



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

Benjamin Radoc^{1,2} 

Received: 6 November 2020 / Revised: 10 March 2024 / Accepted: 15 July 2024 /
Published online: 6 August 2024

© The Author(s), under exclusive licence to Economic Science Association 2024

Abstract

Decision makers typically possess limited knowledge on states of the world so that use of information from past similar experiences is reasonable. This analogical thinking is formalised by case-based decision theory (CBDT). We created a novel experimental setting to validate the predictive power of CBDT versus Bayesian reasoning. Participants encountered a salient but irrelevant cue which a Bayesian decision maker is likely to ignore but a case-based decision maker may use in assessing similarity. We find that although the irrelevant similarity cue was used, the pattern in participants' decisions is neither case-based nor Bayesian. The results suggest that CBDT does not apply in simple decision settings where similarity cues are uninformative.

Keywords Bayesian reasoning · Case-based decision · Gambler's fallacy · Similarity

JEL Classification C91 · D9 · D81

1 Introduction

In decision situations, it is unlikely that a decision maker possesses complete knowledge on all states of the world so that use of information gathered from similar experiences in the past is reasonable. In analogical reasoning, a decision maker may match the features between a base situation and the target problem (Gregan-Paxton & Cote, 2000). If the target problem is perceived as sufficiently similar to the base situation, 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 (Zizzo, 2003).

✉ Benjamin Radoc
bradoc@ateneo.edu

¹ Department of Economics, Ateneo de Manila University, Quezon City, Philippines

² BSP Research Academy, Bangko Sentral ng Pilipinas, Manila, Philippines

Case-based decision theory (CBDT), proposed by Gilboa and Schmeidler (1995), formalises analogical reasoning. Unlike expected utility theory (EUT) where a decision maker is assumed to know all possible outcomes and their associated probabilities, CBDT imposes minimal cognitive demands on a decision maker. *Case-based decisions* are made conditional on one's memory of past actions and the assessment of the similarity between problems.

Central to decision making under CBDT is similarity assessment. Since the concept of similarity is derived from preferences, an individual's similarity function could be unique and may change as she accumulates more relevant experiences (Gilboa & Schmeidler, 2001). However, the similarity function is unspecified in CBDT which makes it difficult to test the predictive power of the theory. Note that there is no explicit assumption that the similarity function is formulated by a decision maker based on the actual or real relevance of stimuli. Even if decision makers are faced with the same situation, individual similarity functions can vary.

Also, Matsui (2000) showed that it is not easy to behaviourally distinguish case-based decisions from decisions consistent with 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.

As an attempt to empirically distinguish CBDT from EUT decisions, we created an experimental setting using colour as a salient similarity cue and allowed a fair chance for either case-based decisions or Bayesian thinking to emerge. In the paper-and-pen experiment, participants played coloured tickets (blue or yellow) which paid earnings based on live random draws (with replacement) from a single mechanical randomiser. The bingo cage, which was visible to all participants throughout the experiment, contained an unknown distribution of white balls marked with either £20 (a hit) or £0 (a miss). Although events during the experiment were colour-coded, colour was not linked in any way to the ticket earnings and was clearly uninformative on the probability of a £20 draw. This means that ignoring colour was an easy strategy for a Bayesian player.

Conditional on the total number of past hits observed, we find that valuations on the coloured tickets appear to be Bayesian (i.e., more hits observed in the past is associated with a higher ticket valuation). Although the irrelevant similarity cue was used, the pattern in participants' decisions is neither case-based nor Bayesian. The results suggest that CBDT does not apply in simple decision settings where similarity cues are uninformative.

When the number of hits on the blue ticket and the yellow ticket are different, a ticket colour with fewer hits was valued more highly than a ticket colour with more hits. This result cannot be explained by either CBDT or Bayesian thinking but is reminiscent of the *gambler's fallacy* or the biased belief that a lottery which had a series of losses was bound to reverse the pattern of its past outcomes (Rogers, 1998). Our conjecture is that the experiment's setup (live random draws with replacement from a mechanical randomiser) and the randomness in the payoffs (with no a priori information given to the participants) triggered a colour-effect leading to the observed pattern in the ticket valuations.

The rest of the paper is organized as follows: Section 2 presents the features of case-based decision theory, Section 3 describes the experimental design, Section 4 states the hypotheses, Section 5 discusses the results, and Section 6 provides further discussion.

2 Case-based decision making

In CBDT, an experience is encoded into memory as a case with three elements: problem p , act a , and result r . When a decision maker faces a new problem q , she scans her memory M for cases encountered in the past and evaluates similarity between q and 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 . 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:

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

Gilboa and Schmeidler (1995) proposed that 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. For experiments specific to act separability, see Bleichrodt et al. (2017) and Radoc et al. (2019).

In decision making, all that CBDT requires is the agent's ability to recall past cases and to evaluate similarity between a new problem and past problems encountered. Since only experienced cases are in memory, CBDT accommodates the possibility for 'structural 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 where the case-based prediction converges with the expected utility from the act (Gilboa & Schmeidler, 2010; Sugden, 2004).

Empirical work with field data testing the validity of CBDT's predictions is limited and there are even fewer experiments. Grosskopf et al. (2015) compared the predictive power of CBDT versus the max-heuristic (i.e., choosing the action with the highest payoff) in a monopoly choice problem. When feedback on the monopolist's profit is delayed, CBDT is a better predictor than the max-heuristic. When immediate feedback is available, CBDT and max-heuristic are equally likely. Ossadnik et al. (2013) induced an environment with structural ignorance and pitted CBDT against four decision criteria: (i) maximin (choosing the option with the highest minimum payoff); (ii) maximax (choosing the option with the single highest outcome); (iii) pessimism-optimism (choosing the option with the maximum weighted value of the lowest and highest outcomes); and (iv) reinforcement learning model (choosing the

option with the highest propensity for selection as a function of the frequency of successful outcomes). They find that CBDT explains participants' choices better than the four criteria. However, case-based decision makers earned less than participants guided by maximax, maximin and pessimism-optimism.

Bleichrodt et al. (2017) conducted an experiment that required participants to choose among hypothetical real-estate investments in the Netherlands where objective probabilities on the price appreciation are unknown. They find that similarity information on past real-estate investments encountered is used as predicted by CBDT. However, the CBDT assumption on *act separability*, which suggests that different acts are stored as separate memories, is violated. Results show that informational value or memory of one property interacts with the memory of an unrelated property.

Although the concept of similarity is central to the decision-making process, CBDT does not provide details on the similarity function. This makes it especially difficult to implement the theory. To deal with this issue, we imposed feature-based similarity (Tversky, 1977) where two objects are considered similar if a salient feature common to the objects matches. This approach 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. Unlike in past CBDT experiments where uncertainty in the outcome space was induced and participants were prodded to pay attention to the similarity across the decision settings, we used a less suggestive approach on similarity to allow a fair opportunity for participants to either use or ignore the similarity cue.

3 Experimental design

We created a novel experimental setting using a salient similarity cue (colour) which participants could either use or ignore. In the between-subjects paper-and-pen experiment, participants encountered coloured tickets (blue or yellow) which paid earnings based on live random draws from a single mechanical randomiser (a white bingo cage). The randomiser containing an unknown distribution of £20 and £0 balls was positioned in front of the room, visible to all participants throughout the experiment.

In all rounds, colour was a highly salient cue to trigger conscious similarity assessment and facilitate recall of past rounds. In addition to coloured tickets and coloured boxes, the “blueness” or “yellowness” of a round was emphasised by the experimenter’s repeated announcement of the round played, and a coloured light bulb illuminating the bingo cage.¹ Since there was only one bingo cage, colour was not linked in any way to the ticket earnings and was clearly uninformative on the probability of a £20 draw (hit). Given the uncertainty in the ticket earnings

¹ A blue (yellow) light bulb directed towards the bingo cage was turned on at the start of every blue (yellow) round.

and colour as the only salient similarity cue, participants had the option to code and retrieve memories of past rounds based on colour as CBDT would predict.

The experiment consisted of two Parts. In Part 1, participants played five *sample* blue rounds and five *sample* yellow rounds to create ten cases in memory.² Unknown to the participants, there was a 20% chance for a hit to occur in any round. Since participants sampled from experience (i.e., earnings were sequentially revealed using live draws from only one bingo cage), the environment imposed a stringent decision setting for case-based decisions to emerge.³

The experiment was set up akin to a *game show* with mechanical logistics and conducted with the aid of a game show assistant. To determine the earnings from the ticket, the assistant drew one ball with replacement from the bingo cage. Participants knew that the bingo cage contained 100 balls, but not the distribution of the balls.⁴ All information in both parts of the experiment was common knowledge among participants in each session.

At the start of each sample round, the assistant randomly picked an envelope containing coloured tickets. After the tickets were distributed, participants indicated on the coloured ticket their expectation of the chance that a £20 ball will be drawn in that round. This task was not incentivised but the pace of each round provided participants the chance to carefully consider their belief on the likelihood of a hit on that round.

After everyone indicated their expectation of a £20 draw, the experimenter gave the signal for the game show assistant to draw a ball. The assistant announced and showed the value written on the ball. Participants then filled in the earnings portion on their ticket and dropped the coloured ticket in an opaque box of the same colour as the ticket. The participants observed the game show assistant return the ball in the bingo cage which had the same 100 balls in each round.

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). Participants knew that their total earnings at the end of the experiment (including a show-up fee of £2) depended only on the outcome of their decision in Part 2. They also knew that the tasks in Part 1 provided information about the distribution of £20 and £0 balls in the bingo cage which was the same randomiser used in Part 2.

² 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 the experiment was close to the median stopping rule of participants in past learning experiments.

³ 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, while *decisions from experience* underweight rare outcomes but improve Bayesian reasoning. In a meta-analysis, Wulff et al. (2018) showed that the *description-experience gap* is sensitive to the structure of the decision problem. The gap is reduced if the choice is between two risky options, but not if the choice is between a risky and a certain option. Also, the gap is larger when the outcome is rarer, especially in problems involving a risky and a certain option.

⁴ 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 sample balls were returned in the bingo cage before Round 1 so participants knew that the bingo cage contained at least one £20 ball.

In Part 2, each participant randomly drew one sealed brown envelope from a bag. An envelope contained either a blue or yellow decision form, and a corresponding coloured ticket. The decision form listed 35 possible offer prices ranging from 20 pence to £20. At each offer price, a participant decided whether she preferred to keep her ticket or to exchange her ticket for money.

Before participants filled in their 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 (Isoni et al., 2011; Plott & Zeiler, 2005), 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 equaled the outcome of her individual draw. Otherwise, she was paid the offer price. Whichever the decision was, each participant came forward for an individual draw which was conducted in the same manner as the sample rounds. The individual draws induced emotional reactions during the experiment; occasional clapping or sighing after each public individual draw was not uncommon. Because 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.

4 Hypotheses

If participants are Bayesian thinkers, the valuations of the blue and yellow tickets will not be statistically different. Since the similarity cue offered irrelevant information on the distribution of the payoffs, it was easy for a Bayesian decision maker to ignore colour and recognise that the payoff distribution of the two ticket colours were identical.

Given the salience of the similarity cue, a decision maker may perceive the two coloured tickets as having separate and distinct payoff distributions. Conditional on observed outcomes, there are two possibilities: (i) more frequent hits on one ticket colour may be used to put a higher valuation on that coloured ticket as CBDT would predict; or (ii) given past hits on one ticket colour, the other ticket colour may be perceived to have a higher likelihood of a hit consistent with the gambler's fallacy, so that a ticket with fewer hits is valued more highly than the other ticket.⁵

⁵ An anonymous referee pointed out that since the experiment was designed similar to a game show and participants showed emotional reactions during the sessions, it is possible that *regret aversion* may play a role in the participants' decisions. While this may explain why ticket valuations may be higher than the expected payoff, regret aversion does not explain why the valuation of blue and yellow tickets would diverge.

Table 1 Summary of switching point (in £) by participant characteristics

	Mean	Median	SD	<i>N</i>
Non-frequent bettor	4.81	4.50	2.95	118
Frequent bettor	5.54	5.00	3.20	58
From the UK	4.59	4.75	2.33	98
From EU	5.03	4.75	2.73	24
From elsewhere	5.89	5.25	4.05	54
Female	4.46	3.50	3.08	89
Male	5.65	5.00	2.90	87

5 Results

Thirty sessions with four to six participants each were conducted at the University of East Anglia's Centre for Behavioural and Experimental Social Science (CBESS) laboratory. Each session involved public draws which were identical for all participants in one session, but different across sessions. Although all participants in a session experienced the same sample rounds in Part 1, a participant encountered only one ticket colour (either blue or yellow) in Part 2. Of the 176 participants, 51% are female, 56% are British, and 41% recently engaged in some form of gambling activity. Average individual earnings were £8.50, ranging from £2 to £22.

We analysed the *switching point* or the offer price at which a participant changed her preference from keeping a coloured ticket to exchanging it for money.⁶ Table 1 summarises the switching points by participant characteristics, irrespective of the number of hits observed in Part 1. In the discussion below, we present Wilcoxon ranksum test for equality of distributions comparing two groups using rank data instead of raw £ values.

Switching point (interchangeably used here with *ticket valuation*) of participants who frequently engaged in betting activities (i.e., at least twice in the past 30 days) is not statistically different from that of the participants who seldom gambled ($z = -1.517$, $p = 0.1292$). Similarly, participants from the UK did not value their coloured tickets significantly lower than players from the rest of the European Union ($z = 0.517$, $p = 0.6051$), and players from elsewhere ($z = 1.480$, $p = 0.1390$). However, male participants had significantly higher switching points ($z = -3.082$, $p = 0.0021$) relative to their female counterparts. This supports the observed higher risk aversion among females found in past studies (Croson & Gneezy, 2009).

Table 2 compares switching points conditional on the total number of hits in Part 1. For either blue or yellow ticket, average ticket valuations are higher when participants observed more hits. The last column shows the calculated valuations for a risk-neutral Bayesian with a prior belief that there is a uniform distribution of the number of £20 balls. The switching points, conditional on the number of

⁶ There were nine participants who selected several switching points. For these participants, the median switching point was used in the data analysis.

Table 2 Switching point (in £) by number of observed hits in Part 1

Number of hits	N	All tickets		Blue ticket		Yellow ticket		Bayesian valuation
		Mean	Median	Mean	Median	Mean	Median	
0	18	3.83	2.40	4.26	3.00	3.40	2.20	1.67
1	59	4.10	3.00	4.42	4.00	3.79	2.70	3.33
2	40	4.81	4.75	4.88	5.00	4.74	3.75	5.00
3	53	6.69	6.00	6.66	6.00	6.72	6.00	6.67
4	6	5.15	5.25	3.97	4.00	6.33	6.00	8.33
All sessions	176	5.05	5.00	5.18	5.00	4.92	4.00	4.72

Table 3 Summary of switching point (in £) by hits and ticket colour

Condition	N	Blue ticket			Yellow ticket			Ranksum test	
		Mean	Median	N	Mean	Median	N	z	p-value
All sessions	176	5.18	5.00	88	4.92	4.00	88	1.39	0.1647
$B = Y$	52	4.59	4.75	26	4.72	3.00	26	0.59	0.5570
$B \neq Y$	124	5.43	5.00	62	5.00	4.25	62	0.14	0.1668
$B < Y$	76	5.42	5.00	38	4.16	3.25	38	2.34	0.0194
$B > Y$	48	5.45	5.00	24	6.34	6.00	24	-0.61	0.5420
$B + Y = 1$	59	4.42	4.00	29	3.79	2.70	30	1.85	0.0643
$B = 1, Y = 0$	18	4.21	3.00	9	4.19	2.40	9	0.71	0.4760
$B = 0, Y = 1$	41	4.52	4.25	20	3.62	3.00	21	1.77	0.0771

hits observed, are higher than the corresponding Bayesian valuations. Also, the mean switching point (£5.05) exceeds the expected payoff of £4; these observations are not unusual. In past studies using valuation tasks to elicit participants' willingness to accept (WTA) an amount of money in exchange for an item they own, WTA has been found to be typically higher than the expected value of the item. 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 & King, 2001) have been offered as explanations for the pattern in WTA amounts. In preference reversal experiments, decision makers' tendency to overvalue a gamble with low probability of winning a large amount is also commonly observed (Seidl, 2002).

CBDT predicts that a decision maker faced with uncertainty will act based on her memory of past similar problems. If colour is used as a similarity cue, experience in blue rounds may be perceived as a different experience from events in yellow rounds. Meanwhile, a Bayesian player is expected to ignore colour because it is uninformative about the likelihood of a hit. Given the experiment's setup, it is apparent that a hit on a blue ticket (B) is also a hit on a yellow ticket (Y), and vice

versa. Table 3 provides a snapshot of ticket valuations disaggregated by the relative number of hits on B and Y .

Unconditional on the number of hits observed in Part 1 (all sessions), we find that the average switching point for the blue ticket is not significantly higher than the switching point for the yellow ticket. In sessions with the same number of hits on B and Y , switching points are also not significantly different.⁷ However, when the number of hits on B and Y differed, there is a significant divergence in ticket valuations. Specifically, the difference is statistically significant only for the $B < Y$ case. This suggests that participants colour-coded events in the experiment despite the irrelevance of colour.

We recognise that ticket valuations capture both the variation in the frequency of hits 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 point of participants conditional on the number of hits observed in Part 1. Consider sessions where $B + Y = 1$. When only one hit was observed, that sample round and its colour is likely to be salient to a case-based decision maker. As shown in Table 3, valuations on B are higher than the valuations on Y but the difference is 'marginally' significant. We separate $B + Y = 1$ into $B = 1, Y = 0$ and $B = 0, Y = 1$ in Table 3. When $B = 1, Y = 0$, B and Y valuations are not statistically different, but this is not the case when $B = 0, Y = 1$. It appears that the marginal difference observed when $B + Y = 1$ is driven by a colour yellow effect. A similar colour effect is observed when $B < Y$ but not when $B > Y$.

If participants choose an act with the highest-similarity weighted outcome as predicted by CBDT, they will tend to assign a higher value on a ticket colour with more hits. We categorised tickets as either "leading" or "lagging" depending on the relative number of hits observed for each ticket colour in Part 1. A leading (lagging) ticket colour in the valuation task has more (fewer) hits in the sample rounds compared to the other ticket colour. Since both ticket colour and the relative number of hits are used in categorising the tickets, a difference in the valuation between leading and lagging tickets also implies colour-coding. Wilcoxon ranksum test results indicate that participants' valuation on a lagging ticket (mean = 5.774) is significantly higher ($z = 2.131, p = 0.033$) than a leading ticket (mean = 4.656) (Fig. 1).

The pattern in switching points is neither Bayesian nor case-based but depicts the gambler's fallacy. This 'error' 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, or when information is presented sequentially (Ayton & Fischer, 2004) similar to our experimental setup. In the context of the experiment, the result implies that a lagging ticket colour is believed to reverse its poor performance and is, therefore, valued more highly versus a leading ticket colour which is unlikely to sustain its previous hit(s). In addition to the gambler's

⁷ To determine if the recency of a hit had an effect on ticket valuation, we used sessions where $B = Y$ and compared switching points conditional on the colour of the last sample round with a hit and the ticket colour played in Part 2. The Wilcoxon ranksum test result does not indicate a statistically significant recency effect ($z = -1.625, p = 0.1041$).

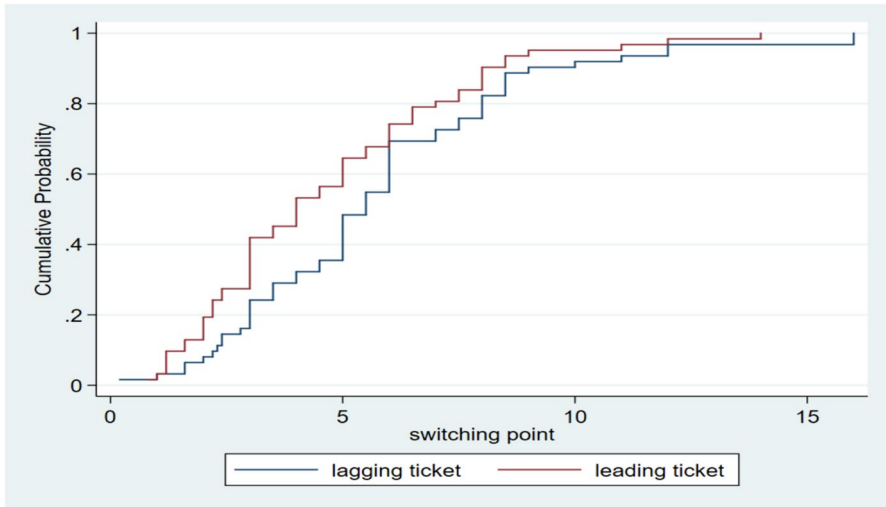


Fig. 1 Distribution of switching points between “leading” and “lagging” tickets

fallacy, our results indicate a colour effect suggesting a difference in participants’ perception of yellow and blue tickets.

6 Discussion

The *ticket experiment* described in this paper was designed to test the predictive power of CBDT versus Bayesian reasoning and focused on validating two predictions of CBDT, namely: (i) decision makers encode and retrieve past experiences using similarity cues; 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 manifest.

The experimental design is similar to another study (*game board experiment*) 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. In the game board experiment (Radoc et al., 2019), participants encountered two coloured boards: a blue board and a yellow board. A coloured board is either high type (H) containing 30 winning boxes or low type (L) with 10 winning boxes. Participants knew that the winning boxes in either board type were pre-drawn. The ticket experiment is similar to the positive correlation treatment of the game board experiment where the two coloured boards were either type H or type L.

Although the ticket experiment is parallel to the positive correlation treatment of the game board experiment, the results do not coincide. The pattern in valuations in the positive correlation treatment of the game board experiment is qualitatively Bayesian. Meanwhile, the pattern in valuations in the ticket experiment is neither

case-based nor Bayesian. The irrelevant similarity cue (colour) is used but in a manner that is consistent with the gambler's fallacy. The results suggest that CBDT does not apply in simple decision settings where similarity cues are uninformative like in the ticket experiment.

Our conjecture on why the gambler's fallacy emerged is that the difference in the uncertainty of a hit in the ticket experiment versus the game board experiment accounts for the divergence in the valuation pattern. 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% for type L and 30% for type H.

Also, given random draws with replacement, the uncertainty of a hit is arguably more salient in the ticket experiment. In each round, whether the coloured ticket is associated with a hit or a miss depends on the live draw of a ball from a mechanical randomiser. Neither the experimenter nor the participants knew in advance whether a hit or a miss will be observed. In the game board experiment, the participants were informed that the winning boxes in the type L and type H boards were pre-drawn so that in each round, randomness was on which box will be opened and not on the value of the box.

Arguably, 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. These suggest that despite CBDT's intuitive appeal, decisions that emerge may be sensitive to the features of the decision setting.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s40881-024-00177-3>.

Acknowledgements I am indebted to Robert Sugden and Theodore L. Turocy for their guidance, and to Dirk Engelmann and anonymous reviewers for their comments and suggestions. I also thank Alasdair Brown, Arvid Hoffmann, and participants of the CBESS-CeDEX-CREED Meeting for their feedback on early versions of this paper. Mike Brock, Natalia Borzino, Bahar Ghezelayagh, Emanuela Lezzi, Allie Mcguire, Melanie Parravano Baro and Jiwei Zheng helped run the experiment sessions, and Mario Rosete III provided invaluable research assistance.

Funding This research was funded by the University of East Anglia's Centre for Behavioural and Experimental Social Science.

Data availability The data supporting the findings of this study are available in the OSF, accessible at <https://osf.io/ku69d>.

Declarations

Conflict of interest Not applicable.

References

- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two facets of subjective randomness. *Memory & Cognition*, 32(8), 1369–1378.
- Becker, G., DeGroot, M., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3), 226–232.

- Bleichrodt, H., Filko, M., Kothiyal, A., & Wakker, P. (2017). Making case-based decision theory directly observable. *American Economic Journal: Microeconomics*, 9(1), 123–151.
- Crosron, R., & Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, 47(2), 448–474.
- Gilboa, I., & Schmeidler, D. (2001). *A Theory of Case-Based Decisions*, Cambridge University Press.
- Gilboa, I., & Schmeidler, D. (2010). *Case-Based Prediction*, World Scientific Publishers.
- Gilboa, I., & Schmeidler, D. (1995). Case-based decision theory. *The Quarterly Journal of Economics*, 110(3), 605–639.
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118(4), 523–551.
- Gregan-Paxton, J., & Cote, J. (2000). How do investors make predictions? Insights from analogical reasoning research. *Journal of Behavioral Decision Making*, 13, 307–327.
- Grosskopf, B., Sarin, R., & Watson, E. (2015). An experiment on case-based decision making. *Theory and Decision*, 79(4), 639–666.
- Hertwig, R., Barron, G., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539.
- Isoni, A., Loomes, G., & Sugden, R. (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.
- Keller, L. R., Segal, U., & Wang, T. (1993). The Becker-DeGroot-Marschak mechanism and generalized utility theories: Theoretical predictions and empirical observations. *Theory and Decision*, 34, 83–97.
- Matsui, A. (2000). Expected utility and case-based reasoning. *Mathematical Social Sciences*, 39, 1–12.
- Ossadnik, W., Wilmsmann, D., & Niemann, B. (2013). Experimental evidence on case-based decision theory. *Theory and Decision*, 75, 211–232.
- Plott, C., & Zeiler, K. (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.
- Radoc, B., Sugden, R., & Turocy, T. L. (2019). Correlation neglect and case-based decisions. *Journal of Risk and Uncertainty*, 59, 23–49.
- Rogers, P. (1998). The cognitive psychology of lottery gambling: A theoretical review. *Journal of Gambling Studies*, 14(2), 111–133.
- Seidl, C. (2002). Preference reversal. *Journal of Economic Surveys*, 16(5), 621–655.
- Sugden, R. (2003). Reference-dependent subjective expected utility. *Journal of Economic Theory*, 111, 172–191.
- Sugden, R. (2004). *Handbook of Utility Theory Volume 2: Extensions*, Kluwer Academic, chapter Alternatives to Expected Utility: Foundations and Concepts.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327–352.
- Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, 144, 140–176.
- Zhao, J., & King, C. L. (2001). A new explanation for the WTP/WTA disparity. *Economics Letters*, 73, 293–300.
- Zizzo, D. J. (2003). Transfer of knowledge in economic decision-making: An overview. *Greek Economic Review*, 22(2), 1–10.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.