 1995; 2010; Sugden 2004).

---

[2] In CBDT, outcomes may also be evaluated against an endogenous aspiration level which serves as the indifference threshold between alternative acts, and determines whether an act yields a satisfactory outcome (Gilboa and Schmeidler 1995; 2001; Simon 1957). Given an aspiration level, a satisficed decision-maker will tend to repeat a previous act without exploring alternative acts (habit formation). Meanwhile, the introduction of an endogenous aspiration level (i.e., where the threshold adjusts according to the number of successes experienced) precludes habit formation. In this thesis, outcomes are evaluated without explicit regard for an aspiration level.

CBDT imposes minimal requirements on a decision-maker's knowledge on the states of the world, and appears as an outlier among decision theories in economics as it veers away from belief formation and Bayesian updating (Sugden 2004; Gilboa and Schmeidler 2010). Case-based reasoning, however, is not uncommon in other fields. In artificial intelligence and cognitive science, cases are useful in deriving benchmark solutions that are adapted to the peculiarities of a new situation (Kolodner 1992). A nearest-neighbour is used as a starting point and the agent adapts to accommodate the current case encountered (Zizzo 2003). As agents implement indexing or the assignment of tags to aid in memory retrieval, they generate classification predictions through instance-based algorithms even if only a limited number of instances are available (Aha, Kibler, and Albert 1991).

The use of analogy between examples facilitates learning (Ross 1987). In psychology, learners compare unfamiliar events against prototypes (instances that vary in the degree to which they share specific properties) and exemplars (specific representations of various instances that are more accessible to individuals than an abstract summary description of an instance). These representations have been found to be sensitive to both context and frequency of occurrence (Smith and Medin 1981; Nosofsky 1988; Nosofsky and Zaki 2002).

Among learning theories in psychology, instance-based learning theory (IBLT) resembles CBDT. IBLT was developed to explain dynamic decision-making – when decision conditions change spontaneously and heuristics become ineffective (Gonzalez, Lerch and Lebiere 2003; Gonzales and Dutt 2011). In IBLT, instances are encoded in memory as a triplet (situation, decision and utility), and new situations are matched against retrieved instances in memory. The probability of retrieval depends on both the recency and frequency of instances. A decision-maker acquires her "blended" knowledge, learns to attribute actions to experienced results, and later allows for generalization.

Despite CBDT's intuitive appeal, empirical work on the validity of its predictions is quite limited. The empirical studies here include the application of CBDT to portfolio selection (Golosnoy and Okhrin 2008), herding behaviour (Krause 2009), home sales and rental price determination (Gayer, Gilboa and Liebermann 2007), and price-capacity

coordination problem for firms (Jahnke, Chwolka, and Simons 2005). Pape and Kurtz (2013) also successfully developed a computer program to implement CBDT.

There are even fewer experiments related to CBDT. Grosskopf, Sarin and Watson (2015) compared the predictive power of CBDT vis-à-vis the max-heuristic (i.e., choosing the action with the highest historical profit). Ossadnik, Wilmsmann, and Niemann (2013) induced an environment with structural ignorance (where it is difficult to ascertain the states of the world) and pitted CBDT against several decision criteria. Meanwhile, Bleichrodt, Filko, Kothiyal, and Wakker (2015) conducted an experiment that required subjects to choose between two real-estate investments given experimentally-induced cases of real-estate properties. Details of these past experiments are briefly described in Chapter 1.

While the concept of similarity is central to the decision-making process, CBDT does not provide details on the similarity function. Since a similarity concept is derived from preferences, it could be unique to each individual (Grosskopf, Sarin and Watson 2015). The similarity function may also change as a decision-maker accumulates more experiences. This makes it especially difficult to implement CBDT as proposed by Gilboa and Schmeidler (1995).

In the studies presented in this thesis, we imposed feature-based similarity where two objects are considered similar if a salient feature common to the objects matched.[3] This

---

[3] Psychology provides alternative representations of similarity. In their survey, Goldstone and Son (2005) categorised similarity models into four: geometric, feature-based, alignment-based and transformational. Geometric or multidimensional scaling models (MDS) measure similarity between a pair of objects as the inverse of their measured distance. The MDS technique is useful in revealing the underlying dimensions used by agents in assessing similarity. Meanwhile, feature-based similarity as a linear combination of the measures of the objects' common and distinctive features shows the degree to which two sets of salient features match each other. (Tversky 1977; Tversky and Gati 2004). Alignment-based models take into account how features correspond with one another, in addition to the process of matching features. Features that serve the same function, i.e., have relational correspondence, are deemed more similar. On the other hand, the transformational approach to similarity assessment involves comparing the relative ease of transforming one representation into another. Objects that require a more

strategy of treating similarity as a binary variable (i.e., two problems are either identical or completely different) and applying across individuals, overcomes the problem of specifying the form of the similarity function.

The three chapters presented in this thesis attempt to validate the predictions of CBDT, i.e., whether decision-makers encode and retrieve past experiences using similarity information, and then choose an act with the highest similarity-weighted outcome. Chapters 1 and 2 are laboratory experiments, while Chapter 3 is an empirical paper using retail investor data.

Chapter 1 reports the results of a paper-and-pencil experiment where colour was used as a salient similarity cue. Participants encountered two coloured tickets (blue or yellow), which paid earnings generated by a single mechanical randomiser. The experimental design allowed a fair chance for either case-based or Bayesian prediction to emerge. While the results indicate that participants categorised their past experience using colour, participants' ticket valuations were inconsistent with CBDT. The evidence suggests that our participants were neither case-based nor Bayesian. Instead, we found an *exhibit* consistent with the gambler's fallacy.

In Chapter 2, we present a computer-based experiment where participants played a two-armed bandit each framed as a coloured game board (blue or yellow). If participants are Bayesian, decisions would correspond to the known game board correlations (positive, independent or negative). Under act separability, participants may neglect correlation information and decisions are inconsistent with Bayesian reasoning. If this holds, each game board is evaluated over a set of different cases and a separate memory for each game board is maintained so that a decision on a blue game board is not influenced by outcomes on the yellow game board. We find that participants' decisions in the positive and independent treatments are qualitatively Bayesian, but decisions in the negative treatment cannot be explained either by Bayesian reasoning or by CBDT. However, evidence suggests that participants systematically used past outcomes in forming expectations, but judgment did not directly translate into the expected decision.

---

complex transformation procedure are considered less similar than objects whose representations are easier to transform.

Chapter 3 presents the results of an empirical study which analyses retail investors' stock purchases vis-à-vis memory of personal trading outcomes. We categorised retail clients' purchase decisions using selected salient similarity concepts: stock name, industry sector, and broker recommendation. Our analysis reveals that retail investors systematically used similarity information and purchased shares similar to stocks they previously traded either at a realised or unrealised gain. The results confirm a significant similarity effect that is not accounted for by an increase in wealth.

Chapter 1

Colour-coded decisions: an experiment on
case-based decision theory

CBDT predicts that under uncertainty, a decision-maker will choose an action in past similar situations with the highest similarity-weighted outcome. We conducted an experiment to compare the predictive power of CBDT and Bayesian reasoning under objective uncertainty. Our experimental design provided a fair chance for either case-based or Bayesian prediction to emerge. However, the results show that our participants were neither case-based nor Bayesian. We found an *exhibit* consistent with the gambler's fallacy, a result that undermines the predictive power of CBDT.

1.1 Introduction[4]

CBDT proposes that a decision-maker's experience is encoded into memory as a case with triple elements: problem, act and result. Before a decision-maker acts, she culls her memory for past cases and evaluates the similarity between the new problem and past problems encountered (Gilboa and Schmeidler 1995; 2010). CBDT predicts that a decision-maker faced with a new problem will act based on her memory of actions and the associated outcomes in past similar situations so that the act with the highest similarity-weighted sum of outcomes is likely to be chosen (Gilboa, Liebermann and Schmeidler 2006).

Case-based decisions, however, may not always be behaviourally distinct from decisions consistent with expected utility theory (EUT). Given a well-specified set of problems, a complete mapping of all possible combinations of problems and actions into outcomes, and a correspondence between conditional belief systems in expected utility models

and similarity functions in CBDT, Matsui (2000) showed that EUT and CBDT yield equivalent behavioural predictions.

Unlike decisions under EUT, case-based decision-making does not rely on an agent's knowledge of all states of the world and therefore imposes minimal cognitive demands on a decision-maker. However, the theory's proponents quickly emphasised that CBDT complements rather than competes against EUT. EUT is plausible under risk, while CBDT performs better under ignorance, i.e., when neither outcomes nor probabilities are known to the decision-maker.

CBDT as analogical thinking is intuitively appealing but its predictions have not yet been extensively validated. While a few empirical studies have been completed (for example, see Jahnke, Chwolka, and Simons 2005; Gayer, Gilboa, and Lieberman 2007; Golosnoy and Okhrin 2008; Krause 2009; Pape and Kurtz 2013), experimental work on CBDT is even more limited. This is not surprising. Apart from the difficulty of behaviourally distinguishing CBDT from EUT (Matsui 2000), CBDT rests on the notion of similarity which is not widely studied in economics.

To our knowledge, only three studies have attempted to experimentally validate the predictions of CBDT. Grosskopf, Sarin and Watson (2015) compared the predictive power of CBDT vis-a-vis the max-heuristic (i.e., choosing the action with the highest historical profit) in an ambiguous monopoly production decision setting with a variable payoff function. In the experiment, participants decided on production choices in several markets described as incomplete sets of market variables (presented as shapes). Each decision was made in a market condition that is independent from the other decisions. When profit information (feedback) was delayed, participants used similarity cues and decisions were consistent with CBDT's prediction; but when past profits were known, participants resorted to the max-heuristic.

Ossadnik, Wilmsmann, and Niemann (2013) pitted CBDT against several decision criteria.[5] To induce an environment with structural ignorance (i.e., when it is difficult to

---

[5] The experiment consisted of 3 runs: A, B, and C. Based on the urn-composition, B is 2/3 similar to A, while C is 6/11 similar to A. Ossadnik, Wilmsmann, and Niemann (2013) compared the conformity of participants' decisions with CBDT against the following alternative decision

ascertain the states of the world), their online repeated choice experiment involved urns containing different ball distributions. Each three-faced ball had three colour fields (red, black, blue) with different numerical values. Participants knew the number of balls in the urn, but not the distribution of the colour-number combinations. Participants' task was to choose a colour (red, black or blue) and accumulate points based on the corresponding colour-number combination on the ball drawn. Results showed that while CBDT performed better at predicting participants' decisions, case-based decision-makers earned less than the players who used alternative decision criteria.

Bleichrodt, Filko, Kothiyal, and Wakker (2015) conducted an experiment that required participants to choose between two real-estate investments given experimentally-induced cases of real-estate properties. Each participant's payoff was revealed a month after the experiment, based on the actual price appreciation of the selected property. While participants' decisions were aligned with the similarity-weighted returns of past investments, results indicated a violation of CBDT's separability axiom.[6] Under act separability, each act is evaluated over a different set of cases so that a decision-maker's memory of cases on one act is disjoint from the memory of cases on other acts. However, in the experiment, they found that the value attached to a real-estate property was influenced by information on other real-estate investments despite clear instructions for participants to treat each case separately.

In these past experiments, ambiguity in the outcome space was induced and participants were prodded to pay attention to the similarity across the decision settings. We designed a one-shot experiment that takes a less suggestive approach on similarity and allowed a fair chance for participants to either use or ignore similarity information.

criteria: i) maximin (very pessimistic: choosing the colour with the highest minimum payoff); ii) maximax (very optimistic: choosing the colour with the single highest outcome); iii) pessimism-optimism (choosing the colour with the maximum weighted value of the lowest and highest outcomes, with a pessimism-optimism index estimated for each individual); and, iv) reinforcement learning model (choosing the colour with the highest propensity for selection as a function of the frequency of successful outcomes separately determined for each run of the experiment).

[6] Although admittedly restrictive, Gilboa and Schmeidler (1995) imposed the separability axiom to "guarantee the additively separable representation" of preferences (p. 614). We attempted to validate act separability in another experiment which is presented as Chapter 2 of this thesis.

With the aim to compare the predictive power of CBDT vis-a-vis Bayesian reasoning, we implemented an environment that effectively induced uncertainty in the outcomes. Our experimental design attempted to validate two predictions of CBDT, namely: i) decision-makers encode and retrieve past experiences using subjective similarity; and ii) agents choose an act with the highest similarity-weighted outcome.

In our two-part experiment, participants played coloured tickets (blue or yellow) which paid earnings based on random draws from a single bingo cage containing an unknown distribution of £20 and £0 balls. While colour was highly salient, it was not linked in any way to the ticket earnings and was clearly uninformative on the probability of a successful draw (£20 payoff). Ignoring colour was an easy strategy for a Bayesian player.

Given uncertainty in the ticket earnings and colour as the only similarity cue, a case-based player is likely to code and retrieve memory of past rounds based on colour. Gilboa and Schmeidler (2001) pointed out that the term similarity in the context of CBDT "should not be taken too literally" (p.36). A salient environmental cue (such as colour) may trigger conscious similarity assessment and facilitate recall of past events. Hence, evidence on participants' systematic use of colour in our experiment is supportive of CBDT, but is inconsistent with Bayesian updating.

Our results provide evidence that participants used colour in coding events during the experiment. Blue and yellow tickets were valued differently in sessions with very few successful draws. However, this does not hold in sessions with more frequent successes. In addition, valuations attached to the coloured tickets were the *opposite* of CBDT's prediction: a ticket colour with fewer successes was valued more highly than a ticket colour with more successes. The results imply that our participants were neither case-based nor Bayesian.

The failure of CBDT in our experiment leads to an interesting *exhibit* (Bardsley, Cubitt, Loomes, Moffatt, Starmer, and Sugden 2010). The pattern in ticket valuations is consistent with the gambler's fallacy or the erroneous belief that a lottery which had a series of losses is bound to reverse the pattern of past outcomes (Rogers 1998). This well-documented misperception arguably results from people's susceptibility to derive

sequential patterns and offer deterministic explanations despite the known uncertainty of outcomes (Gilovich, Vallone, and Tversky 1985; Wood 1992; Sun and Wang 2010).

The rest of the paper is organised as follows: Section 1.2 describes the experimental design; Section 1.3 discusses the results; and Section 1.4 concludes.

1.2 Experimental design

In CBDT, a decision-maker evaluates each act in a similar situation based on the memory of outcomes resulting from the act. In our between-subjects experiment, participants played with coloured tickets (blue or yellow) where ticket earnings in each round were visibly generated by a single mechanical randomiser. After playing several sample rounds to create a set of cases in memory, each participant decided whether she preferred to keep her ticket in the last round of the experiment or to exchange her ticket for money.

The experiment was designed as a one-armed bandit[7] with an added similarity cue. Since earnings in each round were generated by only one bingo machine, a Bayesian player will tend to ignore ticket colour and will perceive each round as playing a one-armed bandit repeatedly. A Bayesian decision-maker will ignore ticket colour and put the same value on a blue and yellow ticket based on the total number of successful draws experienced in the sample rounds. On the other hand, systematic evidence that participants used similarity information (colour) in decision-making is consistent with CBDT; that is, a ticket colour with more successes in the sample rounds will be valued more highly than a ticket colour with fewer successes. This is equivalent to a case-based player treating a one-armed bandit as a two-armed bandit.

The paper-and-pencil experiment, set up akin to a game show, involved coloured tickets (blue and yellow) and a white bingo cage. The use of mechanical logistics provided an

---

[7] Bandit problems are commonly used in studies that attempt to explain decisions to explore or exploit available actions. In an *n*-armed bandit, a player pulls an arm which results in either success or failure. A player then decides which arm to pull in the next round to maximise her total expected payoff (Gittins 1979).

engaging task for participants, emphasised the salience of colour as a similarity cue and the randomness of earnings in each round.

The experiment consisted of two parts. In Part 1, participants played 10 sample rounds[8] to create 10 cases in memory.[9] The number of rounds was close to the median stopping rule of participants in past learning experiments (for example, see Gonzalez, Lerch and Lebiere 2003; Hertwig, Barron, Weber and Erev 2004).

Each session had five blue rounds and five yellow rounds that were randomly ordered, depending on the set of tickets drawn. In addition to coloured tickets and coloured boxes, the "blueness" or "yellowness" of each round was emphasised by the experimenter's repeated announcement of the round played, and coloured light bulbs illuminating the bingo cage.

All sessions were conducted with a lab assistant. To determine the earnings from the ticket at each sample round, the assistant drew one ball (with replacement) from a covered bingo cage containing 100 white balls marked either £20 or £0. Participants knew that the bingo cage contained 100 balls, but not the distribution of the balls.[10] With a 20% probability of a successful draw, £20 was both a rare and salient outcome. At each session, all information in both parts of the experiment was publicly known.

---

[8] Hertwig, Barron, Weber and Erev (2004), and Gonzalez and Dutt (2011) showed that the mode of learning affects the importance attached to a rare event. Decisions from description tend to overweight rare outcomes. Meanwhile, decisions from experience tend to underweight rare outcomes but gave players the opportunity to learn and use base rates. Consequently, trial-by-trial learning brings actual decisions closer to Bayesian reasoning (Hertwig and Ortmann 2001). In our experiment, participants sampled from experience (i.e., earnings were sequentially revealed), and imposed a stringent decision setting for CBDT to work.

[9] Case-based decisions rely on memory of cases so participants were not allowed to take down notes during the sample rounds. With only six participants per session and the synchronized completion of each activity in each round, it was easy for the experimenter to monitor individual participants.

[10] When the instructions in Part 1 were read, participants were given the opportunity to have a close look at the covered bingo cage and two sample balls (£0 and £20). The balls were returned in the bingo cage before Round 1 so participants knew that the bingo cage contained at least one £20 ball.

At the start of each round, participants indicated on the sample ticket their expectation of the chance that a £20 ball will be drawn in that round. After everyone partially filled in the sample ticket, the experimenter gave the signal to draw a ball. The assistant announced and showed the ball drawn. After each draw, participants filled in the earnings on their sample ticket and then dropped the coloured ticket in an opaque box with the same colour as the ticket. Participants were aware that their total earnings at the end of the experiment depended only on the outcome of their decision in Part 2 plus a show-up fee of £2. But participants also knew that the tasks in Part 1 provided information about the distribution of £20 and £0 balls. The synchronized completion of each task following verbal cues from the experimenter and the use of similarly coloured stimuli at each round were implemented to keep the participants engaged throughout the experiment.

Part 2 of the experiment consisted of a task that elicited participants' valuation of a randomly assigned coloured ticket following a BDM mechanism (Becker, DeGroot, Marschak 1964). At the start of Part 2, each participant drew one sealed brown envelope containing a coloured ticket and a coloured decision form from a bag. The decision form listed 35 possible offer prices ranging from 20 pence to £20. At each offer price, each participant decided whether she preferred to keep her ticket or to exchange her ticket for money. Before participants filled in the decision form, one of the participants randomly selected an offer price from a stack of 35 sealed envelopes. To increase the likelihood of truthful willingness-to-accept responses (Plott and Zeiler 2005; Isoni, Loomes, and Sugden 2011), the instructions included an outright statement that the participants' answers on the decision form cannot influence the actual offer price.

The actual offer price was revealed only after all participants submitted their decision form. If a participant decided to keep her ticket at the offer price, her earnings were equal to the outcome of her individual draw. Otherwise, she was paid the offer price. Whichever the decision, each participant came forward for an individual draw conducted in the same manner as the sample rounds. The individual draws successfully induced emotional reactions during the experiment; occasional clapping or sighing after each public individual draw was not uncommon.

While the BDM mechanism assumes that agents are EU maximisers (Keller, Segal and Wang 1993), it was unlikely that participants' knowledge of a forthcoming public announcement of their decision influenced their preference for a ticket colour at specific offer prices.

After the experiment, participants filled in a questionnaire on demographic information, gambling and investment behaviour, and colour preference. The instructions, decision form, and selected experiment photos are in Appendix 1.1, 1.2, and 1.3.

1.3 Results

Participants were recruited from the CBESS participant pool of registered student volunteers. Of the 176 participants, 56% were UK-born students, 14% were from other parts of Europe, and 30% were from elsewhere. Thirty 35-minute sessions with six participants on average were conducted from February to March 2013. Total earnings ranged from £2 to £22, and average earnings were £8.50.

We analysed the *switching point* or the offer price at which a participant changed her preference from keeping a ticket to exchanging it for money.[11]

Result 1.1: Switching points increase with the number of observed successful (£20) draws.

The mean switching point in the full sample was £5.05[12] and the standard deviation (SD) was 3.04. The boxplots in Figure 1.1 depict the distribution of switching points

---

[11] For nine participants who had several switching points, the median switching point was used in the data analysis.

[12] The mean switching point exceeds the expected payoff of £4. In valuation tasks that elicit participants' willingness to accept (WTA) an amount of money in exchange for an item they own (e.g., a coloured ticket), exchange aversion (Sugden 2003) contributes to the commonly observed higher valuation. Zhao and Kling (2001) argued that if an agent is uncertain about the value of a good, and she is asked to give up that good now (and forego the opportunity to learn more about it), WTA could be higher than expected value. People's tendency to overvalue a gamble with low probability of winning a large amount is also commonly observed in preference

(interchangeably used here with ticket valuations), conditional on the number of successful draws experienced in the sample rounds. Excluding outliers, switching points in sessions with three successful draws have the widest range, as shown by the distance between the largest non-outlier (top whisker) and smallest non-outlier (bottom whisker). Notice the pattern of increasing median switching point (middle horizontal bar) as the number of £20 draws in the sample rounds goes up.

We also find that the timing of a successful draw has no significant influence on ticket valuation. We compared the switching point of participants in sessions with only one successful draw in earlier rounds (Round 1 to 5; mean=£4.12, SD=2.63; n=41) against those with a £20 draw in recent rounds (Rounds 6 to 10; mean=£4.07, SD=2.87; n=18). The Wilcoxon ranksum test (Siegel and Castellan 1988) revealed no significant difference in the ranking of switching points (z=-0.380, p=0.704; n=59). We also did not find any clustering effect: switching points for tickets with consecutive successful draws in the same ticket colour were not significantly different from tickets without an adjacent £20 draw (z=-1.333, p=0.182; n=18).

Figure 1.1: Switching points by number of successful draws

Successful draws	0	1	2	3	4
Mean switching point (£)	3.8	4.1	4.8	6.7	5.2
Sample size	18	59	40	53	6
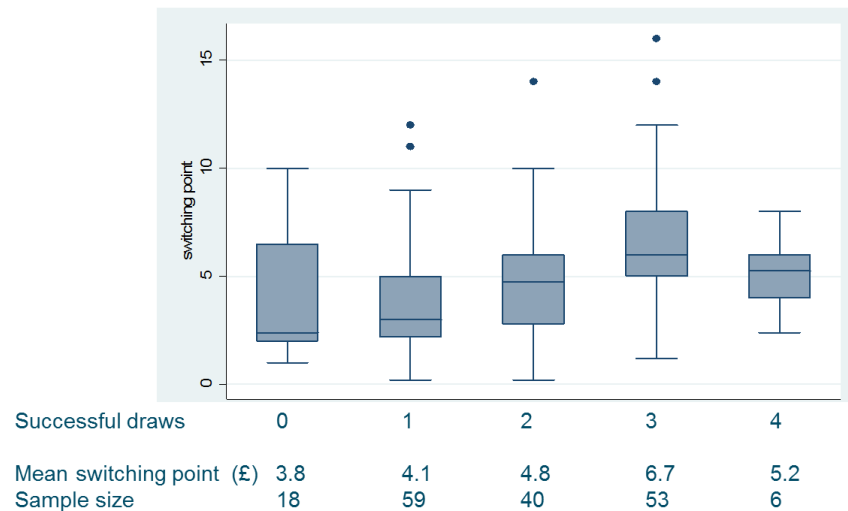
Result 1.2: Participants colour-coded events in the experiment, but not consistently across sessions.

reversal experiments. See Seidl (2002) for a survey on the evidence and explanations for preference reversal.

The experiment was designed to determine if participants coded events into memory using similarity information (colour). We find that average switching point for blue tickets (£5.18) is higher than the valuation for yellow tickets (£4.92). This difference in switching points, however, captures both a colour effect (i.e., colour-coding of events) and the frequency of successful draws [i.e., 60% of successful draws in Part 1 occurred in a yellow round (Y) and 40% in a blue round (B)].

To test for (pure) colour effect, we compared the switching point of participants in sessions with a single successful draw. With only one £20 draw in either a blue or a yellow round, colour was likely to be perceived as a salient cue (Müller, Geyer, Krummenacher and Zehetleitner 2009). The difference in the valuation between blue tickets (mean=£4.42; n=18) and yellow tickets (mean=£3.79; n=41) indicates that participants colour-coded events during the experiment (z=1.850, p=0.064) despite colour's apparent irrelevance in determining ticket earnings. Refer to the cumulative distribution function plotted in Figure 1.2.

However, the observed colour-coding of events is not robust across the subsamples, particularly in sessions with more frequent £20 draws. In the 12 sessions where there were more successful draws in a yellow round than in a blue round (B<Y), the difference in switching points was found to be statistically significant (z=2.339, p=0.019; n=70) but not in the nine sessions where there more successful draws in blue than yellow (B>Y) rounds (z=-0.610, p=0.449; n=54).

Figure 1.2: Switching points by ticket colour

Summary of Switching Points

ticket	n	mean	SD
blue	18	4.42	2.45
yellow	41	3.79	2.88
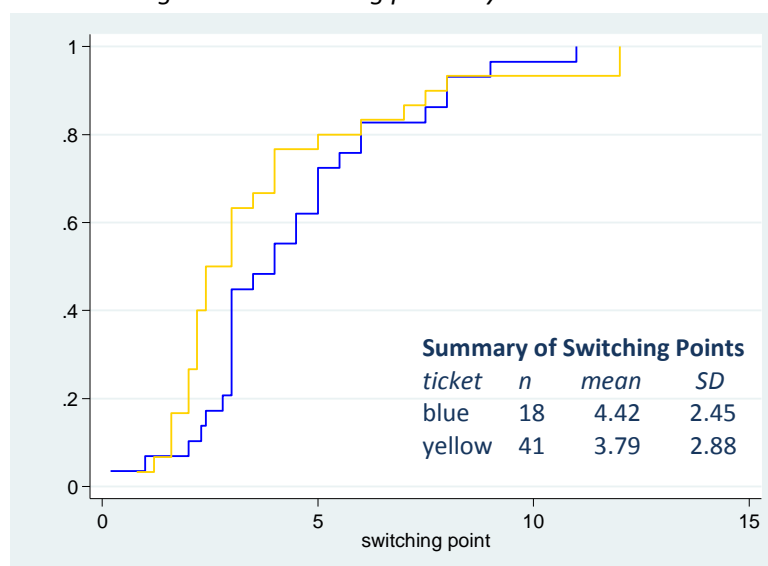
As a control, we compared switching points in the nine sessions where there was an equal number of successful draws in the blue and yellow rounds. In these sessions, the experience in the sample rounds provides no reason to value one ticket more highly than the other. Indeed, colour-coding was not observed in these sessions.[13] The rank-sum test results show that there is no significant difference in subjects' switching point rankings between blue and yellow tickets, whether we include (z=0.587, p=0.557; n=52; Figure 1.3) or exclude (z=-0.086, p=0.931; n=34) sessions with no successful draws.

Figure 1.3: Switching points when B=Y successful draws

Summary of Switching Points			
ticket	*n*	*mean*	*SD*
blue	26	4.59	2.19
yellow	26	4.72	3.61
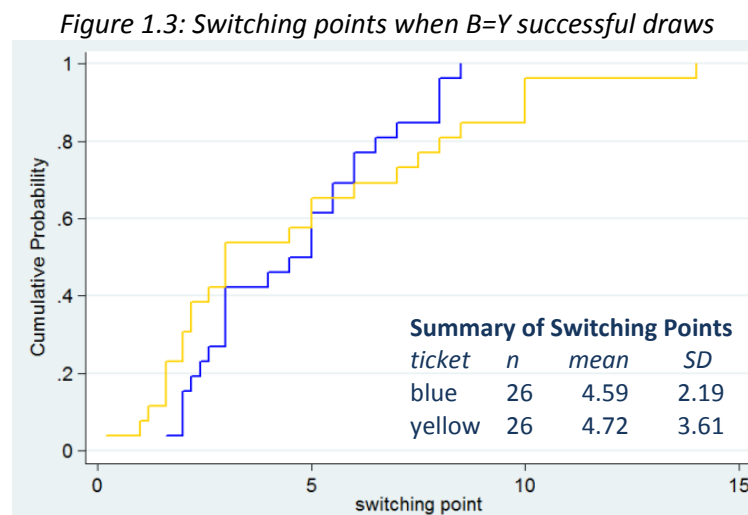
We also controlled for the possible impact of colour preference on ticket valuation in the analysis. Blue was the declared favourite colour by 38% of the participants, while 5% favoured yellow. However, the rank-sum test results reveal no significant difference in switching point (z=-0.103, p=0.918) between participants whose ticket colour in Part 2 of the experiment was similar to their actual favourite colour (n=38) and participants whose ticket colour was not the same as their favourite colour (n=138). While every effort was made to control all salient stimuli in each round during the experiment, there was no certainty on which information participants actually used (or did not use) in the elicitation task.

Result 1.3: The switching point for a "lagging" ticket colour is significantly higher than the switching point for a "leading" ticket colour.

---

[13] Of which, three sessions had no successful draws in either blue or yellow rounds (B+Y=0; n=18); five sessions had one successful draw in the blue and yellow rounds (B+Y=2; n=28), while one session had two successful draws in the blue and yellow rounds (B+Y=4; n=6).

The experiment also attempted to validate whether participants choose an act with the highest-similarity weighted outcome. In our experiment, this translates to participants assigning a higher value on a ticket colour with more successful draws relative to a ticket colour with fewer successes.

To investigate the relationship between switching point and number of successful draws, we categorised tickets as either "leading" or "lagging" depending on the relative number of successful draws; that is, a leading (lagging) ticket colour has more (less) successful draws in Part 1 of the experiment compared to the other ticket colour.[14]

Figure 1.4 shows that participants valued a lagging ticket (mean=£5.77) more highly than a leading ticket (mean=£4.66). The rank-sum test results (z=-2.131, p=0.033; n=124) confirm a statistically significant difference in switching points.

Figure 1.4: Switching points by ticket category
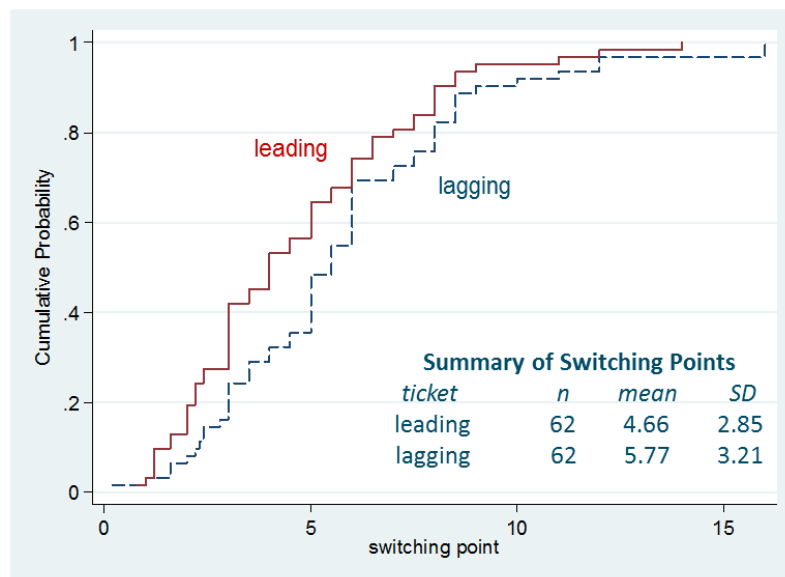
We regressed the logarithm of switching point against number of successful draws on *own* ticket colour, number of successful draws on *other* ticket, ticket colour (0=blue, 1=yellow), and gender (female=0, male=1). Since switching point is truncated at £0.20,

---

[14] Both ticket colour and the relative number of successful draws are used to differentiate leading from lagging tickets. A difference in the valuation between leading and lagging tickets also implies colour-coding.

the dependent variable used was log(switching point). Robust standard errors are reported on Table 1. Notice that the β coefficient on the other ticket colour is higher than own ticket. While the number of successful draws on one's own ticket is positively related to the switching point, the influence of the number of successes on the other ticket appears stronger, all else held constant. However, the difference in the coefficients on *own* ticket and *other* ticket is not statistically significant: $F(1, 171) = 0.60$; $p = 0.4395$.

Table 1: Ordinary least squares regression results
dependent variable: log(switching point)

	coeff.	std.error
male	0.2467**	0.1075
own	0.1660***	0.0623
other	0.2269***	0.0560
yellow ticket	-0.0979	0.1009
constant	0.9801***	0.1228

Number of obs = 176
F(4, 171) = 10.32
Prob > F = 0.0000
R-squared = 0.1582
*** <0.01, **<0.05, *<0.10

Given these results, the failure of CBDT in explaining participants' decisions in our experiment leads to an *exhibit*.[15] Our results are consistent with the gambler's fallacy or the erroneous belief that a lottery which had a series of losses was bound to reverse the pattern of its past outcomes (Rogers 1998). This means that in our experiment, participants expected a lagging ticket colour to reverse its poor performance in Part 2 of the experiment and participants consequently valued a lagging ticket more highly. A leading ticket colour, on the other hand, was not expected to sustain its success and was then given a lower value.[16]

Our experimental design provided a fair chance for either case-based or Bayesian reasoning to emerge. In fact, given the apparent irrelevance of ticket colour in

---

[15] An *exhibit* is defined by Bardsley, Cubitt, Loomes, Moffatt, Starmer, and Sugden (2010) as a "replicable experimental design that reliably produces some interesting result" (p.156).

[16] However, we do not find any significant difference in participants' expectation on the chance of a £20 draw (z=-0.137, p=0.891) between participants who owned a leading ticket (n=62) and those with a lagging ticket (n=62) in Part 2 of the experiment (round 11).

determining outcomes at each draw, ignoring ticket colour was an easy strategy for a Bayesian player. However, the results show that participants were neither case-based nor Bayesian.

While the experiment effectively controlled the set of information presented to the participants, there was no guarantee that participants will not use alternative rules when deciding (Grether 1992). For instance, given the random order of the tickets played and the value of the payoffs, participants may be influenced by the recency effect.

The representativeness heuristic (Kahneman and Tversky 1972) where the likelihood of an uncertain event depends on the similarity of the event to the parent population and its perceived randomness was offered as an explanation for two manifestations of the recency effect (Ayton and Fischer 2004; Croson and Sundali 2005; Sun and Wang 2010): the gambler's fallacy (negative recency) and the hot hand belief (positive recency).[17]

Given the belief that chance is unpredictable but fair, randomness is expected even in small samples. This belief in local representativeness drives the gambler's fallacy. For example, in a fair coin toss, a series of heads is expected to be followed by a tail, as if chance is self-correcting (Kahneman and Tversky 1972) and the coin has "some sort of memory and moral sense" (Rogers 1998, p.119). Estes (1964) argued that individuals'

---

[17] The representativeness heuristic was also offered as an explanation for the hot hand fallacy. The erroneous belief that the same leading outcome will continue (at least in the short run) is formed when a small sample of consecutive wins is misconstrued to be representative of the population. Gilovich, Vallone, and Tversky (1985) suggested that the overestimation of the positive correlation of random (leading) outcomes may be due to memory bias (i.e., a series of successes is more salient than alternating outcomes), while human habit to find sequential patterns is consistent with the tendency to offer deterministic explanations despite the uncertainty of outcomes (Wood 1992; Sun and Wang 2010). A series of wins may be attributed to the agent as being "hot" or remarkably skilful rather than to chance. While empirical evidence does not support the hot hand either in sports or in the stock market, the belief in the hot hand seems prevalent (Ayton and Fischer 2004; Gilovich, Vallone, and Tversky 1985; Offerman and Sonnemans 2004; Wood 1992).

susceptibility to the gambler's fallacy results from generalizing perceived patterns of reversal in real-life outcomes where sampling happens without replacement.

The gambler's fallacy has been found to emerge especially in tasks involving inanimate objects perceived to generate random outcomes, tasks where limited analytical skill is required in decision-making, and information is presented sequentially (Hogarth and Einhorn 1992; Ayton and Fischer 2004; Burns and Corpus 2004; Croson and Sundali 2005; Barron and Leider 2010). While our main research objective was to compare the predictive power of CBDT against Bayesian reasoning, the decision setting in our experiment inadvertently led to the emergence of the gambler's fallacy, hinting on its prevalence (Croson and Sundali 2005).

As an additional result, we found that male participants (mean=£5.65; n=87) had significantly higher (z=-3.082, p=0.002) switching points than their female counterparts (mean=£4.46; n=89)[18], supporting the previously documented higher risk aversion among females. Croson and Gneezy (2009) showed that other than gender difference in terms of risk preference, men also tend to be more overconfident than women in their perceived likelihood of success. Ayton and Fischer (2004) showed that predictions of participants with a high level of self-confidence are consistent with the gambler's fallacy. Now, we ask: are male participants in our experiment more susceptible to the gambler's fallacy than females?

There is no consensus among past empirical studies whether male participants are more vulnerable to the gambler's fallacy. For example, using online state lottery gambling data, Suetens and Tyran (2012) showed that the gambler's fallacy is apparent in men but not in women, while Dohmen, Falk, Huffman, Marklein, and Sunde (2009) found that in a hypothetical coin toss, women are more susceptible to having biased beliefs.

---

[18] Among the demographic variables we regressed against the logarithm of switching points, only gender (0=female, 1=male) was statistically significant: males generally gave higher ticket valuations. However, the gender difference disappears when interacted with the number of successful draws: the interaction terms (i) gender and the number of successful draws in one's own ticket, and ii) gender and the number of successes in the other ticket are found to be statistically insignificant.

Among participants who played a lagging ticket in our experiment, male participants (mean=£6.32, n=35) valued lagging tickets more highly than female participants (mean=£5.07, n=27). Splitting the sample by gender, we find that the difference in the valuations between leading and lagging tickets is not statistically significant, for either the all-female subsample (z=-0.841, p=0.400; n=65) or the all-male subsample (z=-1.471, p=0.141; n=59). We do not have compelling evidence to suggest that there is a difference in vulnerability to the gambler's fallacy between males and females.

1.4 Conclusion

We designed an experiment to test the predictive power of CBDT vis-a-vis Bayesian reasoning. We attempted to validate two predictions of CBDT, namely: i) decision-makers encode and retrieve past experiences using subjective similarity; and ii) agents choose an act with the highest similarity-weighted outcome. Our experimental design, which induced features-based similarity and objective uncertainty, provided a fair chance for either case-based or Bayesian reasoning to emerge.

Unlike earlier CBDT experiments which produced results generally consistent with CBDT, our results show that participants were neither case-based nor Bayesian. The participants used a similarity cue (colour) in the ticket valuations, but their ticket valuations were the opposite of CBDT's similarity-weighted prediction: tickets with fewer successes were valued more highly than tickets with more successes. The results reveal an *exhibit* consistent with the gambler's fallacy, which undermines the predictive power of CBDT.

The pattern in our results is strikingly different from past CBDT experiments. This may suggest that the manner of learning cases matters to subsequent decisions. Our participants sampled from experience while participants in other experiments (Bleichrodt, Filko, Kothiyal, and Wakker 2015; Grosskopf, Sarin and Watson 2015) were shown a description of cases. A decision was made on the same screen and participants did not have to recall cases from memory.

In addition, the pure randomness of ticket outcomes in our experiment (i.e., the live draw of a ball from the bingo cage) may have inadvertently triggered the gambler's fallacy to manifest. However, controlling participants' prior on the probability of a successful outcome, e.g. when participants know that an item played either has a 10% or 30% chance of a successful draw, could reduce the likelihood for the gambler's fallacy to emerge. This variation in the experimental design was implemented in an experiment presented as Chapter 2 of this thesis.

Chapter 2

Correlation neglect and act separability in asset valuation: an experiment

In a two-part experiment designed to validate act separability, participants played two-armed bandits framed as coloured game boards (blue or yellow). If participants are Bayesian, decisions would correspond to the known game board correlations (positive, independent or negative). Under act separability, participants may neglect correlation information and decisions are inconsistent with Bayesian reasoning. If this holds, each game board is evaluated over a set of different cases and a separate memory for each game board is maintained so that a decision on a blue game board is not influenced by outcomes on the yellow game board. We find that participants' decisions in the positive and independent treatments are qualitatively Bayesian, but decisions in the negative treatment cannot be explained either by Bayesian reasoning or by CBDT. However, evidence suggests that participants systematically used past outcomes in forming expectations, but judgment did not directly translate into the expected decision.

2.1 Introduction[19]

When the distribution of outcomes is uncertain, a Bayesian agent uses feedback information gathered from previous experience to modify the prior probability distribution. Updating the conditional distribution of outcomes based on past experience has been proposed as the learning process behind expectations formation, which in turn leads to a predicted action (Cyert and DeGroot 1974).

We created an experimental environment where participants had the opportunity to learn the correlation in the outcomes between two assets. Participants were given the same prior information on the possible distribution of outcomes, and past experience on either asset was informative for the continuous updating of the conditional distribution of outcomes. Actions consistent with the known asset correlation, past

---

experience and revealed expectations favour Bayesian reasoning. On the other hand, an indication that participants did not use correlation information and experience of past outcomes in the experiment supports correlation neglect.

People's tendency to neglect asset correlation has been demonstrated in the literature. For instance, in an asset allocation experiment, Kallir and Sonsino (2009) showed that participants focused their attention on individual asset returns and the resulting portfolio decisions did not take into account return correlations. Eyster and Weizsaecker (2011) also demonstrated that even when equipped with correlation information, participants regarded assets independently and resorted to the *1/n heuristic* (or naïve diversification) when allocating investment funds to individual securities. Similarly, in a hypothetical investment choice experiment, Hedesström, Svedsäter and Gärling (2006) observed that participants focused on individual asset volatility rather than on portfolio volatility. Resulting portfolios were inappropriately diversified and had higher volatility. Gubaydullina and Spiworks (2009) also observed that, in addition to correlation neglect, participants became overly confident about their decisions when they were armed with more information even if the available information was irrelevant to the portfolio decision. This overconfidence is aligned with the observation that retail investors tend to (mis)perceive patterns in market data even if objectively non-existent (De Bondt 1998).

Correlation neglect has been found to be sensitive to the magnitude of the stakes involved in the decisions. In their portfolio experiments, Kroll, Levy and Rapoport (1988) showed that while participants were aware of the correlation in stock returns, correlation information was not reflected in their portfolio choices. However, when the stakes were significantly increased, participants managed to effectively diversify their asset holdings and the resulting portfolio choices were closer to the predictions of mean-variance optimization.[20]

Correlation neglect may be explained by CBDT. Gilboa and Schmeidler (1995; 2001) proposed CBDT to complement EUT. Unlike EUT where a decision-maker is assumed to know all the possible consequences arising from all alternative actions, CBDT is much

---

[20] Under Markowitz's (1952) modern portfolio theory, there is a combination of risky assets that provides the highest expected portfolio return given a minimum variance of returns.

less cognitively demanding and allows for ignorance. A case-based decision-maker chooses from a set of feasible actions based on knowledge gathered from actual decisions in the past. When an agent has no prior experience, no belief about a decision problem is assumed so that under ignorance, where neither outcomes nor probabilities are known, a decision-maker randomly chooses between alternative actions.

In CBDT, an experience is encoded in memory as a case consisting of a problem, act and result. Given a new problem, a decision-maker assesses the similarity of the current situation with past problems encountered, and then recalls the actions taken in those similar situations. The decision-maker then chooses the act with the highest similarity-weighted sum of outcomes. CBDT only requires a decision-maker's ability to encode into and retrieve past experience from memory, and to assess the similarity of the current problem with past situations encountered.

Each act is evaluated over a different set of cases so that a decision-maker's memory of cases on *act A* is disjoint from her memory of cases on *act B*. The act separability axiom proposes that decision-makers maintain separate memories of alternative actions taken in the past. Since an act is evaluated over past outcomes on that act, experience from other acts is not taken into account. Decision predictions under correlation neglect and act separability therefore coincide.

Act separability has an important implication on asset selection and portfolio diversification. If each asset is evaluated in isolation, an investor would fail to account for the correlation between available investment instruments and other assets already in the portfolio. If valid, act separability would imply investors' difficulty in understanding correlations, and by extension, mean-variance analysis commonly used in the construction of efficient portfolios (Markowitz 1952). According to the modern portfolio theory, it is optimal for an investor to choose the portfolio that gives the highest expected return and the minimum variance of returns that can be achieved through portfolio diversification. Effective portfolio diversification does not only mean investing in different assets; it also requires selecting securities with low correlations. Therefore, for mean-variance optimisation to work, investors need to understand and effectively utilise correlation information.

We conducted a two-part experiment that involved two (un)correlated bandit arms and allowed a fair chance for either Bayesian reasoning or case-based reasoning to emerge. Each arm was framed as a coloured game board. The blue game board (B) was always played on the left side of the computer screen, while the yellow game board (Y) was played on the right. The game boards were positively correlated (correlation=1), uncorrelated (correlation=0), or negatively correlated (correlation<0). [21] These correspond to the three treatments in the experiment.

In Part 1 of the experiment, participants experienced sample plays of the blue and yellow game boards to create a history of cases on B and Y. In Part 2, participants were each assigned one game board and engaged in an incentive-compatible task to reveal the minimum amount of money they were willing to accept (WTA) in exchange for the outcome of the play of their game board. If act separability is obeyed in all three treatments, outcomes on B will not influence expected outcomes on Y.

Our approach is different from past studies on correlation neglect and act separability. In earlier studies on correlation neglect, participants were typically informed of the return distributions of assets and their correlations. In the real-estate investment experiment of Bleichrodt, Filko, Kothiyal and Wakker (2015), participants were instructed to treat each property independently.[22] Meanwhile, in our experiment participants were informed of the correlations between the stylised investments (blue

---

[21] Similarity under CBDT ranges from 0 (no similarity) to +1 (perfect similarity). Our experimental design, on the other hand, uses correlation (ρ) equal to +1, 0, and <0. Even if CBDT may not account for negative correlation, under act separability, we expect the divergence in decisions when $\rho<0$ to be at least as large as the observed difference when $\rho=0$.

[22] Our experimental design, which used game board colour to induce feature-based similarity, allowed a test of Bayesianism vis-à-vis act separability. The sample rounds during the experiment were opportunities for participants to learn from *experience* and create cases in memory. Meanwhile, participants in Bleichrodt, Filko, Kothiyal and Wakker's experiment were presented with a *description* of past cases, followed by a choice between two gambles characterised by geographic location and type of dwelling. Participants knew that the past rounds were irrelevant to the current pair of real-estate gambles. The authors also devised a nonparametric method to calculate similarity weights, and tested whether or not act separability is observed, i.e., if a preferred gamble in one round is not influenced by memories encountered in previous rounds.

or yellow game board) but the distributions of game board earnings were unknown. The sample rounds played by the participants were opportunities to create a memory of cases during Part 1 of the experiment. The sample rounds also gave them an idea of the distribution of game board earnings. However, it was up to the participants to decide how to treat the information gathered from each sample round played. By design, our participants' decisions may be determined by either Bayesian or case-based reasoning.

In the positive treatment, game board valuations by colour (blue or yellow) and by outcome (lagging or leading) are not significantly different. These results, while consistent with the Bayesian prediction, are contrary to the findings in the ticket experiment[23] presented in Chapter 1.

Meanwhile, in the independent treatment where the Bayesian and case-based predictions coincide, valuations are significantly different between lagging and leading game boards, and between blue and yellow game boards. Valuations appear to be jointly driven by the number of experienced hits on the game board played and a bias towards yellow.

In the negative treatment, the observed game board outcomes in the sample rounds cannot account for the pattern in valuations; neither Bayesian nor case-based reasoning can explain this result. While evidence suggests that the participants accounted for correlation information in forming expectations, the known correlation was not systematically reflected in the actual decisions.

The rest of the paper is organised as follows. Section 2.2 describes the details of the experimental design. Section 2.3 discusses the results and Section 2.4 concludes.

---

[23] The positive treatment replicates the ticket experiment presented in Chapter 1. In the ticket experiment, participants played two coloured tickets (blue or yellow) which paid earnings based on live random draws from a single bingo cage with an unknown distribution of £20 and £0 balls. While colour was salient in each round, it was uninformative for the purpose of Bayesian updating and was expected to be ignored by a rational player. With a salient colour cue that is likely to trigger conscious similarity assessment and facilitate recall of past events (Gilboa and Schmeidler 2001), a case-based thinker may colour-code and retrieve memories of past experiment rounds using colour despite its irrelevance.

2.2 Experimental Design

The experiment, formulated as a two-armed bandit with added similarity cues, used a mixed design to allow between-subjects and between-treatments comparison. Participants played two saliently distinct bandit arms framed as coloured game boards with both colour and spatial features (i.e., the blue game board was always encountered on the left side of the computer screen, while the yellow game board was consistently played on the right side). Each game board contained 100 numbered boxes (1 to 100) and each box had a pre-determined value of either £0 (losing box) or £20 (winning box). To control prior beliefs, participants knew that forty out of the 200 boxes had a pre-assigned value of £20.

Before the experiment, forty box numbers were pre-drawn (without replacement) from a bingo cage with 100 balls corresponding to the game board boxes. The winning boxes were then distributed in two sealed envelopes in the order the box numbers were drawn: one envelope had thirty winning box numbers while the other had ten. During the experiment, the envelopes were placed in a bag and one participant was asked to pick and mark each envelope as instructed by the experimenter.

There were three treatments corresponding to different correlations between the blue (B) and yellow (Y) game boards. In the *positive treatment*, the same set of winning boxes was assigned to B and Y; the envelope picked by the participant was marked "blue and yellow". While the experimenter continued to read the instructions aloud, the lab assistant programmed the listed winning boxes in z-tree (Fischbacher 2007). Since the game boards were perfectly positively correlated, outcomes on one game board were informative of outcomes on the other game board. This treatment replicated our earlier ticket experiment. In the *negative treatment*, B and Y were assigned different winning boxes. The first envelope picked by the participant was marked "blue" and its contents were programmed into the blue game board, while the second envelope was marked "yellow" and the winning boxes were assigned to the yellow game board. If one game board had thirty winning boxes, participants knew that the other automatically had ten. Outcomes in one game board were also informative of outcomes on the other game board: a winning box revealed implied one less winning box for the other game board. Meanwhile, participants in the *independent treatment*

knew that there were two sets of forty winning boxes, one for each game board. A participant picked one envelope from the first set and marked that envelope "blue" and then picked an envelope from the second set and marked that envelope "yellow". Since there were two sets of winning boxes, the outcomes on one game board were uninformative about the other game board. All envelopes were posted on the wall after the lab assistant programmed the winning boxes into the game boards. Participants were welcome to inspect the envelopes at the end of the experiment.

After the instructions were read aloud, participants were shown a sample screen, a series of questions pertaining to the game boards, and the relationship of the winning boxes on the blue and yellow game boards. If an incorrect answer was given, the participant was instructed to refer to the relevant section on the printed instructions before re-attempting to answer the same set of questions. If a wrong answer was supplied on the second try, the correct answer was provided. In each treatment, about five percent of participants answered at least one question incorrectly after the second attempt.

The experiment had two parts. Part 1 consisted of ten sample rounds where participants played five blue rounds and five yellow rounds presented in random order. Participants in a session encountered the same sample rounds. While no actual earnings were paid in Part 1, participants knew that the ten rounds were opportunities for them to learn as much as they can about the two game boards. At each round, all participants were shown either the blue or yellow game board. After the participants indicated what they thought was the chance that the game board would reveal a winning box ("hit") in that round, the experimenter drew one ball (without replacement) from the bingo cage and announced the number printed on the ball. The lab assistant then inputted the box number into the computer before participants were allowed to click on that box to reveal the value of that box on the game board. If the announced box number had a value of £20, the box on the game board displayed £20 and it was shaded green; otherwise, the box displayed £0 and it was shaded red.

After everyone had clicked on the announced box number, a screen summarising what participants learned about the game boards in that round was displayed for a few seconds before the experiment moved on to the next round. On the summary screen,

participants were reminded of the game board they just played and the value of the box they just clicked. For example, in the positive treatment, if box number 1 (a winning box) was clicked on the blue game board on that sample round, that box was shaded green and displayed £20. On the summary screen, the statement: "1 is a winning box on the blue game board" appeared below the blue game board. On the right side of the summary screen was a statement: "1 is a winning box on the yellow game board." At the end of Part 1, participants had seen the value of 10 boxes.

Throughout Part 1, the computer screen also displayed a header that constantly reminded participants of the correlation between the two game boards. The instructions and sample screens for the three treatments are in Appendices 2.1 to 2.3.

In Part 2, each participant was randomly assigned one game board. Half of the participants played the blue game board, while the other half played the yellow game board. Participants knew that their game board had the same set of winning boxes as the game board they played in Part 1. The game board gave the participants an opportunity to earn money in addition to the £2 show-up fee by deciding either to keep the game board and receive the earnings from the play of that game board, or to exchange the game board for an amount of money (called the offer price).

At the beginning of Part 2, thirty-five sealed envelopes, each containing one possible offer price ranging from 20 pence to £20, were placed in a bag. A participant then picked one envelope and the sealed envelope was posted on the wall. All thirty-five offer prices were listed in a decision form. At each price, participants decided whether they prefer to keep their game board and receive the earnings from the play of that game board, or they prefer to exchange the game board and receive that offer price. This incentive-compatible elicitation task was conducted to determine the minimum amount each participant was willing to accept in exchange for the outcome of the play of their assigned game board (Becker, DeGroot, and Marschak 1964).

Before the individual draws were conducted, all 100 balls were returned in the bingo cage. The experimenter approached each participant, drew one ball from the bingo cage and showed the box number printed on the ball. The participant then clicked on that box to reveal its value. If a participant decided to keep the game board at the offer

price, the participant received the value of the selected box plus the show-up fee. If a participant chose to exchange the play of the game board, the participant received the offer price plus the show-up fee. Immediately after each individual draw, the lab assistant asked the participant to fill in a demographic questionnaire and a receipt form.

If act separability holds, a case-based decision-maker will value a game board based solely on the experience on that game board. Meanwhile, in the valuation task, a Bayesian decision-maker will use correlation information and the number of game board hits. In the positive treatment, a Bayesian participant will ignore game board colour and will focus on the total number of hits on the two game boards. In the negative treatment, a participant will use the difference in the number of game board hits to determine whether that game board likely has ten or thirty winning boxes. Since the game board outcomes are uncorrelated in the independent treatment, the case-based and Bayesian predictions coincide. In the valuation task, participants are expected to use only the observed outcomes on their assigned game board.

2.3 Results

We conducted a total of thirty sessions in March 2014; there were ten sessions for each treatment. Each session had six to eight participants and lasted 45 minutes. Average earnings were £8.40 and ranged from £2 to £22. All 224 participants were recruited from the CBESS subject pool. Fifty-two percent were male, fifty-five percent were native English speakers and fifty-two percent were British.

Participants' *switching point* or the minimum offer price at which a participant changed preference from keeping the game board and receiving the earnings from the play of that game board to exchanging the game board for an amount of money was analysed. Participants' median switching point was recorded among participants with multiple switching points (n=12). The terms switching point, valuation and WTA are used interchangeably throughout this chapter.

Figure 2.1: Unconditional switching points by treatment

	positive	negative	independent
Mean	5.38	5.35	5.41
SD	3.27	3.06	2.92
Sample size	75	75	74
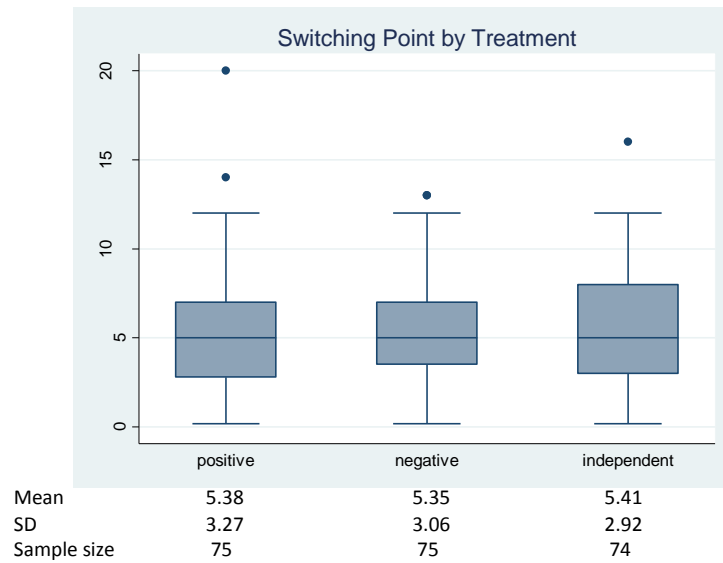
The boxplots in Figure 2.1 summarise unconditional switching points in the positive, negative, and independent treatments. These do not take into account the difference in the number of game board "hits" or the number of revealed winning boxes observed in Part 1 of the experiment. Excluding outliers, the range of switching points (or the distance from the top whisker to the bottom whisker) is similar across the treatments. The common median switching point of £5 (middle horizontal line) is higher than the expected game board payoff of £4, similar to the switching points in the ticket experiment presented in Chapter 1.[24] Meanwhile, the average switching points in the positive, negative, and independent treatments are not statistically different based on the Wilcoxon ranksum test (Siegel and Castellan 1988).[25]

Result 2.1. Switching points in the independent treatment indicate a significant colour effect not found in the positive and negative treatments.

---

[24] In experimental tasks where participants are asked to give up an item they own (e.g., a coloured game board), exchange aversion (Zhao and Kling 2001; Sugden 2003) commonly leads to higher valuation. The tendency to overvalue a gamble with low probability of winning a large amount is also commonly observed in preference reversal experiments (Seidl 2002).

[25] The switching points across the three treatments are slightly higher than the switching points in the ticket experiment presented in Chapter 1 (mean=5.05, n=176), albeit not statistically different based on the ranksum test.

Participants within each session experienced the same game board hits in the sample rounds. As a control, we plotted the cumulative distribution function (CDF) of conditional switching points in sessions where participants saw the same number of hits on the blue and yellow game boards (B=Y). A Bayesian player who has seen the same number of hits on B and Y in the sample rounds is unlikely to value one game board more highly than the other.

Figure 2.2.A plots the CDF of switching points for sessions where B=Y in the positive treatment. While the average valuation is higher for the blue game board, the Wilcoxon ranksum test indicates that participants did not value B and Y differently (z=0.494, p=0.6215, n=23). In the negative treatment (Figure 2.2.C), we also find no significant difference in game board valuations in sessions where B=Y.

However, in the independent treatment, switching points in sessions where B=Y diverge. Figure 2.2.B shows a higher valuation on Y than on B that appears to be driven by a colour bias in favour of the yellow game board (z=-2.070, p=0.0384, n=23).[26] We do not observe this colour bias in either the positive or negative treatment.[27]

Figure 2.2.A: Switching points by game board colour in the positive treatment (B=Y)

---
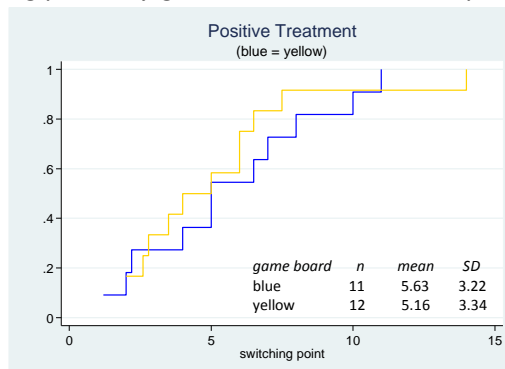
[26] Possibly, this could also be a case of a bias that favours playing right (vs. playing left).

[27] We find that favourite colour does not influence game board valuation. Switching points of participants who played a game board with the same colour as their self-reported favourite colour are not significantly different from the switching points of participants who prefer other colours.

Figure 2.2.B: Switching points by game board colour in the independent treatment (B=Y)
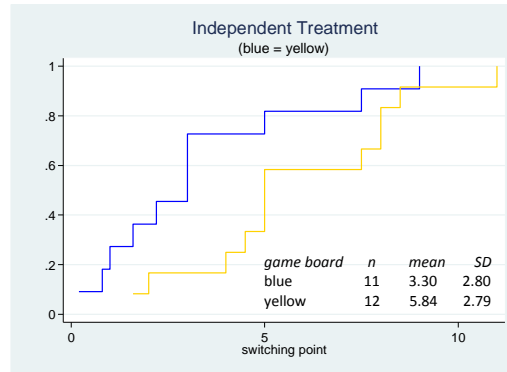
Figure 2.2.C: Switching points by game board colour in the negative treatment (B=Y)
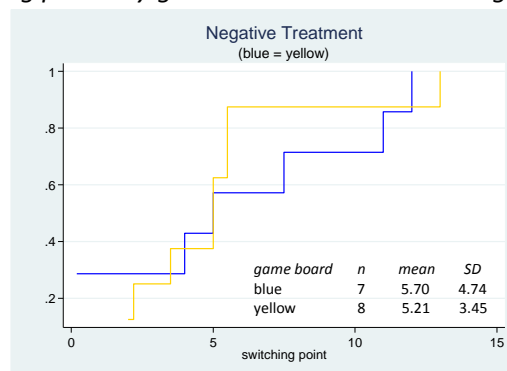
Result 2.2. Switching points between lagging and leading game boards are significantly different in the independent treatment, but not in the positive and negative treatments.

We categorised game boards as either leading or lagging, based on the difference between own-hits and other-hits. A *leading* game board has more hits than the other game board, while a *lagging* game board has fewer hits. In sessions where B and Y have the same number of hits, no game board category is generated.

In the positive treatment, the average switching point for leading game boards (mean=6.03, n=21) is higher than for lagging game boards (mean=4.67, n=24). However, the cumulative distributions depicted in Figure 2.3.A do not show a significant divergence in the valuations; this is confirmed by the ranksum test results (z=-1.062, p=0.2885). Consistent with Result 2.1, we do not have evidence to suggest that participants in the positive treatment valued leading and lagging game boards differently.

This result is starkly different from the findings in our ticket experiment where participants played coloured tickets (blue or yellow) with an unknown probability of a "hit" (earnings of £20). In Chapter 1, we showed that participants' WTAs were the *opposite* of CBDT's prediction: a ticket colour with fewer hits was valued more highly than a ticket colour with more hits. To reduce the likelihood for the gambler's fallacy to emerge in our game board experiment, we set the prior on the likelihood of a hit on the game boards (i.e., the chance of a hit was either ten or thirty out of a hundred) and the winning boxes on each game board were pre-assigned so that the value of a box was not deemed to be purely random (unlike in the ticket experiment). The changes we implemented in our experimental design (vis-à-vis the ticket experiment) appear to have successfully muted the gambler's fallacy.

In the independent treatment, we observe a statistically significant divergence (z=-3.241, p=0.0012) in switching points (Figure 2.3.B) between leading (mean=6.97, n=26) and lagging (mean=4.38, n=26) game boards.

Given the known correlation between game board outcomes in the negative treatment, we would expect a larger divergence between lagging and leading game boards in the negative treatment vis-à-vis the independent treatment. Figure 2.3.C shows that while the cumulative distribution of switching points of the leading game board (mean=5.98, n=31) lies to the right of the lagging game board (mean=4.78, n=28), the divergence is not statistically significant (z=-1.535, p=0.125).[28] Neither Bayesian reasoning nor act separability can explain the pattern in switching points between lagging and leading game boards under the negative treatment.

---

[28] Note that Figures 2.2.C (B–Y<0) and 2.3.C (lagging vs. leading game boards) plot the same data. In 8 out of the 10 sessions in the negative treatment, hits on Y exceeded the number of hits on B. In the two other sessions, B and Y had the same number of hits.

Figure 2.3.A: Switching points by game board type in the positive treatment

Positive Treatment
(lagging vs leading)

game board	n	mean	SD
lagging	24	4.67	2.60
leading	21	6.03	3.98

----- lagging   —— leading
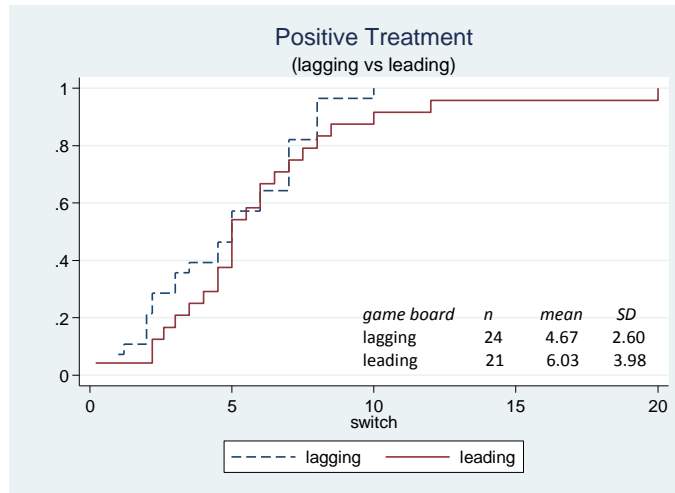
Figure 2.3.B: Switching points by game board type in the independent treatment

Independent Treatment
(lagging vs leading)

game board	n	mean	SD
lagging	26	4.38	2.29
leading	26	6.97	3.20

----- lagging   —— leading
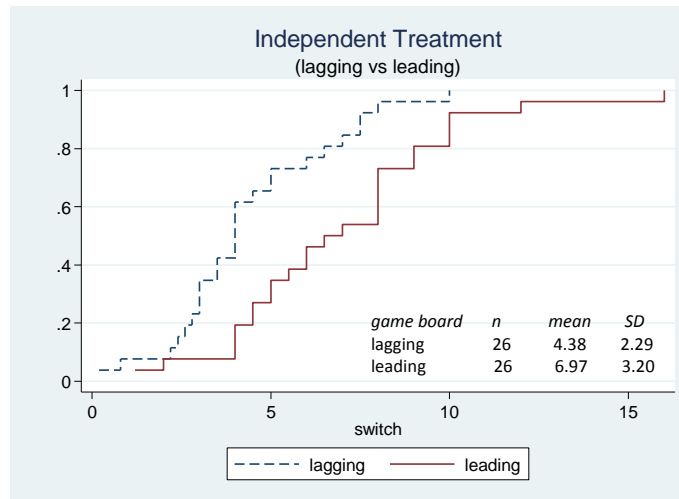
Figure 2.3.C: Switching points by game board type in the negative treatment

Negative Treatment
(lagging vs leading)

game board	n	mean	SD
lagging	28	4.78	2.58
leading	31	5.98	2.61

----- lagging   —— leading
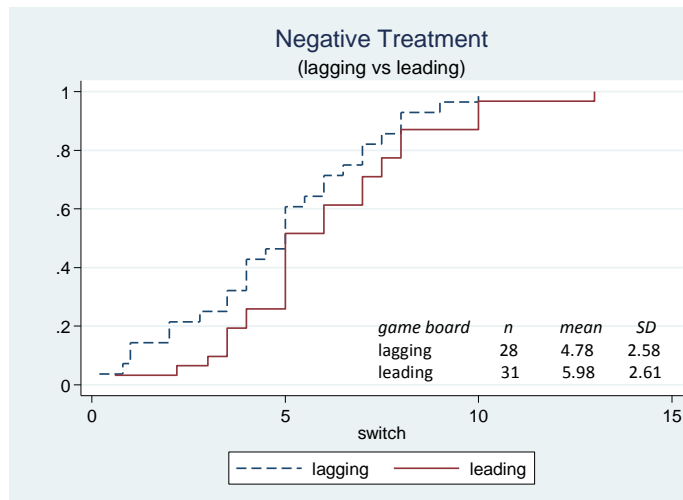
Table 2.1 summarises the regression results on the switching points. Two sets of regression results for each treatment are reported. The regressors in the first equation include *own-hit*, *other-hit*, *game board colour*, and *mother tongue*. In the second equation, regressors are *own plus other*, *own minus other*, *game board colour*, and *mother tongue*. The two equations are linear in its coefficients (switching points are regressed on the same variables), so that the respective constant, parameter for *game board colour*, and parameter for *mother tongue* in the two equations are the same. Given the assigned game board colour in Round 11, *own-hit* refers to the number of £20-boxes revealed on that game board colour during the sample rounds; *other-hit* refers to the number of £20-boxes revealed on the other game board. Game board colour pertains to the game board played in Round 11 (0=blue, 1=yellow), while mother tongue refers to each participant's first language (0=non-English, 1=English). Since the switching points are truncated at £0.20 and are positively skewed, *log(switching point)* is used as dependent variable.

In the positive treatment (column 2), switching points significantly increase with the number of hits on own game board (*own-hit*) and hits on the other game board (*other-hit*). The coefficient on *own-hit* is larger than *other-hit* but the difference is not statistically significant (F=0.38, p=0.5381). Given a perfect positive correlation in game board outcomes, Bayesian reasoning predicts that the switching points will increase with the sum of *own-hit* and *other-hit*, and any difference between *own-hit* and *other-hit* will be ignored. The regression results support the Bayesian prediction.

In the independent treatment, a Bayesian participant will find *own-hit* informative and will tend to ignore past experience on the other game board. If act separability holds, the prediction of CBDT and Bayesian reasoning coincide: *own-hit* matters while *other-hit* does not. Indeed, column 3 shows that the coefficient on *own-hit* is positive and statistically significant while the coefficient on *other-hit* is not statistically different from zero.

In the negative treatment, participants know that if one game board has thirty winning boxes, the other game board automatically has ten (and vice versa). A negative correlation between game board outcomes means that a hit (£20) on the blue game board is a miss (£0) on the yellow game board. Therefore, a Bayesian participant will

account for both *own-hit* and *other-hit* to determine whether the assigned game board in Round 11 is likely to have more hits or fewer hits than the other game board. Consequently, a game board with more hits will be valued more highly than a game board with fewer hits. Meanwhile, act separability implies that a case-based participant will rely only on *own-hit*.

However, Table 2.1 (column 4) shows that none of the regression coefficients on *own-hit*, *other-hit*, the sum or difference between game board hits is statistically different from zero. This is a puzzling result that neither case-based nor Bayesian reasoning can explain.

Table 2.1. Ordinary least squares results
dependent variable: log(switching point)

	Positive		Independent		Negative	
own	0.3401***	(0.0873)	0.2286**	(0.1083)	-0.0811	(0.1978)
other	0.2492**	(0.0985)	-0.0317	(0.1130)	0.0758	(0.1401)
own + other	0.2947***	(0.0571)	0.0985	(0.0841)	-0.0026	(0.1258)
own − other	0.0455	(0.0735)	0.1302*	(0.0720)	-0.0785	(0.1163)
yellow game board	-0.0364	(0.1436)	0.4842***	(0.1652)	0.5443	(0.4697)
native English speaker	-0.4763***	(0.1515)	0.0650	(0.1904)	-0.3273	(0.2040)
constant	1.2327***	(0.1629)	0.9571***	(0.2864)	1.3369***	(0.3848)
Observations	75		75		74	
F-statistic	8.00		4.65		1.62	
p-value	0.0000		0.0022		0.1786	
R-squared	0.2859		0.1559		0.0934	

Robust standard errors are in parentheses. *** <0.01, **<0.05, *<0.10.

Notice that the explanatory power of the fitted model varies significantly across the three treatments, indicating a difference in participants' use of experienced hits in the sample rounds. In the positive treatment, both own-hit and other-hit are statistically significant, while only own-hit is significant in the independent treatment. However, neither game board hits nor colour can explain the variations in switching points in the negative treatment. These results suggest that the game board correlation imposed

during the experiment led participants to treat the information on hits in the sample rounds differently.

Our participants correctly answered the comprehension questions at the start of the experiment. In addition, we show below (under Result 2.3) that participants systematically used the information on hits to form a game board expectation. We therefore rule out confusion as an explanation for the insignificant difference in switching points between game boards under the negative treatment.

Given the known correlation between game board outcomes in the negative treatment, we would expect a larger divergence between lagging and leading game boards in the negative treatment vis-à-vis the independent treatment. Figure 2.3.C shows that while the cumulative distribution of switching points of the leading game board (mean=5.98, n=31) lies to the right of the lagging game board (mean=4.78, n=28), the divergence is not statistically significant (z=-1.535, p=0.125).[29] In addition, as Table 2.1 (column 4) shows, the coefficients on own-hit, other-hit, sum of hits and difference of hits are not statistically different from zero. Neither Bayesian reasoning nor act separability can explain the pattern in switching points between lagging and leading game boards under the negative treatment.

Result 2.3. Participants systematically used game board outcomes to form expectations. However, expectations are not clearly mapped into the game board valuations.

At the beginning of each round, participants indicated on a 10-point scale (with end-points labelled "very low" and "very high") what they thought was the chance that the game board would reveal a hit in that round. To generate a relative measure of each participant's expectation of a hit, we calculated the difference between the self-reported expectation at the start of Round 11 (after the 10 sample rounds), and the expectation at the beginning of Round 1 (before any experience of a hit or a miss). The *expectation difference* captures any change in expectation of a hit, following the

---

[29] Note that Figures 2.2.C (B–Y<0) and 2.3.C (lagging vs. leading game boards) plot the same data. In 8 out of the 10 sessions in the negative treatment, hits on Y exceeded the number of hits on B. In the two other sessions, B and Y had the same number of hits.

experience in the sample rounds. The expectation difference was then re-coded as +1, 0, or -1 based on the sign of the computed values. [30]

Although the solicitation of self-reported expectations was not incentivised, we have strong evidence suggesting that participants systematically used the information on game board hits to form their expectations.

Table 2.2 summarises the ordered logistic regression results using *expectation difference* (ranging from -1 to +1) as dependent variable, and the same explanatory variables shown on Table 2.1. Two sets of regression results for each treatment are shown below. The regressors in the first equation include *own-hit*, *other-hit*, *game board colour*, and *mother tongue*. In the second equation, regressors are *own plus other*, *own minus other*, *game board colour*, and *mother tongue*.

Table 2.2: Ordered logistic regression results
dependent variable: expectation difference

	Positive		Independent		Negative	
own	0.3489	(0.3332)	0.7428**	(0.3662)	0.4309	(0.3789)
other	0.3976	(0.3090)	0.1017	(0.3254)	-0.6750**	(0.3118)
own + other	0.3732**	(0.1893)	0.4222	(0.2791)	-0.1221	(0.2474)
own – other	-0.0243	(0.2597)	0.3205	(0.2052)	0.5529**	(0.2433)
yellow game board	0.4823	(0.4838)	-0.0285	(0.4944)	-0.9030	(0.7554)
native English speaker	0.3463	(0.5110)	0.8022	(0.5407)	-0.5461	(0.4756)
Observations	75		75		74	
Wald chi2	4.36		6.19		6.90	
p-value	0.3593		0.1855		0.1415	
Pseudo R-squared	0.0381		0.0423		0.0482	

Robust standard errors are in parentheses; *** <0.01, **<0.05, *<0.10.

---

[30] Possible individual differences in interpreting a "very low" or "very high" expectation preclude the use of raw self-reported expectation as a cardinal measure. However, re-coded *expectation difference* between the first round and the last round serves as a relative expectation measure (i.e., higher, lower, or no change in expectation).

The primary interest in this section is to determine whether experienced game board hits in the sample rounds influence the change in expectation *(expectation difference)* in the final round. Given game board correlation, and the number of pertinent hits experienced in the sample rounds, evidence suggests that, across the three treatments, the direction in participants' formed expectations was aligned with Bayesian reasoning. In the positive treatment (column 2), although the ordered log-odds on *own-hit* and *other-hit* are positive, the coefficients are not statistically different from zero. However, *own + other hit* is statistically significant: keeping all other variables constant, total number of hits in the sample rounds increases the log-odds of a higher expectation of a hit on one's own game board. This result is consistent with Bayesian reasoning.

In the independent treatment (column 3), we find that participants' expectations are positively and significantly influenced by *own-hit*, but not by *other-hit*. Given uncorrelated game boards, this result supports both act separability and Bayesian reasoning.

Meanwhile, in the negative treatment (column 4), we find that *other-hit* significantly lowers the log-odds of an expected hit, while the positive coefficient on *own–other hit* indicates that more hits on own game board (versus the other game board) increases the log-odds of expecting a hit on the game board played in the last round. These results suggest that participants in the negative treatment formed expectations that are consistent with Bayesian reasoning.

Combining these findings with the OLS results shown on Table 2.1, evidence suggests that participants in the positive treatment systematically used the information on the total number of game board hits during expectation formation and in the valuation task, in a manner consistent with Bayesian reasoning.

In the independent treatment where the game boards are uncorrelated, only *own-hit* was informative. Our results show that participants used *own-hit* in expectation formation and in game board valuation.

In the negative treatment, the difference between *own-hit* and *other-hit* was useful in determining whether a game board was likely to have more hits than the other game

board. We find that *own – other hit* was used in forming expectations consistent with Bayesian reasoning. However, the sample round hits are unable to account for the variation in participants' switching points in the game board valuation task. While participants systematically used the available information on hits to form an expectation, the formed expectation does not directly map into the game board valuation. The results suggest that in the negative treatment, the task of forming a Bayesian expectation based on the revealed correlation between the game boards and the experienced hits in the sample rounds, and mapping these into a game board value, can be cognitively difficult. [31]

Result 2.4. Actual switching points are significantly higher than the Bayesian risk-neutral valuations.

How close are the actual switching points to the Bayesian risk-neutral valuations? We simulated plays of the experiment treatments and estimated the probability that a game board has 30 winning boxes, contingent on the number of hits on B and Y. The probabilities in each treatment were then used to calculate the Bayesian risk-neutral valuations for the blue and yellow game boards.

Figure 2.4 shows violin plots which combine boxplots and kernel density distributions of the valuation difference between actual switching point and the corresponding Bayesian risk-neutral valuation. As shown, valuation difference is more positively skewed in the positive treatment and more peaked than in the independent and negative treatments. Meanwhile, the median deviation from the Bayesian valuation has the highest dispersion in the negative treatment.

Table 2.3 reports the regression results with gap between actual and Bayesian risk-neutral valuation as dependent variable. Controlling for other variables, native English

---

[31] In contrast, Ericson and Fuster (2011) demonstrated that subjects with higher expectations also have higher WTAs. In their experiments, subjects' expectation of keeping an item (mug) was randomly manipulated. Subjects assigned with a higher likelihood of keeping the mug valued the mug 20-30% more than those with a lower expectation. However, subjects' expectations were exogenous, unlike in our experiment where subjects had to cognitively process game board correlation and past information on own-hit and other-hit when forming expectations.

speakers tend to have switching points closer to the Bayesian valuation. In the independent treatment, deviation from the Bayesian valuation is significantly higher on the yellow game board; this is consistent with the colour bias observed earlier.

Figure 2.4: Actual vs Bayesian risk-neutral valuations
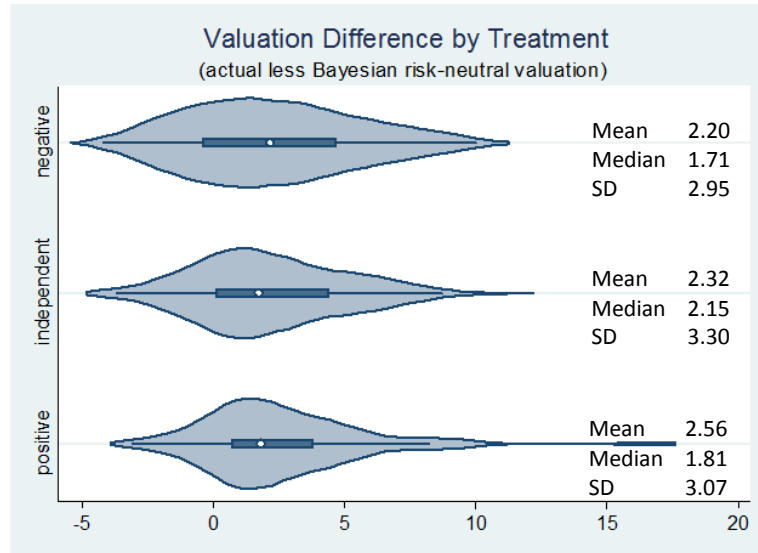
Table 2.3: Ordinary least squares estimation results
dependent variable: actual less Bayesian risk-neutral valuation

	Positive		Independent		Negative	
Own-hit	0.2334	(0.6251)	-0.8317*	(0.4906)	-0.9551*	(0.5681)
Other-hit	-0.9258*	(0.5147)	-0.3246	(0.4040)	0.8887*	(0.4640)
Native English speaker	-1.6728**	(0.6326)	-0.6442	(0.6843)	-1.4034**	(0.6837)
Yellow game board	0.3226	(0.8754)	1.4728**	(0.6658)	-0.5682	(1.4306)
Constant	3.9794***	(0.8052)	3.0081***	(0.9523)	3.4561***	(1.1804)
Observations	75		75		74	
F-statistic	2.34		1.98		11.30	
p-value	0.0635		0.1076		0.0000	
R-squared	0.1539		0.0916		0.2925	

Robust standard errors are in parentheses. *** <0.01, **<0.05, *<0.10

While the coefficient on *own-hit* is not statistically different from zero in the positive treatment, a hit on the other game board brings the switching point closer to the

Bayesian valuation. In the independent treatment, own-hit closes the valuation gap while other-hit does not have a significant influence. Meanwhile, in the negative treatment, the coefficients on own-hit and other-hit are both statistically significant: own-hit reduces the valuation gap while other-hit widens it. Albeit the coefficients are statistically significant only at the 90% level, these still provide indications that relevant game board experience was used by participants in the valuation task (in a different way).

2.4 Conclusion

We designed a two-part experiment to validate act separability. Given past experience of cases, the act separability assumption in CBDT implies that one's memory on *act A* is disjoint from the memory of cases on *act B*. CBDT therefore can explain people's tendency to neglect correlation in decision-making. The experiment allowed a fair chance for either case-based reasoning or Bayesian reasoning to emerge.

In the experiment, participants encountered two-armed bandits framed as coloured game boards. In Part 1, participants played sample rounds as opportunities to learn as much as they can about the game boards. In Part 2, players participated in an incentive-compatible elicitation task to reveal WTA on the assigned game board. The minimum offer price at which a participant changed preference from keeping the game board to exchanging the play of the game board for money was recorded as the participant's switching point. There were three treatments corresponding to different correlations between the blue and yellow game boards: perfect positive correlation (positive treatment), zero correlation (independent treatment), and negative correlation (negative treatment). If act separability holds, participants across the three treatments will only rely on past outcomes from their assigned game board. Meanwhile, Bayesian participants will use correlation information alongside pertinent game board outcomes.

In the positive treatment, we find that game board valuations by colour (blue or yellow) and by outcome (lagging or leading) are not significantly different, consistent with the Bayesian prediction. In the independent treatment, where the Bayesian and case-based predictions coincide, switching points are significantly different between lagging and

leading game boards, and between blue and yellow game boards. Our results suggest that this pattern in the valuations is jointly driven by the number of hits experienced on the assigned game board and a colour bias. However, in the negative treatment, the observed game board outcomes cannot account for the variation in switching points. Neither Bayesian nor case-based reasoning can explain the results.

The results with respect to participants' ability to comprehend correlations are encouraging. Across the treatments, we find that participants systematically used the observed game board outcomes and the correlation information when forming expectations. However, it appears that understanding does not always directly translate into the expected decision. This result has implications on the ease of effectively achieving a well-diversified portfolio among investors.

Forming a Bayesian expectation based on the revealed correlation information and the experienced outcomes in the sample rounds, and mapping out the formed expectation and the information on game board outcomes in a valuation task appears feasible among participants in the positive and independent treatments. However, evidence suggests that among participants in the negative treatment, this task may be cognitively challenging. Arguably, the number of game board hits experienced in the negative treatment could be more difficult to mentally process than the number of hits seen in the positive and independent treatments. In the positive treatment, rational participants are expected to disregard colour and recall the total number of hits in the sample rounds, while participants in the independent treatment only need to recall hits on the game board played in the last round. Meanwhile, in the negative treatment, Bayesian participants need to recall the number of hits on the blue and yellow game boards and then determine which game board is likely to have more winning numbers. However, we do not have actual evidence to support this conjecture. In a post-experiment questionnaire, directly asking subjects to rate the level of complexity of the relationship between game board outcomes, or asking them to answer simple valuation questions for each correlation condition could help address this limitation. Evidence of a significantly greater difficulty in the use of negatively correlated outcomes will support our proposed explanation. In addition, incentivising the elicitation of expectations during the experiment may help reduce the noise in expectations.

Chapter 3

Case-based stock selection

CBDT applied to stock selection predicts that investors rely on the similarity-weighted outcome of their past personal stock trading experience. We categorised retail clients' purchase decisions using selected salient information (stock name, industry sector, and broker recommendation) and compared transactions preceded by trading gains or losses. Our analysis reveals that retail investors systematically use similarity information and tend to purchase shares similar to stocks they previously traded either at a realised or unrealised gain. The results confirm a significant similarity effect that is not accounted for by an increase in wealth.

3.1 Introduction[32]

In stock selection, do investors rely on past personal stock trading successes and apply those successes in deciding to purchase similar stocks? If investors categorise stocks based on observable characteristics and use information on past personal trading experience to purchase similar stocks, such behaviour may be explained by case-based thinking.

Case-based decision theory (Gilboa and Schmeidler 1995) uses the concept of similarity and utility from past experience to explain behaviour. Each experience is stored in memory as a case with triple elements: problem, act, and result. CBDT predicts that given a new problem, a decision-maker will act based on the memory of actions and outcomes in past similar situations. Before deciding, an agent assesses the similarity of the current situation with past problems encountered, and then recalls the actions taken in those similar situations. The theory predicts that a decision-maker will choose a past action in similar situations with the highest similarity-weighted sum of outcomes (Gilboa, Liebermann and Schmeidler 2006). Under CBDT, a past personal trading experience is stored as a case in memory. An investor categorises stocks based on a similarity function and then purchases a stock similar to shares previously traded at a gain. A detailed description of the model is presented in Section 3.2.

We analysed the daily stock transactions of retail investors in the Philippines to determine if investors use their past personal stock trading experience and similarity information in the stock selection problem consistent with CBDT's prediction. Since we cannot know the similarity function of investors, we imposed a subset of salient and readily available stock information (stock name, industry sector, and broker recommendation) as bases for feature-based stock similarity assessment (Tversky 1977; Tversky & Hutchinson 1986; Tversky and Gati 2004; Goldstone and Son 2005). These stock similarity concepts are described in Section 3.3.

This paper aims to contribute to the empirical validation of CBDT in a stock market setting and to the literature on short-term stock trading using a unique dataset and its focus on stock similarity and retail investors' personal stock trading history. Given the uncertainties surrounding stock market outcomes, the stock selection problem among retail investors provides an appropriate setting to validate the predictive power of CBDT. According to Riesbeck and Schank (1989), the strategy used in decision-making depends on problem difficulty. Decision-makers faced with difficult problems, such as stock selection, rely on case-based reasoning while those working on less complicated problems (where rules on problem solving are clear) derive solutions through rules-based reasoning.

Our analysis focuses on the influence of past personal stock trading experience and selected stock similarity concepts. Meanwhile, past studies on momentum trading or trend-chasing, where the strategy entails buying past winners and selling past losers (Jegadeesh and Titman 1993), are based on the naïve use of public information.[33] Other studies have investigated the influence of personal investment experience. For instance, Barber and Odean (2000) showed that retail investors are influenced by past trading gains and losses experienced on the same stock. In this study, however, we considered

---

[33] Trend-chasing may be a rational trading strategy. Rational speculators who recognise the presence of positive feedback traders (those who buy assets when prices rise and sell when prices fall) opt to buy more stocks to further drive up the price (thus creating a self-feeding bubble) and then sell at an even higher price before positive feedback traders realise how far the market price has risen vis-à-vis fundamental value (De Long, Shleifer, Summers, and Waldman 1990).

different notions of stock similarity and analysed whether an investor's own past trading experience is applied to an objectively similar stock.

Reinforcement learning[34] (Sutton and Barto 1998) leads to repeating behaviour associated with pleasurable outcomes, and avoiding those causing pain or regret (Strahilevitz, Odean and Barber 2011). This type of learning is consistent with habit-formation under CBDT. However, case-based decision-making entails an additional process that involves the assessment of problem similarity. In the context of this study, both reinforcement learning and CBDT imply the use of past experience on the same stock. Given alternative similarity concepts, CBDT further suggests that investors rely on past personal experience on similar stocks.

Unlike expected utility theory, CBDT does not require decision-makers to know alternative courses of action or all possible outcomes associated with each action.[35] This means that a case-based investor is not expected to account for counterfactual events[36]

[34] Reinforcement learning has been used to explain investment decisions. For instance, Choi, Laibson, Madrian, and Metrick (2009) showed that American investors who had rewarding experiences from saving tend to increase their savings rate more than those who had less rewarding experiences. Kaustia and Knüpfer (2008) also showed that Finnish investors who subscribed to an initial public offering (IPO) and experienced subsequent high returns were likely to subscribe to future IPOs of other firms. Meanwhile, Hoffmann and Post (forthcoming) showed that while individual investors' past personal stock trading returns influenced their return expectation, market volatility experienced had no effect on their risk perception.

[35] Matsui (2000) demonstrated the mathematical equivalence between CBDT and Bayesian reasoning; behaviour consistent with CBDT may also be explained by Bayesian reasoning which poses difficulty in disentangling the two. For instance, if an investor has formed a belief that stocks in a certain industry sector (similar stocks) can replicate the gains of a previously traded stock belonging to that same sector, the tendency to purchase stocks similar to a past gainer may be explained by either CBDT or Bayesian reasoning. However, if the formed belief is not plausible given the lack of a clear pattern in past trading outcomes, a pattern in the results showing that investors tend to purchase stocks similar to past gainers may be better explained by CBDT.

[36] Counterfactual thinking involves a comparison of actual past events with imagined alternative states (Epstude and Roese 2008). In stock trading, an investor re-evaluates a recent sell decision by comparing the transacted price and the price movement following the transaction (Strahilevitz, Odean and Barber 2011). For example, an investor will avoid the purchase of a past

during decision-making (e.g., market outcomes on stocks not traded by that investor). We did not consider counterfactual events in the analysis, which distinguishes this study from Strahilevitz, Odean and Barber (2011).

We find that retail investors tend to use past personal experience on a similar stock in their purchase decisions. This result holds across the three similarity concepts considered, and trading outcome categories (realised vs. unrealised gain/loss). Our findings confirm a significant similarity effect that is not accounted for by an increase in wealth.

The rest of the paper is organised as follows: Section 3.2 presents the CBDT model; Section 3.3 discusses the similarity concepts used in this study; Section 3.4 describes the data; Section 3.5 presents the results; and Section 3.6 concludes.

3.2 CBDT model

CBDT (Gilboa and Schmeidler 1995; Gilboa, Lieberman and Schmeidler 2006) argues that an experience is encoded in memory as a case with triple elements: problem p, act a, and result r. When an agent faces a new problem q, she scans memory M for problems encountered in the past and evaluates the similarity with q, given similarity function s. The decision-maker recalls the act a taken and the corresponding outcome r at each similar problem stored in M. Given problem q, memory M, similarity function s, and utility function $u(r)$, available acts $a' \in A$ are ranked based on the similarity-weighted sum of utilities from each act: $U(a') = U_{q,M}(a') = \sum_{(p,a=a',r)\in M} s(p,q)u(r)$. In the stock selection problem, each past personal trading experience is stored as a case in memory. An investor categorises stocks based on a similarity function s. Conditional on personal stock trading memory, CBDT predicts that a typical retail investor is more likely to purchase a stock similar to a past gainer than any other stock.

---

realised gainer if the subsequent market price is higher than the actual selling price. Although the sale of the shares generated realised profits, the gain would have been higher had the investor waited a little longer before selling.

We constructed a dataset based on details of retail investors' stock transactions. For each stock x and period t, investor i has a net stock position $p(i, x, t)$, i.e., number of shares on stock x held at period t. For each (i, x, t), we define a change in i's stock position $\Delta(i, x, t) \in buy, sell, hold$. That is, trading experience $\Delta(i, x, t) = buy$ results in an increase in p, trading experience $\Delta(i, x, t) = sell$ leads to a lower p, while $\Delta(i, x, t) = hold$ results in no change in p.

Given $\Delta(i, x, t)$, we define trading outcome $r(i, x, t) \in \{positive, negative, zero, null\}$ which could be a realised or an unrealised gain/loss, zero or null. A gain or loss $r(i, x, t)$ is recorded in memory only if a previous act $a(i, x, t, s) = buy$ was made by an investor. A realised gain/loss is calculated from a buy-and-sell transaction pair, while an unrealised gain/loss is determined from a buy-and-hold transaction pair.

Realised trading outcome per share refers to difference between the selling price and weighted average cost. The difference is multiplied by the number of shares sold to determine the gain/loss for a sell transaction. Meanwhile, unrealised gain/loss per share is determined by comparing the weighted average cost and intraday prices adjusted for trading costs. A paper gain is incurred if both high and low intraday prices are above the average buying cost; it is a paper loss if both high and low intraday prices are below the average buying cost. If the average buying cost lies between the high and low intraday prices, no paper gain/loss is recorded for that stock. Paper gain/loss per stock is calculated by multiplying gain/loss per share and the number of outstanding shares. This procedure follows Odean (1998), Barber and Odean (2000), and Strahilevitz, Odean and Barber (2011). The distributions of gains and losses are described in Appendix 3.1.

For each (i, x, t), we define memory $M(i, x, t)$ separately for realised and unrealised trading outcomes. It is possible to have a long series of stock trading outcomes. However, we assume that an investor who is deciding to purchase a stock relies on the most recent personal trading gains or losses incurred on a similar stock within the past 30 calendar days; that is, M consists of the most recent entry in $r(i, x, t)$ provided that the outcome has been experienced in one of the periods $t - 30, \dots, t - 1$.

Given stock x and similarity concept $s \in$ *{same stock, same sector, same recommendation*[37]*},* there is a set of stocks $S(x, s)$ which is similar to x. Based on (i, x, t) and similarity concept s, we define $m(i, x, t, s)$ as i's memory of the outcome from trading a stock similar to x at period t. $m(i, x, t, s)$ is a function of $M(i, x', t)$ for all stocks x' that are similar to x (i.e., elements of $S(x, s)$). Given similarity concept s, an investor who faces problem (i, x, t, s) scans her memory $m(i, x, t, s)$ and recalls act $a(i, x, t, s) = \{buy, not \ buy\}$.

If multiple memories of trading gains or losses on similar stocks are incurred on a particular day, either a net gain or net loss is recorded depending on the difference between the number of gainers and losers on that day. Among similar stocks, if the number of gainers is equal to the number of losers, we assume that the memory of gains and losses cancel out and no history is recorded among those stocks on that day. Under *same stock* similarity, an investor recalls the most recent case that involves stock x. If the similarity concept is either *same sector* or *same recommendation*, an investor recalls the most recent case within S but not the same as x.

We analysed investors' decisions separately for each similarity concept, and for each trading outcome category (realised vs. unrealised), and investigated the relationship between $m(i, x, t, s)$ and act $a(i, x, t, s)$. Conditional on personal stock trading memory, CBDT predicts that a typical retail investor is more likely to purchase a stock similar to a past gainer than any other stock. Given similarity concept s, $r(i, x, t) > 0$ stored in $m(i, x, t, s)$, investor i is more likely to purchase a stock in $S(x, s)$ than a stock that is not an element of S. If $r(i, x, t) < 0$, a case-based investor will tend to buy a stock not in $S(x, s)$.

---

[37] Under the *same recommendation* similarity concept, x and y are deemed similar if stock x is currently considered for purchase (as recommended by the broker), while stock y was bought when it had a *recent* "buy" recommendation from the broker (i.e., the number of days elapsed between the date of the stock purchase and the "buy" recommendation report is not more than 14 calendar days).

3.3 Stock Similarity

CBDT gives similarity assessment a central function in decision-making that allows an agent to transfer knowledge acquired from past experience (Gregan-Paxton and Cote 2000; Zizzo 2003) to a present problem. Characteristics of objects may be plotted on a multidimensional plane and similarity may be assessed as the inverse of the measured distance between two objects. While humans tend to subconsciously extract parametric properties from objects, only similarity judgments are accessible. It has also been found that qualities are perceived as nominal features rather than continuous properties; this is consistent with feature-based models where similarity is measured as the degree to which two sets of salient features match each other (Tversky 1977; Tversky & Hutchinson 1986; Tversky and Gati 2004).

Investors have been observed to use only a subset of available information given their limited attention and processing capability (Hirshleifer, Lim and Teoh 2009). In this paper, we imposed feature-based similarity among stocks based on salient and readily available information, including: (i) stock name, (ii) industry sector, and (iii) broker recommendation.

Stock name. Each stock is represented by a unique alphanumeric symbol. Two stocks are similar if they have the same alphanumeric symbol. A-shares, B-shares, and preferred shares issued by the same company are considered similar.

Industry sector corresponds to the products and services that comprise the bulk of the firm's revenues as categorised by the Philippine Stock Exchange (PSE).[38] PSE categorises listed stocks across twenty-two sectors; the number of stocks by sector is summarised on Table 3.1. Stocks that belong to the same sector are considered similar. For example, given $s = same\ sector$, and a recent purchase of x, $S(x, s)$ consists of all stocks under the same industry category as x, but excluding x.

---

[38] The history and disclosures of listed companies can be accessed online at edge.pse.com.ph.

Table 3.1: Number of traded securities by industry sector

Sector	Number of Stocks	% of all buy transactions	Sector	Number of Stocks	% of all buy transactions
Banks	14	0.9	Media	6	5.4
Casinos and Gaming	6	1.8	Mining	22	5.8
Chemicals	7	2.7	Oil	8	6.3
Construction	10	1.8	Other Financials*	13	6.7
Education	4	1.2	Other Industrials**	2	7.2
Electrical Equipment	3	2.7	Other Services	6	7.6
Energy, Power & Water	14	3.1	Property	39	8.1
Food and Beverages	13	3.6	Retail	1	1.3
Holding Firms	43	4.0	Small & Medium Enterprises	1	1.0
Hotel and Leisure	3	4.5	Telecommunications	5	9.4
Information Technology	7	4.9	Transportation	7	9.9

*Other Financials include brokerage, insurance, leasing, fund management firms
**Other Industrials such as manufacturers of paper, cosmetics

Broker's recommendation is salient information for investors. Malmendier and Shanthikumar (2007) found that despite analysts' possible conflicts of interest, small investors are enthusiastic about recommended stocks especially those with a strong buy recommendation.

As part of the broker's service to its clients, research on a subset of stocks traded in the PSE is made available through the broker's website and sent to each client's nominated email address. From January 2007 to December 2010, the broker issued 748 recommendations on 37 stocks (16% out of the 234 transacted securities) which are either prepared quarterly or when the need to update a recommendation arises. 70% of the recommendation reports issued a "buy" rating, 11% had a "sell" rating, and 19% reported a "hold" rating. The 37 stocks covered by the broker's research account for 44% of investors' transactions.

Using broker recommendation as a similarity concept, each stock is categorised as either recommended (i.e., recently recommended for purchase) or not recommended (i.e., with a sell/hold rating or no recommendation). A recommendation is deemed recent if 0 to 14 calendar days[39] have elapsed between the report's release date and

---

[39] While only a handful of investors may seek investment advice (Bhattacharya, Hackethal, Kaesler, Loos and Meyer 2011), those who do are found to immediately act on a broker's recommendation (Michaely and Womack 2005). For clients who actively trade and track their

the transaction date. Here we assume that two weeks is sufficient time for investors to act on a broker's recommendation.[40] Under this similarity concept, $S(x, s)$ consists of stocks with a "buy" rating at the time shares were purchased, but excluding x.

3.4 Data

We analysed the daily transactions of 620 individual equity investors from January 2007 to December 2010.[41] The raw data was provided by one of the largest online stockbrokers in the Philippines. Only clients with an existing portfolio by 01 January 2007 were included in the dataset. Details on each transaction include the date, reference number[42], stock name, number of shares, buy/sell price, and fees paid. Daily intraday and closing stock prices were downloaded from *Thomson Reuters Datastream*.

Among the 620 retail investors, 77.6% are male and average age is 45 years. Median portfolio size is Php378,828 (£5,411). This amount falls between the 9th and 10th decile of the annual median family income in the Philippines.[43] Median transaction value is Php71,985 (£1,028) and the median number of buy transactions and sell transactions per investor is 122 and 106, respectively. This translates to 2.4 buy transactions and 2.1

---

equity portfolio's performance, two weeks allow sufficient time to act on the broker's recommendation.

[40] We tracked the timing of investors' purchase of broker-recommended stocks relative to the release date of the recommendation report. Purchases after the release of the recommendation (0–14 calendar days) reflect how quickly investors act on the buy-recommendation, while purchases before the report's release indicate investors' use of other sources of information. We find that while a significant number of broker-recommended purchases occurred in $t + 0, t + 1$ and $t + 2$, a substantial number of purchases were also recorded in $t - 1$. More details are found in Appendix 3.2.

[41] Appendix 3.3 shows a screenshot of the Philippine Stock Exchange Index from January 2007 to December 2010.

[42] The raw data provided by the broker does not include a time-stamp for each of the transactions. However, the order of the reference numbers reflects the chronology of each investor's stock transactions.

[43] Based on the 2006 and 2009 results of the Family Income and Expenditures Survey; https://psa.gov.ph/old/data/pressrelease/2011/ie09frtx.html.

sell transactions per investor per month. Annual median portfolio turnover is 8% or a portfolio holding period of 150 months. We have no information on clients' total wealth, or whether they maintain an equity portfolio with other brokerages.

History of trading gains and losses

The broker's online client interface is a matrix of stock positions that displays the stock's alphanumeric symbol, stock name, total number of outstanding shares, number of uncommitted shares (available for sale), average cost per share, current market price, market value, portfolio share in percent, and unrealised gain/loss (in pesos and in percent). Beside each stock in the portfolio is a "buy" or "sell" action choice. Purchase orders on other stocks may be made on another screen.

Also displayed on-screen is the portfolio's current total market value, the change in portfolio value from the previous day, and the portfolio's total paper gain/loss. Clients also have online access to the history of their recent sell transactions showing sales proceeds[44] and corresponding realised gain or loss; this complements an investor's memory of trading outcomes. Appendix 3.4 shows a screenshot of the client interface.

We reconstructed the daily portfolio of each investor to calculate the realised/unrealised gain or loss on each stock. Based on the portfolio on 31 January 2007, we worked backwards to determine the shares held by each client as of 31 December 2006. Using this portfolio as starting point, we calculated the daily weighted average cost per share of each security, adjusted for trading costs, until 31 December 2010. As described in Section 3.2, a realised gain/loss is calculated from a buy-and-sell transaction pair, while an unrealised gain/loss is determined from a buy-and-hold transaction pair. With 234 available stocks[45] and 976 trading days, there are 208,000 buying opportunities for each client i. Based on the history of trading gains and losses

---

[44] Net of fees and taxes (trading costs), such as: broker's commission (0.25% of gross transaction amount), value-added tax (12% of broker's commission), PSE transaction fee (0.005% of gross transaction amount), clearing fee (0.01% of gross transaction amount), and sales tax (0.5% of gross selling amount).

[45] A few stocks had an initial public offering after January 2007, while others were de-listed from the stock exchange before December 2010.

retained in memory as described in Section 3.2 and given a similarity concept s presented in Section 3.3, a buy/no buy decision at each opportunity is either: i) similar to a gainer, ii) similar to a loser, or iii) dissimilar.

Propensity to purchase similar/dissimilar stocks

Given memory of a gain for each trading day we counted the number of: (i) similar stocks purchased [SPG]; (ii) similar stocks not purchased [SNPG]; (iii) dissimilar stocks purchased [DPG]; and (iv) dissimilar stocks not purchased [DNPG]. On the other hand, given memory of a trading loss, we counted the number of: (v) similar stocks purchased [SPL]; (vi) similar stocks not purchased [SNPL]; (vii) dissimilar stocks purchased [DPL]; and (viii) dissimilar stocks not purchased [DNPL]. For instance, an investor experienced a recent gain from trading a mining stock. On the next trading day, shares of *another* mining company and a property firm were purchased. SPG under same-sector similarity refers to the purchase of another mining firm, while DPG refers to the purchase of property shares. The other available mining stocks (excluding the stock traded earlier) not purchased by the investor are recorded as SNPG, while non-mining stocks not purchased are counted under DNPG.

The frequency counts are used to calculate the following ratios, conditional on the memory of gains or losses on a similar/dissimilar stock.

(i)   Given memory of a trading gain, similar stocks purchased (PSG) relative to opportunities

$$\frac{SPG}{SPG + SNPG}$$

(ii) Given memory of a trading loss, similar stocks purchased (PSL) relative to opportunities

$$\frac{SPL}{SPL + SNPL}$$

(iii) Given memory of a trading gain, dissimilar stocks purchased (PDG) relative to opportunities

$$\frac{DPG}{DPG + DNPG}$$

(iv) Given memory of a trading loss, dissimilar stocks purchased (PDL) relative to opportunities

$$\frac{DPL}{DPL + DNPL}$$

A comparison between PSG and PSL may be confounded by wealth effect, i.e., a purchase decision may be influenced by the increase in portfolio size arising from a past trading gain. Instead, we compared two proportions: the ratio of PSG to PSL, and the ratio of PDG to PDL. Among similar stocks, $S(x, s)$, PSG/PSL indicates the probability that an investor will purchase a past gainer relative to the likelihood of purchasing a past loser. Meanwhile, among dissimilar stocks, i.e., not in $S(x, s)$, PDG/PDL reflects the chance that a dissimilar stock is purchased after a past trading gain (on a stock in $S(x, s)$) than if a past trading loss is incurred. Conditional on a past trading gain, a comparison between PSG/PSL and PDG/PDL reveals the influence of a personal trading outcome on a similar stock; that is, a statistically larger PSG/PSL relative PDG/PDL indicates a significant similarity effect consistent with CBDT.

Hypotheses

A Bayesian investor is unlikely to rely on past personal trading experience on a similar stock. Given the transparency of transactions in the stock exchange, an individual investor's past personal stock trading experience does not contain additional information vis-à-vis publicly known stock market results. In addition, if a series of stock price movements is assumed to have no memory of outcomes [under the random-walk hypothesis[46]] (Fama 1965; Fama, Fisher, Jensen and Roll 1969), past stock performance

---

[46] In financial markets, the arrival of news (i.e., market surprises) may be staggered which may lead to asset price dependence. For example, see Solnik (1973), Lo and MacKinlay (1988), Frennberg and Hansson (1993), Chang and Ting (2000). Fama (1965) argued that such price dependence is not sustainable in the presence of savvy investors who will readily take advantage of any market mispricing. Fama showed that while there is evidence of a serial correlation in stock price changes, on average, the observed price dependence is consistently close to zero. This means that active trading that is based on price dependence will not significantly outperform a naïve buy-and-hold strategy and potential profits from frequent trading may be easily eroded by commissions paid to brokers (Fama 1970). In addition, Fama (1998) showed

is not expected to recur[47] and the recent performance of a stock is unlikely to systematically spill-over to other similar stocks. Among Bayesian investors, PSG/PSL will not be significantly different from PDG/PDL.

Meanwhile, CBDT argues that agents rely on past personal experience on similar problems when making decisions. Case-based investors are not expected to account for market results beyond those they actually experienced. In stock selection, this suggests that investors find past personal trading experience informative when deciding to purchase a similar stock; that is, a case-based investor tends to put more weight on one's own experience vis-à-vis publicly known market outcomes. Given memory of trading outcomes on a similar stock, a case-based investor is likely to rely on one's personal history of trading gain/loss on similar stocks. Among case-based investors, PSG/PSL will be significantly higher than PDG/PDL.

3.5 Results

The conditional ratios described above are separately calculated for realised and unrealised trading outcomes and across similarity concepts: stock name, industry sector and broker recommendation. Given a pair of PSG/PSL and PDG/PDL ratios for each investor, we can determine whether PSG/PSL is significantly different from PDG/PDL for the same individual. The Wilcoxon matched-pairs signed-ranks test (Siegel and Castellan 1988) considers both the magnitude and direction of the difference between PSG/PSL and PDG/PDL. A significantly higher PSG/PSL vis-à-vis PDG/PDL suggests that past personal trading outcome on a similar stock influences purchase decisions.

---

that over a long-term horizon, overreaction to market news is as common as underreaction. If investors overreact to information, past performance is overweighted while the possibility for mean-reversion is underweighted (i.e., past winners tend to be future losers, and vice versa). Meanwhile, investors who underreact take a long time to respond to new information.

[47] This challenges "technical analysis" where emerging price patterns are used to predict the near-term direction of an asset's price. Doubt on the predictive power of technical analysis primarily arises from professional investment managers' inability to consistently beat benchmark indexes (Malkiel 2005). Hoffman and Shefrin (2014) also showed that investors who use technical analysis earn lower returns than other investors.

Result 1. Retail investors use past personal trading experience on a stock, whether realised or unrealised, when deciding to repurchase that same stock.

Table 3.2 summarises investors' propensity to repurchase the same stock conditional on personal trading history. On average, the chance that an investor will repurchase the same stock within 30 calendar days following a realised gain (5.44%) is higher than the probability of repurchasing that same stock after a realised loss (3.07%). Similarly, the propensity to repurchase the same stock within 30 calendar days after a recent paper gain (62.93%) is higher than the likelihood of repurchasing a past paper loser (48.44%). The ratios for past paper gainers (losers) are larger than the ratios for realised gainers (losers); this pattern suggests a typical investor's tendency to accumulate shares of the same stock before eventually shorting a position. Also, notice that PDG and PDL are lower than PSG and PSL given a larger universe of dissimilar stocks.

Table 3.2: Average propensity to repurchase same stocks

	Past Realised Gains or Losses	Past Paper Gains or Losses
PSG	0.0544	0.6293
PSL	0.0307	0.4844
Aggregate PSG/PSL	1.77	1.30
PDG	0.0079	0.0220
PDL	0.0054	0.0223
Aggregate PDG/PDL	1.46	0.99

While a higher PSG relative to PSL reflects investors' propensity to repurchase a stock similar to a past gainer, the similarity effect captured by this comparison is confounded by an increase in wealth (especially in the case of a realised gain.) Instead, we compare the ratios PSG/PSL and PDG/PDL. Table 3.2 shows that aggregate PSG/PSL is higher than aggregate PDG/PDL for either realised or unrealised trading outcomes. This result points to a similarity effect consistent with CBDT; that is, given a similarity concept and a trading gain on a stock, the likelihood of repurchasing that same stock is higher than the likelihood of purchasing a different stock.

PSG/PSL and PDG/PDL ratios for each individual investor are plotted in Figure 3.1.A, excluding those who did not engage in a same-stock repurchase within the 30 calendar day window. The median PSG/PSL ratio indicates that a past similar realised gainer is 1.08 times more likely to be repurchased than a past similar realised loser, while a

different stock is less likely to be purchased after a gain than after a loss on another stock (PDG/PDL=0.86).

Refer to Figure 3.1.A. The x-axis refers to the ratio of similar purchases after a gain to similar purchases after a loss (PSG/PSL) or the propensity to buy a similar stock conditional on trading gains or losses, while the y-axis is the ratio of dissimilar purchases after a gain to dissimilar purchases after a loss (PDG/PDL) or the propensity to buy a dissimilar stock given trading gains or losses. Each data point represents a pair of ratios for one investor. Conditional on a trading gain, investors who lie on the 45-degree line are deemed indifferent between similar and dissimilar stocks (n=36). Meanwhile, investors who are further away from the origin tend to purchase either the same or different stock after a gain (given an increase in wealth).

Notice the bounded area on the lower left of the chart (PSG/PSL<1 and PDG/PDL<1). Investors who appear in this region (n=66) are more likely to purchase a stock after a realised trading loss. Among these investors, 36 are equally likely to buy either the same or different stock. Those below the 45-degree line tend to repurchase the same stock (n=20), while those above that line are more likely to buy a different stock (n=10).

On the other hand, investors who are likely to make a purchase after a realised gain appear in the outer region where PSG/PSL>1 and PDG/PDL>1 (n=62). Those who are below the 45-degree line tend to repurchase the same stock (n=41) while investors who lie above that line are more likely to purchase a different stock (n=21).

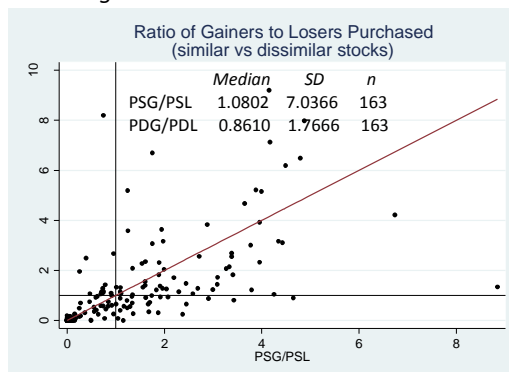*Figure 3.1.A: Same-stock repurchase given realised gains or losses*

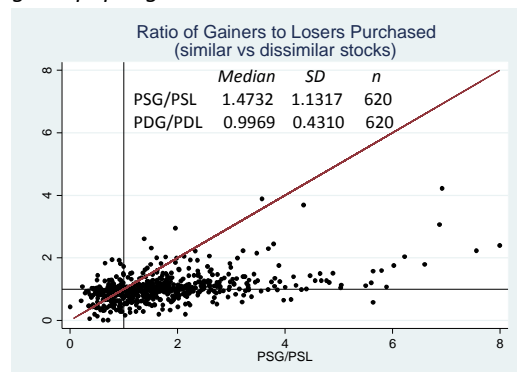*Figure 3.1.B: Same-stock repurchase given paper gains or losses*

Figure 3.1.A indicates that there are more investors who lie below the 45-degree line than above it. The Wilcoxon matched-pairs signed-ranks test confirms this observation; we find that individual level PSG/PSL and PDG/PDL are significantly different (z=3.577, p=0.0003, n=163). Keeping the investors who appear on the outer region of Figure 3.1.A yields a similar pattern in the signed-ranks test results (z=1.946, p=0.0517, n=62). Conditional on trading gains, the results suggest that a typical investor is more likely to purchase the same stock than a different stock. Apart from confirming a similarity effect, this also implies that investors tend to repurchase the same stock just a few days after taking profit on that stock.[48]

Figure 3.1.B captures the pattern in investors' purchases conditional on a recent unrealised trading outcome. The median PSG/PSL ratio indicates that a stock with a past paper gain is 1.47 times more likely to be repurchased than if a paper loss on that stock is incurred, while a different stock is equally likely to be purchased after a gain or after a loss on another stock (PDG/PDL = 0.99).

Among investors who tend to purchase a stock after a loss (i.e., those who appear within the bounded region in Figure 3.1.B, n=125), there are slightly fewer who lie below the 45-degree line (n=56) compared to those who lie above that line (n=69). Among those in the outer region (i.e., investors who tend to purchase after a gain, n=264), there are significantly more individuals who are likely to repurchase the same stock [appear below the 45-degree line; n=236] than those who tend to purchase a dissimilar stock [appear above the 45-degree line; n=28]. The signed-ranks test result shows that PSG/PSL is significantly higher than PDG/PDL (z=16.657, p=0.0000, n=620). We find a similar pattern in the results if we only include investors who appear on the outer region of Figure 3.1.B (z=13.068, p=0.0000, n=264).

---

[48] Retail investors seem to take advantage of the volatility in the Philippine Stock Exchange Index (past 20 trading days, annualised) by repurchasing stocks. On months with heightened volatility (>15), the average number of same-stock repurchases is 43, while periods with less volatility (<10) showed an average of 34 repurchases. In addition, investors who tend to repurchase a stock previously sold at a gain seem to have substantial gains: average 30-day unrealised return on a repurchased stock is 1.08%, although median return is zero.

Consistent with the similarity effect predicted by CBDT, these results indicate that a typical retail investor tends to purchase additional shares of the same stock, suggesting that clients do not systematically attempt to lower the average purchase cost of their stock holdings, i.e., buying additional shares at a lower price to reduce average cost per share.[49]

Result 2. Evidence suggests that retail investors apply their past personal stock trading experience to other same-sector stocks.

Under the same-sector similarity concept, a new stock is considered similar to a recently bought stock if that new purchase is under the same industry sector, but is not the same stock previously traded. On aggregate, the chance that an investor will purchase a same-sector stock within 30 calendar days following a realised gain on another stock (21.06%) is higher than if a realised loss is experienced (16.20%). Similarly, the average propensity to purchase a same-sector stock following a paper gain (16.14%) is higher than if a paper loss is incurred (13.96%). PSG higher than PSL, as shown on Table 3, may be explained by a similarity effect (i.e., tendency to purchase a similar stock after a gain), or may be due to higher purchasing power after a realised trading gain (wealth effect), or both.

Given same-sector similarity, we find that aggregate PSG/PSL is higher than aggregate PDG/PDL, for either realised or unrealised trading outcomes. This implies that investors tend to apply their past personal trading experience to another stock belonging to the same sector so that a typical investor is more likely to purchase a same-sector stock than shares of a stock categorised under another industry.

Table 3.3: Average propensity to purchase same-sector stocks

	Past Realised Gains or Losses	Past Paper Gains or Losses
PSG	0.2106	0.1614
PSL	0.1620	0.1396
Aggregate PSG/PSL	1.30	1.16
PDG	0.0440	0.0461
PDL	0.0405	0.0523
Aggregate PDG/PDL	1.09	0.88

---

[49] Portfolio managers (in the Philippines) commonly advise their clients to "buy on dips" to take advantage of market corrections. Evidence indicating that investors increase their stock position following a paper gain suggests that clients do otherwise.

Huang (2013) found similar results. Following trading gains on a stock in excess of a market benchmark, unsophisticated investors are likely to purchase another stock belonging to the same industry group. Huang's results suggest that investors engage in categorical learning using industry as basis for grouping. This is consistent with the view that given limited attention and cognitive capacity (Peng and Xiong 2006) investors use their past personal experience when deciding to purchase another stock belonging to the same industry.

Figure 3.2.A plots individual-level PSG/PSL (same-sector stocks) against PDG/PDL (different-sector stocks) conditional on realised trading outcomes. The median PSG/PSL ratio indicates that a same-sector stock is 1.30 times more likely to be purchased after a realised trading gain than after a loss, while a different-sector stock (PDG/PDL) is 1.16 times more likely to be purchased after a realised gain (on another stock) than if a realised loss was experienced.

Among investors in the outer region in Figure 3.2.A (n=297), we find that there are more individuals who lie below the 45-degree line (n=176) than those that appear above that line (n=121). Meanwhile, among those who appear within the bounded area, there are about the same number of investors who are below the 45-degree line (n=48) and those who are above that line (n=45). There are only five investors with equal conditional ratios. The signed-ranks test confirms a statistically significant difference between PSG/PSL and PDG/PDL (z=3.226, p=0.0013, n=572). Including only the investors who appear on the outer region of Figure 3.2.A, we find a similar pattern in the signed-ranks test result (z=2.244, p=0.0249, n=297). Given a realised gain on a stock, a typical retail investor tends to purchase shares of another stock that belongs to the same industry sector; this is evidence of a similarity effect consistent with CBDT.

Figure 3.2.B shows the scatterplot of PSG/PSL and PDG/PDL given unrealised trading outcomes on same-sector stocks. The median PSG/PSL (1.14) is higher than PDG/PDL (0.81), and the difference between the ratios is statistically significant based on the signed-ranks test (z=15.522, p=0.0000, n=601). Among investors who appear inside the bounded region, we find more individuals below the line (n=133) than those above it (n=71). There are 30 investors who are indifferent between similar and dissimilar stocks.

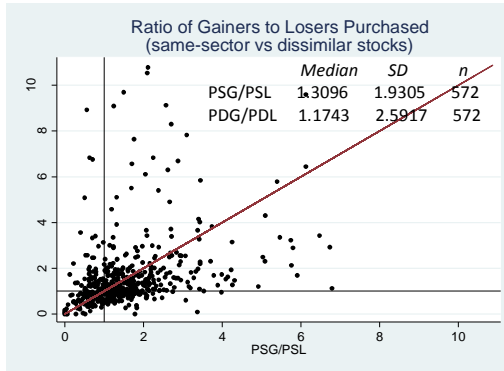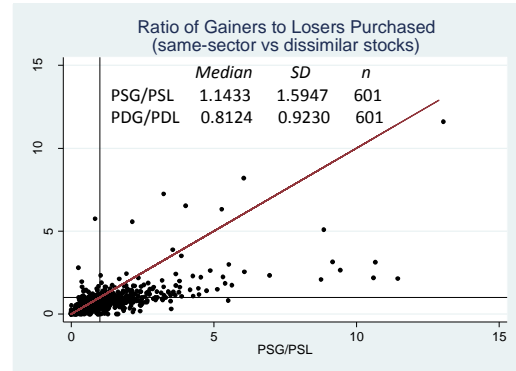*Figure 3.2.A: Same-sector purchase given realised gains or losses*

Figure 3.2.B: Same-sector purchase given paper gains or losses

The stronger tendency to purchase same-sector stocks among investors in the outside region in Figure 3.2.B (z=8.501, p=0.0000, n=148) is also reflected in the number of data points below the 45-degree line (n=127) relative to those above that line (n=21). Retail investors who use past personal experience in decisions to purchase same-sector stocks may imply that they hold a poorly diversified equity portfolio. Similar to what has been documented among US retail investors (Barber and Odean 2000), we find that the median number of stocks held in individual portfolios at the end of each month is four, vis-à-vis the Philippine Stock Exchange Index consisting of 30 (large-capitalization) stocks, the 22 sector categories enumerated on Table 3.1, and the 234 stocks traded by the same investors at some point during the sample period.

Figure 3.3: Median number of transactions disaggregated by sector similarity
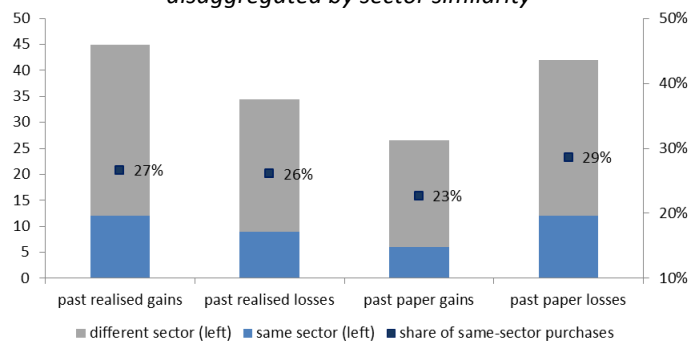
To emphasise investors' tendency to continue investing in familiar stocks (or shares they personally traded in the past), Figure 3.3 summarises the median number of same-sector and different-sector purchases following either trading gains or losses. While there are more different-sector purchases in absolute terms, about a quarter of all

purchases are same-sector stocks. In fact, this tendency to purchase same-sector stocks holds even in the case of past trading losses, and as indicated by the PSG and PSL ratios on Table 3.3. This observation complements past studies showing that investors have a tendency to invest in local and familiar stocks, i.e., buying shares in firms they work for, or firms that are geographically close to them (Barber and Odean 2013). Stock familiarity, and in this case – the past personal experience on a same-sector stock, influences purchase decisions.

Result 3. Past personal experience of a realised or unrealised gain on a broker-recommended stock spills over to a similar stock.

Consider *stock A* that is recommended for purchase by the broker and actually purchased by an investor. Does the experience on *stock A* spill over to another stock that has a buy recommendation from the broker? On aggregate, the chance that an investor will purchase a broker-recommended stock following a realised gain on another broker-recommended stock is higher (36.68%) than if a realised loss is experienced (28.60%). Meanwhile, given paper gains on a broker-recommended stock, the likelihood of a similar purchase (30.44%) is also higher than if a past paper loss is incurred (27.54%). Parallel to the similarity effect shown on Tables 3.2 and 3.3, under same-recommendation similarity, the calculated ratios on Table 3.4 suggest that past personal trading experience on a broker-recommended stock spills over to another stock that is similarly recommended for purchase.

Table 3.4: Average propensity to purchase broker-recommended stocks

	Past Realised Gains or Losses	Past Paper Gains or Losses
PSG	0.3668	0.3044
PSL	0.2860	0.2754
Aggregate PSG/PSL	1.28	1.11
PDG	0.0196	0.0205
PDL	0.0188	0.0228
Aggregate PDG/PDL	1.04	0.90

Figure 3.4.A plots PSG/PSL (broker-recommended) against PDG/PDL (not recommended) for each individual investor following realised gains or losses. The median PSG/PSL ratio (1.46) is greater than the median PDG/PDL (1.09); that is, a broker-recommended similar to a past realised gainer is 1.46 times more likely to be purchased than a broker-recommended loser. Meanwhile, a stock not recommended

for purchase is 1.09 times more likely to be bought after a realised gain (on a recommended stock) than if a loss is incurred. The scatterplot indicates that among investors in the outer region (n=309), there are more who lie below the 45 degree line (n=212) relative to those who appear above that line (n=97). Among those within the bounded region, we find slightly more individuals below the line (n=43) than those above it (n=32). We find that the difference between PSG/PSL and PDG/PDL is statistically significant for the entire sample (z=9.166, p=0.0000, n=578), and if we include only the investors who appear in the outer region (z=8.618, p=0.0000, n=309).

Figure 3.4.A: Broker-recommended purchase given realised gains or losses
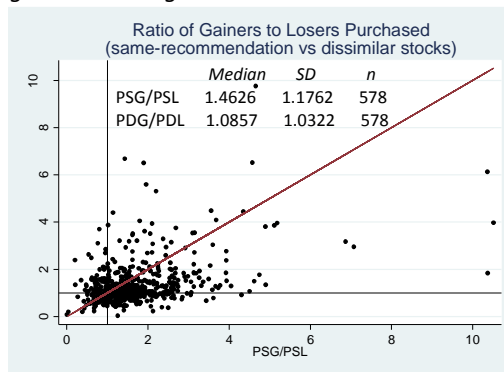
Figure 3.4.B: Broker-recommended purchases given paper gains or losses
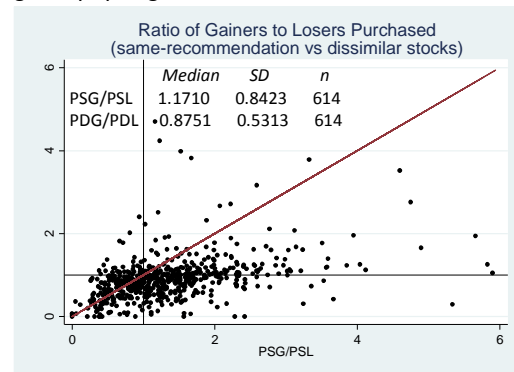
Figure 3.4.B depicts conditional ratios given a recent paper gain or loss. The median PSG/PSL ratio (1.17) among broker-recommended stocks is also higher than PDG/PDL (0.88) for non-recommended stocks. Among investors who appear on the outer region (n=174), notice that there are more who appear below the 45-degree line (n=146) than above it (n=28); this indicates a higher likelihood to purchase similar stocks than dissimilar stocks. Meanwhile, among investors who tend to purchase following a paper loss, we find that more individuals tend to buy a broker-recommended stock (n=112) relative to those who buy non-recommended shares (n=83). In our sample, only five investors are indifferent between broker-recommended and non-recommended stocks.

The difference between PSG/PSL and PDG/PDL is statistically significant for either realised (z=9.166, p=0.0000, n=578) or unrealised (z=14.428, p=0.0000, n=614) trading outcomes. If we exclude investors who have a stronger likelihood of purchasing a stock given a trading loss (i.e., those who are not in the outer region), the higher tendency to buy a similar stock persists for either realised (z=6.670, p=0.0000, n=309) or unrealised

(z=8.618, p=0.0000, n=174) trading outcomes. These results suggest that personal trading experience on a broker-recommended stock, whether realised or unrealised, spills over to another broker-recommended stock.

Result 4. The regression results confirm that investors apply past personal trading experience to other similar stocks.

For each investor, stock purchases were chronologically arranged by transaction date and by transaction reference number. We separately estimated panel probit regression models (Wooldridge 2002) across the similarity concepts and trading outcome categories with buy decision as dependent variable (1=similar stock purchased, 0=dissimilar stock purchased). Below is a description of the explanatory variables:

- *past trading outcome* = recent realised/unrealised gain or loss prior to the stock purchase.
- *portfolio change* = difference between the beginning portfolio value on the day of the stock purchase and the beginning portfolio value on the previous trading day. The difference serves as proxy for change in wealth.
- *portfolio size* = beginning portfolio value on the day of stock purchase.
- *recommended* = whether a stock has had a recent "buy" recommendation from the broker (1) or otherwise (0).
- *bull market* = whether the stock transaction occurred on a rising market (1) or a falling market (0). Using 20 percent change in the PSE index's previous peak or trough as basis for a change in market condition, the identified bull market in the dataset is from 1 January until 9 October 2007; the bear market is from 10 October 2007 until 10 October 2008, and the next bull begins on 11 October 2008.
- *index change* = day-on-day percent change in the PSE Index.
- *frequency* = total number of transactions (buy or sell) from January 2007 to December 2010.
- *age* = in years as of September 2012.
- *male* = dummy for sex (1=male, 0=female).

Table 3.5: Panel probit regressions: purchases given past realised gains or losses

	(1) same-stock	(2) same-sector	(3) same-recommendation
past gains or losses	3.61E-06	4.80E-07	1.89E-06
	(8.01)**	(11.23)**	(12.00)**
recommended	-1.36E-02	1.48E-01	
	-0.23	(9.92)**	
portfolio change	-1.19E-09	-3.60E-10	-4.70E-10
	-1.01	-0.69	-0.27
portfolio size	-2.58E-09	5.30E-10	1.32E-09
	-1.58	-1.19	-0.93
frequency	1.22E-04	-7.58E-05	-1.41E-04
	(3.94)**	(4.39)**	(3.65)**
bull market	-0.10153	0.13388	-0.07386
	-1.54	(6.61)**	-1.86
index change	-0.25397	0.57739	0.85651
	-0.44	(3.46)**	(2.26)*
male	-0.02398	-0.04220	-0.18155
	-0.31	-1.47	(2.91)**
age	4.30E-03	6.42E-04	4.90E-03
	-1.51	-0.64	(2.28)*
N	4,481	57,240	10,999

* $p<0.05$; ** $p<0.01$

Table 3.5 reports the marginal effect of the regressors on the likelihood of a similar stock purchase given realised trading outcomes, i.e., the percentage point change in the predicted probability of the dependent variable given an (instantaneous) change in the covariate while keeping all other variables equal to their means (and the dummy variables equal to 1).

The marginal effect of a past realised trading gain/loss is positive and statistically significant under the same-stock regression; that is, an infinitesimal gain on a stock increases the probability of a same stock repurchase relative to a dissimilar stock. This supports our earlier result showing that retail investors use past personal experience on the same stock in their repurchase decisions. We find a similar pattern in the marginal effect of past realised trading gain/loss under the same-sector regression and the same-recommendation regression. Consistent with CBDT's prediction, these findings alongside our nonparametric results suggest that past realised trading outcomes spill over to similar stocks.

Table 3.6 summarises the marginal effect of various explanatory variables on stock purchases preceded by unrealised trading gains or losses. Similar to the results presented on Table 3.5, the marginal effect of a past paper gain is positive and

statistically significant across the regression equations. The marginal effects imply that an infinitesimal paper gain on a stock encourages the purchase of a similar stock more than the purchase of a dissimilar stock.

Table 3.6: Panel probit regressions: purchases given past paper gains or losses

	(4) same-stock	(5) same-sector	(6) same-recommendation
past gains or losses	2.48E-04	8.35E-04	8.80E-04
	(3.71)**	(12.38)**	(6.62)**
recommended	0.22506	0.15509	
	(15.14)**	(9.37)**	
portfolio change	-3.30E-10	-7.00E-11	1.90E-09
	-0.65	-0.13	-0.76
portfolio size	-2.00E-10	7.70E-10	-4.21E-09
	-0.43	-1.61	(2.36)*
frequency	1.93E-04	-2.15E-05	-7.80E-05
	(9.61)**	-1.18	-1.72
bull market	0.05474	0.06494	0.00710
	(3.03)**	(2.99)**	-0.17
index change	-1.41066	0.89139	2.52734
	(8.84)**	(4.60)**	(5.63)**
male	0.07314	0.02368	-0.15280
	(2.38)*	-0.72	(1.99)*
age	0.00378	0.00255	0.00774
	(3.54)**	(2.30)*	(2.97)**
N	60,109	47,261	9,668

* $p<0.05$; **$p<0.01$

Let us consider the other explanatory variables. The marginal effect of a broker's buy recommendation is generally positive. While statistically insignificant in regression (1), we find that a "buy" recommendation has a positive influence on similar-stock purchases in the other models. For instance, on Table 3.6, a "buy" recommendation increases the probability of a same-stock repurchase and same-sector stock purchase by 22.5 and 15.5 percentage points, respectively, all other variables kept equal to their means and the dummy variables equal to 1. These results are aligned with past studies which showed that investors are enthusiastic about recommended stocks (Malmendier and Shanthikumar 2007) and act on changes in stock recommendation (Michaely and Womack 2005).

Change in portfolio value is used in the probit regressions as proxy for change in wealth (or change in available funds for investment). A positive and statistically significant marginal effect indicates the presence of wealth effect. However, the marginal effects

in the regressions reveal that a change in portfolio value is generally not statistically different from zero.

Are retail investors' purchases influenced by the prevailing market trend or recent market rallies or corrections (trend-chasing)? Same-sector regressions (2) and (5) show striking results supportive of trend-chasing. We find a statistically higher likelihood of purchasing a same-sector stock given a previous day PSE increase and an indication of bull market trend.

For regressions (3) and (6), we find that a gain in the PSE index increases the probability of a similar (same-recommendation) stock purchase, but the prevailing market trend does not seem to matter. While the marginal effect of the proxy indicators on market performance (bull market and index change) are statistically significant in regression (4), the direction is inconsistent vis-à-vis a trend-chasing story. In addition, the marginal effects of these variables are not significantly different from 0 in regression (1).

Investors who frequently trade tend to be naturally more experienced than investors who transact less frequently. On Table 3.5, frequency is statistically significant across the three equations (although the marginal effects have varying signs), while Table 3.6 shows that the marginal effect of frequency matters only under same-stock similarity. In regressions (1) and (4), under same-stock similarity, the marginal effect is positive and statistically significant. This implies that an additional transaction is more likely to be a trade on the same stock than a different stock. However, the marginal effect of frequency is in the opposite direction in regressions (2) and (3). A higher frequency of trades tends to increase the likelihood of purchasing a dissimilar stock more than a same-sector/same-recommendation stock. These results suggest that while experienced investors (i.e., individuals who trade more frequently) are more likely to repurchase the same stock than a dissimilar stock, this tendency does not extend to same-sector and same-recommendation purchases.

Another interesting result common on Tables 3.5 and 3.6 is the marginal effect of being male in regressions (3) and (6): male investors tend to veer away from a broker-recommended stock, but tend to repurchase the same stock as indicated under regression (4). These results support the findings of Barber and Odean (2001) where

they showed that males (single males in particular) are strongly biased towards their own assessment of an asset and are less influenced by the valuations of other people so that overconfidence leads them to trade excessively. This overconfidence arises from self-serving attribution bias or taking too much credit for one's successes (Hoffmann and Post 2014) which men are found to be more susceptible (Gervais and Odean 1998) especially in highly uncertain tasks with unclear feedback (Barber and Odean 2001) like stock selection.

3.6 Conclusion

Using daily stock transactions of retail investors in the Philippines, we analysed whether past personal trading experience on a stock is applied to an objectively similar stock (stock name, industry sector, or broker-recommendation). Unlike past studies on individual stock trading, this paper determines the influence of realised or unrealised trading gains or losses personally experienced, excluding counterfactual events or market outcomes after a decision to trade has been executed.

CBDT proposes that a decision-maker who encounters a new problem recalls past acts taken in similar cases actually experienced, and then chooses the act with the highest similarity-weighted sum of outcomes. In the stock selection problem, CBDT suggests that given a similarity concept and a recent history of trading outcomes, a retail investor is likely to use past personal trading experience when deciding to purchase shares of a similar stock.

We attempted to empirically validate the similarity effect predicted by CBDT. Our analysis reveals that retail investors systematically use the three stock similarity concepts considered in this paper (stock name, industry sector and broker's recommendation). Across these similarity concepts, we find that investors apply past personal trading experience to other similar stocks. The results confirm a significant similarity effect that is not accounted for by an increase in wealth.

As an extension to this study, we can explore whether counterfactual market events are applied to objectively similar stocks. For example, if an investor had realised gains from

trading Bank A shares, but the subsequent stock price had gone up *t period* after the shares have been sold, does the personal experience of the counterfactual event affect the likelihood of purchase of other bank shares? Also, if multiple memories of trading gains or losses on similar stocks are incurred on a particular day, either a net gain or net loss is recorded depending on the difference between the number of gainers and losers on that day. As a possible refinement to the treatment of multiple memories on a single day, the net gains or losses (in pesos) could be used to determine if a positive or negative experience is stored in memory.

Conclusion

Case-based decision theory predicts that based on a similarity concept and past personal experience, a decision-maker will choose a past action in similar situations with the highest similarity-weighted sum of outcomes. In two laboratory experiments, and one empirical paper using real market data, we attempted to validate whether decision-makers: i) encode and retrieve past experiences using similarity information; ii) choose an act with the highest similarity-weighted outcome; and iii) maintain separate memories of alternative actions taken in the past (act separability). Unlike past related research and the limited number of experimental and empirical work on CBDT, the studies presented in this thesis pitted the theory against Bayesian reasoning. Our results seem to suggest that the success of CBDT in predicting behaviour is influenced by the complexity of the decision situation, and the pure randomness of outcomes personally experienced in the past.

In a paper and pencil ticket experiment, we induced features-based similarity (using colour) and objective uncertainty (using live draws from a bingo cage). Participants played sample rounds to learn about the likelihood of a successful draw (£20) on a coloured ticket. Since ticket outcomes were randomly generated by a single randomiser, it was an easy strategy for a Bayesian player to ignore ticket colour and rely on the total number of successful draws in the valuation task. Our results show that participants were neither case-based nor Bayesian. While evidence suggests that participants used a similarity cue (colour), valuations on the randomly assigned coloured ticket (blue or yellow) in the last round were the opposite of CBDT's similarity-weighted prediction: a ticket with fewer successes (lagging ticket) was valued more highly than a ticket with more successes (leading ticket) – a result consistent with the gambler's fallacy. Relative to other experiments where CBDT's similarity-weighted prediction emerged, our results suggest that the manner of learning cases (Gonzalez, Lerch and Lebiere 2003; Hertwig, Barron, Weber and Erev 2004), and the randomness of outcomes may matter to subsequent decisions.

In a two-part computer experiment, we investigated whether participants' decisions are consistent with CBDT's act separability axiom. Given past experience of cases, act separability suggests that memory on *act A* is disjoint from the memory of cases on *act*

B. During the experiment, participants encountered two-armed bandits framed as coloured game boards (blue or yellow) with known correlation (positive, negative, or independent). Similar to the ticket experiment, participants played sample rounds as opportunities to learn as much as they can about the game boards before the incentive-compatible valuation task in the final round. If act separability holds, participants across the three treatments will only rely on past outcomes from their assigned game board. Meanwhile, Bayesian participants will use correlation information alongside pertinent game board outcomes.

The positive treatment replicated the ticket experiment, but controlled for *a priori* knowledge on the game board outcomes. Participants in the game board experiment knew that winning numbers in each game board are pre-determined, and a game board could either have a 10% or 30% chance of a successful draw. This alteration in the pure randomness of outcomes effectively muted the gambler's fallacy. Consistent with the Bayesian prediction, we find that game board valuations by colour (blue or yellow) and by outcome (lagging or leading) are not significantly different.

In the independent treatment, where the Bayesian and case-based predictions coincide, valuations are significantly different by colour and by outcome. Meanwhile, in the negative treatment, we find no significant difference in game board valuations despite participants' knowledge that if one game board has 30 winning numbers, then the other game board automatically has only 10 winning numbers. Neither case-based nor Bayesian reasoning can explain this result.

Across the three treatments, evidence suggests that participants systematically used the observed game board outcomes and the correlation information when forming expectations. While participants in the positive and independent treatments succeeded in mapping their formed expectations into game board decisions, this task appears cognitively difficult among participants in the negative treatment.

CBDT appears to be more successful in predicting behaviour in more complex decision situations (Riesbeck and Schank 1989) like stock selection. We analysed the daily stock transactions of retail investors in the Philippines to determine if investors use their past personal stock trading experience and similarity information (stock name, industry

sector, and broker recommendation) in selecting stocks to purchase, in a manner consistent with CBDT's prediction. Indeed, we find that retail investors tend to use past personal experience on a similar stock in their purchase decisions. This result holds across the three similarity concepts considered, and trading outcome categories (realised vs. unrealised gain/loss). Our findings confirm a significant similarity effect that is not accounted for by an increase in wealth.

The three studies presented in this thesis contribute to the limited CBDT literature, and provide tests on Bayesian reasoning. Given the varying success of CBDT across the studies, further work may help verify if the emergence of case-based decisions is indeed influenced by the nature and complexity of the decision setting.

References

Agnew, Julie R. and Lisa R. Szykman (2005). Asset Allocation and Information Overload: The Influence of Information Display, Asset Choice, and Investor Experience. *The Journal of Behavioral Finance*, 6(2): 57-70.

Aha, David W., Dennis Kibler, and Marc K. Albert (1991). Instance-Based Learning Algorithms. *Machine Learning,* 6: 37-66.

Ayton, Peter and Ilan Fischer (2004). The Hot Hand Fallacy and the Gambler's Fallacy: Two Faces of Subjective Randomness. *Memory & Cognition*, 32(8): 1369-1378.

Barber, Brad M. and Terrance Odean (2000). Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. *Journal of Finance*, LV: 773-806.

_____ (2001). Boys will be boys: gender, overconfidence and common stock investment. *The Quarterly Journal of Economics*, 116(1): 261-292.

_____ (2013). The Behavior of Individual Investors. In George Constatinides et al (eds), *Handbook of Economics of Finance, Volume 2.* Elsevier Publishing.

Bardsley, Nicholas, Robin Cubitt, Graham Loomes, Peter Moffatt, Chris Starmer and Robert Sugden (2010). *Experimental Economics: Rethinking the Rules*. Princeton University Press.

Barron, Gregory M. and Stephen Leider (2010). The Role of Experience in the Gambler's Fallacy*. Journal of Behavioral Decision Making*, 23(1): 117-129.

Becker, Gordon, Morris DeGroot and Jacob Marschak (1964). Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, 9(3): 226-232.

Bhattacharya , Utpal, Andreas Hackethal, Simon Kaesler, Benjamin Loos, and Steffen Meyer (2012). Is Unbiased Financial Advice to Retail Investors Sufficient? *Review of Financial Studies*, 25(4): 975-1032.

Bleichrodt, Han, Martin Filko, Amit Kothiyal and Peter P. Wakker (2015). Making Case-Based Decision Theory Directly Observable. Available at: http://people.few.eur.nl/wakker/pdf/cbdt.pdf.

Burns, Bruce D. and Bryan Corpus (2004). Randomness and Inductions from Streaks: "Gambler's Fallacy" Versus "Hot Hand". *Psychonomic Bulletin & Review*, 11(1): 179-184.

Buschena, David (2003). Similarity Models as an Approach to Behavior Violating Expected Utility. *Greek Economic Review*, 22(2): 79-98.

Chang, Kuo-Ping and Kuo-Shiuan Ting (2000). A variance ratio test of the random walk hypothesis for Taiwan's stock market. *Applied Financial Economics*, 10(5): 525-532.

Choi, James J., David Laibson, Brigitte C. Madrian and Andrew Metrick (2009). Reinforcement Learning and Savings Behavior. *The Journal of Finance*, LXIV(6): 2515-2534.

Croson, Rachel and James Sundali (2005). The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos. *The Journal of Risk and Uncertainty*, 30(3): 195-209.

Croson, Rachel and Uri Gneezy (2009). Gender Differences in Preferences. *Journal of Economic Literature*, 47(2): 448-474.

Cyert, Richard M. and Morris H. DeGroot (1974). Rational Expectations and Bayesian Analysis. *Journal of Political Economy*, 82(3): 521-536.

De Bondt, Werner F.M. (1998). A Portrait of the Individual Investor. *European Economic Review*, 42: 831-844.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldmann (1990). Positive Feedback Investment Strategies and Destabilizing Rational Speculation. *The Journal of Finance*, XLV(2): 379-395.

Dohmen, Thomas, Armin Falk, David Huffman, Felix Marklein, and Uwe Sunde (2009). Biased Probability Judgment: Evidence of Incidence and Relationship to Economic Outcomes from a Representative Sample. *Journal of Economic Behavior & Organization*, 72: 903-915.

Epstude, Kai and Neal J. Roese (2008). The Functional Theory of Counterfactual Thinking. *Personality and Social Psychology Review*, 12(2): 168-192.

Ericson, Keith M. Marzilli and Andreas Fuster (2011). Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments. *The Quarterly Journal of Economics*, 126: 1879-1907.

Estes, William K. (1964). Probability Learning. In Arthur W. Melton (ed), *Categories of Human Learning*. New York: Academic Press.

Eyster, Erik and Georg Weizsaecker (2011). Correlation Neglect in Financial Decision-making. DIW Berlin Discussion Paper No. 1104. Available at http://dx.doi.org/10.2139/ssrn.1735339.

Fama, Eugene F. (1965). The Behavior of Stock-Market Prices. *Journal of Business*, 38(1): 34-105.

_____ (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2): 383-417.

_____ (1998). Market efficiency, long-term returns, and behavioural finance. *Journal of Financial Economics*, 49: 283-306.

Fama, Eugene, Lawrence Fisher, Michael Jensen and Richard Roll (1969). The Adjustment of Stock Prices to New Information. *International Economic Review*, 10 (1): 1-21.

Fischbacher, Urs (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10(2): 171-178.

Frennberg, Per and Bjorn Hansson (1993). Testing the random walk hypothesis on Swedish stock prices: 1919-1990. *Journal of Banking & Finance*, 17(1): 175-191.

Gayer, Gabrielle, Itzhak Gilboa and Offer Lieberman (2007). Rule-Based and Case-Based Reasoning in Housing Prices. *The B.E. Journal of Theoretical Economics,* 7(1): 1-35.

Gervais, Simon and Terrance Odean (2001). Learning to be overconfident. *Review of Financial Studies*, 14(1): 1-27.

Gilboa, Itzhak and David Schmeidler (1995). Case-Based Decision Theory. *The Quarterly Journal of Economics*, 110(3): 605-639.

_____ (2001). *A Theory of Case-Based Decisions*. Cambridge University Press.

_____ (2010). Case-Based Predictions: Introduction. In Eric Maskin (ed), *Introduction to Case-Based Prediction*. World Scientific Publishers.

Gilboa, Itzhak, Offer Lieberman, and David Schmeidler (2006). Empirical Similarity. *The Review of Economics and Statistics,* 88(3): 433-333.

Gilovich, Thomas, Robert Vallone and Amos Tversky (1985). The Hot Hand in Basketball: On the Misperception of Random Sequences. *Cognitive Psychology*, 17: 295-314.

Gittins, John C. (1979). Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society Series B*, 41(2): 148-177.

Goldstone, Robert and Ji Yun Son (2005). Similarity. In K. Holyoak and R. Morrison (eds), *Cambridge Handbook of Thinking and Reasoning*. Cambridge University Press.

Golosnoy, Vasyl and Yarema Okhrin (2008). General Uncertainty in Portfolio Selection: A Case-Based Decision Approach. *Journal of Economic Behavior & Organization*, 67: 718-734.

Gonzalez, Cleotilde, Javier F. Lerch, and Christian Lebiere (2003). Instance-based Learning in Dynamic Decision Making. *Cognitive Science*, 27: 591-635.

Gonzalez, Cleotilde and Varun Dutt (2011). Instance-Based Learning: Integrating Sampling and Repeated Decisions from Experience. *Psychological Review*, 118(4): 523-51.

Gregan-Paxton, Jennifer and Jane Cote (2000). How Do Investors Make Predictions? Insights from Analogical Reasoning Research. *Journal of Behavioral Decision Making*, 13: 307-327.

Grether, David M. (1992). Testing Bayes Rule and the Representative Heuristic: Some Experimental Evidence. *Journal of Economic Behavior and Organization,* 17: 31-57.

Grosskopf, Brit, Rajiv Sarin, and Elizabeth Watson (2015). An Experiment on Case-Based Decision Making. *Theory and Decision,* 79(4): 639-666.

Guerdjikova, Ani (2007). Case-Based Learning with Different Similarity Functions. *Games and Economic Behavior*, 63: 107-132.

Gubaydullina, Zulia and Markus Spiworks (2009). Portfolio Diversification: An Experimental Study. Volkswirtschaftliches Seminar Diskussionsbeitraege Nr. 140, Universitaet Goettingen. Available at http://www.econstor.eu/bitstream/10419/31969/1/598709924.pdf.

Hedesström, Ted Martin, Henrik Svedsäter and Tommy Gärling (2006). Covariation Neglect Among Novice Investors. *Journal of Experimental Psychology: Applied*, 12(3): 155-165.

Hertwig, Ralph, Greg Barron, Elke U. Weber and Ido Erev (2004). Decisions from Experience and the Effect of Rare Events in Risky Choice. *Psychological Science*, 15(8): 534-539.

Hertwig, Ralph and Andreas Ortmann (2001). Experimental practices in economics: A methodological challenge for psychologists. *Behavioral and Brain Sciences*, 24: 383-451.

Hirshleifer, David A., Sonya Seongyeon Lim and Siew Hong Teoh (2009). Driven to Distraction: Extraneous Events and Underreaction to Earnings News. *Journal of Finance*, LXIV(5): 2289-2325.

Hoffmann, Arvid O. I. & Thomas Post (2014). Self-Attribution Bias in Consumer Financial Decision-Making: How Investment Returns Affect Individuals' Belief in Skill. *Journal of Behavioral and Experimental Economics*, 52: 23-28.

Hoffmann, Arvid O. I. & Hersh Shefrin (2014). Technical Analysis and Individual Investors. *Journal of Economic Behavior and Organization*, 107: 487-511.

Hoffmann, Arvid O. I. & Thomas Post. (forthcoming) How Return and Risk Experiences Shape Investor Beliefs and Preferences. *Accounting and Finance*. http://onlinelibrary.wiley.com/doi/10.1111/acfi.12169/abstract

Hogarth, Robin M. and Hillel J. Einhorn (1992). Order Effects in Belief Updating: The Belief Adjustment Model. *Cognitive Psychology*, 24: 1-55.

Huang, Xing (2013). Industry Investment Experience and Stock Selection. Available at http://dx.doi.org/10.2139/ssrn.1786271.

Isoni, Andrea, Graham Loomes, and Robert Sugden (2011). The Willingness to Pay–
Willingness to Accept Gap, the "Endowment Effect," Subject Misconceptions, and
Experimental Procedures for Eliciting Valuations: Comment. *The American
Economic Review*, 101: 991-1011.

Jahnke, Hermann, Anne Chwolka and Dirk Simons (2005). Coordinating Service-Sensitive
Demand and Capacity by Adaptive Decision Making: An Application of Case-Based
Decision Theory. *Decision Sciences*, 36(1): 1-32.

Jegadeesh, Narasimhan and Sheridan Titman (1993). Returns to Buying Winners and
Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance*,
48(1): 65-91.

Kahneman, Daniel and Amos Tversky (1972). Subjective Probability: A Judgment of
Representativeness. *Cognitive Psychology*, 3: 430-454.

Kallir, Ido and Doron Sonsino (2009). The Neglect of Correlation in Allocation Decisions.
Southern Economic Journal, 75(4): 1045-1066.

Kaustia, Markku and Samuli Knüpfer (2008). Do Investors Overweigh Personal
Experiences? Evidence from IPO Subscriptions. *Journal of Finance*, 63(6): 2679-
2702.

Keller, L. Robin, Uzi Segal and Tan Wang (1993). The Becker-DeGroot-Marschak
mechanism and generalized utility theories: Theoretical predictions and empirical
observations. *Theory and Decision*, 34(2): 83-97.

Kolodner, Janet L. (1992). An Introduction to Case-Based Reasoning. *Artificial
Intelligence Review,* 6: 3-34.

Krause, Andreas (2009). Learning and Herding using Case-based Decision Theory. *IEEE
Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans,*
39(3): 662-669.

Kroll, Yoram, Haim Levy and Amnon Rapoport (1988). Experimental Tests of the Separation Theorem and Capital Asset Pricing Model. *The American Economic Review*, 78(3): 500-519.

Leland, Jonathan W. (1994). Generalized Similarity Judgments: An Alternative Explanation for Choice Anomalies. *Journal of Risk and Uncertainty*, 9: 151-172.

Lo, Andrew W. and A. Craig MacKinlay (1988). Stock market prices do not follow random walks: evidence from a simple specification test. *The Review of Financial Studies*, 1(1): 41-66.

Malkiel, Burton G. (2003). The Efficient Market Hypothesis and Its Critics. *Journal of Economic Perspectives*, 17(1): 59-82.

_____ (2005). Reflections on the Efficient Market Hypothesis: 30 Years Later. *The Financial Review*, 40: 1-9.

Malmendier, Ulrike and Devin Shanthikumar (2007). Are Small Investors Naïve About Incentives? *Journal of Financial Economics*, 85(2): 457-489.

Markowitz, Harry (1952). Portfolio Selection. *The Journal of Finance*, 7(1): 77-91.

Matsui, Akihiko (2000). Expected Utility and Case-Based Reasoning. *Mathematical Social Sciences*, 39: 1-12.

Michaely, Roni and Kent L. Womack (2005). Brokerage Recommendations: Stylized Characteristics, Market Responses, and Biases. In Richard Thaler (ed), *Advances in Behavioral Finance II*, Princeton University Press and Russell Sage Foundation.

Müller, Hermann J., Thomas Geyer, Joseph Krummenacher and Michael Zehetleitner (2009). Attentional Capture by Salient Color Singleton Distractors Is Modulated by Top-Down Dimensional Set. *Journal of Experimental Psychology*, 35(1): 1-16.

Nosofsky, Robert M. (1988). Similarity, Frequency, and Category Representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(1): 54-65.

Nosofsky, Robert M. and Safa R. Zaki (2002). Exemplar and Prototype Models Revisited: Response Strategies, Selective Attention, and Stimulus Generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(5): 924-940.

Odean, Terrance (1998). Are Investors Reluctant to Realize Their Losses? *The Journal of Finance*, LIII(5): 1775-1798.

Offerman, Theo and Joep Sonnemans (2004). What's Causing Overreaction? An Experimental Investigation of Recency and the Hot-hand Effect. *Scandinavian Journal of Economics*, 106(3): 533-553.

Ossadnik, Wolfgang, Dirk Wilmsmann and Benedikt Niemann (2013). Experimental evidence on case-based decision theory. *Theory and Decision*, 75(2): 211-232.

Pape, Andreas Duus and Kenneth J. Kurtz (2013). Evaluating Case-Based Decision Theory: Predicting Empirical Patterns of Human Classification Learning. *Games and Economic Behavior*, 82: 52-65.

Peng, Lin and Wei Xiong (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80: 563-602.

Plott, Charles R. and Kathryn Zeiler (2005). The Willingness to Pay/Willingness to Accept Gap, the Endowment Effect, Subject Misconceptions and Experimental Procedures for Eliciting Valuations, *The American Economic Review*, 95(3), 530-545.

Riesbeck, Christopher K. and Roger C. Schank (1989). *Inside Case-Based Reasoning*. New Jersey: Lawrence Erlbaum Associates, Publishers.

Rogers, Paul (1998). The Cognitive Psychology of Lottery Gambling: A Theoretical Review. *Journal of Gambling Studies*, 14(2): 111-133.

Ross, Brian H. (1987). This is Like That: The Use of Earlier Problems and the Separation of Similarity Effects. *Journal of Experimental Psychology*, 13(4): 629-639.

Rubinstein, Ariel (1988). Similarity and Decision-making under Risk (Is There a Utility Theory Resolution to the Allais Paradox?). *Journal of Economic Theory*, 46: 145-153.

Seidl, Christian (2002). Preference Reversal. *Journal of Economic Surveys*, 16(5): 621-655.

Siegel, Sidney and N. John Castellan (1988). *Nonparametric Statistics for the Behavioral Sciences*, 2nd ed. McGraw-Hill.

Simon, Herbert A. (1957). A Behavioral Model of Rational Choice. In: *Models of Man*. John Wiley & Sons, Inc.

Smith, Edward E. and Douglas L. Medin (1981). *Categories and Concepts*. Harvard University Press.

Solnik, Bruno H. (1973). Note on the Validity of the Random Walk for European Stock Prices. *Journal of Finance*, 28(5): 1151-59.

Strahilevitz, Michal Ann, Terrance Odean, and Brad M. Barber (2011). Once Burned, Twice Shy: How Naïve Learning, Counterfactuals, and Regret Affect the Repurchase of Stocks Previously Sold. *Journal of Marketing Research*, XLVIII (Special Issue 2011): S102–S120.

Suetens, Sigrid and Jean-Robert Tyran (2012). The Gambler's Fallacy and Gender. *Journal of Economic Behavior & Organization*, 83(1): 118-124.

Sugden, Robert (2003). Reference-dependent subjective expected utility. *Journal of Economic Theory*, 111: 172-191.

_____ (2004). Alternatives to Expected Utility: Foundations and Concepts. In Salvador Barbera and Christian Seidl (eds), *Handbook of Utility Theory Volume 2: Extensions*. Kluwer Academic.

Sun, Yanlong and Hongbin Wang (2010). Gambler's Fallacy, Hot Hand Belief, and the Time of Patterns. *Judgment and Decision Making*, 5(2): 124-132.

Sutton, Richard S. and Andrew G. Barto (1998). *Reinforcement Learning: An Introduction*. The MIT Press.

Tversky, Amos (1977). Features of Similarity. *Psychological Review*, 84(4): 327-352.

Tversky, Amos and Itamar Gati (2004). Studies of Similarity. In Eldar Shafir (ed), *Preference, Belief, and Similarity: Selected Writings of Amos Tversky*. Bradford Books.

Tversky, Amos and J. Wesley Hutchinson (1986). Nearest Neighbor Analysis of Psychological Spaces. *Psychological Review*, 93(1): 3-22.

Wood, Gordon (1992). Predicting Outcomes: Sports and Stocks. *Journal of Gambling Studies*, 8(2): 201-222.

Wooldridge, Jeffrey M. (2002). *Econometric Analysis of Cross Section and Panel Data*. The MIT Press.

Zhao, Jinhua and Catherine L. Kling (2001). A new explanation for the WTP/WTA disparity. *Economics Letters*, 73: 293-300.

Zizzo, Daniel John (2003). Transfer of Knowledge in Economic Decision-making: An Overview. *Greek Economic Review*, 22(2): 1-10.

APPENDICES

Instructions (Part 1)

Welcome to today's experiment on decision-making. Today's session will begin shortly. Before we start, I have a few friendly reminders. First, to help us keep the lab neat and tidy, we ask that you do not eat or drink in the lab. Also, we ask that you turn off completely your mobile phones and other devices, as they may not be used during today's session. Please refrain from talking to other participants during the experiment. Finally, in the event that you should need to use the toilet during today's session, you will of course be permitted to do so, but it will delay the session while we wait for you. You may wish to take advantage of this opportunity to visit the toilets which are located down the corridor on the left.

A copy of the instructions is on your desk. Please follow along as I read through the instructions. If you have a question, please raise your hand and I will come and answer it privately.

On your desk is a consent form. Please read the form and sign it now.

In this experiment, you will make a decision that involves coloured tickets and this bingo cage. A coloured ticket entitles you to a draw from this bingo cage which contains balls with amounts of money written on them. The bingo cage contains 100 balls. Each of the 100 balls has either £0 or £20 written on it. You will not know how many balls of each kind there are.

Each of you will earn £2 for participating in today's session. You will have the opportunity to earn an additional amount of money which will depend on a decision you will make and on chance. You will receive your earnings before you leave today. I will now describe the tasks within the experiment.

This experiment has two parts. Part 1 consists of 10 rounds, while Part 2 has one round.

Your earnings will depend on the outcome of your decision in Part 2.

In Part 1, the 10 rounds are samples that will familiarize you with the bingo cage and will give you an idea of the possibility of drawing a ball marked £20.

At the start of each round, I will pick one envelope from a bag. Each envelope contains either a set of blue tickets or a set of yellow tickets, one ticket for each of you. The colour of the tickets will determine whether we are playing a Blue Round or a Yellow Round. I will then distribute the tickets. Pictures of the sample tickets are shown below.

Once you have your ticket, you will write down the round number and what you think is the chance that a £20 ball will be drawn in that round.

I will then draw a ball from the bingo cage. I will rotate the bingo cage, draw one ball, and show the amount written on the ball. If the ball drawn has £20 on it, your ticket will be worth £20.

If the ball drawn has £0 on it, your ticket will be worth £0.

At the bottom of your ticket is a sentence that reads: If I owned this ticket, my earnings would be £__.

After the draw, you will record in the blank the outcome of the draw.

On your desk are two coloured boxes. If your ticket is blue, you will drop it in the blue box; if it is yellow, you will drop it in the yellow box. I will then put the ball back in the bingo cage.

Therefore, the number of £20 and £0 balls remains unchanged for every draw.

In Part 2 of the experiment, you will be given ownership of either a blue or a yellow ticket, just like the sample tickets. Your ticket will entitle you to one draw from the bingo cage. This draw will be conducted in the same manner as in Part 1 using the same bingo cage. I will describe Part 2 in more detail after we complete Part 1.

Before we begin Part 1, are there any questions?

Instructions (Part 2)

We have now completed Part 1. I will now describe the task in Part 2.

In Part 2 of the experiment, you will be given ownership of either a blue or a yellow ticket, just like the sample tickets. Your ticket will entitle you to one draw from the bingo cage. This draw will be conducted in the same manner as in Part 1 using the same bingo cage.

Part 2 has one round where each of you will come forward for an individual draw from the bingo cage. Before the draw is made, you will have the opportunity to exchange your ticket for an amount of money. I will now describe this opportunity.

Each of you will now pick an envelope from this bag. Leave the envelope on your desk and open it only when I tell you to do so.

You may now open your envelope. Your envelope contains your coloured ticket and your decision form. Now write on your ticket what you think is the chance that when you come forward, you will draw a £20 ball.

Your ticket gives you the chance to earn money either by keeping your ticket and receiving the amount from your draw or exchanging your ticket for an amount of money.

Now, look at your decision form. At the top right of your decision form is a space for your participant number. Your participant number is the station number where you are seated. Please fill in the space now. Fill in the rest of the form only when I tell you to do so.

I am going to offer a price in exchange for your ticket. Here is a bag containing 35 envelopes. Each envelope contains one of 35 possible prices ranging from 20p to £20. Each price is listed on your decision form. I will now ask one of you to draw one envelope.

I will now post the envelope on the board. I will open it only after everyone has completed the decision form. The price in the envelope will be the price I will offer which we will call offer price.

[continued over the page]

Now, look at your decision form. You now have the opportunity to exchange your ticket for the offer price posted on the board. Listed on the decision form are possible offer prices that may be in the envelope. Think of each price individually and carefully consider whether you prefer to keep your ticket and receive the amount from your draw or you prefer to exchange your ticket and receive the offer price. At each price, you will tick the appropriate box to indicate which you prefer.

I will collect your decision form when you have completed filling it in. When I have collected everyone's decision form, I will open the envelope posted on the board to reveal the offer price. If at that price, you indicated that you preferred to exchange your ticket, you will receive the offer price. If at that price, you indicated that you preferred to keep your ticket, you will be entitled to the earnings from your draw.

You will each come forward for an individual draw. Here is how we will conduct your individual draw. When I announce your participant number, you will come forward for your individual draw. I will draw a ball for you in the same manner as in Part 1 using the same bingo cage. You will then record the outcome of the draw on your ticket.

If you decided to keep your ticket at the offer price, your earnings will be the outcome of your draw plus your participation fee. If you decided to exchange your ticket at the offer price, your earnings will be the offer price plus your participation fee.

After your individual draw, I will fill in the bottom part of your form, and I will hand you a receipt form. Fill in the receipt form but sign it only after you are actually paid.

Before we begin Part 2, are there any questions?

Participant No. _____

Decision Form

You have the opportunity to exchange your ticket for the amount of money in the envelope posted on the board. Below is a list of possible offer prices that may be in the envelope. Think of each price individually and carefully consider whether you prefer to keep your ticket and receive the amount from your draw or you prefer to exchange your ticket and receive the offer price. At each price, tick the appropriate box to indicate which you prefer.

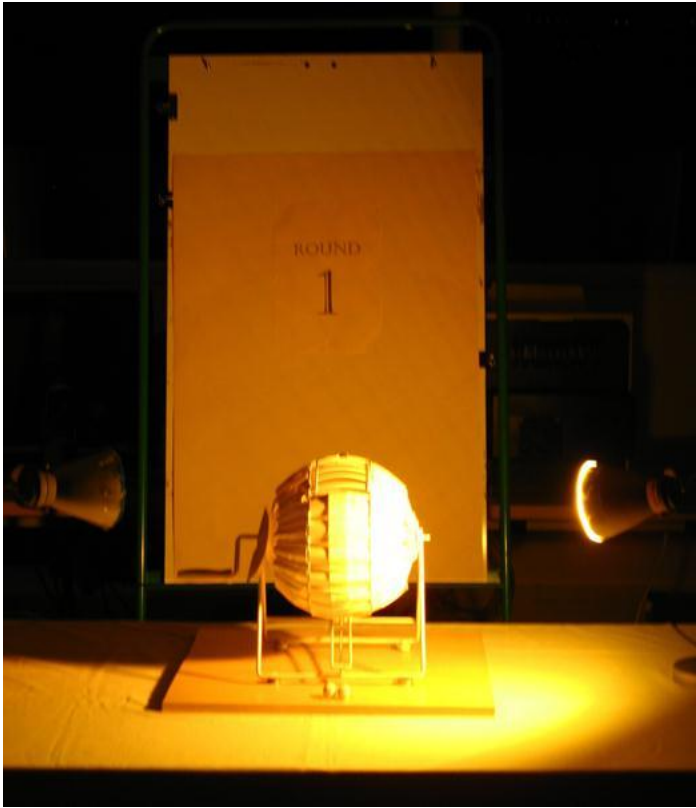|		
If the price offered is 20p	☐ I will keep my ticket	☐ I will exchange my ticket for 20p
If the price offered is 40p	☐ I will keep my ticket	☐ I will exchange my ticket for 40p
If the price offered is 60p	☐ I will keep my ticket	☐ I will exchange my ticket for 60p
If the price offered is 80p	☐ I will keep my ticket	☐ I will exchange my ticket for 80p
If the price offered is £1.00	☐ I will keep my ticket	☐ I will exchange my ticket for £1.00
If the price offered is £1.20	☐ I will keep my ticket	☐ I will exchange my ticket for £1.20
If the price offered is £1.40	☐ I will keep my ticket	☐ I will exchange my ticket for £1.40
If the price offered is £1.60	☐ I will keep my ticket	☐ I will exchange my ticket for £1.60
If the price offered is £1.80	☐ I will keep my ticket	☐ I will exchange my ticket for £1.80
If the price offered is £2.00	☐ I will keep my ticket	☐ I will exchange my ticket for £2.00
If the price offered is £2.20	☐ I will keep my ticket	☐ I will exchange my ticket for £2.20
If the price offered is £2.40	☐ I will keep my ticket	☐ I will exchange my ticket for £2.40
If the price offered is £2.60	☐ I will keep my ticket	☐ I will exchange my ticket for £2.60
If the price offered is £2.80	☐ I will keep my ticket	☐ I will exchange my ticket for £2.80
If the price offered is £3.00	☐ I will keep my ticket	☐ I will exchange my ticket for £3.00
If the price offered is £3.50	☐ I will keep my ticket	☐ I will exchange my ticket for £3.50
If the price offered is £4.00	☐ I will keep my ticket	☐ I will exchange my ticket for £4.00
If the price offered is £4.50	☐ I will keep my ticket	☐ I will exchange my ticket for £4.50
If the price offered is £5.00	☐ I will keep my ticket	☐ I will exchange my ticket for £5.00
If the price offered is £5.50	☐ I will keep my ticket	☐ I will exchange my ticket for £5.50
If the price offered is £6.00	☐ I will keep my ticket	☐ I will exchange my ticket for £6.00
If the price offered is £6.50	☐ I will keep my ticket	☐ I will exchange my ticket for £6.50
If the price offered is £7.00	☐ I will keep my ticket	☐ I will exchange my ticket for £7.00
If the price offered is £7.50	☐ I will keep my ticket	☐ I will exchange my ticket for £7.50
If the price offered is £8.00	☐ I will keep my ticket	☐ I will exchange my ticket for £8.00

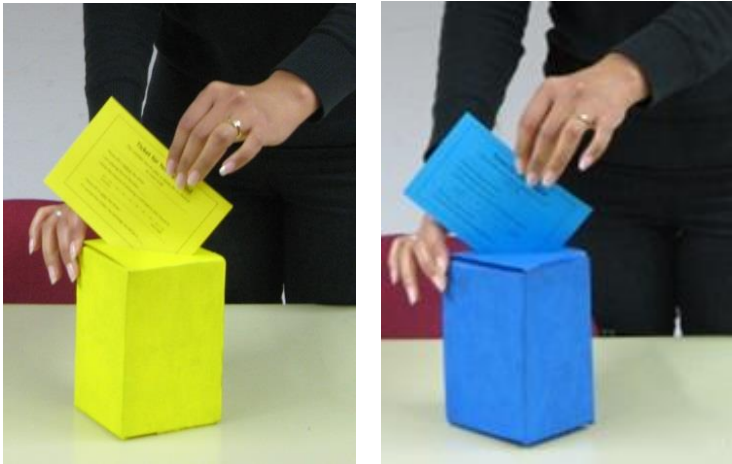|		
If the price offered is £8.50	☐ I will keep my ticket	☐ I will exchange my ticket for £8.50
If the price offered is £9.00	☐ I will keep my ticket	☐ I will exchange my ticket for £9.00
If the price offered is £10.00	☐ I will keep my ticket	☐ I will exchange my ticket for £10.00
If the price offered is £11.00	☐ I will keep my ticket	☐ I will exchange my ticket for £11.00
If the price offered is £12.00	☐ I will keep my ticket	☐ I will exchange my ticket for £12.00
If the price offered is £13.00	☐ I will keep my ticket	☐ I will exchange my ticket for £13.00
If the price offered is £14.00	☐ I will keep my ticket	☐ I will exchange my ticket for £14.00
If the price offered is £16.00	☐ I will keep my ticket	☐ I will exchange my ticket for £16.00
If the price offered is £18.00	☐ I will keep my ticket	☐ I will exchange my ticket for £18.00
If the price offered is £20.00	☐ I will keep my ticket	☐ I will exchange my ticket for £20.00

TO BE FILLED IN BY THE EXPERIMENTER

	Earnings
The offer price was £_____. At this price, you ○ kept your ticket. ○ exchanged your ticket for the offer price. *(If you exchanged your ticket for the offer price, write the offer price on the second column. Otherwise, leave it blank.)*	£
The outcome of your individual draw was £_____. *(If you kept your ticket, write the outcome of your individual draw on the second column. If you exchanged your ticket for the offer price, leave it blank.)*	£
Participation Fee	£ 2
Total	£

Appendix 1.3: Selected Experiment Photos

Sample yellow round

Coloured tickets and boxes

Part 1 Instructions

Please follow along as I read through the instructions. If you have a question, please raise your hand and I will come to answer your question privately.

Each of you will earn £2 for participating in today's session. You will have the opportunity to earn an additional amount of money which will depend on a decision you will make and on chance. You will receive your earnings before you leave today.

This experiment involves two coloured game boards: a blue game board and a yellow game board. These game boards have been set up on your computer. Towards the end of the experiment, you will make a decision that involves one of the game boards.

This experiment has two parts. Part 1 consists of ten rounds, while Part 2 has one round. Your earnings will depend on the outcome of your decision in Part 2.

In Part 1, the ten rounds are samples that will give you the opportunity to learn as much as you can about the game boards. Each game board contains 100 boxes numbered from 1 to 100. Each box is either a winning box or a losing box. A winning box has a value to you of £20, while a losing box has a value to you of £0.

At the start of the experiment, you will not know the values of the boxes. You will not even know how many winning boxes there are on each game board. Each game board has either ten or thirty winning boxes.

In each of the ten rounds, the value of a different box will be revealed. At the end of Part 1, you will have seen the value of ten boxes.

We will now fix how many winning boxes there are on each game board, and what those box numbers are.

Here is a bingo cage containing 100 balls numbered from 1 to 100. Before today's experiment, we picked forty balls from this bingo cage. We put thirty of these numbers into one envelope and the other ten numbers into a different envelope.

Here are the two envelopes. I will now put the envelopes in this bag.

I will now ask one of you to come forward. Please pick one envelope from the bag but do not open the envelope. Please write "blue and yellow" on that envelope. Now, take the other envelope from the bag and write "unused" on that envelope.

The numbers inside this envelope marked "blue and yellow" are the numbers of the winning boxes for both the blue and yellow game boards. This means that the winning boxes are the same for the two game boards. Depending on which envelope was picked, each game board may have ten winning boxes, or it may have thirty winning boxes.

My assistant will open the envelope and then programme the winning boxes into the game boards. I will post the envelopes on the wall after we finish reading the instructions. You are welcome to inspect the envelopes at the end of the experiment.

I will now describe the computer screens you will encounter in each round.

At the start of each round, your computer will display either the blue game board or the yellow game board.

Look at the sample screens on the next page.

[continued over the page]

Sample Round No. 1
The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 30 out of 100   10 out of 100 10 out of 100

I think the chance that this game board will reveal a winning box in this round is

very low ⊙⊙⊙⊙⊙⊙⊙⊙⊙ very high

confirm

I AM PLAYING THE BLUE GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

I think the chance that this game board will reveal a winning box in this round is

very low ⊙⊙⊙⊙⊙⊙⊙⊙⊙ very high

confirm

I AM PLAYING THE YELLOW GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

The top screen is the screen you will see if you are playing the blue game board. You will always play the blue game board on the left side of your screen. The bottom screen is the screen you will see if you are playing the yellow game board. You will always play the yellow game board on the right side of your screen.

[continued over the page]

Sample Round No. 1
The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 30 out of 100   10 out of 100 10 out of 100

Here is a picture of the top portion of the sample screens shown on the previous page. It shows the possible sets of winning boxes on the blue and yellow game boards. Remember that the set of winning boxes depends on the contents of the envelope marked "blue and yellow".

On the top-left of the sample screen is the case where both game boards have thirty winning boxes. This possibility occurs if the envelope that was marked "blue and yellow" contains thirty winning numbers. On the top-right of the screen is the case where both game boards have ten winning boxes. This other possibility occurs if the envelope that was marked "blue and yellow" has ten winning numbers. These two possibilities are equally likely.

In each round, you will indicate on your screen what you think is the chance that the game board will reveal a winning box in that round. You will then click on the button labelled confirm. A sample of that section on the screen is highlighted on the game board below.

I will then draw a ball from the bingo cage. I will announce the box number printed on the ball and you will click on that box to open the box and reveal whether it is a winning box or a losing box. A winning box is shaded green and displays £20, while a losing box is shaded red and displays £0.

[continued over the page]

Let us consider two examples: a case with a winning box and a case with a losing box. Below is a sample screen you will see if the announced box number is a winning box.

Suppose that we had revealed the value of box number 5 in sample round number 1. We are in sample round number 2 where you are playing the blue game board.

I announce that the box number for this round is "1". You click on that box to open the box and reveal whether it is a winning box or a losing box. Since box number 1 is a winning box in this example, it displays £20 and it is shaded green.

Notice that box number 5 is shaded grey. The boxes with values revealed in the previous rounds are shaded grey. These boxes will not be opened again in the remaining rounds in Part 1.

Your computer will then show a screen that summarises what you learned about the game boards in that round. A sample of that screen is shown on the next page.

Sample Round No. 2
The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

I AM PLAYING THE BLUE GAME BOARD

£20	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

1 is a winning box on the blue game board

1 is a winning box on the yellow game board

The left side of the screen reminds you that you just played the blue game board in sample round number 2. It also reminds you that box number 1 on the blue game board is a winning box. Because box number 1 is a winning box, it is shaded green and it displays £20 on the blue game board.

Now look at the right side of the screen. Because the winning boxes are the same for the two game boards, box number 1 is also a winning box on the yellow game board.

Now suppose instead that box number 1 is a losing box on the blue game board. Below is a sample screen you will see in this situation.

I announce that the box number for this round is "1". You click on that box to open the box and reveal whether it is a winning box or a losing box. Since box number 1 is a losing box in this example, it displays £0 and it is shaded red.

Your computer will then show a screen that summarises what you learned about the game boards in that round. A sample of that screen is shown on the next page.

[continued over the page]

Sample Round No. 2
The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

I AM PLAYING THE BLUE GAME BOARD

£0	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

1 is a losing box on the blue game board

1 is a losing box on the yellow game board

The left side of the sample screen reminds you that box number 1 is a losing box on the blue game board, therefore, it is shaded red and it displays £0.

Now look at the right side of the sample screen. Because the winning boxes are the same for the two game boards, box number 1 is also a losing box on the yellow game board.

Part 2 has one round where each of you will be given one play of either the blue game board or the yellow game board. The game board will have the same set of winning boxes programmed into it as in Part 1.

You will then have the opportunity to earn an amount of money, in addition to your participation fee of £2. The additional amount of money you will earn will depend on the outcome of your decision in Part 2.

At the beginning of Part 2, I will put back all balls drawn in Part 1 so that it is possible to draw a box number that was opened in Part 1.

I will describe Part 2 in more detail after we complete Part 1.

Are there any questions?

Part 2 Instructions

We have completed Part 1. I will now describe the task in Part 2.

Part 2 has one round where each of you will play either the blue game board or the yellow game board. Each coloured game board has the same set of winning boxes programmed into it as in Part 1. Recall that the set of winning boxes on the blue and yellow game boards are the box numbers listed on the envelope marked "blue and yellow".

Your game board gives you the chance to earn money either by keeping your game board and receiving the earnings from your play of it, or exchanging your game board for an amount of money.

Before we begin, my assistant will now put back in the bingo cage all balls drawn in Part 1.

This bingo cage now contains 100 balls.

Each of you will have an individual draw from this bingo cage. The ball I draw will determine the box number to be opened on your game board. Because all the balls selected during Part 1 have been returned in the bingo cage, it is possible for me to draw any box number from 1 to 100, including the boxes that were opened in Part 1.

Please click on the button labelled continue.

Your computer screen now displays your game board. Indicate what you think is the chance that your game board will reveal a winning box when we conduct your individual draw. After you have done so, click on the button labelled confirm.

I am going to offer a price in exchange for your game board.

Here is a bag containing thirty-five envelopes. Each envelope contains one of thirty-five possible prices ranging from 20p to £20. Each price is listed on a decision form that will be shown on your computer screen.

I will now ask one of you to draw one envelope from this bag but do not open the envelope. The price in the envelope will be the price I will offer in exchange for your game board. We will call this the offer price.

My assistant will post the envelope on the wall. I will open the envelope only after everyone has submitted their decision form.

Look again at your computer screen. Your decision form gives you the opportunity to exchange the result of your play of your game board for the offer price posted on the wall. Listed on the decision form are all the possible offer prices that may be in the envelope. Think of each price individually. At each price, carefully consider whether you prefer to keep your game board and receive the earnings from your play of it, or you prefer to exchange your game board and receive that offer price. For each price, click on the appropriate button to indicate which you prefer.

After everyone has submitted their decision form, I will open the envelope posted on the wall to reveal the offer price. I will announce the offer price and my assistant will input the offer price into the computer. Your computer will then remind you of your decision at that offer price.

I will then go to each of you for your individual draw. I will draw one ball from the bingo cage and I will show you the box number printed on the ball. You will then click on that box as you did in Part 1 to open the box and reveal the value of that box to you.

I will then return the ball before conducting the individual draw for the next participant. All 100 balls will be in the bingo cage when we conduct your individual draw.

If you decided to keep your game board at the offer price, you will receive the earnings from your play of your game board plus your participation fee. If you decided to exchange your game board at the offer price, your earnings will be the offer price plus your participation fee.

Before we begin Part 2, are there any questions?

Part 1 Instructions

Please follow along as I read through the instructions. If you have a question, please raise your hand and I will come to answer your question privately.

Each of you will earn £2 for participating in today's session. You will have the opportunity to earn an additional amount of money which will depend on a decision you will make and on chance. You will receive your earnings before you leave today.

This experiment involves two coloured game boards: a blue game board and a yellow game board. These game boards have been set up on your computer. Towards the end of the experiment, you will make a decision that involves one of the game boards.

This experiment has two parts. Part 1 consists of ten rounds, while Part 2 has one round. Your earnings will depend on the outcome of your decision in Part 2.

In Part 1, the ten rounds are samples that will give you the opportunity to learn as much as you can about the game boards. Each game board contains 100 boxes numbered from 1 to 100. Each box is either a winning box or a losing box. A winning box has a value to you of £20, while a losing box has a value to you of £0.

At the start of the experiment, you will not know the values of the boxes. You will not even know how many winning boxes there are on each game board. Each game board has either ten or thirty winning boxes.

In each of the ten rounds, the value of a different box will be revealed. At the end of Part 1, you will have seen the value of ten boxes.

We will now fix how many winning boxes there are on each game board, and what those box numbers are.

Here is a bingo cage containing 100 balls numbered from 1 to 100. Before today's experiment, we picked forty balls from this bingo cage. We put thirty of these numbers into one envelope and the other ten numbers into a different envelope. Here are the two envelopes. I will now put the envelopes in this bag.

We returned all the balls in the bingo cage and then picked another set of forty balls. We put thirty of these numbers into one envelope and the other ten numbers into a different envelope. Here are the two other envelopes. I will now put the envelopes in this other bag.

I will now ask one of you to come forward. Please pick one envelope from the first bag but do not open the envelope. Please write "blue" on that envelope. Now, take the other envelope from the bag and write "unused for blue" on that envelope.

Here is the second bag. Now pick one envelope and then write "yellow" on that envelope. Now, take the other envelope from the bag and write "unused for yellow" on that envelope.

The numbers inside the envelope marked "blue" are the numbers of the winning boxes for the blue game board. The numbers inside the envelope marked "yellow" are the numbers of the winning boxes for the yellow game board. Since there are two sets of forty winning boxes, one for the blue game board and another for the yellow game board, the winning boxes for the two game boards may be the same or may be different. Depending on which envelopes were picked, each game board may have either ten or thirty winning boxes.

My assistant will open the envelope and then programme the winning boxes into the game boards. I will post the envelopes on the wall after we finish reading the instructions. You are welcome to inspect the envelopes at the end of the experiment.

I will now describe the computer screens you will encounter in each round.

[continued over the page]

At the start of each round, your computer will display either the blue game board or the yellow game board. Look at the sample screens below.

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 | 30 out of 100   10 out of 100 | 10 out of 100   30 out of 100 | 10 out of 100   10 out of 100 | 30 out of 100

I think the chance that this game board will reveal a winning box in this round is

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

confirm

I AM PLAYING THE BLUE GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 | 30 out of 100   10 out of 100 | 10 out of 100   30 out of 100 | 10 out of 100   10 out of 100 | 30 out of 100

I think the chance that this game board will reveal a winning box in this round is

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

confirm

I AM PLAYING THE YELLOW GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

The top screen is the screen you will see if you are playing the blue game board. You will always play the blue game board on the left side of your screen. The bottom screen is the screen you will see if you are playing the yellow game board. You will always play the yellow game board on the right side of your screen.

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 | | 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

Here is a picture of the top portion of the sample screens shown on the previous page. It shows you the possible sets of winning boxes on the blue and yellow game boards. Remember that the sets of winning boxes depend on the contents of the envelopes marked "blue" and marked "yellow".

Starting from the top-left, the first possibility is that the blue and yellow game boards both have thirty winning boxes. This possibility occurs if the envelopes marked "blue" and marked "yellow" both contain thirty winning numbers. The second possibility is that the blue and yellow game boards both have ten winning boxes. This possibility occurs if the envelopes marked "blue" and marked "yellow" both contain ten winning numbers. The third possibility is that the blue game board has thirty winning boxes while the yellow game board has ten winning boxes. This possibility occurs if the envelope that was marked "blue" contains thirty winning numbers and the envelope marked "yellow" has ten winning numbers. The fourth possibility is that the blue game board has ten winning boxes while the yellow game board has thirty winning boxes. This possibility occurs if the envelope that was marked "blue" contains ten winning numbers and the envelope marked "yellow" has thirty winning numbers. These four possibilities are equally likely.

In each round, you will indicate on your screen what you think is the chance that the game board will reveal a winning box in that round. You will then click on the button labelled confirm. A sample of that section on the screen is highlighted on the game board below.

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 | | 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

I think the chance that this game board will reveal a winning box in this round is

very low　○○○○○○○○○　very high

confirm

I AM PLAYING THE YELLOW GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

[continued over the page]

I will then draw a ball from the bingo cage. I will announce the box number printed on the ball and you will click on that box to open the box and reveal whether it is a winning box or a losing box. A winning box is shaded green and displays £20, while a losing box is shaded red and displays £0.

Let us consider two examples: a case with a winning box and a case with a losing box. Below is a sample screen you will see if the announced box number is a winning box.

Suppose that we had revealed the value of box number 5 in sample round number 1. We are in sample round number 2 where you are playing the blue game board.

I announce that the box number for this round is "1". You click on that box to open the box and reveal whether it is a winning box or a losing box. Since box number 1 is a winning box in this example, it displays £20 and it is shaded green.

Notice that box number 5 is shaded grey. The boxes with values revealed in the previous rounds are shaded grey. These boxes will not be opened again in the remaining rounds in Part 1.

Your computer will then show a screen that summarises what you learned about the game boards in that round. A sample of that screen is shown on the next page.

Sample Round No. 2
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 | | 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ very high

I AM PLAYING THE BLUE GAME BOARD

£20	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

1 is a winning box on the blue game board

1 may be a winning box or may be a losing box on the yellow game board

The left side of the screen reminds you that you just played the blue game board in sample round number 2. It also reminds you that box number 1 on the blue game board is a winning box. Because box number 1 is a winning box, it is shaded green and it displays £20 on the blue game board.

Now look at the right side of the screen. Since there are two sets of forty winning boxes, one for the blue game board and another for the yellow game board, the winning boxes for the two game boards may be the same or may be different. This means that while box number 1 is a winning box on the blue game board, it may be a winning box or may be a losing box on the yellow game board.

[continued over the page]

Now suppose instead that box number 1 is a losing box on the blue game board. Below is a sample screen you will see in this situation.

I announce that the box number for this round is "1". You click on that box to open the box and reveal whether it is a winning box or a losing box. Since box number 1 is a losing box in this example, it displays £0 and it is shaded red.

Your computer will then show a screen that summarises what you learned about the game boards in that round. A sample of that screen is shown on the next page.

The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 | | 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

I AM PLAYING THE BLUE GAME BOARD

£0	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

1 is a losing box on the blue game board

1 may be a winning box or may be a losing box on the yellow game board

The left side of the sample screen reminds you that box number 1 is a losing box on the blue game board, therefore, it is shaded red and it displays £0.

Now look at the right side of the sample screen. Recall that there are two sets of forty winning boxes, one for the blue game board and another for the yellow game board. This means that the winning boxes for the two game boards may be the same or may be different. In the example, while box number 1 is a losing box on the blue game board, it may be a winning box or may be a losing box on the yellow game board.

Part 2 has one round where each of you will be given one play of either the blue game board or the yellow game board. The game board will have the same set of winning boxes programmed into it as in Part 1.

You will then have the opportunity to earn an amount of money, in addition to your participation fee of £2. The additional amount of money you will earn will depend on the outcome of your decision in Part 2.

At the beginning of Part 2, I will put back all balls drawn in Part 1 so that it is possible to draw a box number that was opened in Part 1.

I will describe Part 2 in more detail after we complete Part 1.

Are there any questions?

Part 2 Instructions

We have completed Part 1. I will now describe the task in Part 2.

Part 2 has one round where each of you will play either the blue game board or the yellow game board. Each coloured game board has the same set of winning boxes programmed into it as in Part 1. Recall that the set of winning boxes on the blue game board are the box numbers listed on the envelope marked "blue" while the set of winning boxes on the yellow game board are the box numbers listed on the envelope marked "yellow".

Your game board gives you the chance to earn money either by keeping your game board and receiving the earnings from your play of it, or by exchanging your game board for an amount of money.

Before we begin, my assistant will now put back in the bingo cage all balls drawn in Part 1.

This bingo cage now contains 100 balls.

Each of you will have an individual draw from this bingo cage. The ball I draw will determine the box number to be opened on your game board. Because all the balls selected during Part 1 have been returned in the bingo cage, it is possible for me to draw any box number from 1 to 100, including the boxes that were opened played in Part 1.

Please click on the button labelled continue. Your computer screen now displays your game board. Indicate what you think is the chance that your game board will reveal a winning box when we conduct your individual draw. After you have done so, click on the button labelled confirm.

I am going to offer a price in exchange for your game board.

Here is a bag containing thirty-five envelopes. Each envelope contains one of thirty-five possible prices ranging from 20p to £20. Each price is listed on a decision form that will be shown on your computer screen.

I will now ask one of you to draw one envelope from this bag but do not open the envelope. The price in the envelope will be the price I will offer in exchange for your game board. We will call this the offer price.

My assistant will post the envelope on the wall. I will open the envelope only after everyone has submitted their decision form.

Look again at your computer screen. Your decision form gives you the opportunity to exchange the result of your play of your game board for the offer price posted on the wall. Listed on the decision form are all the possible offer prices that may be in the envelope. Think of each price individually. At each price, carefully consider whether you prefer to keep your game board and receive the earnings from your play of it, or you prefer to exchange your game board and receive that offer price. For each price, click on the appropriate button to indicate which you prefer.

After everyone has submitted their decision form, I will open the envelope posted on the wall to reveal the offer price. I will announce the offer price and my assistant will input the offer price into the computer. Your computer will then remind you of your decision at that offer price.

I will then go to each of you for your individual draw. I will draw one ball from the bingo cage and I will show you the box number printed on the ball. You will then click on that box as you did in Part 1 to open the box and reveal the value of that box to you.

I will then return the ball before conducting the individual draw for the next participant. All 100 balls will be in the bingo cage when we conduct your individual draw.

If you decided to keep your game board at the offer price, you will receive the earnings from your play of your game board plus your participation fee. If you decided to exchange your game board at the offer price, your earnings will be the offer price plus your participation fee.

Before we begin Part 2, are there any questions?

Part 1 Instructions

Please follow along as I read through the instructions. If you have a question, please raise your hand and I will come to answer your question privately.

Each of you will earn £2 for participating in today's session. You will have the opportunity to earn an additional amount of money which will depend on a decision you will make and on chance. You will receive your earnings before you leave today.

This experiment involves two coloured game boards: a blue game board and a yellow game board. These game boards have been set up on your computer. Towards the end of the experiment, you will make a decision that involves one of the game boards.

This experiment has two parts. Part 1 consists of ten rounds, while Part 2 has one round. Your earnings will depend on the outcome of your decision in Part 2.

In Part 1, the ten rounds are samples that will give you the opportunity to learn as much as you can about the game boards. Each game board contains 100 boxes numbered from 1 to 100. Each box is either a winning box or a losing box. A winning box has a value to you of £20, while a losing box has a value to you of £0.

At the start of the experiment, you will not know the values of the boxes. You will not even know how many winning boxes there are on each game board. Each game board has either ten or thirty winning boxes.

In each of the ten rounds, the value of a different box will be revealed. At the end of Part 1, you will have seen the value of ten boxes.

We will now fix how many winning boxes there are on each game board, and what those box numbers are.

Here is a bingo cage containing 100 balls numbered from 1 to 100. Before today's experiment, we picked forty balls from this bingo cage. We put thirty of these numbers into one envelope and the other ten numbers into a different envelope.

Here are the two envelopes. I will now put the envelopes in this bag.

I will now ask one of you to come forward. Please pick one envelope from the bag but do not open the envelope. Please write "blue" on that envelope. Now, take the other envelope from the bag and write "yellow" on that envelope.

The numbers inside this envelope marked "blue" are the numbers of the winning boxes for the blue game board, while the numbers inside this envelope marked "yellow" are the numbers of the winning boxes for the yellow game board. This means that the winning boxes are different for the two game boards. Depending on which envelopes were picked, one game board will have ten winning boxes, and the other game board will have thirty winning boxes.

My assistant will open the envelope and then programme the winning boxes into the game boards. I will post the envelopes on the wall after we finish reading the instructions. You are welcome to inspect the envelopes at the end of the experiment.

I will now describe the computer screens you will encounter in each round.

At the start of each round, your computer will display either the blue game board or the yellow game board.

Look at the sample screens on the next page.

[continued over the page]

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100   10 out of 100   10 out of 100   30 out of 100

I think the chance that this game board will reveal a winning box in this round is

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

confirm

I AM PLAYING THE BLUE GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100   10 out of 100   10 out of 100   30 out of 100

I think the chance that this game board will reveal a winning box in this round is

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

confirm

I AM PLAYING THE YELLOW GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

The top screen is the screen you will see if you are playing the blue game board. You will always play the blue game board on the left side of your screen. The bottom screen is the screen you will see if you are playing the yellow game board. You will always play the yellow game board on the right side of your screen.

Sample Round No. 1
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 | 10 out of 100   10 out of 100 | 30 out of 100

Here is a picture of the top portion of the sample screens shown on the previous page. It shows the possible sets of winning boxes on the blue and yellow game boards. Remember that the sets of winning boxes depend on the contents of the envelopes marked "blue" and marked "yellow".

On the top-left of the screen is the case where the blue game board has thirty winning boxes and the yellow game board has ten winning boxes. This possibility occurs if the envelope that was marked "blue" contains thirty winning numbers and the envelope marked "yellow" has ten winning numbers. On the top-right of the screen is the case where the blue game board has ten winning boxes and the yellow game board has thirty winning boxes. This other possibility occurs if the envelope that was marked "blue" has ten winning numbers and the envelope marked "yellow" has thirty winning numbers. These two possibilities are equally likely.

In each round, you will indicate on your screen what you think is the chance that the game board will reveal a winning box in that round. You will then click on the button labelled confirm. A sample of that section on the screen is highlighted on the game board below.

I will then draw a ball from the bingo cage. I will announce the box number printed on the ball and you will click on that box to open the box and reveal whether it is a winning box or a losing box. A winning box is shaded green and displays £20, while a losing box is shaded red and displays £0.

[continued over the page]

Let us consider two examples: a case with a winning box and a case with a losing box. Below is a sample screen you will see if the announced box number is a winning box.

Suppose that we had revealed the value of box number 5 in sample round number 1. We are in sample round number 2 where you are playing the blue game board.

I announce that the box number for this round is "1". You click on that box to open the box and reveal whether it is a winning box or a losing box. Since box number 1 is a winning box in this example, it displays £20 and it is shaded green.

Notice that box number 5 is shaded grey. The boxes with values revealed in the previous rounds are shaded grey. These boxes will not be opened again in the remaining rounds in Part 1.

Your computer will then show a screen that summarises what you learned about the game boards in that round. A sample of that screen is shown on the next page.

Sample Round No. 2
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ very high

I AM PLAYING THE BLUE GAME BOARD

£20	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

1 is a winning box on the blue game board

1 is a losing box on the yellow game board

The left side of the screen reminds you that you just played the blue game board in sample round number 2. It also reminds you that box number 1 on the blue game board is a winning box. Because box number 1 is a winning box, it is shaded green and it displays £20 on the blue game board.

Now look at the right side of the screen. Because the winning boxes are different for the two game boards, box number 1 cannot be a winning box on the yellow game board. This means that box number 1 on the yellow game board is a losing box.

[continued over the page]

Now suppose instead that box number 1 is a losing box on the blue game board. Below is a sample screen you will see in this situation.

Sample Round No. 2
The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ ⃝ very high

I AM PLAYING THE BLUE GAME BOARD

£0	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

I announce that the box number for this round is "1". You click on that box to open the box and reveal whether it is a winning box or a losing box. Since box number 1 is a losing box in this example, it displays £0 and it is shaded red.

Your computer will then show a screen that summarises what you learned about the game boards in that round. A sample of that screen is shown on the next page.

[continued over the page]

The winning boxes depend on the envelopes drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 10 out of 100 | | 10 out of 100 | 30 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ○ ○ ○ ○ ○ ○ ○ ○ ○ very high

I AM PLAYING THE BLUE GAME BOARD

£0	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

1 is a losing box on the blue game board

1 may be a winning box or may be a losing box on the yellow game board

The left side of the sample screen reminds you that box number 1 is a losing box on the blue game board, therefore, it is shaded red and it displays £0.

Now look at the right side of the sample screen. Recall that there are forty winning boxes, and the winning boxes are different for the blue and yellow game boards. Since box number 1 is a losing box on the blue game board, box number 1 may be a winning box or may be a losing box on the yellow game board.

Part 2 has one round where each of you will be given one play of either the blue game board or the yellow game board. The game board will have the same set of winning boxes programmed into it as in Part 1.

You will then have the opportunity to earn an amount of money, in addition to your participation fee of £2. The additional amount of money you will earn will depend on the outcome of your decision in Part 2.

At the beginning of Part 2, I will put back all balls drawn in Part 1 so that it is possible to draw a box number that was opened in Part 1.

I will describe Part 2 in more detail after we complete Part 1.

Are there any questions?

Part 2 Instructions

We have completed Part 1. I will now describe the task in Part 2.

Part 2 has one round where each of you will play either the blue game board or the yellow game board. Each coloured game board has the same set of winning boxes programmed into it as in Part 1. Recall that the set of winning boxes on the blue game board are the box numbers listed on the envelope marked "blue" while the set of winning boxes on the yellow game board are the box numbers listed on the envelope marked "yellow".

Your game board gives you the chance to earn money either by keeping your game board and receiving the earnings from your play of it, or by exchanging your game board for an amount of money.

Before we begin, my assistant will now put back in the bingo cage all balls drawn in Part 1.

This bingo cage now contains 100 balls.

Each of you will have an individual draw from this bingo cage. The ball I draw will determine the box number to be opened on your game board. Because all the balls selected during Part 1 have been returned in the bingo cage, it is possible for me to draw any box number from 1 to 100, including the boxes that were opened played in Part 1.

Please click on the button labelled continue. Your computer screen now displays your game board. Indicate what you think is the chance that your game board will reveal a winning box when we conduct your individual draw. After you have done so, click on the button labelled confirm.

I am going to offer a price in exchange for your game board.

Here is a bag containing thirty-five envelopes. Each envelope contains one of thirty-five possible prices ranging from 20p to £20. Each price is listed on a decision form that will be shown on your computer screen.

I will now ask one of you to draw one envelope from this bag but do not open the envelope. The price in the envelope will be the price I will offer in exchange for your game board. We will call this the offer price.

My assistant will post the envelope on the wall. I will open the envelope only after everyone has submitted their decision form.

Look again at your computer screen. Your decision form gives you the opportunity to exchange the result of your play of your game board for the offer price posted on the wall. Listed on the decision form are all the possible offer prices that may be in the envelope. Think of each price individually. At each price, carefully consider whether you prefer to keep your game board and receive the earnings from your play of it, or you prefer to exchange your game board and receive that offer price. For each price, click on the appropriate button to indicate which you prefer.

After everyone has submitted their decision form, I will open the envelope posted on the wall to reveal the offer price. I will announce the offer price and my assistant will input the offer price into the computer. Your computer will then remind you of your decision at that offer price.

I will then go to each of you for your individual draw. I will draw one ball from the bingo cage and I will show you the box number printed on the ball. You will then click on that box as you did in Part 1 to open the box and reveal the value of that box to you.

I will then return the ball before conducting the individual draw for the next participant. All 100 balls will be in the bingo cage when we conduct your individual draw.

If you decided to keep your game board at the offer price, you will receive the earnings from your play of your game board plus your participation fee. If you decided to exchange your game board at the offer price, your earnings will be the offer price plus your participation fee.

Before we begin Part 2, are there any questions?

Appendix 2.4: Quiz screenshots

Screenshot of questions answered by participants in the positive treatment

WELCOME TO THE EXPERIMENT!

You will now be shown two sample screens, each followed by a series of questions.

We will begin Part 1 once all participants have answered all questions correctly.

Raise your hand if you have a question and an experimenter will come to answer your question privately.

Click on the button labelled next to begin.

[next]

Imagine that you just completed Sample Round No. 1, 2 and 3.

You are in Sample Round No. 4.

You are playing the yellow game board where the value of box number 75 is revealed.

Click on the button labelled next to continue.

[next]

Sample Round No. 4
The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 | **30 out of 100**   **10 out of 100** | **10 out of 100**

Please answer the following questions.

I thought the chance that this game board would reveal a winning box in this round was

very low   ○ ○ ○ ○ ○ ○ ○ ○ ○   very high

1. Refer to the top portion of the screen. If the yellow game board has ten winning boxes, how many winning boxes does the blue game board have?

- Ten. The blue and yellow game boards have the same winning boxes
- Thirty. The blue and yellow game boards have different winning boxes
- Ten or Thirty. The blue and yellow game boards may have the same or different winning boxes

I AM PLAYING THE YELLOW GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

2. Refer to the grid of 100 boxes on the right. Why is box number 75 shaded green on the yellow game board?

- box number 75 is a losing box on the yellow game board
- box number 75 is a winning box on the yellow game board
- box number 75 was played in a previous round

3. Why are box numbers 5, 10 and 30 shaded grey on the yellow game board?

- boxes 5, 10 and 30 have a value of £0 on the yellow game board
- boxes 5, 10 and 30 have a value of £20 on the yellow game board
- the values of boxes 5, 10 and 30 were revealed in previous rounds

4. Which statement is TRUE?

- The blue and yellow game boards have the same winning boxes.
- The blue and yellow game boards have different winning boxes
- The blue and yellow game boards may have the same or different winning boxes

Please click on the button labelled check answers.   [check answers]

Sample Round No. 4
The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

30 out of 100 | **30 out of 100**   **10 out of 100** | **10 out of 100**

I thought the chance that this game board would reveal a winning box in this round was

very low   ○ ○ ○ ○ ○ ○ ○ ○ ○   very high

1. Refer to the top portion of the screen. If the yellow game board has ten winning boxes, how many winning boxes does the blue game board have?

- Ten. The blue and yellow game boards have the same winning boxes
- Thirty. The blue and yellow game boards have different winning boxes
- Ten or Thirty. The blue and yellow game boards may have the same or different winning boxes

INCORRECT! Refer to page 3 of the instructions.

I AM PLAYING THE YELLOW GAME BOARD

1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

2. Refer to the grid of 100 boxes on the right. Why is box number 75 shaded green on the yellow game board?

- box number 75 is a losing box on the yellow game board
- box number 75 is a winning box on the yellow game board
- box number 75 was played in a previous round

INCORRECT! Refer to page 4 of the instructions.

3. Why are box numbers 5, 10 and 30 shaded grey on the yellow game board?

- boxes 5, 10 and 30 have a value of £0 on the yellow game board
- boxes 5, 10 and 30 have a value of £20 on the yellow game board
- the values of boxes 5, 10 and 30 were revealed in previous rounds

INCORRECT! Refer to page 4 of the instructions.

4. Which statement is TRUE?

- The blue and yellow game boards have the same winning boxes.
- The blue and yellow game boards have different winning boxes
- The blue and yellow game boards may have the same or different winning boxes

INCORRECT! Refer to page 3 of the instructions.

Please click on the button labelled next to try again.   [next]

The winning boxes depend on the envelope drawn earlier.
The following sets of winning boxes are equally likely.

| 30 out of 100 | 30 out of 100 | | 10 out of 100 | 10 out of 100 |

I thought the chance that this game board would reveal a winning box in this round was

very low ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ ⊙ very high

1. Refer to the top portion of the screen. If the yellow game board has ten winning boxes, how many winning boxes does the blue game board have?

- ⊙ Ten. The blue and yellow game boards have the same winning boxes
- ○ Thirty. The blue and yellow game boards have different winning boxes
- ○ Ten or Thirty. The blue and yellow game boards may have the same or different winning boxes

CORRECT!

2. Refer to the grid of 100 boxes on the right. Why is box number 75 shaded green on the yellow game board?

- ○ box number 75 is a losing box on the yellow game board
- ⊙ box number 75 is a winning box on the yellow game board
- ○ box number 75 was played in a previous round

CORRECT!

3. Why are box numbers 5, 10 and 30 shaded grey on the yellow game board?

- ○ boxes 5, 10 and 30 have a value of £0 on the yellow game board
- ○ boxes 5, 10 and 30 have a value of £20 on the yellow game board
- ⊙ the values of boxes 5, 10 and 30 were revealed in previous rounds

CORRECT!

4. Which statement is TRUE?

- ⊙ The blue and yellow game boards have the same winning boxes.
- ○ The blue and yellow game boards have different winning boxes
- ○ The blue and yellow game boards may have the same or different winning boxes

CORRECT!

I AM PLAYING THE YELLOW GAME BOARD

| 1	11	21	31	41	51	61	71	81	91
2	12	22	32	42	52	62	72	82	92
3	13	23	33	43	53	63	73	83	93
4	14	24	34	44	54	64	74	84	94
5	15	25	35	45	55	65	75	85	95
6	16	26	36	46	56	66	76	86	96
7	17	27	37	47	57	67	77	87	97
8	18	28	38	48	58	68	78	88	98
9	19	29	39	49	59	69	79	89	99
10	20	30	40	50	60	70	80	90	100

Please click on the button labelled next to conclude the quiz.

[next]

You have completed the quiz.

Please wait while the other participants complete the questions.

Appendix 3.1: Distribution of Trading Gains and Losses

Figure 3.B: Distribution of realised gains and losses per transaction
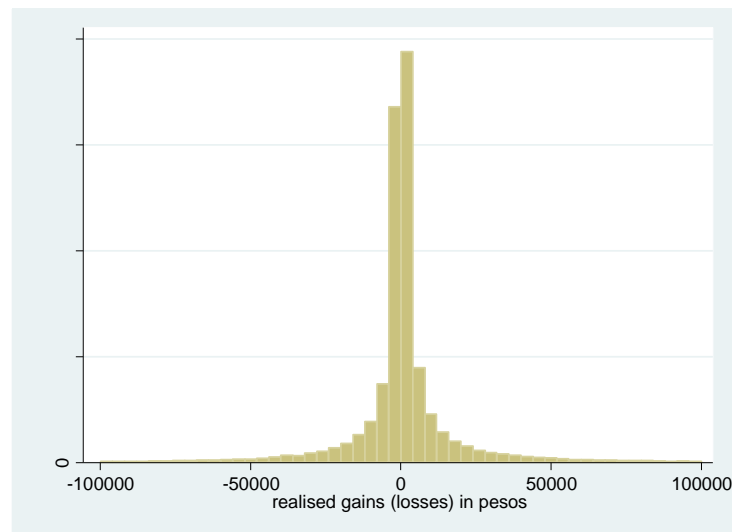
48% of buy-and-sell transactions resulted in realised gains, while 52% incurred realised losses. Average earnings amounted to Php3,816.40 (approximately £54.52); this is equivalent to 1.9% of the mean transaction amount (Php197,743). Meanwhile, median earnings equal to Php208.51 (approximately £2.98), equivalent to 0.3% of the median transaction amount (Php71,985).

Figure 3.C: Distribution of unrealised gains and losses per transaction
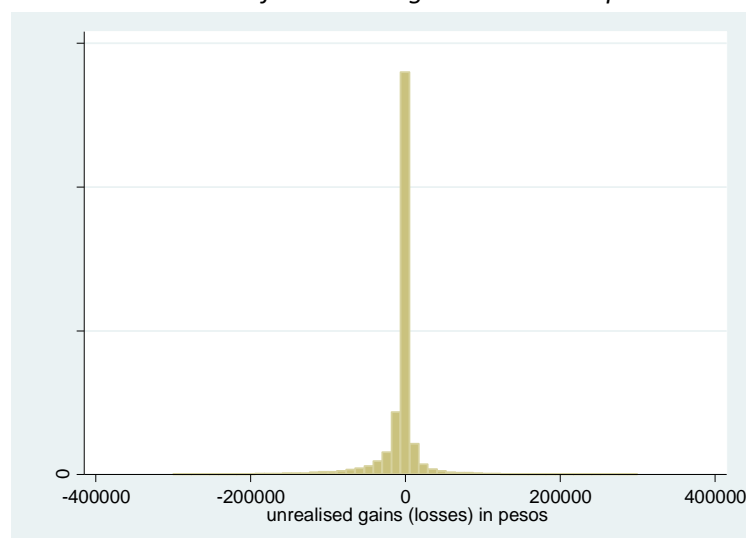
30% of buy-and-hold transactions resulted in unrealised gains, while 70% resulted in unrealised losses. Mean paper losses amounted to -Php6743.50 (approximately -£96.34) equivalent to 3.4% of the average transaction amount. Median paper losses were -Php458.57 (-£6.55) equivalent to 0.6% of the median transaction amount.

Appendix 3.2: Broker Recommendation

Forty-four percent of investors' buy transactions are on stocks covered by the broker's research. The chi-square test results below, which compare the observed vs expected transaction counts, confirm that retail investors tend to purchase recommended stocks, and veer away from securities not recently recommended for purchase.

Table 3.A: Stock purchases by broker recommendation

action	recommended	not recommended	Pearson chi2, p-value
buy	17,464	1,701	7.8e+03, 0.000
not buy	4,657,273	110,207,182	6.1e+07, 0.000

How quickly do investors act on the broker's recommendation? Figure 3.A summarises the number of buy transactions with a buy-recommendation, disaggregated by the number of days following/preceding the release of the broker's report. Purchases after the release of the recommendation (0-14 calendar days) reflect how quickly investors act on the buy-recommendation, while purchases before the report's release indicate investors' use of other sources of information.

Figure 3.A: Purchase transactions by recommendation date
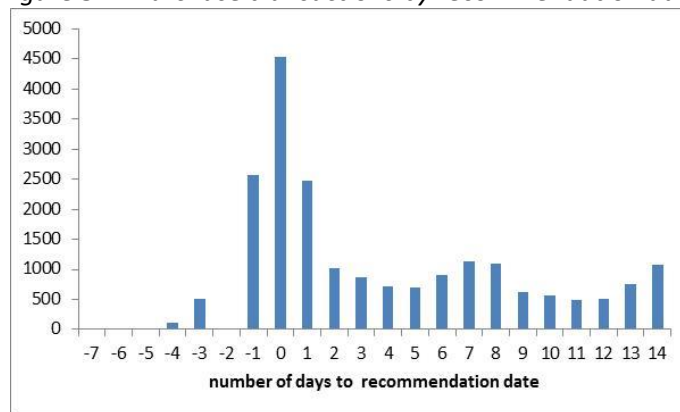
*Broker's recommendation report is dated on day 0.

Among purchases preceded by a buy-recommendation, 26% occurred on the day of the report's release date, while 15% were made one day after the recommendation. Surprisingly, there were about the same number of buy transactions one day before the recommendation date. It is likely that the broker's analysts have openly shared their "stock picks" with the media or during investor briefing events prior to the official release of the recommendation report.

Appendix 3.3: Historical Philippine Stock Exchange Index (2007-2010)

Source: Bloomberg

Appendix 3.4: Client interface

PSEi 7,180.34	ALL	FIN	IND	HLDG	PRO	SVC	M-O	Trades	Value	Market	Need Help?
0.27% 19.07	0.35%	0.73%	0.39%	-0.05%	0.66%	0.35%	1.23%	41,936	6,643,593,385	Closed...	Click Here

Home	Quotes	Trade	Research	StreetSmart	Sep. 17 5:36:16 AM	Log-Off
Enter Order	View/Modify Order	Trading History	Portfolio	O-H Order	O-H View/Cancel	EIP Scheduler

ACCOUNT SUMMARY *Dividends, deposits & withdrawal postings may be viewed through your Monthly Account Ledger statements-click link below*

Cash Balance

Actual Balance	Buying Power
8,632.92	8,632.92

- Print Account Summary
- Print Monthly Account Ledger
- View IPO Request Status
- View Tender Offer
- View Rights Offer

- View Monthly Account Ledger (current)
- View Monthly Account Ledger (historical)
- Activate I Deactivate E-Delivery of Confirmation
- Transaction Invoice Daily I Historical

Stock Position

Stock Code: [   ] [Search]

Action	Stock Code	Stock Name	Portfolio %	Market Price	Average Price	Total Shares	Uncommitted Shares	Market Value	Gain / Loss	%Gain/ Loss
BUY I SELL	BEL	BELLE CORPORATI	37.75	5.3400	5.7168	3,500	3,500	18,541.41	-1,467.39	-7.33%
BUY I SELL	EDC	ENERGY DEVELOPM	58.41	7.1400	5.5231	4,050	4,050	28,687.11	6,318.55	28.25%
BUY I SELL	MEGW1	MEGAWORLD CORP.	3.84	3.2000	0.8244	600	600	1,887.71	1,393.07	281.63%
				TOTAL TRADE VALUE:				49,116.24		
				DAY CHANGE:		0.00%			0.00	
				GAIN/LOSS:		14.56%			6,244.24	

Your Total Account Equity Value is 57,749.16

*Market Value, Day Change, Gain/Loss and Total Account Equity are NET values.

Stock Related Dividend and News
BEL
● Latest Dividend: YEAR: YEAR STOCK: STOCK CASH: CASH SPLITS: SPLITS RIGHTS: RIGHTS EX-DATE: EX-DATE RECORD DATE: RECORD PAYABLE: PAYABLE
● 09/15/2014: C04898: Belle Corp. clarification of news reports (PSE)
● 09/05/2014: C04795: Belle Corp. reply to exchange's query (PSE)
EDC
● Latest Dividend: YEAR: 2014 CASH: 0.10 EX-DATE: 12-Mar-2014 RECORD DATE: 17-Mar-2014 PAYABLE: 10-Apr-2014
● 09/10/2014: C04846: EDC press release on environmental management ISO certified - 1 (PSE)

Appendix 3.5: PSG/PSL and PDG/PDL subsample comparison

The following charts compare the influence of past trading outcome on a similar stock after controlling for market conditions or selected investor characteristics. In most of the subsamples, we find a significantly stronger tendency for investors to purchase shares of a similar stock given a history of either realised or unrealised gains on a similar stock. These support the results presented in Chapter 3.

I.   Figures 3.D and 3.E show the number of investors with a higher or lower PSG/PSL ratio (relative to PDG/PDL) across subsamples, given a history of gains/losses on the same stock. Not included in the bar charts are clients who are indifferent between similar and dissimilar stocks (i.e., PSG/PSL ≈ PDG/PDL).

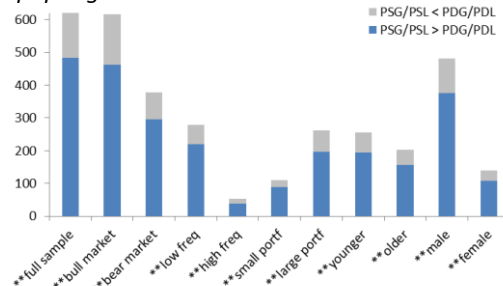*Figure 3.D: Repurchase of same stock given realised gain or loss*
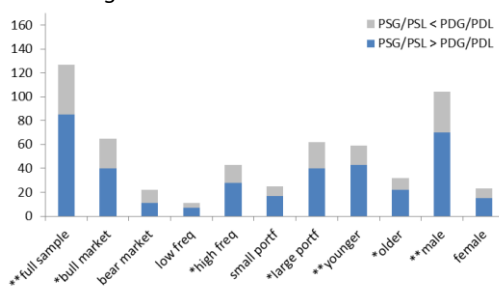
Figure 3.E: Repurchase of same stock given paper gain or loss

*p<0.05; **p<0.01

Figure 3.D shows that given a realised trading gain on a stock, there are relatively more investors who are likely to purchase that same stock than a different stock. This holds across the subsamples except in a bear market[1] ($z=0.467$, $p=0.6405$, $n=30$), among low frequency investors, i.e., those with total transactions of 208 or less ($z=0.813$, $p=0.4160$, $n=21$), and small portfolio account holders, i.e., clients with an average portfolio size of Php100,000 [approx. £1400] or less ($z=1.518$, $p=0.1289$, $n=35$). Meanwhile, Figure 3.E shows that given a paper gain on a stock, there are relatively more investors who are likely to purchase additional shares of that same stock than shares of another stock. This pattern holds across all the

---

[1] Using 20 percent change in the stock market index's previous peak or trough as basis for a change in market condition, the identified bull market in our dataset is from 1 January until 9 October 2007; the bear market is from 10 October 2007 until 10 October 2008, and the next bull begins on 11 October 2008.

subsamples, supporting our previous result which suggests that retail investors tend to accumulate shares of a stock given past successes on that same stock.

II. Given a history of gains/losses on a same-sector stock, we also compared individual-level PSG/PSL and PDG/PDL ratios in selected subsamples. The bar charts in Figures 3.F and 3.G show the number of investors with a higher or lower PSG/PSL ratio vis-a-vis PDG/PDL. In Figure 3.F, while there is evidence that past realised outcomes in same-sector stocks are useful to investors, this observation does not hold across all subsamples, particularly during a bear market, among low frequency traders, clients with large portfolios, older investors and female clients. Meanwhile, Figure 3.G shows more compelling results. Evidence suggests that investors apply their unrealised trading gains in decisions to purchase a same-sector stock.

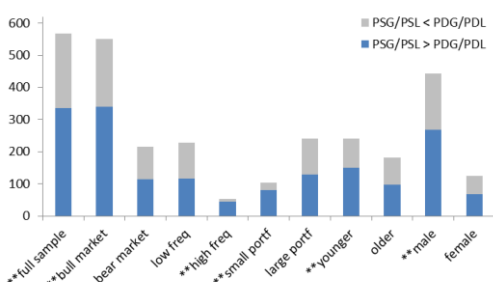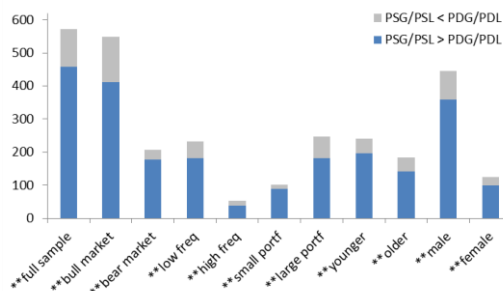*Figure 3.F: Purchase of same-sector stock given realised gain or loss*

Figure 3.G: Purchase of same-sector stock given paper gain or loss

*p<0.05; **p<0.01

III. Given a history of gains/losses on a broker-recommended stock, we compared individual-level PSG/PSL and PDG/PDL ratios in selected subsamples. These are summarised in Figures 3.H and 3.I.

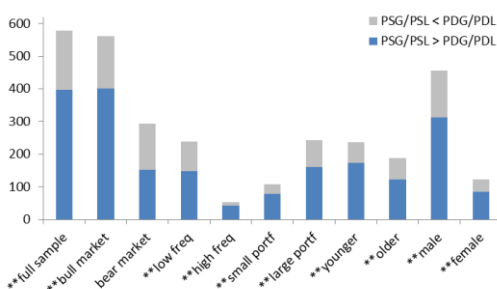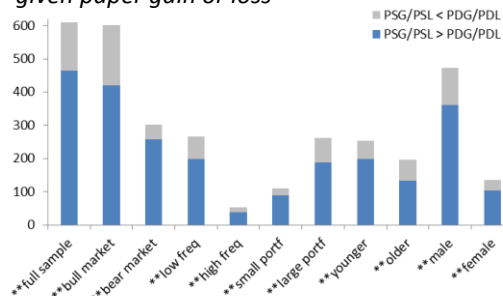*Figure 3.H: Purchase of recommended stock given realised gain or loss*

Figure 3.I: Purchase of recommended stock given paper gain or loss

*p<0.05; **p<0.01

Across the subsamples, results suggest that past experience on broker-recommended stocks, whether realised or unrealised, are applied to other broker-recommended stocks, which confirms investors' use of broker recommendation as basis for categorising stocks (recommended or not recommended), and that success on a broker-recommended stock spills over to other similar stocks.