Artif Intell Rev (2007) 27:209–222
DOI 10.1007/s10462-008-9083-4

Computational modelling of switching behaviour in repeated gambles

Jiaying Zhao · Fintan J. Costello

Published online: 13 September 2008

Abstract We present a computational model which predicts people's switching behaviour in repeated gambling scenarios such as the Iowa Gambling Task. This Utility-Caution model suggests that people's tendency to switch away from an option is due to a utility factor which reflects the probability and the amount of losses experienced compared to gains, and a caution factor which describes the number of choices made consecutively in that option. Using a novel next-choice-prediction method, the Utility-Caution model was tested using two sets of data on the performance of participants in the Iowa Gambling Task. The model produced significantly more accurate predictions of people's choices than the previous Bayesian expected-utility model and expectancy-valence model.

Keywords Switching · Repeated gambles · Iowa gambling task · Emotion

1 Introduction

Theories of economic and behavioural decision making have traditionally assumed that humans are fundamentally rational creatures. For example, expected utility (EU) theory (Bernoulli 1967) assumes that decisions are made by assessing the EU and the weighting of these utilities of each outcome, and that the option with the highest EU will be chosen. Following the same logic, rational choice theory assumes that the decision maker will always choose the most preferred alternative (e.g. Mas-Collel et al. 1995). However, in the past 20 years research has consistently shown the limits of human rationality and systematic errors in decision making processes (Simon 1982; Kahneman et al. 1982).

Most studies on decision making have typically used single-choice scenarios. For example, the decision maker has to choose either winning \$450 *for sure or a* 0.5 *chance of winning* \$1,000 (Kahneman and Tversky 1979). However, in real life people often make a number

J. Zhao · F. J. Costello (✉)
School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland
e-mail: fintan.costello@ucd.ie

of sequentially related choices. When making repeated choices people tend not to stick with one option but to switch among available choices (Coombs and Meyer 1969; Keren and Wagenaar 1987) For example, Ratner and Herbst (2005) find that people switch away from options they believe are most profitable due to emotional responses to the possible negative outcomes of those options. Investigating such switching behaviour can offer a more in-depth view of the decision process involved in repeated choices.

In this paper we attempt to clarify the following questions: what causes people to switch among choices? what factors are involved in such switching? and how rational is people's switching behaviour? We address these questions by fitting a computational model to people's switching behaviour in a repeated gambling task. The Iowa Gambling Task (IGT) designed by Bechara et al. (1994) is used for the current investigation, as firstly the task has been well established in the decision literature and secondly it provides a paradigm for investigating switching behaviour in repeated gambles. The focus of this paper is to show how (healthy) people switch among their choices in repeated gambles, and is not however to provide an optimal gambling strategy.

2 The Iowa gambling task

The IGT paradigm is designed in a way that resembles real-world contingencies (Bechara et al. 1994). The task consists of four decks of cards (labelled as A, B, C, and D), each containing 40 cards. Every card in the four decks gives an immediate and definite win but some of the cards also generate an unpredictable loss. The IGT represents risks in real-word problems in terms of contingencies of wins and losses, and the gambler is told to choose a number of cards and to maximise his/her final reward.

More specifically, in deck A every card gives a definite \$100 win but 50% of the cards give an average \$250 loss. This means that the average net payoff per card in deck A is a \$25 loss ($0.5*100 - 0.5*(250 - 100)$). In deck B, every card gives a definite \$100 win but there are 10% of the cards, each of which also gives a \$1250 loss. This means that the average net payoff per card in deck B is a \$25 loss ($0.9*100 - 0.1*(1250 - 100)$). In other words, decks A and B are disadvantageous as choosing decks A and B would eventually lead to an overall loss in the long run. In deck C, every card gives a definite \$50 win and there are 50% of the cards, each of which also gives an average \$50 loss. Thus, the average net payoff per card in deck C is a \$25 gain ($0.5*50 - 0.5*(50 - 50)$). In deck D, every card gives a definite \$50 win and there are 10% of the cards, each of which also gives a \$250 loss. That is, the average net payoff per card in deck D is a \$25 gain ($0.9*50 - 0.1*(250 - 50)$). In other words, decks C and D are advantageous since choosing decks C and D would lead to an overall gain in the long run. Participants are instructed to choose a number of cards (typically 100) and are told to switch among decks as they like with the aim of maximising their overall payoff. The payoff schedule for each deck is pre-arranged but is not explicitly told to participants: participants learn these schedules over the course of the task.

3 Preliminary observations on IGT

In this section we present our preliminary observations from two datasets obtained by past researchers on the IGT.

3.1 Switching tendency in the IGT

Our interest in the IGT was aroused by preliminary examinations of the IGT data provided by Bishara et al. (2006) and Stocco and Fum (2006) which showed a surprising trend of people switching away from a deck after a win or several wins on that deck. In other words, people switched away from an option when the previous choice did not involve any losses. It is striking because it seems counterintuitive for people to abandon a winning option. This pattern of switching on wins also seems counterintuitive as switching away from a winning deck is contradictory to what EU theory proposes.

People's switching tendency (probability of switching) was examined in two ways. Firstly, the overall switching tendency was investigated for the two datasets. Figure 1 shows the overall switching tendency observed in 31 healthy participants in Bishara et al.'s (2006) and 75 healthy participants in Stocco and Furn's (2006) studies. The switching tendency was calculated as the number of switches (n) divided by the number of choices made.

It can be seen that the overall switching tendency decreased in both cases, which means that as the task proceeded people switched less, suggesting that a preference for specific deck(s) was formed as they gradually learned about the four decks. This result is consistent with epsilon-first strategy in the multi-armed bandit algorithm (Robbins 1952), in which the gambler pulls a number of levers in a single K-slot machine to maximize his total reward in a series of trials. The epsilon-first strategy assumes that people explore all options randomly at the beginning (exploration phase) and then select the lever of highest estimated reward

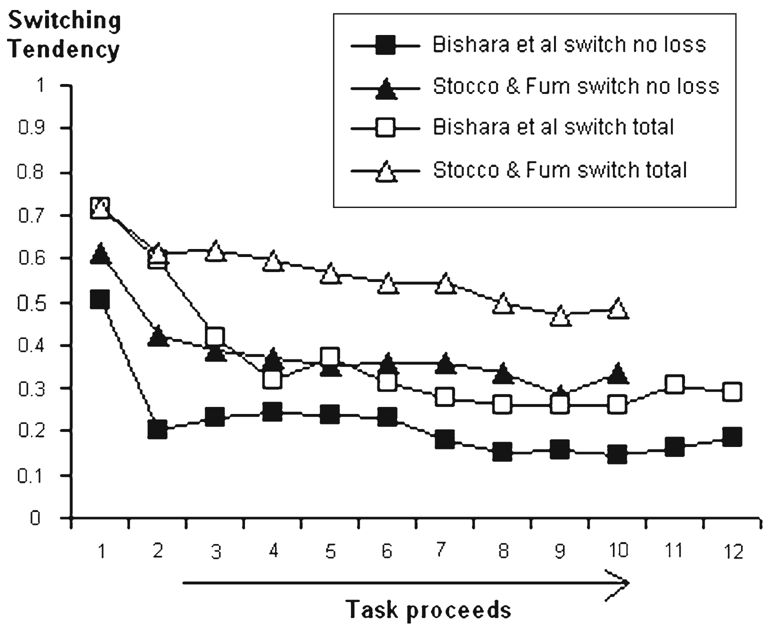

Fig. 1 The graph shows the overall switching tendency and the switching tendency when the last choice did not involve any losses as the IGT task proceeds in Bishara et al.'s and Stocco and Fum's datasets. The x axis shows the task process in which every 10 choices are grouped into one block. In Bishara et al.'s study 120 choices were made, therefore there were 12 blocks of choices; whereas in Stocco and Fum's study 100 choices were made, resulting in 10 blocks. Thus, the switching tendency was calculated as the number of switches (n) divided by 10 per block

at later stages (exploitation phase) in which their choices are less random (Even-Dar et al. 2002).

Moreover, it was also found that people tend to stay at the advantageous decks even though a loss has occurred (around 76% of the times they stayed after getting a loss for advantageous decks); and people tend not to stay at the disadvantageous decks when a loss occurs (around 12% of the times they stayed after a loss for disadvantageous decks).

Secondly, cases of switching away from a deck when the previous choice in that deck involved a win were examined in the two datasets (see Fig. 1). This switching tendency was calculated as the number of switches (n) when the last choice involved a win divided by the number of choices made. Figure 2 shows the proportion of switching when the previous choice did not involve any losses as the IGT task proceeds, which is denoted the number of switches when previous choice involved a win divided by the overall number of switches made.

In Bishara et al.'s data people switched away from a deck 26.06 times on average (SD = 15.52) when the previous choice on that deck did not involve any losses, which accounted for 55.27% of the total switches they made. In Stocco and Fum's data people switched away 37.91 times on average (SD = 17.08) when there was no loss in the previous choice for that deck, which accounted for 66.47% of the total switches they made. In other words, in both cases more than half of the times people switched away from a deck were when no loss occurred in the immediately preceding selection. From Fig. 2 it can also be seen that the switching tendency when no loss involved also decreased as the task proceeded. More specifically, the tendency decreased sharply in the first 20 choices, which suggested randomness in people's choices in the beginning, and it decreased relatively moderately in later stages.

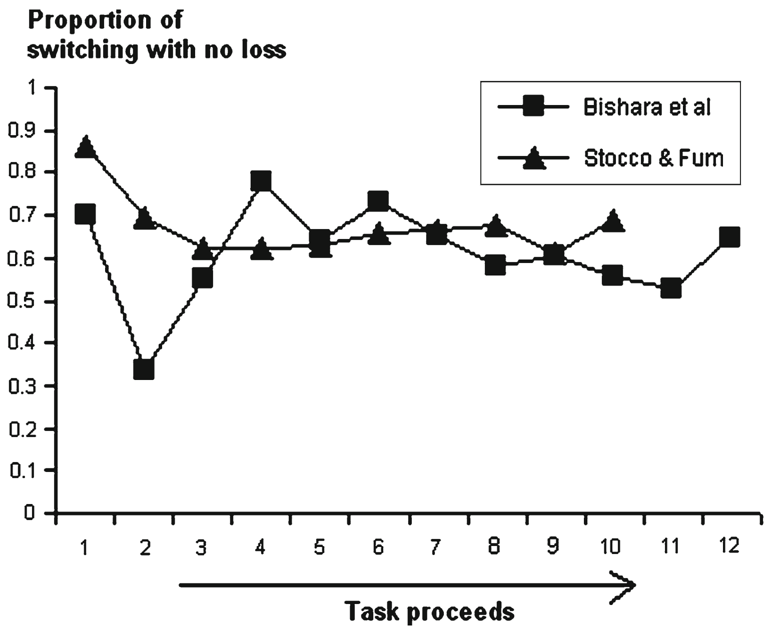

Fig. 2 The graph shows the proportion of switching when the previous choice did not involve any losses as the IGT task proceeds. This proportion was calculated as the number of switches with no loss divided by the overall number of switches. The x axis shows the task process in which every 10 choices are grouped into one block. There were 12 blocks in Bishara et al.'s data and 10 blocks in Stocco and Fum's data

3.2 Violations of decreasing/increasing absolute risk aversion

Economists have suggested that the degrees of risk aversion and the wealth one has are related (Arrow 1965): that the absolute amount of wealth a person is willing to expose to risk increases as his/her wealth increases (a phenomenon referred to as 'decreasing absolute risk aversion'). Economists also coin the term 'increasing absolute risk aversion' to refer to the phenomenon that the absolute amount of wealth a person is willing to expose to risk decreases as his/her wealth increases. In other words, a person's degree of risk aversion is related to the current wealth owned.

In our case we can regard a person as risk-prone if his/her tendency to switch away from a winning deck increases, and risk-averse if switching away from a winning deck decreases. To see whether the current total wealth and the switching tendency are related, the current sum of money in the IGT and the tendency of switching away when no loss occurs were correlated for each participant in Bishara et al. and Stocco and Fum datasets. The correlation coefficient was then compared to 0 to see whether the overall correlation was positive or negative. It was found that the mean coefficient was −0.11 (SD = 0.31) and was no different from 0 ($t = -1.965$; df = 30; $p > 0.05$) for Bishara et al. dataset. For Stocco and Fum dataset it was found that the mean coefficient was −0.09 (SD = 0.39) and was also not different from 0 ($t = -1.883$; df = 74; $p > 0.05$). In other words, in both cases the current wealth and the degree of risk aversion were not related. Thus, it is safe to assume that people's tendency to switch away from a deck when no loss has occurred previously is not due to their decreasing risk aversion due to increases in the current total wealth they have earned. In the next section we describe various other models of people's behaviour in the IGT task, and examine their ability to account for this switching behaviour.

4 Current models on the IGT

To date, only a handful of cognitive models have been developed to account for the results in the IGT (Levine et al. 2005; Fum and Stocco 2004). These models have primarily focused on people's overall preference for specific decks but ignored the exact details of how people choose and how they switch among the four decks. This section presents two prominent computational models proposed by Busemeyer and Stout (2002).

4.1 Bayesian-expected utility model

Details of the Bayesian-EU model are not described here (see Busemeyer and Stout 2002). This model claims that a deck should be chosen if it has the maximum EU. This model successfully captures the rational aspect of people's choices. That is, a win increases the EU of a deck and suggests that people should stay with that deck. In contrast, a loss decreases the EU of the deck and suggests people might switch to another deck if it has the maximum EU. However, there are two limitations with this model. Firstly, this model does not have a memory mechanism for previous choices. That is, the model makes a decision based on the last choice. This means that if a large loss is encountered (for example, in deck B), the EU decreases dramatically and it may drop to so low that the deck is unlikely to be chosen again. However, people still choose deck B quite frequently despite a previous large loss, because there is 90% chance of winning in that deck. Levine et al. (2005) suggest that people's preference for deck B may be due to their tendency to underweight rare events

(Barron and Erev 2003). This model fails to capture this aspect of people's choices. Secondly, in terms of switching behaviour this model only predicts that the deck with the maximum EU should be chosen. It cannot explain the times when people choose a deck which does not have the maximum EU, for example, when people switch away from a deck when no loss occurs.

4.2 Expectancy-valence learning model

Details of the expectancy-valence learning model are not described here (see Busemeyer and Stout 2002). This model assumes that people integrate the wins and losses experienced on each choice into a single affective reaction called a valence, and the deck with the highest expected valence should be chosen. According to Busemeyer and Stout (2002), this model gives the best fit for the IGT data. However, this model has several limitations. Firstly, the model only considers the valence generated by the wins and losses, but not the probability of wins and losses. Secondly, if a large loss is encountered the expected valence drops significantly, and so does the probability of choosing that deck, thus the deck may be never chosen again. Thirdly, as with the previous model this model cannot explain cases were people choose a deck which does not have the maximum valence; that is, the model cannot capture people's tendency to switch when no loss is involved.

5 Rational or irrational?

It seems counterintuitive that people switch away from a deck even though their most recent choice from that deck did not involve any losses. From the EU point of view, obtaining a win from a deck has two functions. Firstly, a win increases the utility of that deck. If the utility of the deck is increased and the utility was already the highest (that being the reason why the deck was chosen in the first place), then the likelihood of leaving the deck is decreased. Secondly, getting a win decreases the ratio of losses to wins in a deck, and thus increases the estimate of winning in the future from that deck. Expected utility theory thus predicts that if someone's previous card from a deck was a win, they will pick again from that deck and will not switch away. People's observed tendency of switching away when no loss occurs thus presents a paradox from the perspective of expected-utility theory, a paradox that none of the rational accounts (Bayesian-EU model and expectancy-valence model) can explain.

However, we cannot conclude that people are behaving irrationally just because their observed switching behaviour cannot be accounted for by the rational theories. The rationality of their switching is debatable since people undertaking the IGT are never completely certain about the probability of winning or losing, or in other words, their knowledge about their options is never complete. There are some other reasons for such switching that we have not considered. The question that needs to be answered is: if no loss occurs then why do people switch away? There are several explanations. Some authors have indicated that people making repeated choices among a limited set of alternatives often switch away from optimal options to other options from which they expect to derive less pleasure simply for the sake of variety (Kahn et al. 1997; Ratner et al. 1999). Levine et al. (2005) suggest that even when people have developed a distinct preference for specific deck(s), they might occasionally choose other decks to 'try their luck', or switch to other decks simply due to fatigue or boredom.

In this paper we argue that the reason for switching away from a deck when no previous loss occurs is due to a caution factor which represents people's nervousness of obtaining a loss on the next choice. This view is consistent with the gambler's fallacy in which people mistakenly believe that a successive sequence of the same event occurring increases the probability of a different event occurring (Kahneman and Tversky 1972). For example, if a person tosses a coin ten times and all gets heads, the person tends to mistakenly believe that a tail is more likely (on the next toss) than if the person tosses six times and gets heads. It should be noted, however, that in the gambler's fallacy the events are random and independent (the fallacy consists in believing that a repeated sequence of heads, for example, increases the subsequent probability of a tail, whereas in reality the occurrence of a sequence of heads tells us nothing about whether or not a tail will occur). The IGT, by contrast, is one where the cards which participants see give them useful information about the decks before them. As participants go through the task they gain more knowledge about the decks and will therefore gradually learn about the payoff schedule for each deck as the task goes on. In a gambler's fallacy task seeing ten heads in a row tells you nothing about what the next coin toss will be. In the IGT, however, seeing ten winning cards from a given deck tells you that winning cards are frequent in that deck, and so the next card is more likely to be a win.

Why, then, do people switch away after a win in a deck? We postulate that a participant who has won on ten consecutive cards in a deck may believe that a loss is more likely on the eleventh card, and therefore he/she might switch away from that deck. Thus, this factor serves as a caution for not sticking with an option for too long; and it is puzzling in a sense that people switch away regardless of the wins and losses associated with that deck. We also propose that, the longer people have stayed at a deck continuously, the more nervous they are about getting a loss on the next card in that deck, and thus the more likely they are to switch away from that deck. Therefore, the reason for switching away is mainly due to two factors: one is obtaining a loss, represented by a utility factor in which the decision to switch away is determined by the utilities of the decks; and the other is nervousness about obtaining a loss, represented by a caution factor which depends on the number of continuous cards drawn from the deck so far. We further propose that if people were to switch away from a deck, they would switch to a deck which gives the largest wins and also the highest probability of winning.

6 A computational model of switching behaviour

We propose a Utility-Caution model which examines task factors in people's choices in the IGT. Moreover, the model predicts whether people should switch away from or stay at a deck at a given point of the task, and also predicts the deck to which people will switch when switching takes place. There are two task factors that determine switching. The utility factor (F_1) describes the contingencies of winning and losing and the caution factor (F_2) represents the number of choices people have made in a deck consecutively.

6.1 Task factors

For F_1 it is assumed that the probability of switching away is an increasing function of the probability of losing multiplied by the amount of losses and a decreasing function of the probability of winning multiplied by the amount of wins. The estimated probability of a loss is the same as the one in the Bayesian-EU model, which is an estimated probability $P_D(t)$ given that deck D is chosen on trial t using a beta prior distribution updating rule:

$$P_D(t) = \frac{f_D(t) + f(0)}{n_D(t) + n(0)} \tag{1}$$

where $f_D(t)$ is the number of cards producing a loss experienced by choosing deck D up to and including trial t; $n_D(t)$ is the total number of trials that deck D was chosen; and $f(0)$ and $n(0)$ reflect the prior estimates before any experience.

The amount of losses $L[D(t)]$ is instantiated as the average of total losses experienced by choosing deck D up to and including time t and is denoted:

$$L[D(t)] = \frac{\sum_{j=1}^{N} L[D]}{N} \tag{2}$$

where N is the total number of cards chosen so far in deck D. The probability of winning is simply $1 - P_D(t)$, which means the probability of cards producing a net win. The amount of wins $R[D(t)]$ is instantiated the same way as the losses, which is the average of total wins experienced by choosing deck D up to and including time t:

$$R[D(t)] = \frac{\sum_{j=1}^{N} R[D]}{N} \tag{3}$$

where N is the total number of cards chosen so far in deck D. Therefore, F_1 is denoted:

$$\mathrm{F}_1 = \frac{P_D(t) \times L[D(t)]}{(1 - P_D(t)) \times R[D(t)]} \tag{4}$$

It is predicted that F_1, the utility factor is positively correlated with the switching probability. That is, the larger the probability of losing and the average losses compared to the probability of winning and the average wins of the deck, the more likely people are to switch away from that deck on the next choice. If people do switch (out of their own will or if the preferred deck is exhausted), they would choose the deck with the maximum $(1 - P_D(t)) \times R[D(t)]$.

F_2, the caution factor, is denoted n/N, where n is the number of consecutive choices from the same deck and N is the number of total choices from that deck. It is assumed that F_2 will be positively correlated with the switching probability, since the longer people stay in a deck consecutively, the more likely they are to switch away from that deck. The probability of switching away from a deck is a multiple regression predictor equation computed as follows. Multiple regression produces the best fit with F_1 and F_2 as the predictor variables, and switching probability as the dependent variable.

$$S = \mathrm{a} \times \mathrm{F}_1 + \mathrm{b} \times \mathrm{F}_2 \tag{5}$$

6.2 Multiple regression

To determine the regression coefficients in the Utility-Caution model, the dataset by Bishara et al. in which a total of 120 cards were chosen for each of the 31 participants was used. A Java program was written to conduct the assessment. The switching probability was calculated as a probability function which was the number of switches in every N choices (N could range from 2 to the total number of choices, and in this case $N = 10$). For example, if there were 7 times of switching in the first 10 choices, the switching probability of the first 10 choices was 0.7. The values of F_1 and F_2 were calculated for each choice for each participant. To match the switching probability data the values of F_1 and F_2 in every N choices were averaged.

In Fig. 3 it is seen that the switching probability decreased in the task, which means that as people have gained more experience with the four decks they gradually develop a preference for a specific deck(s), therefore they do not switch as much as they used to. The values of

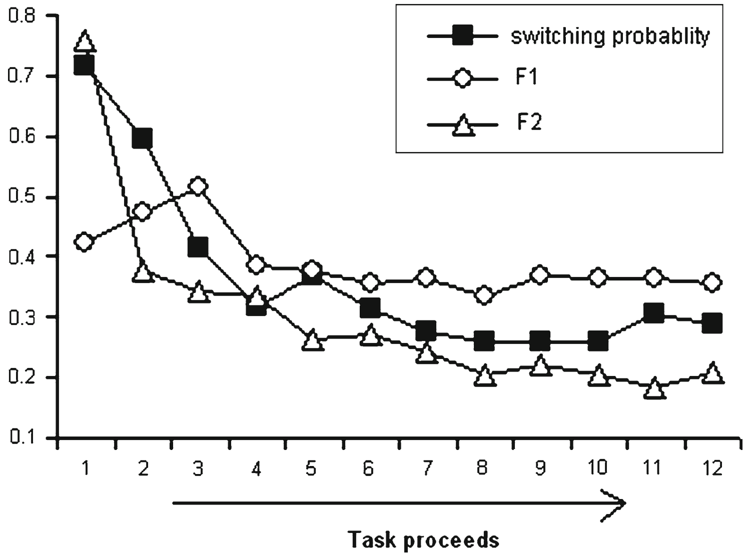

Fig. 3 The graph shows the switching probability and the values of F1 (the utility factor) and F2 (the caution factor) as the task proceeds for Bishara et al.'s dataset. The x axis shows the task process in which the 120 choices in the IGT are broken down to 12 blocks of 10 choices

F_1 and F_2 followed the same trend as the switching probability, which gave good support to the expectation that the coefficients a and b would both be positive. In fact, F_1 values and switching probability were positively correlated ($r = 0.658$, $p < 0.05$) and F_2 values and switching probability were also positively correlated ($r = 0.905$, $p < 0.01$), which supported the previous predictions.

Multiple regression analysis was run across all individual participants in the data of Bishara et al. The two coefficients were compared to 0 using one-sample t-test. It was found that the mean of coefficient a was 0.146 and was significantly above 0 ($t = 2.456$; $p < 0.05$). It was also found that the mean of coefficient b was 0.151 and was also significantly above 0 ($t = 2.243$; $p < 0.05$). These results confirm the model's predictions.

7 Assessment

Since the three computational models all operate under certain learning mechanisms, it is safe to assume that no model can work without seeing a number of choices first. That is, the models cannot make predictions without having a certain amount of experiences. We propose that a fair way to assess these models is by using a novel next-choice-prediction method, in which we give the models the first n choices made by the participant and then ask the model to predict the $n+1$ choice. In this way the model makes the prediction based on the same experiences shared with the participant. n can vary from 1 to $N-1$ (N is the total number of cards to be picked) and the model can at most make $N-1$ predictions. Prediction accuracy is calculated as the number of accurate predictions divided by the total number of predictions.

7.1 Assessing the three models

The Bishara et al. dataset was used for assessing the three models using the next-choice-prediction method. For the Utility-Caution model, the coefficients ($a = 0.146$, $b = 0.151$) were put into the equation: $S = \text{a} \times \text{F}_1 + \text{b} \times \text{F}_2$. The threshold for switching was set to 0.5, arbitrarily. This means that if the S value was below 0.5, then the model should stay with the same deck as before; if the S value was equal to or greater than 0.5, the model should switch to another deck with the maximum $(1 - P_D(t)) \times R[D(t)]$. The threshold for switching was also tested, using 0.1, 0.3, 0.5, 0.7, 0.9, and it was found that 0.5 gave the highest prediction accuracy.

It was found that the average correct predictions were 75.90 (SD = 18.16) out of 119 predictions (see Fig. 4), which means that the model could predict accurately people's choices 63.8% of the time. Since the criterion for judging accuracy is 25% (1 in 4 decks), the prediction accuracy for each participant was compared to 0.25 (the chance level). It was found that the accuracy of Utility-Caution model was significantly higher than 25% ($t = 14.148$; df = 30; $p < 0.01$).

The Bayesian-EU model and the expectancy-valence learning model were also tested using next-choice-prediction method (the reported parameters in Busemeyer and Stout (2002) for the two models were used). Figure 4 shows the prediction accuracy of the three models. The Bayesian-EU model gave an average of 47.13 (SD = 20.15) accurate predictions out of 119 total predictions made, which means that it predicted accurately 39.6% of the times. The expectancy-valence learning model gave an average of 48.45 (SD = 21.93) accurate predictions out of 119 predictions, which means that it predicted accurately 40.7% of the times. While the accuracies for the two competing models remained stable during the task,

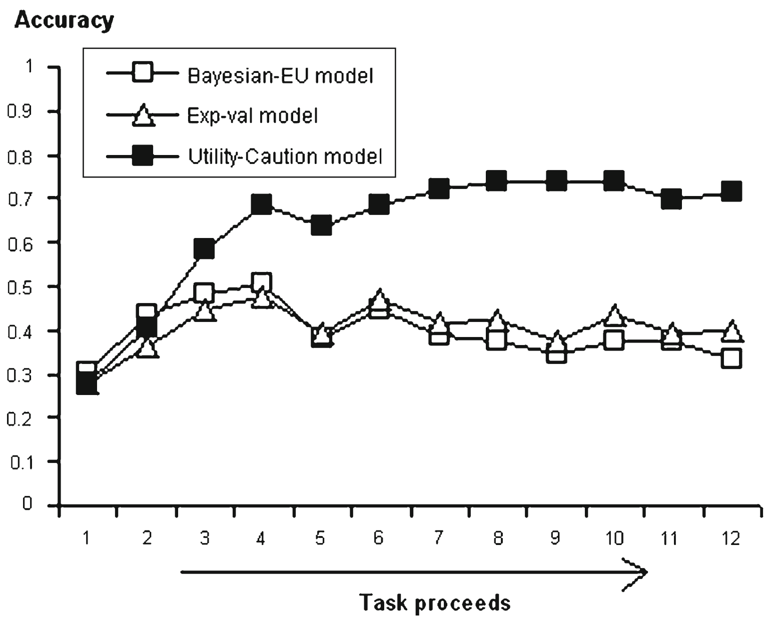

Fig. 4 The prediction accuracy for the three models as the task proceeds in Bishara et al.'s dataset. The x axis shows the task process in which the 120 choices in the IGT are broken down to 12 blocks of 10 choices. The y axis shows the accuracy of predictions the three models made as compared to the choices people made

the accuracy for our model increased consistently and remained at 70%. The three model prediction accuracies for each participant were compared and it was found that our model gave significantly more accurate predictions than the other models (F = 20.169; df = 2; $p < 0.01$). Thus, we conclude that our Utility-Caution model was superior to the Bayesian-EU model and the expectancy-valence model.

7.2 A further assessment

Results obtained in the Bishara et al.'s dataset suggest that our Utility-Caution model performs better than the previous Bayesian-EU model and expectancy-valence model. To investigate the Utility-Caution model's performance relative to the other two models in other contexts, we used the Stocco and Fum's dataset which involved a different set of people's responses in the IGT. The Stocco and Fum's dataset is slightly different from the Bishara et al.'s, as it involves significantly more participants (75 participants in Stocco and Fum's dataset, as opposed to 31 in Bishara et al.'s data) and also uses a modified IGT task which operates under a different payoff schedule. By applying our model to this alternative dataset the reliability and validity of our model can be demonstrated.

For the Utility-Caution model the coefficients (a = 0.146, b = 0.151) were put into the regression $S = a \times F_1 + b \times F_2$ and the model was assessed using the next-choice-prediction method. In Stocco and Fum's dataset a total of 100 cards were chosen for each of the 75 participants. It was found that the model's average correct predictions were 62.36 (SD = 8.70) out of 99 predictions, which means that the model could predict people's choices accurately at 62.99% of the times. The accuracy was significantly higher than 25% ($t = 15.46$; df = 74; $p < 0.01$), which means that the accuracy level was well above chance level. This result is consistent with the accuracy level (63.8%, see Fig. 4) earlier in the dataset of Bishara et al.

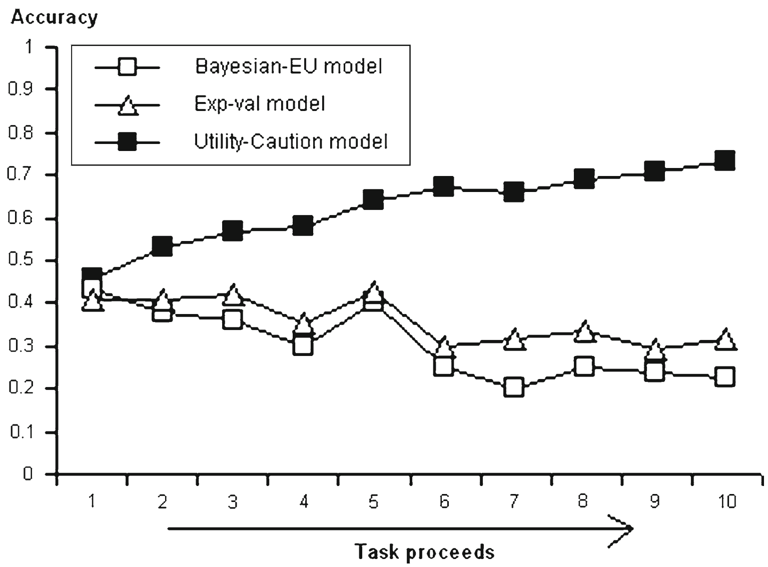

Fig. 5 Prediction accuracy for the three models as the task proceeds in Stocco and Fum's dataset. The x axis shows the task process in which the 100 choices in the IGT are broken down to 10 blocks of 10 choices. The y axis shows the accuracy of predictions the three models made as compared to the choices people made

The Bayesian-EU model and the expectancy-valence model were also tested. Figure 5 shows the prediction accuracy of the three models during the task. It is seen that the accuracy of our model persisted fairly consistently above 0.6 and gradually increased in the task, whereas the other models showed a consistent decrease in accuracy. The Bayesian-EU model predicted on average 30.34 (SD = 8.30) choices accurately out of 99 predictions, and the expectancy-valence model predicted 35.84 (SD = 5.27) choices accurately. The three model prediction accuracies for each participant were compared and it was found that our model predicted significantly more accurately than the other two models ($F = 40.342$; $df = 2$; $p < 0.01$). The superiority of the model is consistent with the previous findings. Overall, the high consistency between the results from Bishara et al. and the results from the Stocco and Fum shows a high reliability of the model.

8 Discussion

This paper makes two contributions to the literature on people's choices in repeated gambling tasks. One contribution is the next-choice-prediction method used in this study. For useful cognitive modelling it is necessary for a model to have some experiences to make future predictions. A good way to examine the accuracy of a model is to give the model the sequence of choices and responses seen by each individual participant, and ask the model to predict the next choice or behaviour for that participant. The model's accuracy can then be assessed by comparing its predicted choice with that made by the participant.

The second contribution made by this paper is a computational model that successfully accounts for people's switching behaviour among a finite set of options in repeated gambling scenarios. The model claimed that people's tendency to switch away from an option is in direct proportion to a utility factor, which represents the probability and the amount of losses experienced compared to wins of an option, and a caution factor, which represents the number of choices made consecutively in an option. The results obtained so far showed that our Utility-Caution model produced more accurate predictions of people's actual choices than the Bayesian-EU model and the expectancy-valence learning model, not only in Bishara et al.'s but also in Stocco and Fum's datasets. The reason to this superior performance was that our model incorporated a novel factor (the caution factor) which has not been identified in the existing literature and could account for people's switching behaviour more precisely.

The utility factor reflects the rational part of people's decision for switching. For the advantageous decks the denominator was consistently larger than the nominator, which decreased the value of the factor, making it unlikely to switch away. This explained the fact that people tend to stay at the advantageous decks even when an occasional loss was occurred in those decks. Conversely, for the disadvantageous decks the nominator was consistently larger than the denominator, thus it was more likely to switch away, especially when a loss occurred. This explained the fact that people tend not to stay at the disadvantageous decks when a loss occurs.

The caution factor suggested that people tend to switch away from a deck not only because it involves losses or gives the least payoff as traditional decision models have claimed, but also because people have stayed long enough in this option. It provided an account for the observed tendency that more than half of the times that people switched away from a deck were when no loss occurred on the previous choice. Thus, regardless of the advantageousness of the decks, this factor serves as a caution for not staying with a deck for too long. It also implies that people are prudent and vigilant beings in that they are wary of possible future losses or damages lying ahead.

One might question why people abandon the winning option in the task, a tendency contrary to the rational choice theory and the EU theory. Gigerenzer (2007) has provided an explanation. He argues that much of our reasoning, decisions, and judgments are based on our unconscious intelligence, intuitions or gut feelings, whose underlying reasons we are not fully aware of. We can argue that it is our intuition that we switch away from a winning option in which we have chosen successively for the nervousness or fear of losing soon. However, people nonetheless behave according to this rule of thumb.

Although the Utility-Caution model can predict with above 60% accuracy on average, there is still 30% that the model cannot predict. We suggest that the remaining percentage is due to the randomness inherent in people's choices, especially during the beginning of the task when people have no experiences about their options.

Our model suggests that the longer people have stayed at a deck, the more nervous (afraid) they feel about getting a loss in that deck. This idea is not tested in the current analysis, since with the available datasets we have no access to people's emotional states during the task. Thus, for future work we suggest that the emotional states can be examined. In fact, we propose that there are two emotion factors (E_1 and E_2), each of which corresponds to the task factors (F_1 and F_2). A further experiment on the IGT is to be conducted with two purposes: firstly, our model can be again validated in this experiment; secondly, the two emotion factors can be examined.

In this experiment E_1 represents the level of happiness of a participant at a given point in the gambling task. It is assumed that the happier people are about their choices, the less likely they are to switch away. Thus, E_1 is predicted to be negatively correlated with the switching probability. A second factor E_2 represents the level of nervousness about their choices. It is assumed that the more nervous people are about their choices, the more likely they are to switch away from that deck. Thus, E_2 is predicted to be positively correlated with the switching probability. The multiple regression predictor equation is used in the emotion model: $S = \mathrm{c} \times E_1 + \mathrm{d} \times E_2$. The coefficient c is predicted to be negative and the d to be positive.

In conclusion, the Utility-Caution model proposed in this paper suggests that people's tendency to switch away from an option is due to the losses produced by choosing that option compared to gains and also to the number of choices made in that option consecutively. This model showed better performances than the Bayesian-EU model and the expectancy-valence learning model.

Acknowledgements The authors wish to thank the School of Computer Science and Informatics at UCD, Dr. Bishara and Prof. Busemeyer, and Dr. Stocco for their help and support in this study. The second author's contribution to this work was supported by the FP6NEST Programme of the European Commission (ANALOGY: Humans—the Analogy-Making Species: STREP Contr. No 029088).

References

Arrow KJ (1965) Aspects of the theory of risk-bearing. Hahnsson Foundation, Helsinki

Barron G, Erev I (2003) Small feedback-based decisions and their limited correspondence to description-based decisions. J Behav Decis Mak 16:215–233

Bechara A, Damasio H, Tranel D, Anderson SW (1994) Insensitivity to future consequences following damage to human prefrontal cortex. Cognition 50:7–15

Bernoulli D (1967) Exposition of a new theory on the measurement of risk. Gregg Press, England (Original work published in 1738)

Bishara AJ, Pleskac TJ, Fridber DJ, Yechiam E, Lucas J, Busemeyer JR, Fin PR, Stout JC (2006) Models of risky decision-making in Marijuana and stimulant users. Unpublished Manuscript

Busemeyer JR, Stout JC (2002) A contribution of cognitive decision models to clinical assessment. Psychol Assess 14:253–262

Coombs CH, Meyer DE (1969) Risk-preference in coin-toss games. J Math Psychol 6:514–527

Even-Dar E, Mannor S, Mansour Y (2002) PAC bounds for multi-armed Bandit and Markov decision processes. In: Fifteenth annual conference on computational learning theory (COLT), pp 255–270

Fum D, Stocco A (2004) Memory emotion and rationality: an ACT-R interpretation for gambling task results. In: Schunn CD, Lovett MC, Lebiere C, Munro P (eds) Proceedings of the sixth international conference on cognitive modelling. Lawrence Erlbaum, Mahwah, NJ, pp 211–216

Gigerenzer G (2007) Gut feelings: the intelligence of the unconscious. Penguin Books Ltd, London

Kahn BE, Ratner RK, Kahneman D (1997) Patterns of hedonic consumption over time. Mark Lett 8:85–96

Kahneman D, Tversky A (1972) Subjective probability: a judgment of representativeness. Cogn Psychol 3:430–454

Kahneman D, Tversky A (1979) Prospect theory: an analysis of decisions under risk. Econometrica 47:263–291

Kahneman D, Slovic P, Tversky A (1982) Judgment under uncertainty: heuristics and biases. Cambridge University Press, Cambridge, England

Keren G, Wagenaar WA (1987) Violation of utility theory in unique and repeated gambles. J Exp Psychol Learn Mem Cogn 13:387–391

Levine DS, Mills B, Estrada S (2005) Modeling emotional influences on human decision making under risk. In: Proceedings of internal joint conference on neural networks. IEEE Press, Montreal, pp 1657–1662

Mas-Collel A, Whinston M, Green JR (1995) Microeconomic theory. Oxford University Press, New York

Ratner RK, Herbst KC (2005) When good decisions have bad outcomes: the impact of affect on switching behavior. Organ Behav Human Decis Process 96:23–37

Ratner RK, Kahn BE, Kahneman D (1999) Choosing less-preferred experiences for the sake of variety. J Consumer Res 26:1–15

Robbins H (1952) Some aspects of the sequential design of experiments. Bull Am Math Soc 55:527–535

Simon HA (1982) Models of bounded rationality. MIT Press, Cambridge

Stocco A, Fum D (2006) Memory and emotion in the gambling task: the case for independent processes. In: R Sun, N Miyake (eds) Proceedings of the 28th annual conference of the cognitive science society. Lawrence Erlbaum Associates, Mahwah, NJ, pp 2192–2197