האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

REVERSAL OF RISKY CHOICE IN A GOOD VERSUS A BAD WORLD

By

EINAV HART, YAAKOV KAREEV, and JUDITH AVRAHAMI

Discussion Paper # 619 Aug 2012

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

Reversal of Risky Choice in a Good versus a Bad World

Einav Hart, Yaakov Kareev, and Judith Avrahami

The Center for the Study of Rationality, Hebrew University of Jerusalem, Israel

Corresponding Author:

Einav Hart

The Center for the Study of Rationality

Hebrew University of Jerusalem

Jerusalem, 91904

ISRAEL

e-mail: einav.hart@mail.huji.ac.il

Abstract

In many situations one has to choose between risky alternatives, knowing only one's past experience with those alternatives. Such decisions can be made in more – or less – benevolent settings or 'worlds'. In a 'good world', high payoffs are more frequent than low payoffs, and vice versa in a 'bad world'. In two studies, we explored whether the world influences choice behavior: Whether people behave differently in a 'good' versus a 'bad' world. Subjects made repeated, incentivized choices between two gambles, one riskier than the other, neither offering a sure amount. The gambles were held equivalent in terms of their expected value, differing only in variance. Worlds were manipulated both between- and within-subject: In Study 1, each subject experienced one world – good, bad or mediocre; in Study 2, each subject experienced both a good and a bad world. We examine the aggregate pattern of behavior (average choice frequencies), and the dynamics of behavior across time. We observed significant differences in the aggregate pattern: In a good world, subjects tended to choose the riskier alternative, and vice versa in a bad world. The pattern of the dynamics, i.e., the transitions from round to round, were best explained by a reaction to the counterfactual reward: When the unchosen alternative yielded a better payoff, the tendency to subsequently choose it was higher. We compared these two patterns to the predictions of three types of models: Reinforcement learning, regret-based and disappointment-based models. Behavior was in line only with the predictions of regret-based models.

Keywords: risky choice; decision from experience; regret; counterfactual outcome

Introduction

Say you are contemplating which route – particularly, which of two roads – you want to take back to your house this afternoon. The two roads are usually pretty similar, neither is longer than the other, nor much slower. However, one of the roads tends to be more capricious: Sometimes there's a roadblock which delays you, and in other times it is totally vacant and you get home quicker than average. The other road is pretty regular, and the variance in driving times is small. Let's assume you have to get home from work each day, and so you make this decision regularly. All the while, you hear traffic reports, letting you know how well you would do on both roads. What would influence you more when making the choice each day: The time it took, or maybe the report on the other road? Would it be a comparison of the two, or the minimal time it could have taken? Maybe you would think about the expected, average time? Or would you care about the variance in driving times on both roads?

After making this trip several times, you know that at certain times of the day, both roads are more congested (say, in rush hours) than at others, regardless of the roads' usual ups and downs. Does this congestion affect your decisions? That is, would you choose one road over the other at 1pm, and make the opposite choice at 5pm, even though neither road prevails (on average), at any time?

We're pretty sure you, much like our subjects, had to make this type of decision many times – not only between roads, but between grocers in the market, coffee stands, radio stations, stocks and so forth. The main difference is that our subjects encountered a simulated version of these choices. More specifically, we ran two experiments which entail repeated decisions between two risky alternatives (gambles). In the next section, we describe in detail the meaning of the above phrase: What counts as risky alternatives, and which aspects of behavior can be illuminated by looking at repeated decisions.

In this paper, we look at choices made in various 'worlds' – settings pertaining to the overall "goodness" (i.e., expected value) of the alternatives: In a 'good world', high payoffs are more frequent than low payoffs, and vice versa in a 'bad world'. In our example above, a bad world would be one in which all roads are, in general, congested, and a good world would be one in which there are only few cars on both roads. We examine whether preferences for the alternative with larger variance, over the alternative with smaller variance, differ in the various worlds. This allows identifying a previously overlooked factor in decision making: The probability of receiving a good outcome.

If one assumes that only the experienced rewards – reinforcements, or driving times – are those which impact decisions, then the frequencies of choosing each alternative should be independent of the overall goodness of both alternatives (i.e., their expected value); in other words, the worlds should play no role. Specifically, when the alternatives' expectancy is the same, one should be indifferent between them, regardless of said expectancy. However, one could be motivated not only by the values received, but, say, by their relative position within the alternative chosen. One could be disappointed to get the lower, rather than the higher, possible payoff in the gamble, even if one gains money either way. If such disappointments motivate us, the world does come into play. It can be easily seen that there are more disappointments in worse worlds than in better worlds. Moreover, and less obvious, the relative number of low (i.e., disappointing) outcomes in the different alternatives changes with the change in worlds: This number is higher the worse the world; thus, choice frequencies will vary depending on the world.

We attempt to explain two aspects of behavior in our experiments. First, we look at the overall (aggregate) behavior in each world: How many times was each alternative chosen, out of all choices made. Crucially, we explore whether the world impacts these frequencies, and how. Second, we look at the dynamics of choices over time: How the choices were distributed, when examined round by round; whether and how the previous choice determines its successor.

Our settings model many situations, as people are ubiquitously asked to make decisions between alternatives which are essentially equivalent to each other – as in the examples above. These decisions are made repeatedly, in more and less benevolent settings. Therefore, it is highly important (and equally interesting) to find out what people indeed choose, which factors impact their choices, and first and foremost, whether the "goodness" of the world matters. We conclude the paper with a discussion of the findings and their implications for decision making.

Previous Studies

Decisions among several gambles, each having several rewards with certain probabilities, are termed Risky Decisions (Knight, 1921). However, the probabilities are not necessarily stated explicitly – they may be learned through past experiences with the alternatives and their outcomes (e.g., the quality of the coffee in different stands, the congestion, or our liking of radio stations' playlists). In such cases, no one tells us the probability of success, or, in fact, the possible rewards; we learn these through repeated exposure to the same 'gambles' and through our payoffs and their respective frequencies.

These problems are thus termed 'Decisions from Experience', contrasted with problems in which there is an explicit description of the alternatives (Barron & Erev, 2003; Hertwig, Barron, Weber & Erev, 2004; Hertwig & Erev, 2009).[1]

Interestingly, in most studies of choice between gambles, one of the alternatives is in fact a sure amount, not a gamble (e.g., Camilerri & Newell, 2009; Erev & Barron, 2005; Grosskopf, Erev & Yechiam, 2006; Haruvy & Erev, 2002; Hertwig & Erev, 2009; Hertwig et al., 2004; Weber, Shafir, & Blais, 2004; there are a few exceptions, such as Barron & Erev, 2003, and Rakow & Miler, 2009). Several studies (e.g., Allais, 1953; Kahneman & Tversky, 1979) have shown that certainty has a special status, in that people will tend to choose the sure amount over a gamble, even to the point of a somewhat significant loss. Importantly, this choice of a higher probability of winning over a higher expected value is highly dependent on how close the probability is to one (i.e., certainty). Hence, it seems that one cannot readily generalize from choices made between a risky alternative and a sure amount. Furthermore, one might argue that in reality, one rarely encounters a truly safe and sure payoff – making our setting more realistic in this sense. We thus include two gambles, each of them having more than one possible reward with a positive, non-negligible probability. In this we actualize a "Risky" gamble and a "Safer" – rather than completely 'safe' – gamble. The riskiness of our gambles is derived from their variance and spread, rather than from their expected value. The expected value of the gambles is kept equal (again, unlike most studies such as those noted above). In this, we aim to isolate variability and risk from value considerations.

Repetition of choices is of course necessary for the discussion of decisions from experience, as there is no meaning to a single choice among unknown alternatives. That is, decision makers have to rely on the feedback regarding their choices to construct the problem by themselves. Repetitions seem to enable more than just discerning the decision problem: They offer the opportunity of studying how experience impacts choice over repeated trials. In fact, recent outcomes have been shown to be strong determinants of choice behavior (Avrahami & Kareev, 2010, 2011; Hertwig, Barron, Weber & Erev, 2006; Weber et al., 2004; Hogarth & Einhorn, 1992; Stewart, Chater, & Brown, 2006; Thaler & Johnson, 1990), and so does the feedback given after each choice (Camerer & Ho, 1999;

[1] Decisions from experience have been shown to be consistently and qualitatively different from decisions based on description, in which the choice set is presented explicitly to decision-makers (Hertwig et al., 2004; Hertwig & Erev, 2009; Rakow & Newell, 2009). For example, if shown the two options of {3 with probability 1} and {4 with probability 0.8, 0 with probability 0.2}, people will tend to choose the sure option of 3 points. This preference is often reversed when making the decision from experience rather than from explicit description (e.g., Barron & Erev, 2003).

Cheung & Friedman, 1999; Cooper, Garvin & Kagel, 1997; Daniel, Seale & Rapoport, 1998; Erev & Roth, 1998; Selten & Buchta, 1998).

Interestingly, as we pointed out in the introduction, the influence of outcomes – observed or expected – is not limited to the realized payoffs; it includes the effect of foregone payoffs, those which would have been received had one made a different choice (Avrahami, Güth & Kareev, 2005; Avrahami & Kareev, 2011; Camerer & Ho, 1999; Daniel et al., 1998; Erev & Barron, 2005; Ert & Erev, 2007; Grosskopf et al., 2006; Yechiam & Busemeyer, 2005; 2006; Yechiam & Rakow, 2011; Zeelenberg, Beattie, van der Pligt & de Vries, 1996; Zeelenberg & Pieters, 2004). Outcomes – both realized and counterfactual – change one's affective state and thus one's subsequent choice (Avrahami & Kareev, 2010, 2011; Bechara, Damasio, Damasio & Anderson, 1994; Loewenstein, Weber, Hsee & Welch, 2001; Shiv, Loewenstein, Bechara, Damasio & Damasio, 2005).

Building on this evidence, we explore behavior over time – choice dynamics, i.e., changes between consecutive rounds – as well as the aggregate behavior (i.e., frequencies of each gamble being chosen). Many studies have examined overall choice frequencies, comparing it to various benchmarks, like that of Prospect Theory (Kahneman & Tversky, 1979) and those derived from the proposition of underweighting of small probabilities (Barron & Erev, 2003; Erev & Barron, 2005), as well as various learning models (Busemeyer & Stout, 2002; Camerer & Ho, 1999; Rakow & Miler, 2009; Yechiam & Rakow, 2011). Choice dynamics have rarely been looked at, or contrasted with theoretical benchmarks (exceptions include Gonzalez & Dutt, 2011; Erev, Ert, Roth et al., 2010; Hills & Hertwig, 2010). The current investigation aims to correct this: We examine which features of the presented outcomes influence choice behavior in these two aspects. Particularly, we test the predictions of three classes of models: Reinforcement-Learning models (e.g., Erev & Roth, 1998; Roth & Erev, 1995); those hinging on the disappointments felt by subjects – particularly a Win—Stay; Lose—Shift model, with the low reward in each gamble defined as a 'loss', and the high reward termed a 'win' (WSLS; Levine, 1966; Matsen & Nowak, 2004; Nowak & Sigmund 1993); those hinging on the notion of regret, i.e., the comparison of the counterfactual and the actual rewards – particularly, a Regret—Shift; No Regret—Stay model (Regret-Shift, an extreme version of the model presented by Avrahami & Kareev, 2011). The models and their predictions are described in detail below. Comparing behavior in the various worlds allows testing for these models' compatibility qualitatively, as well as helps enlighten the role of the gambles' "goodness".

Models and Predictions

Whether one assumes an influence of the received payoff, the disappointment (or lack thereof) in receiving it, or the regret one feels (or doesn't) after the choice, these models take an affective, outcome-based approach to decision-making. That is, unlike various "Subjective Expected Utility" theories, affective models do not require that decision-makers have an explicit or an implicit representation of the relevant probabilities. Rather than presume calculations and deliberation, choices are assumed to be based on reactions to outcomes and payoffs (Loewenstein et al., 2001). Hertwig (2011) categorizes such models as "associative learning models and outcome heuristics".

Associative learning models are those in which choice propensities or frequencies change according to the realized outcomes (e.g., Barron & Erev, 2003; Denrell, 2007; Erev & Barron, 2005). The Reinforcement model we consider falls in that category. Conversely, "outcome heuristics" describe the choice dynamics as well as the aggregate behavior in terms of responses to outcomes, and largely ignore probabilities. Importantly, they take into account only the outcome of the previous choice, and one's response to it. As noted in the previous section, this assumption is corroborated by studies showing recency effects for the prior decision, and those showing that affective responses to a decision's outcomes influence choosing its successor (Avrahami & Kareev, 2010, 2011; Erev & Roth, 1998; Hertwig et al., 2006; Loewenstein et al., 2001). Interestingly, Hertwig (2011) points out, several of these heuristics may in fact appear as if settling the tradeoff between outcomes and probabilities – but without actually weighting the former by the latter. The other two models we consider, namely WSLS and Regret-Shift, fall under this category. These models are both very simple and parsimonious, each assuming an effect of a single factor: WSLS assumes it is the received payoff vis-à-vis the alternative payoff in the chosen gamble; Regret-Shift posits that the comparison is between the received payoff and the forgone payoff, in the unchosen gamble.

Returning to the congested roads example, Reinforcement predicts that one gradually learns which road is better – the faster the drive, the more likely one is to pursue that road in the future. Over time, one will gravitate towards the better road (if there is one, or towards indifference). WSLS posits that one examines the road taken, and repeats it the next day only if it was better than a specific benchmark, e.g., half an hour. Anything faster than that is great, and will lead to choosing the road again; anything longer than that is counted as a 'loss' or 'disappointment', and to the subsequent abandonment of that road (at least for a day). Lastly, the Regret-Shift model presumes that one considers travel time at the road taken, as well as the travel time at the road not taken – and if the non-taken road would have been better, then that is indeed what one will subsequently choose.

As it turns out, the aforementioned models differ in the predicted pattern for either overall choice pattern or choice dynamics, or both. We use the following notation in delineating the predictions:

Risky Gamble = {H_R with probability p_R; L_R with probability $(1-p_R)$}

Safer Gamble = {H_S with probability p_S; L_S with probability $(1-p_S)$}

H_i is the high reward in gamble i, and L_i is the low reward. We set the values of H_S and L_S, the rewards of the safer gamble, to be within the range (H_R, L_R). Gamble S is termed the safer gamble, in that the variance in its rewards is smaller than that of gamble R. As aforesaid, the expected value (EV) of the two gambles was kept equal. Moreover, we operate within the gains domain: All the possible outcomes are positive. We do not include zero as a possible payoff in any of the gambles, unlike many of the previous explorations.

The 'world' is determined by p_R, the probability of a high reward in the risky gamble R (p_S is set by p_R as well, due to setting $EV_R=EV_S$). In what we term the 'neutral world', $p_R = p_S = 0.5$; in a 'bad world', $p_S < p_R < 0.5$; in a 'good world', $p_S > p_R > 0.5$.

Overall Choice Frequencies

The models' predictions coincide for the case of the neutral world; in this case, all models predict that each gamble will be chosen in half of the rounds. The predictions qualitatively diverge for good and bad worlds, in the following way.

Regret- and disappointment-based models both predict biases towards one of the gambles, depending on the world, i.e., the EV of both gambles – but with opposite directions. Disappointment-based models posit that the frequency of choosing each gamble decreases with the probability of losing in it (i.e., receiving the low payoffs L_R or L_S). The smaller the probability of winning (that is, the larger the probability of disappointment), the smaller the probability of choosing that gamble. In a bad world, this makes the safe gamble worse than the risky gamble (P(disappointment in gamble R) = $1-p_R$ < P(disappointment in gamble S) = $1-p_S$) – and vice versa in a good world. Specifically, in the WSLS model, the probability of choosing the riskier gamble R is the ratio between the probability of losing in gamble S, and the sum of all probabilities of losing = $(1-p_S) / (1-p_R+1-p_S)$. This ratio is smaller than 0.5 in a good world, and larger than 0.5 in a bad world.[2]

[2] Depending on the specific probabilities used, a similar prediction may arise from a model assuming underweighting of small probabilities.

Conversely, a regret-based model predicts that the riskier gamble be chosen with a probability proportional to the probability of its high reward – as H_R is the highest possible payoff. In case it was realized, one has all the reasons to regret a different decision (of gamble S), regardless of his own payoff – or no reason to regret making a choice of gamble R, again regardless of the outcome of gamble S. Thus, the better the world, the less opportunity to regret choosing gamble R and more opportunity to regret choosing gamble S. Particularly, Regret-Shift predicts that the aggregate proportion of the riskier gamble choices is in fact equal to the probability of receiving its high reward. One will repeat a choice of gamble R – or switch to it – when its high reward is realized, which happens at a frequency equal to p_R. This probability is larger than 0.5 in a good world, making the riskier gamble preferred; the safer gamble is preferred in a bad world, in which p_R is smaller than 0.5.

As aforesaid, reinforcement models assume that when the gambles' EVs are the same, they should be chosen with equal probability: Neither choice is more reinforced than the other – regardless of the 'world' or the gambles' respective variances. Therefore, this model conjectures a 50-50 split of choices between the gambles.[3]

To sum up, for the situation we propose (and test experimentally), three leading models predict distinct, different patterns of aggregate choices in the different worlds.

Choice Dynamics over Time

Behavior following a choice of the risky gamble R is pretty straightforward (at least according to the models above). If one received L_R, then one is likely to subsequently choose the safer gamble S: L_R is not only the low payoff in gamble R, it is also the minimal reward. If one received H_R, then one should definitely be encouraged to stay, regardless of what one considers, since it is the highest possible reward. Thus, behavior following a risky choice cannot differentiate between the aforementioned models. In contrast, the models differ considerably in their predictions concerning choice dynamics following a choice of the safer gamble.

As noted, Regret-Shift assumes that choices are affected by both the obtained (actual) payoff, and the counterfactual reward, in the last round. Specifically, if the unchosen gamble's payoff is higher than one's actual payoff, then one would shift; otherwise, choice persists. Because one will shift when the unchosen gamble yields a better payoff, one will shift away from the choice of gamble S after H_R is realized in gamble R: The obtained payoff (whether it is H_S or L_S) is lower than the counterfactual. One will repeat the choice of gamble

[3] This conjecture assumes risk neutrality. Without this assumption, the propensities of choices may be different. However, regardless of risk-preference, choice frequencies should be independent of world.

S after L_R is realized in gamble R, as the obtained payoff (whether it was H_S or L_S) was higher than the counterfactual.

WSLS also assumes a dependence of a subsequent choice on the last round's outcome – but only of the actual, received payoff. According to WSLS, after losing in a gamble (receiving its lower reward), one will choose the other gamble; one will repeat the previous choice if its high payoff was realized. That is, after receiving H_S in the safe gamble, one should repeat this choice, and *not* shift – regardless of the outcome of the risky gamble R. The opposite would happen should gamble S yield the low payoff of L_S. In this case, WSLS posits one should subsequently shift to gamble R. Hence, the WSLS prediction in these two cases is opposite that of Regret-Shift.

Reinforcement supposes a more gradual learning of the probabilities and EV: Outcomes have a diminishing effect across rounds. When the EV of the two gambles is equal, as in our setting, the model predicts indifference. As only the realized payoff in one's chosen gamble determines choice (and correspondingly, the tendency to shift or stay with the current choice) – the larger this reward is, the more it reinforces one's choice. Following a choice of gamble S, the prediction of reinforcement learning models is in the same direction of that of WSLS, but not as extreme. Table 1 summarizes the models' various predictions for both the overall frequencies and the step-to-step dynamics.

Table 1

Predictions regarding overall (aggregate) behavior and choice dynamics, as derived from the Regret-Shift, WSLS and Reinforcement-Learning models

	Overall: P(choose riskier gamble)	**Dynamics:** Choice following a safer gamble choice
Regret-Shift	p_R	Depends on the previous payoff in the unchosen gamble R
WSLS	$(1-p_S)/(1-p_R+1-p_S)$	Depends on the previous payoff in the chosen gamble S
Reinforcement	0.5	Depends on the size and history of rewards in the chosen gamble.

Current Experiments

Two studies explore repeated, incentivized choices between two gambles, comparing behavior in different worlds. Our experiments lie within the realm of repeated decisions from experience, with two major novelties, as discussed above: The choices being between two gambles (without a sure amount), with the "goodness" or EV of both gambles changed together – defining the 'world' subjects play in.

Subjects face a choice between two alternatives, each representing a simple gamble of two values. These values determine the corresponding payoffs. The gambles' properties are unknown a-priori, and are experienced – learned, in a sense – throughout the experiment rounds. The outcome of both gambles, chosen and unchosen, are shown after making the choice. The two gambles are of equal expected value, but differ in their riskiness, in the sense that one gamble has a bigger spread and variance than the other. Importantly, the variance is always larger than zero, making both gambles "risky". In Experiment 1, we contrast choices made in three worlds (good, neutral and bad), manipulated between subjects. In Experiment 2, we look at a similar choice structure, but here subjects experience both the good and the bad worlds, in two separate blocks (the world is manipulated within-subject). In both experiments, we examine aggregate behavior and choice dynamics over time. We discuss the compatibility of the various models to the choice behavior in our settings.

Experiment 1

Experiment 1 examines choices between a riskier and a safer gamble, in three worlds (between-subjects): A bad world, a neutral world and a good world. Further, we introduce two variants of the "safer" gamble S, in order to test whether the values or the exact relationship between the gambles plays a role, or whether the choices are affected (mostly) by the relative riskiness of the riskier and safer alternatives.

Method

Design. Subjects were asked to make 100 consecutive choices between two "boxes", each representing a gamble with two possible rewards. These rewards were realized with constant probabilities, not told to the subjects. The riskier gamble (gamble R) outcomes were set at 90 and 10. Gamble S had two variants – gamble S1 outcomes were {70, 30} and gamble S2 outcomes were {80, 20}.

The probabilities define the 'world' variable. Three worlds were introduced: In the 'bad world', p_R was 0.3, and p_S was either 0.1 or 0.233 in gambles S1 and S2, respectively; in the 'neutral world', $p_R = p_{S1} = p_{S2} = 0.5$; in the 'good world' – p_R was 0.7 and p_S was either 0.9 or 0.767 in gambles S1 and S1, respectively.[4] The gambles' expected values in the bad, neutral and good world were 34, 50 and 66, respectively.

There were thus six experimental conditions: three worlds (Bad, Neutral, Good) * two gamble-S versions (S1, S2). All of the conditions were between-subjects.

Subjects' winnings were accumulating, and their payoff was their sum of points, multiplied by a known exchange rate. The exchange rate of points to NIS (New Israeli Shekel, worth approximately 0.27 USD) was set such that the expected payoff is approximately equal in all conditions: In the bad world, every 230 points equaled 1 NIS in the neutral world, 330 points = 1 NIS; in the good world, 440 points = 1 NIS.

Subjects. 144 students at the Hebrew University participated in the study (mean age = 25; 72 females) and were randomly assigned to the six conditions. Subjects received monetary compensation, hinging on their payoffs in the task. Payoffs ranged between 12-20 NIS, with an average of 15.11 NIS.

Apparatus and Procedure. Subjects sat in front of a computer in a quiet room for the duration of the experiment (around 10 minutes). The instructions were on the screen: Subjects were told that they would make 100 repeated choices between two unmarked buttons. They were told that each of these had fixed values with fixed probabilities, and that the values will appear on both buttons after making a choice. The value on the chosen button constitutes the earning for that round, and the value on the alternative box is what one could have received. The earnings accumulate over the 100 rounds and are multiplied by a known ratio, constituting the subjects' final payoff in NIS.

Results

Overall frequencies. We first look at the aggregate choice frequencies of the two gambles, in each of the six conditions (by world and version of gamble S), averaged by subject. An ANOVA revealed a significant effect of world ($F(2,144)=34.37$, $p<.001$, $\eta_p^2=.332$): The probability of choosing the riskier gamble R was significantly lower in the bad world than in the neutral world, which was, in turn, significantly lower than that probability in the good world ($M_{Bad}=0.38$; $M_{Neu}=0.47$; $M_{Good}=0.67$). All pairwise contrasts between worlds, across gamble S variants, were significant (all p's<.02). There was also an interaction effect of world

[4] p_S was determined by the value of p_R, as $EV_R = EV_S$.

and gamble S variant ($F(1,144)=4.41$, $p=.014$, $\eta_p^2=.060$). There was a significant and large positive correlation between the observed proportions and the predictions of the Regret-Shift model ($r=0.91$, $p=.018$). The correlation with the WSLS prediction was negative, but not significant ($r=-0.61$, $p=.194$). The direction and the strength of the correlations both point to the Regret-Shift model as providing the best account of the data. Figure 1 shows this pattern.

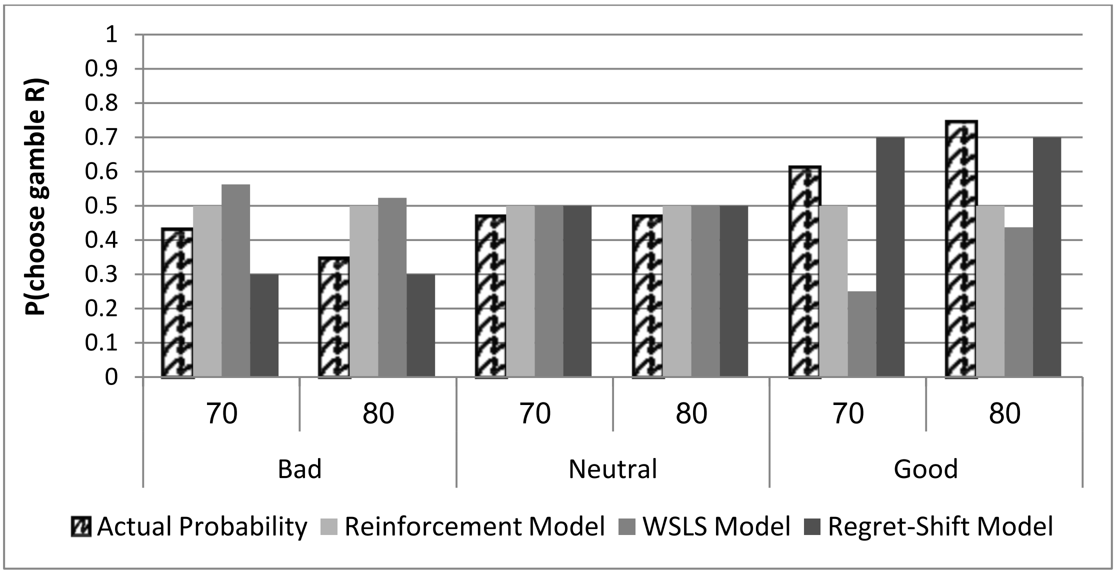

Figure 1. Probability of gamble R choices by World and variant of gamble S, compared with choice probabilities as predicted by theoretical models: Reinforcement Learning models, WSLS, and Regret-Shift.

Dynamics over time. A binary "repetition" variable was calculated, reflecting whether the current choice is the same as its predecessor. As aforesaid, all models predict a shift after one chose the riskier gamble and received the low payoff, and repetition following receiving the high payoff in the riskier gamble. Indeed, when examining repetition following a riskier choice of gamble R, the most prominent influence is that of the realized reward in that gamble. This was the strong result of a logistic regression model, clustered by subject ($n=7326$ data points, 144 subjects): The tendency to stay was larger after receiving the high reward, H_R, rather than the low one, L_R ($Z=11.96$, $p<.001$). In addition, there was a significant positive effect of world ($Z=3.26$, $p=.001$) and a negative effect of the realized reward in gamble S ($Z=-3.52$, $p<.001$).

However, as the models diverge only in their predictions following a safer choice, fitting the data following a choice of the riskier gamble does not differentiate between the models. For the purpose of distinguishing between the three models, we look at behavior (repetition) following a safer choice. We thus analyze the binary variable of repetition, conditional on a previous choice of the safer gamble, by a logistic regression model as above (n=7074 data points). We ran a comprehensive model encompassing all the above variables – the realized reward in gamble R, gamble S reward, and the two between-subjects variables of world and version of gamble S – as well as several interactions that are of interest.

Both gambles' payoffs are significant predictors of the probability to repeat the safer choice: A high (counterfactual) reward in gamble R decreased the probability of repeating the choice of gamble S (Z=-10.47, p<.001), while a high realized payoff in gamble S increased this probability (Z=3.15, p=.002). Here too, the world had a significant effect (Z=-3.46, p=.001): The better the world, the less subjects tended to repeat their previous safer-gamble choice. The statistics are presented in Table 2. Looking at effect sizes, it is clear that the tendency to repeat gamble S is most influenced by what one might have received had one chosen gamble R – as predicted by the Regret-Shift model.

Table 2

Predicting repetition following a choice of the safer gamble – full regression model

(standardized variables)

Factor	Beta (S.E.)	Exp(B)	Z	Sig
Realized R	-0.560 (0.053)	0.571	-10.47	.000
Realized S	0.104 (0.033)	1.110	3.17	.002
Value of Reward	0.009 (0.084)	1.009	0.11	.912
World	-0.287 (0.083)	0.750	-3.46	.001
RealR * RealS	0.055 (0.035)	1.057	1.57	.117
Value * RealS	0.007 (0.033)	1.008	0.22	.823
RealR * RealS* Value	-0.035 (0.033)	0.966	-1.05	.292
World * RealR	0.069 (0.055)	1.071	1.24	.216
World * RealS	-0.052 (0.066)	0.949	-0.79	.430
World * Value	-0.148 (0.082)	0.863	-1.80	.072

Note. N=7074 data points, clustered for 144 subjects. Model fit: -2log likelihood = 8576.31; R^2=0.078.

One might ask whether there were subjects who tended more (or less) to repeat their previous actions, regardless of outcomes; whether subjects were differentially influenced by inertia. We included each subject's propensity to repeat, calculated as the proportion of repetitions across trials for each subject minus their tendency to choose either alternative, squared, i.e., for choosing it twice (P(repetition) - P(gambleR)^2 - P(gambleS)^2). This did not significantly alter the aforementioned effects, even as its main effect was highly significant (Z=10.59, p<.001). The addition also significantly increased the predictive power of the model (R^2 increased to 0.139 for the comprehensive model).

Discussion

Subjects were indeed affected by the world they played in: In a good world, the riskier gamble was very much preferred; in a bad world, the safer gamble was chosen more often. This bias follows the direction predicted by Regret-Shift, which posits that the comparison between the actual and counterfactual rewards is the critical factor making the decisions.

Regret-Shift appears to provide the best explanation for the observed data regarding the step-to-step basis as well – the transitions between consecutive rounds: Choices were strongly affected by the reward in the risky gamble. Following the choice of the safer gamble, when the outcome of the risky gamble constituted the counterfactual outcome, this counterfactual outcome was the best predictor of the decision whether or not to repeat the choice of the safer gamble. Our interpretation is that this may be so because this outcome determines whether or not the previous choice was cause for regret. Interestingly, the step-to-step dynamics were also significantly affected by the world. This effect is not predicted by any of the surveyed models, and may indicate an effect of inertia, or an influence of subjects' mood or perception of the situation. It may also suggest some refinement to the Regret-Shift model. In the next experiment, we let subjects experience both a good and a bad world, in order to substantiate and further understand this effect, and test explanations such as the latter, which assume lasting effects of the world on choice behavior.

Experiment 2

We set to replicate the results of Experiment 1 with different parameter values, and more importantly, when implementing the 'worlds' within-subject, rather than between-subjects. In this, we aim to further explore how the world affects subjects' choices, when they are now faced with both extremes in turn: We examine if and how subjects react to the change in environment; whether behavior in a specific world depends on it being the first- or second-experienced environment, as well as whether there is a difference between an earlier and a later change.

All subjects experienced both a good and a bad world (defined, as before, by the probabilities of high payoffs) in a randomized block order. The probabilities of the good payoffs (p_R and p_S) changed at some predetermined round, its occurrence unbeknownst to the subjects. Again, the various models described earlier have diverging predictions, enabling comparison and testing of suitability.

Method

Design. Subjects played a similar game to that described in Experiment 1, with two minor differences. First, the values of the rewards were changed, as were the probabilities: The gambles' rewards were set {9,3} for gamble R, and {7,5} for gamble S; the good world was defined as one in which p_R=0.633 and p_S=0.9, and in the bad world, p_R=0.367 and p_S=0.1. The EVs were thus 6.8 and 5.2 in the good world and bad world, respectively – again, equal between gamble R and gamble S. Second, subjects played the game for 120 rounds (instead of 100 as in Experiment 1). The exchange rate from points to NIS was 70 points = 1 NIS (= 0.27 USD).[5]

The main difference is that in Experiment 2 we introduced a change between the good and bad worlds. This change occurred at one of three points: After the 40th round, after the 60th, or after the 80th. It could be either from the good world to the bad one, or vice versa. Subjects were not informed of the change. There were thus six between-subject conditions: Two orders (Bad-Good, Good-Bad) * three change-rounds (41, 61, 81).

Subjects. 132 Hebrew University students (mean age=25.5; 65 females) participated in the experiment for monetary compensation. The condition in which the change occurred after the 60th round was run first, with 65 subjects (mean age=25.2; 32 females). The other two conditions, in which the change occurred after the 40th or 80th rounds, were run in parallel, and included 68 subjects (34 subjects in each condition, randomly assigned; mean age=25.8; 33 females).

The payoff subjects received was contingent upon their performance, and ranged between 9-11 NIS (1 NIS = 0.27 USD), with an average of 9.95.[6] These earnings were very similar across conditions.

Apparatus and Procedure. The procedure was similar to that of Experiment 1. There were two differences in the instructions shown to the subjects: First, the experiment lasted for 120 rounds (instead of 100); second, they were not told anything about the probabilities

[5] The exchange rate was the same for all conditions since all of the subjects experience both a good and a bad world, and thus their expected earnings (as well as actual earnings) are relatively similar, almost irrespective of the change round – unlike the different conditions in Experiment 1.
[6] Earnings in Experiment 2 were lower than those of Experiment 1. This is due to the fact that in Experiment 1, subjects filled a questionnaire which extended the duration of the experiment. In both Experiment 1 and 2, subjects payoff was about twice the hourly student wage, relative to the duration of the experiment.

of the possible values (compared to Experiment 1, in which they were told that the probabilities were fixed). Importantly, subjects were not told that there would be a change in the probabilities (world) at some point in the experiment.

Results

Overall frequencies. The frequency of risky choices was calculated for the two blocks for each subject, and submitted to a Repeated-Measures ANOVA, with the block order (good world to bad, or vice versa) and round of change, as independent variables. The biggest influence was that of world, as expressed in the interaction between the block and the order ($F(1,122)=47.98$, $p<.001$, $\eta_p^2=.282$): In both blocks, subjects were more likely to choose the risky gamble R in the good world rather than in the bad world ($M_{Good}=0.57$, $M_{Bad}=0.46$). It is worth noting that the round of change did not have a significant effect on choice ($p>.3$), nor did it interact with any of the other variables (p's>.09). Clearly, a constant prediction of 50-50 split cannot explain the observed patterns. We thus compare (via Pearson's r) the choice frequencies to the two remaining predictions, of the Regret-Shift and WSLS models: $r_{WSLS}=-0.811$, $p=.001$; $r_{Regret}=0.811$, $p=.001$. That is, the models are an equally good fit to the data – but the direction is that predicted by the Regret-Shift model (opposite of the WSLS prediction). The predictions and the observed pattern, averaged across the round of change, are presented in Figure 2.

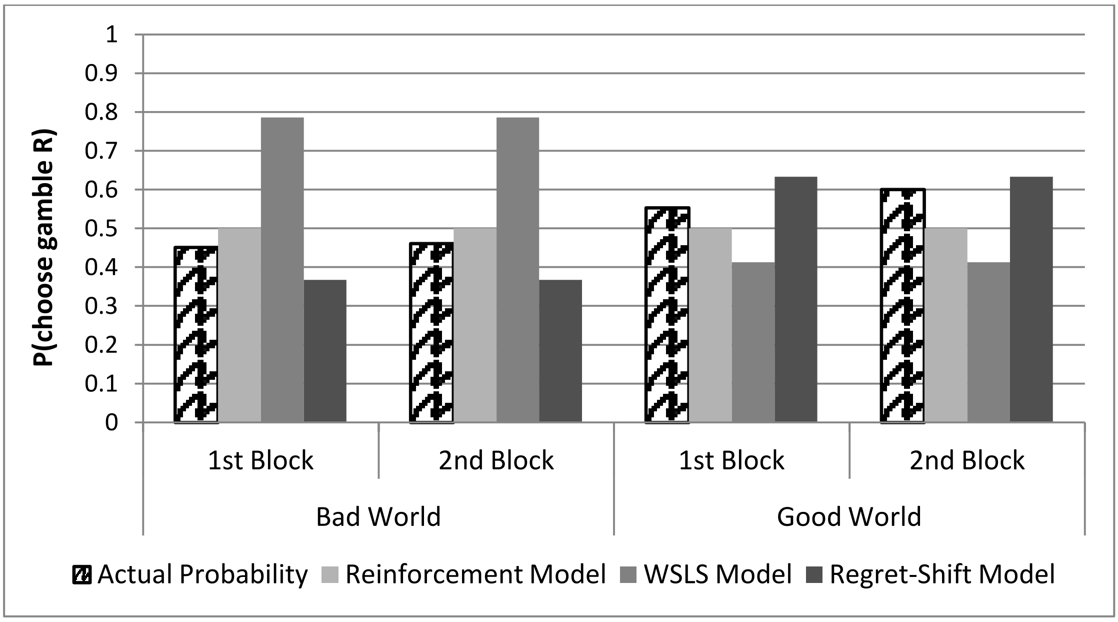

Figure 2. Probability of gamble R choices by world and block (before or after the change), compared with choice probabilities as predicted by the theoretical models: Reinforcement, WSLS, and Regret-Shift.

Dynamics over time. The analysis was equivalent to that described for Experiment 1. When looking at repetition following a riskier choice of gamble R via a logistic regression model (clustered by subject; n=8130 for 132 subjects), the realized reward had a strong effect (Z=13.59, p<.001): Receiving the high payoff led to a higher probability of repeating the riskier choice, as predicted by all models. No other significant effects were observed.

Next, as in Experiment 1, we turn to look at repetitions of choices of the safer gamble S, and submit this binary variable to a separate logistic regression (clustered by subject; n=7710 data points). The results are in Table 4. We observe a large significant negative effect of the reward in the unchosen gamble R (Z=-13.11, p<.001), and a significant positive effect of one's received reward in gamble S (Z=4.23, p<.001). That is, seeing a counterfactual high reward decreases the likelihood of repeating the safer choice; at the same time, receiving a high payoff in the safer gamble slightly modulates this effect. The world had a significant main effect, similar to that previously found: In the good world, compared with the bad world, subjects tended more to shift to the riskier gamble after choosing the safer one (Z=-3.52, p<.001). Only one other interaction, that between the counterfactual and the round of change, was significant (Z=-2.24, p=.025). However, its magnitude was small, and its direction difficult to interpret by either of the models.

The results again point to the suitability of the Regret-Shift model, rather than the other proposed models: The largest influence on choice shifts between subsequent rounds is that of the counterfactual reward of gamble R. This tentative conclusion is strengthened by the absence of a main effect of block. Subjects did not choose differently (or shift more or less) in the second compared with the first block, nor were they influenced by the timing of the change in worlds. These suggest that learning models in general are not very well suited to the choice dynamics; more specifically, that the reaction to the counterfactual reward is indeed short-lived (as predicted by Regret-Shift) and not an accumulation of such reactions.

We also ran a regression in which we include the repetition tendency, or inertia, of each subject. This tendency is calculated, as in Experiment 1, as the subject's frequency of repetitions, minus their probability of choosing gamble R or gamble S twice in a row (that is, the respective probabilities, squared). Adding this variable increased the explanatory power of the comprehensive model ($Z_{stay\_tendency}$=10.61, p<.001; R^2 increased to 0.165), yet it did not affect any of the other variables' influence on choice.

Table 4

Predicting repetition following a choice of the safer gamble – full regression model (standardized variables)

Factor	Beta (S.E.)	Exp(B)	Z	Sig
Realized R	-0.820 (0.062)	0.440	-13.11	.000
Realized S	0.191 (0.045)	1.210	4.23	.000
World	-0.194 (0.055)	0.824	-3.52	.000
RealR * RealS	0.049 (0.047)	1.050	1.04	.299
World * RealR	-0.126 (0.218)	0.882	-0.58	.563
World * RealS	-0.004 (0.215)	0.996	-0.02	.984
World * RealR * RealS	0.013 (0.043)	1.013	0.29	.771
Block (before/ after change)	0.071 (0.054)	1.074	1.32	.185
Change Round	0.066 (0.090)	1.068	0.73	.465
Block * ChangeRnd	-0.015 (0.060)	0.985	0.24	.812
Block * RealR	-0.191 (0.158)	0.826	-1.21	.227
Block * RealS	0.056 (0.215)	1.058	0.26	.794
Block * RealR * RealS	-0.020 (0.046)	0.980	-0.43	.664
Block * World	0.061 (0.085)	1.063	0.72	.474
Block * World * RealR	0.049 (0.231)	1.050	0.21	.832
Block * World * RealS	0.120 (0.185)	1.127	0.65	.518
ChangeRnd * RealR	-0.125 (0.056)	0.882	-2.24	.025
ChangeRnd * RealS	-0.000 (0.046)	1.000	-0.01	.995
ChangeRnd * RealR * RealS	-0.029 (0.049)	0.971	-0.59	.553
World * ChangeRnd	0.011 (0.060)	1.011	0.18	.857
World * Block * ChangeRnd	-0.001 (0.089)	0.999	-0.02	.987
ChangeRnd * World * RealR	0.003 (0.004)	1.003	0.86	.392
ChangeRnd * World * RealS	0.001 (0.003)	1.001	0.32	.746
ChangeRnd * Block * RealR	0.001 (0.003)	1.001	0.32	.746
ChangeRnd * Block * RealS	-0.002 (0.004)	0.998	-0.61	.542
ChangeRnd * World * Block * RealR	-0.001 (0.004)	0.999	-0.20	.845
ChangeRnd * World * Block * RealS	-0.002 (0.003)	0.998	-0.66	.507

Note. N=7710 data points, clustered for 132 subjects. Model fit: -2log likelihood = 8946.99; R^2=0.114.

Discussion

The results of Experiment 2 echo those of Experiment 1. Subjects' overall choice rates diverge from strict indifference between the alternatives (as presumed by Reinforcement models): The better the world (higher probability of receiving a high payoff), the larger the rate of risky-gamble choices. This pattern is in the direction predicted by Regret-Shift, and is opposite to that predicted by WSLS. In this experiment, we observe the same pattern even when subjects experience both a good and a bad world in turn. There is almost no influence of the order in which these are experienced – suggesting subjects are influenced not by general mood or state of mind, but rather by the game characteristics. It appears that the influence of the counterfactual reward is short-lived rather than accumulating, as predicted by Regret-Shift.

Findings regarding choice dynamics also support Regret-Shift, and are quite similar to those found in Experiment 1 – even with the addition of the change in worlds. The differences between consecutive choices were found to be significantly dependent on the reward in the unchosen gamble; specifically, the counterfactual reward in the riskier option was very influential in determining whether to repeat a previous choice of the safer option. There were almost no lasting effects of the change in the game: Neither order effects nor timing effects were observed.

General Discussion

If you were in the situation we opened our paper with, and had to choose between the two roads, of course you would care about the average time it takes to drive them (their expected value). This seems very intuitive. What is somewhat less predictable is that you would make different decisions in regard to this variable: That you would tend to make different choices when both roads are more, compared with less, benevolent – even though their relative standing remains the same. That is, when choosing which route to take, it seems that you would tend to drive the more risky (in terms of variance in driving time) road in cases in which all the roads are pretty empty, but refrain from doing so in rush hours; in that case, you would usually choose the road which has less variance.

This type of behavior refutes the claim that all people look at is the actual payoff they receive, and that they are always indifferent between choices of equivalent expected value. Our findings point to the consideration of variance when making decisions – and that the impact of variance depends on the world (the expected value). We show that people operate according to regret-based models: They look at what they could have gotten in the alternative choice, and make subsequent amendments in their decisions. This is demonstrated in both aggregate behavior – choice frequencies – and in the step-to-step dynamics of choices.

Looking at the probability of repeating a previous choice, we again find an effect of the world, as well as the intuitive influence of the received payoff. However, both of these effects are overshadowed by the influence of the counterfactual reward – the payoff which could have been received had one made the alternative choice. Specifically, when one could have gotten a better payoff making a different choice, then one subsequently tends to shift to that choice, almost regardless of the payoff of one's chosen gamble. Moreover, the effect of the unchosen gamble seems not only to not deteriorate over time, but is actually increasing.

An interesting stipulation of our results is that they may be thought of as disentangling the often confused or confounded feelings of regret and disappointment. In the current setting, disappointment occurs when one receives the low payoff in one's chosen gamble, while regret occurs when one receives a low payoff relative to the counterfactual reward in the alternative gamble. These are orthogonal in the case of a safer gamble choice – and the analyses clearly point to regret being the operating force in our setting.

The effect of world on choice repetition is unaccounted for by any of the surveyed models. The influence of world may appear to reflect inertia – a repetition tendency. However, it is unlikely: If this were true, then there should be a difference between the same world when it is played in the first compared with the second block in Experiment 2, in which the world was manipulated within-subject. Such a difference was not observed. In fact, subjects seemed to adjust quite quickly to the change even though they were not informed of it. Another possible explanation is that the world affects subjects' perception of the situation, causing them to take more risks in environments they believe to be good. This is somewhat in accordance with previous research, showing that positive affect led subjects to optimism when evaluating the probability of winning – and to more risky choices when the potential losses are small (Isen, Nygren & Ashby, 1988; Isen & Patrick, 1983; Nygren, Isen, Taylor & Dulin, 1996). It should be noted that since the time of change did not affect subjects' behavior, in order to believe the above conjecture, one has to assume a quick grasp of the worlds' 'goodness'.

In our studies, we disconnect the riskiness factor – the gambles' variance and spread – from the gambles' expected value; both gambles have equal EV. This condition is a rarity in studies examining choice between gambles – in most cases, the riskier alternative also has a higher EV (e.g., Camilerri & Newell, 2009; Erev & Barron, 2005; Grosskopf et al., 2006; Haruvy & Erev, 2002; Hertwig & Erev, 2009; Hertwig et al., 2004). This enables differentiation between the proposed models. We strengthen previous findings in the literature regarding the impact of foregone, counterfactual payoffs (Avrahami et al., 2005; Avrahami & Kareev, 2011; Camerer & Ho, 1999; Daniel et al., 1998; Erev & Barron, 2005; Ert & Erev, 2007; Grosskopf et al., 2006; Yechiam & Busemeyer, 2005; 2006; Yechiam & Rakow, 2011; Zeelenberg et al., 1996; Zeelenberg & Pieters, 2004): We demonstrate these effects in a repeated choice situation in which subjects' experiences are those that help them both understand and react to the environment. We further show that this influence is quite immediate and does not require aggregation of response tendencies over time.

A limitation of the current study is that all of the possible payoffs were strictly positive – subjects could not lose money, or step out empty-handed (even from a single round). It is possible that different behavioral patterns will emerge – either in the overall frequencies or in the choice dynamics – when considering losses rather than gains, or any combination of the two. At the very least, choices between two gambles in the losses domain is, in our opinion, an avenue which warrants future research. It would also be interesting to repeat the research with alternatives that differ in their expected values.

Epilogue

Granted, the experiments described in this paper will probably not help you know, think or decide what to do when faced with gambles or risky choices. However, it seems that it doesn't much matter. It appears that most choices between gambles can be explained or accounted for without looking ahead or far back, or assuming some calculation of any of the relevant parameters. Just respond to what happens – or better yet, to what could have happened.

Acknowledgment

Work reported in this paper was supported by Israel Science Foundation Grant 2011/121 to Yaakov Kareev.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. *Econometrica, 21*, 503–546.

Avrahami, J., Güth, W., & Kareev, Y. (2005). Games of competition in a stochastic environment. *Theory and Decision, 59*, 255-294.

Avrahami, J., & Kareev, Y. (2010). Detecting change in partner's preferences. *Discussion Paper Series #557,* Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem.

Avrahami, J., & Kareev, Y. (2011). The role of impulses in shaping decisions. *Journal of Behavioral Decision Making*, *24*, 515-529.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16,* 215-233.

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

Busemeyer, J., & Stout, J.C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, 14, 253-262.

Camerer, C., & Ho, T.H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica, 67,* 827-874.

Camilleri, A.R., & Newell, B.R. (2009). The role of representation in experience-based choice. *Judgment and Decision Making, 4*, 518–529.

Cheung, Y.W., & Friedman, D. (1998). A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior and Organizations, 35*, 263-280.

Cooper, D., Garvin, S., & Kagel, J. (1997). Signalling and adaptive learning in an entry limit pricing game. *Rand Journal of Economics, 28*, 662-683.

Daniel, T.E., Seale, D.A., & Rapoport, A. (1998). Strategic play and adaptive learning in sealed bid bargaining mechanism. *Journal of Mathematical Psychology, 42*, 133–166.

Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review*, *114*, 177-187.

Erev, I., & Barron, G. (2005). On adaptation, maximization and reinforcement learning among cognitive strategies. *Psychological review, 112*, 912-931.

Erev, I., Ert, E., Roth, A.E., Haruvy, E., Herzog, S.M., Hau, R., Hertwig, R., et al. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15-47.

Erev, I., & Roth, A. (1998). Prediction how people play games: Reinforcement learning in games with unique strategy equilibrium. *American Economic Review, 88*, 848-881.

Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making, 20*, 305-322.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*, 523-551.

Grosskopf, B., Erev, I., & Yechiam, E. (2006). Forgone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory, 34*, 285-302.

Haruvy, E., & Erev, I. (2002). On the application and interpretation of learning models. In: Zwick, R., & Rapoport, A. (Eds.), *Experimental Business Research*. Kluwer Academic Publishers.

Hertwig, R. (2011). The psychology and rationality of decisions from experience. *Synthese*.

Hertwig, R., Barron, G., Weber, E.U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.

Hertwig, R., Barron, G., Weber, E., & Erev, I. (2006). Risky prospects: When valued through a window of sampled experiences. In: Fiedler, K., & Juslin, P. (Eds.), *Information sampling and adaptive cognition* (pp. 72-91). Cambridge, England: Cambridge University Press.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13,* 517-523.

Hills, T.T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science, 21,* 1787-1792.

Hogarth, R.M., & Einhorn, H.J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology, 24*, 1-55.

Isen, A.M., Nygren, T.E., & Ashby, A.F. (1988). Influence of positive affect on the subjective utility of gains and losses: It is just not worth the risk. *Journal of Personality and Social Psychology, 55*, 710 –717.

Isen, A.M., & Patrick, R. (1983). The effect of positive feelings on risk-taking: When the chips are down. *Organizational Behavior and Human Decision Processes, 31*, 194 –202.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263-291.

Knight, F.H. (1921). *Risk, uncertainty and profit*. Boston: Houghton Mifflin.

Levine, M. (1966). Hypothesis behavior of humans during discrimination learning. *Journal of Experimental Psychology, 71*, 331–338.

Loewenstein, G.F., Weber, E.U., Hsee, C.K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin, 127*, 267–286

Matsen, F.A., & Nowak, M.A. (2004). Win-stay, lose-shift in language learning from peers. *PNAS*, 101,18053-18057.

Nowak, M.A., & Sigmund, K.A. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, *364*, 56-58.

Nygren, T.E., Isen, A.M., Taylor, P.J., & Dulin, J. (1996). The influence of positive affect on the decision rule in risk situations: Focus on outcome (and especially avoidance of loss) rather than probability. *Organizational Behavior and Human Decision Processes, 66,* 59-72.

Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with non-stationary payoffs. *Memory and Cognition, 37*, 985-1000.

Roth, A.E., & Erev, I. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economics Behavior, 8*, 164-212.

Selten, R., & Buchta, J. (1998). Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions". In: Budescu, D., Erev, I., & Zwick, R. (Eds.), *Games and Human Behavior*. Hillsdale, NJ.

Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., & Damasio, A. (2005). Investment Behavior and the Dark Side of Emotion. *Psychological Science, 16*, 435-439.

Stewart, N., Chater, N., & Brown, G.D.A. (2006). Decision by sampling. *Cognitive Psychology, 53*, 1-26.

Thaler, R.H., & Johnson, E.J. (1990). Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. *Management Science, 36,* 643-660.

Weber, E.U., Shafir, S., & Blais, A.R. (2004). Predicting risk-sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

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.

Yechiam, E., & Busemeyer, J.R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making, 19,* 1-16.

Yechiam, E., & Rakow, T. (2011). The effect of foregone payoffs on choices from experience: An individual level modeling analysis. *Experimental Psychology*, *13*, 1-13.

Zeelenberg, M., & Pieters, R.G.M. (2004). Consequences of regret aversion in real life: The case of the Dutch postcode lottery. *Organizational Behavior and Human Decision Processes, 93*, 155-168.

Zeelenberg, M., Beattie, J., van der Pligt, J., & de Vries, N.K. (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes, 65*, 148-158.