ATENEO DE MANILA UNIVERSITY
Department of Economics
ATENEO CENTER FOR ECONOMIC RESEARCH AND DEVELOPMENT

Bandit with similarity information

BY

Benjamin Radoc

Working Paper No.2020-02

March 31, 2020

ADMU Econ-ACERD working papers are preliminary versions of research papers that have not been published in a journal or vetted through a peer-review process. The purpose of circulation is to elicit comments and suggestions. The contents or opinions expressed in the paper are the author(s)' sole responsibility and do not necessarily reflect the views of Ateneo de Manila University and ACERD. Any errors of commission or omission are the author(s)' and should not be attributed to any of the above.

Rm. 409 4/F, Dr. Ricardo & Rosita Leong Hall,
Ateneo de Manila University
Loyola Heights, Quezon City 1108, Philippines
Telephone: (632) 8426 6001 Local 5221 or 5222
Telefax: (632) 8426 5661

Bandit with similarity information*

Benjamin Radoc†
Philippine Competition Commission and Ateneo de Manila University

March 2020

Abstract

In problem situations, it is unlikely that a decision maker possesses knowledge on all states of the world so that use of information gathered from past similar experience is plausible. This analogical thinking is formalised by case-based decision theory (CBDT) which suggests that under uncertainty, a decision maker acts based on her memory of past actions and the associated outcomes in past similar situations. A unique experimental setting that makes similarity information in the problem situation salient was created, while providing a fair chance for either Bayesian or case-based decisions to emerge. Consider a two-armed bandit with similarity information. If the similarity cue offers irrelevant information on the payoff distribution of the two arms, it is easy for a Bayesian decision maker to ignore the cue and recognise that the distribution of payoffs of the two arms are identical. But if the similarity cue triggers a decision maker to perceive the two arms as separate, there are two possibilities: (i) more frequent positive payoffs on one arm may be used to put a higher valuation on a similar arm consistent with the prediction of CBDT, or (ii) given past positive payoffs on one arm, the other arm may be valued lower if the subsequent likelihood of success is perceived as lower. The experiment results suggest that although participants in general correctly updated their expectation of a positive payoff based on past experience, the pattern in the decisions shows that participants systematically used the irrelevant similarity cue in a manner that cannot be explained by CBDT but consistent with the gambler's fallacy or the biased belief that the pattern of past outcomes will reverse.

Keywords: Bayesian reasoning, case-based decisions, gambler's fallacy, expectation, similarity

JEL Classifications: C91, D9, D81

*The author is indebted to Robert Sugden and Theodore L. Turocy for their guidance; to the University of East Anglia's Centre for Behavioural and Experimental Social Science (CBESS) for financial support; and to Jiwei Zheng, Emanuela Lezzi, Bahar Ghezelayagh, Natalia Borzino, Allie Mcguire, and Mike Brock for helping run the experiment sessions. All errors are the responsibility of the author.

†Corresponding author: `bradoc@ateneo.edu`

1 Introduction

In a bandit problem, a player encounters a machine that generates random payoffs as she simultaneously attempts to optimise earnings and gain additional information for each action taken (Robbins (1952)), similar to a slot machine where a player decides to pull any one of k levers (or *arms*) in a sequence of trials. For each action, the player receives information useful in updating the expected distribution of payoffs of that arm. After a series of plays, the player decides either to learn more about the distribution of payoffs offered by the other arms (called *exploration*) or to keep choosing the same arm presumed to give the best payoff (called *exploitation*). The challenge of balancing exploration and exploitation in real-life contexts contributed to the proliferation of research on solving different variants of the bandit problem. For example, a contextual bandit is a type of a multi-armed bandit that offers side information to help a player infer the rewards of the other unobserved arms. In a contextual bandit, similarity in the context for different arms is useful in predicting the rewards for arms that are perceived to be similar (Auer (2002)).

In this study, a ”ticket experiment” was conducted akin to a bandit where participants encountered a salient similarity cue that may or may not be perceived as relevant. By design, the experiment allows a comparison of the predictive power of Bayesian reasoning under objective uncertainty and case-based decision theory (CBDT) as proposed by Gilboa and Schmeidler (1995).

In problem situations, it is unlikely that a decision maker possesses knowledge on all states of the world so that use of information gathered from similar experiences in the past is plausible. If analogical reasoning is applied, a decision maker may match the surface and structural features 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. Similarity assessment therefore allows the transfer of knowledge acquired from past experience to a present problem (Gregan-Paxton and Cote (2000), Zizzo (2003)).

Case-based decision theory formalises analogical reasoning and imposes minimal cognitive demands on a decision maker. In CBDT, an 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 (which may also include hypothetical or counterfactual cases) and evaluates the similarity between the new problem and past problems encountered. Given a new problem, the decision maker will act based on her memory and will choose an act in past similar situations with the highest similarity-weighted sum of outcomes. This means that if a past problem is perceived as similar to a new problem, a decision maker will adopt the same action taken in the past. In the ticket experiment, an environment was created where it is possible for participants to learn the probabilities from experience. This is a departure from the ”structural ignorance” decision environment under CBDT where outcomes and probabilities are not given nor can be easily derived from experience (Gilboa and Schmeidler (2001). In this experiment, we test whether CBDT is also predictive of decisions under uncertainty.

Central to decision making under CBDT is a *similarity function* that is learned from experience by the decision maker. Since the concept of similarity is derived from preferences, it could be unique to each individual and the similarity function may also change as a decision maker accumulates more relevant experiences (Gilboa and Schmeidler (2001)). The situation of an unspecified similarity function makes it especially difficult to test CBDT. Also, Matsui (2000) showed that it is difficult to behaviourally distinguish case-based decisions 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, EUT and CBDT yield equivalent behavioural predictions.

It is therefore not surprising that most experimental tests on case-based decision theory did not compare the predictive power of CBDT and EUT. Such experiments include: (i) an ambiguous monopoly production decision setting that compared the prevalence of case-based decision making with the *max-heuristic* or choosing the action with the highest historical profit (Grosskopf et al. (2015); (ii) an environment with "structural ignorance" where neither outcomes nor probabilities are known. In this experiment, Ossadnik et al. (2013) pitted CBDT against several decision criteria such as *maximin* (choosing the act with the highest minimum payoff), *maximax* (choosing the act with the single highest outcome), *pessimism-optimism* (choosing the act with the maximum weighted value of the lowest and highest outcomes), and *reinforcement learning model* (choosing the act with the highest propensity for selection as a function of the frequency of successful outcomes); and (iii) an experiment where participants chose between two real-estate investments based on induced cases of real-estate properties. Payoffs were based on the actual market price appreciation of the selected property one month after the experiment (Bleichrodt et al. (2013)).

In these past experiments, uncertainty in the outcome space was induced and participants were prodded to pay attention to the similarity across the decision settings. In this study, we take a less suggestive approach on similarity, thereby allowing a fair opportunity for participants to either use or ignore similarity information. In the experiment, we imposed feature-based similarity (Tversky (1977)) where two objects are considered similar if a salient feature that is common to the objects matched. This strategy of treating similarity as a binary variable (i.e., two problems are either identical or completely different) overcomes the problem of specifying the form of the similarity function.

In the *ticket experiment*, participants encountered two arms (blue ticket and yellow ticket) with payoffs determined by random draws (with replacement) from a single mechanical randomiser. The possible outcomes (either GBP 20 or GBP 0) are known, but not the distribution of outcomes.

Participants may infer the distribution of the outcomes from personal experience on a limited but equal number of sample rounds on either arm. Given the experimental environment, information learned about one arm also applies directly on the other arm so the situation can be reduced to a one-armed bandit. In this situation, the contextual cue (colour) offers irrelevant information on the distribution of the payoffs so that it is easy for a Bayesian decision maker to ignore the similarity information and recognise that the distribution of payoffs of the two arms are identical. Alternatively, given the salience of the similarity cue a decision maker may perceive the two arms as having separate and distinct payoff distributions. If the experiment setup is perceived as a two-armed bandit, the distribution of valuations of the two arms may not be identical.

Among participants who use the similarity information as a relevant cue, there are two possibilities conditional on observed outcomes: (i) more frequent GBP 20 draws (or ”hits”) on one ticket colour may be used to put a higher valuation on a coloured ticket relative to the other ticket; or (ii) given past hits on one ticket colour, the other ticket colour may be perceived to have a higher likelihood of a hit to even out the payoffs of the two coloured tickets.

The first scenario is consistent with reinforcement learning where behavior associated with a known pleasurable outcome is repeated while an action causing pain or regret is avoided (Strahilevitz et al. (2011)). When a decision is required in a situation where outcomes are ambiguous or cannot be easily constructed by the decision maker (such as the unknown distribution of payoffs in the mechanical randomiser used in the experiment), CBDT predicts that a decision maker will act based on her memory of past actions and the associated outcomes in past similar situations. Applying analogical reasoning, if an action taken in one problem has a history of actual successes, that action is likely to be repeated when a *similar* problem situation is encountered.

Meanwhile, the second scenario is a manifestation of the classic gambler’s fallacy or the erroneous belief that an arm which had a series of losses is bound to ”self-correct” or reverse the pattern of past outcomes (Rogers (1998)). This well-documented misperception results from people’s susceptibility to derive sequential patterns from experience and offer deterministic explanations for events despite the objective uncertainty of outcomes (Sun and Wang (2010)). Outside experimental settings, evidence on the gambler’s fallacy has been observed in refugee asylum court decisions, loan application reviews, and Major League Baseball umpire pitch calls (Chen et al. (2016)).

In past studies, evidence of the gambler’s fallacy has been largely inferred from actual decisions. Without a direct measure of decision makers’ beliefs, we can only infer the prevalence of the gambler’s fallacy and reasons for its occurrence (Farmer et al. (2017)). In the ticket experiment, at the start of each round participants indicated on a Likert scale their expectation of a chance that a GBP 20 ball will be drawn in that round. Since the experiment was designed as a game show, step-by-step instructions were announced by the experimenter so that participants had sufficient time to complete each task in a synchronised manner. The filling in of the Likert scale was not incentivised but the pace of each round provided participants the chance to carefully consider and indicate on the coloured ticket their belief on the likelihood of a hit on that round. Presumably, the revealed expectation takes into account the past hits and misses observed in the past rounds. The design of the *ticket experiment* is parallel to the design of a *game board experiment* (Radoc et al. (2019)) which was conducted to validate the *act separability axiom* under CBDT. The 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 during decision making. If applied to asset selection and portfolio diversification, the axiom implies that an investor would fail to account for the correlation between available investment instruments and other assets already in the portfolio. In the between-subjects experiment, participants encountered two coloured game boards: a blue game board and a yellow game board. A high type board (H) contained thirty pre-drawn winning boxes while a low type board (L) had ten pre-drawn winning boxes. The game board experiment had three treatments. In the independent treatment, the winning boxes in the two game boards were pre-drawn independently. In the negative correlation treatment, one board was always of type H and the other of type L, while in the positive correlation treatment, either both boards were type H, or both boards were type L.

However, the two experiments differed in two ways. In the ticket experiment the likelihood of a hit was unknown to the participants. In the game board experiment, participants had information on the objective probability of a hit: 10 percent for type L and 30 percent for type H. Also, given random draws with replacement, the uncertainty of a hit is more salient in the ticket experiment. In each round, whether the coloured ticket is associated with a hit (GBP 20) or a miss (GBP 0) depends on the live draw of a ball from a mechanical randomiser. Neither the experimenter nor the participants know in advance whether a hit or a miss will be observed. In the game board experiment the participants are informed that the winning boxes in the type L and type H boards are pre-drawn so that in each round, randomness was on which box will be opened and not on the value of the box.

In both experiments, an observed hit in a preceding round is associated with a higher *expectation* of a hit in the current round whether the same colour or a different colour is played. This suggests that participants used the experience of a hit in a manner consistent with Bayesian thinking. However, results in the *valuation task* differed significantly between the two experiments despite the similarity in the design.

In the ticket experiment, there is evidence that participants did not ignore colour despite its irrelevance. Valuations attached to blue and yellow tickets diverged especially among participants who observed very few hits in the sample rounds. Specifically, valuations attached to the coloured tickets were the opposite of CBDT's prediction: a ticket colour with fewer hits was valued more highly than a ticket colour with more hits. Although the pattern in revealed expectations is supportive of Bayesian thinking (i.e., more hits observed is associated with a higher expectation of a hit), the subsequent pattern in ticket valuations are reminiscent of the gambler's fallacy. This was not the case in the game board experiment. We conjecture that the ticket experiment's setup (live random draws with replacement from a mechanical randomiser) and the randomness in the payoffs (with no a priori information provided to the participants) account for the pattern in the valuations. Sun and Wang (2010) suggested that human habit to find sequential patterns is consistent with the tendency to offer deterministic explanations despite the uncertainty of outcomes.

The remainder of the paper is organised as follows: Section 2 describes the experimental design, Section 3 presents the results, and Section 4 concludes.

2 Experimental design

We created a unique experimental setting that makes the similarity cue salient while providing a fair chance for either Bayesian or case-based decisions to emerge. Participants played with coloured tickets (blue or yellow) which paid earnings based on random draws from a single mechanical randomiser (a white bingo cage) containing an unknown distribution of GBP 20 (hit) and GBP 0 (miss) balls. Colour was made a highly salient environmental cue to trigger conscious similarity assessment and facilitate recall of past events. However, colour was not linked in any way to the ticket earnings and was clearly uninformative on the probability of a GBP 20 draw. Ignoring colour was therefore an easy strategy for a Bayesian player.

Meanwhile, given the uncertainty in the ticket earnings and colour as the only similarity cue, participants have the option to code and retrieve memories of past rounds based on colour. Evidence supportive of CBDT are participants' systematic use of colour in the experiment and if higher ticket valuations are in the same direction as the past payoffs. The between-subjects paper-and-pencil ticket experiment was set up akin to a game show. The use of mechanical logistics provided an engaging task for participants and emphasised both 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 to create 10 cases in memory.[1] In the experiment, there was an equal but randomly ordered blue and yellow rounds depending on the set of tickets used for that round. 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.

[1]Since case-based decisions rely on memory of past cases participants were not allowed to take down notes during the sample rounds. Also, the number of sample rounds in our design is close to the median stopping rule of participants in past learning experiments.

Hertwig et al. (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, while *decisions from experience* underweight rare outcomes but improve Bayesian reasoning. In our experiment, since participants sampled from experience (i.e., earnings were sequentially revealed), the environment imposed a more stringent decision setting for case-based decisions to emerge.

All sessions were conducted with the aid of a game show assistant. To determine the earnings from the ticket at each sample round, the assistant drew one ball with replacement from the bingo cage containing 100 white balls marked either GBP 20 or GBP 0. Participants knew that the bingo cage contained 100 balls, but not the distribution of the balls.[2] With a 20% probability that a GBP 20 ball is drawn at each round, a hit was both a rare and salient outcome. All information in both parts of the experiment was common knowledge among participants in each session.

At the start of each round, participants indicated on the *sample ticket* their expectation of the chance that a GBP 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 the ball drawn and the ball was shown to the participants. After each draw, participants filled in the earnings portion on their sample ticket and then dropped the coloured ticket in an opaque box of the same colour as the ticket.

Participants knew 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 GBP 2. They also knew that the tasks in Part 1 provided information about the distribution of GBP 20 and GBP 0 balls in the bingo cage. The synchronised completion of each task following verbal cues from the experimenter and the use of similarly coloured stimuli at each round kept participants engaged throughout the experiment.

Part 2 of the experiment consisted of a task that elicited participants' valuation of a coloured ticket following a BDM mechanism (Becker et al. (1964)). At the start of Part 2, each participant randomly drew one sealed brown envelope containing either a blue or yellow ticket (similar to Part 1) and a blue or yellow decision form from a bag. The decision form listed 35 possible offer prices ranging from 20 pence to GBP 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 their decision forms, 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 et al. (2011)), the instructions included an outright statement that the participants' answers on the decision form cannot influence the actual offer price.

[2]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 (GBP 0 and GBP 20). The balls were returned in the bingo cage before Round 1 so participants knew that the bingo cage contained at least one GBP 20 ball.

The actual offer price was revealed only after all participants submitted their decision forms. If a participant decided to keep her ticket at the offer price, her earnings equal the outcome of her individual draw. Otherwise, she was paid the offer price. Whichever the decision, each participant came forward for an individual draw which was conducted in the same manner as the sample rounds. The individual draws clearly 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 expected utility maximisers (Keller et al. (1993)), it was unlikely that participants' knowledge of a forthcoming public announcement of their decision influenced their preference for a ticket colour at the various offer prices.

3 Results

Thirty sessions with an average of six participants per session were conducted at the University of East Anglia's Centre for Behavioural and Experimental Social Science laboratory. Participants in a session experienced the same sample rounds, but encountered only one ticket colour (either blue or yellow) in the valuation task. Of the 176 participants, 51% are female, 56% are British, and 41% recently engaged in some form of gambling activity. Average individual earnings was GBP 8.50, ranging from GBP 2 to GBP 22.

In the 10 sample rounds, participants observed either a GBP 20 draw (hit) or GBP 0 draw (0) which allowed them to learn the distribution of payoffs. The sample rounds were followed by a valuation task where each participant owned one coloured ticked and a decision form showing a list of offer prices. At each offer price, a participant decided whether she preferred to keep her ticket or to exchange her ticket for money.

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.[3] On average, switching point (interchangeably used here with *ticket valuation*) is higher among participants who observed more hits in the sample rounds.

Notice that the mean switching point exceeds the expected payoff of GBP 4; this observation is not unusual. In past studies with valuation tasks eliciting participants' willingness to accept (WTA) an amount of money in exchange for an item they own, WTA is typically higher than the expected value. Given uncertainty about the true value of a good, exchange aversion (Sugden (2003)) and the consequence of giving up that good and foregoing the opportunity to learn more about it (Zhao and King (2001)) have been offered as possible explanations for the pattern in WTA amounts. In preference reversal experiments, decision makers' tendency to overvalue a gamble

[3] There were nine participants who selected several switching points. For these participants, the median switching point was used in the data analysis.

Table 1: Switching point by number of hits

Number of hits observed	Mean switching point	SD	N
0	3.828	2.771	18
1	4.102	2.676	59
2	4.810	2.903	40
3-4	6.532	3.039	59
All sessions	5.049	3.045	176

with low probability of winning a large amount is also commonly observed (Seidl (2002)).

Result 1. Participants colour-coded events in the experiment despite the irrelevance of colour.

When faced with uncertainty, CBDT predicts that a decision maker will act based on her memory of past similar problems. If participants use colour as a similarity cue, experience in blue rounds may be perceived as different from the experience in yellow rounds. In the experiment, colour was salient but uninformative so that a Bayesian player is expected to ignore it. Given only one mechanical randomiser that is visible to the participants throughout the experiment, a hit on a blue ticket (B) is also a hit on a yellow ticket (Y), and vice versa.

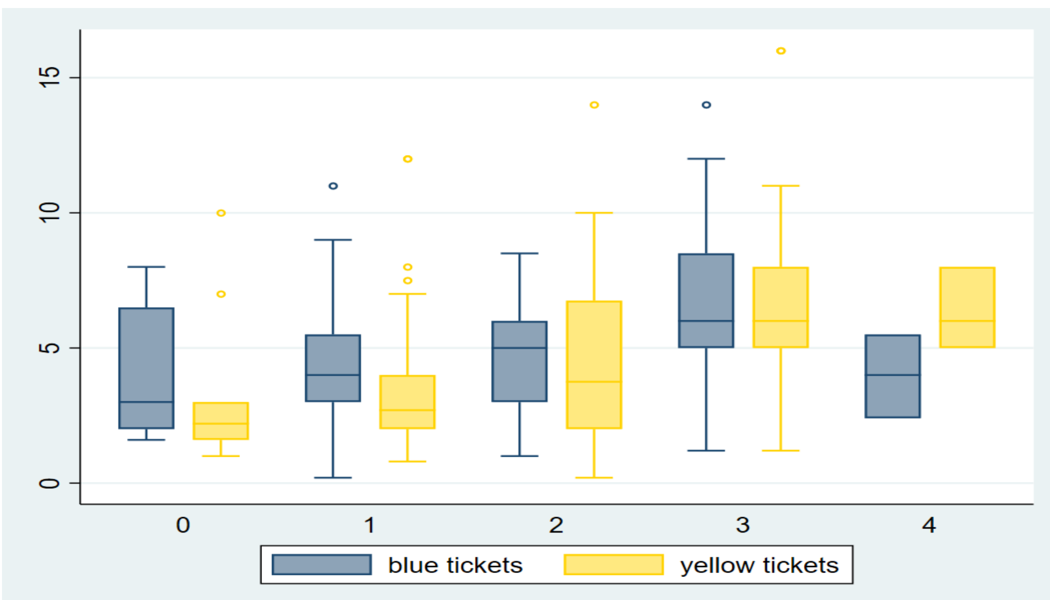

Figure 1: Distribution of switching points by colour and total number of hits observed in the sample rounds.

In Figure 1, the median switching points suggest a difference in WTA between B and Y. Disregarding the number of hits observed in the sample rounds, we find that average switching point for blue tickets (GBP 5.18, n=88) is higher than the switching point for yellow tickets (GBP 4.92, n=88). The observed difference in valuation, although not statistically significant (z=1.389, p=0.1647, n=176), captures both the variation in the frequency of hits in the coloured rounds[4] and a colour effect (where participants considered blue rounds and yellow rounds as dissimilar experiences).

To validate a pure colour effect, we compared the switching points of participants conditional on the number of hits observed in the sample rounds. Figure 2 shows the distribution of ticket valuations conditional on hits. Switching points in B and Y are not significantly different in sessions with the same number of hits on B and Y. The ranksum test results confirm that such is the case in sessions where one hit each occurred in the blue sample rounds and in the yellow sample rounds (mean on B=4.719, mean on Y=4.588, z=0.462, p=0.6443, N=28), and in sessions where no hit is observed in any of the sample rounds (mean on B=5.427, mean on Y=5.003, z=0.842, p=0.3996, N=18).

In sessions where only one hit is experienced in the sample rounds, a case-based decision maker is likely to perceive colour as a salient cue especially in the round associated with a hit. Figure 2 indicates a difference in the valuation between blue and yellow tickets. Based on the ranksum test results, the observed divergence is statistically significant when B=0 and Y=1 (mean on B=4.515, mean on Y=3.624, z=1.768, p=0.0771, N=41), but not when B=1 and Y=0 (mean on B=4.211, mean on Y=4.189, z= 0.713, p=0.476, N=18).

Although ticket colour as a similarity cue was selectively used by the participants, this result is not necessarily inconsistent with CBDT since we do not have information on the similarity function used by the participants during the experiment.[5] While every effort was made to control all salient stimuli in each round during the experiment, there was no certainty as to which information participants actually used (or did not use) in the elicitation task.

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. If a participant's favourite color is the same as the ticket colour played in the valuation task, the colour effect may be more significant than for players valuing a ticket that is different from their favorite colour. Comparing the switching point of participants in sessions with an equal number of hits on B and on Y, the rank-sum test results reveal no significant difference in valuation (z=0.233, p=0.8157, N=52) whether the ticket colour played is a participant's favourite colour (N=13) or not (N=39).

The timing of a successful draw was also found not to have a significant influence on ticket valuation. We compared the switching point of participants in selected sessions with one hit each

[4]In Part 1 of the experiment, 60% of GBP 20 draws occurred in a yellow round and 40% of hits in a blue round.

[5]Thanks to Chris Starmer for pointing this out.

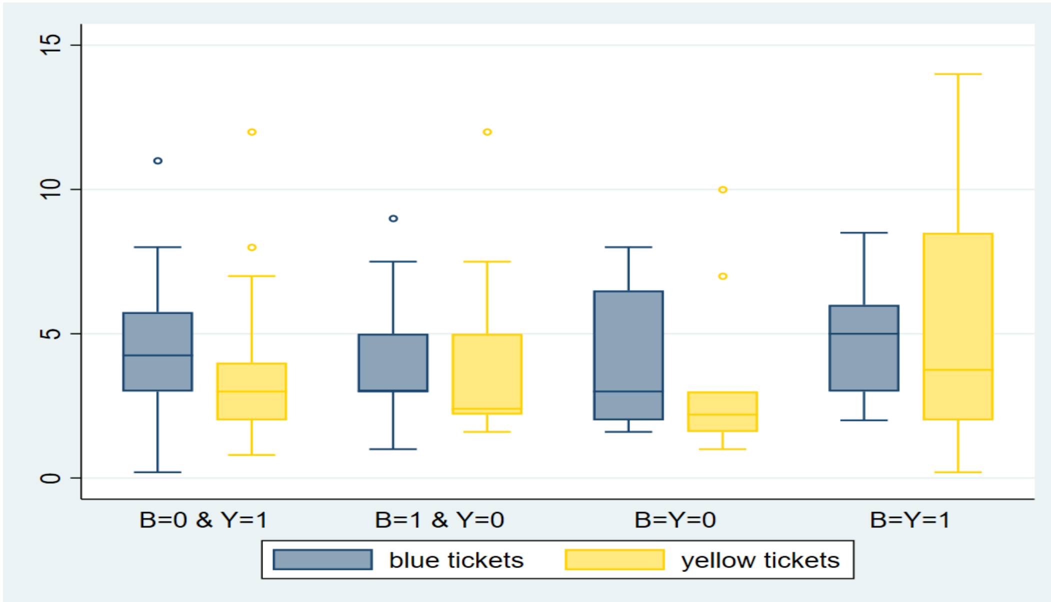

Figure 2: Distribution of switching points by colour conditional on the difference in the number of hits observed in the sample rounds.

on B and Y. In four of these sessions, the hit on B occurred on a latter sample blue round (8th to 10th sample round) while the hit on Y was experienced in a middle yellow round (4th to the 7th sample round). If recency in the observed hits matters, the valuation on blue tickets may be significantly higher than for yellow tickets. However, there is no indication of a recency effect (z=1.058, p=0.2901, N=22).

Result 2. Participants systematically put a *higher value* on a coloured ticket with *fewer hits* than a ticket with more hits.

The ticket experiment was designed to validate whether participants choose an act with the highest-similarity weighted outcome. This translates to assigning a higher value on a ticket colour with more hits relative to a ticket colour associated with fewer hits.

In the analysis, tickets are categorised as either ”leading” or ”lagging” depending on the relative number of hits observed in the sample rounds for each ticket colour. A leading (lagging) ticket colour in the valuation task had more (fewer) hits in the sample rounds compared to the other ticket colour. Since both ticket colour and the relative number of hits were used in categorising the tickets, a difference in the valuation between leading and lagging tickets also implies colour-coding.

As confirmed by the ranksum test results (z=2.131, p=0.033; n=124), participants systematically valued a lagging ticket (mean=5.774) more highly than a leading ticket (mean=4.656). This pattern in the ticket valuations is surprising. Rather than having results that are either Bayesian or

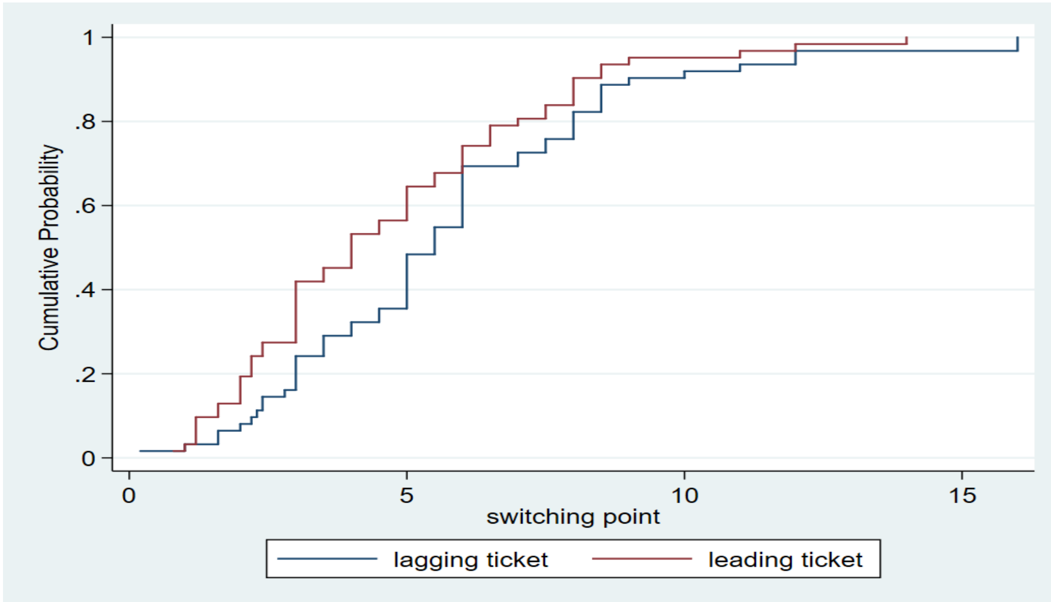

Figure 3: Distribution of switching points between "leading" and "lagging" tickets.

case-based, the pattern in the ticket valuations is 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 as if it has some sort of memory and moral sense (Rogers (1998)). In the context of our experiment, this implies that if a lagging ticket colour is *believed* to reverse its poor performance in the valuation task, it is likely to be valued more highly than a leading ticket colour which is unlikely to sustain its previous hit/s.

Human cognitive constraints have been offered as explanations to the gambler's fallacy. Experiencing a long sequence of events necessarily provides more information than a shorter sequence. However, limited attention and short-term memory constrain decision makers' perceptions about the randomness of events (Warren and Hahn (2009). This means that even if a long series of information shows an equal chance between two outcomes, events may be perceived in terms of "chunks" rather than in its entirety (Farmer et al. (2017)).

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 when information is presented sequentially (Hogarth and Einhorn (1992); Ayton and Fischer (2004)). While the main research objective is to compare the predictive power of CBDT against Bayesian reasoning, the decision setting in the ticket experiment inadvertently led to the emergence of the gambler's fallacy.

Note that the gambler's fallacy observed in the ticket experiment was *not* observed in a parallel game board experiment (Radoc et al. (2019)). The *game board experiment* was designed to validate the *act separability axiom* proposed under CBDT. The axiom suggests that if decision makers code experiences based on a similarity function, then it is implied that people maintain separate memories of alternative actions taken in the past. If applied to portfolio allocation, this suggests that an investor would ignore the correlation between available assets. In the game board experiment, participants encountered two coloured game boards (blue and yellow), which were either the high type (H) or the low type (L). Participants knew that H contained 30 winning boxes, L had 10 winning boxes, and that the winning boxes have been pre-drawn. There were three treatments. In the *independent treatment* the winning boxes in the two game boards were pre-drawn independently. In the *negative correlation treatment*, if one board is of type H and the other is necessarily of type L, while in the *positive correlation treatment*, either both game boards were type H, or both game boards were type L. The reader will notice that the positive correlation treatment in the game board experiment is similar to the ticket experiment but with important differences.

In the ticket experiment, no *a priori* on the probability of a hit was provided, while in the game board experiment participants had information on the objective probability of a hit: 10 percent for type L and 30 percent for type H. In the ticket experiment, a hit in each round was determined based on a live random draw of a ball with replacement. In the game board experiment, since participants knew that the winning boxes have been pre-drawn, randomness was on which box will be opened and not on the value of the box. The conjecture is that the difference in the uncertainty of a hit in the ticket experiment vis-a-vis the game board experiment accounts for the divergence in the pattern of valuation results.

Result 3. Ticket valuations are biased towards the gambler's fallacy despite evidence of Bayesian beliefs.

Given the pattern in the valuation of lagging and leading tickets, it is reasonable to suspect that the participants' revealed expectations are consistent with such valuations.

In past studies, evidence of the gambler's fallacy has been largely inferred from actual decisions. Without a direct measure of decision makers' beliefs, we can only infer the prevalence of the gambler's fallacy and reasons for its occurrence (Farmer et al. (2017)).

In the ticket experiment, at the start of each round participants indicated on a Likert scale their expectation of a chance that a GBP 20 ball will be drawn in that round. Since the experiment was designed as a game show, step-by-step instructions were announced by the experimenter so that participants had sufficient time to complete each task in a synchronised manner. The filling in of the Likert scale was not incentivised but the pace of each round provided participants the chance to carefully consider and indicate on the coloured ticket their belief on the likelihood of a hit on that round. Presumably, the revealed expectation takes into account the hits and misses observed in the past rounds.

The summary statistics in Table 2 suggest that generally an observed hit in a preceding round

Table 2: Summary of current expectation of a hit on a round

total hits	expectation of a hit conditional on cumulative hits observed			
	median (in sample rounds)	mean (in sample rounds)	median (in last round)	mean (in last round)
0	3	3.7408	2	2.8333
1	3	3.5873	2	2.6610
2	4	4.4206	3	3.3750
3	4	4.0735	3	4.0566
4	5	4.7500	5	5.1667

is associated with a higher *expectation* of a hit in the current round whether the same colour or a different colour is played. Although it appears that the relationship between an incremental hit and a change in expectation is nonlinear, the relationship is non-negative. If the biased belief consistent with the gambler's fallacy is supported, we would expect a negative relationship between a past incremental hit and a lower expectation in a new and different coloured round.

We also test if the expectation of a GBP 20 draw in the valuation task is different in sessions with only one hit observed in the sample rounds. For these sessions, the ranksum test results indicate that there is no significant difference in the expectation of a hit between blue and yellow tickets (z=-0.769, p=0.442, N=124) and between leading and lagging tickets (z=0.424, p=0.671, N=59).

These results suggest that participants used the experience of a hit reminiscent of Bayesian thinking. It is therefore surprising for valuations to be biased despite non-biased beliefs.

4 Concluding remarks

The ticket experiment conducted for this study was designed to test the predictive power of CBDT vis-a-vis Bayesian reasoning, and focused on validating two predictions of CBDT, namely: i) decision makers encode and retrieve past experiences using similarity information; and ii) decision makers choose an act with the highest similarity-weighted outcome. The decision setting induced objective uncertainty which provided a fair chance for either case-based or Bayesian reasoning to emerge.

The design of the ticket experiment is parallel to the positive treatment of a game board experiment (Radoc et al. (2019)) conducted to validate CBDT's separability axiom. In both experiments, a past incremental hit on either colour is associated with a higher expectation on the chance of a hit in a new round. However, the results in similar valuation tasks do not coincide. The conjecture is that the difference in the uncertainty of a hit in the ticket experiment vis-a-vis the game board experiment accounts for the divergence in the pattern of valuation results.

In the game board experiment, there was no evidence of colour-coding events. The pattern in the results of the valuation task is reminiscent of Bayesian thinking: a hit on one colour was associated with a higher valuation on either colour. In the ticket experiment, participants used colour selectively and valuations were the opposite of CBDT's similarity-weighted prediction: tickets with fewer hits were valued more highly than tickets with more hits. Despite CBDT's intuitive appeal, the results imply that case-based decisions may be sensitive to the features of the decision setting.

References

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.

Peter Ayton and Ilan Fischer. The hot hand fallacy and the gambler's fallacy: Two facets of subjective randomness. *Memory & Cognition*, 32(8):1369–1378, 2004.

Gordon Becker, Morris DeGroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3):226–232, 1964.

Han Bleichrodt, Martin Filko, Amit Kothiyal, and Peter Wakker. Making case-based decision theory directly observable. Working paper, 2013.

Daniel Chen, Tobias Moskowitz, and Kelly Shue. Decision making under the gambler's fallacy: evidence from asylum judges, loan officers, and baseball umpires. *The Quarterly Journal of Economics*, 131(3): 1181–1242, 2016.

George D. Farmer, Paul A. Warren, and Ulrike Hahn. Who "believes" in the gambler's fallacy and why? *Journal of Experimental Psychology: General*, 146(1):63–76, 2017.

Itzhak Gilboa and David Schmeidler. Case-based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639, 1995.

Itzhak Gilboa and David Schmeidler. *A Theory of Case-Based Decisions*. Cambridge University Press, 2001.

Cleotilde Gonzalez and Varun Dutt. Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118(4):523–551, 2011.

Jennifer Gregan-Paxton and Jane Cote. How do investors make predictions? insights from analogical reasoning research. *Journal of Behavioral Decision Making*, 13:307–327, 2000.

Brit Grosskopf, Rajiv Sarin, and Elizabeth Watson. An experiment on case-based decision making. *Theory and Decision*, 79(4):639–666, 2015.

Ralph Hertwig, Greg Barron, Elke Weber, and Ido Erev. Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8):534–539, 2004.

Robin M. Hogarth and Hillel J. Einhorn. Order effects in belief updating: The belief adjustment model. *Cognitive Psychology*, 24:1–55, 1992.

Andrea Isoni, Graham Loomes, and Robert Sugden. 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, 2011.

L. Robin Keller, Uzi Segal, and Tan Wang. The becker-degroot-marschak mechanism and generalized utility theories: Theoretical predictions and empirical observations. *Theory and Decision*, 34:83–97, 1993.

Akihiko Matsui. Expected utility and case-based reasoning. *Mathematical Social Sciences*, 39:1–12, 2000.

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

Charles Plott and Kathryn Zeiler. 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, 2005.

Benjamin Radoc, Robert Sugden, and Theodore L. Turocy. Correlation neglect and case-based decisions. *Journal of Risk and Uncertainty*, 59:23–49, 2019.

Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.

Paul Rogers. The cognitive psychology of lottery gambling: A theoretical review. *Journal of Gambling Studies*, 14(2):111–133, 1998.

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

Michal Ann Strahilevitz, Terrance Odean, and Brad M. Barber. Once burned, twice shy: How naive learning, counterfactuals, and regret affect the repurchase of stocks previously sold. *Journal of Marketing Research*, XLVIII:S102–S120, 2011.

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

Yanlong Sun and Hongbin Wang. Gambler's fallacy, hot hand belief, and the time of patterns. *Judgment and Decision Making*, 5(2):124–132, 2010.

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

Paul A. Warren and Ulrike Hahn. Perceptions of randomness: Why three heads aer better than four. *Psychological Review*, 116:454–561, 2009.

Jinhua Zhao and Catherine L. King. A new explanation for the wtp/wta disparity. *Economics Letters*, 73: 293–300, 2001.

Daniel John Zizzo. Transfer of knowledge in economic decision-making: An overview. *Greek Economic Review*, 22(2):1–10, 2003.

5 Appendix

5.1 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 familiarise 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?

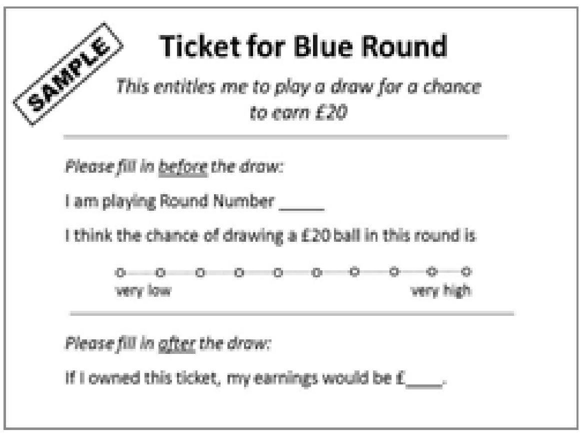

SAMPLE

Ticket for Blue Round

This entitles me to play a draw for a chance to earn £20

Please fill in before the draw:

I am playing Round Number _____

I think the chance of drawing a £20 ball in this round is

very low … very high

Please fill in after the draw:

If I owned this ticket, my earnings would be £____.

SAMPLE

Ticket for Yellow Round

This entitles me to play a draw for a chance to earn £20

Please fill in before the draw:

I am playing Round Number _____

I think the chance of drawing a £20 ball in this round is

very low … very high

Please fill in after the draw:

If I owned this ticket, my earnings would be £____.

5.2 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.

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?