

Observed Variability and Values Matter: Toward a Better Understanding of Information Search and Decisions from Experience

KATJA MEHLHORN^{1*}, NOAM BEN-ASHER¹, VARUN DUTT² and CLEOTILDE GONZALEZ¹

¹*Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA*

²*School of Computing and Electrical Engineering, School of Humanities and Social Sciences, Indian Institute of Technology, Mandi, India*

ABSTRACT

The search for different options before making a consequential choice is a central aspect of many important decisions, such as mate selection or purchasing a house. Despite its importance, surprisingly little is known about how search and choice are affected by the observed and objective properties of the decision problem. Here, we analyze the effects of two key properties in a binary choice task: the options' observed and objective values, and the variability of payoffs. First, in a large public data set of a binary choice task, we investigate how the observed value and variability relate to decision-makers' efforts and preferences during search. Furthermore, we test how these properties influence the chance of correctly identifying the objectively maximizing option, and how they affect choice. Second, we designed a novel experiment to systematically analyze the role of the objective difference between the options. We find that a larger objective difference between options increases the chance for correctly identifying the maximizing option, but it does not affect behavior during search and choice. Copyright © 2013 John Wiley & Sons, Ltd.

KEY WORDS decisions from experience; information search; bounded rationality; maximization; payoff variability; memory

In many important real-life decisions, we seek out information about different possibilities before making a choice. For example, most people would not purchase a house without looking at several possibilities or marry a partner without having gotten to know him/her first. While the question of information search was often sidestepped in classical decision literature, significant progress has been made toward a better understanding of search and choice in recent years. For example, we have learned that people seem to generally search for little information before making a choice (for an overview, see Hau, Pleskac, & Hertwig, 2010), and that search can be affected by several characteristics of the decision maker and the choice ecology (for an overview, see Lejarraga, Hertwig, & Gonzalez, 2012). Relatively little is known, however, about how search is related to properties of payoffs that are actually observed during search, and how these properties affect subsequent choice.

COSTS OF INFORMATION SEARCH VERSUS ACCURACY OF CHOICE

A general finding in the decision-making literature is that people tend to search for "little" information before making a consequential choice between different options (e.g., Hau et al., 2010; Hertwig & Pleskac, 2010). This behavior can be advantageous. For example, relying on smaller amounts of information reduces the explicit (e.g., monetary) and implicit (e.g., cognitive) costs of information search (Hau, Pleskac, Kiefer, & Hertwig, 2008), reduces demands on working-memory capacity (Kareev, 2000; Rakow, Demes,

& Newell, 2008), tends to amplify differences between the options and thereby renders choice easier (Hertwig & Pleskac, 2008, 2010), and maximizes the time available for other decisions (Vul, Goodman, Griffiths, & Tenenbaum, 2009). As suggested by the statistical law of large numbers, however, using less information can come at a price: The chance of forming a valid representation of a decision problem's objective payoff structure strongly decreases with less and less information gathered (Brehmer, 1980; Fiedler, 2000; Hau et al., 2008; Johnson, Budescu, & Wallsten, 2001). Specifically, people are likely to under-experience low-probability payoffs with smaller samples and thereby risk getting a wrong impression about the options' objective values (Fox & Hadar, 2006; Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Pleskac, 2010). As a result, they might make suboptimal decisions, because the samples do not allow them to correctly identify the objectively maximizing option.

This trade-off between the costs of searching and the accuracy of choice has inspired many studies that investigated how much information should optimally be acquired (e.g., Fiedler & Kareev, 2006; Gittins, 1979; Hertwig & Pleskac, 2010; Kareev, 2005; Vul et al., 2009; Wallsten, Budescu, Erev, & Diederich, 1997). Maybe most prominently, the Gittins index (Gittins, 1979) predicts for each point in time which option should be selected to make an "optimal" decision, given the exact history of experiences. Unfortunately, such prescriptive solutions often require vast computational resources and make strong and not always tenable assumptions about the environment and the decision maker's goals (Cohen, McClure, & Yu, 2007). Here, we evaluate how peoples' actual experiences during information search are related to search effort and how they, in turn, affect choice. Specifically, we investigate the effects of two key ecological properties of information search and choice: the variability and value of the payoffs.

*Correspondence to: Katja Mehlhorn, Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA. E-mail: katjam@cmu.edu

TWO KEY ECOLOGICAL PROPERTIES: VARIABILITY AND VALUE OF PAYOFFS

To understand how observed payoff variability and values might affect information search and choice, one first needs to consider the decision-maker's goals. A considerable amount of research has been devoted to how people make choices based on their experiences (e.g., Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Busemeyer & Townsend, 1993; Erev & Barron, 2005; Gonzalez & Dutt, 2011; Hertwig, in press; Tsetsos, Chater, & Usher, 2012). Building on the classic ideas of prospect theory (Kahneman & Tversky, 1979), decision-maker's goals during choice are generally assumed to be related to the *value* and the *variability* of observed payoffs. When deciding between options, people tend to choose the one where they observed higher values (e.g., Busemeyer, 1985; Gonzalez & Dutt, 2012; Tsetsos et al., 2012). Variability in the observed payoffs can affect this value maximization in different ways. Some have concluded that overall it reduces peoples' reliance on their experiences, by moving choice toward a random choice (the "payoff variability effect": Busemeyer & Townsend, 1993; Erev & Barron, 2005; Myers & Sadler, 1960). Others have argued that variability moves preferences away from the risky option, resulting in risk aversion (e.g., shown in the "hot stove effect": Denrell & March, 2001). Recent evidence suggests that payoff variability can cause risk aversion as well as risk seeking behavior, depending on whether high or low outcomes are more salient in the choice environment (Tsetsos et al., 2012).

People's goals during information search are less well understood than during choice (Gonzalez & Dutt, under review). It has been argued that during information search, people explore the different options at random with the single purpose of gathering information (Erev, Ert, & Yechiam, 2008; also see Cohen et al., 2007). This idea of random search is contradicted by evidence from Hills and Hertwig (2010, 2012), who found that people consistently followed one of two distinct exploratory strategies during information search, and by evidence from Gonzalez and Dutt (2011, 2012), who found that the alternation between options during search decreased over time. But could search also be affected by the properties of the observed outcomes? Recent studies suggest this might be the case. Rakow and Newell (2010) proposed that even if people do not receive the observed outcomes' rewards during search, they might use the outcomes to test their hypotheses about the different options. Such hypotheses might concern the observed variability and value of the outcomes. Lejarraga et al. (2012) found that people tend to search more in options where they observed variable outcomes, relative to options with stable outcomes, suggesting a preference for risky options during information search. Gonzalez and Dutt (2011, 2012) showed that people tend to search more in options with higher observed values than in options where lower values were observed, suggesting a confirmatory search for the maximizing option.

While there is thus some evidence for the possible effects of observed variability and value, open questions remain about whether and how these two factors jointly influence search and choice. On the basis of the literature described

earlier, we will focus on the following four questions. First, is the total amount of information search related to the variability and value of the observed outcomes? Second, are peoples' preferences during information search related to the observed variability and value? Third, does the observed variability and value affect the chance of correctly identifying the objectively maximizing option? Fourth, how do variability and value, as observed during search, affect final choice? In a first step toward answering these questions, we reanalyze a large public dataset of a binary choice task. Subsequently, we present a novel experiment designed to test the possible role of the objective difference between the options.

REANALYSIS OF THE TECHNION PREDICTION TOURNAMENT DATASET

We reanalyzed a dataset from the Technion Prediction Tournament (TPT, Erev et al., 2010), which is especially suited for this purpose because it is a large and representative dataset, using the sampling paradigm (Hertwig et al., 2004). In this paradigm, search and choice are separated into two distinct phases. In an initial sampling phase, participants are free to explore the different options without explicit costs or consequences for as long as they want. Only after they feel satisfied with the obtained information, they proceed to the decision phase where they make a consequential choice between the options (Figure 1).

The set of problems used in the TPT was large and representative of a broad diversity of decision situations. Each of 80 participants solved 30 out of 120 choice problems. All problems consisted of a safe option yielding a fixed medium outcome on every draw, and a risky option yielding a high outcome with some probability p , and a low outcome with the complementary probability $1-p$, as shown in Figure 1. The problems were generated with a problem selection algorithm, which ensured that for each participant equal

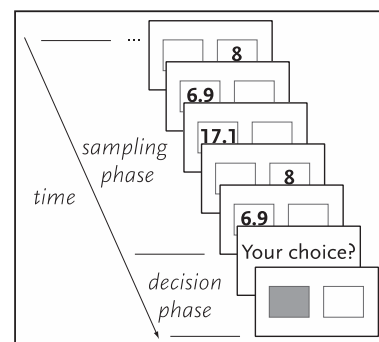


Figure 1. The sampling paradigm as used in the TPT dataset (Erev et al., 2010). During the sampling phase, participants obtained information about two options through a self-structured and self-terminated search without explicit costs. In the decision phase, they selected the option from which a single draw would be made for determining the problems' payoff. In each problem, one option was "safe" with a medium payoff and a probability of 1 (8 in the example), while the other option was "risky" with a high payoff with probability p and a low payoff with probability $1-p$, (17.1 with $p = .1$ and 6.9 with $p = .9$ in the example)

proportions of problems were from one of three payoff domains (positive, negative, and mixed outcomes) and from one of three probability ranges (high, medium, and low probability of the high risky outcome).¹ Most importantly, peoples' experiences from information search will not always correctly reflect the objective underlying values of the problem, as discussed in the introduction. To illustrate this, consider the example shown in Figure 1. The objective underlying value of the risky option is *lower* than the value of the safe option (7.92 vs. 8). However, on the basis of the actually observed outcomes, the value of the risky option appears *higher* than the value of the safe option (10.3 vs. 8). Thus, the observed information points the participant away from the objectively maximizing option.

Dependent measures and data classification

We use four dependent measures to answer the four questions outlined earlier: (1) the total number of samples taken from both options, as a measure of the amount of information searched for; (2) the proportion of samples from the risky option, as indication of participants' preferences during search; (3) the proportion of cases where the observed information correctly indicates the objectively maximizing option, as indication of the validity of the observed information; and (4) the proportion of final choices from the risky option, as indication of participants' preferences during choice.

To systematically evaluate the combined effects of the observed variability and value on these dependent measures, we classified the data based on participants' experiences in the sampling phase. In each decision problem, participants either experienced variability or not, by observing two or only one of the outcomes in the risky option (one or two risky outcomes). The observed value of the risky option (i.e., the average observed outcome) was either lower or higher than the safe option's value (lower or higher). As shown in Table 1, the data was distributed approximately equally across the resulting four categories.

Results

Is the total amount of search related to observed variability and value?

To investigate the total amount of search, we used the number of samples taken from both options by each participant in each problem. As expected, participants took relatively few samples (*average median* = 12.0, *SD* = 6.6). Figure 2 shows a positive relationship between sample size and observed variability. The average median sample size was higher in cases where both risky outcomes were observed than when only one outcome was observed (Wilcoxon signed ranks test: $z = -3.37$, $p < .001$). Unlike variability, the observed values were not related to sample size ($z = -.36$, $p = .719$).

The positive relationship between sample size and the number of observed outcomes in the risky option agrees with the

results from Lejarraga et al. (2012). However, as discussed in their paper, the causality behind such a relationship is unclear. On the one hand, it is a statistical fact that large samples result in a higher probability of observing the two outcomes in the risky option. Thus, it might simply be the case that participants who sampled more were more likely to have observed both outcomes. But it is also possible that observing both outcomes prompts participants to increase their search effort and take more samples.

To disentangle these possibilities, we examined the sample size in all cases where at least k samples were taken, as a function of whether or not participants observed both risky outcomes within the first k samples (for this analysis, we collapsed across participants and observed values). As shown in Table 2 for four different values of k , sample size tended to be *smaller* after observing both risky outcomes than after observing only one of the outcomes. A complete analysis of all possible values of k shows that sample size was significantly smaller (as indicated by $p < .05$ in Mann–Whitney U -tests) after observing both risky outcomes for $k = 5$ to $k = 14$, as well as for $k = 28$, $k = 30$, and $k = 31$. Most importantly, there is no value of k for which the sample size was *larger* after observing both risky outcomes. Table 2 furthermore shows the relationship between samples sizes and the chance to observe payoff variability. Out of all 1907 cases in which at least 4 samples were taken, only 19% observed both risky outcomes within the first 4 samples. This proportion increased at a diminishing rate to .33, .43, and .48 within the first 8, 12, and 16 samples, respectively.

These results suggest that the positive relationship between sample size and observed outcome variability as shown in Figure 2 is due to the increased chance of observing both risky outcomes with larger samples sizes, rather than due to an increased search effort after observing both risky outcomes. Instead, observing both risky outcomes tended to even have a *negative* effect on sample size: Participants were inclined to sample less after observing both risky outcomes within the first k samples than after observing only one of the outcomes. As we will discuss later, this behavior might be in anticipation of payoff variability. Participants may have learned over the course of the experiment that one of the options had two outcomes while the other had only one, and therefore started to search for these outcomes through repeated problems that shared the same structure (see also Lejarraga et al., 2012).

Are preferences during search related to observed variability and value?

To test preferences during search, we calculated the proportion of samples from the risky option (p_{Risky}) for each participant and problem. Overall, p_{Risky} was above chance (*mean* $p_{Risky} = .55$, *SD* = .08, $t(70) = 6.01$, $p < .001$), indicating a general preference for the risky option during sampling. Figure 3 shows the average p_{Risky} in each of the four categories. As confirmed by a repeated-measures ANOVA using the factors of observed variability and value, participants sampled the risky option more often when both risky outcomes were

¹The data and a detailed description of the methods can be found at <http://tx.technion.ac.il/~erev/Comp/Comp.html>.

Table 1. The four categories formed by observed variability and value, together with the number (proportion) of the 2102 total participant and problem combinations (cases) used in our analysis. Each cell has data from all 71 analyzed participants

		Observed value of the risky compared with the safe option	
		Lower (.51)	Higher (.49)
Number of observed outcomes in the risky option	One (.55)	Only <i>one outcome</i> of the risky option was observed. The observed value of the risky option was <i>lower</i> than the observed value of the safe option. 581 cases (.28)	Only <i>one outcomes</i> of the risky option was observed. The observed value of the risky option was <i>higher</i> than the observed value of the safe option. 568 cases (.27)
	Two (.45)	The <i>two outcomes</i> of the risky option were observed. The observed average value of the risky option was <i>lower</i> than the observed value of the safe option. 498 cases (.24)	The <i>two outcomes</i> of the risky option were observed. The observed average value of the risky option was <i>higher</i> than the observed value of the safe option. 455 cases (.22)

Note: Note that in 31 cases, which are not shown in the table and will not be used in our analyses, participants observed the same values in the risky and the safe option. Note also that 8 of the 80 Technion Prediction Tournament (TPT) participants did not contribute cases to all four categories and are therefore excluded from the table and our analyses. Of those 8 participants, 5 did not observe the second risky outcome in any of the problems and 2 did not observe it in either problems with higher or lower risky values; 1 participant did not observe lower risky values in problems with one risky outcome. The data of 1 additional participant is missing from the TPT dataset, leaving us with data of 71 participants for all analyses.

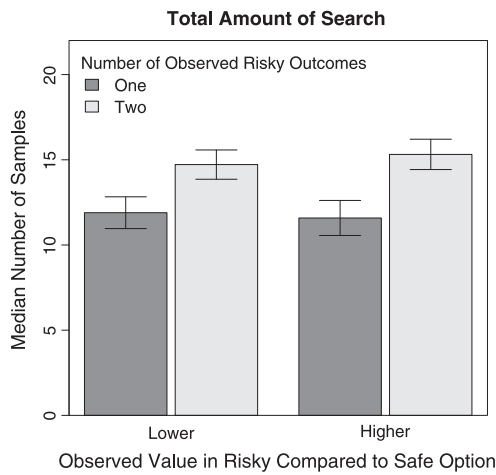


Figure 2. Average ($\pm 1 SE$) median sample size, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option)

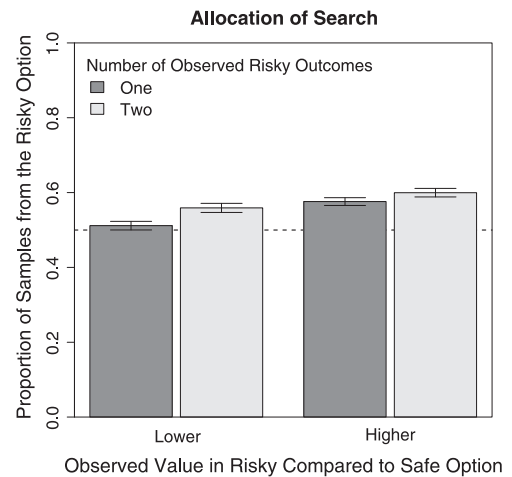


Figure 3. Mean ($\pm 1 SE$) proportion of samples from the risky option, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option). The dashed line indicates chance level

Table 2. Number (proportion) of cases and median sample size in cases where at least k samples were taken, as a function of whether one or both risky outcomes were observed during the initial k samples. Also shown are the results of Mann–Whitney U -tests comparing the sample size between cases with one and two outcomes at each k

k	Outcomes observed in the risky option	Number of cases (proportion)	Median Sample size	Mann–Whitney U -tests
4	One	1542 (.81)	12	$p = .068$
	Two	365 (.19)	11	
8	One	953 (.67)	16	$p < .001$
	Two	474 (.33)	14	
12	One	569 (.57)	20	$p = .001$
	Two	431 (.43)	18	
16	One	367 (.52)	22	$p = .359$
	Two	339 (.48)	22	

observed than when only one was observed, $F(1,70) = 14.47$, $p < .001$. Participants also sampled the risky option more often when observing higher risky values than when observing

lower risky values, $F(1,70) = 42.19$, $p < .001$. There was no interaction between the two factors, $F(1,70) = 2.14$, $p = .148$. Additional t -tests against chance level (.5) further verify that $pRisk_{y}$ did not differ from chance in the “lower risky value and one-observed-risky-outcome” category (leftmost bar in Figure 3; $t(70) = 1.00$, $p = .319$). $pRisk_{y}$ was above chance in the remaining three categories (all p -values $< .001$).

These results suggest that participants’ preferences during search are related to the value of observed outcomes. They sampled more from the risky option when higher values were observed in this option. Search also was related to the variability of observed outcomes. Participants sampled more from the risky option, when they observed both risky outcomes than when observing only one risky outcome. Interestingly, even in cases where only one of the risky outcomes was observed, we find no preference for the safe option (dark gray bars in Figure 3). In the “higher risky value” category, we find a preference for the risky option, which is well explained by confirmatory search for the

maximizing option. However, in the “lower risky value” category, confirmatory search would predict a preference for the safe option. The lack of such a preference might be another result of participants anticipating variability. Because people received many problems with the same structure in the TPT, they might have learned this structure over the course of the experiment, allowing them to “eyeball” the risky option from the observed outcomes (Lejarraga et al., 2012).

To test whether or not preferences during search indeed changed over the course of the TPT experiment, we compared $pRisky$ between the first and the last two of the 30 problems for cases where those problems were in the “lower risky value and one-observed-risky-outcome” category (for this analysis, we collapsed across participants). $pRisky$ increased significantly from .43 ($SD = .12$) in the first two, to .53 ($SD = .13$) in the last two problems, $t(67) = -3.42, p = .001$. This suggests that participants did indeed learn over the course of the experiment to expect more than one outcome in the risky option, which in turn triggered additional search in this option.

How do variability and value affect the chance of identifying the objectively maximizing option?

To assess the chance of correctly identifying the objectively maximizing option, we coded each participant and problem on whether or not the observed values pointed in the same direction as the objective values. Overall, the correspondence between observed and objective values was above chance level ($mean\ proportion = .58, SD = .11, t(70) = 6.03, p < .001$). As shown in Figure 4, correspondence was greater if both risky outcomes were observed than if only one of the outcomes was observed, $F(1,70) = 52.23, p < .001$. There was no main effect of the observed value, $F(1,70) < 1$, and no interaction, $F(1,70) < 1$. T -tests against chance level further illustrate the relevance of observed variability. The correspondence between observed and objective values was significantly




Figure 4. Mean ($\pm 1 SE$) proportion of cases in which the observed values point in the same direction as the objective values, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option). The dashed line indicates chance level

above chance when both risky outcomes were observed (both p -values $< .001$) but did not differ from chance otherwise (both p -values $> .2$).

How do the variability and value, as observed during search, affect choice?

To assess preferences in the final choice, we coded each participant and problem on whether the risky or the safe option was selected in the final choice. Overall, we find a lower proportion of risky final choices than safe choices ($mean\ pRisky = .46, SD = .12, t$ -test against chance level: $t(70) = -2.42, p = .018$). Figure 5 shows that choices strongly depended on the values that were observed during sampling. Participants were less likely to choose the risky option after observing lower risky values than after observing higher risky values, $F(1,70) = 938.10, p < .001$. Observed outcome variability did not have a main effect on the final choice, $F(1,70) < 1$. Instead, we find an interaction between the two factors, $F(1,70) = 40.16, p < .001$: When both risky outcomes were observed during sampling, participants relied less on the observed values than when only one of the risky outcomes was observed.

Discussion

Our results support the idea that during information search, people test hypotheses about the options' payoffs (Rakow & Newell, 2010). Search was related to payoff variability. By controlling for observed variability in the initial samples, we were able to show that *anticipated*, rather than observed, variability seemed to drive the amount of information search. This idea of searching for anticipated variability was further supported by participants' preferences during search, as participants learned to search more in the risky option over the course of the experiment. Search was also related to the observed values, as participants sampled more from the option with higher observed values than from the option with lower values. Such a preference for options with higher




Figure 5. Mean ($\pm 1 SE$) proportion of final choices for the risky option, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option). The dashed line indicates chance level

values is strikingly similar to behavior in repeated choice paradigms where every decision is consequential (Gonzalez & Dutt, 2011). Given that participants did not receive actual rewards during search, their preferences for the option with higher values might be an instance of the confirmation bias, where people search for information that confirms their current hypothesis while disregarding other evidence (Nickerson, 1998; Wason, 1960).

The results further shed light on how observed problem properties influence the chance of correctly identifying the objectively maximizing option, and how they affect choice. In more than half of all cases, participants stopped sampling before they observed both risky outcomes, thereby reducing the correspondence between observed and objective maximization to chance level. In cases where both risky outcomes were observed, the chance to identify the maximizing option increased to 68%, which is still far from perfect. Despite the overall low validity of their experiences, participants strongly relied on these experiences for their consequential choice. Reliance on experience was stronger in cases where only one risky outcome was observed than in cases with both risky outcomes, but it was not affected by the observed values. Thus, in the TPT dataset, observed variability generally moved choice preferences toward random choice, rather than moving them away from the risky option. As we will show in the experiment in the succeeding text, this move toward random choice might be related to increased difficulty when trying to assess an option's value from variable outcomes (cf. Tsetsos et al., 2012).

EXPERIMENT: EFFECTS OF THE OBJECTIVE DIFFERENCE BETWEEN VALUES

In the analyses of the TPT experiment, we investigated the effects of *observed* properties of the problem. One open question is how those observed properties are related to the *objective* payoff structure. As described earlier, the algorithm used to generate the TPT problems ensured that different underlying properties, such as the payoff domain and the probability of the high risky payoffs, were equally distributed across the decision problems. However, the differences between the options' objective expected values (EV differences) were relatively small in all problems (absolute EV differences ranged from .004 to 4.52, *median* = .24). Such small EV differences are typical for the investigation of risky choice, because they make it possible to show the effects of decision biases. However, they might have affected our results in several ways.

On the one hand, the possible gain from correctly selecting the objectively maximizing option is relatively low in problems with small EV differences. Participants may realize this and consequently invest less effort in search and choice. On the other hand, small EV differences make it more difficult to correctly identify an objectively maximizing option from sampled information. Participants might realize this and consequently invest *more* effort in information search. To systematically test these possible effects, we designed an

experiment where the objective EV difference between the options was manipulated.

Methods

Participants

A total of 320 participants were recruited on Amazon Mechanical Turk. Out of those, 26 participants were excluded from the analysis because they took less than two samples in at least one of the problems. The remaining 294 participants had a mean age of 27.1 years ($SD = 8.4$, range = [18; 66]), and 25% of them were women. Access for participation was restricted to US IP addresses.

Material and procedure

We generated 16 different decision problems in which the EV difference was manipulated, and other relevant problem characteristics (the domain: gain vs. loss; the objectively maximizing option: safe vs. risky; and the frequency of the high risky outcome: rare vs. frequent) were counterbalanced. To manipulate the EV difference between the options, we adjusted the value of the safe option as far as possible without creating deterministically dominant options (i.e. the two risky outcomes always bracketed the safe outcome [Table 3]). As in the TPT dataset, each problem consisted of a safe option yielding a fixed medium outcome on every draw, and a risky option yielding a high outcome with some probability p , and a low outcome with the complementary probability $1-p$. To prevent learning across the experiment, each participant solved only two of the problems. The two problems were assigned in a pseudorandom, balanced manner, ensuring that both problems had the same EV difference (small or large) but also that they were from different domains (gain and loss). Within those constraints, all problems were paired together equally often. The order of presentation of the two problems, as well as the placement of the options on the screen (right or left), was counterbalanced across participants.

Participants were instructed that they would be playing two games; in each of which they could explore the possible outcomes of the two options before making a final choice between the options. The gains/losses obtained in this final choice would be added to/subtracted from their base payment of 10 cents. Every 10 dollars of "game money" would be worth 1 US cent, resulting in an additional gain of up to 25 cents. To control participants' expectations of the problems' payoff structure, they were informed that: "One of the alternatives is associated to a sure outcome. The other alternative is associated to two different possible outcomes, each of which occurs with a certain probability."

Dependent measures and data classification

We used the same dependent measures as in our analyses of the TPT dataset, along with one additional measure: the difference between observed values, which will be explained in the succeeding text. We collapsed across participants in all analyses, because each participant solved only two problems.

Table 3. Manipulated factors and the 16 different decision problems as used in the experiment. In each problem, one option was “safe” with a medium payoff and a probability of 1. The other option was “risky,” with a high payoff and with probability, p , which is shown in the table. The second risky payoff was 0 in all problems and occurred with the complementary probability, $1-p$

Problem characteristics				Decision problems			
Size of EV difference	Domain	Objectively maximizing option	Frequency of high outcome	Safe option	Risky option	EV risky	EV difference (R-S)
Small	Gain	Safe	Rare	\$24.5 (1)	\$120 (.2)	24	-.5
			Frequent	\$24.5 (1)	\$30 (.8)	24	-.5
		Risky	Rare	\$23.5 (1)	\$120 (.2)	24	.5
			Frequent	\$23.5 (1)	\$30 (.8)	24	.5
	Loss	Safe	Rare	\$-23.5 (1)	\$-120 (.2)	-24	-.5
			Frequent	\$-23.5 (1)	\$-30 (.8)	-24	-.5
		Risky	Rare	\$-24.5 (1)	\$-120 (.2)	-24	.5
			Frequent	\$-24.5 (1)	\$-30 (.8)	-24	.5
Large	Gain	Safe	Rare	\$29.5 (1)	\$120 (.2)	24	-5.5
			Frequent	\$29.5 (1)	\$30 (.8)	24	-5.5
		Risky	Rare	\$18.5 (1)	\$120 (.2)	24	5.5
			Frequent	\$18.5 (1)	\$30 (.8)	24	5.5
	Loss	Safe	Rare	\$-18.5 (1)	\$-120 (.2)	-24	-5.5
			Frequent	\$-18.5 (1)	\$-30 (.8)	-24	-5.5
		Risky	Rare	\$-29.5 (1)	\$-120 (.2)	-24	5.5
			Frequent	\$-29.5 (1)	\$-30 (.8)	-24	5.5

As shown in Table 4, we classified the data based on the observed variability and values separately for problems with small and large objective EV differences.

Results

Does the EV difference affect the total amount of search?

As shown in Figure 6, sample size did not differ between problems with small (*median*=4.0, *quartiles*=[2, 8]) or large (*median*=4.0, *quartiles*=[2, 8.75]) EV differences (Mann-Whitney *U*-Test: $z = -.35, p = .724$). As in the TPT, sample size was greater when both risky outcomes were observed than when only one outcome was observed (small EV difference: $z = -12.97, p < .001$; large EV difference: $z = -12.93, p < .001$); and sample sizes did not differ as a function of the options’ observed values (both $p > .6$).

Does the EV difference affect preferences during search?

Figure 7 shows participants’ preferences during search, indicated by the average proportion of samples from the risky option (p_{Risky}). Preferences during search were not affected by the EV difference, as confirmed by an ANOVA with the

factors of EV difference, observed variability, and observed value, which showed no significant main effect or interactions involving EV difference (all $p > .14$). As in the TPT, participants sampled more often from the risky option when both risky outcomes were observed than when only one outcome was observed, $F(1580) = 71.58, p < .001$; and when observing higher risky values than when observing lower risky values, $F(1580) = 34.49, p < .001$. As expected, because of the absence of learning opportunities over the course of the experiment, we find a significant interaction between observed variability and value, $F(1580) = 34.76, p < .001$. If no variability was observed in the risky option, sampling preferences were affected by observed values only. This conclusion is supported by *t*-tests against chance level in each category, which show that p_{Risky} was below chance in the “lower risky value and one-observed-risky-outcome” category (leftmost bars in both plots; both $p < .001$). As in the TPT, p_{Risky} was above chance level in all other categories (all $p < .036$).

Does the EV difference affect the chance of identifying the objectively maximizing option?

Figure 8 shows the chance of correctly identifying the objectively maximizing option, indicated by the proportion

Table 4. The four categories formed by observed variability and observed value in the risky option, together with the total number (proportion) of cases falling into each of the categories. Separately for the 298 cases with small EV difference and the 290 cases with large EV difference

Number of observed outcomes in the risky option		Observed value of the risky compared with the safe option		
		Lower (.50)	Higher (.50)	
Small EV difference	One (.70)	100 cases (.34)	108 cases (.36)	
	Two (.30)	47 cases (.16)	43 cases (.14)	
	<i>Large EV difference</i>		Lower (.52)	Higher (.48)
	One (.68)	102 cases (.35)	95 cases (.33)	
Two (.32)	50 cases (.17)	43 cases (.15)		




Figure 6. Median (± 1 SE) sample size, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option), for problems with small (left) and large (right) EV differences




Figure 7. Mean (± 1 SE) proportion of samples from the risky option, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option), for problems with small (left) and large (right) EV differences. The vertical line indicates chance level




Figure 8. Proportion (± 1 SE for proportions) of cases in which the observed values point in the same direction as the objective values, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option), for problems with small (left) and large (right) EV differences. The dashed lines indicate chance level

of cases where the objective and observed values pointed in the same direction. As expected, correspondence was significantly less for problems with small EV difference (.50) than for problems with large EV difference (.64; $\chi^2 = 10.32$,

$p = .001$). Additional χ^2 -tests against chance level in each category illustrate the combined importance of EV difference and observed payoff. Correspondence was above chance in problems with large EV difference when both risky outcomes

were observed (light gray bars in the right plot; both $p < .001$). Correspondence was at chance level in all other categories (all $p > .15$).

Does the EV difference affect choice?

Figure 9 shows participants' preferences during choice, indicated by the proportion of choices for the risky option. Proportions did not differ between problems with small (.53) and large EV differences (.51; $\chi^2 < 1$). Instead, choices depended on the values that were *observed* during sampling. As in the TPT, participants chose the risky option less often after observing lower risky values than after observing higher risky values, $\chi^2 = 74.01$, $p < .001$. The extent of this reliance on observed values was again reduced by observed variability. As confirmed by χ^2 -tests against chance in each category, proportions deviated from chance when only one risky outcome was observed (dark gray bars; all $p < .001$). They did not when both risky outcomes were observed (light gray bars; all $p > .22$).

The results replicate the interplay between observed payoff variability, the validity of observed values, and participants' reliance on their experience as found in the TPT dataset: Although observed variability *increased* the correspondence between observed and objective values, it *decreased* participants' reliance on those experiences during choice. This apparent contradiction might have a simple explanation. Observing both risky outcomes could reduce the average difference between the options' observed values, thereby making it more difficult to distinguish between them. To test this assumption, we calculated the absolute difference between the options' observed values for each participant and problem. We find that the absolute observed difference is indeed considerably smaller if both risky outcomes were observed ($mean = 10.0$, $SD = 11.8$), than if only one risky outcome was observed ($mean = 21.0$, $SD = 21.5$; $t(564) = 7.98$, $p < .001$).

The analysis of observed values also shows why participants' search and choice behavior was not affected by the size of the EV difference. The absolute difference between *observed* values did not differ between problems with a small

EV difference ($mean = 17.5$, $SD = 19.7$) and problems with a large EV difference ($mean = 17.6$, $SD = 19.7$, $t(586) < 1$), showing that participants' average experiences did not differ between problems with small or large EV differences. For small and large objective EV differences, the absolute *observed* difference was significantly higher than the absolute *objective* difference (t -test against .5 for small EV differences: $t(297) = 14.89$, $p < .001$; t -test against 5.5 for large EV differences: $t(289) = 10.47$, $p < .001$). This amplification of objective differences in participants' experiences is caused by the relatively small sample sizes (see Hertwig & Pleskac, 2010, for a statistical proof of this effect).

Discussion

The results of this experiment replicate the effects from the TPT dataset and extend our conclusions in several important ways. First, they further clarify the role of variability and value on peoples' preferences during search. In the TPT data, we found a general tendency toward more search in the risky option, even if no payoff variability was observed. Our analyses suggested that this tendency results from participants' learning to anticipate payoff variability in the risky option based on observed outcomes. In the current experiment, participants solved too few problems to learn how to anticipate payoff variability. Consequently, and as expected in cases without observed payoff variability, the general preference for the risky option was replaced by a preference for the option with higher experienced values.

Second, we illustrated one of the possible reasons behind the payoff variability effect (Busemeyer & Townsend, 1993; Erev & Barron, 2005; Myers & Sadler, 1960): After observing variable payoffs, differences between the options' values were considerably smaller than after observing stable payoffs. Smaller observed differences made it more difficult to distinguish between the options and thereby might have reduced reliance on observed values. This interpretation is consistent with the results from Hertwig and Pleskac (2010), which showed that participants reported increasing difficulty choosing between two options as the observed difference between those options decreased.




Figure 9. Proportion (± 1 SE for proportions) of final choices for the risky option, separately for observed variability (one versus two risky outcomes) and value (lower versus higher values in the risky compared with the safe option), for problems with small (left) and large (right) EV differences. The dashed lines indicate chance level

Third, we clarified the effects of the objective difference between the options' values. Although larger objective EV differences increased the chance to correctly identify the objectively maximizing option, they did not affect how much information participants searched for or which options they preferred during search and choice. A reason behind this lack of effects on participants' behavior is suggested by our analysis of the observed difference between the options. The absolute difference between *observed* values did not differ between problems with small or large objective EV differences, showing that participants' experiences did not reflect the true difference between objective values well.

CONCLUSIONS AND NEW CHALLENGES

To summarize, our results contribute to a growing literature on decisions from experience, by showing the respective contributions of two key ecological properties to information search and choice: the value and variability of observed payoffs. Overall, participants preferred the option with higher *observed values* during search as well as during choice. Observed values were not related to the total amount of searched information or to the chance of correctly identifying the objectively maximizing option. *Observed variability* was positively related to participants' efforts and preferences during search. Contrary to a recent speculation (Lejarraga et al., 2012), our results suggest that this relationship is largely driven by the increased chance to observe variability in larger samples, rather than by increased search following observed variability. Despite resulting in an increased chance to identify the objectively maximizing option, observed variability decreased participants' reliance on observed values when choosing between the options. Finally, we find that a larger *objective difference* between the options increased the chance to identify the objectively maximizing option, but it did not affect participants' search or choice behavior, because it was not well reflected in the information obtained during sampling.

Often, peoples' performance in experiential choice is evaluated with respect to a normative underlying payoff structure, which is unknown to the decision maker. Our results illustrate why this approach can be misleading. People make decisions based on their own individual experiences, which can strongly deviate from objective definitions of the decision problems. In more than half of all analyzed cases, participants did not observe the payoff variability in the risky option, thereby changing the nature of the choice task into a choice between two seemingly safe options. In the remaining cases, both risky outcomes were observed; however, even in those cases, observed values did not always allow for identifying the objectively maximizing option. As has been argued by others (e.g., Fiedler, 2000; Hadar & Fox, 2009; Hau et al., 2010; Rakow et al., 2008), this discrepancy between the objective and observed properties of a decision problem is likely an important contributor to decision biases that have been proposed in the literature.

Many real-life search and choice situations are a lot more complex than the simplified sampling paradigm used here. One source of added complexity is the variability of observed

payoffs. In the current paper, variability was defined in the simplest possible way, by comparing options with one versus two observed outcomes. In real life, payoff variability is related not only to the number of outcomes in an option but also to their values, the similarity between the outcomes, and their relative frequencies (cf. Hau et al., 2010; Leland, 2012; Tsetsos et al., 2012). The respective contribution of those different aspects of variability to search and choice are not yet well understood and their investigation will be a fruitful topic of future research. As demonstrated by our results, such an investigation will not only need to consider how the objective characteristics of those aspects are reflected in decision-makers' actual experiences but also to what extent behavior is due to the anticipation, rather than the experience of those characteristic.

Also the relationship between experienced and objective information can be more complex than in the sampling paradigm. The only way to gather information in the sampling paradigm is through experience, by sampling from the options' outcomes, while in many real-life situations, personal experiences can be supplemented by more objective external sources of information. For example, when buying a car, we would not only use previous experiences with different types of cars but also be likely to consult some kind of (hopefully) objective information. How are those different sources of information integrated, especially when subjective and objective information diverge? Evidence suggests that even in the presence of objective descriptions, subjective experiences still have a major impact at the behavioral (Lejarraga & Gonzalez, 2011) and even at the neuronal level (Walsh & Anderson, 2011).

Finally, the separation of a search phase, where decision makers obtain information, and a choice phase, where decision makers obtain rewards, is less straightforward in real life than in the sampling paradigm. Consider again the example of mate selection. Although the "ultimate reward" might be to marry someone for life, the experiences leading to this choice may be rewarding (and consequential) as well. This fuzzy boundary between search and choice is closely related to the "exploration-exploitation trade-off," which has been described as one of the major challenges of decision research (Cohen et al., 2007), and it might be responsible for some of the identified similarities between people's behavior during search and choice (see also Gonzalez & Dutt, 2011).

ACKNOWLEDGEMENTS

This research is supported by the National Science Foundation Award number: 1154012 to Cleotilde Gonzalez. The authors would like to thank Hau-yu Wong and Jason Harman for providing insightful comments and Ido Erev for providing the data that was reanalyzed here.

REFERENCES

- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. DOI: 10.1037/0033-295X.113.4.700

- Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica*, 45, 223–241. DOI: 10.1016/0001-6918(80)90034-7
- Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(3), 538–564. DOI: 10.1037/0278-7393.11.3.538
- Busemeyer, J., & Townsend, J. (1993). Decision field theory: A dynamic cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432–459. DOI: 10.1037/0033-295X.100.3.432
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942. DOI: 10.1037/0033-295X.111.4.939
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5), 523–538. DOI: 10.1287/orsc.12.5.523.10092
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112(4), 912–931. DOI: 10.1037/0033-295X.112.4.912
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., et al. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47. DOI: 10.1002/bdm.683
- Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, 21(5), 575–597. DOI: 10.1002/bdm.602
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, 107(4), 659–676. DOI: 10.1037/0033-295X.107A659
- Fiedler, K., & Kareev, Y. (2006). Does decision quality (always) increase with the size of information samples? Some vicissitudes in applying the law of large numbers. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4), 883–903. DOI: 10.1037/0278-7393.32.4.883
- Fox, C. R., & Hadar, L. (2006). “Decisions from experience” = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, 1(2), 159–161.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indexes. *Journal of the Royal Statistical Society Series B-Methodological*, 41, 148–177. DOI: 10.1.1.295.4422
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118(4), 523–551. DOI: 10.1037/a0024558
- Gonzalez, C., & Dutt, V. (2012). Refuting data aggregation arguments and how the instance-based learning model stands criticism: A reply to Hills and Hertwig (2012). *Psychological Review*, 4(119), 893–898. DOI: 10.1037/a0024558
- Gonzalez, C., & Dutt, V. (2013). Manuscript submitted for publication. Concurrence of Exploration and Maximization in Experiential Decisions from Sampling.
- Hadar, L., & Fox, C. R. (2009). Information asymmetry in decision from description versus decision from experience. *Judgment and Decision Making*, 4(4), 317–325.
- Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making*, 23(1), 48–68. DOI: 10.1002/bdm.665
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21(5), 493–518. DOI: 10.1002/bdm.598
- Hertwig, R. (in press). Decisions from experience. In G. Keren & G. Wu (Eds.), *Blackwell handbook of decision making*. Oxford, UK: Blackwell.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539. DOI: 10.1111/j.0956-7976.2004.00715.x
- Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 209–236). Oxford: Oxford University Press.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, 115(2), 225–237. DOI: 10.1016/j.cognition.2009.12.009
- Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, 21(12), 1787–1792. DOI: 10.1177/0956797610387443
- Hills, T. T., & Hertwig, R. (2012). Two distinct exploratory behaviors in decisions from experience: Comment on Gonzalez and Dutt (2011). *Psychological Review*, 119(4), 888–892. DOI: 10.1037/a0028004
- Johnson, T. R., Budescu, D. V., & Wallsten, T. S. (2001). Averaging probability judgments: Monte Carlo analyses of asymptotic diagnostic value. *Journal of Behavioral Decision Making*, 14(2), 123–140. DOI: 10.1002/bdm.369
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292. DOI: 10.2307/1914185
- Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review*, 107(2), 397–402. DOI: 10.1037/TO33-295X.107.2.397
- Kareev, Y. (2005). And yet the small-sample effect does hold: Reply to Juslin and Olsson (2005) and Anderson, Doherty, Berg, and Friedrich (2005). *Psychological Review*, 112(1), 280–285. DOI: 10.1037/0033-295X.112.1.280
- Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, 116(2), 286–295. DOI: 10.1016/j.obhdp.2011.05.001
- Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, 124(3), 334–342. DOI: 10.1016/j.cognition.2012.06.002
- Leland, J. W. (2012). Equilibrium selection, similarity judgments, and the “nothing to gain/nothing to lose” effect. *Journal of Behavioral Decision Making*, online first. DOI: 10.1002/bdm.1772
- Myers, J. L., & Sadler, E. (1960). Effects of range of payoffs as a variable in risk taking. *Journal of Experimental Psychology*, 60(5), 306–309. DOI: 10.1037/h0042499
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 1–46. DOI: 10.1037/1089-2680.2.2.175
- Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes*, 106(2), 168–179. DOI: 10.1016/j.obhdp.2008.02.001
- Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, 23(1), 1–14. DOI: 10.1002/bdm.681
- Tsetsos, K., Chater, N., & Usher, M. (2012). Saliency driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109(24), 9659–9664. DOI: 10.1073/pnas.1119569109
- Vul, E., Goodman, N. D., Griffiths, T. L., & Tenenbaum, J. B. (2009). One and done? Optimal decisions from very few samples. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the cognitive science society* (pp. 66–72). Austin, TX: Cognitive Science Society.
- Wallsten, T. S., Budescu, D. V., Erev, I., & Diederich, A. (1997). Evaluating and combining subjective probability estimates. *Journal of Behavioral Decision Making*, 10(3), 243–268. DOI: 10.1002/bdm

10.1002/(SICI)1099-0771(199709)10:3<243::AID-BDM268>3.0.CO;2-M

Walsh, M. M., & Anderson, J. R. (2011). Learning from delayed feedback: Neural responses in temporal credit assignment. *Cognitive, Affective, & Behavioral Neuroscience*, *11*, 131–143. DOI: 10.3758/s13415-011-0027-0

Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, *12*(3), 129–140. DOI: 10.1080/17470216008416717

Authors' biographies:

Katja Mehlhorn received her PhD in Cognitive Science from the University of Groningen, the Netherlands and is a postdoctoral research fellow at the Dynamic Decision Making Lab at Carnegie Mellon University.

Noam Ben-Asher received his PhD in Human Factors Engineering from Ben-Gurion University in Beer Sheva, Israel and is a postdoctoral research fellow at the Dynamic Decision Making Lab at Carnegie Mellon University.

Varun Dutt is an Assistant Professor, School of Computing and Electrical Engineering and School of Humanities and Social Sciences, Indian Institute of Technology, Mandi, India. He is also the Knowledge Editor of the English daily, Financial Chronicle.

His current research interests are in environmental decision making and modeling human behavior.

Cleotilde (Coty) Gonzalez received her PhD in Management Information Systems from Texas Tech University. She is an associate research professor in the department of Social and Decision Sciences and the director of the Dynamic Decision Making Lab at Carnegie Mellon University.

Authors' addresses:

Katja Mehlhorn Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA

Noam Ben-Asher Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA

Varun Dutt School of Computing and Electrical Engineering, School of Humanities and Social Sciences, Indian Institute of Technology, Mandi, India

Cleotilde Gonzalez Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, USA