

An Agent for the Prospect Presentation Problem

Amos Azaria
Dept. of Computer Science
Bar Ilan University, Israel

Ariella Richardson
Dept. of Industrial Engineering
Jerusalem College of
Technology, Israel

Sarit Kraus
Dept. of Computer Science
Bar Ilan University, Israel

ABSTRACT

Evaluating complex propositions that are composed of several lotteries is a difficult task for humans. Presentation styles can affect the acceptance rate of such proposals. We introduce an agent that chooses between two presentation methods, while aspiring to maximize proposal acceptance. Our agent uses decision theory in order to model human behavior and uses the model to select the presentation which maximizes its expected outcome. We examine several decision theories, and use machine learning to adapt them to our domain. We perform an extensive evaluation of our agent in comparison to other baseline agents and show that presentation can indeed affect the acceptance rate of propositions and that the agent we propose succeeds in selecting beneficial presentations.

Categories and Subject Descriptors

I.2.m [Computing Methodologies]: ARTIFICIAL INTELLIGENCE—*Miscellaneous*

Keywords

Automated Agents, Human Persuasion, Prospect Theory

1. INTRODUCTION

Many systems are designed to encourage users to accept beneficial proposals. Complex propositions often confuse the user and make the decision non-trivial. It is well known that problem presentation [23, 11, 14] may have an impact on the human decision-making process. In our work we consider beneficial proposals which are composed of several gains or losses, that are associated with varying probabilities and must be accepted or rejected together. We will compare two possible presentation methods (for the possible outcomes and their associated probabilities) for each proposal: a separate presentation and a combined presentation. Many real life situations resemble our problem.

Our first example is a medical system which assists a doctor or patient in deciding whether to use a certain medication. The medication is associated with one or more benefits, such as curing the infection or reducing pain, and also with

several side effects, such as headache, nausea, rash or an allergic reaction. These outcomes have varied significance; for instance, a headache might be slightly unpleasant whereas an allergic reaction could be life threatening. Each of these outcomes is also associated with a certain probability; for example, the probability of overcoming the infection may be 90% while the side effect of a headache might occur in 20% of the patients, whereas an allergic reaction might only become evident in 0.5% of the patients. The expected overall reaction to the medication must be positive (otherwise it is more harmful than helpful). In order to decide whether to use a medication, all of the potential benefits and the risk of side effects must be evaluated together. Combining the various components is associated with a cost, since it is unclear how to quantify a headache compared to a rash. Different people may associate different values with each benefit or side effect. Therefore assigning values to the components using a joined metric would require some effort, such as questioning many people on their preferences, and therefore impose a cost.

Another example is an investment adviser who is trying to build an investment portfolio for one of his customers. Some stocks have a higher risk but also offer an opportunity to receive greater interest, while on the other hand one can invest in a bond with a lower risk and a lower interest level. Most people combine different stocks and bonds, combining different levels of risk. The investment adviser's primary goal is to get the customer to invest her money, therefore he would like to show the portfolio to the customer using the most appealing presentation. Should the investment adviser show the expected probability and value of revenue (or loss) for each stock, or should he try to combine all stocks in a single chart which presents the total investment?

Our last example is a travel agent who would like to promote the sales for a specific vacation package. Every day of a multiple-day vacation has some probability of rain or heat load (the strength of the rain or heat load may also vary). The travel agent wants to show the customer the probabilities for rain on each of the planned days. How should the travel agent present these probabilities (while his goal is to sell the package)? Should he present them for every day separately, or should he combine them all in one chart?

In order to determine how to present complex proposals, we propose an automated agent that utilizes behavioral economic theory. A *prospect* is a lottery (possibly with several outcomes, where each outcome has its own probability) [19], and a *simple prospect* is a prospect with some probability p to gain or lose some amount x and gain or lose

Appears in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, Lomuscio, Scerri, Bazzan, Huhns (eds.), May, 5–9, 2014, Paris, France.

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

nothing otherwise. The problems we study in this paper are composed of several simple prospects. These prospects must either **all** be accepted or **all** be rejected (there is no option to accept a partial set of prospects) and the system gains from accepted proposals. In the medical system we described earlier, the benefit of *cure infection* with a probability of 90% (and 10% of not curing the infection) is an example of a prospect, and similarly so is the side effect of acquiring a *headache* with a probability of 20% (and 80% of not resulting in a headache). All of the prospects must be selected or rejected together since a patient either takes the medication or does not. The agent we propose must decide whether to present the proposal in a separate method, as is, or in a combined method, combining all of the simple prospects into a single (more complex) prospect. For example, in the medical system the separate method would list all of the separate prospects. In combined presentation combining the curing prospect with the headache prospect results with a 72% chance that infection will be cured without the headache side effect appearing, an 18% chance that the infections will be cured and a headache will appear, an 8% that the infection will not be cured and no headache will appear and finally a 2% chance that the infection will not be cured and that a headache will appear. Note that in the combined method all of the probabilities add up to 100%. We assumed, and show experimentally, that presenting the problem as a set of simple prospects or as a combined prospect is not necessarily equivalent and can affect people’s choices. Thus the automated agent will determine when to use a separate presentation and when to use a combined presentation in order to encourage the users to accept the propositions.

When several prospects are proposed, the issue of bracketing arises. Read et al. [21] introduced the term “Choice Bracketing” to mean the grouping of choices. It has been shown that when people face several choices in which each choice has several options, people tend to treat these choices separately rather than treating them as a single decision. Our agent must take this into account when considering whether to use a combined or separate presentation for a problem, since similarly to what has been shown on separated choices, people might also treat each prospect separately even if these prospects are part of a group.

Behavioral economic theory describes the decision processes that people use when deciding whether to accept a prospect or reject it. The most significant theories in this field are the Expected Utility Hypothesis [13], the Prospect Theory [19] and the Cumulative Prospect Theory [26]. We embed these theories into our agent in order to model the expected human choice that will be made for a given set of prospects, in order to determine if the separate or combined presentation should be used.

We introduce the Prospect Presentation Problem, along with its formal description. This problem requires selecting whether to represent the multiple prospects in a separate presentation or a combined presentation, while maximizing a system’s utility. We propose a method for modeling and handling grouped separate prospects. We use different decision process models and settings in order to compose an agent that is capable of solving the Prospect Presentation Problem. We demonstrate the efficiency of the agent, in choosing the better presentation method, using an extensive experimental evaluation.

2. BACKGROUND

In this section, we first provide a very brief scan in the field of automated human persuasion and then, since our proposed agent integrates behavioral economic theory, we describe theories upon which we rely. Various methods have been proposed based on technology for persuading humans into taking certain actions. Most common methods do not involve agents at all [12]. Some of the prior work involving agents has attempted to influence people’s decision making by providing them advice [5, 4], by offering monetary incentives [6], by removing some of the options available to them [15, 24] or by revealing partial information to them by showing probabilities on the state of the world [3].

2.1 Expected Utility Hypothesis

The Expected Utility Hypothesis (EUH) was initiated by Bernoulli in 1738 [13]. Under this hypothesis, people have a utility function, u , which associates any possible total wealth with some utility. People use this function when deciding whether to accept or reject a lottery simply by maximizing their expected utility. For example, a person with a current total wealth of $\$W$ facing a prospect (lottery), P , with a probability of p to win $\$x$ and a probability of $1 - p$ to lose $\$y$, will compare the expected utility from accepting the offer:

$$U(P) = u(W + x) \cdot p + u(W - y) \cdot (1 - p) \quad (1)$$

with the expected utility from rejecting the offer, which is simply $u(W)$. The person will accept the lottery if the former is greater and reject it otherwise. A common utility risk averse function, suggested by Bernoulli himself, is the log function:

$$u(X) = \log(X) \quad (2)$$

2.2 Prospect Theory

The Prospect Theory was presented by Kahneman and Tversky in [19] and later refined to the Cumulative Prospect Theory (CPT) in [26]. The Prospect Theory is based on three principles. The first is that people do not take into account their total wealth when accepting or rejecting an uncertain opportunity (as suggested by the expected utility hypothesis [13]), but rather use their current wealth as a baseline, and will be happy if they win an amount and become upset if they lose an amount. The second principle is loss aversion, where people hate losing more than they like winning. The third principle is that people have a subjective representation of probabilities and do not interpret probabilities fully rationally, but rather use their own decision weights when deciding whether to reject or accept a gamble. In his book, Kahneman [18] (p.314) gives the following examples: The decision weight that corresponds to a 90% chance is 71.2%, while the decision weight that corresponds to a 10% chance is 18.6%. According to these examples, people are likely to prefer a guaranteed outcome of \$80 than a gamble with a 90% chance of winning \$100, since the latter is only worth \$71.2 to them. Tversky and Kahneman elicited these weights by sequentially asking subjects to choose between a specific lottery and many different guaranteed outcomes. The equivalent to the given lottery for a certain subject was set to the average between the greatest rejected guaranteed outcome and the smallest accepted guaranteed outcome [26]. However, these decision weights

depend on people’s personalities, their wealth, culture and the scope of the payoff in question. The cumulative prospect theory determines the value for any prospect based on its possible outcomes and the probability of each of its outcomes. Given a prospect P which is composed of T possible ordered outcomes (as defined by Tversky and Kahneman), $\{x_1, x_2, \dots, x_T\}$, and the first t are negative outcomes, i.e. $x_1 < x_2 < \dots < x_t < 0 \leq x_{t+1} < \dots < x_T$. Each outcome is associated with some probability $p(x)$. The value of the prospect is given by the following formula:

$$U(P) = \sum_{i=1}^t v(x_i) \cdot \left(w\left(p_i + \sum_{j=1}^{i-1} p(j)\right) - w\left(\sum_{j=1}^{i-1} p(j)\right) \right) + \sum_{i=t+1}^T v(x_i) \cdot \left(w\left(p_i + \sum_{j=i+1}^T p(j)\right) - w\left(\sum_{j=i+1}^T p(j)\right) \right) \quad (3)$$

where $v(x)$ stands for the value function and w is the weighting function (the decision weight function described above). Both of these functions must be non-decreasing, and $w(0) = 0, w(1) = 1$. v is negative for losses and positive for gains, and $v(0) = 0$. The intuition behind this formula is that every possible output’s value is assumed to have an impact which is proportionate to the marginal affect that its accumulated probability has on the weighting function. For example, given a prospect P' with three possible outcomes, \$2, \$3 and \$10 (no negative outcomes) with probabilities of $p(\$2) = 0.1, p(\$3) = 0.7, p(\$10) = 0.2$, the value of the prospect is given by:

$$U(P') = v(\$10) \cdot w(0.2) + v(\$3) \cdot (w(0.9) - w(0.2)) + v(\$2) \cdot (w(1) - w(0.9))$$

Note that due to the nature of w and v , the value of P' is at least $v(\$2)$. This complies with the fact that the prospect guarantees a win of at least \$2.

Tversky and Kahneman suggested the value function:

$$v(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\mu(-x)^\beta & \text{if } x < 0 \end{cases} \quad (4)$$

where α, β and μ are parameters, and the weighting function is:

$$w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1-\gamma}} \quad (5)$$

where γ is a parameter used for positive payoffs and is replaced by a different parameter, δ , for negative payoffs. Several studies try to estimate parameters for the Prospect Theory [17, 22, 16], however most studies try to maximize the likelihood of the results obtained by each subject individually. This approach could not be applied in our work since we build a model based on a group of users and apply this model to new users (for whom we have little or no data). Building a model for each and every user will not allow us to generalize the model to new subjects.

2.3 Logit Quantal Response

Stochastic decision-making (logit quantal response) suggests that when humans need to make a choice among several options, the option yielding the greater utility to the user is more likely to be chosen than the option yielding a lesser utility. However, an option yielding a lesser utility may occasionally be chosen. Logit quantal response is

very common in the literature which aims to model human behavior, [7, 10, 4, 20], and suggests that when a person confronts several possible choices $\{C_1, C_2, \dots, C_\ell\}$ the probability that the person will chose a specific choice C_i is given by:

$$p(C_i) = \frac{e^{\lambda U(C_i)}}{\sum_{j=1}^{\ell} e^{\lambda U(C_j)}} \quad (6)$$

where $U(C)$ is the utility associated with choice C and λ is a parameter.

2.4 Bracketing

“Choice Bracketing”, termed by Read et al. [21], designates the grouping of individual choices together into sets. “Broadly Bracketing” indicates that the decision-maker takes all choices into account when making his decision, while “Narrow Bracketing” indicates that the decision-maker isolates each choice from all other choices. When humans face a broad spectrum of topics, where each topic consists of several options, they usually make a decision on each topic separately. A classic experiment that illustrates narrow bracketing was done by Tversky and Kahneman [25]. They asked their subjects the following question:

”Imagine that you face the following pair of concurrent decisions. First examine both decisions, then indicate the options you prefer:

Choice I. Choose between:

- A. A guaranteed gain of \$240.
- B. A 25% chance to gain \$1000 and a 75% chance to gain nothing.

Choice II. Choose between:

- C. A guaranteed loss of \$750.
- D. A 75% chance to lose \$1000 and a 25% chance to lose nothing.”

Since people tend to be risk averse with a positive payoff and risk seeking with a negative payoff, a large majority of subjects (73%) chose both A and D. Only 3% of the subjects chose B and C. Combining A and D yields a 25% chance to gain \$240 and a 75% chance to lose \$760. However, combining B and C dominates this with a 25% chance to gain \$250 and a 75% chance to lose \$750. Tversky and Kahneman performed an additional experiment in which they presented only the following combined choices to the subjects, that is:

Choose between:

- A+C. A guaranteed loss of \$510.
- A+D. A 25% chance to gain \$240 and a 75% chance to lose \$760.
- B+C. A 25% chance to gain \$250 and a 75% chance to lose \$750.
- B+D. A 6.25% chance to gain \$1000, a 56.25% chance to lose \$1000 and a 37.5% chance to gain or lose nothing.”

This time, not a single subject chose the dominated option (A+D). This experiment demonstrates that people tend to treat each decision on its own and do not combine the choices, unless they are explicitly combined for them.

3. PROSPECT PRESENTATION PROBLEM

The *prospect presentation problem* is a decision problem for a system and is defined as follows. We first define a set of simple prospects $s = \{P_1, P_2, \dots, P_{k_i}\}$; recall that a simple prospect P is composed of a probability P_p of gaining or losing a certain amount, P_x . In the prospect presentation problem, a system has a set of n sets of simple prospects, $S = \{s_1, s_2, \dots, s_n\}$. Each set $s \in S$ must be offered to h human clients. Each of the human clients may either accept the set of prospects s (and participate in the lotteries associated with the prospects) or reject it. Each of the sets, s , may be presented to the human clients in two different presentation modes $m(s)$; the presentation mode may either be *separate*, which indicates that the set of prospects are presented separately (as they are), or the presentation mode may be *combined*, which indicates that the prospects are *validly* combined into a single prospect. The separate presentation mode of s is denoted by s^s and the combined presentation mode of s is denoted by s^c . The probability for any possible outcome must be identical in both s^s and s^c . A cost $c(m(s))$ may be applied to the system and may depend on the method of presentation. The system gains a utility of $1 - c(m(s))$ every time a human client accepts a set of prospects. The human clients are assumed to follow a stochastic decision policy, in which, given a set of prospects, s , and a presentation mode, m , $p(s, m)$ determines the probability that the humans will accept the set of prospects. The *prospect presentation problem* is intended for the system to determine for each of the sets of prospects, s , its presentation mode, $m(s)$, in order to maximize:

$$\sum_{s \in S} p(s, m(s)) \cdot h \cdot (1 - c(m(s))) \quad (7)$$

4. AGENT FOR PROSPECT PRESENTATION PROBLEM

In this section we introduce an **Agent for the Prospect Presentation Problem (APPP)**. Section 4.1 describes how APPP calculates the combined presentation for a set of prospects and solves the prospect presentation problem. However, this solution relies on a component that accurately models human decision policy. We therefore propose several alternatives for APPP's composition of the human model in Section 4.2.

4.1 Solving the Prospect Presentation Problem

The first component of the APPP agent is described in Algorithm 1 and handles the task of efficiently (linear in output length) calculating a combined presentation for a set of prospects s^c . The algorithm receives a set of simple prospects and outputs a hash map with all of the possible outcomes (as keys) and probabilities (as values). The algorithm iterates via all prospects. In every iteration the algorithm takes the previous iteration's result and doubles it, once assuming that the current prospect obtained its outcome and once assuming that the current prospect did not yield its outcome. For example, consider a set of simple prospects in which one prospect has a 25% chance to win \$37 and another prospect has a 60% chance to lose \$10, i.e. $s = \{(0.25, \$37), (0.60, \$-10)\}$. In the first iteration, with the prospect (0.25, \$37), there will be only two possible

outcomes, 0 with a probability of 0.75 and \$37 with a probability of 0.25. In the second (and last) iteration, with the prospect (0.60, \$-10), the algorithm will first assume that the (negative) outcome was not obtained (with a probability of 0.40), and thus will have two possible outcomes: \$37 with a probability of $0.25 \cdot 0.40 = 0.10$ and \$0 with a probability of $0.75 \cdot 0.40 = 0.30$. Then the algorithm will add two additional outcomes, assuming that the outcome of the second prospect (\$-10) was obtained (with a probability of 0.60): $\$27 - \$10 = \$17$ with a probability of $0.25 \cdot 0.60 = 0.15$ and $\$0 + \$-10 = -\$10$ with a probability of $0.75 \cdot 0.60 = 0.45$.

Algorithm 1 Calculation of the combined presentation for a set of prospects.

Input: s - A set of simple prospects $s = \{P_1, P_2, \dots, P_k\}$, where $P_i = (P_i^p, P_i^x)$.
Output: s^c - A hash map holding all possible outcomes as keys and their associated probabilities as values.

- 1: $s^c[0] \leftarrow 1$
- 2: **for** each prospect P in s **do**
- 3: $s^{c'} \leftarrow s^c$
- 4: clear s^c
- 5: **for** each *outcome* in $s^{c'}$ **do**
- 6: $s^c[\textit{outcome}] \leftarrow s^{c'}[\textit{outcome}] \cdot (1 - P^p)^1$
- 7: $s^c[\textit{outcome} + P^x] \leftarrow s^{c'}[\textit{outcome}] \cdot P^p$
- 8: **end for**
- 9: **end for**
- 10: return s^c

The second component of the APPP agent is described in Algorithm 2. This is the procedure that solves the prospect presentation problem. The input for Algorithm 2 is a set of prospect sets, a cost function and a human decision policy. The output is a determination policy for each prospect set. The determination policy determines whether to use the *separate* or the *combined* method for each prospect set. This algorithm simply iterates via all sets of prospects and calculates, for each of the sets, which of the presentation methods is more profitable for the system, i.e. whether $p(s, \textit{separate}) \cdot (1 - c(\textit{separate}))$ is larger than $p(s, \textit{combined}) \cdot (1 - c(\textit{combined}))$ or vice versa. Were the human decision policy to be known,

Algorithm 2 Procedure for solving the prospect presentation problem.

Input: A set composed of prospect sets S . A cost function $c(m(s))$, and a human decision policy $p(s, m(s))$.
Output: A determination policy, m , for each of the prospect sets.

- 1: **for** each problem s in S **do**
- 2: **if** $p(s, \textit{separate}) \cdot (1 - c(\textit{separate})) > p(s, \textit{combined}) \cdot (1 - c(\textit{combined}))$ **then**
- 3: set $m(s) \leftarrow \textit{separate}$;
- 4: **else**
- 5: set $m(s) \leftarrow \textit{combined}$;
- 6: **end if**
- 7: **end for**

the algorithm would have fully solved the prospect presentation problem. However, in real life, a system agent facing the prospect presentation problem is not likely to have access to the human decision policy, $p(s, m(s))$. Therefore, the

major concern for an agent facing the prospect presentation problem is to accurately model the human decision policy.

4.2 Decision Policy Modeling in APPP

An agent facing the prospect presentation problem does not have specific information about the human user, and therefore must use a general model for modeling human decision policy. While several theories may be considered for building this model, the Expected Utility Hypothesis and Cumulative Prospect Theory definitely stand out as being significant in their field. We therefore decided to embed each of these theories in APPP’s decision policy model, but it is possible to embed other theories into the agent should one want to. However, in both the expected utility hypothesis and the cumulative prospect theory, a human’s response does not depend on the form of the presentation of the problem. This implies that both presentation methods (*combined* and *separated*) would yield the same probability that the user will accept a given lottery regardless of the lottery, i.e. for any s , $p(s, \textit{combined}) = p(s, \textit{separate})$. This assumption is clearly inappropriate for our work (and is shown to be false in the results section). Therefore, each of the theories requires a slight modification when considering the *separate* representation method, by taking the bracketing effect (see section 2.4) into account.

We use the following subsections to describe in detail how each of the theories can be embedded into APPP. All of the methods we test need to set some parameters, therefore APPP requires training data. The data set, ψ , is composed of a set of tuples $\langle s, m(s), d \rangle$, in which s is a set of simple prospects presented to a human user, $m(s)$ determines whether the set of prospects were presented in the *combined* method or in the *separated* method, and d is a boolean, indicating whether or not the user decided to participate in the lottery. In order to accurately model human decision-making, it is essential to assume stochastic decision-making, since the agent is required to evaluate the probability that a user will accept a lottery. Not assuming stochastic decision-making would mean that a lottery (possibly depending on its presentation method) would either be accepted by everyone or rejected by everyone, i.e. $p(s, m(s)) \in \{0, 1\}$. (It is not required to assume that every individual actually uses stochastic decision-making, but that the crowd as a whole can be modeled as using stochastic decision-making.)

APPP assumes logit quantal response and thus relies on Equation 6. Recall that in the prospect presentation problem, the user needs to choose between participating in a prospect (or a set of simple prospects in the *separate* presentation method) or not. Thus, the user must actually choose between the lottery and the value of not participating in it, denoted by $U(\textit{null})$. In EUH, $U(\textit{null}) = u(\$W)$, where W is the person’s initial wealth, and in CPT, $U(\textit{null}) = v(0) = 0$. Therefore, given a lottery L , according to Equation 6, the probability that a user will accept that lottery is:

$$p(L) = \frac{1}{1 + e^{\lambda(U(\textit{null}) - U(L))}} \quad (8)$$

4.2.1 Modeling using Prospect Theory with Learned Parameters

¹In 6 and 7, if $s^c[\textit{outcome}]$ or $s^c[\textit{outcome} + P^x]$ already have a value, increment that value by $s^{c'}[\textit{outcome}] \cdot (1 - P^p)$ or $s^{c'}[\textit{outcome}] \cdot P^p$ respectively.

The first model we consider for APPP’s decision policy is the cumulative prospect theory. We use cumulative prospect theory, CPT, (see section 2.2) in order to evaluate $U(L)$ for a user facing the *combined* method (a single prospect). However, as mentioned above, when considering the *separated* method (a set of simple prospects), APPP deviates from CPT. For the instances in the data set using the *combined* method of presentation, in which the user is only presented with a single prospect (s^c), APPP calculates $U(s^c)$ according to Equation 3 (and Equations 4 and 5 for the value and weighting functions respectively). APPP assumes that people who face the separated method evaluate each prospect separately and then combine all values together to receive the total value of the lottery². This assumption is based on the bracketing effect (see section 2.4), which suggests that people treat each problem separately, and thus we assume that they will evaluate each prospect separately. Formally, given a set of prospects, s , presented in the *separated* method, the value of the set of prospects is given by:

$$U(s^s) = \sum_{P \in s} U(P) \quad (9)$$

APPP searches for parameters $\alpha, \beta, \mu, \gamma, \delta$ and λ that minimize the mean squared error (MSE) between $p(s, m(s))$ and the actual fraction of users in (ψ) who accepted lottery s ($d = \textit{true}$) out of all of those who were shown that lottery using $m(s)$ as the presentation method. Once APPP has set the parameters to use for the decision policy, it can be used to determine how to present a new set of prospects. Given a set of prospects, s , and a presentation mode, $m(s)$; if the presentation mode is *combined*, APPP uses Equations 8 and 3 and parameters $\alpha, \beta, \mu, \gamma, \delta$ and λ . If the presentation mode is *separated*, APPP uses Equations 8, 9 and 3 and the parameters ($\alpha, \beta, \mu, \gamma, \delta$ and λ) to evaluate the probability that the user will accept the associated lottery.

4.2.2 Modeling using Prospect Theory with Original Parameters

We also investigated another possible model; it is the same model described in 4.2.1, but rather than learning the parameters $\alpha, \beta, \mu, \gamma$ and δ , APPP uses the parameters proposed by Kahneman and Tversky. That is, $\alpha = 0.88, \beta = 0.88, \mu = 2.25, \gamma = 0.61$ and $\delta = 0.69$. APPP only uses the data to evaluate λ .

4.2.3 Modeling using Expected Utility Hypothesis

As the Expected Utility Hypothesis is very well known, we also show how to embed it into APPP. APPP, based on EUH, uses Equations 1 and 2³. As in the model assuming CPT, we must determine how to account for sets of prospects presented in the *separated* method. However, using EUH, APPP may not simply use the exact same approach as when using CPT, since in EUH, $U(P)$ includes the initial wealth. Therefore, if we were simply to add up the utilities from all of the prospects (using Equation 9), we would end-up adding the initial wealth several times. We therefore propose

²This approach in an environment of only simple prospects is actually similar to the assumptions made by the prospect theory rather than CPT.

³In practice we also tried two different power functions, but they yield the exact same results as the log function (presented in Section 5.2).

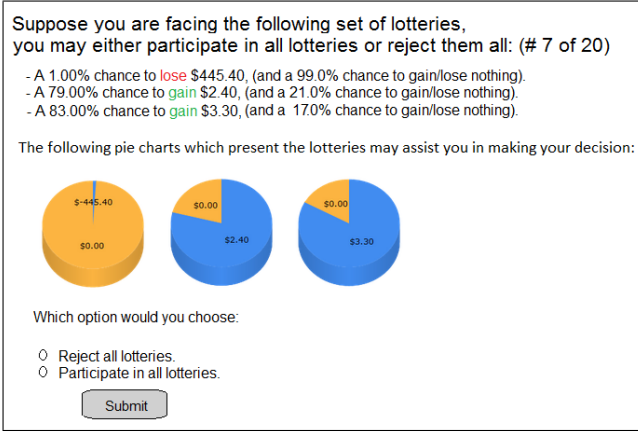


Figure 1: A subject facing a set of prospects presented separately.

a simple modification, using the following equation:

$$U(s^s) = u(W) + \sum_{P \in s} (U(P) - u(W)) \quad (10)$$

APPP based on EUH, searches for parameters W and λ that minimize the mean squared error (MSE) between $p(s, m(s))$ and the actual fraction of users in (ψ) who accepted lottery s ($d = true$) out of all those who were shown that lottery using $m(s)$ as the presentation method. Given a set of prospects, s , and a presentation mode, $m(s)$; if the presentation mode is *combined*, APPP uses Equations 8 and 1 and parameters W and λ . If the presentation mode is *separated*, APPP uses Equations 8, 10 and 1 and the parameters (W and λ) to evaluate the probability that the user will accept the associated lottery.

5. EVALUATION

5.1 Experimental Setup

We ran our experiments using Amazon’s Mechanical Turk (AMT) [1]. We constructed a total of 120 sets of simple prospects, with k (the number of simple prospects in a set) varying between 3 and 5. The upper boundary of 5 was chosen since the number of possible outcomes is exponential in the number of k ($2^5 = 32$) and we didn’t want to present too many possible outcomes to the subjects. Each simple prospect (P) had a random probability (P_p) between 1% and 100% (only whole probabilities) and a random expected outcome ($P_p \cdot P_x$) between $-\$15.00$ and $+\$15.00$. For the human decision not to be trivial, every set of prospects had at least one prospect with a negative outcome and one with a positive outcome. For ethical reasons, since we did not want to encourage traditional gambling, we ensured that all gambles had a positive expected utility. Therefore, a player that tries to maximize her expected outcome should have accepted all gambles. Thus, the subjects were not urged into a gamble which was not good for them. We recruited a total of 612 participants. 58.1% of the subjects were males and 41.9% were females. Subjects’ ages ranged from 18 to 73, with a mean of 31.6, a median of 29 and a standard deviation of 11.0. All subjects were residents of the USA. The subjects were paid 30 cents to participate in the experiment. Each subject was presented with 20 sets of prospects,

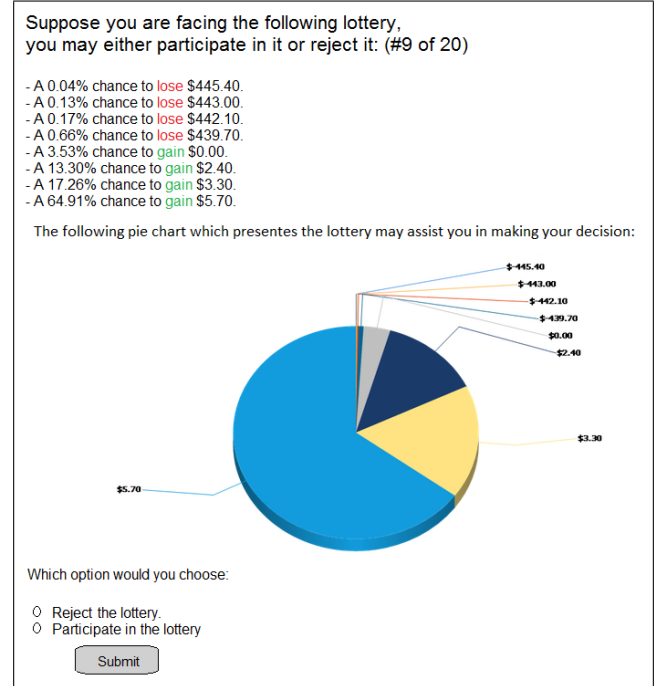


Figure 2: A subject facing a set of prospects presented in the combined mode.

half in their original form and half as combined prospects. This resulted in an average of approximately 50 instances for each of the 120 sets of prospects for each of the two different presentation modes. The subjects were given the following instructions: ”Suppose you are facing the following lottery / set of lotteries, you may either participate in it / in all lotteries or reject it / them all”. (The exact text depended on the mode of presentation of the current set of prospects.) To enhance comprehension, we also provided the subjects with pie-charts presenting the prospects (as used in [2]). The following explanation was provided: ”The following pie chart(s) which present(s) the lottery / lotteries may assist you in making your decision.” Figure 1 presents an example of a screen-shot for a subject facing a set of prospects presented separately, and Figure 2 presents an example of a screen-shot for a subject facing the same set of prospects presented in the combined mode. Note that there are 3 prospects in the separate mode, and thus $2^3 = 8$ in the combined mode. We set the cost function to 0 for the separated method and to 0.15 for the combined method. This setting was chosen since we assume that the problem is provided as the separate prospects and therefore some cost is associated with presenting a combined prospect to the user. A cost of 0.15 to the *combined* mode, equalizes the performance of the *combined* mode to the *separate* mode.

5.2 Results

We ran APPP using 10-fold cross-validation on the data. In every fold APPP was trained on 108 sets of prospects, and tested on the remaining 12 sets of prospects. For each of these sets, APPP had to decide whether to present the prospects separately or combined. Figure 3 presents the performance of APPP in three modes. The APPP refers to the agent using the learned parameters as described in Section

| Agent | Acceptance rate | Combined presentation | Average cost | Score | Number of parameters |
|-------------------|-----------------|-----------------------|--------------|--------------|----------------------|
| APPP | 42.8% | 37.5% | 0.0266 | 0.401 | 6 |
| APPP-KT | 39.6% | 16.7% | 0.009 | 0.387 | 1 |
| Combined | 44.4% | 100% | 0.066 | 0.378 | 0 |
| Separate/APPP-EUH | 37.8% | 0% | 0 | 0.378 | 0 – 2 |

Table 1: Average performance of APPP compared to the other agents

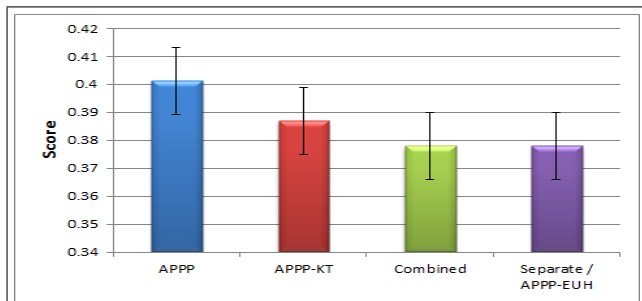


Figure 3: Average score obtained by each of the methods.

4.2.1, APPP-KT refers to the agent that uses the parameters described by Kahneman and Tversky as explained in Section 4.2.2 and APPP-EUH is the agent described in 4.2.3. These versions of the agent are compared to an agent that always presents the *combined* mode and an agent that always presents the *separate* mode. The black bars represent the 95% confidence interval using a paired t-test. APPP significantly ($t(119) = 2.3645$, $p < 0.01$ using a paired t-test on the score obtained from every set of prospects) outperformed all other methods, and yielded an increase of 6% in the average score over the two baselines. Recall that APPP may only control the method of presentation to the users (and not the actual lotteries), therefore this achievement is very impressive. Table 1 provides additional details regarding the acceptance rate of APPP and the baseline agents. As can be observed by the table, the combined mode enjoyed a much greater acceptance rate than the separate mode. This clearly demonstrates that the presentation mode has an impact on the human acceptance rate, justifying our initial assumptions. However, recall that presenting the combined method requires some additional effort and is therefore assumed to be associated with some cost. If this cost is reduced, the *combined* mode becomes much more appealing, and vice versa, as the cost increases, the *separated* mode becomes more appealing. It is not surprising that APPP yields a slightly lower acceptance rate than the combined agent, since, as mentioned, many more subjects accept the combined mode than the separated mode, and APPP chose to present the combined mode only in 37.5% of the sets. APPP-EUH did not present the combined mode in any of the sets. This was not because it assumed that using the separate mode is more appealing to the user, but because it was not willing to pay the cost associated with presenting the combined mode (if the cost was totally removed, APPP-EUH would present the combine mode in 94% of the sets). Recall that the expected outcome on all lotteries was always positive, yet on average people still accepted less than 50% of the lotteries.

Regarding the pie-charts, 79.1% of the subjects said that the pie-charts helped them and only 19.6% said that they didn't help them. Only 6 subjects (less than 1%) said that they did not understand the pie-charts. While the total average of participation in the lotteries was 41.5%, females participated on average in only 37.9%, which is significantly lower ($p < 0.01$ using ANOVA test) than males who participated on average in 44.1% of the lotteries. This indicates that females are more risk averse than males (this finding was also present in many studies such as [8]). We also found that young people, up to age 29 (which was our median), are significantly ($p < 0.01$) more likely to participate in the lottery than people aged 30 and above (44.0% vs. 37.7%). Risk aversion also seems to increase with education, as subjects with only a high school education (49.6%) participated in 43.3% of the lotteries, while subjects with a bachelor's or a master's degree or a PhD participated in only 39.7% of the lotteries (these results differ statistically with $p < 0.05$, however, these results may be explained by the fact that the younger subjects tend to have lower education). Interestingly, while the gap between the males and females participating in the lotteries was smaller when the prospects were combined (6.93% vs. 5.62%), the gap between the young subjects and the older ones was almost doubled in size when the presentation mode was *combined* (4.29% vs. 8.31%). This finding may encourage future work on a personalized agent based on demographic data alone.

5.3 Discussion and Future Work

In this study we used $k \leq 5$ (the number of simple prospects in a set). It remains questionable, what will happen as the number of prospect selection problems increases. In such a case, the pie chart may become very complex as the number of outcome options increases exponentially. One option is to group together similar outcomes, but still it is unclear how to do so. Another interesting question is what happens for lower values of k . Lower values of k should enable a better evaluation of the problem and a better comprehension of the agent's advice; it remains to be seen how this affects the users' interaction with the agent. We chose the pie charts (combined and separate) in order to enhance comprehension of the proposed prospects. The users provided positive feedback to these charts. It would be interesting to study how much of a role they played in the user acceptance rates. Enhancing APPP to include more presentation modes or other visual enhancements would be an interesting extension of this study. Using CPT with learned parameters was shown to outperform other methods we tested. In the future we would like to consider modifying the functions suggested by Tversky and Kahneman to see if it is possible to further improve the APPP's performance. Another possible extension would be to study how to apply alternative modeling theories such as the priority heuristic [9].

6. CONCLUSION

In this paper we introduced the *prospect presentation problem*, in which users are presented with sets of prospects that must be accepted or rejected as a group (such as an investment portfolio). We refer to a system that gains a positive utility when clients accept the prospects (decide to invest). We defined the *prospect presentation problem* as determining, for each set of prospects, which presentation mode to use in order to maximize the system's utility. We presented the APPP agent that solves the *prospect presentation problem* and chooses between representations of complex beneficial proposals. APPP uses Cumulative Prospect Theory (CPT) in order to model human decision-making and uses this model to select the better representation method for a set of prospects. We demonstrated that fixing the presentation method to always present the prospects separately, or always present them combined, results in a lower score than using APPP to select the presentation mode based on the human model and a given set of prospects. We investigated several variations on decision process models and found that using CPT results in the most successful agent for selecting the presentation mode. Furthermore, learning the parameters for CPT from sets of prospects presented in the two presentation modes resulted in a better agent than simply using the parameters elicited by Kahneman and Tversky in their original study. We show that the combined method yields much higher acceptance rates, and thus, if the cost associated with presenting the combined method is very low, it may be more beneficial for the system to always use the combined mode. If the cost is not negligible, we suggest using APPP for determining which sets to present separately and which to present combined.

7. REFERENCES

- [1] Amazon. Mechanical Turk services. <http://www.mturk.com/>, 2014.
- [2] A. Azaria, S. Kraus, and A. Richardson. A system for advice provision in multiple prospect selection problems. In *RecSys*, 2013.
- [3] A. Azaria, Z. Rabinovich, S. Kraus, and C. V. Goldman. Strategic information disclosure to people with multiple alternatives. In *AAAI*, 2011.
- [4] A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and Y. Gal. Strategic advice provision in repeated human-agent interactions. In *AAAI*, 2012.
- [5] A. Azaria, Z. Rabinovich, S. Kraus, C. V. Goldman, and O. Tsimhoni. Giving advice to people in path selection problems. In *AAMAS*, 2012.
- [6] Amos Azaria, Yonatan Aumann, and Sarit Kraus. Automated agents for reward determination for human work in crowdsourcing applications. *Autonomous Agents and Multi-Agent Systems*, pages 1–22, 2013.
- [7] Michael H Birnbaum and Alfredo Chavez. Tests of theories of decision making: Violations of branch independence and distribution independence. *Organizational Behavior and Human Decision Processes*, 71(2):161–194, 1997.
- [8] Lex Borghans, Bart H.H. Golsteyn, James J. Heckman, and Huub Meijers. Gender differences in risk aversion and ambiguity aversion. Working Paper 14713, National Bureau of Economic Research, 2009.
- [9] Eduard Brandstätter, Gerd Gigerenzer, and Ralph Hertwig. The priority heuristic: making choices without trade-offs. *Psychological review*, 113(2):409, 2006.
- [10] Nathaniel D Daw, John P O'Doherty, Peter Dayan, Ben Seymour, and Raymond J Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876–879, 2006.
- [11] Maier Fenster, Inon Zuckerman, and Sarit Kraus. Guiding user choice during discussion by silence, examples and justifications. In *ECAI*, 2012.
- [12] Brian J Fogg. Persuasive technology: using computers to change what we think and do. *Ubiquity*, 2002(December):5, 2002.
- [13] M. Friedman and L. J. Savage. The expected-utility hypothesis and the measurability of utility. *The Journal of Political Economy*, 60(6):463–490, 1952.
- [14] Y. Gal, B. Grosz, S. Kraus, A. Pfeffer, and S. Shieber. Agent decision-making in open mixed networks. *Artificial Intelligence*, 174(18):1460–1480, 2010.
- [15] Chen Hajaj, Noam Hazon, David Sarne, and Avshalom Elmalech. Proceedings of the twenty-seventh aaai conference on artificial intelligence. In *AAAI*, pages 401–408, 2013.
- [16] David W Harless and Colin F Camerer. The predictive utility of generalized expected utility theories. *Econometrica*, 62:1251–1289, 1994.
- [17] Glenn Harrison and E. Rutstrom. Expected utility theory and prospect theory: one wedding and a decent funeral. *Experimental Economics*, 12:133–158.
- [18] Daniel Kahneman. *Thinking, fast and slow*. Allen Lane, London, 2011.
- [19] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- [20] T. Nguyen, R. Yang, A. Azaria, S. Kraus, and M. Tambe. Analyzing the effectiveness of adversary modeling in security games. In *AAAI'13*, 2013.
- [21] Daniel Read, George Loewenstein, and Matthew Rabin. Choice bracketing. *Journal of Risk and Uncertainty*, 19(1):171–197, December 1999.
- [22] J. Rieskamp. The probabilistic nature of preferential choice. *Journal of Experimental Psychology*, 34:1446–1465, 2008.
- [23] Shawn W Rosenberg, Lisa Bohan, Patrick McCafferty, and Kevin Harris. The image and the vote: The effect of candidate presentation on voter preference. *American Journal of Political Science*, pages 108–127, 1986.
- [24] David Sarne, Avshalom Elmalech, Barbara J. Grosz, and Moti Geva. Less is more: restructuring decisions to improve agent search. In *Proc. of AAMAS*, pages 431–438, 2011.
- [25] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, January 1981.
- [26] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, October 1992.