Rewards and Punishments in Iterated Decision Making: An Explanation for the Frequency of the Contingent Event Effect

Antonio Napoli (antonio.napoli@phd.units.it)
Danilo Fum (fum@units.it)
Dipartimento di Psicologia, Università degli Studi di Trieste
via S. Anastasio, 12, I-34134 Trieste, Italy

Abstract

Iterated decision making can be studied in laboratory using situations, like the Iowa Gambling Task (IGT), in which participants face repeatedly the same decision problem getting feedback after each choice. In the paper we focus on a recurring finding in experiments carried out with the IGT, the frequency of the contingent event effect—i.e., the fact that people consistently prefer options associated with rare losses, independently of their attractiveness, expected value and loss magnitude—that has not yet received a satisfactory explanation. An experiment reveals that the effect relies on simply experiencing rewards and punishments, not being influenced by the net outcome (loss or win) to which they are associated, and a computational model, implemented in the ACT-R cognitive architecture, corroborates the idea that punishments and losses on one hand, and rewards and wins on the other, play the same functional role in determining the participants' behavior in IGT.

Keywords: Iterated decision making; Reinforcement learning; Iowa Gambling Task; ACT-R; Feedback.

Introduction

Iterated decision making relies on the regulation of behavior according to its consequences. This process is characterized by three steps (Ahn, Busemeyer, Wagenmakers, & Stout, 2008): (1) the choice of a possible option and the execution of the associated action, (2) the encoding of the action consequences, (3) the integration of the consequences in a format that allows options comparison. Iterated decision making can be simulated in laboratory using the so-called *multi-armed bandit* tasks (Sutton & Barto, 1998) in which participants face repeatedly the same decision problem and get a numerical reward after each choice. Behavior in multi-armed bandit tasks is usually modeled by Reinforcement Learning models in which agents, requested to maximize their expected total reward over a given number of trials, learn about the structure of the environment by taking into account the reward associated with each choice. In the paper we will adopt Reinforcement Learning to explain the results obtained in a particular multi-armed bandit task, the Iowa Gambling Task—henceforth, IGT (Bechara, Damasio, Damasio, & Anderson, 1994). Our models will be based on the ACT-R cognitive architecture (Anderson, 2007) which provides the resources for the steps (1) and (3) of the decision making process described above, and we try to figure out how step (2) is carried out.

The IGT has been proposed as a simulation of real life decision making in the way it factors reward, punishment and outcome uncertainty (Bechara et al., 1994). The IGT involves four decks of cards. Participants repeatedly choose a card at a time from one of the decks. Each time a card is turned, it allows participants to gain a given amount of money, but sometimes the card forces them to give up some money, too; therefore, while all cards contain a reward, only some cards contain a punishment. Two card decks (let's call them A and B) feature high wins per card ($100) but they yield also higher losses so that, by choosing them, participants lose more money than they win. These decks are referred to as "bad decks". The remaining decks (C and D) give rise to small gains ($50) but even smaller losses, so that it is profitable to choose cards from them. These decks are referred to as "good decks". Generally participants, after being initially attracted by the dangerous bad decks featuring high wins and higher losses, gradually shift their preferences toward the good ones, a result which has been replicated by most IGT studies (Dunn, Dalgleish, & Lawrence, 2006). So, according to the standard interpretation, participants' behavior can be explained by a conflict between two deck features: their *attractiveness*, i. e., the amount of money each cards allows immediately to win—which drives the participants choices in the first trials—and the long term *expected value*, i. e., the net amount of money gained or lost— which drives them in the subsequent trials.

In recent years a growing number of researchers have been suggesting that this interpretation of the IGT is unsatisfactory (see Dunn et al. (2006) for a critical review of the literature). In the present paper we will focus on a recurring finding in the experiments carried out with the IGT which has not yet received a satisfactory explanation. This finding has been termed the "frequency of the contingent event effect" by Fum, Napoli, and Stocco (2008) and the "prominent deck B phenomenon" by Chiu et al. (2008) and refers to the fact that people consistently prefer the decks associated with rare losses—to the point that the bad-but-rare-loss deck B which gives raise to a small number of losses is consistently preferred to the good-but-frequent-loss deck C—independently of their attractiveness, expected value and loss magnitude. Even if the theoretical interpretations of the phenomenon put forward by the two research groups are similar, they differ in some important details.

Frequency of the contingent event

Traditionally, the performance in the IGT has been recorded by subtracting the number of bad deck selections from the good ones (the so-called Good−Bad index). In the original version of the IGT (see Table 1), for every block of ten cards, decks A and C originate five money losses while decks B and D give rise to only one.

Table 1: Deck matrices of the original Iowa Gambling Task

Card	A		B		C		D	
#	**Rew**	**Pun**	**Rew**	**Pun**	**Rew**	**Pun**	**Rew**	**Pun**
1	+100	0	+100	0	+50	0	+50	0
2	+100	0	+100	0	+50	0	+50	0
3	+100	-150	+100	0	+50	-25	+50	0
4	+100	0	+100	0	+50	0	+50	0
5	+100	-300	+100	0	+50	-75	+50	0
6	+100	0	+100	0	+50	0	+50	0
7	+100	-200	+100	0	+50	-25	+50	0
8	+100	0	+100	0	+50	0	+50	0
9	+100	-250	+100	-1250	+50	-75	+50	-250
10	+100	-350	+100	0	+50	-50	+50	0
EV	Bad		Bad		Good		Good	

Rew: Reward. Pun: Punishment. EV: Expected Value. Punishments which do not result in a net loss are evidenced in gray.

Because A and B are the bad decks and C and D are the good ones, any possible effect of the number of losses is confounded with that of the deck quality, as expressed by their expected value. In recent years researchers have started to present the analytical data for each deck and evidence has been growing about the "frequency effect", i.e. the functional role that the frequency of money losses could play in addiction (or in opposition) to the effects of decks' attractiveness and expected value.

To understand which deck features exert the most important effect on IGT, Fum et al. (2008) manipulated the decks pay-off matrices in three different experimental conditions. In all the conditions the decks attractiveness and the loss frequency were kept the same as in the original IGT, while their expected values were manipulated. The first condition replicated the setting of the original IGT. In the second condition the expected value of the decks was zeroed, so that the amount of money participants were expected to win in the long run for each deck was identical to that they were expected to lose. In the third condition the two decks with frequent punishments (A and C) were good while the decks with less frequent punishments (B and D) were the bad ones; in this case loss frequency and expected value were put in opposition for each deck.

Two findings were particularly significant: (1) the number of selections from each deck was almost the same in all the conditions, and (2) participants showed a strong preference for the low frequency loss decks, even in the condition in which these decks were bad. In the same study, the IGT task was carried out in a scenario in which participants always lost money when they turned a card while the contingent event was represented by a win, a variant originally developed by Bechara, Tranel, and Damasio (2000). Similar results were obtained with the same pattern of choices in all the conditions and a strong preference for the decks originating a higher number of wins. The fact that participants chose the same number of cards from all decks despite the change in their expected value means that this feature plays a small or no functional role in determining their choices. The fact that participants preferred the decks with a small number of losses (or those with a high number of wins) means that the frequency effect is both independent from and much stronger than the effect of the other two features. This effect was termed "the contingent event effect".

An important empirical finding remains, however, unexplained by the contingent event effect and it is constituted by the fact that, when this effect is confounded with that of the expected value, a preference for the economically advantageous decks (a "goodness" effect) is normally found which indicates that the frequency of the contingent event cannot cover the whole story in the IGT. Stocco, Fum, and Napoli (2009) hold the idea that participants' behavior in this task is guided by a dual process. The first one is a low-level emotion-based mechanism which is sensitive to punishment (or reward) frequency, while the second one, high-level and based on the analysis of the monetary outcomes, is sensitive to the decks' expected value. Even if the former is normally the most important factor in guiding participants' choices, the latter may sometimes enter into play being responsible for the goodness effect.

A different explanation for the goodness effect which devaluates the deck's expected value has been put forward by Chiu et al. (2008). In order to understand their proposal it is necessary, however, to introduce some terminological distinctions.

From now on, we will discriminate between a punishment and a loss, on one hand, and between a reward and a gain, on the other. A *punishment* is an event that happens every time participants turn a card that makes them give away money. So, for example (see Table 1), in card #3 of deck C, after having earned \$50 you are forced to give \$25 back, and this is a punishment. A *loss* is a particular kind of punishment in which the amount of money lost is higher than that won; so, in card #3 of deck A, you win \$100 but you are forced to refund \$150, and this constitues a loss. All losses are therefore punishments, but not vice versa. In the same vein, in the variant IGT in which every card turn makes you lose some money, a *reward* is a contingent event in which you earn some money while a *gain* is a reward in which the amount of money gained is higher than that lost.

Chiu et al. (2008) argue that the process driving participants' behavior in the IGT is sensitive to loss (in the sense we have just defined) frequency. Some cards in deck C (evidenced in gray in Table 1) present a punishment which is not a loss, as for example the card (+\$50, -\$25), whose outcome is a net gain of \$25. Every block of 10 cards, deck C contains on average 6.25 gains, 2.5 standoffs and 1.25 losses, deck D contains 9 gains and 1 loss, deck A contains 5 gains and 5 losses, and deck B contains 9 gains and 1 loss. Therefore, taken together, the good decks (C and D) present a total of 15.25 gains, 2.5 standoffs and 2.25 losses, whereas the bad decks

(A and B) present 14 gains and 6 losses. According to Chiu et al. (2008), the lower number of losses in the good decks explains the participants' preference for them. These authors also propose their own version of the IGT, the Soochow Gambling Task (henceforth, SGT), in which every punishment is always a loss, thus eliminating the "ambiguous" Deck C. In SGT the bad decks have a high number of wins, while the good decks have a high number of losses. Results show that participants choose more cards from the former than from the latter type of decks, and this corroborates the idea that their behavior is more sensitive to losses than to expected value.

The proposals of the two research groups differ in two respects: the first one is that Fum et al. (2008) assume that participants avoid all kind of punishments, while according to Chiu et al. (2008) they avoid only punishments which result in a net loss. The second, which is strictly tied to the first, is that according to Stocco et al. (2009), the goodness effect is due to an understanding of the decks' expected value, while according to Chiu et al. (2008) the goodness effect is due to the lower number of losses in the deck C. In this paper we present an experiment which tries to distinguish between the two proposals by addressing the (possible) different effects of punishments and losses.

The Experiment

A first idea for discriminating between the above mentioned positions is to compare the choices made from two different kind of decks that, while sharing the same expected value, provide the same number of punishments but a different number of losses. So, the first deck should give rise to a given number of losses (which are all punishments) while the second should originate the same number of punishments of which, however, only some are losses. According to Chiu et al. (2008), participants should prefer the latter kind of deck while, according to Fum et al. (2008), participants should choose the same number of cards from the two decks.

A second way of discriminating between the hypotheses would take into account the specific format of the information provided during the experiment, i.e., the feedback received after each choice. In the original IGT, participants received a "double feedback" stating separately the amount of money provided by the default and the contingent event (which could be possibly null). In a "single feedback" task (such as the SGT) each card turn informs only about the net amount of money lost or gained. According to Chiu et al. (2008), participants should exhibit the same pattern of choices both in a Single and in a Double feedback task, while, according to Fum et al. (2008), participants should modify their behavior whenever the manipulation changes the number of punishments in one or more decks.

In the experiment we contrasted the participants' behavior in a variant of the IGT featuring both a Double feedback and a Single feedback condition. In the Double condition all the decks (A, B, C and D) provided the same punishment frequency (5 every 10 cards), but for two of the decks (A and C) all the punishments were losses (giving thus 5 losses every 5 punishments) while the remaining decks (B and D) provided only 1 loss every 5 punishments (see Table 2).

Table 2: Deck matrices of the Double Feedback - Standard Frame condition.

Card	**A**		**B**		**C**		**D**	
#	**Rew**	**Pun**	**Rew**	**Pun**	**Rew**	**Pun**	**Rew**	**Pun**
1	+90	0	+90	0	+90	0	+90	0
2	+110	-300	+110	-25	+110	-125	+110	-25
3	+120	-250	+120	-1050	+120	-175	+120	-550
4	+90	0	+90	0	+90	0	+90	0
5	+100	-250	+100	-50	+100	-150	+100	-50
6	+110	0	+110	0	+110	0	+110	0
7	+120	-150	+120	-50	+120	-150	+120	-50
8	+100	0	+100	0	+100	0	+100	0
9	+80	0	+80	0	+80	0	+80	0
10	+80	-300	+80	-75	+80	-150	+80	-75
EV	Bad		Bad		Good		Good	

Rew: Reward. Pun: Punishment. EV: Expected Value. Punishments which do not result in a loss are evidenced in gray.

In the Single condition we used the same pay-off matrices of the Double condition but we presented participants only the net amount of money won or lost. This resulted in a different effect for the punishment cards which were losses and those which were not. In fact, a card such as (+$100, -$75) in the Double condition would become a (+$25) card in the Single one, thus resulting in a non-loss card. On the other hand, a card such as (+$100, -$300) would become a (-$200) card in the Single condition, giving thus rise to a net loss. As a result, decks B and D, which presented 1 loss every 5 punishments in the Double condition, had 1 loss every 10 cards in the Single condition, while the decks C and D, which had 5 losses every 5 punishments in the Double condition, presented 5 losses every 10 cards in the Single condition (see Table 3).

To control for the other features, all decks had the same attractiveness, so participants gained on average $100 every time they turned a card. The expected value was balanced instead: there was one good deck and one bad deck among the ones with high loss frequency, and one good deck and one bad deck among the ones with low loss frequency.

We ran both feedback conditions in two different frames: a Standard condition, which we just described and in which each card turn originated as default event a win and the contingent event was represented by punishments as in the original IGT scenario presented in Bechara et al. (1994), and a Reversed condition, in which participants always got a punishment when they turned a card and the contingent event was represented by rewards, as in Bechara et al. (2000). In the Reversed condition all the decks had the same reward frequency but differed in the number of gains; the effect of attractiveness and expected value was controlled in the same way as in the Standard condition.

Table 3: Deck matrices of the Single Feedback - Standard Frame condition.

Card #	**A** **Payoff**	**B** **Payoff**	**C** **Payoff**	**D** **Payoff**
1	+90	+90	+90	+90
2	-190	+85	-15	+85
3	-130	-930	-55	-430
4	+90	+90	+90	+90
5	-150	+50	-50	+50
6	+110	+110	+110	+110
7	-30	+70	-30	+70
8	+100	+100	+100	+100
9	+80	+80	+80	+80
10	-220	+5	-70	+5
EV	Bad	Bad	Good	Good

Please note that the "Payoff" column results from the sum of "Reward" and "Punishment" columns of Table 2.

Method

Participants. Eighty-eight participants (40 males) were recruited from students enrolled at the University of Trieste, in Italy. They were aged between 19 and 28 years (M= 19.9, SD= 3.7). The participants were randomly assigned to the experimental conditions. We excluded from the analyses those participants who, in some condition, turned a number of cards from a deck that differed by 3 SDs, or more, from the average number of choices made for that deck. Eight participants satisfied this criterion and were discarded.

Experimental Design. The experiment followed a 2x2 between subjects design with Feedback (Single vs. Double) and Frame (Standard vs. Reversed) as main factors.

Materials. Deck features are summarized in Table 2 and Table 3. Note that in all the conditions A and B were the bad decks while C and D were the good ones, and that B and D were those decks in which a possible frequency effect should show up since they provided low-frequency losses in the Standard condition and high-frequency gains in the Reversed condition.

Procedure. Experimental sessions were held individually. Participants played a computer-based implementation of the IGT. Decks were visually presented in the lower part of a 15 in LCD screen, and participants used a mouse to point and select the deck they had chosen. Immediately after each card selection, the amount of money obtained through the default event (and possibly through the contingent one) was displayed in the upper half of the screen. The running total of money was coarsely indicated by a colored bar in the uppermost part of the screen that was updated after each selection. In each experimental condition participants had to perform 100 card selections.

Results and Discussion

Table 4 reports the average number of choices made from each deck in the different experimental conditions.

Table 4: Means (and Standard Deviations) of deck choices in the four experimental conditions.

Condition	Deck A	Deck B	Deck C	Deck D
Double-Reversed	**21.06** (7.99)	**22.94** (5.03)	**26.71** (9.3)	**29.29** (9.48)
Double-Standard	**22.45** (8.18)	**23.65** (9.33)	**23.35** (9.68)	**28.55** (12.56)
Single-Reversed	**19.59** (5.82)	**25.86** (8.08)	**24.18** (10.03)	**30.36** 10.57)
Single-Standard	**17.62** (6.4)	**28.95** (12.24)	**19.57** (6.87)	**33.86** (13.24)

We analyzed the participant's performance on two synthetic indices: P, which measures the tendency to choose according to the expected value and is calculated by (C+D)−(A+B), and Q, which measures the tendency to choose according to the frequency of the contingent event. Q is calculated by (B+D)−(A+C) and it measures the preference for decks with low loss frequency in the Standard condition and decks with high gain frequency in the Reversed condition (see: Stocco et al. (2009)). We monitored the participants' behavior throughout the experiment by analyzing the two indices in successive blocks of 20 choices each. We ran a mixed design ANOVA both on P and Q, using Feedback (Single vs. Double) and Frame (Standard vs. Reversed) as between factors, and Blocks (from 1-20 to 81-100) as within factors.

As for P, the ANOVA didn't reveal any significant difference for the two factors nor for the blocks. The interaction between Blocks and Feedback resulted marginally significant ($F(4,304)=2.39$, $p=0.51$) and was caused by the low number of selections from good decks in the first block made by participants in the Single condition in comparison to those in the Double one. Since there was no main effect of any factor, we collapsed the value of P at the end of the experiment across all conditions. A t-test on this value revealed that participants chose more cards from the good decks than from the bad ones ($M=8.8$, $t(79)=3.44$, $p<.001$).

As for Q, the effect of Feedback ($F(1,76)=8.15$, $p<.01$), of Blocks ($F(4,304)=4.72$, $p<.01$) and the Blocks x Frame interaction ($F(4,304)=3.6$, $p<.01$) resulted statistically significant, while the Blocks x Feedback interaction was only marginally significant ($F(4,304)=2.1$, $p=.081$). We also performed two t-tests on the value of Q at the end of the experiment separately for the Single and Double Feedback conditions collapsing the Standard and Reversed Frame. The results were significant for the Single condition ($M=18.89$, $t(42)=5.32$, $p<.0001$) but not for the Double condition (M=4.43, t(36)=1.19, p=.24).

The analyses show thus that there was a frequency effect only in the Single condition but not in the Double one. As explained in the previous section, according to Chiu et al. (2008), participants were expected to be influenced by the frequency of the contingent event in both cases, while according to Fum et al. (2008) the effect should only be present in the Single feedback. The results support our hypothesis that participants try to avoid all kind of punishments and not just the ones which result in a net loss (and are sensible to any reward and not only to wins). Because the matrices of the decks in the Single feedback condition were obtained directly from those used in the Double one, this result cannot be attributed to possible different values employed in the two conditions. On the other hand, because the SGT did not directly contrast Single vs. Double feedback, the results obtained by Chiu et al. (2008) could depend critically on the specific values used in their matrices. This experiment also suggests that participants, being sensible to the difference between Single and Double feedback, take separately into account the value of both the default and contingent event and do not rely only on the net value of each trial.

The analyses, by highlighting a goodness effect in all the conditions, show that participants are somehow sensible to the expected value of the decks, too. However, if they had really understood which decks were the good ones, they would have consistently chosen them. This did not happen because in no condition the (good) deck C was chosen more frequently than the (bad) deck B, a result that is compatible with the "prominent deck B phenomenon" normally found in traditional IGT.

The difference between the results of our experiment and those obtained with the SGT by Chiu et al. (2008) demonstrate that participants' behavior cannot be easily ascribed to the effect of a single feature. Participants could behave differently when dealing with decks which have similar qualitative features but that vary in their numerical values. Therefore, an understanding of their performance would require the use of cognitive models capable of making any feature effect an emergent property of their parameters providing thus an explanation for the influence of the qualitative features.

Modeling the results

In discussing the models of the IGT used by previous researchers, Ahn et al. (2008) identified three general assumptions: "First, an individual's evaluation of the positive and negative payoffs can be represented by a unidimensional utility function. Second, expectations about payoffs for each deck are learned on the basis of the experienced utilities on each trial. Third, these expectancies determine the choice probabilities for selecting each deck on each trial" (p. 1393). As a consequence, any model for this task, and similar iterated decision making problems, will employ at least three different functions: (1) an evaluation function to assess the payoff associated with each choice, (2) a learning function to upgrade the expectancies concerning the expected payoff of each option, (3) a selection function to choose on each trial a particular option on the basis of its expected payoff.

By adopting an architectural approach to modeling, the problem of identifying the functions necessary to replicate human performance in the task of interest is facilitated because some of these are considered as resources provided directly by the architecture. In particular, ACT-R (Anderson, 2007) makes available, by default, both a learning and a selection function. The former is given by the linear equation proposed by Bush and Mosteller (1955):

$$U_i(n) = U_i(n-1) + \alpha[R_i(n) - U_i(n-1)] \tag{1}$$

where:
U_i is the utility associated with option i
n is the current time step, with $n-1$ indicating the previous one
R_i is the reward associated with option i,
and α is a parameter regulating the learning rate.

The second equation is given by:

$$P_i = \frac{e_j^U/s}{\sum_i e_j^U/s} \tag{2}$$

and determines the probability P that a given option i will be selected among the j possible options. This probability is a function of the value U (the utility, in ACT-R parlance) of the particular option compared to the sum of all the possible option values, while s is a noise parameter, analogous to the temperature of Boltzmann machines, that introduces some kind of nondeterminism in the selection process.

By having two of the three main modeling problems solved by the architecture, we concentrated on the evaluation function used to assess the outcome of each card choice. Traditionally (Ahn et al., 2008; Yechiam & Busemeyer, 2005) two different kind of functions have been employed.

The first one, called the *expectancy function* by Ahn et al. (2008), computes a weighted average of the rewards and punishment associated with the chosen option in each trial. This function can be expressed as following:

$$v(t) = (1-W)\cdot rew(t)^\gamma - W\cdot pun(t)^\gamma \tag{3}$$

with $rew(t)$ and $pun(t)$ indicating the value of the reward and punishment at time t, respectively, while γ is a parameter that determines the curvature of the evaluation function, and W denotes the differential weight participants place on losses over gains.

An alternative evaluation rule is provided by the so called *prospect function* (Ahn et al., 2008) expressed by:

$$v(t) = \begin{cases} net(t)^\gamma & : \text{ if } net(t) \geq 0 \\ -\lambda|net(t)|^\gamma & : \text{ if } net(t) < 0 \end{cases} \tag{4}$$

with $net(t)$ indicating the net outcome, i.e. the difference between the default and contingent event, and λ representing a loss aversion parameter.

The two functions are similar according to several features: they both assume a nonlinear evaluation of the monetary outcome and both weight losses differently from gains. The most important difference between them is constituted by how they take into account the default and contingent event. The expectancy function assess them separately before combining them into a scalar value; the prospect function, on the other hand, assumes that decision makers process directly the net outcome. The two functions can thus be considered as implementing the different assumptions held by Fum et al. (2008) and Chiu et al. (2008), respectively, and we used them to implement two different computational models through which we tried to replicate the empirical results. We ran a series of 500-run simulation trials with a large range of parameters and the results we obtained were quite straightforward.

Both functions are able to capture the frequency of the contingent event effect as revealed in the Single feedback condition but the prospect function, taking into account only gains and losses, is not sensitive to the effect of rewards and punishments, which also play a critical role in determining the participants' behavior in IGT, and therefore gives raise in the Double feedback condition to an effect that is absent in the experimental data. Table 5 reports the best performing models employing the expectancy (with parameters W=0.05 and $\gamma = 0.15$) and the prospect functions (with parameters $\lambda = 0.1$ and $\gamma = 0.1$) respectively. While these models have grossly similar synthetic measures of fit (for instance, RMSE= 2.35 for the expectancy and RMSE= 3.23 for prospect; chi-squared= 3.56 (p = .99) for the expectancy and chi-squared= 6.87 (p = .96) for the prospect) the prospect model fails to replicate the participants' performance by providing predictions that fall out of the 95% confidence intervals in four data points.

Table 5: Means of deck choices by the two models. The predictions which fall out of the confidence intervals are evidenced in grey.

		Deck			
Condition	**Model**	**A**	**B**	**C**	**D**
DR	Expectancy Function	25.01	24.63	24.79	25.58
	Prospect function	20.86	28.4	21.00	29.74
DS	Expectancy Function	24.55	25.03	24.93	25.5
	Prospect function	21.23	29.21	20.31	28.76
SR	Expectancy Function	20.35	28.1	20.79	30.77
	Prospect function	20.82	28.49	20.99	29.7
SS	Expectancy Function	20.22	29.59	19.84	30.36
	Prospect function	20.31	29.43	20.71	29.56

DR: Double-Reversed. DS: Double-Standard. SR: Single-Reversed. SS: Single-Standard.

Conclusions

In the paper we proposed an explanation for the frequency of the contingent event phenomenon which lies beneath the fact that people are attracted by options that are associated with the most frequent positive, and the less frequent negative, outcomes. A fundamental problem, deriving from the fact that the IGT is grounded on a conflict between the value of the default event (which codes the immediate attractiveness of an option) and the contingent one (which represents the options' long term expected value) is to establish whether this phenomenon is caused by any positive or negative outcome independently of its magnitude or, on the contrary, it is triggered by the net result deriving from the two events. The findings of our experiment corroborate the former hypothesis and the simulation results indicate that only a model sensible to rewards and punishments, and capable of analyzing them separately, can replicate the empirical data.

References

Ahn, W.-Y., Busemeyer, J. R., Wagenmakers, E.-J., & Stout, J. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive Science*, *32*(8), 1376–1402.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.

Bechara, A., Damasio, A., Damasio, H., & Anderson, S. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*, 7–15.

Bechara, A., Tranel, D., & Damasio, H. (2000). Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions. *Brain*, *123*, 2189–2202.

Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning*. New York: Wiley.

Chiu, Y.-C., Lin, C.-H., Huang, J.-T., Lin, S., Lee, P.-L., & Hsieh, J.-C. (2008). Immediate gain is long-term loss: Are there foresighted decision makers in the Iowa Gambling Task? *Behavioral and Brain Functions*, *4*, 13.

Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: a critical evaluation. *Neuroscience and Biobehavioral Reviews*, *30*(2), 239–71.

Fum, D., Napoli, A., & Stocco, A. (2008). Somatic markers and frequency effects: Does emotion really play a role on decision making in the Iowa Gambling Task? In *Proceedings of 30th Annual Conference of the Cognitive Science Society* (pp. 1203–1208).

Stocco, A., Fum, D., & Napoli, A. (2009). Dissociable processes underlying decisions in the Iowa Gambling Task: a new integrative framework. *Behavioral and Brain Functions*, *5*, 1.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: Oxford University Press.

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*, 387–402.