

When Suboptimal Rules

Avshalom Elmalech

Bar-Ilan University,
Israel
elmalea@cs.biu.ac.il

David Sarne

Bar-Ilan University,
Israel
sarned@cs.biu.ac.il

Avi Rosenfeld

Jerusalem College of Technology,
Israel
rosenfa@jct.ac.il

Eden Shalom Erez

Independent Researcher,
Israel
edenerez@gmail.com

Abstract

This paper represents a paradigm shift in what advice agents should provide people. Contrary to what was previously thought, we empirically show that agents that dispense optimal advice will not necessarily facilitate the best improvement in people's strategies. Instead, we claim that agents should at times suboptimally advise. We provide results demonstrating the effectiveness of a suboptimal advising approach in extensive experiments in two canonical mixed agent-human advice-giving domains. Our proposed guideline for suboptimal advising is to rely on the level of intuitiveness of the optimal advice as a measure for how much the suboptimal advice presented to the user should drift from the optimal value.

Introduction

Previous research on advice provisioning focused on providing the optimal advice as the assumption was this advice would definitely be accepted, thus yielding the best performance (Power and Sharda 2009; Bharati and Chaudhury 2004). However, while designing agents to dispense fully rational advice is inherently logical, previous work has shown that people often do not adhere to rigid models of rationality. Research in psychology behavioral economics, and artificial intelligence has established that people are boundedly rational (Kahneman 2000; Rosenfeld and Kraus 2009; Hajaj et al. 2013) and showed that their decision-making may be influenced by other factors and that people are at times biased towards certain conclusions (Blackhart and Kline 2005; Drake 1993; Hajaj, Hazon, and Sarne 2014). Similarly, as our work shows, optimal advice given by agents is not always accepted by the people they supposedly are supporting. Please note that the use of the term "optimal" is for describing the "best" advice the system can come up with. However, the person may not recognize this advice as such (e.g. she believes the agent is untruthful) and we thus must consider when agents should deviate from dispensing fully rational advice, and if so, by how much, such that the advice will be more likely to be accepted.

This work presents a new approach for advice provisioning. We posit that in cases where the solution is not intuitive, people will often not accept the optimal advice. In these

types of tasks, an advising agent should intentionally provide suboptimal advice instead. Thus, the agent should reason about the tradeoff between the probability its advice will be accepted and the value the person's acceptance of this advice will bring the system. In cases where the *expected utility* of suboptimal advice is higher than that of the optimal one, it should be presented instead.

Focusing in giving suboptimal advice contrasts with previous works that attempted to overcome people's inability to recognize what is optimal in a given decision setting. For example, some work attempted to convince people of the correctness of the optimal solution, e.g., through providing them a series of manipulations, though with not much success (Grosskopf, Bereby-Meyer, and Bazerman 2007; Elmalech, Sarne, and Grosz 2014). Others attempted to teach people how to optimally solve the decision problem (Lesser 1999; Elmalech, Sarne, and Agmon 2014). Our approach, on the other hand, does not require any overhead from the advisee's side. A third approach considered self interested advice agent which intentionally withheld advice (Buntain et al. 2012; Azaria et al. 2014; Rosenfeld and Kraus 2015; Azaria et al. 2013; 2012), or customized it for the people (Hazon, Lin, and Kraus 2013), to maximize the advice agent's reward. However, we consider a collaborative environment with a shared goal. This work investigates the efficiency of the solution presented on real people, in contrast to other work that measured their solution using agents to simulate (to some extent) people's behavior (Mash, Lin, and Sarne 2014; Chalamish, Sarne, and Lin 2012; Elmalech and Sarne 2014).

To support our claims, we provide extensive empirical results from two canonical domains: a "Company Valuation Game", (aka. the "Takeover Game") originally presented by Samuelson and Bazerman (1985), and a "Birthday Game", (also known as the "Birthday Paradox") previously developed by Ball (1914). Both domains have optimal non-intuitive solutions and previous studies have shown that participants consistently do not make the optimal decisions in them (Selten, Abbink, and Cox 2005; Voracek, Tran, and Formann 2008). The analysis of the results validates that indeed people's strategies in these two domains are far from optimal and demonstrates that people overall reject the optimal advice. Furthermore, we show that agents giving our proposed suboptimal advice facilitated significantly better

performance that those giving the optimal advice.

Sub-Optimal Advice Giving Approach

As previous work into bounded rationality would imply, and as our results support, people do not always accept an agent’s optimal advice. Consequently, and in contrast to all previous research known to us, we claim that a category of problems exists where the optimal solution is non-intuitive and thus less likely to be accepted by people. In these cases, suboptimal advice should be given. More formally, we model the decision making process as follows. Assume that agent A can either provide the optimal advice AD^* with a utility of V^{AD^*} if adopted by the user, or sub-optimal advice AD' with utility $V^{AD'}$. We further assume that A can autonomously choose between which advice to dispense (either AD^* or AD') to a person X . In theory, the system’s utility will be maximized when AD^* is chosen as $V^{AD^*} > V^{AD'}$. However, in practice, when providing any advice $AD \in \{AD^*, AD'\}$, it is possible that X will prefer following AD' , which she mistakenly believes to be optimal, instead. Therefore, the expected utility from providing an advice AD is given by $P(AD) \cdot (V^{AD} - V^{AD''})$ where $P(AD)$ is the probability the advised subject will adopt the advice given and $V^{AD''}$ is the utility from what she would have chosen to follow if not receiving the advice in the first place. The system’s expected-utility-maximizing advice is thus given by $\text{argmax}_{AD} P(AD) \cdot V^{AD}$. Thus, in practice, the agent should reason about the relative probability X will accept both advices.

This solution concept poses several challenges, such as intelligently generating suboptimal advices that will seem appealing to people in different decision settings, or predicting the probability a given advice (either optimal or a sub-optimal one) will be accepted by the advisee. In addition, the above calculation is a bit naive as it assumes X would have used AD'' even if not given the advice AD . In reality, of course, it is possible that giving the advice AD will result in changing what X considers to be optimal, and this also needs to be modeled. The design we propose in this paper attempts to bypass these complexities—while we do not fully attempt to solve for $\text{argmax}_{AD} P(AD) \cdot V^{AD}$ (or any of its more complicated versions) we do follow the principle of giving a slightly suboptimal advice that is more likely to be accepted by the advisee.

Our general design relies on the level of intuitiveness of the optimal advice in a given decision situation. The greater the non-intuitiveness, the greater the drift in the advice generated towards what might seem to be intuitive (and thus highly appealing) to people in general. The limitation of this design is that it only fits domains where the suboptimal advice’s utility monotonically decreases with respect to that advice’s distance from the optimal one, and where the distance parameter is well defined. Nevertheless, as we show in the following section, both canonical domains we consider for our experiments satisfy these conditions. A second limitation of this approach is that it requires the system designer to hypothesize what constitutes highly intuitive advice, towards which the proposed suboptimal advice should drift.

This can be done by collecting data of people’s decisions in similar decision situations when no advice is given. Other alternatives include interviewing people or simply relying on common sense and/or psychological biases reported in literature. Yet, it seems that the question of “what is intuitive” to people is unavoidable and should be addressed in any future architecture of suboptimal advising agents.

Domains Descriptions

Company Valuation (Samuelson and Bazerman 1985)

In this game there are two players: a seller of a company and a buyer interested in buying it. The true value of the company, denoted by p_{seller} , is privately held by the seller. The buyer’s best assessment is that the company’s worth is equally likely to be any number (i.e., uniformly distributed) between 0 and 100. The buyer knows that she can improve the value of the company by factor x . The buyer needs to come up with a “take it or leave it” offer to the seller. Thus, if the buyer’s offer for the company, denoted O_{buyer} , is above p_{seller} , the seller will accept it, and the buyer’s profit will be $x \cdot p_{seller} - O_{buyer}$. Otherwise the offer will be rejected and the buyer’s profit will be zero. The buyer’s goal in this game is to maximize her expected benefit.

Assuming the seller is willing to accept the offer for O_{buyer} , the buyer updates his assessment to conclude that p_{seller} is uniformly distributed between 0 and O_{buyer} , with a rational expectation of $\frac{O_{buyer}}{2}$. The expected profit of the buyer in this case is thus: $x \cdot \frac{O_{buyer}}{2} - O_{buyer} = \frac{O_{buyer}(x-2)}{2}$, meaning that the buyer’s optimal strategy in this game solely depends on the value x : for $x < 2$ the optimal strategy is to offer zero, whereas the optimal strategy for $x > 2$ is to offer 100. This is illustrated in Figure 1(a). While the game is easy to understand and its optimal strategy is relatively easy to compute, the solution is found to be highly non-intuitive for people, and indeed prior work reports substantial deviation in people’s offers from the optimum in experiments (Samuelson and Bazerman 1985). The main non-intuitiveness in the optimal solution is that although the company is known to be worth more to the buyer than to the seller, the optimal solution within the range $1 < x < 2$ is not to purchase it at all, and for $x > 2$ to make an offer that is the maximum worth of the company to the seller. In particular, the sharp transition from offering 0 to 100, exactly at $x = 2$, is confusing to people.

Birthday (Ball 1914)

In this game the participant needs to calculate the probability that in a class of N students, there will be at least two students celebrating their birthday on the same day. The correct answer is calculated as follows. First we calculate the probability $p(\bar{N})$ that all N birthdays are different: $p(\bar{N}) = 1 \cdot \frac{364}{365} \cdot \frac{363}{365} \cdots \frac{366-N}{365} = \prod_{i=1}^N \frac{366-i}{365}$. The probability that at least two students celebrating their birthday on the same day, denoted $p(N)$ is thus $1 - p(\bar{N})$. Figure 1(b) depicts the probability $p(N)$ as a function of the number of students in the class, N . As in Company Valuation game, this game is easy to explain, yet in contrast it is difficult for people to calculate the optimal strategy, as the mathemati-




Figure 1: The optimal advice; (a) Company Valuation, and (b) Birthday game.

cal truth in this case contradicts people’s intuition (Voracek, Tran, and Formann 2008). The non-intuitiveness of $p(N)$ in this case may derive from people’s tendency to estimate probability by “assessing availability”, or using “associative distance” (Tversky and Kahneman 1973). For example, one may estimate the probability of at least two people sharing the same birthday by counting the people she knows that share the *same birthday as hers*.

Advice Generation in Birthday Game and Company Valuation Game

Following our general approach that was outlined above in the “advice giving approach” section, we designed an advising agent for each of the two games that is capable of providing suboptimal advice. For the Company Valuation game we hypothesized that the greater the distance of the value of x from $x = 2$, the more intuitive the optimal solution would seem to people, where proposing the optimal solution for $x < 1$ or $x > 3$ is fully intuitive. This is straightforward, as the lower the value of x , the smaller the expected worth of the company to the buyer, making it more likely the person will accept low offers. Specifically, our advising agent provided the advice $O_{buyer} = 50 \cdot (x - 1)^3$ for $1 < x < 2$ (and 0 for $x < 1$), and $O_{buyer} = 50 \cdot \sqrt[5]{x - 2} + 50$ for $2 < x < 3$ (and 100 for $x > 3$), though any other functions following the above monotonicity principle are applicable.¹

For the Birthday game we hypothesized that misunderstanding $p(N)$ is rooted in the rapidness of change in the probability as the number of students (N) increases—the larger the rate of change, the less intuitive the optimal advice is to people. Therefore, our advising agent was designed to rely on the first derivative of $p(N)$ as a measure for the amount of deviation required from the true answer. Specifically, we used the function $p(N)^t = p(N) - 6 \cdot N \cdot dp(N)/dN$. Here, the coefficient 6 was set arbitrarily to determine the maximal deviation allowed from the optimal advice.

¹Since the goal of this paper is to provide a proof of concept we did not focus on finding the function that would provide the “best” performance and arbitrarily set what seemed to “make sense”. Also, our preference of a function with a slightly lesser aggressive convergence to the optimal value for $x > 2$ (compared to $x < 2$) is rooted in the different extent of non-intuitiveness of the optimal advice in both cases — as discussed in the previous section, the optimal advice for $x < 2$ is less intuitive for people.

Experimental Design

We implemented both games such that participants could interact with the system using a relatively simple graphical interface, thus facilitating interactions with a variety of people. Participants were recruited and interacted through Amazon Mechanical Turk (AMT) (AMT 2010).² Each participant took part in one out of three treatments: (a) playing the game without an advisor; (b) playing the game with the optimal advisor; and (c) playing the game with our suboptimal advisor. From the user-interface point of view, the advice was presented to participants by a virtual advisor that outputted its advice (e.g., the amount to offer for the Company Valuation or the probability in the Birthday game) though the participant was free to input any value as a response. Overall, we had 50 different participants for each treatment in each game (300 participants overall). Participants differed in age (21-60) and gender (37% men and 63% women). Each participant received thorough instructions of the game rules and her goal in the game. Participants were told of the compensation structure, which was composed of a show-up fee (the basic “HIT”) and a bonus which was linear in the participant’s performance in the experiment in order to encourage thoughtful participation. For the Company Valuation game, performance was measured as the average profit in all games played. For the Birthday game, performance was measured as the average absolute deviation from $p(N)$ in the games played.

After the instructions step, participants were asked to engage in practice games which they were encouraged to repeat until stating that they understood the games’ rules and had developed a game strategy (with a strict requirement for playing at least two practice games). Then, participants had to correctly answer a short quiz before continuing to the experiments whose results we logged for this study.

The experiment protocol for the Company Valuation game was as follows. Participants were asked to play 10 times the buyer side of the game, each time with a different (randomly selected) x value from the set $\{1, 1.2, 1.4, 1.6, 1.8, 2.2, 2.4, 2.6, 2.8, 3\}$, i.e., five values for which the optimal strategy is 0 and five for which it is 100. The decision to require participants to play 10 games was made primarily in order to push them to use their expected-benefit maximizing strategy. It has been shown that in repeated-play settings people’s strategies asymptotically approach the expected monetary value (EMV) strategy as the number of repeated plays increases (Wedell 2011; Klos, Weber, and Weber 2005; Keren and Wagenaar 1987; Barron and Erev 2003). In particular, the probability a person will prefer the option associated with the highest expected value is substantially greater than in single-play settings (Montgomery and Adelbratt 1982). To avoid any learning and carryover effects, participants were told that on each iteration they will be facing a different seller and a different company offered for sale. For the same reason, we did not inform participants of their result after each game was played, and they discovered the value of the company to the

²For a comparison between AMT and other recruitment methods see (Paolacci, Chandler, and Ipeirotis 2010).

seller and their corresponding profits in all games through a summary table which was only visible after completing the experiment.

The experiment protocol for the Birthday game was similar, except that each participant had to play only 5 games, differing in the number of students in the class (N , randomly selected from the set $\{17, 25, 38, 51, 70\}$ with corresponding true probabilities of $\{31.5\%, 56.8\%, 86.4\%, 97.4\%, 99.9\%\}$). For each game, the participant was informed of the value of N and was asked to input the probability of having at least two people celebrating their birthday on the same day.

Results and Analysis

Company Valuation Results

Figure 2(a) depicts the average offer made by participants in the Company Valuation game for different values of x . The two other graphs of the figure show the distribution of offers made, for $x < 2$ and $x > 2$ values, according to six classes of offers: the first is the optimal offer (0 and 100, depending on x), the second is of near-optimal offers (within a distance of 10 from the optimal offer) and the remaining are of classes progressively far from optimal offers. These results align with those reported by Samuelson and Bazerman (1985) who experimented in this domain with $x = 1.5$ settings, hence further validating our experiments—while they reported that 38% of their MBA students offered 60\$ and more, the analysis of our data for $x = 1.6$ reveals that 36% of the participants offered 60\$ and more.

As can be observed from the figure, participants' offers were mostly far from optimal for all x values, with a greater tendency to provide non-optimal offers for $x < 2$ values. Only 3% of the participants (averaging over the different x conditions) used the optimal offer. The size of near-optimal class was also small, with a slight advantage to the case where $x > 2$. Peoples' greater tendency to provide a non-optimal offer for $x < 2$ values may suggest that while the solution of both cases is non-intuitive, the $x < 2$ case is less intuitive than the $x > 2$ one, hence supporting the use of a non aggressive convergence to the optimal strategy as reported in our advice generation design. Furthermore, the graph supports our hypothesis that the greater the distance from $x = 2$, the closer we get to the optimal solution (and implicitly, the greater the intuitiveness of the optimal solution). This is exemplified by the decrease in the average offer as x decreases, for $x < 2$, and the increase in the average offer as x increases, for $x > 2$. It is also nicely evidenced in the distribution of offers: with both the increase in the value of x (for the case of $x < 2$) and the decrease in the value (for $x > 2$) we observe an increase in the percentage of participants that use the optimal or the near-optimal strategies. This increase becomes even sharper when switching to the two extremes at $x = 1$ and $x = 3$.

Figure 3 compares the offers made by participants when given an optimal and suboptimal advices and those made by participants when not receiving any advice. The first (left-most) graph in the figure depicts the average offer made by participants, according to the treatment used (none, op-

timal and suboptimal), as a function of the value x , similar to the graph given in Figure 2. In fact, the curve for the case where no advice is provided is the same as the one given in 2(a). The two other graphs of the figure comparatively show the distribution of offers made between the three treatments. Here, in contrast to 2, we present the aggregative data for all $x < 2$ and $x > 2$ values, due to space considerations. This figure shows that providing the optimal advice facilitates an improvement in the average offer made. Yet, from the distribution of offers we observe that the offers made are still relatively far from optimal: for the case of $x < 2$ only 35% of participants adopted the optimal or near-optimal strategy (middle graph). This combination of a relatively large improvement in the average offer and a relatively low percentage of participants adopting the optimal offer suggests that while giving the optimal advice does not lead people to an optimal offer, it does push them towards better ones. From Figure 3(b) we see the improvement in the percentage of participants that offered optimal or near-optimal amounts was almost entirely at the expense of the classes offering very far from optimal (offers greater than 50). For the case of $x > 2$ (right graph), we observe a similar pattern—a transition from far-from-optimal offers to optimal and near-optimal ones—though in this case the improvement is of a greater extent: almost 43% of the offers made are optimal or near-optimal and the improvement is at the expense of all other offer classes, i.e., of offers smaller than 90. The larger improvement in the $x > 2$ case is attributed the fact that the optimal strategy is more intuitive here compared to the $x < 2$ case as discussed above.

The suboptimal advice yielded an even greater improvement, both in the average offer made and the distributions of offers, compared to those achieved with providing the optimal advice: 64% of the offers made for $x > 2$ and 40% of the offers made for $x < 2$ were optimal and near-optimal. As expected, the improvement is mostly noticeable in the percentage of the near-optimal offers. Interestingly, for $x > 2$ we do observe an improvement in the percentage of participants that used the optimal strategy, despite receiving the suboptimal advice. These mostly correspond to the case of $x = 3$, for which the suboptimal advice equals the optimal one, according to our design, where 52%(!) of the participants used the optimal offer. However, even in other $x > 2$ cases we see that 16 – 30% (increasing as x increases) of participants used the optimal offer despite being provided a slightly lower one. We believe this is because, similar to the explanation given above, even the rejected advice has the power to change a person's strategy. In this case, therefore, people who gave the optimal offer were affected by the suboptimal advice in a way that made them realize that the optimal strategy is indeed to offer as much as possible.

One interesting finding in Figure 3(a) is that there is a difference between the average offer when using the optimal and suboptimal advices, both for the case of $x = 1$ and $x = 3$ (and similarly a difference between the offers distributions in these cases) In these two cases our suboptimal advice was the same as the optimal advice (as discussed in the previous section). Seemingly, this should have led to the same offers. Yet, the results reflect a substantial change in




Figure 2: Participants' offers in Company Valuation, for different x values, with no advice: (a) average offer; (b) distribution of offers for $x < 2$; (c) distribution of offers for $x > 2$.




Figure 3: Participants' offers in Company Valuation, for different x values, with different treatments: (a) average offer; (b) distribution of offers for $x < 2$; (c) distribution of offers for $x > 2$.

the offers made, with better results noted when suboptimal advice was given. We believe that the explanation for this is that participants who received suboptimal advices tended to increase their trust in the agent over time, compared to when given the optimal advice. Therefore, when the virtual advisor presented advice for the cases of $x = 1$ and $x = 3$, more people chose to adopt it. Overall, the differences between the offers in all three treatments were found, when using ANOVA, to be statistically significant ($p < 0.001$), indicating that using our suboptimal advice results in better offers being made.




Figure 4: Average expected profit for different x values.

Last, we present Figure 4 that depicts the average expected profit as a function of the advice given (none, optimal and suboptimal) for different x values. The expected profit of using an offer O_{buyer} is calculated as the product of the probability the offer is accepted and the expected benefit if accepted, i.e.: $\frac{O_{buyer}}{100} \cdot \frac{O_{buyer}(x-2)}{2}$. This measure is of much importance, since it is possible that despite following a strategy that is very different from the optimal one, the expected achieved performance is very close to the one achieved with latter. Therefore, this measure shows that suboptimal advice

facilitates people not only acting in a way that is closer to optimal, but their performance substantially improves. From the figure we see that indeed this is the case with our suboptimal advising in the Company Valuation game. For $x < 2$ the expected profit through using the optimal offer is zero, hence any offer which is not optimal leads to a negative expected benefit. With $x > 2$ any positive offer results in a positive expected profit. Overall, the differences between the expected benefit in all three treatments were found, when using ANOVA, to be statistically significant ($p < 0.001$), indicating that using our suboptimal advice in the Company Valuation game substantially increased expected profit.

Birthday Game Results

Figure 6 depicts the average performance of participants in the Birthday game for different values of N using the three treatments (no advice, optimal and suboptimal advice). Each of the three graphs relate to a different performance measure. The left and middle graphs present the participants' average answer and the absolute distance between the participants' answer and the correct answer, respectively. The observations made based on these graphs are similar to those made in the Company Valuation case: (a) when given no advice, participants' answers are very different from the true answer, demonstrating the non-intuitiveness of the correct answer. The most substantial deviation from the true answer is reported for $N = 38$ and $N = 51$, indicating that people find it difficult to believe that $p(N)$ is so great with such a moderate group of students. (b) When given the optimal advice, people generally provide substantially better answers, compared to when required to estimate the probability themselves, yet are still far from optimal. (c) When the agent provided our suboptimal advice, participants managed to further improve




Figure 5: Average absolute error in Birthday, under: (a) no advice; (b) optimal advice; and (c) suboptimal advice.




Figure 6: Birthday's results under different treatments: (a) average answer; (b) average absolute error; and (c) average inefficiency reduction.

the accuracy of their answers.

The right graph depicts the inefficiency improvement of using the suboptimal advice compared to the optimal one. To this end, the inefficiency of a given method is given by its absolute distance from the true answer. The inefficiency improvement thus measures the inefficiency decrease (in percentages) when agents use suboptimal advice rather than the optimal one, supporting the performance improvement claim of our method. All of the above reported findings are generally consistent across values of N . Similar to our analysis of the Company Valuation results, we present the distribution of the absolute difference between the correct answer and participants' answers to the Birthday game, for the three treatments according to the N value used (Figure 5). Here we use four classes: exact value, nearly-exact (within an absolute distance of 7), close answer (within an absolute distance of 7 – 20) and completely false answer. From the figure we see a pattern similar to the one observed in the Company Valuation results. In particular we see that while many choose to adopt the optimal advice (22% – 42%), many others (28% – 38%) still provide an answer that is very far from the correct one (an absolute distance of more than 20). With the suboptimal advice, while very few manage to calculate the correct answer (as this is not provided to them), many choose to adopt the nearly-exact answer, and even more change their answer to one that is relatively close to optimal—only 16% – 32% are still providing an answer that is very far from the correct one. Furthermore, even in cases where the suboptimal advice is not accepted as is, participants tend to use a more accurate answer. Here, once again, the differences between the absolute errors obtained in all three treatments were found, when using ANOVA, to be statistically significant ($p < 0.001$).

Conclusions and Future Work

The results reported in this paper suggest that domain advice agents should at times intentionally provide people with suboptimal advice instead of automatically providing the optimal one. We see much innovation in these findings, as the approach they conclusively support is fundamentally different from the traditional advice provisioning approach (Lesser 1999; Grosskopf, Bereby-Meyer, and Bazerman 2007). To the best of our knowledge, there has not been any attempt to date to empirically demonstrate that our approach of generating suboptimal advice can facilitate better results, or to suggest domains where it can be suitable. In this paper we make the first step towards the development of suboptimal-based advising agents by providing a proof-of-concept for the success of this paradigm in two canonical domains, as well as the proposed design principle of relying on the intuitiveness of the optimal advice as the main measure for deciding the extent to deviate from it.

We are currently studying several extensions for future work. First, we hope to create learning agents to predict if an individual is more or less likely to accept optimal advice. This will facilitate the development of adaptive agents that can decide, based on very few initial interactions whether to stick with the optimal advice or to switch to suboptimal ones. Furthermore, we agree that people do at times consistently accept optimal advice, especially in less complex domains. We hope to further study how to categorize different problems so we can predict a-priori the likelihood people will accept the optimal advice. We also aim to utilize aspects of trust-building in repeated-advice-provisioning interactive settings and additional designs for producing suboptimal advice.

Acknowledgement

This work was partially supported by ISF grant 1083/13.

References

- AMT. 2010. Amazon mechanical turk. <http://www.mturk.com/>.
- Azaria, A.; Rabinovich, Z.; Kraus, S.; Goldman, C. V.; and Gal, Y. 2012. Strategic advice provision in repeated human-agent interactions. In *Proc. of AAAI*.
- Azaria, A.; Hassidim, A.; Kraus, S.; Eshkol, A.; Weintraub, O.; and Netanel, I. 2013. Movie recommender system for profit maximization. In *RecSys*.
- Azaria, A.; Kraus, S.; Goldman, C. V.; and Tsimhoni, O. 2014. Advice provision for energy saving in automobile climate control systems. In *IAAI*.
- Ball, W. W. R. 1914. *Mathematical recreations and essays*. MacMillan.
- Barron, G., and Erev, I. 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making* 16(3):215–233.
- Bharati, P., and Chaudhury, A. 2004. An empirical investigation of decision-making satisfaction in web-based decision support systems. *Decision support systems* 37(2):187–197.
- Blackhart, G. C., and Kline, J. P. 2005. Individual differences in anterior EEG asymmetry between high and low defensive individuals during a rumination/distraction task. *Personality and Individual Differences* 39(2):427–437.
- Buntain, C.; Golbeck, J.; Nau, D.; and Kraus, S. 2012. Advice and trust in games of choice. In *Privacy, Security and Trust (PST), 2012 Tenth Annual International Conference on*, 157–158. IEEE.
- Chalamish, M.; Sarne, D.; and Lin, R. 2012. The effectiveness of peer-designed agents in agent-based simulations. *Multiagent and Grid Systems* 8(4):349–372.
- Drake, R. A. 1993. Processing persuasive arguments: Discounting of truth and relevance as a function of agreement and manipulated activation asymmetry. *Journal of Research in Personality* 27(2):184–196.
- Elmalech, A., and Sarne, D. 2014. Evaluating the applicability of peer-designed agents for mechanism evaluation. *Web Intelligence and Agent Systems* 12(2):171–191.
- Elmalech, A.; Sarne, D.; and Agmon, N. 2014. Can agent development affect developers strategy? In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Elmalech, A.; Sarne, D.; and Grosz, B. J. 2014. Problem restructuring for better decision making in recurring decision situations. *Autonomous Agents and Multi-Agent Systems* 1–39.
- Grosskopf, B.; Bereby-Meyer, Y.; and Bazerman, M. 2007. On the robustness of the winners curse phenomenon. *Theory and Decision* 63(4):389–418.
- Hajaj, C.; Hazon, N.; Sarne, D.; and Elmalech, A. 2013. Search more, disclose less. In *AAAI*.
- Hajaj, C.; Hazon, N.; and Sarne, D. 2014. Ordering effects and belief adjustment in the use of comparison shopping agents. *Proc. of AAAI*.
- Hazon, N.; Lin, R.; and Kraus, S. 2013. How to change a group's collective decision? In *Proc. of IJCAI*, 198–205.
- Kahneman, D. 2000. A psychological point of view: Violations of rational rules as a diagnostic of mental processes (commentary on stanovich and west). *Behavioral and Brain Sciences* 23:681–683.
- Keren, G., and Wagenaar, W. A. 1987. Violation of utility theory in unique and repeated gambles. *Journal of experimental psychology. Learning, memory, and cognition* 13(3):387–391.
- Klos, A.; Weber, E. U.; and Weber, M. 2005. Investment decisions and time horizon: Risk perception and risk behavior in repeated gambles. *Management Science* 51(12):1777–1790.
- Lesser, L. M. 1999. Exploring the birthday problem with spreadsheets. *The Mathematics Teacher* 407–411.
- Mash, M.; Lin, R.; and Sarne, D. 2014. Peer-design agents for reliably evaluating distribution of outcomes in environments involving people. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 949–956. International Foundation for Autonomous Agents and Multiagent Systems.
- Montgomery, H., and Adelbratt, T. 1982. Gambling decisions and information about expected value. *Organizational Behavior and Human Performance* 29(1):39 – 57.
- Paolacci, G.; Chandler, J.; and Ipeirotis, P. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision Making* 5(5):411–419.
- Power, D. J., and Sharda, R. 2009. Decision support systems. *Springer Handbook of Automation* 1539–1548.
- Rosenfeld, A., and Kraus, S. 2009. Modeling agents through bounded rationality theories. In *IJCAI 2009*, 264–271.
- Rosenfeld, A., and Kraus, S. 2015. Providing arguments in discussions based on the prediction of human argumentative behavior. In *AAAI*.
- Samuelson, W. F., and Bazerman, M. H. 1985. Negotiation under the winners curse. *Research in experimental economics* 3:105–38.
- Selten, R.; Abbink, K.; and Cox, R. 2005. Learning direction theory and the winners curse. *Experimental Economics* 8(1):5–20.
- Tversky, A., and Kahneman, D. 1973. Availability: A heuristic for judging frequency and probability. *Cognitive psychology* 5(2):207–232.
- Voracek, M.; Tran, U. S.; and Formann, A. K. 2008. Birthday and birthmate problems: Misconceptions of probability among psychology undergraduates and casino visitors and personnel 1. *Perceptual and motor skills* 106(1):91–103.
- Wedell, D. H. 2011. Evaluations of single-and repeated-play gambles. *Wiley Encyclopedia of Operations Research and Management Science*.