Journal of Behavioral and Experimental Finance 12 (2016) 23–39

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Finance

journal homepage: www.elsevier.com/locate/jbef

Full length article

Real options in the laboratory: An experimental study of sequential investment decisions

Ryan O. Murphy[a], Sandra Andraszewicz[b,*], Simon D. Knaus[c]

[a] University of Zürich, Blümlisalpstrasse 10, 8006 Zürich, Switzerland
[b] ETH Zürich, Scheuchzerstrasse 7, 8092 Zürich, Switzerland
[c] University of St. Gallen, Bodanstrasse 6, 9000 St. Gallen, Switzerland

ARTICLE INFO

Article history:
Received 28 January 2016
Received in revised form
26 August 2016
Accepted 29 August 2016
Available online 16 September 2016

Keywords:
Real options
Sequential choice
Multistage risky decision making
Dynamic decision making
Prospect theory

ABSTRACT

Many real-life risky decisions in finance and management are dynamic and decision policies can be adapted as uncertainty is reduced by the arrival of new information. In this type of situation, called a *real options* problem, a decision maker must choose how much of his finite resources to invest in a dynamic risky environment. In two laboratory experiments, we test a well-defined decision problem with the central characteristics of a real options framework and do so in such a way that it is amendable to formal modeling. We find that people choose differently than the expected value maximizing policy, consistent with risk aversion and non-linear probability weighting. We conclude that although real options analysis is useful as a normative valuation method, its recommendations are sometimes contrary to people's innate tendencies when making risky choices and this counterintuitiveness should be considered when implementing real options analysis in training and practice.

1. Introduction

Decision makers must sometimes choose whether to invest valuable resources in an ongoing and evolving prospect or abandon the prospect given the arrival of new information. Take as one prototypical example a firm that has the option to invest money in a nascent start-up company. This start-up has an uncertain future and thus an uncertain chance of paying off as an investment. The potential return for the venture capital (VC) firm is proportional to the amount of money it cumulatively invests in the start-up. If the expected return is sufficiently high for the start-up, the VC firm should invest money; conversely if the expected return is low for the start-up, the VC firm should refrain from investing its resources. One constraint that makes this type of problem nontrivial is that the VC firm is limited in that it must make these investment decisions sequentially over discrete periods, and moreover it is restricted in how much money it can invest at any one time. Dynamic risky decision situations (Edwards, 1962) like this are referred to as *real options* problems, a term first coined by Myers (1977). In these contexts, a decision agent must repeatedly choose, over multiple distinct periods, if it is sensible to invest in ongoing and evolving prospects, and if so, how much to invest. Conversely, a decision agent can choose to abandon the prospect and not invest resources. Real option problems are characterized by sequential dynamic decisions, where information about the success of these decisions is revealed over time, past decisions are irreversible, and future decision strategies can be adapted depending on newly available information. In this paper, we develop a real options decision task, derive the optimal normative policy, experimentally examine people's risky choices, compare the observed behavior to the normative policy, and comparatively evaluate several descriptive models of the observed behavioral patterns emerging from the choices.

1.1. Literature review

Real options contexts are widely encountered in real world settings including entrepreneurial decision making (McGrath, 1997, 1999), labor economics (Jacobs, 2007; Hogan and Walker, 2007), mineral and oil exploration (Babajide, 2007), research and development (Rogers et al., 2002; Schwartz, 2004; Sereno, 2010), environmental technology investments (Cortazar et al., 1998), and venture capital investing (Hsu, 2008). As such, the topic has been addressed by several disciplines and from a variety of different angles. At the macro level, these situations are often related to

* Corresponding author.
E-mail addresses: ryan.murphy@econ.uzh.ch (R.O. Murphy), sandraszewicz@ethz.ch (S. Andraszewicz), sknaus@gmx.ch (S.D. Knaus).

http://dx.doi.org/10.1016/j.jbef.2016.08.002
2214-6350/

strategy and corporate finance. Bowman and Hurry (1993) points out that real options are essential components of a firm's overall strategic capacity, an assertion that multiple other studies echo (e.g., McGrath, 1997, 1999; Sirmon et al., 2007; Dess et al., 2003). Klingebiel and Adner (2015) conducted qualitative research in which 28 practitioners from eleven firms in diverse industries were interviewed to evaluate the performance advantage in product innovation when firms applied a real options approach. They found that sequential investments, consistent with real options analysis, resulted in better firm performance. This supports the notion that in a dynamic choice environment a firm can capitalize on flexibility and this yields added value. Although real options may confer a strategic advantage to firms, it is not always observed in industry and markets. For example, Quigg (1993) featured a real options pricing model that considered the value of waiting different time intervals to invest in real estate. Results showed the existence of persistent market inefficiencies in that observed prices systematically exceeded the model-implied values. Howell and Jägle (1997) found similar results when they investigated the valuation of options by managers in nine major British companies and report a general overvaluation of options relative to the normative (i.e., expected value maximizing) pricing policy. These studies provide evidence of the strategic importance of real options problems for firms, but they do not address the micro foundations of individual decision making, namely individuals' abilities and propensities to follow the dictates of an optimal decision policy in a dynamic setting.

Other studies highlight the relationship between a real options context and individual risky decision making. In a micro setting Miller and Shapira (2004) asked decision makers to specify the price for selling or buying a call or a put option for simple binary lotteries. In that task, individual participants generally undervalued options (with regard to the expected payoffs), but overestimated expected losses for selling a put. Further, Yavas and Sirmans (2005) reported that in a real options problem individuals generally invested too early and thus failed to realize the full value of the flexibility granted by their options. In another study (Wang et al., 2009), participants had to trade commodities produced by their "factory", depending on the commodities' price changes that developed over time. The results indicated that participants did not follow an optimal policy but rather exhibited two distinct kinds of decision biases: first, not incorporating the expected price, and second, exhibiting a general insensitivity to the termination date. Oprea et al. (2009) presented somewhat contrasting results from a laboratory experiment where participants had to decide when to invest in a risky option. In this setting, participants eventually learned to wait and enact their decision when uncertainty was sufficiently resolved. However, this near optimal behavior was not observed at the beginning of the study, but only in the last block of the experiment.

Overall, these results indicate that, although the real options valuation approach is widely applicable and does bestow value, it may be difficult for people to implement properly. Previous work suggests innate behavioral tendencies that are contrary to the normative dictates of real options analysis and expected value maximization in dynamic contexts. However, the distribution, persistence, and structure of these biases is not clear. Additionally, previous approaches to real options research do not allow for formal modeling, which can be a useful tool for discerning the actual decision policies of decision makers.

To better understand the behavioral propensities and the cognitive biases operating as people make these kinds of dynamic risky choices, we develop a general real options problem that is amenable to both formal modeling and laboratory experimentation. It has a clear and simple structure that can be easily implemented in an incentive compatible experimental setting. This dynamic decision problem retains the central conflict that is at the core of the real options vignette presented previously and in many other problems found in the wild. In these settings a decision maker faces, over multiple discrete stages, the option of trading a certain alternative of real value, for a risky option with potentially greater value, all within the context of a dynamic environment with updating information about the probability of gain.

1.2. Motivation for the current paper

The major research questions for this paper are twofold: First, we wish to compare individuals' sequential investment decisions in a set of real option decision problems to the normative (e.g., a risk neutral expected value maximizing) decision policy. Do individuals make decisions (generally) consistent with this normative solution, and thus maximize their expected earnings? Alternatively, do individuals' decisions deviate significantly from what is optimal to these ends, thus diminishing potential earnings. And if there are significant deviations, what, if any, decision patterns do individuals exhibit in these real option problems?

Second, assuming we find some systematic non-expected value maximizing behavioral results,[1] can the emergent choice patterns be sensibly modeled? Would simple utility theory account for the stylized facts? Or would non-expected utility theory involving both risk aversion and probability weighting (considered concurrently) be necessary to describe the results? Or would a different characterization invoking different heuristic models better capture the pattern of empirical findings? In order to adequately address this second set of questions, a host of different choice models will be competitively compared for goodness-of-fit and out-of-sample predictive power in accounting for the observed decisions. The data used to address these empirical questions result from two laboratory experiments using a tractable multistage real options decision problem explained below.

Lastly, we are motivated to promote more attention to the study of dynamic risky decision making in general. Many interesting and important decision problems from the real world have a sequential structure, where the decision maker is called upon to make a series of choices with updated information and where her later options depend on her previous choices. This class of problems has a long history (e.g., Edwards, 1962) but unfortunately remains understudied given its potential richness. To be sure, simple static gambles, which currently prevail in the research landscape, are useful building blocks—but there is a great deal more to understanding risky decision making in people than can possibly be uncovered using only one-shot gambles.

In the following sections in the paper, we formally define and solve a well structured real options problem, identifying the expected value maximizing decision policy. Then, we analyze and report the behavior resulting from the experiments, using the normative solution as a benchmark. Thereafter, we develop, contrast, and explore different behavioral choice models and evaluate how well they correspond to the pattern of empirical results. This includes testing whether expected utility maximization is sufficient to describe the results. Lastly, we conclude with a discussion about some persistent biases we find and offer suggestions for countering these tendencies when considering real options analysis.

[1] The descriptive inadequacy of expected value maximization is well known (see for example Camerer, 1995, Wakker, 2010, Barberis, 2013, and Fox et al., 2015 for general overviews).

1.3. A stylized real options decision problem

This paper started with an example of a nascent start-up company with an uncertain future. This is just one multistage decision scenario and there are of course many other situations which have at their heart a conflict between accepting a less valuable but certain alternative in lieu of foregoing it and choosing instead a risky future option. Consider yet another example, that of an employer recruiting a new employee. The employer is contemplating investing money, time, and effort into training the new employee. The likelihood of the employee succeeding (learning new tasks and performing well in her new position) is unknown, but the uncertainty about her prospects is reduced over time as the employer learns more about her capabilities and inclinations. Nonetheless, the employer must decide upfront if and how much money and effort to invest toward the employee's training. However, the employer can revisit this arrangement periodically (e.g., every time the contract is renewed), based on new information related to the likelihood of the employee's success. Employees with a high estimated probability of success warrant ongoing investments, whereas employees with an inauspicious future would warrant abandonment (i.e., removal from the position). These examples can be mapped to a general dynamic decision problem which we develop more precisely next.

1.3.1. Verbal description of the experimental problem

Consider a decision maker (DM) facing the following situation: The DM starts with an endowment of 36 dollars. A fair die will be rolled 6 times, and a running sum of the identically and independently distributed (*i.i.d.*) results from each of the rolls will be recorded. Before every roll, the DM can invest integer values between 0 and 6 dollars inclusive, but no more at any one time. The money the DM invests "stays on the table" until after the 6th roll of the die. The DM is paid double whatever has accumulated on the table after the 6th roll if and only if the sum of rolls equals at least 21. If the sum of the rolls does not equal at least 21, then the money invested and accumulated on the table vanishes, returning to the house. In either case, the DM keeps whatever portion of her endowment she chooses not to invest. Thus, a DM could, at the extremes, double her endowment or lose it all. The DM could also risk nothing and guarantee a payoff equal to her starting endowment by always choosing to invest 0. All of these conditions are known and understood by the decision maker (e.g., there is no ambiguity in the decision context).

Table 1 shows an example of a simple real options decision problem played out. In the first stage of the problem, the DM invested 4 dollars. The roll of the die in the first stage was 5. This is a relatively good roll and puts the DM closer to the goal of reaching a cumulative sum of at least 21 over the 6 stages (one roll per stage). In the second stage, the DM chose to invest 6 dollars. The roll on the second stage was 4. The DM is now in a situation where the goal is 12 away with 4 rolls remaining. The DM chose to invest 5 dollars in the third stage. The roll in this stage was only 1. Subsequently, the DM acted more conservatively and invested nothing in the fourth stage. This process of a DM making an investment choice in a stage followed by the roll of the die continued. By the last stage, the DM had invested a total of 23 dollars. The running sum of rolls reached the goal of at least 21 in the last stage, yielding a "win" and a payoff that doubles the total amount invested by the DM. In this particular round, the DM earned $(23 \cdot 2) + 13 = 59$. The 23 was the total amount invested which was then doubled because the goal was met, and the 13 was the amount not invested and left over from the endowment.

Although this is a simplified decision problem, it retains several fundamental features underlying real options problems encountered in the real world. The DM has to sequentially choose between a sure thing and a risky option. It is sensible for the DM to invest in the risky option provided the expectation is sufficiently high. More information about the probability of the prospect paying off is revealed over time. In tandem, with the gradual revelation of information, the DM has to choose periodically if and how much money to invest in the risky prospect. As all of this happens over discrete stages of the decision task, uncertainty is reduced for the DM by the revelation of more information (and possibly by the last stages of the task, the outcome may be certain), and choices once made are irreversible.

Table 1
An example of a simple real options decision problem played out.

Stage	Amount invested	Roll of the die	Running sum of rolls
1	4	5	5
2	6	4	9
3	5	1	10
4	0	4	14
5	3	5	19
6	5	3	22
End	Total: 23	Total: 22	Outcome: WIN
Payoff: $(36 - 23) + (2 \cdot 23) = 13 + 46 = 59$			

1.3.2. Formal description of the problem

The decision task consists of T discrete stages. At the beginning of each stage t, the DM invests $I_t \in \{0, 1, \ldots, L\}$ dollars. L is the constraint upon how much the DM can invest in any stage, and clearly the sum of investments cannot exceed the DM's endowment E. Thus the DM is constrained also such that $\sum_{t=1}^{T} I_t \leq E$. After the DM has chosen I_t, there is an *i.i.d.* discrete uniform random variable R_t drawn from a well-defined set, in this case $\{1, 2, 3, 4, 5, 6\}$. The value of R_t is shown to the DM after he has made his investment decision I_t, and this ends stage t. This process of investing money and then observing a random value repeats for T stages. Let X_t be defined as the sum of the random variables R_t from stage 1 to t. Thus $X_t = \sum_{i=1}^{t} R_i$, where R_t is the outcome of the roll in stage t. Therefore, X_t is the cumulative sum of rolls up to and including stage t. If after T stages (*cf.* $\max(t) = T$), X_T is equal to or greater than the goal G, then the indicator function $\mathbb{1}_{\{X_T \geq G\}}$ yields 1; otherwise it yields 0. The DM's payoff for the decision task is:

$$\Pi = \left(E - \sum_{t=1}^{T} I_t\right) + \left(\mathbb{1}_{\{X_T \geq G\}} \cdot 2 \cdot \sum_{t=1}^{T} I_t\right). \tag{1}$$

Throughout the decision task, the DM knows: the endowment E, the probability distribution of R_t (i.e., here R_t is uniformly distributed over $\{1, \ldots, 6\}$), the value of the goal G, the current stage t, the realized random values $R_1, \ldots, R_t$, the sum of draws so far X_t, and the payoff function for the task. The optimal normative policy[2] for this problem dictates that the DM should invest 6 units whenever the probability of reaching the goal is greater than 0.5.[3] According to this policy, an EV maximizing DM should never invest "part-way" but would invest all or nothing—either 0 or 6. This directive stands in stark contrast to heuristics, such as naïve diversification. Naïve diversification, a $1/N$ heuristic, would lead a DM to put at risk half of her available resources in each stage and keep the other half of the endowment (for a complete discussion the reader is referred to Section 5.2).

2 The derivation of the optimal policy is presented in Appendix D.

3 There is indifference in the case when $P[X_T \geq G|X_{t-1}] = 0.5$. For this edge condition, any investment amount is equivalently attractive to the DM. To obtain a clear deterministic rule, we set the optimal policy to invest whenever the strict inequality is satisfied $P[X_T \geq G|X_{t-1}] > 0.5$.

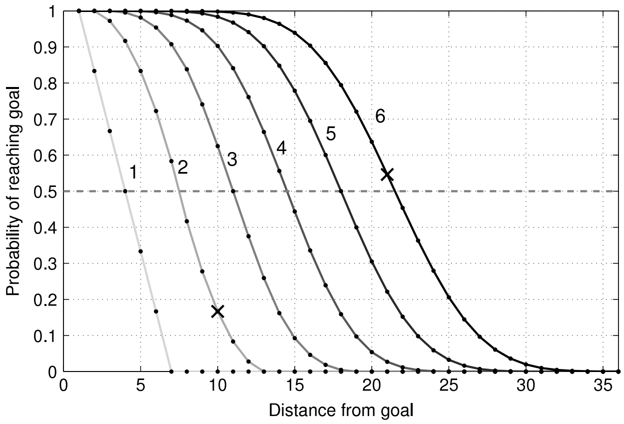

Fig. 1. Probabilities of winning given the distance from goal and number of rolls remaining. The different lines correspond to the number of rolls remaining; the state with 6 rolls remaining is on the far right, the state with 1 roll remaining is on the far left, and the intermediate lines are ordered sequentially therein. Each point corresponds to a potential overall outcome of the sum of rolls. Two ×'s provide illustrative examples: At the start of a round, a DM could have 6 draws remaining and be 21 away from the goal. The probability of success in that first stage of the decision task is 0.546. If the DM had only 2 rolls remaining and the running sum was 10 less than the goal, then the probability of winning is 0.167.

The normative decision policy here depends exclusively upon the probability of success (i.e., the probability that the sum of rolls is greater than or equal to the goal) in each decision stage. The probability of eventual success is wholly contingent on two factors: the *distance from the goal* $(G - X_i)$ and the *number of rolls remaining* $(T - t)$. Fig. 1 shows the relation between the distance from the goal, remaining rolls, and the probability of reaching the goal. The probabilities show the full enumeration of all states in the decision problem. Tangentially, this particular representation of the decision problem nicely illustrates the central limit theorem—the greater the number of remaining stages, the more closely the curve approximates 1 minus the CDF of a Gaussian distribution.

2. Experiment 1: choosing investments over stages

An experiment was designed and conducted using the real options investment problem described above in order to test the descriptive accuracy of the normative model. This multistage dynamic decision problem was presented to participants and iterated over 80 independent rounds, where each round contained 6 stages.

2.1. Participants

Forty-two participants were recruited to participate in "decision making research". These individuals were recruited from a university based in Singapore and were primarily business school undergraduates. All of the research participants were fluent in English, and their ages ranged between 20 and 25 years. Men and women were equally represented in the sample. Participation in the research was voluntary and incentive compatible. No deception was used in this research.

2.2. Procedure

The research was conducted in a dedicated behavioral laboratory with 16 personal computers arranged in private cubicles. Participants were brought into the laboratory in groups of 12–16 and given a brief overview of the research task. After the introduction, participants drew random numbers to determine which cubicles they would occupy for the duration of the experiment. Once participants entered the main laboratory room they were not allowed to communicate with each other.

Participants were provided with comprehensive written instructions (see Appendix B) describing the task and the interface of the computer program that was used to administer the experiment. The instructions described the sequential decision problem in non-technical language, and the values of the parameters of the problem were all presented explicitly. Experimenters were available to answer any of the participants' questions privately; few questions were asked. Once the participants verified they were confident in their understanding of the task and ready to start, they performed 80 rounds of the decision task. At the beginning of each round, the participants were endowed with 36 experimental dollars. In each of the 6 stages of a round, a participant could invest between 0 and 6 experimental dollars. The goal G for a particular round was randomly drawn from the set $\{20, 21, 22, 23\}$. The four goal values were counterbalanced such that each participant had 20 rounds with each goal value. Information about past investments, remaining endowment, rolls, sum of rolls, current goal, round and stage number were provided on the computer screen and updated in real time. Appendix A shows the user interface presented to the participants. At the end of each round, the participants were informed whether the goal was reached and what their earnings were. Rounds were independent; earnings from one round did not carry over into subsequent rounds.

The vast majority of participants completed the experiment between 40 and 55 min. After completion of 80 rounds, participants drew a card from a deck of 80 shuffled cards, numbered from 1 to 80. Their drawn card's number determined the single payment round. A constant exchange factor of 1/3 was used to control for the magnitude of potential earnings; for example, earnings of 72 experiment dollars resulted in a payment of 24 Singapore dollars (1.00 SGD = 0.80 USD at the time of the experiment, but the Singapore dollar has a comparatively high index of purchasing power). Participants were informed of this remuneration scheme at the beginning of the experiment and thus had an incentive to do their best in all of the experimental rounds, as all rounds were equally likely to determine their compensation. The participants were paid for their participation privately.

2.3. General results

Fig. 2 shows all investment decisions aggregated over all participants, plotted as a function of the objective probability of winning.[4] The darker the areas in the figure, the higher the frequency of each invested amount. The estimated logistic fit indicates that overall participants invested higher amounts when the probability of winning was greater. However, the normative prediction that investments will only take values of 0 or 6 can be refuted, as only 55% of investment decisions occurred at these levels.

The most common investment value was 0 observed 38% of the time, followed by 6 with the choice frequency of 17%. The next most common investment was 3; this middle investment alternative occurred 15% of the time. Next, the number of investments at the lower investment alternatives 1 and 2 (20% of the total) occurred twice as frequently as investments at the higher alternatives 4 and 5 (10% of the total). This finding is consistent with an asymmetric probability weighting function as instantiated by prospect theory (Tversky and Kahneman, 1992; Prelec, 1998) and will be discussed in more detail in Section 5.1.

[4] The raw data from these experiments are available from the authors upon request.

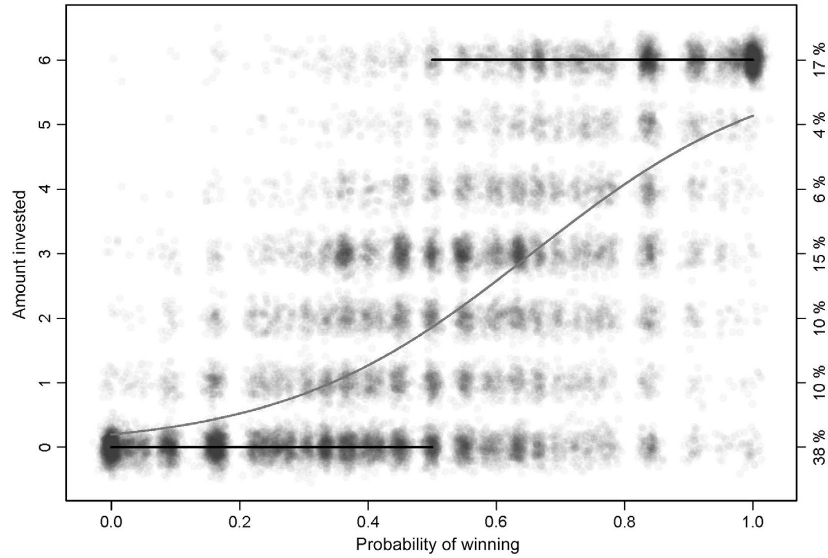

Fig. 2. The frequency of each investment decision (integer values ranging 0–6) depending on the probability of winning, with the initial endowment of 36 dollars. Each point corresponds to a particular investment decision. The *x*-axis is the objective probability of reaching the goal when a particular investment was made. The *y*-axis corresponds to the amount invested. The points are jittered to circumvent stacking. The percentages listed on the right indicate the relative frequency of choices of the different investment levels over all participants. The curved line shows a logistic fit to the data. The optimal investment policy, the flat lines, is a limit case (with a slope parameter approaching infinity) of the logistic fit. The logistic regression was fit to the data using OLS. The probabilities were mapped onto [−1, 1] and investments linearly rescaled to [0, 1]. The plotted curved line shows the back-transformed fit.

Table 2
Frequencies distribution of bias values over each stage expressed as percentages. These data are from Experiment 1 where the endowment is 36.

	Stage						
	1	2	3	4	5	6	Total
−6	11.40	7.17	9.97	6.64	6.91	2.98	7.51
−5	4.11	3.30	4.91	3.36	3.51	1.37	3.43
−4	5.75	6.67	7.38	4.32	3.30	1.10	4.75
−3	18.58	9.29	8.45	6.43	5.27	2.26	8.38
−2	4.35	6.43	5.81	4.61	3.33	1.73	4.38
−1	1.61	2.89	4.76	4.26	3.63	1.55	3.12
0	19.05	31.82	36.47	49.75	62.82	77.97	46.31
1	5.54	10.39	9.68	8.54	4.61	3.42	7.03
2	6.61	9.02	5.92	5.00	2.86	1.88	5.22
3	17.48	7.92	4.50	4.47	1.99	2.35	6.45
4	3.51	3.01	1.04	1.13	0.45	0.95	1.68
5	0.68	0.80	0.48	0.48	0.42	0.39	0.54
6	1.34	1.28	0.63	1.01	0.89	2.05	1.20
Total	100	100	100	100	100	100	100

Note: The shaded columns highlight the structure of the diminishing bias levels across stages. The prominent investment of 3 (resulting in a bias of ±3) decreases as the decision task progresses over stages.

2.4. Bias index

A bias index is defined as *actual amount invested—optimal amount that should have been invested under the EV maximizing policy*. This index was computed for each participant's investment choice given the respective state of the decision problem. The bias value of 0 indicates an optimal investment decision, whereas positive values indicate overinvestment and negative values indicate underinvestment. For example, a bias value of −6 indicates that the DM invested nothing in a particular stage of the task when it would have been EV maximizing to invest 6.

Table 2 and Fig. 3 show average proportions of bias values in each round, across all participants. The figure and the table highlight a few major findings. First, across all stages, 46% of DMs' choices were consistent with the optimal policy. Second, underinvesting was more prevalent than overinvesting (32% vs. 22%). Third, biases are most pronounced in the early stages of the decision task, when subjective uncertainty is the greatest. In the first stage of the decision task, DMs are investing sub-optimally 81% of the time, commonly choosing to invest 3. This tendency persists until at least stage four, where roughly half of all investment choices are optimal. By the last stage of the task, DMs are investing optimally nearly 78% of the time, and the tendency for naïve diversification has almost entirely disappeared. An illustration of this tendency can be found in Fig. F.2 where the plot of Fig. 2 is separated into stages 1–6.

2.5. Earnings

Following the optimal decision policy, a DM would expect to earn 44.46 experimental dollars per round given the random virtual die rolls from this experiment. A static decision policy of never investing would obviously yield earnings of exactly 36, and the static policy of investing 6 in every round would have yielded average earnings of 36.40 but resulted in extreme variance (as all the resulting earnings per round would have been either 0 or 72 experimental dollars). Empirically, participants on average earned 41.08 experimental dollars across all rounds.[5] The return on endowment (ROE; *ROE* = (amount won − endowment)/endowment) was 14.1% whereas the optimal investment yields a ROE of 23.5%, indicating significant underperformance (one sample t-test: $t(41) = -10.56$, $p < 0.001$). We observed a substantial variance in individual earnings across participants ($\sigma = 15.52$), where the middle 50% (i.e., second and third quartile of participants by experimental earnings) of participants earned on average between 30 and 54 experimental dollars across all rounds.

3. Experiment 2: choosing investments over stages with a constrained budget

3.1. Motivation

In order to more rigorously test the persistence of decision biases (i.e., systematic deviations from the optimal policy), we

[5] Data analysis is conducted on all results from the experiment, regardless of whether a round was actually chosen as the payment round.

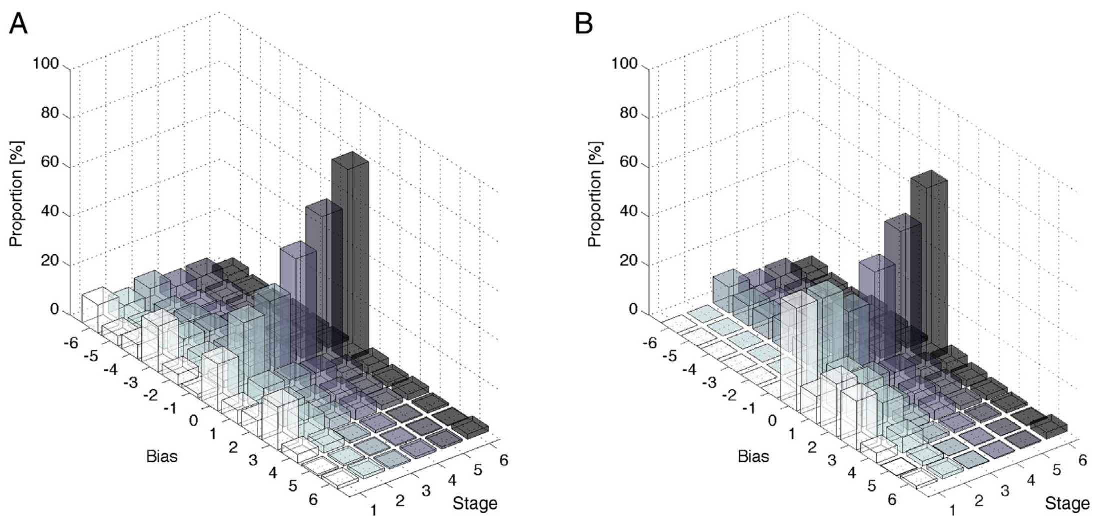

Fig. 3. Bivariate distribution of bias levels over the different stages of the choice task. (A) 36-condition: In the first two stages, a substantial number of investments have a bias of 3 or −3 (about 39% of all choices in this stage). This results from DMs' propensity to invest 3 in the initial stage of a round. However, as stages continue and the objective probability moves toward 0 or 1, DMs are more likely to move away from investing 3 and make optimal investments, as can be seen by the larger mass with a bias of 0 at the last two stages. (B) 24-condition (see Section 3.1): A substantial number of investments that have a positive bias (about 58% of all choices in these stages). This is the result of DMs' propensity to invest in the initial stages of a round, even though this is a dominated decision policy. However, in the further stages, the objective probability moves toward 0 or 1 and the DMs more often make optimal investments, as can be seen by the larger proportion of choices with a bias of 0 in the last two stages.

developed a second version of the real options investment task that may compel a DM to "think twice" before investing money. In this experiment, DMs were endowed with only 24 dollars from the start of each round, resulting in a more constrained budget. The DMs could no longer exploit the full range of investments over all stages but in some stages would necessarily have to invest less than 6 units. All other parameters and methods here were the same as in Experiment 1. As a result of the reduced endowment, the normative policy is to *never* invest in the first 2 stages of the round, and then to invest 6 experimental dollars in stages 3–6 if and only if the probability of success is strictly greater than 0.5; otherwise a DM should invest nothing. This means that non-investing in the first two stages is a dominant policy, due to the reduced budget, and is invariant with the realized probability of reaching the goal. As before, the normative DM would invest only 0 or 6 dollars, never in-between values. And again, the optimal policy stands in stark contrast to the $1/N$ diversification policy and other choice heuristics.

3.2. Participants and procedure

Forty-one people were recruited to participate in "decision making research". None of these individuals had taken part in Experiment 1. The procedure for this experiment was identical to that used before (see Section 2.2) but with the change in the endowment level ($E = 24$) described above. The exchange rate for this experiment was increased to 1/2, thus maintaining the range of potential payments from 0 to 24 Singapore Dollars, which was the same as in Experiment 1.

3.3. Bias index

Bias values are displayed in Table 3.[6] The majority of investments in the first stage were different from 0, which is inconsistent with the optimal policy. DMs displayed a strong propensity to overinvest in early stages, primarily investing lower values from 1 to 3, even though these investment choices were dominated given the guaranteed option to invest in later stages with potentially less subjective uncertainty. Compared to Experiment 1, the total bias decreased by only about 1.5%.

Table 3
Frequencies distribution of bias values over each stage expressed as percentages. These data are from Experiment 2 (i.e., where the endowment is 24).

	Stage						
	1	2	3	4	5	6	Total
−6	–	–	10.52	6.55	6.80	3.75	4.60
−5	–	–	8.17	4.30	4.21	2.23	3.15
−4	–	–	8.32	5.79	4.82	1.86	3.47
−3	–	–	10.37	8.57	6.68	3.63	4.88
−2	–	–	7.20	6.68	4.39	2.56	3.47
−1	–	–	2.23	3.51	3.11	1.89	1.79
0	41.83	41.77	32.32	44.33	56.89	70.40	47.92
1	11.25	16.16	10.03	8.78	5.70	4.51	9.41
2	21.55	17.20	5.37	5.43	4.05	2.77	9.40
3	19.88	15.12	3.54	4.05	1.77	1.95	7.72
4	4.02	6.31	1.55	1.34	0.98	1.25	2.58
5	0.24	1.95	0.15	0.34	0.15	0.46	0.55
6	1.22	1.49	0.24	0.34	0.46	2.74	1.08
Total	100	100	100	100	100	100	100

Note: Each column shows the distribution of bias across investment stages (summing to 100% per column).

3.4. Earnings

Following the optimal decision policy, a DM would expect to earn 30.87 experimental dollars given the random stimuli from this condition. Empirically however, participants on average earned 28.59 experimental dollars, which was significantly less than they could have earned (one sample t-test: $t(40) = -9.02$, $p < 0.001$). The observed return on endowment was 19.1% and for the optimal policy it was 28.6%, yielding a slightly higher ROE than in Experiment 1. Again, there were substantial individual differences in earnings between participants; the standard deviation of earnings was 12.54 with the middle 50% of DMs earning between 19 and 40 experimental dollars.

[6] Blanks in Table 3 result from the impossibility of *underinvesting* given the optimal choice of 0 in the first two stages. In these instances the only possible bias is overinvesting, which happens about 58% of the time in the first two stages.

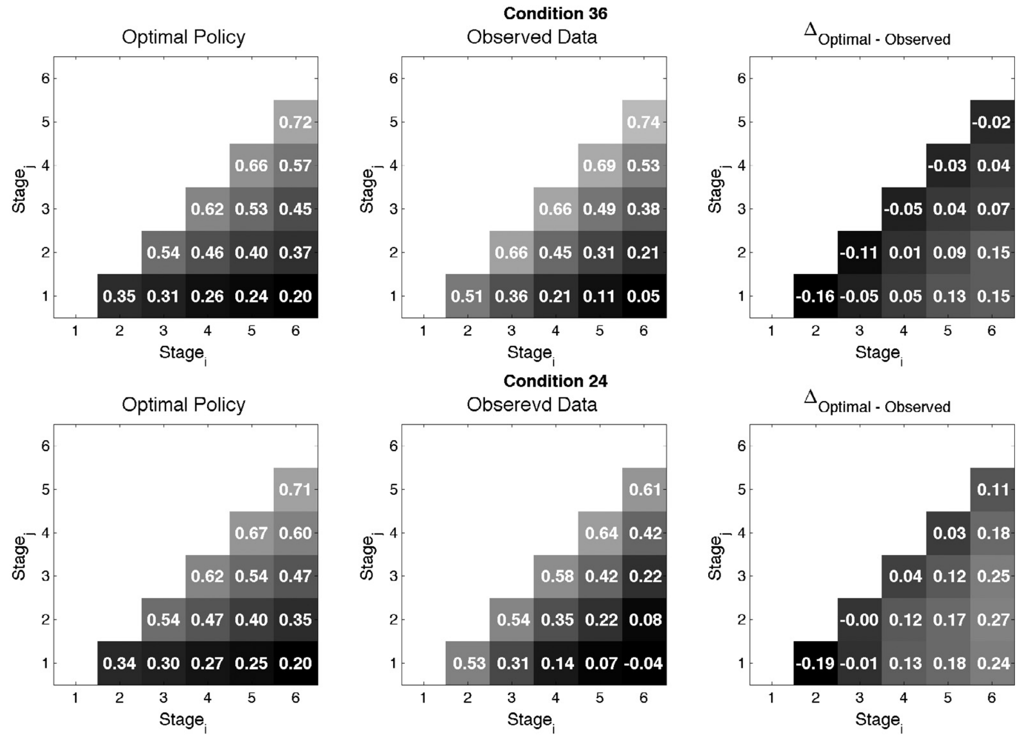

Fig. 4. Correlations between investment amount in stage i and stage j for the optimal policy, and separately for observed investment behavior in the 36-condition (upper row) and 24-condition (lower row). The right most sub-figures show the difference between these two matrices. The darker the color, the weaker the correlation.

4. Sequential dependence and learning

4.1. Sequential dependence

We turn our attention here to the temporal dependence of the investments across stages. We first examine the correlations between stages i and j, where $i, j \in \{1, 2, 3, 4, 5, 6\}$. We compare lag correlations between EV maximizing investment choices (see Section 1.3.2 and Appendix D), and the correlations among the observed choices from participants. Fig. 4 shows these correlations and the difference between them (third column). What can be inferred from Fig. 4 is that the larger the distance between two stages the weaker the correlation between investment choices. For both the optimal policy and the observed choices, we observe the highest correlations between consecutive stages (which is displayed in the diagonal of the matrices in each panel).

An idealized DM would exhibit a positive serial correlation in investment choices between stages, even though the draws are *i.i.d*. This is a perhaps non-obvious consequence of the structure of the decision problem, and sensible heuristics (Gigerenzer and Gaissmaier, 2011) could take advantage of this information structure. If a DM has reason to invest in stage i, he will likely have reason again to invest in stage $i + 1$. This is due to the fact that if the probability of reaching the goal is greater than 0.5 in stage i, it will be more likely to still be above 0.5 in stage $i + 1$ than not. So a degree of investment "inertia" is expected, even from a perfect EV maximizing decision maker. These resulting lag correlations provide the proper baseline for determining if participants exhibited more or less inertia than is warranted by the normative decision policy. The results show that on average, the correlations between stages are higher for the optimal policy than for the observed data (36-condition: $M_{Optimal-Observed} = 0.01$, 24-condition: $M_{Optimal-Observed} = 0.08$), but this difference is significant only for the 24-condition ($t(14) = 3.53$, $p < 0.005$). These findings indicate a weaker sequential dependence between observed choices than would be expected. The sequential dependence is even weaker with the budget restriction (as imposed by the 24-condition), however the effect is not large. The weaker sequential dependence can be in part attributed to DMs' use of in-between investment values.

Table 4
Proportions of over- and underinvestments, given overinvestment (underinvestment) in the previous stage.

i	36			24		
	+/+	−/−	p-value	+/+	−/−	p-value
2	54.0%	55.7%	0.35	–	83.9%	
3	71.0%	47.5%	0.00	–	30.5%	
4	54.5%	51.0%	0.06	60.4%	52.0%	0.00
5	53.7%	36.1%	0.00	54.7%	38.2%	0.00
6	27.4%	27.9%	0.84	27.3%	32.8%	0.01

Here B_i is the observed bias in stage i. "+/+" is defined as $P(B_i > 0 | B_{i-1} > 0)$ and "−/−" as $P(B_i < 0 | B_{i-1} < 0)$. The reported p-value is a two-sided proportion test of the null hypothesis that these conditional probabilities are equal (i.e., "+/+" = "−/−"). The results show that consecutive overinvesting is more prominent than consecutive underinvesting.

Further, to more fully analyze investment "inertia", we computed the lag correlation of bias scores, examining to what degree overinvesting (underinvesting) is conditional on having overinvested (underinvested) in the preceding stage. As shown in Table 4, consecutive overinvesting is more prominent than consecutive underinvesting and more pronounced in mid-stages such that we may conclude that indeed the correlation patterns found in Fig. 4 are mostly due to consecutive overinvesting. This in turn, may suggest that participants – when in doubt about the true probability of success – exhibit a "hot hand" tendency (Gilovich et al., 1985; Croson and Sundali, 2005). Of course, this is not a comprehensive test of the "hot hand" phenomenon, however, the higher than expected correlations between two consecutive stages suggests that DMs are more inclined to stay with a current (overoptimistic) investment policy. In this setting, the "hot hand" would imply investing more when the roll in the previous stage was high (i.e., higher probability of reaching the goal). In these experiments, the participants exhibit a persistent bias that is especially prominent in the early stages (i.e., stages 2 and 3). The bias of over- and underinvesting decreased with stages consistent with resolving subjective uncertainty about reaching the goal.

4.2. Learning

The large number of repetitions of the experimental decision tasks, in conjunction with the clear feedback, provided DMs a chance to learn and perhaps adapt and improve their decision policies. We investigate possible learning in three ways.

First, we analyzed the change of the DMs' efficiency in earnings. Efficiency is quantified as the possible earnings resulting from the optimal decision policy minus the mean earnings of experimental dollars over rounds. This is computed for each DM and earnings' efficiency is regressed on rounds. For the 36-condition, the intercept equals -4.30 and the slope 0.023 ($p < 0.005$), with an explained variance of $R^2 = 0.11$, $F(1, 78) = 9.77$, $p < 0.005$. This indicates that in the early rounds of the task participants are underperforming by about 4 experimental dollars relative to the optimal policy, but by the end of the 80 rounds this underperformance is less and only about 2.5 dollars. For the 24-condition, a regression of earnings over rounds did not yield a significant relationship (slope $\beta = 0.0025$, $p = 0.62$, $R^2 = 0.0031$, $F(1, 78) = 0.25$) indicating no systematic learning in this condition. Generally we see little to no learning that affects earnings, and choice behavior is not clearly converging to the optimal decision policy.[7]

Second, DM's tendency to evenly diversify investments over stages is remarkably robust, even when doing so is not in their best fiscal interest. In the 36-condition, participants tended to follow a non-optimal heuristic: In the first 20 rounds of the experiment, DMs invested 3 in the first stage 38% of the time. For the last 20 rounds of the experiment, DMs still invested 3 in the first stage 34% of the time. Even with full feedback and clear incentives to perform well, participants' tendency to "risk half" remained largely unchanged as indicated by a non-significant result from a paired sample t-test ($t(41) = 1.17$) comparing frequency of 3-investment in the first 20 and the last 20 rounds. In conclusion, the learning effects here are weak, indicating that generally people use consistent policies over time and adapt little over the iteration of repeated rounds.

Third, we computed the proportion of participants in the 24-condition who invested zero in the first two stages. Recall from Section 3.1 that in this condition, not investing in the first two stages is the EV maximizing policy. For the 24-condition, in the first 20 rounds for the first stage, 31% of investment decisions were at the 0. This increased to 47% for the last 20 rounds. However, only 12% of DMs did not invest in the first two stages for the last 20 rounds of the experiment. Conversely, 46% of DMs invested in either the first or second stage for *every round* for the last quarter (20 rounds) of the experiment. On average, participants invested 0 in the first stage in only 16 out of 80 rounds.

[7] The decision problems here have a relatively flat maxima, meaning that the marginal decrement for non-optimal choice behavior is comparatively small. This feature could induce behavior from people that is "close enough" to optimal, but not motivate them to expend additional cognitive efforts to improve payoffs further. Different payoff structures could be implemented experimentally to investigate this. For example generally larger payoffs could be used. Moreover, a payoff structure with a relatively greater decrement for non-optimal choices could be built into a extensions of the real options problem. Both variations may induce greater learning from decision makers but this remains an open question.

5. Model testing

5.1. Major styled facts

Three main results of this experiment deserve particular attention. First, approximately half of all DMs' choices are inconsistent with expected value maximization (*cf.* Appendix D). Second, the most common non-optimal investment was 3, the "risk half" alternative. Third, the distribution of non-optimal investment decisions is skewed such that non-optimal investments of less than 3 are much more frequent than non-optimal investments greater than 3. In other words, we observe participants investing 1 or 2 about twice as often as they invest 5 or 6.

5.2. Descriptive decision models

To provide some insights into the non-EV maximizing choice tendencies displayed by decision makers, we fit and contrast several decision models of risky choice. We use the following models: probability matching (Estes, 1950), expected utility theory (von Neumann and Morgenstern, 1944), prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992), and the $1/N$ diversification heuristic (Benartzi and Thaler, 2001). Other simple (zero-parameter) models are also considered simply to provide additional benchmarks. These include the following models: sure bets only, uniformly random investments, and no investments ever.

Probability matching posits that the attractiveness of an option is proportional to its estimated probability of yielding a positive outcome. In this case, a DM would invest proportionally to the estimated likelihood of success in any stage. For example, if in a given stage the probability of reaching the goal was 0.33, the DM would invest about 1/3 of the possible amount; in this case that would be 2 experimental dollars. The probability matching policy is a linear investment policy over probabilities, whereas the optimal policy is a strict step function. Probability matching has been considered and discussed extensively as a descriptive account of sequential choice behavior (Estes, 1950; Bush and Mosteller, 1955; Estes and Suppes, 1959; Vulkan, 2000).

We consider **expected utility** theory and recognize that it is nested as a unique parameterization of prospect theory (PT) but treat it separately here for clarity. Expected utility theory accounts for risk aversion by positing a concave utility function for money, but it assumes no probability weighting. Herein, we use a power function to represent utility.

Prospect theory posits that a decision maker's utility of money $u(\cdot)$ is a power function that is concave in the domain of gains. Concurrently, PT asserts that a decision maker acts as if they overestimate small probabilities and underestimate large probabilities. The objective probability of a choice outcome is transformed according to a weighting function $\pi(\cdot)$ (see Appendix C for details). Prospect theory is a prominent decision framework that could account for the aggregate results from our experimental task well because it accounts for risk-averse choices, as we commonly observe from the range of investments over rounds. Additionally, it incorporates a non-linear distortion of the objective probability of winning that is consistent with the fact that participants overinvested when the probability of winning was below 0.5. In the experiments outlined in the preceding sections, all of the choices occur in the domain of gains. Thus, we model the utility derived from an investment I_t in a particular stage t as

$$U(I_t) = \pi(p_t) \cdot u(2 \cdot I_t) + u(6 - I_t), \tag{2}$$

where p_t is the objective probability of winning in that stage, $u(\cdot)$ is the power utility function shaped by parameter α, and π is the probability weighting function with one parameter γ, as

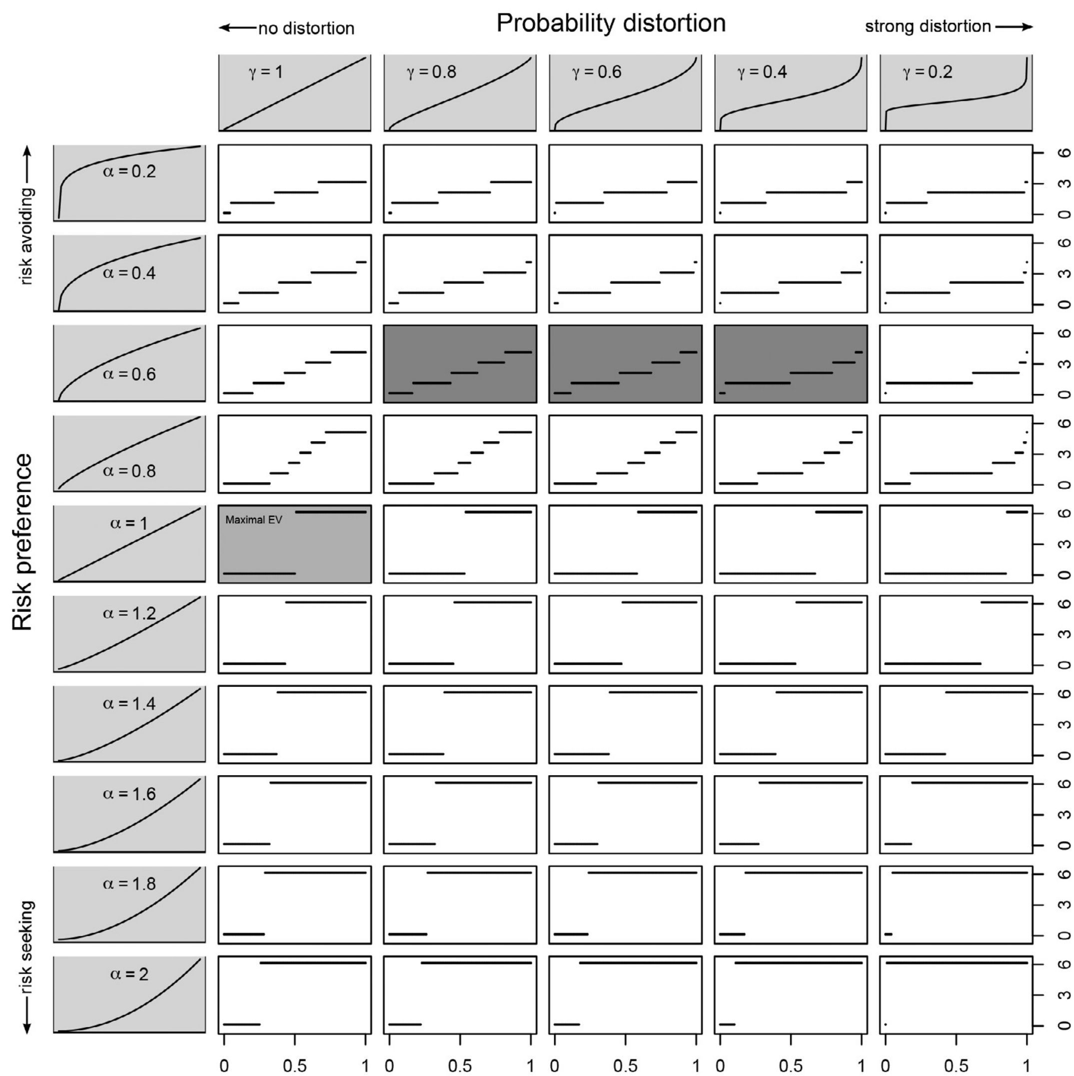

Fig. 5. Each subplot shows the utility maximizing investment policy for a given objective probability of reaching the goal. These utility maximizing policies are conditional on two descriptive parameters, namely risk preference (α) and probability weighting (γ). The expected value maximizing policy is a special case where $\alpha = 1$ and $\gamma = 1$. What is noteworthy is that investing "part-way" (1–5 inclusive) is consistent with risk aversion (i.e., $\alpha < 1$). PT also predicts investing middle values for risk averse DMs. Whereas for risk-loving DMs (i.e., $\alpha \geq 1$), PT predicts investing either 0 or 6, similar to an EV maximizing DM. Moreover, PT's probability weighting ($\gamma < 1$) yields asymmetry in the predicted investment levels over the objective probabilities. Investing 1 and 2 increases in frequency as γ decreases below 1. Taken together, the concave value function and the nonlinear probability weighting function from prospect theory predict the systematic non-optimal investment pattern of choices observed from the experiments. The estimated parameters in Experiment 1 are also consistent with parameter estimates from other empirical work (Booij et al., 2010; Murphy and ten Brincke, in press).

specified by Prelec (1998). Consistent with the approach of Thaler and Johnson (1990), Eq. (2) posits that risky and certain outcomes are valued separately.

For any objective probability of winning, we compute the subjectively "optimal" investment as defined by prospect theory. The utility maximizing investment, assuming utility is derived according to (2), is shown in Fig. 5 where α and γ are assigned different values over common ranges of the parameters. This figure maps combinations of risk aversion and probability weighting to the resulting predicted investments. The revenue maximizing policy defined in Appendix D is just a special case where both parameters equal 1. The empirical results reported in Sections 2–5, particularly Fig. 2, exhibit a non-normative pattern that is better reflected by the shaded subplots that accommodate both risk aversion and probability weighting. Prospect theory can account for the three stylized facts summarized above. DMs tend to: avoid all or nothing investments; invest the middle amount (e.g., 3); more frequently investment small amounts (e.g., 1–2) than large amounts (e.g., 4–5).

The $1/N$ **diversification heuristic**, or naïve diversification as coined by Benartzi and Thaler (2001), would lead a DM to spread his investments equally among the investment options available (see also Read and Loewenstein, 1995).[8] In the 36-condition, the investment should always be 3 using this heuristic: the endowment divided by the number of stages, then that quotient divided by two (in words, for each stage, the DM would keep half their resources safe and risk the other half). Analogously, in the 24-condition, the investment should be 2 over all six stages. However, one could also imagine an alternative $1/N$ investment policy for the 24-condition where the DM never invests in the first 2 stages, having realized that investments then are dominated. This would result in investing 0 in stages 1 and 2 and then investing 3 in each of the remaining stages 4–6. We consider all of these versions of the $1/N$ diversification heuristic when examining experimental data.

[8] Even Nobel laureate Harry Markowitz has reportedly expressed a preference to invest following a $1/N$ heuristics (Zweig, 2008, p. 4) in spite of his knowledge of more sophisticated methods of choice optimization.

5.3. Flexible models

Well established choice models such as prospect theory or utility theory posit their underlying functions a priori. However, there are many possible functions that could account for the pattern of choices observed in practice. Relaxing the assumption of the fixed shapes of the utility function and probability weighting function could potentially result in models with improved predictive power and perhaps give some better insight into the "real" structure of people's preferences. A major motivation to include highly flexible and agnostic models in this paper is to make a "horserace" between the different accounts of non-EV maximizing choice behavior. Only considering two models at a time and showing that one model outperforms another, does not provide convincing evidence for the general quality of a model in understanding risky choice. We strive for a more stringent examination that involves a wide variety of different possible models.

Here, we employ approaches from statistical learning to account for potential choice patterns which may not captured by the other proposed models. These statistical learning models are completely agnostic and know nothing but the data (i.e., supervised learning). Further, these models are non-parametric because they make no assumptions about the shape of the underlying functions. In particular, we apply two types of models, where the first one is nested in the second. First, a regression spline of the investment on the probability of winning, later denoted simply as "spline". And second we consider generalized additive models (abbreviated GAMs; see Hastie and Tibshirani, 1990).

Spline predicts investment in each stage as a function of the objective probability of winning:

$$\hat{I}_t = s(p_t), \tag{3}$$

where $s(\cdot)$ denotes a smoothing 'spline'. A spline is a piecewise polynomial of the Nth order such that $s(x) = \beta_N x^N + \beta_{N-1}x^{N-1} + \cdots\cdots + \beta_0 x^0$ for $\beta_N, \ldots, \beta_0 \in \mathbb{R}$. For every interval between two out of K knots, $s(x)$ can have different values of $\beta_N, \beta_{N-1}, \ldots, \beta_0$. Therefore, each kth interval is described by a different *basis function*. A spline has a continuous first and second derivative at every knot. The more knots, the better the predictive power of the spline. To avoid overfitting the residual sum of squares (RSS) of a spline, there is a function that minimizes the penalized residual sum of squares ($RSS(f, \lambda)$) which consists of two components: RSS measuring the closeness of the prediction to the data, and a penalization criterion of the function's curvature. Therefore, conceptually $RSS(f, \lambda) = RSS + \lambda \cdot \textit{penalization}$, where λ is called the smoothing parameter. $RSS(f, \lambda)$ represents a tradeoff between the model's predictive power and its complexity. The higher the value of the model's degree of freedom, the greater the curvature of the spline's function. We allowed for the automatic selection of the smoothing parameters. For further information on the smoothing spline statistical learning method refer to Hastie and Tibshirani (1990).

Generalized additive models (Hastie and Tibshirani, 1990) are an extension of the regression spline method. Whereas splines possess only one predictor, GAMs can have multiple predictors. A GAM fits to each predictor a nonlinear regression function and explains the dependent variable as the sum of smooth ("nonparametric") functions, i.e., if $\hat{y} = \hat{\beta}_1 x_1 + \cdots + \hat{\beta}_l x_l$ is the multivariate regression prediction, then $\hat{y} = \hat{f}_1(x_1) + \cdots + \hat{f}_k(x_l)$ is the GAM prediction where $\hat{f}_1, \ldots, \hat{f}_k$ are smooth functions and x_1 to x_l are any predictors considered relevant. The estimated model fit is measured with the penalized sum of squares ($PRSS$) that is dependent on $f_1, \ldots, f_k$. GAMs can include, smoothing, linear and non-linear predictors, which gives full flexibility in constructing a model.

We consider two GAMs applicable for this experimental task. The first GAM, denoted GAM(p_t, R_{t-1}), has two predictors of investment—the probability of winning (p_t) and preceding roll (R_{t-1}):

$$\hat{I}_t = s(p_t) + \beta_0 + \beta_1 R_{t-1}. \tag{4}$$

Second model includes four theoretically sound predictors, namely the probability of winning (p_t), preceding roll (R_{t-1}, i.e., $R_{t-1} \in \{1, 2, 3, 4, 5, 6\}$), value of the goal ($G$, i.e., $G \in \{20, 21, 22, 23\}$) and stage (t, i.e., $t \in \{1, 2, 3, 4, 5, 6\}$):

$$\hat{I}_t = s(p_t) + \beta_0 + \beta_1 R_{t-1} + \beta_2 G + \beta_3 t. \tag{5}$$

In the spline and both GAMs, we expect that $s(p)$ could account for the curvature of the probability weighting. However, the main advantage of this nonparametric method is that it makes no a priori assumption about the shape of the probability weighting function. Predictors G, t and R_{t-1} are introduced to the model as linear predictors for reasons of parsimony. All of this information (i.e., reaching the goal value, preceding roll, and stage), apart from the objective probability of eventually winning, were provided to the participants on the computer screen.

5.4. Model comparison

A natural question that arises when contrasting models is what is a fair measure of quality when making comparisons between very different representations. Classically a criterion such as variance explained or a similar measure is used to assess the quality of a model. However, when comparing models of different classes, the variance explained might no longer be as useful as a measure. Thus, we employ the concept of a training set and a test set[9] and compare the predictions of the models out-of-sample on a test set. The need for using out-of-sample comparisons is a consequence of these models being very different in their structure and in their statistical complexity, hence, when using an in-sample criterion (such as variance explained) one should correct for the complexity and flexibility of the model (Myung, 2000).[10]

As a performance indicator we chose the mean squared error (MSE, the difference of the squared deviation of the actual investment from the predicted investment), percentage incorrectly predicted (%W, which is defined as 100% minus the fraction of observations where the predicted investment and the actual investment coincide), and percentage incorrectly predicted within 1 dollar (%W ± 1, 100% minus the fraction of observations where the predicted investment and actual investment are no further than 1 dollar apart).

Using the training dataset, we fit (i.e., tuned) five models: expected utility theory, prospect theory, a spline, and two GAMs. The training dataset consisted of 70% of the randomly selected data from the full dataset. Next, we measured the models' predictions fits to the test data, which corresponded to the remaining 30% of the data. The full dataset was aggregated over all participants and consisted of 20,154 observations.[11] Additionally, we measured the fit of seven models that have no free parameters: PT with typical fixed parameters $\alpha = 0.8$, $\gamma = 0.65$ (Booij et al., 2010;

[9] This in-sample and out-of-sample approach is commonly used in statistical learning. See for example Hastie et al. (2009) for more details.

[10] In the classical regression case this entails employing the adjusted R^2 or, and more generally, a *penalized information criterion*.

[11] The reason for the missing 6 choices is that for one of the subjects, the experimental program terminated at the start of her 80th round. She was the only subject for whom this happened and we elected to keep her data included in the analysis as it was almost complete.

Table 5
Parametric coefficients of GAMs.

Predictor	Condition 36		Condition 24	
	β	t-value	β	t-value
GAM				
Intercept	−1.31 (0.30)*	−4.32	−1.59 (0.32)*	−4.91
R_{t-1}	−0.04 (0.01)*	−4.20	0.00 (0.01)*	0.11
G	0.15 (0.01)*	11.21	0.12 (0.01)*	8.39
t	−0.06 (0.01)*	4.65	0.20 (0.02)*	12.63
GAM (p_t, R_{t-1})				
Intercept	2.25 (0.04)*	−4.32	1.91 (0.04)*	47.74
R_{t-1}	−0.02 (0.01)	−4.20	0.02 (0.01)	1.78

Note: Numbers in brackets represent standard errors.
* $p < 0.001$.

Glöckner and Pachur, 2012), probability matching, $1/N$ heuristic, optimal investment, sure bets, random investment, and a never invest policy, to the test dataset and to the full dataset. GAMs and spline also have no free parameters, but the spline functions have to be estimated in the course of statistical training. All model fits are presented in Table 6.

The estimated parameters of prospect theory were $\alpha = 0.69$ and $\gamma = 0.87$ in the 36-condition ($\alpha = 0.69$, $\gamma = 0.52$ in the 24-condition). The estimated parameter of the expected utility function is $\alpha = 0.67$ in the 36-condition and $\alpha = 0.64$ in the 24-condition.[12] Risk aversion by itself is helpful in accounting for the pattern of behavioral results but alone is not sufficient to capture the overall pattern.

We fit spline, GAM and GAM (p_t, R_{t-1}) to the training data. The estimated degrees of freedom of the smoothing parameters $s(p_t)$ were 5.39, 5.7 and 6.31 in the 36-condition, and 5.67, 7.41 and 8.37 in the 24-condition.[13] This indicates that in the 24-condition the estimated splines had stronger curvature. In both conditions, a simple spline had the fewest degrees of freedom and similar fit as the GAMs. When comparing GAM to GAM (p_t, R_{t-1}), we find the latter have greater degrees of freedom and worse fit than GAM, which indicates that the linear predictors G and t lead to lower curvature of the $s(p_t)$ function and improvement in the predictive performance of the model. Also, as shown in Table 5, the preceding roll was not a significant predictor in GAM (p_t, R_{t-1}) and had a very low coefficient in GAM. In both GAMs, $s(p_t)$ was a significant predictor (GAM: $F = 1396, p < 0.001$, GAM (p_t, R_{t-1}): $F = 1397, p < 0.001$ in the 36-condition and GAM: $F = 967, p < 0.001$, GAM(p_t, R_{t-1}): $F = 990, 5, p < 0.001$ in the 24-condition). In spline, GAM and GAM (p_t, R_{t-1}) the estimated $s(p_t)$ was a monotonically increasing concave function indicating overall underweighting of the objective probabilities. This is in line with the probability weighting function of prospect theory, as described by Kahneman and Tversky (1979).

Considering the results of Table 6 we find, unsurprisingly, that methods which were tuned to the data work better than methods which "see the data for the first time". However, two results are particularly noteworthy: First, prospect theory, if allowed to be adjusted for the degree of concavity of the utility function and the non-linear weighting of the probability weighting, i.e., tuned, works almost as well as the completely flexible and wholly agnostic GAM in both the 24- and 36-conditions. Second, for a true ex-ante methodology, PT does best in class.[14] The reason why PT works fairly well (or, vice versa, why a GAM only works slightly better than PT) is explained by the fact that the GAM fits the investment choices in a very similar way as predicted by PT. In essence, the GAM "rediscovers" prospect theory. The agreement of the most general GAM model with PT is as high as 86% for the 36-condition and 74% for the 24-condition.[15] Moreover, all nonparametric methods used here show high agreement with PT. Conversely, they fail to support heuristics such as the $1/N$ rule. This result can be inferred from Table 6 which highlights that risk-aversion and probability weighting are necessary to adequately account for the observed behavior.

To address the point why tuned prospect theory works better than the original prospect theory parameters we refer to the well-known fact that the estimates of the parameter values in PT (e.g., the concavity of the utility function and the probability weighting) vary significantly in the literature, which has been surveyed and collected in Booij et al. (2010). Thus parameters are expected to differ at least in some degree from experiment to experiment. Nonetheless, median "off the shelf" parameters from PT produce a good first approximation of the empirical results here.

6. Concluding discussion

Flexibility in making decisions is valuable. Decision makers can benefit from the capacity to delay making commitments and further by being able to radically change their investment levels between distinct stages of dynamic decision problems. Flexibility facilitates the efficient allocation of capital and allows DMs to be responsive in a dynamic environment, using as much new information as possible before committing limited resources and thus avoiding unproductive exposure to risk. In spite of these benefits, our results show that DMs have pronounced tendencies to squander flexibility and not fully capitalize upon options as a source of value. Instead of greatly changing their investment levels between stages, as is sometimes dictated by the normative decision policy, evidence shows that DMs select "middle of the road" options (cf. Bearden et al., 2006) and avoid extremes in their sequential investments, especially when subjective uncertainty is highest. DMs' non-optimal propensity to avoid extremes in these decision contexts is consistent with risk aversion (committing half way as a means of mitigating risk), and the observed downward shift of investments is consistent with probability weighting. We also find that DMs make irrevocable commitments of resources before they should. The result of these tendencies is a substantial reduction in earnings and significantly lower returns on investments than would be secured under the optimal decision policy, which fully exploits the available flexibility.

Naïve diversification has its benefits as a choice heuristic, and in many risky environments it serves (although inexactly) to mitigate risk. However, diversification is not a panacea, as some risky environments require nimble adjustments of investment policies with the arrival of new information. The tendency of DMs to uniformly spread resources across stages may seem sensible at

12 Note that these parameters are not unique: Given the stepwise nature of the problem, as illustrated in Fig. 5, there are ε-neighborhoods of these parameters which have equally good fits. The reported parameters here are approximately in the middle of the best fitting ε-neighborhood. Moreover, Fig. 5 remains a useful representation as for each of the combinations shown there is a unique, utility maximizing investment policy.

13 For estimation of spline and GAMs, we used default R-functions *smooth.spline* and *gam* that correspond to the functions described in Hastie et al. (2009). Both methods use a Gaussian smoothing kernel. GAMs were estimated using a general cross validation optimization method. No initial values for the estimation parameters were used.

14 Except for the hard criterion %W where the non-superiority of PT can be attributed to the high proportion of investing at the extreme values.

15 Note that a GAM model might fit any kind of curve as the optimal investment, also, should the data require it, a highly convoluted one. The agreement can also be seen by visual inspection of the fitted GAM model which exhibits almost identical shapes as shown in Fig. 5. The complete coincidence analyses are presented in Appendix E and found in Table E.1.

Table 6
Model fits.

		36 condition						24 condition					
		Test dataset			Full dataset			Test dataset			Full dataset		
		MSE	%W	%W ± 1	MSE	%W	%W ± 1	MSE	%W	%W ± 1	MSE	%W	%W ± 1
Trained	(a) GAM (all)	**2.56**	0.67	0.30				**2.92**	**0.63**	0.40			
	(b) GAM (p, *prec*)	2.58	**0.64**	**0.29**				2.96	0.67	0.33			
	(c) Spline	2.70	0.69	0.33				2.98	0.64	0.33			
	(d) Expected utility theory	2.73	0.72	0.33				3.08	0.65	0.34			
	(i.T) Prospect theory (tuned)	2.72	0.72	0.32				2.96	0.67	0.33			
Untrained	(i) Prospect theory	**2.86**	0.69	**0.31**	**2.87**	0.69	**0.31**	**3.17**	0.63	**0.33**	**3.18**	0.62	**0.33**
	(ii) Probability matching	3.46	0.69	0.37	3.52	0.68	0.37	3.94	0.63	0.40	3.95	0.63	0.40
	(iii.1) 1/N	5.70	0.85	0.69	5.68	0.85	0.69	5.36	0.80	0.64	5.39	0.79	0.64
	(iii.2) 1/N (2 only)							4.17	0.87	0.62	4.19	0.87	0.62
	(iv) Optimal investment	6.65	**0.54**	0.43	6.83	**0.54**	0.44	5.58	**0.53**	0.42	5.63	**0.53**	0.42
	(v) Sure Bets	8.17	0.57	0.47	7.85	0.56	0.46	5.67	0.53	0.41	5.60	0.53	0.40
	(vi) Random investment	9.41	0.85	0.64	9.65	0.85	0.65	7.96	0.77	0.58	8.02	0.76	0.58
	(vii) No investment	10.08	0.62	0.52	9.85	0.62	0.52	7.37	0.58	0.45	7.34	0.58	0.45

Note: Models which have been trained on parts of the data should only be evaluated on a separate part of the dataset (i.e., the test dataset). The best performing model for a given criterion and condition is highlighted in bold typeface. There are two versions of the $1/N$ model: one assuming an investment of 3 in all stages of the 36-condition, and an investment of 0 in the first two stages and 3 in the remaining stage of the 24-condition, and second assuming an investment of 2 in all stages. The second version is applicable only in the 24-condition.

first, may be easy to justify, and may have an intuitive appeal. This heuristic is obviously easy to understand and easy for DMs to implement. Nonetheless, this static tendency squanders value and in the long run undermines earnings. Furthermore, the non-optimal behaviors we observed in our experiments were largely impervious to learning, even with clear incentives and explicit feedback. Biased decisions persistently occurred across all eighty rounds and showed little sign of abating.

The successful implementation of a real options (EV maximizing) policy depends on a responsiveness that most people did not display, even in a well controlled, incentivized, and plainly structured setting. However, when limiting a DM's freedom by imposing a budget constraint, we find slightly better performance, suggesting that some DMs are able to grasp the deep structure of the decision task and adjust their behavior accordingly. But overall the results call into question untrained DMs' ability to successfully navigate real options problems in more complex and demanding environments, especially those in the real world contexts highlighted above. When training decision makers to use real options valuation methods, it may be prudent from a pedagogical standpoint to also have students (e.g., MBA, undergraduates) make risky choices like the well-defined problems above so as to highlight the tension between the optimal policy and the innate behavioral tendencies most people exhibit. Having a simplified decision model with a clear solution can serve as a starting point for better understanding preferences, risk, and the consequences of different decision policies. This experience may allow students to align their intuitions with normative dictates and develop a deeper understanding and appreciation of the value of flexibility in dynamic choice environments.

The results also show that forcing DMs to reflect upon the decision task more thoroughly (as in the limited budget condition) may indeed lead to improved performance. This fact may serve as a starting point for future research examining the extent to which DMs can be supported and nudged toward acting in accordance with the optimal policy—specifically, what decision aids could be introduced that encourage DMs to take full advantage of decision flexibility in an environment and thereby allow them to fully capitalize on the value of a dynamic decision context. Given the importance and magnitude of such real option decision situations outside the laboratory, it of utmost relevance to understand the behavior of DMs in a real options context such that non-optimal behavior, which diminishes the overall gain, can be remedied.

In sum, this paper investigates people's investment decisions in a stylized real options setting, which allows for testing behavior in a laboratory environment while preserving the defining characteristics of a real options problem: the task is a sequential decision problem with a Markovian structure, demanding that DMs repeatedly make tradeoffs between sure things and risky prospects. Many real world problems share these properties. Our experimental results show that DMs behave in non-optimal ways by routinely overinvesting in the early stages of the decision task and underinvesting in the late stages of the task. These results can be captured by prospect theory, requiring both its utility curvature and non-linear probability weighting. This substantiates the applicability of prospect theory to dynamic decision contexts, which is a useful extension for a theoretical framework designed for static decision tasks. In this case, PT can account for non-optimal middle choices when it assumes that a risk-averse DM strongly undervalues high outcomes ($\alpha < 1$) and that the DM strongly overweights small probabilities and underweights small probabilities of outcomes ($\gamma < 1$). However, PT will always predict choosing 0 as more likely than choosing 6. This is not the case for the generalized additive models (GAMs), which do not pre-define the shape of the probability weighting function. In this task, GAMs predict uniform underweighting of the probabilities, which results in a better prediction of middle values here.

Additionally, we demonstrate the applicability and usefulness of comparative evaluations with GAMs. The results show how highly flexible, parameter-free methods can be constructively used to simultaneously evaluate a wide range of decision-making models. In this instance results provides support for prospect theory by "rediscovering" its central features. Lastly, our experiments indicate that in real options settings, people behave inconsistently with the expectation maximizing decision policy and exhibit persistent and systematic departures from optimal behavior. These behavioral tendencies can best be captured by cognitively inspired models positing risk aversion and probability weighting. Lastly, people's innate tendencies when making risky choices are often contrary to the dictates of real options analysis, and this incongruity could help explain, in part, the mixed success of real options analysis in real world applications.

Appendix A. User interface screen shot

See Fig. A.1.

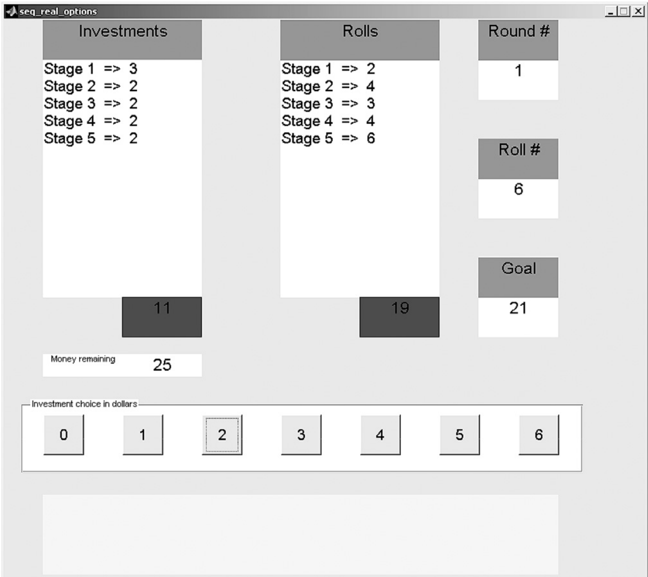

Fig. A.1. A screen shot of the decision task's user interface. Decision makers used the mouse to record their choices and the screen updated relevant information in real time.

Appendix B. Instructions from the experiment

Instructions text from Experiment 1

The present experiment is concerned with the decision making process of choosing to invest money in a risky option as uncertainty about the prospect decreases over time. The instructions for the experiment are rather simple. If you follow them carefully and make good decisions you can expect to earn a decent amount of money. Payments for this experiment typically range between S\$5 and S\$24 and the research session typically lasts between 45 and 60 min. In this decision task more skilled decision makers earn more money. Hence, it is worth paying attention to the instructions, thinking carefully, and doing your best. Your earnings for the session will be paid to you in cash when the experiment is over. A national research foundation has provided the funds to finance this study.

The decision task

In this task, you—the decision maker starts with an endowment of 36 dollars. You will have the chance to invest this money in six different stages. In each stage, the computer will roll a perfectly fair die (each of the values 1–6 are equally likely) and keep a running sum of the results from each of the rolls. Before every roll, you can invest between \$0 and \$6 from your endowment, but no more on any one stage. The money you invest is held until the end of the 6 rolls. If the sum of the rolls is equal to or greater than the goal of 21, then the amount that has been invested will be doubled and paid back to you—the decision maker. If the sum of the rolls is not at least the goal of 21, then all the money that was invested disappears. In any case, any money that is not invested is kept by the decision maker at the of the round.

Screen shots and an example

Following are screen shots of the decision task. We will use these here to go over an example of one round of this decision task.

At the start of a round, a decision maker sees the task screen:

Screen Shot [16]

First, the decision maker has to decide how much to invest on the first stage. None of the rolls has happened yet and the goal is 21.

[16] The actual screen shots used in the instructions are excluded here to save space. The full instructions with all the example screen shots are available upon request.

In this example, the decision maker invests \$3. In order to register this choice, the 3 box is clicked on by the decision maker.

Screen Shot

After the 3 box is clicked upon, the computer "rolls the die" and randomly draws a number from between 1 and 6 with uniform probability, just like rolling a perfectly fair die. The result from the first roll happens to be 4. After seeing this, the decision maker decides to invest more money, and so chooses 5 on the next stage of the task. The decision maker does so by clicking the 5 box.

Screen Shot

The computer then rolls the die and the result is 3, for a running total of 7. Remember the goal is 21. The running total of the die rolls must reach the goal in order for the investments to payoff. Otherwise, whatever is invested will be lost.

Screen Shot

Next we can see the decision maker invested \$5. So far the decision maker has invested \$13 in total (leaving \$23 left of the initial endowment).

Screen Shot

On the 4th stage, the decision maker invests nothing. The computer still rolls the die and the running total of rolls reaches 13.

Screen Shot

On the 5th stage, the decision maker invests only \$1. The computer then rolls the die and this results in a 3. The running total of rolls is 16 now, which is far away from the goal of 21. It is unlikely (but not impossible) that the goal will be reached in this round.

Screen Shot

On the last stage, the decision maker invests nothing. The last roll is a 3, which brings the running total to 19. After the last roll of the die, the round ends. The goal has not been reached, therefore whatever money that was invested is lost and the decision maker earns the remainder of their endowment that they did not invest. On this round, the decision maker invested \$14. This investment did not payoff and was lost. Thus, the decision maker earned the reminder of their endowment, which in this case was \$22. Formally, their payoff is computed: $(14 \times 0) + 22 = 22$.

Screen Shot

Conversely, if the running total of rolls had reached the goal of 21, the decision maker would have earned $(14 \times 2) + 22 = 50$. If the goal is reached or exceeded, then the total amount invested is multiplied by 2 and is returned to the decision maker.

Iterations over the experimental session

You will be presented with a total of 80 rounds like the one described above. Each of the rounds is independent. On each round you will start with an endowment of 36.

Repeating the task will give you a chance to participate in this decision problem multiple times and develop intuitions about what makes for good decisions in this environment. The 80 rounds will be identical except of the goal will change between values of 20, 21, 22, or 23. The value of the goal will be clearly displayed at the beginning and during each round.

There is no deception used in this experiment. There is nothing hidden nor is there any trickery in this decision problem. The Research Center can verify this fact as can your professor. If you would like to see the source code for the decision task, talk to one of the experimenters after the session and you can review the source code for the decision task. Please also feel free to ask any questions about this decision task or about the research in general. We are happy to answer any questions you might have about this research project.

Table D.1
Explanation of notation used in the formal description of the problem and the derivation of the optimal policy. The notation is listed in the alphabetic order.

Symbols	Explanation
E	Initial endowment (e.g., 36)
G	Goal of the sum of rolls (e.g., 21)
I_t	Investment at time t
L	Per stage investment limit (e.g., 6)
Π	Payoff of the decision task that is dependent on $\sum_{t=1}^{T} R_t$ and $\sum_{t=1}^{T} I_t$
R_t	Outcome of a die roll at stage t, such that $R_t \in \{1, 2, 3, 4, 5, 6\}$
$t \in \{1, \dots, T\}$	Discrete decision stages, such that $T = 6$
X_t	Sum of rolls at stage t, such that $t \leq T$ and $t \in \{1, 2, 3, 4, 5, 6\}$
X_T	Sum of rolls at the end of the game
$\{I_1^*, I_2^*, I_3^*, I_4^*, I_5^*, I_6^*\}$	Optimal decision policy; a set of 6 single investments.
$\mathbb{1}_{\{X_T \geq G\}}$	Indicator function yielding 1 if the goal is reached, 0 otherwise

You will not be paid for every round in this experiment; this is simply too costly for us to do. Instead, you will be paid for one randomly selected round. At the end of the experiment, you will randomly draw a number between 1 and 80 from a deck of cards. That number will determine your payment round. You will be paid 1/3 of whatever you earned on the chosen round. It is clearly in your best interest to do your best on every trial. You can earn and be paid in cash between S$0 and S$24 by participating in this research.

If you wish to take notes during the experiment, please feel free to do so on the back of these instructions or other scratch paper. Please do not use a calculator or computer during the experiment. Also, there should be no talking or interactions between research subjects during this research.

If you have any questions about this experiment please ask the research administrator now. Also feel free to raise your hand at anytime during the experiment. The research administrator will come to you and privately answer your queries.

Good luck and thank you for participating in this research project.

Appendix C. Mathematical model specifications

Prospect theory

The utility function of a choice outcome x is defined as a power function of the objective value of outcome x:

$$u(x) = x^{\alpha}. \tag{C.1}$$

Depending on the value of α, the utility function could be convex or concave. While estimating α, we did not constrain it. In the experiments presented in this paper, possible outcomes occurred only in the domain of gains. Therefore, we did not include parameter β as originally described in Kahneman and Tversky (1979).

The probability weighting function $\pi(p_t)$ we used is the one parameter version from Prelec (1998):

$$\pi(p_t) = \exp[-(-\ln(p_t))^{\gamma}], \tag{C.2}$$

where p_t is the objective probability of winning at time t as described in Fig. 1. The expected prospect theory utility of an investment choice I at time t is defined as:

$$U(I_t) = \pi(p_t) \cdot (2 \cdot I_t)^{\alpha} + (6 - I_t)^{\alpha} + (1 - \pi(p_t)) \cdot 0 \cdot (I_t)^{\alpha}. \tag{C.3}$$

Therefore, (C.3) reduces to:

$$U(I_t) = \pi(p_t) \cdot (2 \cdot I_t)^{\alpha} + (6 - I_t)^{\alpha}. \tag{C.4}$$

The choice predicted ($\hat{I}_t$) by the model is simply the investment option that maximizes the prospect theory utility:

$$\hat{I}_t = \arg\max_{I_t} \; U(I_t). \tag{C.5}$$

Expected utility theory

The expected utility theory is a nested model of prospect theory. The expected utility is defined as in (C.4), where $\pi(p_t) = p_t$.

Probability matching

In each stage of a round, the predicted choice equals the nearest lowest integer of the product of the distorted probability and 7:

$$\hat{I}_t = \lfloor p_t \cdot 7 \rfloor. \tag{C.6}$$

In the special case where $p_t = 1$, $\hat{I}_t$ is set equal to 6. Therefore, $\hat{I}_t \in \{0, 6\}$. The higher the probability of winning, the more money a decision maker should invest. The decision strategy is updated in every stage due to the dynamic nature of the probability of winning.

Appendix D. Formal derivation of the optimal decision policy

The optimal decision policy for this problem can be found using dynamic programming. For the problem here, we restrict our attention to the case where there are $T = 6$ stages and the endowment E is fixed at 36 units such that the only binding constraint is the per stage limit $M = 6$ and not the total endowment E. This problem resembles a classical Markov decision process, implicating Bellman's principle of optimality (Bellman, 1957) as a means of obtaining a solution. However, the problem at hand is slightly different in the sense that there are neither payoffs at intermediate stages nor do the actions of the DM affect the transition probabilities. Hence, our problem is a special instance of a Markov decision process that is straightforward to address analytically. Table D.1 summarizes the notation used in this section.

Assuming risk-neutrality, the problem a DM faces is defining an optimal policy such that Π, as defined in (1) has the maximum expected value. Consequently, a DM's optimal policy, $\{I_1^*, \dots, I_6^*\}$, is such that:

$$\{I_1^*, \dots, I_6^*\} = \arg\max\{I_1, \dots, I_6\} \, \mathrm{E}[\Pi]. \tag{D.1}$$

Then, as in the classic Bellman case, any optimal policy $\{I_1^*, \dots, I_6^*\}$ cannot be optimal if the investment I_t^* is not optimal in each stage t. To this end, we define expectation at time t to be conditional on previous rolls and previous investments:

$$\mathrm{E}_t[\cdot] := \mathrm{E}[\cdot | R_{t-1}, I_{t-1}, \dots, R_1, I_1]. \tag{D.2}$$

Now, let $f(X_T)$ be a random variable that depends on the outcome of X_T, where X_T is the sum of rolls at the end of the game. X_T and $I_1, \dots, I_t$ are independent for any t because the investments do not influence the sum of rolls and the sum of rolls at time t is updated after the investment I_t is made. That makes $f(X_t)$ dependent only on the preceding rolls. Therefore, the expectation of $f(X_T)$ is a

Table E.1
Coincidence matrices for prediction models.

	(a)	(b)	(c)	(i.T)	(ii.T)	(i)	(ii)	(iii.1)	(iii.2)	(iv)	(v)	(vi)	(vii)
A: Coincidence matrix for condition 36													
(a) GAM (all)	100												
(b) GAM (prec., proba.)	87	100											
(c) Spline	84	94	100										
(i.T) Prospect theory (tuned)	86	97	92	100									
(ii.T) Probability matching (tuned)	59	60	58	61	100								
(i) Prospect theory	70	68	73	66	47	100							
(ii) Probability matching	54	55	50	56	53	48	100						
(iii.1) 1/N	42	43	42	43	50	39	42	100					
(iii.2) 1/N (2 only)													
(iv) Optimal investment	50	46	51	45	46	56	61	37	0	100			
(v) Sure bets	50	46	51	45	46	56	52	33	0	71	100		
(vi) Random investment	43	43	43	43	43	43	43	43	9	43	43	100	
(vii) No investment	49	46	51	45	40	56	45	33	0	65	94	43	100
B: Coincidence matrix for condition 24													
(a) GAM (all)	100												
(b) GAM (prec., proba.)	85	100											
(c) Spline	85	93	100										
(i.T) Prospect theory (tuned)	74	80	79	100									
(ii.T) Probability matching (tuned)	52	52	58	45	100								
(i) Prospect theory	60	65	62	76	26	100							
(ii) Probability matching	32	30	29	16	41	15	100						
(iii.1) 1/N	15	15	17	16	28	12	19	100					
(iii.2) 1/N (2 only)	18	17	25	18	35	11	17	0	100				
(iv) Optimal investment	24	28	21	19	12	24	29	6	0	100			
(v) Sure bets	24	27	21	19	12	24	18	0	0	53	100		
(vi) Random investment	14	14	14	14	14	14	15	14	14	15	14	100	
(vii) No investment	18	20	15	19	6	24	12	0	0	47	94	14	100

Note: Both tables report the percentage of any two investment models predicting the same investments on the full dataset. The labels correspond to the ones introduced in Section 5.2 and Table 6.

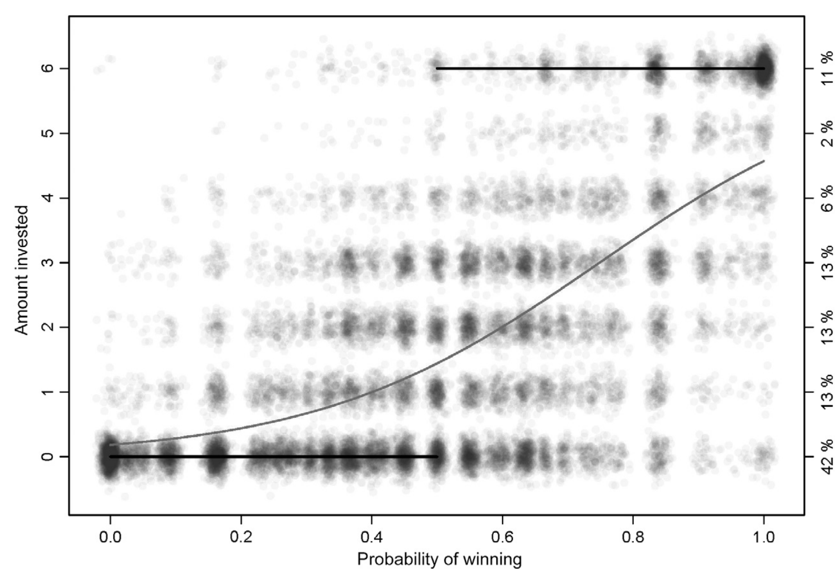

Fig. F.1. The amount invested over the probabilities of win for an endowment of 24 dollars. The percentages given at the right axis indicate the relative frequency of choices of the different investment levels. The optimal decision policy for this problem is a step function with two levels: for probability of winning between 0 and 0.5, the optimal investment is 0, whereas if the probability of winning is greater than 0.5, the optimal investment should be 6. For the instances where probability of winning equals exactly 0.5, the normative decision maker would be indifferent across all investment options.

conditional expectation dependent on rolls at $t-1 \ldots, 1$, such that

$$E_t[f(X_T)] = E[f(X_T)|R_{t-1}, \ldots, R_1]. \quad \text{(D.3)}$$

Consequently, the optimal policy solution in any stage t is defined by the arguments that, in each stage, it gives the highest expected outcome from the investment conditional on the sum of preceding rolls, which results in the following simplification of the problem:

$$\begin{aligned} I_t^* &= \arg\max_{I_t} E_t[\Pi] = \arg\max_{I_t} \sum_{i=1}^{t} E_t[(2 \cdot \mathbb{1}_{\{X_T \geq G\}} - 1)I_i|X_{t-1}] \\ &= \arg\max_{I_t} E_t[2 \cdot \mathbb{1}_{\{X_T \geq G\}} - 1|X_{t-1}]I_t. \end{aligned} \quad \text{(D.4)}$$

Therefore, if the goal is reached, $\mathbb{1}_{\{X_T \geq G\}} = 1$, (D.4) reduces to $\arg\max_{I_t} \sum_{i=1}^{t} E_t[2|X_{t-1}]I_t$, whereas when the goal is not reached

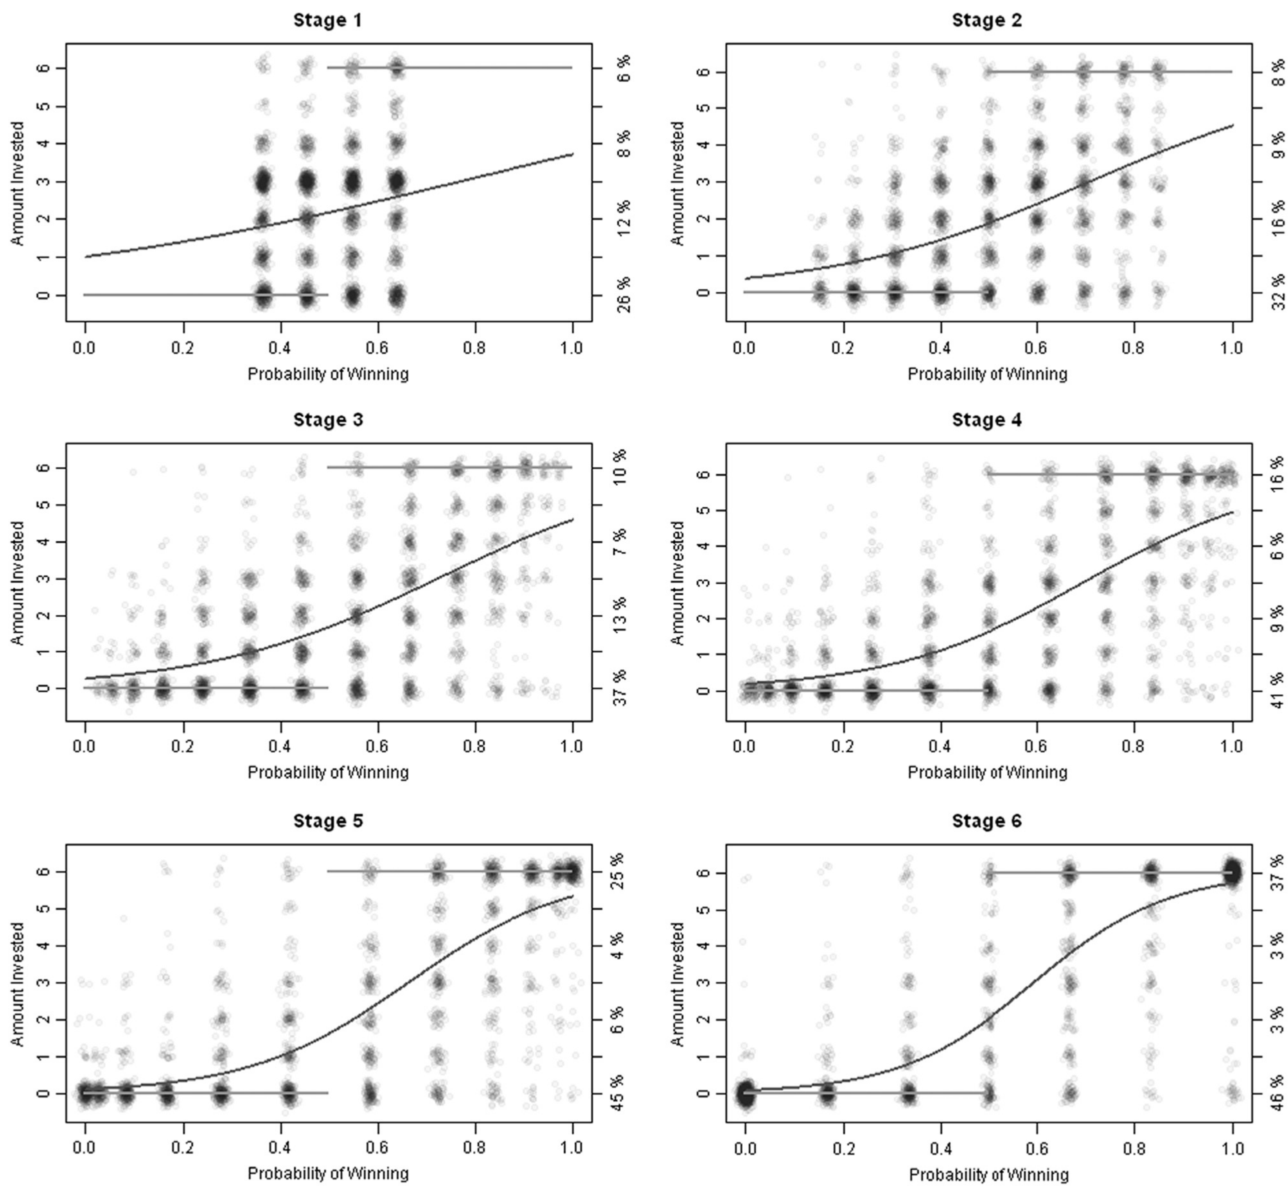

Fig. F.2. The plots shown are similar to the illustration presented in Fig. 2 but here are split across stages. DMs tend to choose closer to the optimal decision policy the closer they move toward the final roll, resulting in a more pronounced *S*-shaped regression fit as the stage number increases.

and $\mathbb{1}_{\{X_T \geq G\}} = 0$, (D.4) reduces to $\arg\max_{I_t} \sum_{i=1}^{t} \mathrm{E}_t[-1|X_{t-1}]I_i$. By combining (D.1) and (D.4), we define the optimal I_t^* that maximizes the payoff, given the sum of rolls in the preceding stage:

$$\begin{aligned} I_t^* &= \arg\max_{I_t} \mathrm{E}_t[2 \cdot \mathbb{1}_{\{X_T \geq G\}} - 1|X_{t-1}]I_t \\ &= \arg\max_{I_t} \big(2 \cdot \Pi[X_T \geq G|X_{t-1}] - 1\big)I_t \end{aligned} \tag{D.5}$$

by the fact that $\mathrm{E}[\mathbb{1}_A|B] = \mathrm{P}[A|B]$, where A and B are measurable sets, we conclude that

$$I_t^* = \begin{cases} 6 & \text{if } 2 \cdot \mathrm{P}[X_T \geq G|X_{t-1}] - 1 > 0 \\ & \Leftrightarrow \mathrm{P}[X_T \geq G|X_{t-1}] > \dfrac{1}{2} \\ 0 & \text{if } \mathrm{P}[X_T \geq G|X_{t-1}] - 1 \leq 0 \\ & \Leftrightarrow \mathrm{P}[X_T \geq G|X_{t-1}] \leq \dfrac{1}{2}. \end{cases} \tag{D.6}$$

Appendix E. Coincidence measures

See Table E.1.

Appendix F. Additional figures

See Figs. F.1 and F.2.

Appendix G. Supplementary data

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.jbef.2016.08.002.

References

Babajide, A., 2007. Real options analysis as a decision tool in oil field developments (Ph.D. thesis), Massachusetts Institute of Technology.

Barberis, N., 2013. Thirty years of prospect theory in economics: A review and assessment. J. Econ. Perspect. 27 (1), 173–195.

Bearden, J.N., Rapoport, A., Murphy, R.O., 2006. Experimental studies of sequential selection and assignment with relative ranks. J. Behav. Decis. Mak. 19 (3), 229–250.

Bellman, R., 1957. Dynamic Programming. Princeton University Press, Princeton, NJ, USA.

Benartzi, S., Thaler, R., 2001. Naïve diversification strategies in defined contribution saving plans. Amer. Econ. Rev. 91 (1), 79–98.

Booij, A.S., van Praag, B.M.S., van de Kuilen, G., 2010. A parametric analysis of prospect theory's functionals for the general population. Theory and Decision 68 (1–2), 115–148.
Bowman, E.H., Hurry, D., 1993. Strategy through the option lens: An integrated view of resource investments and the incremental-choice process. Acad. Manag. Rev. 18 (4), 760–782.
Bush, R.R., Mosteller, F., 1955. Stochastic Models for Learning. John Wiley & Sons, Inc., Oxford, England.
Camerer, C., 1995. Individual decision making. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, pp. 587–703.
Cortazar, G., Schwartz, E.S., Salinas, M., 1998. Evaluating environmental investments: A real options approach. Manage. Sci. 44 (8), 1059–1070.
Croson, R., Sundali, J., 2005. The gambler fallacy and the hot hand: Empirical data from casinos. J. Risk Uncertain. 30 (3), 195–209.
Dess, G., Ireland, R., Zahra, S., Floyd, S., Janney, J., Lane, P., 2003. Emerging issues in corporate entrepreneurship. J. Manag. 29 (3), 351–378.
Edwards, W., 1962. Dynamic decision theory and probabilistic information processing. Hum. Factors 4, 59–73.
Estes, W., 1950. Toward a statistical theory of learning. Psychol. Rev. 57 (2), 94–107.
Estes, W.K., Suppes, P., 1959. Foundations of Linear Models. Stanford University Press, Stanford, CA, USA.
Fox, C., Erner, C., Walters, D., 2015. Decision under risk: From the field to the laboratory and back. In: The Wiley Blackwell Handbook of Judgment and Decision Making. pp. 43–88.
Gigerenzer, G., Gaissmaier, W., 2011. Heuristic decision making. Annu. Rev. Psychol. 62, 451–482.
Gilovich, T., Vallone, R., Tversky, A., 1985. The hot hand individual basketball: On the misperception of random sequences. Cogn. Psychol. 17 (3), 295–314.
Glöckner, A., Pachur, T., 2012. Cognitive models of risky choice: Parameter stability and predictive accuracy of prospect theory. Cognition 123 (1), 21–32.
Hastie, T., Tibshirani, R., 1990. Generalized Additive Models, first ed. Chapman and Hall/CRC, London, England.
Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction, second ed. Springer Verlag, New York, NY, USA.
Hogan, V., Walker, I., 2007. Education choice under uncertainty: Implications for public policy. Labour Econ. 14 (6), 894–912.
Howell, S., Jägle, A., 1997. Laboratory evidence on how managers intuitively value real growth options. J. Bus. Financ. Account. 24 (7 & 8), 915–935.
Hsu, Y.-W., 2008. Staging of venture capital investment: A real options analysis. Small Bus. Econ. 35 (1), 265–281.
Jacobs, B., 2007. Real options and human capital investment. Labour Econ. 14 (6), 913–925.
Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47 (2), 263–291.
Klingebiel, R., Adner, R., 2015. Real options logic revisited: The performance effects of alternative resources allocation regimes. Acad. Manag. 58 (1), 221–241.
McGrath, R.G., 1997. A real options logic for initiating technology positioning investments. Acad. Manag. Rev. 22 (4), 974–996.
McGrath, R.G., 1999. Falling forward: Real options reasoning and entrepreneurial failure. Acad. Manag. Rev. 24 (1), 13–30.
Miller, K., Shapira, Z., 2004. An empirical test of heuristics and biases affecting real option valuation. Strateg. Manage. J. 25 (3), 269–284.
Murphy, R.O., ten Brincke, R.H., 2016. Hierarchical maximum likelihood parameter estimation for cumulative prospect theory: Improving the reliability of individual risk parameter estimates. Manag. Sci. (in press).
Myers, S., 1977. Determinants of corporate borrowing. J. Financ. Econ. 5 (2), 147–175.
Myung, I.J., 2000. The importance of complexity in model selection. J. Math. Psych. 44 (1), 190–204.
Oprea, R., Friedman, D., Anderson, S., 2009. Learning to wait: A laboratory investigation. Rev. Econom. Stud. 76 (1), 1103–1124.
Prelec, D., 1998. The probability weighting function. Econometrica 66 (3), 497–527.
Quigg, L., 1993. Empirical testing of real option-pricing models. J. Finance 48 (2), 621–640.
Read, D., Loewenstein, G., 1995. Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. J. Exp. Psychol.: Appl. 1 (1), 34.
Rogers, M.J., Gupta, A., Maranas, C.D., 2002. Real options based analysis of optimal pharmaceutical research and development portfolios. Ind. Eng. Chem. Res. 41 (25), 6607–6620.
Schwartz, E.S., 2004. Patents and R&D as real options. Econ. Notes 33 (1), 23–54.
Sereno, L., 2010. Real options valuation of pharmaceutical patents: A case study. Available at SSRN 1547185.
Sirmon, D., Hitt, M., Ireland, D., 2007. Managing firm resources in dynamic environments to create value: Looking inside the blakc box. Acad. Manag. Rev. 32 (1), 273–292.
Thaler, R.H., Johnson, E.J., 1990. Gambling with the house money and trying to break veen: The effects of prior outcomes on risky choice. Manage. Sci. 36, 643–660.
Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. J. Risk Uncertain. 5 (4), 297–323.
von Neumann, J., Morgenstern, O., 1944. Theory of Games and Economic Behavior. Princeton University Press, Princeton, NJ, USA.
Vulkan, N., 2000. An economist's perspective on probability matching. J. Econ. Surv. 14 (1), 101–118.
Wakker, P., 2010. Prospect Theory: For Risk and Ambiguity. Cambridge University Press.
Wang, A., Bernstein, A., Chesney, M., 2009. An experimental study on real option strategies. Swiss Finance nstitute Research Paper Series, 9(50). Available at SSRN 1534506.
Yavas, A., Sirmans, C., 2005. Real options: Experimental evidence. J. Real Estate Financ. Econ. 31 (1), 27–52.
Zweig, J., 2008. Your Money and Your Brain: How the New Science of Neuroeconomics Can Help Make You Rich. Simon and Schuster, New York, NY, USA.