Don't Stop 'Til You Get Enough: Adaptive Information Sampling in a Visuomotor Estimation Task

Mordechai Z. Juni[1], Todd M. Gureckis[1], Laurence T. Maloney[1,2]
[1]Department of Psychology, NYU [2]Center for Neural Science, NYU
6 Washington Place, New York, NY 10003 USA
{mjuni, todd.gureckis, laurence.maloney}@nyu.edu

Abstract

We investigated how subjects sample information in order to improve performance in a visuomotor estimation task. Subjects were rewarded for touching a hidden circular target based on visual cues to the target's location. The cues were 'dots' drawn from a Gaussian distribution centered on the middle of the target. Subjects could sample as many cues as they wished, but the potential reward for hitting the target decreased by a fixed amount for each additional cue requested. The subjects' objective was to balance the benefits of increased information against the costs incurred in acquiring it. We compared human performance to ideal and found that subjects sampled more cues than dictated by the optimal stopping rule that tries to maximize expected gain. We contrast our results with recent reports in the literature that subjects typically under-sample.

Keywords: decision making, information sampling, optimal stopping, adaptive cue-combination, value of information.

Introduction

A critical challenge facing human decision makers is balancing the potential advantage gained by gathering (i.e., sampling) information against the time, energy, or money spent collecting it. For example, Stigler (1961) analyzed the economic costs of prolonging a search for a better price on a commodity. Since a relatively cheap price is easily obtained after a brief search, the cost of exhaustive search will often not offset the increased savings of finding the cheapest price. Instead, consumers should search only so long as the expected savings from finding a cheaper price are enough to offset the costs in continuing to search. Similar ideas are reflected in the mathematical literature on optimal stopping and optimal search (Wald, 1945a,b; Arrow et al., 1949; Stone, 1989), and feature prominently in the study of animal foraging (Stephens & Krebs, 1986). Less is known, however, about how effective humans are in trading off the costs and benefits of additional information, or how their performance varies across decision environments.

Optimal information sampling behavior was a topic of interest in psychology in the 1960s (Green et al., 1964; Edwards, 1965; Tversky & Edwards, 1966; Wendt, 1969; Rapoport & Tversky, 1970), and this question has returned to prominence recently (Hertwig et al., 2004; Hau et al., 2008; Gureckis & Markant, 2009; Vul et al., 2009; Hertwig & Pleskac, 2010). The critical issue in all this past research has been a comparison of human sampling behavior to that of an optimal decision maker. A common finding in the more recent studies is that participants often collect *less* information than needed in order to maximize expected gain (i.e., they "under-sample").

However, a key limitation of recent work in this area (e.g., Hertwig et al., 2004) is that the cost of additional information was not precisely specified and, as a consequence, it was difficult to determine the optimal decision strategy for the task at hand.

The goal of the present study was to examine information-sampling behavior in a simple visuomotor task where the cost of additional information was made explicit, and the optimal decision strategy was amenable to mathematical analysis. The subject had to estimate the location of an invisible target on a monitor, and touch it to earn rewards. Participants sampled cues that provided information about the location of the hidden target. Each cue was sampled from a bivariate Gaussian distribution centered on the target and, the more cues subjects had, the better their localization (Tassinari et al., 2006). However, each additional cue reduced the potential reward for hitting the target by a fixed amount. Participants had to balance the benefits of additional information (more cues) against the costs required to collect it. In the analysis below, we analyze our task and show how to compute the optimal number of samples to request in order to maximize expected gain. We then compare ideal performance to the performance of human subjects.

Our experiment departs from previous work on information sampling in three key ways. First, unlike a number of recent analyses (e.g., Hertwig et al., 2004; Hau et al., 2008), we made the costs of collecting information explicit. Second, our novel decision making task involved accumulating evidence to guide a single, continuous, reaching movement (as opposed to, for example, making a one-off decision between multiple, discrete choice options). Finally, in our task, all the sampled cues were simultaneously present on the screen, limiting the secondary task demands placed on participants (e.g., keeping recent samples in memory). Prior work with similar tasks has shown that subjects are close to optimal in their ability to integrate such cues to guide action (Battaglia & Schrater, 2007).

In contrast to the recent findings and emphasis on under-sampling in the decision making literature (e.g., Hertwig & Pleskac, 2010; Vul et al., 2009), we find that people systematically *over-sample* information. In our analysis, we rule out a number of possible explanations for why this might be the case. Ultimately, our results appear consistent with a type of risk-aversion wherein participants are biased against uncertain outcomes (Tversky & Kahneman, 1981).

Prior work on adaptive information sampling

In a classic study, Tversky & Edwards (1966) had people perform a simple probability learning task requiring them to guess which of two lights would light up on each of 1000 trials. Correct guesses were rewarded and incorrect guesses punished, but the observer did not receive immediate feedback. To learn about the event probabilities, participants were given the option to, at any point, forgo guessing on a trial and observe the outcome instead. Optimal performance in the task entails observing a certain number of trials at the start of the experiment to learn the relative probabilities of each event, and then selecting on the rest of the trials the more frequent of the two events (Wald, 1947). However, participants in this study greatly oversampled (preferring on average around 300 observation trials compared to the optimal strategy of sampling around 30 trials). One likely explanation is that the participants mistakenly thought that the underlying reward probabilities were non-stationary (changing across time), and would therefore return intermittently to observe more outcomes to track changes in the relative probabilities.

In contrast to Tversky & Edwards (1966), recent work on adaptive information sampling has focused on tasks were participants are forced to first sample information from various alternatives, only to utilize that information in a later decision phase (cf. Hertwig et al., 2004; Weber, Shafir, Blais, 2004; Hau et al., 2008; Gureckis & Markant, 2009). The key dependent measure of interest in such studies is how much information people collect before stopping and making a decision. A striking finding in this literature is how *little* information people collect before making a decision. For example, in Hau et al. (2008) subjects were presented with two decks of cards and given as much time as they wanted to sample freely the payout distribution of each deck before making a final decision as to which deck to choose from to be rewarded. Interestingly, subjects sampled a median of around 11 cards across the two decks before making their final decision. This very low level of search was often insufficient to accurately assess the expected gain of the choices (see also Rapoport & Tversky, 1970). A similar preference for less rather than more information is observed in naturalistic choice scenarios as well (Todd, 2007). For example, in high stakes choices such as marriage, it has long befuddled social economists that individuals report dating relatively few people prior to marriage (Miller & Todd, 1998).

The frequent reports of under-sampling has led to a number of recent papers arguing that collecting small samples may actually be advantageous in certain decision environments. For example, Hertwig & Pleskac (2010) present an analysis showing how the limited samples participants take in certain tasks may actually help to amplify small differences between payoff functions and may enable relatively effective choice behavior in particular environments (meeting the criterion of "satisficing" laid out by Simon, 1956). However, one limitation of these analyses is that, due to the complexity of the empirical tasks, it is mathematically challenging to define the optimal rule for terminating information-sampling. In addition, access to information is often free (in terms of money) but entails unspecified costs associated with time or effort. In the present study we attempt to quantify and control both of these variables to help us better define the normative standard against which to judge human performance.

Overview and model of the present experiment

In the current experiment, subjects performed a task similar to those used by Battaglia & Schrater (2007) and Tassinari et al. (2006). The goal on each trial was to touch a hidden circular target on a touch-screen, akin to throwing a dart at a dartboard. The target's location changed from trial to trial and was cued by dots drawn from a Gaussian distribution that was centered on the middle of the target. To increase the probability of hitting the target, subjects were given the option to sample dots one at a time at a set cost-per-dot. With increased dots sampled, subjects had a greater chance of hitting the target, because the variance of the sampling distribution decreases with increased sample size. Indeed, previous work with such tasks has shown that subjects correctly interpret the arrival of more dots as reducing uncertainty about the target's location (Battaglia & Schrater, 2007; Tassinari et al., 2006). But in the present task, the expected benefit of more information comes at the cost of reducing the points awarded for hitting the target. Do subjects know when to stop sampling and plan an action?

How much information is enough? Specifying optimal sampling behavior

In order to analyze behavior in the task, we begin by defining the behavior of an optimal subject who samples information with the goal of maximizing expected gain. The variance of the underlying Gaussian distribution from which the cues were sampled, denoted σ_S^2, remained constant throughout the experiment. To maximize expected gain, optimal subjects must minimize their estimation variance by using the mean (center of gravity) of the sample as the estimate of the target's location. The variance of the sample mean, denoted σ_S^2/n, depends critically on sample size, denoted n. Furthermore, the optimal subject must take into account other sources of variability such as one's ability to precisely specify a location using the experimental apparatus. We refer to the aggregate of all the other sources of variability as "adjustment" variability and denote it σ_A^2. These two sources of variance give rise to the subject's total experimental estimation variance:

$$\sigma^2(n) = \sigma_S^2/n + \sigma_A^2. \tag{1}$$

The amount of adjustment variability (σ_A^2) is a latent parameter to our model that we could estimate empirically.

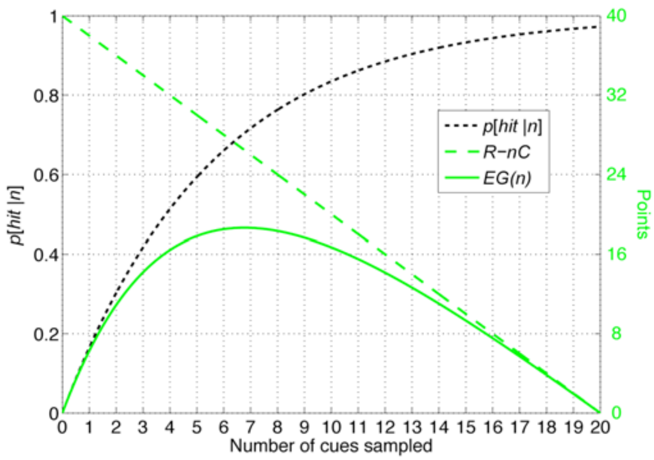

Figure 1: The expected gain function *EG(n)* is the product of the probability of hitting the target with *n* samples, and the gain earned for hitting the target (which takes into account the costs of the samples). Probability is plotted on the left vertical scale, gain and expected gain on the right.

In defining our optimal decision maker, we assume that adjustment variability was zero for simplicity.

Given $\sigma^2(n)$, we can compute the probability of hitting the circular target as a function of sample size as follows:

$$p\left[hit \mid n\right] = \iint_T \phi\left(\vec{0}, \Sigma(n)\right)\, dx\, dy\,, \qquad (2)$$

where the region of integration T is the invisible circular target, and $\phi\left(\vec{0}, \Sigma(n)\right)$ denotes the probability density function of a bivariate Gaussian distribution centered on $\vec{0} = (0,0)$ with covariance

$$\Sigma(n) = \begin{bmatrix} \sigma^2(n) & 0 \\ 0 & \sigma^2(n) \end{bmatrix}. \qquad (3)$$

Given $p[hit \mid n]$, we can compute the expected gain as a function of sample size as follows:

$$EG(n) = p\left[hit \mid n\right]\left(R - nC\right), \qquad (4)$$

where R is the initial point value of the target, and C is the fixed cost to the target value that is incurred for each cue that is sampled. (Note that in the current experiment subjects were never penalized for missing the target, and so both R and nC must be mediated by $p[hit \mid n]$. Subjects were allowed to sample additional cues only so long as $R - nC > 0$. All the figures will be cut off where $EG(n)$ starts to dip below zero.)

Figure 1 shows an example of the expected gain function $EG(n)$ elicited in Equation 4, using the actual $p[hit \mid n]$ of the current experiment (assuming $\sigma_A^2 = 0$), and the potential gain $(R - nC)$ used in one condition of the current experiment (the "low stakes, low cost" condition).

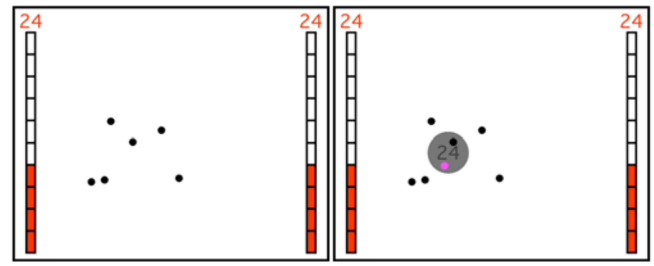

Figure 2: Schematic display of the "high risk, high cost" condition. Left panel: Subject has sampled 6 cues (i.e., dots) and thus the potential reward for hitting the hidden grey target has been reduced from 60 points to 24 points. Right panel: The purple response dot shows the subject's response (made after touching the display and making fine adjustments) and the revealed target. Since the purple dot is within the target, the trial would be coded as a hit. The reward for this trial was 24 points. Had the subject's final setting been outside the target, the center of the target would show 0 (no reward).

Note that the experimental design gives rise to a single-peaked expected gain function (cf. Hertwig & Pleskac, 2010, Fig. 3). The ideal subject would continue sampling on each trial until $EG(n)$ peaks and then attempt to hit the target using the mean of the sample as their estimate. This strategy, irrespective of the adjustment variability σ_A^2, will always lead to the maximum expected gain. In the following experiment, we computed the $EG(n)$ curve and its maximum for two different decision environments and compared this normative standard to the number of samples taken by participants in the task.

Experiment

The cover story for the estimation task was a simple game where the goal was to collect points by hitting an invisible dartboard (using one's finger instead of darts). To gain information as to the location of the dartboard, subjects could only observe the end-points of darts thrown by another shooter who could see the target and who would be aiming for the center of the target. The subject could use the outcomes of the other shooter's attempts as a guide to the location of the target.

Each subject alternated between blocks of two different conditions: "low stakes, low cost" and "high stakes, high cost". In the first condition, the initial reward for hitting the target was 40 points, and decreased by 2 points per cue. In the second condition the initial reward for hitting the target was 60 points, and decreased by 6 points per cue.

For illustration, in Figure 2 (left panel), the subject has sampled six cues on a high stakes trial and must decide whether to sample a seventh cue or attempt to hit the target based on the sample of size $n = 6$. In the right panel, the subject has successfully hit the target (the purple dot was visible to the participant and represented the response). The reward for that trial is: 60 points - (6 cues x 6 points) = 24 points. If the subject had missed the target then there would have been no reward or penalty for that trial.

Subjects were obligated to sample at least one cue per trial, and they were allowed to continue sampling one cue at a time so long as the value of the potential reward $(R - nC)$ would not be reduced to zero. Thus, they were limited to sampling 9 cues in the high stakes condition and 19 cues in the low stakes condition.

These two conditions and the properties of the stimuli were carefully chosen so that the expected gain of the ideal subject (with $\sigma_A^2 = 0$, and who stops sampling optimally) would be approximately the same for the two conditions (18.55 points per trial). The main difference between conditions was in their expected gain functions. The high stakes condition (red) had very steep curvature at its maximum and peaked at 4.05 cues, while the low stakes condition (green) had shallower curvature at its maximum and peaked at 6.77 cues (see Figure 3).

Subjects Eight subjects at New York University participated in the experiment. None were aware of the purpose of the experiment and each was paid $10 per hour for their participation, plus a potential monetary bonus.

Apparatus Stimuli were displayed on a vertically mounted 338 mm by 270 mm touch-screen LCD in a dimly lit room. The monitor was set at a resolution of 1280 pixels by 1024 pixels (1 pixel = 0.26387 mm) with a 60-Hz refresh rate. Subjects were asked to seat themselves at a comfortable distance and adjust the height of the chair so that they could perform the experimental task with ease. The experiment was programmed and run using MATLAB and the Psychtoolbox libraries (Brainard, 1997; Pelli, 1997).

Stimuli The hidden target was a grey circle with radius = 12.67 mm. The cues were small white dots with radius = 1.056 mm. The cues were drawn from a Gaussian distribution (SD = 21.11 mm) that was centered on the target[1].

At the far ends of the screen there were two vertical reward bars that decreased in height with each additional cue sampled. Additionally, there was a number at the top of each bar indicating how many points would be awarded for hitting the target. To help subjects remember which condition they were in, green and red bars indicated the "low stakes, low cost" and "high stakes, high cost" condition, respectively.

Design Each subject ran in a practice session followed by the experimental session. The practice session consisted of 30 trials for each condition in alternating blocks of 10 trials, for a total of six practice blocks and 60 practice trials. The actual experiment consisted of 100 trials for each condition in alternating blocks of 25 trials, for a total of 8 experimental blocks and 200 experimental trials. The ordering of the conditions was randomly assigned and counterbalanced between subjects. However, the ordering was kept constant between the practice session and the experimental session. Hence, a subject assigned to start with the high stakes condition would start both the practice and the experimental sessions with a high stakes block. Since the task was self-paced, subjects' participation time (including practice) ranged from 41 min to 74 min with an average of 62 min.

Monetary bonus Points earned in the task were converted into bonus money at a rate of five cents per point. This means that the maximum potential reward for hitting the target was 38 x $0.05 = $1.90 per low stakes trial, and 54 x $0.05 = $2.70 per high stakes trial. In order to maintain motivation throughout the task, subjects were informed that they would receive a monetary bonus on 5% of the trials, by randomly choosing five trials from each condition at the end of the experiment. The total expected monetary bonus of the most optimal ideal subject was $9.28.

Procedure Subjects were asked to use their dominant hand throughout the experiment. The experimenter stayed in the room during the practice blocks to explain the display and encourage subjects to use the practice trials to explore and observe the outcome of different decision strategies.

At the start of each block, subjects were shown an instruction screen providing explicit information as to the initial point value of the target (40 points or 60 points) and the cost per sample (2 points or 6 points). In addition, each condition ("low stakes, low cost" or "high stakes, high cost") was associated with a particular color scheme for the display elements (green or red). Lastly, subjects were required to confirm the appropriate cost per sample by pressing "2" or "6" before the block would begin.

At the start of each trial, the screen was completely black except for the colored reward bars and numbers at the far ends of the screen indicating the current reward for hitting the target. To sample a cue, subjects pressed the space bar causing a white dot to appear on the screen. Concurrent with the appearance of the cue, the reward bars and numbers would decrease to indicate the lower reward available for hitting the target. If subjects wanted more samples, they simply hit the space bar repeatedly until they were ready to reach for the hidden target.

To hit the invisible target, subjects simply touched the screen with their finger causing a small purple response dot to appear. Subjects were allowed to adjust their response by moving the purple dot with their finger or by pressing the arrow keys. Once satisfied with their response, subjects pressed the space bar to receive feedback. During the feedback phase of each trial, the hidden target would appear in grey along with all the white cues sampled on that trial.

[1] These properties entail that approximately 16.5% of the cues landed directly on the target. So if an extreme risk taker would choose to sample only one cue on each trial and use that as the estimate of the location of the target, then they would hit the target and get the maximum reward only 16.5% of the time—a strategy that in the current experiment is far from optimal.

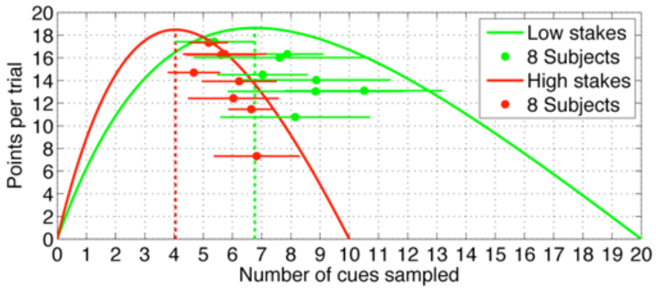

Figure 3: Results: Points per trial earned by subjects is plotted versus the number of samples taken (error bars mark the standard deviation of the number of samples taken by each subject) for the "low stakes, low cost" (green) and "high stakes, high cost" (red) conditions. The expected gain curve for each condition is plotted in the corresponding color, with corresponding dashed lines marking the number of samples needed to maximize expected gain. For this figure, we plot the corresponding expected gain curves for $\sigma_A^2 = 0$ (see Figure 4 and Discussion).

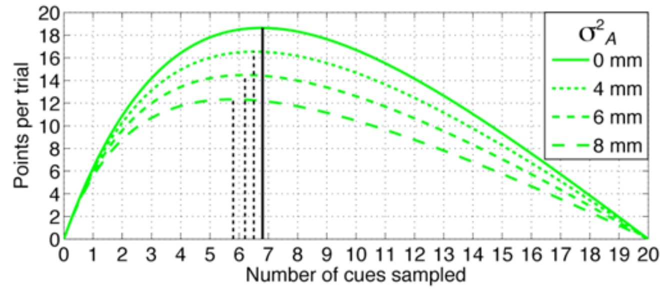

Figure 4: The expected gain function *EG(n)* changes with increased adjustment variability σ_A^2. Notice that as adjustment variability increases, one would need to sample fewer and fewer cues to have the maximum expected gain (indicated by the black solid and dashed lines for different degrees of adjustment variability).

If the small purple response dot was within the target, the trial was counted as a hit and the number of points awarded for that trial would appear at the center of the target. If, however, the small purple response dot was not within the target, then the trial was coded as a miss and a 0 appeared at the center of the target. A second space bar press began the next trial.

Finally, at the end of both the practice session and the experimental session, subjects were given feedback as to how they performed on each of 10 randomly selected trials and how much bonus money they received as a result. This feedback at the end of the practice session (which was not actually paid out) served to give subjects a rough sense of the range of actual bonuses possible.

Results

On average, subjects collected 14.06 points per trial (SD=2.19), which is 76% of the maximum possible expected gain of 18.55 points per trial.

Figure 3 shows each subject's mean and standard deviation of the number of cues sampled for each of the two conditions. For illustration, we place individual data points at heights that correspond to their average gain per trial. For the same sampling behavior, some subjects where better able to successfully hit the target and collect more points. This is primarily due to each individual's adjustment variability σ_A^2. (However, see Discussion for an explanation of why this does not affect our results.)

All subjects correctly sampled more cues in the low stakes condition (M=8.04, SD=1.5) than the high stakes condition (M=5.87, SD=0.72), $t(7)=4.54$, $p<.003$. The solid curves show the respective low stakes and high stakes expected gain functions for the perfect, ideal subject whose adjustment variability $\sigma_A^2 = 0$. The expected gain for this ideal sampler is maximized when sampling 6.77 cues in the low stakes condition and 4.05 cues in the high stakes condition.

Statistical tests back up the intuition shown in Figure 3 that almost all subjects were risk-averse and sampled more cues than dictated by the optimal stopping rule that tries to maximize expected gain. The exceptions were one subject who was risk-seeking and under-sampled in the low stakes condition (M=5.38, SD=1.36), $t(99)=-11.9$, $p<.001$, and one other subject who was not significantly different from the optimal stopping rule in the low stakes condition (M=7.03, SD=1.55), $t(99)=0.19$, $p>.05$.

Discussion

Our discussion takes the form of a set of questions and explores possible alternative explanations of our results.

Question 1: In our model, we assumed that adjustment variability σ_A^2 was 0. Could the observed oversampling be due to subjects' adjustment variability?

Figure 4 shows how the expected gain function $EG(n)$ of an ideal sampler changes as the adjustment variability increases. Notice first that as σ_A^2 increases, one must sample fewer and fewer cues in order to have the maximum expected gain. Thus, accounting for increased adjustment variance would lead to a decrease in the number of samples taken and could not account for the pattern of oversampling that we found. Different individual settings of σ_A^2 might explain why some subject's data fall right beneath the curves in Figure 3 (due to low σ_A^2), while others fall lower (due to high σ_A^2). Note that none of the subjects' data in Figure 3 is higher than the expected gain curves for $\sigma_A^2 = 0$.

Question 2: Could sub-optimal decision-making reflect an inability of participants to discern small differences around the peak of the utility function (i.e., the "flat maxima phenomena")?

Failure to distinguish regions around the peak should look like an unsystematic tendency to both over- and under-sample. Only one subject under-sampled, and only in one condition. The systematic tendency to over-sample in our experiment indicates that subjects were sensitive to the expected gain differences that accompanied changes in the number of cues sampled.

Question 3: Why might people oversample?

We conjecture that over-sampling is the result of a form of risk-aversion (Tversky & Kahneman, 1981). Risk-averse subjects are willing to pay money to reduce the variability of their rewards. They accept a smaller expected gain per trial, but the variation in gain from trial to trial is reduced as well. In our task, risk-aversion implies that subjects will systematically collect *more* information than is optimal for maximizing expected gain. The additional information offers them a higher probability of hitting the target, but at the cost of a reduced expected reward for doing so.

Acknowledgments

LTM was supported by Grant EY019889 from the National Institutes of Health, TMG was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract D10PC20023. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government. The authors thank Tim Pleskac for a helpful discussion in the development of the paper.

References

Arrow, K., Blackwell, D., & Girshick, M. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica, Journal of the Econometric Society,* 17(3), 213-244.

Battaglia, P.W., & Schrater, P.R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *Journal of Neuroscience*, 27(26), 6984-6994.

Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision,* 10, 443-446.

Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology,* 2, 312-329.

Green, E.P., Halbert, M.H., & Minas, J.S. (1964). An experiment in information buying. *Journal of Advertising Research*, 4, 17-23.

Gureckis, T.M., & Markant, D. (2009). Active learning strategies in a spatial concept learning game. In Taatgen, N., van Rijn, H., Schomaker, L., & Nerbonne, J. (Eds.) *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 3145-3150). Austin, TX: Cognitive Science Society.

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, 493-518.

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

Hertwig, R, & Pleskac, T.J. (2010). Decisions from experience: Why small samples? *Cognition*, 115, 225-237.

Miller, G., & Todd, P.M. (1998). Mate choice turns cognitive. *Trends in Cognitive Science,* 2(5), 190-198.

Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision,* 10(4), 437-442.

Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5, 105-120.

Simon, H. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129-138.

Stephens, D., & Krebs, J. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy,* 69(3), 213-225.

Stone, L. (1989). Theory of optimal search. Military Applications Section. *Operations Research Society of America*. Arlington, VA: ORSA Books.

Tassinari, H., Hudson, T.E., & Landy, M.S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *Journal of Neuroscience*, 26(40), 10154-10163.

Todd, P.M. (2007). How much information do we need? *European Journal of Operational Research,* 117(3), 1317-1332.

Tversky. A., & Edwards. W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, 71(5), 680-683.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453-458.

Vul, E., Goodman, N.D., Griffiths, T.L., & Tenenbaum, J.B. (2009). One and done? Optimal decisions from very few samples. In Taatgen, N., van Rijn, H., Schomaker, L., & Nerbonne, J. (Eds.) *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 66-72). Austin, TX: Cognitive Science Society.

Wald, A. (1945a). Sequential method of sampling for deciding between two courses of action. *Journal of the American Statistical Association*, 277-306.

Wald, A. (1945b). Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics,* 16(2), 117-186.

Wald, A. (1947). *Sequential analysis.* New York: Wiley.

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

Wendt, D. (1969). Value of information in decision. *Journal of Mathematical Psychology*, 6, 430-443.