



Rivals in the dark: How competition influences search in decisions under uncertainty



Nathaniel D. Phillips^{a,*}, Ralph Hertwig^a, Yaakov Kareev^b, Judith Avrahami^b

^a Max Planck Institute for Human Development, Berlin, Germany

^b The Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem, Israel

ARTICLE INFO

Article history:

Received 24 September 2013

Revised 8 June 2014

Accepted 10 June 2014

Keywords:

Decisions under uncertainty

Competition

Information search

Decisions from experience

ABSTRACT

In choices between uncertain options, information search can increase the chances of distinguishing good from bad options. However, many choices are made in the presence of other choosers who may seize the better option while one is still engaged in search. How long do (and should) people search before choosing between uncertain options in the presence of such competition? To address this question, we introduce a new experimental paradigm called the competitive sampling game. We use both simulation and empirical data to compare search and choice between competitive and solitary environments. Simulation results show that minimal search is adaptive when one expects competitors to choose quickly or is uncertain about how long competitors will search. Descriptively, we observe that competition drastically reduces information search prior to choice.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Whether the question is what to eat, where to live, or with whom to mate, decisions are often made under competitive conditions. This holds for species ranging from humans to hermit crabs. Arguably choosier than humans are about their housing, hermit crabs are always on the look-out for new and better shells. Because the abdomen of a hermit crab is extremely vulnerable, hermit crabs need find suitable seashells to protect their vital organs in order to pass their genes on to the next generation. When a solitary crab encounters an empty shell, it thoroughly inspects the potential new home. The crab will meticulously explore the outer surface of the shell looking for holes and weak points. It will then insert its vulnerable

abdomen into the shell opening to see whether the potential new home is a good fit. If the shell passes this thorough inspection, the crab may decide to discard its current shell and exchange it for the new one. However, when a *group* of crabs simultaneously encounters an empty shell, each individual crabs' search process is dramatically truncated. In this competitive situation, the crab nearest to the shell will make a split-second decision on whether or not to take it based on a brief visual inspection alone (Rotjan, Chabot, & Lewis, 2010).

Swap the hermit crab for a human and the shell for a television on a clearance rack, and intuition suggests that human behavior may be similar to that of hermit crabs'. On a slow shopping day, the leisurely shopper can take his time deciding whether or not to buy the television. He can thoroughly examine the television's attributes, look up expert reviews on his smartphone, or take advantage of the wisdom of crowds by soliciting advice from friends on a social networking site. However, on a frantic shopping day like Black Friday, the same shopper is likely to behave very differently. Surrounded by dozens of other eager

* Corresponding author. Address: Max Planck Institute for Human Development, Center for Adaptive Rationality (ARC), Lentzeallee 94, 14195 Berlin, Germany. Tel.: +49 30 824 06 475.

E-mail address: phillips@mpib-berlin.mpg.de (N.D. Phillips).

shoppers, he might spend only a few moments looking at the television before deciding to grab it before someone else does. Why might competition reduce pre-decisional search so dramatically? What costs and benefits do organisms reap by reducing their search efforts in the presence of competition? What factors in choice options and the social environment affect good search rules? In this paper, we seek to provide initial answers to these questions using a new experimental paradigm that we call the *competitive sampling game*.

Organisms rarely have complete and certain information about options before making even the most consequential choices; instead, they must make choices in the darkness of uncertainty. To shed light on the available options, they must learn about those options' possible outcomes and their associated probabilities through an exploratory search process (Real, 1991). Most people go on dates before proposing marriage, vacationers research and compare hotels before deciding where to stay, and hermit crabs inspect new shells before making a move. After a period of exploration, organisms *exploit* an option by making a long-term consequential choice. Exploration and exploitation represent two diametric goals associated with choice, namely, gathering information about options (exploration) versus consuming an option (exploitation) based on current information (Cohen, McClure, & Yu, 2007). Although exploration provides organisms with more information, it can come at costs in the form of money, time, or lost opportunities. There is thus a tradeoff between exploration and exploitation: If you search too little, you might struggle to distinguish good from bad options. If you search too much, you may suffer from excessive search costs.

In solitary choice situations, the exploration–exploitation tradeoff has been extensively studied both theoretically (Brezzi & Lai, 2002; Gittins, 1979; Gittins, 1989) and empirically (Gans, Knox, & Croson, 2007; Groß et al., 2008), mostly in “multi-armed bandit” problems in which individuals attempt to maximize their payoffs from multiple gambles with initially unknown reward distributions. However, previous research on the exploration–exploitation tradeoff has largely ignored a real-world search cost that dramatically changes how organisms behave: the impact of competition during search. Although search affords more information about available options, it also increases the risk that good option(s) will be taken by competitors.

In this article, we research how competition affects pre-decisional exploration from a descriptive as well as a normative perspective. The essence of what we study concerns supply and demand. In a solitary environment, the “supply,” that is, the number of options available to choose from, is stable. It cannot be affected by the actions of others. Hence, a solitary decision maker can engage in extensive exploration, allowing her to carefully separate good from bad options at leisure before making a consequential choice. In contrast, in a competitive environment, “demand” increases and the danger lurks that competitors will claim desirable options, leaving the thoroughly exploring decision maker with an inferior option set to choose from. With the increased tension between exploration

and exploitation driven by competition, decision makers might be best advised to choose as soon as they detect an option that is likely to be good enough. But when does that moment come? Does search under competition indeed become as truncated as the crab's shell search and the shopper's television search suggest and, if so, how good or bad are the resulting choices? To address these questions, we take advantage of an experimental tool that has recently been used to study the process of search in a range of solitary choice situations (Erev & Barron, 2005; Hertwig, Barron, Weber, & Erev, 2004; Weber, Shafir, & Blais, 2004): the sampling paradigm from research on decisions from experience (Hertwig & Erev, 2009). In this paradigm, participants explore options with a priori unknown underlying probability distributions before deciding between them (exploration before exploitation). In the present research, we pit a solitary variant of this paradigm against a novel competitive variant that we call the *competitive sampling game*.

1.1. Decisions from experience

In the sampling paradigm, a solitary player learns about (i.e. explores) options with a priori unknown payoff distributions that differ in value by sampling outcomes for as long as she wishes, without financial cost. When ready, she chooses (i.e. exploits) her preferred option on the basis of her sampling experience. This final choice then results in a real financial consequence, such as a random payment drawn from the option's payoff distribution. Since the information decision-makers gain through sampling reduces uncertainty about options and increases the likelihood of choosing good over bad options, a key measure in the sampling paradigm is how long people search for information before making a choice. Given that sampling has no cost other than time, one might expect solitary choosers to sample extensively, but previous research shows that protracted search is not the norm. Across studies, participants have generally been found to take between 11 and 19 draws, or about 7 ± 2 samples per option before making a final choice between two gambles (for a review, see Hertwig, *in press*). Researchers have proposed several reasons why people do not search extensively in solitary choice: small sample statistics can be quite accurate where differences are large enough to matter (Johnson, Budescu, & Wallsten, 2001), frugal search reduces choice difficulty (Hertwig & Pleskac, 2010), short-term maximization goals prompt limited search (Wulff, Hills, & Hertwig, 2014), short-term memory constrains information use, and opportunity costs mount as search continues (Hertwig, *in press*).

1.2. The Competitive Sampling Game (CSG)

In this paper we introduce a competitive variant of the sampling paradigm called the competitive sampling game. In the game, players choose between two options realized as urns on the computer screen. Each urn contains 100 virtual balls, with each ball bearing a number. The distribution of numbers in an urn dictates its value. Before making a final consequential choice, players have the

opportunity to learn about the distribution of numbers in each urn by drawing random balls (with replacement), one at a time, from either urn as often as and in any order they wish at no financial cost. When a player decides to stop sampling and chooses an urn, she receives the expected value of the distribution of numbers in her chosen urn. In the *solitary* condition, players play alone, as in the sampling paradigm (see Hertwig et al., 2004). In the *competition* condition, they play in pairs. Each player samples independently but at the same rate, meaning that all players see the same number of samples. As long as both players wish to continue sampling, they both do so. As soon as one or more players decide to stop searching and choose an option, all sampling stops and the choosing phase begins. Players receive the option of their choice following the rule of first come, first served. If only one player, the “chooser,” decides to stop sampling and make a final choice, that player obtains her chosen option. This forces the other player, the “receiver,” to accept the remaining option. If both players simultaneously stop sampling then one of two outcomes can occur: If players want different options, each player gets the option of his or her choice. If both players want the same option, the options are randomly assigned to each player.

The competitive sampling game is akin to “games of timing” (Dutta & Rustichini, 1993), in which two players independently decide when to stop a game and seize a reward while the reward either increases (preemption games) or decreases (war-of-attrition games) over time. An example of a preemption game is the “grab-the-dollar” game, in which two players have the option of either grabbing the money on a table or waiting for an additional period, during which the pot increases by one unit (Park & Smith, 2008). The players’ dilemma is that they want both to wait for a larger pot and to be the one claiming the money. The competitive sampling game has the nature of a preemptive race; here, the value of the options becomes clearer over time, but the first person to terminate sampling can decide which option to exploit. Nonetheless, it differs fundamentally from previous games of timing in that players face uncertainty not only about the other’s behavior but also about what is at stake—that is, the distribution of each option’s outcomes. In other words, the competitive sampling game is a competitive social game (Hertwig, Hoffrage, & the ABC Research Group, 2013), representing situations in which organisms need to trade off exploration of the quality of options for earlier exploitation in order to reduce the risk of the best option being snatched away by a competitor. In the following sections, we address the normative question of how much search is optimal in different variants of the competitive sampling game, and then describe the results of an experimental study.

2. How should accuracy and opportunity be traded off: a simulation study

How should decision makers adjust exploration efforts between solitary and competitive environments? To answer this question, we began by making the following

assumptions about the choice ecology, sampling rules, decision rules, and social environment, respectively. Let us emphasize here that our conclusions regarding good sampling sizes in the game will be contingent on these assumptions. In the discussion, we turn to alternative, more elaborate assumptions and more complex environments.

2.1. Choice ecology

Each game presents players with two options with two-outcome payoff distributions. Each distribution has a positive and a negative outcome that occur with complementary probabilities. Positive (O^+) and negative (O^-) outcomes are drawn from uniform distributions ranging from 0 to +100 and –100 to 0, respectively. The probability of the positive outcome $p(O^+)$ is drawn from a uniform distribution with support $[0, 1]$, while the probability of the negative outcome $p(O^-)$ is set to $1 - p(O^+)$. We define the option in the pair with the higher expected value as the H option, and we define the performance of a strategy as its likelihood of obtaining the H option.¹ For our analyses, we generated 10,000 pairs of payoff distributions and averaged expected strategy performance across all pairs.

2.2. Sampling rules

All players use a “fixed- N ” sampling rule, where N represents the player’s planned sampling size. A player with a fixed planned sampling size N will elect to continue sampling until the $N + 1$ sampling round, at which point he will stop search and choose. Players distribute their samples equally,² between the two options except where N is odd, in which case the player will allocate one additional sample to a randomly chosen option. Strategies with small N values dictate little exploration prior to exploitation, while those with large N values mandate extensive exploration. We calculated expected performance for fixed- N strategies with planned sampling sizes ranging from 1 to 50.

2.3. Decision rules

Given a pair of players, the player with the smaller planned sampling size is the *chooser* in the game and gets to decide which option to take. The player with the larger planned sampling size is the *receiver* and automatically receives the remaining option that was not chosen by their competitor. Choosers choose the option that has the highest observed sample mean (i.e., highest mean reward). This rule has been proposed in the context of n -armed bandit

¹ We discuss other reasonable performance measures in Section 5. For now, we note that for the two-gamble case, the probability of obtaining option H is similar, if not identical, to other performance measures such as the probability of outperforming one’s competitor. Additionally, assuming that distributions are not heavily skewed, the probability of obtaining option H in most cases should be very similar to the average expected reward.

² This assumption is made for simplicity. Although, on average, people draw roughly equal samples from both options in solitary decisions from experience (see Hertwig et al., 2004 Fig. 1), there is also evidence that sampling effort is impacted by factors such as the variability of outcomes encountered during search (Lejarraaga, Hertwig, & Gonzalez, 2012).

problems as a computationally simple method for estimating the values of actions such as the play of one of a slot machine's levers, and for using the estimates to select an action (Sutton & Barto, 1998). Unlike in n -armed bandit problems, in the sampling paradigm as studied here, the outcomes in the sampling stage inform the players about the value of an option but do not yet represent actual rewards (that is, sampling is exogenous; Denrell, 2007).

In the case that players sample only once, and thus observe an outcome from one option only, they use the following decision rule: If the sample is positive, *choose that option*, if the sample is negative, *choose the other option*. This rule is in the spirit of the win-stay, lose-shift strategy that has been shown to be effective in repeated games environments (e.g., Nowak & Sigmund, 1993). We refer to our variation of this strategy as the *take-good-enough, otherwise-shift strategy*. It specifies that a player will take an observed option if the sample was satisfying (any positive value in our simulation), otherwise he will reject the observed option and take the alternative, unobserved option.³

2.4. Social environments

The cost of sampling in competitive environments is likely to depend on the probability that a desirable option will be scooped up by a competitor. For this reason, we expect that the performance of sampling strategies will depend on the specific social environment an organism is in. To measure how competitors' decision speed affects the performance of different levels of search, we simulated choice performance in four different social environments. Mathematically, we defined social environments in terms of the probability distribution of opponents' sampling sizes. We generated four social environments: a *slow* environment in which competitors tend to have large sampling sizes and thus require extensive information before making a decision; a *fast* environment in which competitors tend to have small sampling sizes and are primarily motivated to not let good options slip away; an *uncertain* environment where competitors vary equally between small, medium, and large sampling sizes; and an *as-if solitary* environment that consisted of searchers taken from the original Hertwig et al. (2004) study on solitary decisions from experience. This environment is called as-if solitary because competitors behave as if they are in a solitary environment. To the extent that competition will likely reduce search, the as-if solitary environment represents one possible distributional upper bound on how long individuals will search under competition.

2.5. Simulation results

We show mean performance results across 10,000 randomly generated pairs of payoff distributions drawn from

³ Importantly, this rule assumes no uncertainty aversion (Ellsberg, 1961), in that players do not hesitate to take a completely unobserved option when an observed option is found to contain negative outcomes. In the empirical study, we test this assumption (and find evidence against it). However, for the purposes of simplicity we maintain the assumption in determining good sampling sizes in the simulation.

the aforementioned choice ecology (see Appendix A for details). Again, the benchmark used to assess the performance of a sampling rule is the probability that an agent obtains the option with the higher expected value in a gamble pair (the H option). We present the simulation results in three sections. First, we contrast expected outcomes for agents who are choosers versus agents who are receivers in a game as a function of the number of sampling rounds in that game. Second, we demonstrate an imbalance in the costs of oversampling versus undersampling. Finally, taking into account this imbalance in costs, we derive the best search length for each social environment.

To what extent does being the chooser (i.e., the one whose planned sampling size is smallest) increase the likelihood of obtaining the H option? Is it always good to be the chooser in a game or is it sometimes better to be the receiver and allow a competitor to choose? To answer this question, we calculated the probability that an agent obtains the H option given that he ends the game as the chooser across sampling rounds 1–50.⁴ In other words, assuming the game lasts for x sampling rounds, what is the probability that an agent obtains the H option if he or she is the chooser in the game? Fig. 1 shows the expected outcomes for choosers compared to receivers across rounds 1–50. Recall that, as our implementation of the competitive sampling game requires that the receiver take the option not chosen, the probability that the receiver obtains the H option is simply the complement of the probability that the chooser chose it.

We draw two main conclusions from the data presented in Fig. 1. First, across all sampling rounds, choosers are always expected to obtain the H option with probability greater than .50. Because receivers receive the H option with the complement of the chooser's probability, receivers always obtain the H option with a probability less than .50. No matter how few samples one takes, the expected outcome of being a chooser is always better than the expected outcome of being a receiver. Second, the probability that a chooser obtains the H option increases monotonically with additional sampling rounds, but with marginally decreasing gains. In other words, the gain in information a chooser gets from an additional sample in early sampling rounds is larger than the gain in later sampling rounds. From these two findings, it follows that it is always better to have more sampling rounds *as long as* one ends the game as the chooser. In Fig. 1, this means that an agent should try to get as far to the right on the choosing line as possible without being "scooped" and dropping down to the increasingly negative receiving line.

These findings also allow us to construct an optimal sampling rule for an omnipotent player who knows how long her competitor plans to sample. If the omnipotent player knows that his competitor has a fixed and known planned sampling size of n_c , then his best sampling size is $n_c - 1$ (or 1 when $n_c = 1$). Of course, most people are not omnipotent and do not have perfect knowledge of their

⁴ Hertwig and Pleskac (2010) conducted a very similar simulation that paralleled ours (in the case where a player is always the chooser). Our results are virtually identical.

opponent's planned sampling size. Instead, we suspect that people will base their sampling decisions on their expectations of their competitors' behavior, where expectations are defined as a probability distribution over sampling sizes. In other words, players could ask themselves: "How likely is my competitor going to stop search on each sampling round? Once a player has these expectations, she then needs to take into account the costs of over- versus undersampling. If the costs of undersampling (deciding too quickly) are larger than the costs of oversampling (deciding too slowly), then a player should err on the side of sampling too much relative to her expectations of her opponent. On the other hand, if the costs of undersampling are smaller than the costs of oversampling, then the player should shorten her search relative to her expectations of her opponent.

To determine the relative costs of over- versus under-sampling in the competitive sampling game, we calculated the expected choice performance of an agent given all combinations of planned sampling sizes from 1 to 10 that an agent and his or her competitor might implement. These data are presented in Fig. 2, where the *x*-axis represents an agent's planned sampling size and the *y*-axis represents a competitor's planned sampling size.

Consistent with our previous conclusion, Fig. 2 shows that for any competitor's planned sampling size n_c , an agent's best planned sampling size, n_a , equals $n_c - 1$ (or 1 when $n_c = 1$). For example, if a competitor's planned sampling size is 8, the best planned sampling size for the agent is 7, with an 83% chance of obtaining option *H*. However, consider the cost of over- versus undersampling against this competitor. If the agent undersamples by 2, with a planned sampling size of 5, he will still be the chooser in the game and have an 80% chance of obtaining option *H*—a drop of only 3 percentage points relative to the best possible outcome. If, on the other hand, the agent oversamples by 2, with a planned sampling size of 9, he will be the receiver in the game and will have only a 16% chance of obtaining option *H*—a plunge of 67 percentage points.

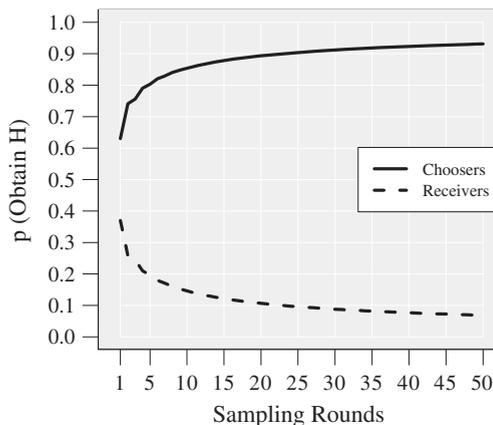


Fig. 1. Mean probability that an agent obtains the *H* option given that the game stops at a specified sampling round across 10,000 gambles. The choosing line shows the probability for the chooser, the player with the larger planned sampling size, while the receiving line shows the probability for the receiver, the player with the smaller planned sampling size.

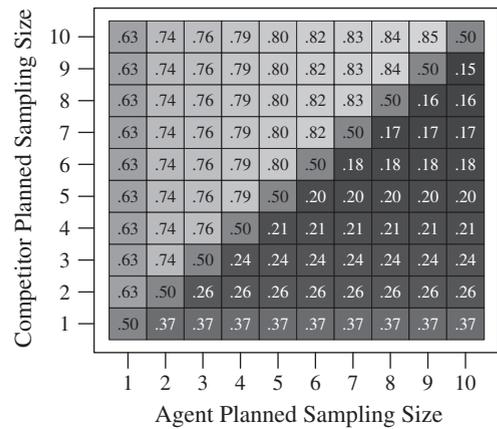


Fig. 2. Probability that an agent obtains the *H* option given the agent's planned sampling size and the competitor's planned sampling size. The figure shows mean values over 10,000 randomly generated two-outcome gambles.

These results show that the cost of oversampling is much larger than the cost of undersampling: it is always better to undersample (by any amount) and keep the chooser advantage, than to oversample (by even one sample) and suffer the receiver disadvantage.

Given that it is better to err on the side of undersampling versus oversampling, how should an individual behave in different social environments? In other words, how little should one sample before making a decision given certain expectations about the behavior of competitors? To answer this, we calculated a player's expected probability of obtaining the *H* option given his or her planned sampling size within each of the four social environments (see Appendix B for details). In the slow environment, competitors had relatively large planned sampling sizes (mean of 30), following a bounded, discretized normal distribution with a standard deviation of 5. In the fast environment, competitors had relatively small planned sampling sizes (mean of 3.33) following a right-skewed distribution. In the uncertain environment, competitors had—with equal probability—any planned sampling size from 1 to 50 (mean of 25.5). Finally, in the as-if solitary environment, competitors had a right-skewed distribution of sampling sizes (mean of 18). Fig. 3 (left panel) shows the probability mass functions for each of these social environments.⁵

The right panel of Fig. 3 shows the expected probability than an agent obtains the *H* option given the planned sampling size within a specific social environment. In a slow environment, agents with a planned sampling size of 18 did best, with an 88.2% chance of obtaining *H*. In contrast, in an uncertain environment, the best planned sampling size was 6, with a 75.8% chance of obtaining *H*. In a fast

⁵ For the purposes of the prescriptive analyses, these distributions can represent either the variability in the behavior of one's opponent from one game to another, or the variability in the behavior of an entire population of individual competitors. Assuming that an opponent's sampling rule in each game is an independent, random sample from its parent distribution, the mathematics are the same whether we attribute variability to inter- or intra-individual causes.

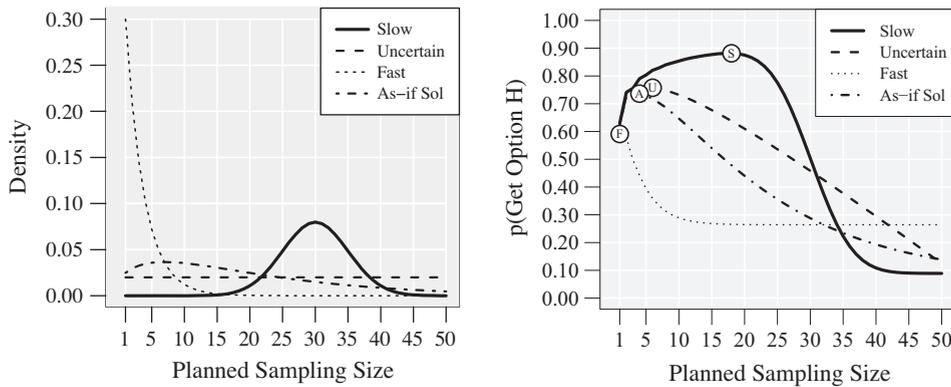


Fig. 3. Left panel: Probability of an agent encountering a competitor with a given planned sampling size in the slow, fast, uncertain, and as-if solitary environments (see text). Right panel: Results from the agent-based simulation averaged across 10,000 randomly generated decision problems. The x -axis indicates the planned sampling size of an agent, and the y -axis shows the expected probability of obtaining the higher expected value option (H) as a function of the planned sampling size in the four social environments. While lines are continuous, the underlying data is discrete.

environment, agents with a planned sampling size of just 1 did best, obtaining the H option 59% of the time. In the as-if solitary environment, a planned sampling size of 4 proved best, with a 73.8% chance of obtaining H . Finally the only sampling size that ensures that one will obtain option H with probability no less than .50, regardless of the behavior of one's competitor, is a sample size of just 1 using the take-good-enough, otherwise-shift strategy.

These results show that players should dramatically reduce exploration in a competitive context when they lack clear information of their competitors' intentions. Consider for illustration the uncertain social environment, in which a competitor is equally likely to stop anywhere between 1 and 50 samples (see Fig. 3, left panel). If an omnipotent player knew exactly how long her competitor planned to sample, then the best strategy would simply be to take one fewer samples than her opponent. However, in this uncertain social environment, the exact planned sampling size of the opponent is unknown. What happens if a player plans to sample one less than the *expected* planned sampling size of her competitor? In this uncertain social environment where the expected planned sampling size of a competitor is 25.5, this rule would dictate a planned sampling size of 25. Reference to Fig. 3 (right panel) shows that this planned sampling size constitutes dramatic oversampling, as the best planned sampling size for this social environment is only 6—less than one-fourth of the competitor's expected sampling size. In this example, planning to sample one round less than the expected sampling size of one's uncertain competitor leads to oversampling by over 19 rounds. The reason behind this dramatic effect of uncertainty about the other agent's actions on the best planned sampling size is the disproportionate costs of over- versus undersampling.

To conclude, adaptive sampling in the competitive sampling game depends on expectations of one's competitors. If competitors value accuracy highly and consequently represent a slow social environment, decision makers can afford to gather more information. Yet there is considerable asymmetry in the costs of over- versus underestimating competitors' need for accuracy. In our choice ecology, underestimating one's competitor's sampling size, no matter by what degree, will always ensure that one will be the

chooser and thus more likely than not to obtain option H . On the other hand, overestimating one's competitor's sampling size, no matter by what degree, will always ensure that one will be the receiver and thus more likely than not to obtain the short end of the stick (i.e., the lower expected value option, L). For these reasons, we find that it is better to err on the side of underestimating the competitor's sampling size and minimizing the risk of being scooped.

These simulations show that competition presents a substantial additional cost of search. Consequently, we expect that real people will search much less in a competitive compared to a solitary context. But how much more restricted will it be? Will real people competing with others decrease their search in a magnitude prescribed by our simulations? Or will people be reluctant to decrease pre-decisional search so dramatically? To answer this question, we conducted an empirical study on the competitive sampling game and compare the search behavior of people participating in solitary to competitive games.

3. An empirical investigation of the competitive sampling game

3.1. Method

A total of 180 students from the University of Basel participated in the study.⁶ They received a flat fee of CHF 7.50 (approximately \$8.12 at the time) for their participation, as well as a bonus contingent on their winnings in the game. The mean bonus across both experimental conditions was CHF 1.18 (approximately \$1.26) with a standard deviation of CHF 1.19. Participants completed the study in groups of four, each on a separate computer. They received no information about the choice ecology prior to beginning the task. All players began by playing three practice games without financial consequences to familiarize themselves with the experimental interface (see Appendix C for practice game parameters). They were then presented with five decision tasks. Each decision task contained two, two-outcome

⁶ Gender data were not recorded due to a programming error.

Table 1
Choice ecology.

Gamble set	Task 1		Task 2		Task 3		Task 4		Task 5	
	H	L	H	L	H	L	H	L	H	L
1	(37, 0.44, -17): 6.76	(25, 0.22, -11): -3.08	(29, 0.36, -13): 2.12	(43, 0.35, -20): 2.05	(33, 0.46, -15): 7.08	(49, 0.28, -23): -2.84	(55, 0.41, -26): 7.21	(55, 0.28, -26): -3.32	(37, 0.44, -17): 6.76	(37, 0.26, -17): -2.96
2	(43, 0.43, -20): 7.09	(29, 0.24, -13): -2.92	(49, 0.35, -23): 2.20	(33, 0.35, -15): 1.80	(25, 0.5, -11): 7.00	(37, 0.26, -17): -2.96	(55, 0.41, -26): 7.21	(55, 0.28, -26): -3.32	(37, 0.44, -17): 6.76	(37, 0.26, -17): -2.96
3	(55, 0.41, -26): 7.21	(55, 0.28, -26): -3.32	(37, 0.44, -17): 6.76	(37, 0.26, -17): -2.96	(29, 0.48, -13): 7.16	(43, 0.27, -20): -2.99	(49, 0.42, -23): 7.24	(33, 0.25, -15): -3.00	(25, 0.36, -11): 1.96	(37, 0.35, -17): 1.90
4	(37, 0.44, -17): 6.76	(25, 0.22, -11): -3.08	(29, 0.36, -13): 2.12	(43, 0.35, -20): 2.05	(37, 0.44, -17): 6.76	(55, 0.28, -26): -3.32	(49, 0.42, -23): 7.24	(49, 0.28, -23): -2.84	(33, 0.46, -15): 7.08	(33, 0.25, -15): -3.00
5	(33, 0.46, -15): 7.08	(33, 0.25, -15): -3.00	(49, 0.42, -23): 7.24	(49, 0.28, -23): -2.84	(55, 0.35, -26): 2.35	(37, 0.35, -17): 1.90	(43, 0.43, -20): 7.09	(29, 0.24, -13): -2.92	(25, 0.5, -11): 7.00	(37, 0.26, -17): -2.96
6	(55, 0.41, -26): 7.21	(37, 0.26, -17): -2.96	(25, 0.36, -11): 1.96	(37, 0.35, -17): 1.90	(29, 0.48, -13): 7.16	(43, 0.27, -20): -2.99	(49, 0.42, -23): 7.24	(49, 0.28, -23): -2.84	(33, 0.46, -15): 7.08	(33, 0.25, -15): -3.00
7	(43, 0.43, -20): 7.09	(43, 0.27, -20): -2.99	(33, 0.46, -15): 7.08	(49, 0.28, -23): -2.84	(29, 0.48, -13): 7.16	(29, 0.24, -13): -2.92	(55, 0.35, -26): 2.35	(37, 0.35, -17): 1.90	(37, 0.44, -17): 6.76	(25, 0.22, -11): -3.08
8	(55, 0.41, -26): 7.21	(37, 0.26, -17): -2.96	(49, 0.35, -23): 2.20	(33, 0.35, -15): 1.80	(25, 0.5, -11): 7.00	(37, 0.26, -17): -2.96	(43, 0.43, -20): 7.09	(43, 0.27, -20): -2.99	(29, 0.48, -13): 7.16	(29, 0.24, -13): -2.92
9	(25, 0.36, -11): 1.96	(37, 0.35, -17): 1.90	(49, 0.42, -23): 7.24	(33, 0.25, -15): -3.00	(29, 0.48, -13): 7.16	(29, 0.24, -13): -2.92	(43, 0.43, -20): 7.09	(43, 0.27, -20): -2.99	(37, 0.44, -17): 6.76	(55, 0.28, -26): -3.32
10	(55, 0.41, -26): 7.21	(37, 0.26, -17): -2.96	(29, 0.36, -13): 2.12	(43, 0.35, -20): 2.05	(33, 0.46, -15): 7.08	(49, 0.28, -23): -2.84	(37, 0.44, -17): 6.76	(37, 0.26, -17): -2.96	(25, 0.5, -11): 7.00	(25, 0.22, -11): -3.08
11	(49, 0.35, -23): 2.20	(33, 0.35, -15): 1.80	(43, 0.43, -20): 7.09	(29, 0.24, -13): -2.92	(37, 0.44, -17): 6.76	(37, 0.26, -17): -2.96	(25, 0.5, -11): 7.00	(25, 0.22, -11): -3.08	(37, 0.44, -17): 6.76	(55, 0.28, -26): -3.32
12	(49, 0.42, -23): 7.24	(33, 0.25, -15): -3.00	(55, 0.35, -26): 2.35	(37, 0.35, -17): 1.90	(29, 0.48, -13): 7.16	(43, 0.27, -20): -2.99	(37, 0.44, -17): 6.76	(37, 0.26, -17): -2.96	(25, 0.5, -11): 7.00	(25, 0.22, -11): -3.08

Note: Gamble sets used in both the solitary and the competitive conditions. Rows correspond to the 12 different combinations of decision tasks. *H* represents the higher expected value option, and *L* represents the lower expected value option within each decision task. Each option is a discrete, two-outcome random variable with one positive and one negative outcome that occur with complementary probabilities. Values in parentheses are the value and the probability of the positive outcome, and the value of the negative outcome, for each gamble. The value outside the parentheses is the expected value of the gamble.

gambles, each with one positive and one negative outcome occurring with complementary probabilities. The gamble sets were constructed such that in certain pairs the options differed in expected value and in others they did not; likewise, in certain pairs the options differed in range and in others they did not (see Appendix C for a full description of how gamble parameters were selected). Three different orders of each of the 12 gamble sets were created, resulting in 36 unique experimental sessions (see Table 1). Location of the urns on the screen was randomly determined for each decision task and on each run.

At the outset of each decision task, participants saw two options represented visually as opaque urns. They were told that each urn contained 100 virtual balls, each of which was worth a (not necessarily unique) number of points. Participants were informed that they would be rewarded with one-tenth of the average value of all the balls in the urn they chose (or were allocated). Each of the participants ($n = 36$) assigned to the solitary condition completed one of the 36 unique experimental sessions. These participants could sample from the urns as many times as they wished before making a final choice. Having made a final choice of an urn in a decision task, they moved onto the next task. The other 144 participants played each decision task in the *competition condition*. At the beginning of each task, they were paired randomly with one of the other three participants. This pairing was done independently between tasks. Players were not told which person (of the three) they were playing against in each decision task.

Every decision task, in both the solitary and competition conditions, began with one mandatory sampling round. On every subsequent sampling round, each player indicated whether he or she wanted to sample from an urn or to make a final choice. These decisions were made privately and were only revealed to both players after both had made a sampling or choice decision. If both wanted to take a sample, they were asked to click on an urn and viewed a randomly sampled outcome from that urn. Players could see which urn the other player sampled from, but could not see the outcome the other player observed. If, after observing a sample, both players wanted another sample then another sampling round began. If one player decided to make a final choice (becoming the “chooser”⁷), she then selected the urn she wanted and her choice was recorded. Subsequent to the chooser’s choice, the other player (the “receiver”) was informed that her competitor had made a choice and that he must take the remaining non-chosen urn. If both players made a choice on the same sampling round, one of two outcomes was possible: If the two players chose different urns, they each received the urn of their choice. If both players chose the same urn, the two urns were randomly assigned to the players. After final choices were made and players learned which urn they received, they were randomly paired again and the next decision task began. The random pairing was done independently of prior rounds, so a player could play the same opponent on sequential games. Participants did not receive immediate

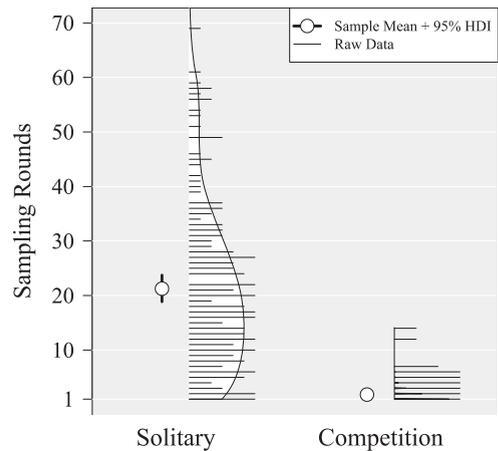


Fig. 4. Distribution of sampling rounds across all decision tasks and individuals, separately for the solitary and the competition conditions. One sampling round value of 100 in the solitary condition is not displayed. The 95% HDI interval for the solitary condition is plotted ‘behind’ the solitary sample mean.

task-by-task feedback on how much money they won from their chosen urns. At the end of the session, participants were informed how much money they had earned across the five decision tasks and were paid accordingly.

4. Results

We used Bayesian graphical modeling for all inferential statistics. Bayesian posterior densities were calculated using the R2jags package in R. Posterior densities were calculated with uninformative uniform priors, 10,000 iterations, a burn-in value of 1000, and no thinning. We conducted Bayesian hypothesis tests using Bayes Factors calculated using the Savage–Dickey method for nested model comparisons. We used the conventions developed by Jeffreys (1961) to determine the categorical degrees of strength indicated by Bayes Factors (BF). All raw data and complete code are available in our online [supplementary materials](#). In the following sections, we first report on search and then on choice, comparing both behaviors in the competition condition relative to the solitary condition.

4.1. How drastically do people restrict explorative behavior under competition?

We measure exploration efforts by the number of sampling rounds tasks lasted prior to a choice. For solitary games, this is simply the number of samples the player took. For competitive games, this is the number of sampling rounds that occurred prior to the first choice. Fig. 4 presents the distribution of sampling rounds across all decision tasks in the solitary and competition conditions. In the solitary condition, the median number of sampling rounds was 18 (mean of 21.05; 95% highest density interval [HDI]: 18.78, 23.67).⁸ In the competition condition, in

⁷ Players were not explicitly given the “chooser” and “receiver” labels in the experiment.

⁸ The mean for the distribution of sampling sizes in the solitary condition was calculated from the 95% HDIs for the p and r parameters in the negative binomial distribution.

contrast, the median number of sampling rounds was 1 (mean of 1.82; 95% HDI: 1.61, 2.08).⁹ The difference in sample means was 19.27 (95% HDI: 17.00, 21.92) and provides extremely strong evidence against the hypothesis that the means of the two distributions were the same ($BF > 100$). Thus, the amount of sampling by participants who were competing for resources was dramatically lower than that of participants who were searching alone.¹⁰

4.2. How much does very restricted search compromise decision quality?

To see how the restricted search in the competition condition affected decision quality, we calculated how often choosers in the competition condition chose the H option relative to receivers and to participants in the solitary condition. These results are presented in Fig. 5. In the solitary condition, players chose the H option in 71% (95% HDI: 64%, 78%) of decision tasks. This result constitutes extremely strong evidence for the hypothesis that participants in the solitary condition were more likely than chance to choose the H option ($BF > 100$). Next we analyze the outcomes for choosers and receivers in the competition condition. In cases where players chose the same option simultaneously, one player was randomly assigned to be the chooser (and obtained the option both chose) and the remaining player was assigned to be the receiver (and obtained the alternative option). Choosers obtained the H option in 58% of decision tasks (95% HDI: 53%, 63%)—fewer than in the solitary condition ($BF = 12.36$, strong evidence), but more than would be expected by chance alone ($BF = 8.90$, moderate evidence).

⁹ The mean for the distribution of sampling sizes in the competition condition was calculated from the 95% HDIs for the p parameter in the geometric distribution.

¹⁰ The fact that sampling rounds decreased in competition relative to solitary conditions, is necessary, but not sufficient evidence that competition reduced individual sampling decisions. The reason for this lies in how sampling rounds are defined. Under competition, sampling rounds are defined at the level of pairs of participants rather than individual participants. Because sampling rounds are restricted by the behavior of the fastest player in a pair, we would expect a decrease in sampling rounds in the competitive task relative to the solitary task even if players did not change their sampling rules. For example, if two players employ fixed- N rules of 5 and 10, respectively, across solitary and competitive games, the average number of sampling rounds in the solitary games would be 7.5, while the average number of sampling rounds under competition would (always) be 5. To test whether or not this shrinkage effect could explain the different distributions of sampling rounds between solitary and competition conditions, we generated all possible pairs of sampling rounds from the solitary game and calculated the minimum sampling round number from each pair. This represented the expected distribution of sampling rounds in the competition condition if behavior was the same as in the solitary condition. The median number of sampling rounds in this distribution was 11 (mean of 12.61, 95% HDI: 12.45, 12.76). The difference in the mean sampling rounds between this distribution and the observed distribution for the competition condition was 10.76 (95% HDI: 10.49, 11.06), thus offering extremely strong evidence against the hypothesis that the means of the two distributions are the same ($BF > 100$). We conclude that the difference between the mean number of sampling rounds in the competition condition and the solitary condition was due not only to the rules of the competitive game (i.e., the fastest player determines sample size) but also to the fact that competition per se shifted the balance from exploration to exploitation.

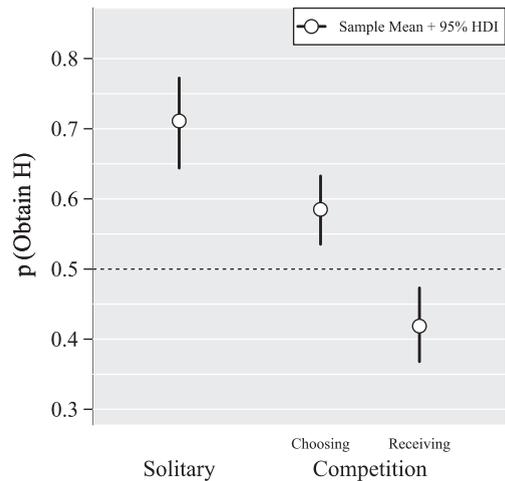


Fig. 5. Proportions of tasks where players obtained the higher expected value (H) option in the solitary and the competition conditions. Error bars represent 95% highest density intervals for the population probability.

Thus, the reduced information available to fast choosers in the competitive condition indeed reduced their choice performance relative to solitary choosers. Nevertheless, as choosers took the better option at above chance level, receivers obtained option H in only 42% (95% HDI: .37, .47) of cases¹¹ This is consistent with our simulation analysis (Fig. 3, right panel) showing that fast choosing is advantageous under competitive conditions.

4.3. How did players make choices based on minimal information?

The following analyses focus on the competition condition only. To further analyze the specific decisions that produced the distribution of sampling rounds in the competition condition, we looked at how quickly players terminated sampling. For each of the 720 decision tasks (144 participants \times 5 decision tasks each), we recorded the total number of sampling rounds that occurred in the task, and whether players were choosers or receivers (i.e., two choosers or one chooser and one receiver). We found that in 32% (227 of 720) of all cases, participants were choosers who decided to choose immediately after the first sampling round. Of these choices, 88% (200 of 227) were consistent with the take good enough, otherwise-shift heuristic. The remaining 12% (27 out of 227) either chose an option with an observed negative value, or did not choose an option with an observed positive value. In 22% (159 of 720) of all cases, participants were receivers after the first sampling round, because they opted to continue sampling while their competitors decided to choose. Participants made it to the second round in only 46% (334 of 720) of cases. Of these 334 participants, 34% (113 of 334) were

¹¹ Receivers did not receive option H at a percentage exactly equal to one minus the percentage that choosers wanted option H (which was 57%). This is due to the effects of simultaneous choosing by choosers. In games where both players simultaneously chose option H , both wanted option H , but only one player got it.

Table 2
Distribution of participant-level proportion of decision tasks ending with a choice.

0/5	1/5	2/5	3/5	4/5	5/5
2 (1.4%)	16 (11.1%)	34 (23.6%)	45 (31.3%)	30 (20.8%)	17 (11.8%)

choosers who decided to stop sampling, 25% (82 of 334) were receivers, and 41% (139 of 334) “survived” to the third sampling round. This analysis demonstrates that, while not all choosing decisions were made after one round (as prescribed by the analysis of the fast environment; Fig. 3), people indeed drastically confined their information search, and to a greater extent than prescribed by the uncertain and solitary search environments.

What makes people decide to choose after just one sample versus to continue sampling? Our data suggest the valence of the first sample influences this decision. Of the 261 cases in which a player experienced a positive outcome in the first sample, he or she stopped sampling and chose immediately in 130 cases (50%; 95% HDI: 44%, 56%). Of the 459 cases in which a player experienced a negative outcome in the first sample, he or she stopped sampling in only 97 cases (21%; 95% HDI: 18%, 25%). Thus, players were more willing to immediately choose an option with a known positive outcome than they were to reject an option with a known negative outcome and take a completely uncertain option ($BF > 100$, extreme evidence)—perhaps a manifestation of aversion to ambiguity (Ellsberg, 1961).

4.4. How closely did players pursue a single strategy?

Next, we examined how consistent individual participants were in their behavior across tasks. In the competitive sampling game, some players could always try to be the chooser, while others might require so much pre-decisional information that they always defer the choice to their competitor. To see if people had stable choosing versus receiving outcomes, we calculated the percentage of all 5 games that each participant was a chooser. Table 2 reports the distribution of these percentages across individuals. We found little evidence that people’s search strategies resulted in stable outcomes across decision tasks. Only 19 individuals (13.2%) were either always choosers (17) or always receivers (2), while the remaining 86.8% ended some tasks as the chooser and other tasks as the receiver.

Next, we determined whether and how individuals changed their behavior across the tasks. For example, if a player ended up as a receiver in one task, did she decrease her sampling in order to increase the chance of being a chooser in the subsequent task? In contrast, did the chooser take the liberty of sampling a little more in the next task? To answer this question, we calculated the marginal probability that a player was the chooser each task in addition to the probability of being the chooser conditional on his or her status in the previous task (i.e., chooser versus receiver). If behavior in a game changes as a function of the outcome in the previous game, we would expect differences in the conditional probability of choosing when a player was a chooser in the previous game compared to

being a receiver in the previous game. Table 3 reports the results.

We begin by looking at the marginal probability of choosing across decision tasks. If players adjust their explorative efforts downward with each round, the probability that they end games by choosing (rather than receiving) should increase and converge toward 1. We did not find substantial evidence for this, as the probability of choosing oscillated between around 57% and 61% across tasks. Next, we looked at the probability of choosing, conditioned on the outcome of the previous task. We found no evidence that being a receiver in one round prompted less search and a higher probability of being the chooser in the next task. In other words, we do not find clear evidence that players changed their behavior across tasks based on experience. One possible explanation for this finding is that, although players were told whether they were the chooser or the receiver in any given round, they did not receive immediate feedback on the direct monetary consequences of their behavior. Without immediate consequential feedback, players may not have had sufficient information to adjust their sampling strategies.

5. General discussion

We designed the competitive sampling game to extend the sampling paradigm used in research on decision from experience (see Hertwig & Erev, 2009) to competitive decision tasks under uncertainty. The task enables researchers to investigate how people adapt their exploration efforts in response to the simultaneous presence of uncertainty about nature (i.e., the parameters of the payoff distribution) and uncertainty about the social environment of competitors (i.e., their desire for accuracy versus their desire to beat their competitors to the punch). In our initial research using the game, we found that competition dramatically affects the exploration–exploitation tradeoff. In a simulation analysis, we found that sampling sizes as low as 1 or 2 can be best in certain environments when people compete with others for advantageous options. Empirically, our participants showed dramatically reduced search in competitive task compared to solitary one.

5.1. Varying the number of players and options

In our experiment and simulations, we contrasted a solitary task with one player and two options, with a competitive task with two players and two options. One could wonder whether it is merely the presence of competition, or the ratio between competitors and the number of available options that drives the need for speedy decisions.¹² A moment’s reflection makes it clear that good sampling rules

¹² We thank Jonathan Nelson for pointing this out.

Table 3
Choosing behavior across sequential decision tasks.

Rounds/ decision task	$p(\text{choosing})$	$p(\text{choosing} \text{previouschoosing})$	$p(\text{choosing} \text{previousreceiving})$
1	82 (56.9%)		
2	83 (57.6%)	48 (58.5%)	35 (56.4%)
3	87 (60.4%)	51 (61.4%)	36 (59.0%)
4	83 (57.6%)	54 (62.1%)	29 (50.9%)
5	89 (61.8%)	55 (66.3%)	34 (55.7%)

are likely to depend both on the number of options available, and the number of competitors: If there are many options available relative to players, then one can take more time for sampling, knowing that even if all other players make quick decisions, there will likely be many good options remaining. To find out the relationship between the number of competitors and number of options on good decision speed, we ran a supplementary set of simulations in which we systematically manipulated three factors: number of players (from 2 to 5), number of options (from 2 to 5) and the speed of competition (fast, uncertain, and slow). For each factor combination, we simulated the performance of an agent using a fixed sampling size of 1 through 15. In contrast to our previous simulations, we now define performance as the average expected reward the agent obtained across all simulations.¹³ Details of the simulation are in [Appendix E](#).

We present the main results of the simulation in [Fig. 6](#). Each plot shows an agent's planned sampling size on the horizontal axis, and the agent's expected reward on the vertical axis. Each line corresponds to a different social environment, mirroring three of the four (excluding the as-if solitary condition) from our earlier simulations (see left panel of [Fig. 3](#)). Plots in each column refer to different numbers of players, from 2 to 5, while plots in each row refer to different numbers of options, from 2 to 5. The top left graph (2 players and 2 options) replicates the same social and environmental structure as our initial simulation. Within each plot, the sampling size that maximizes an agent's expected rewards for each social environment is highlighted with an enlarged point.

We briefly summarize 4 key results from [Fig. 6](#). First, holding the number of players constant, as the number of options increases the best sampling sizes tend to increase. Additionally, the expected reward given the best sampling size increases as well. This means that the more options there are, the longer one should search, and the better one's expected outcomes will be. Second, holding the number of options constant, as the number of players increases the best sampling sizes tend to decrease. Both of these results support our prior intuitions: the degree to which people should reduce sampling in the presence of competition depends on the number of options available and the number of other players.

¹³ We use expected reward instead of the probability of obtaining the highest expected value option for two reasons. First, organisms frequently want to obtain good options, not necessarily the best option. Second, as the number options increases, the probability that any player will discover and take the highest expected value option will necessarily decrease. This makes it more difficult to compare performance between option number conditions.

Next, we look at the effect of the player/option ratio on performance. If the ratio of players to options remains the same, does the absolute number of players and options matter? We find that indeed, there is a substantial effect. Consider games where the player: option is 1:1. Here, we find that as the absolute number of players and options increases, the risk involved with taking large samples increases in a fast environment, but decreases in a slow one. To see this, compare the expected rewards of having a large (15) sampling size in the 2:2 game (top left panel) compared to the 5:5 game (bottom right panel). In the 2:2 game, extensive search against slow competition leads to an expected reward of around 20, while the same level of search against fast opponents leads to an expected loss of around -10 . Here, the difference in expected rewards against slow and fast opponents is 30. Now consider the 5:5 game. Here, the expected reward against slow opponents *increases* to around 35, while the expected loss against fast opponents *decreases* to around -15 . Now, the difference in expected rewards between competitors is 50, an increase of 66% in the range of potential outcomes compared to the smaller 2:2 game. This means that when the absolute number of players and options increases, while keeping the player to options ratio 1:1, both the potential benefits one can gain using extensive search against slow competition increases while the potential losses one can suffer against fast competition increases. In other words, the more players and options are in the game, the more risk one runs (with 'risk' defined as the difference between the expected reward with the best sample size and with the largest sample size) by extensive search, in fast and uncertain environments. Independently of this effect, the main result from our previous analyses still hold—the faster you expect your opponents to decide, the faster you should decide, regardless of how many options and how many players are in the game.

Next, we consider games where there are more players than options. These games are akin to real-world problems such as house-hunting and mate-search where there may be more 'buyers' than 'sellers.' For example,¹⁴ in Beijing, men outnumber women, and thus (heterosexual) men find themselves in a competitive game with more players than options ([Jacobs, 2011](#)). Assuming that options cannot be shared among players, these games necessitate that some players will leave without an option of their own. In our simulation, we assumed that these players receive neither rewards nor losses from leaving empty-handed. However, one can easily imagine real world decisions where this assumption does not hold. If you are competing with others for one of three open positions at a company, it might be much worse to get no job at all than to get a random (or even the worst) job of the three. Similarly, in mate-selection, leaving empty-handed could very well be the worst possible outcome from an evolutionary perspective. To incorporate this cost, one could assign a fixed negative loss for players that leave the game empty-handed, with larger losses representing domains where it is especially bad to leave with no option (e.g.; mate-search). While we do not run these

¹⁴ We thank Jonathon Nelson for providing this example.

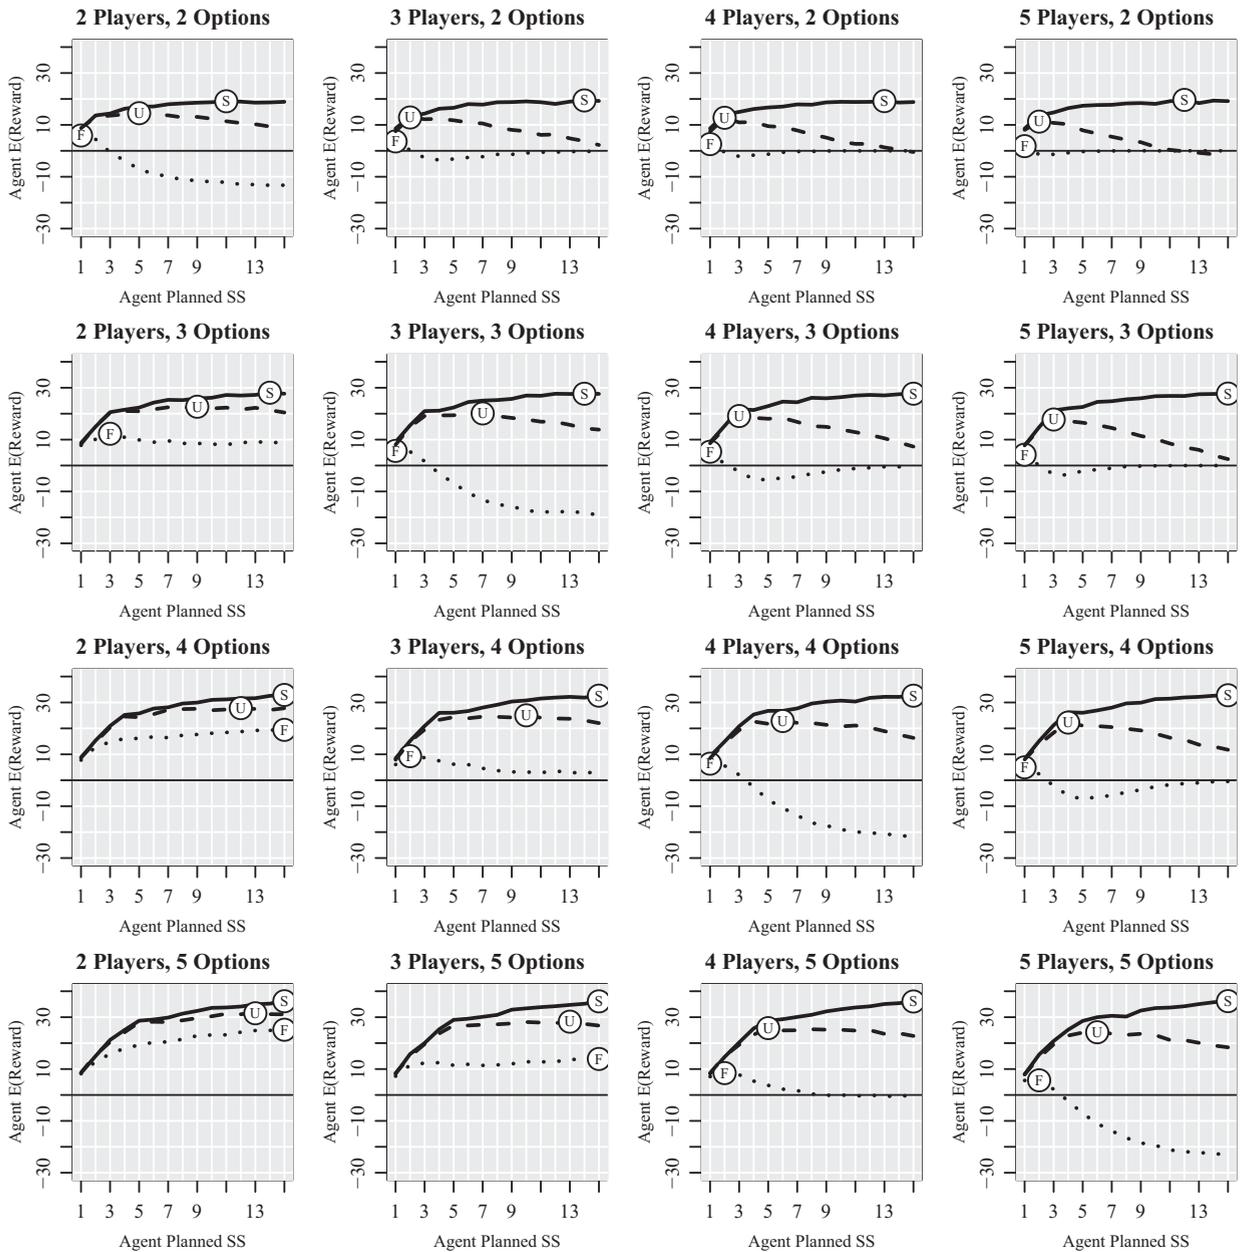


Fig. 6. Simulation results depicting the expected reward (vertical axes) of an agent in the competitive sampling game given a specified sampling size (horizontal axes). Separate plots in each column correspond to different numbers of players in the game, while separate plots in each row correspond to different numbers of options in the game. Solid, dashed, and dotted lines refer to the slow, uncertain, and fast social environments respectively.

simulations, our prediction for the effect is clear: as the cost of leaving empty-handed increases, sampling sizes should decrease.

5.2. Alternative search and decision rules

In our simulations, we assumed that players used a “fixed-*N*” sampling rule. That is, players were assumed to have a fixed sampling size threshold that they had to reach before making a decision. Additionally, we assumed that players distribute their samples equally between options. We chose to limit our analyses to this simple class of

search rules as a starting point for exploring the effects of different exploration efforts on performance. Of course, the fixed-*N* sampling rule plus equal allocation constitutes just one of many possible search rules people are likely to use. For example, one promising, more complex, class of search rules are those that compare sample statistics with an information threshold in order to decide whether to stop or to continue sampling. These rules have been found to be promising both normatively and descriptively (e.g., [Busemeyer & Rapoport, 1988](#)). While we cannot claim that fixed-*N* sampling rules represent either the best approach to the competitive sampling game or that most

participants employed them, we do predict that fixed- N rules will mimic the behavior of more complex search rules. For example, a fixed-1 rule should behave very similarly to a rule that says “Always try to choose before your opponent” or “Choose after the difference in sample means is greater than ε (where ε is a small threshold).” Similarly, a Fixed-30 rule will behave similarly to a conservative rule such as: “Choose when the expected probability of choosing the best option is greater than p (where p is a large probability).” Here is our point: to the extent that small fixed- N values mimic search rules that make do with minimal information before making a decision, and large fixed- N rules mimic those that require extensive information, our conclusion that competition should and does reduce the amount of information people require before making a decision should hold.

Finally, we did not deal with agents’ expectations concerning their competitor’s search behavior. Rather, we took the competitor sampling size distributions as given, and determined which sampling size best responds to them. We do suspect that real people in competitive tasks (in our empirical study and in the real world) try to predict the sampling size of their opponents and, through some iterated process of strategic thinking (e.g., Ho, Camerer, & Weigelt, 1998) settle on a sampling size. Future research should model the processes by which a person generates expectations of other player’s behavior, translates those expectations into a decision rule, and after gaining experiences, updates their expectations.

5.3. When will very frugal search fail?

Our simulation results are based on aggregation across specific distributions of gambles and are only valid within those distributions. It is clear that other gamble distributions can lead to very different results. One important factor is how much the gamble distributions are favorable for decisions based on small samples; if options are ‘unfriendly’ to small samples, then our previous conclusions will not hold. The gamble distributions in our stimuli and empirical study, were indeed small-sample-friendly. When averaged over 10,000 simulated pairs of gambles, we found that a sample size of 1 results in an expected rate of obtaining the H option of greater than .50 (see Fig. 1). Thus, we created an environment where a sample of size 1 using the take-good-enough, otherwise-shift rule was, on average, sufficient. However, this success rate does not generalize to any gamble environment. In Appendix D, we show that only gamble pairs where the sum of the probability of obtaining a positive outcome from option H and the probability of obtaining a negative outcome from option L is greater than 1 guarantees an expected probability of choosing the H option that is greater than .50 (see proof in Appendix D). We label gamble pairs that satisfy this condition as “one-sample favorable.” In our simulations, the proportion of gamble pairs that were one-sample favorable was .787; in these gambles, the probability of choosing option H using one sample was .688. The remaining portion of gamble pairs that were not one-sample favorable was thus .213; in these gambles, the probability of choosing option H using one sample .414. This result

highlights the fact that the accuracy of decisions based on very small samples will depend on the specific distributions encountered by agents.

One of the most important findings from early research on decisions from experience was that, in experience-based decisions, low-probability (rare) events appear to receive less impact than they deserve in light of their objective probability (Hertwig et al., 2004). This effect is, among other factors (Hertwig & Erev, 2009), caused by the fact that people do not search long enough to experience rare events often enough or at all during search. Formally, in environments where the ranks of *most* samples from each option diverge greatly from the true rank of options’ long-term average values, choosing based on a small sampling size can lead to a small probability of obtaining the H option. In other words, in gambles where small samples (e.g., a single date with a potential mate, a glance at a TV on sale) produce data that are inconsistent with an option’s long-term value (e.g., a disastrous first date with Mr. or Ms. Right, a paid celebrity endorsement of a low-quality product), frugal predecision sampling can lead people to choose poor options. For example, consider a payoff distribution that delivers +1 with probability .9 and –100 with probability .1 and thus has an expected value of –9.1. Small samples are unlikely to reveal the rare but large negative outcome of –100, making the distribution look advantageous to most agents that inspect it only briefly. This suggests that fast choosers in competitive environments involving rare events run the risk of choosing options that appear beneficial in the short term but have detrimental long-term consequences resulting from rare but impactful negative events (e.g., “black swans”; Taleb, 2007). Indeed, in such environments, a player could even benefit from competing against others who are “tricked” into grabbing options with apparent short-term gains, but actual long-term losses.

6. Conclusion

Our findings suggest that competition shifts the balance between exploration and exploitation in an uncertain choice environment: Faced with the threat of being outpaced in the process of making a decision, people dramatically reduce search. As our results show, this is a smart thing to do in ecologies in which competitors can be expected to choose quickly, and modal samples are good indicators of an option’s value. Although exploitation means forgoing the benefits of exploration that can be enjoyed in solitary situations, those who seize the first-mover advantage do better than those who do not in many (but, of course, not all) competitive situations. It is a proverbial truth that you should “look before you leap” (see also Savage, 1954/1972, p. 16). In our competitive environment, it emerged that a quick peek before leaping was very helpful—but that more extensive looking permitted the competitor to leap first and gain an edge.

Acknowledgments

This research was supported by the Swiss National Science Foundation project number 100014_129572 and by

the Israel Science Foundation grant 121/11. We thank the members of the Center for Adaptive Rationality (ARC), Arend Hintze, Julian Marewski, and Jonathan Nelson for many helpful comments. We also are grateful to Valerie M. Chase and Susannah Goss for editing the manuscript. We thank Daniel Lowengrub of the Hebrew University of Jerusalem for programming the competitive sampling game. The first author thanks D.S.P, J.D.T, and V.C.H. for inspiring him throughout the preparation of this manuscript.

Appendix A

A.1. Calculating the expected probability of obtaining option H given a choice ecology and social environment

We calculated a player’s expected probability of obtaining the option with the higher expected value (*H*) given its planned sampling size using Eq. (1):

$$p(H|n_i) = p(H|n_c > n_i) \cdot p(n_c > n_i) + p(H|n_c < n_i) \cdot p(n_c < n_i) + p(H|n_c = n_i) \cdot p(n_c = n_i) \quad (1)$$

Eq. (1) represents the weighted sum of three possible scenarios that differ with respect to the relationship between the player’s planned sampling size (n_i) and the planned sampling size of the competitor (n_c), that is, whether the player faces an opponent with a smaller, larger, or the same planned sampling size.

The first half of the first term in Eq. (1) corresponds to the probability of obtaining the *H* option given that the competitor will sample longer than the player, and, consequently, the probability that the player will obtain *H* equals that of choosing the *H* option given n_i samples. Because choices are based on sample means, this equals the probability that the order of the sample means from the two options matches the order of the population means. In this case, the option with the higher sample mean will also be the option with the higher population mean and the player will choose the better option. Formally:

$$P(H|n_c > n_i) = p(\bar{x}_H > \bar{x}_L|n_i) \quad (2)$$

Note that this calculation is specific to the choice ecology under consideration. For two outcome payoff distributions such as those used here, this can be calculated directly by comparing the results of two binomial distributions. The second half of the first term is the weight given to this outcome, defined as the probability of encountering an opponent with a larger sample size than the player’s sample size.

The remaining two terms in Eq. (1) follow the same logic. When the competitor samples less than the player, the probability that the player obtains the *H* option is the probability that the competitor will *not* choose the *H* option. This equals the probability that the opponent observes sample means

whose order is not equal to the true order of population means and can be calculated as follows:

$$p(H|n_c < n_i) = p(\bar{x}_H < \bar{x}_L|n_c) \quad (3)$$

Finally, the third term in Eq. (1) represents the expected outcome when both players have the same sampling size. This is set to .5 and is independent of the sampling distributions of payoff distributions *H* and *L* (=lower expected value distribution):

$$p(H|n_c = n_i) = .5 \quad (4)$$

Appendix B

B.1. Distributions of planned sampling sizes for competitive social environments

In the *fast* environment *F*, the probability that a randomly sampled agent has a planned sampling size n_k is given by a geometric distribution with $p = .3$, ranging from 1 to 50 and normalized to sum to 1:

$$f(F = n_k) = \frac{(1 - .3)^{n_k-1} \cdot .3}{\sum_{i=1}^{50} (1 - .3)^{i-1} \cdot .3} \quad n_k = 1, 2, \dots, 50$$

In the *slow* environment *S*, the probability that a randomly sampled agent has a planned sampling size n_k is a reflected version of *F* around the point $n_k = 25.5$:

$$f(S = n_k) = \frac{(1 - .3)^{50-n_k-1} \cdot .3}{\sum_{i=1}^{50} (1 - .3)^{i-1} \cdot .3} \quad n_k = 1, 2, \dots, 50$$

In the *uncertain* environment *U*, the probability that a randomly sampled agent has a planned sampling size n_k is given by the discrete, uniform distribution with bounds at 1 and 50:

$$f(U = n_k) = \frac{1}{50} \quad n_k = 1, 2, \dots, 50$$

In the *as-if solitary* environment *A*, the probability that a randomly sampled agent has a planned sampling size n_k is given by a negative binomial distribution with $p = .071$ and $r = 1.59$. The distribution is bounded from 1 to 50:

$$f(A = n_k) = \frac{\binom{n_k + 1.59 - 1}{n_k} (1 - .071)^{1.59} \cdot .071^{n_k}}{\sum_{i=1}^{50} \binom{1 + 1.59 - 1}{i} (1 - .071)^{1.59} \cdot .071^i} \quad n_k = 1, 2, \dots, 50$$

Appendix C

C.1. Properties of the practice games

Practice 1		Practice 2		Practice 3	
<i>H</i>	<i>L</i>	<i>H</i>	<i>L</i>	<i>H</i>	<i>L</i>
(32, .458, -13)	(39, .352, -18)	(47, .417, -25)	(42, .349, -21)	53, .407, -24)	(35, .458, -17)

C.2. How gamble parameters were selected

We started out with four binary-valued options, labeled A1–A4. Their values were (−17, 37), (−20, 43), (−23, 49), and (−26, 55), for options A1 through A4, respectively. Each of these options had three versions, differing in expected value; the expected values were high (EV = 7), medium (EV = 2), or low (EV = −3). The three EVs were obtained by modifying the probabilities of the two outcomes. Additionally, each A option had a corresponding B option, for which the two values spanned a smaller range (hence, smaller variance); the range of the B option was about 2/3 that of the corresponding A option. We then constructed 12 gamble sets, each comprising five decision tasks (displayed in Table 1). One involved a choice between an A option with an EV of 7 and a B option with an EV of −3 (e.g., A1High and B1Low); a second involved a choice between another A option and its corresponding B option, in which the A option had the low EV and the B option the high EV (e.g., A3Low and B3High); a third involved a choice between another A and B pair in which the two were both of the medium value, with EV = 2 (e.g., A2Medium and B2Medium); a fourth involved a choice between two (large variance) A options belonging to the same set but differing in value (e.g., A4High and A4Low); a fifth involved a choice between two (small variance) B options belonging to the same set but differing in value (e.g., B4High and B4Low). There were 12 such sets, and we used them all.

Appendix D

D.1. Conditions that make decision tasks one-sample favorable

Consider an environment containing two gambles (options) H and L , where $E(H) > E(L)$. Assume a player selects a gamble at random and draws a random sample. Let the random variable S represent the selected option where $S \in \{H, L\}$. Let the random variable $X \in R$ be outcome drawn be the outcome drawn from the selected gamble. Finally, let the random variable $C \in \{H, L\}$ be the chosen option.

Consider a player using a one-sample search and decision rule: (1) Select an option at random and draw one sample. (2) If the sample value is positive, choose the selected option. If the sample value is negative, chose the unselected option. From the law of total probability, the probability that player will choose option H can be written as the sum of the probabilities of two disjoint events:

$$p(C = H) = p(S = H)p(X > 0|S = H) + p(S = L)p(X < 0|S = L)$$

Because options are selected at random,
 $p(S = H) = p(S = L) = .50$:

$$p(C = H) = .5p(X > 0|S = H) + .5p(X < 0|S = L)$$

Moving terms around

$$p(C = H) = .5(p(X > 0|S = H)) + p(X < 0|S = L)$$

It follows that for $p(C = H)$ to be greater than .50, $p(X > 0|S = H) + p(X > 0|S = L)$ must be greater than 1.0.

Appendix E

E.1. Second simulation procedure

We simulated the performance of agents with varying fixed sampling sizes playing the competitive sampling game against varying numbers of competitors, number of options, and competitor speed. The key parameters we varied were: N.Players (2, 3, 4, 5, 6): the number of competitors in the game. N.Options (2, 3, 4, 5, 6): the number of options (gambles) in the game. Competition.Speed (Slow, Uniform, Fast): the decision speed of competitors. This created 72 simulation classes. For each simulation class, we simulated the decision performance of 15 agents playing the competitive sampling game, each using a fixed sample size of 1–15. We aggregated each agent's performance over 5000 stochastic factors: (1) the outcome distributions within each of the (N.Options) options and (2) the specific stopping rules of its (N.Competitors) competitors. Each option represented a discrete, two-outcome gamble with one positive and one negative outcome, each occurring with complementary probabilities. For each of the options, we drew a positive outcome from $Unif(0, 100)$ and a random negative outcome from $Unif(-100, 0)$. We then drew the probability of the positive outcome ($p+$) from $Unif(0, 1)$ and set the probability of the negative outcome ($p-$) to $1 - p+$. We constructed the probability mass function for each option independently of other options. For each of the competitors, we drew a sample size from its corresponding decision speed distribution (Slow, Uniform, or Fast). These distributions corresponded to those in Appendix B

Each game proceeded as follows: agents sampled equally from options until the first player reached its sampling size. That agent then choose the option with the highest observed sample mean (with ties broken at random). In the case where two agents stopped at the same time, one of two outcomes could occur: If they wanted different options, they each got their desired option. If they wanted the same option, then the desired option was randomly given to one agent, and the remaining agent then attempted to take is next most desired option. After all agents who stopped on that round received an option, the game continued with the remaining players and options. At the end of each game, each agent got the expected value of its chosen option. In the games where there were more players than options, if a player ends the game with no option (because all options were taken by other players), then it received a reward of 0.

Appendix F. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cognition.2014.06.006>.

References

- Brezzi, M., & Lai, T. L. (2002). Optimal learning and experimentation in bandit problems. *Journal of Economic Dynamics and Control*, 2(1), 87–108.

- Busemeyer, J. R., & Rapoport, A. (1988). Psychological models of deferred decision making. *Journal of Mathematical Psychology*, 32(2), 91–134.
- 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*, 362(1481), 933–942.
- Denrell, J. (2007). Adaptive learning and risk taking. *Psychological Review*, 114(1), 177–187.
- Dutta, P. K., & Rustichini, A. (1993). A theory of stopping time games with applications to product innovations and asset sales. *Economic Theory*, 3(4), 743–763.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75, 643–669.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112(4), 912–931.
- Gans, N., Knox, G., & Croson, R. (2007). Simple models of discrete choice and their performance in bandit experiments. *Manufacturing and Service Operations Management*, 9(4), 383–408.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B*, 41, 148–177.
- Gittins, J. C. (1989). *Multi-armed bandit allocation indices*. New York, NY: Wiley.
- Groß, R., Houston, A. I., Collins, E. J., McNamara, J. M., Dechaume-Moncharmont, F. X., & Franks, N. R. (2008). Simple learning rules to cope with changing environments. *Journal of the Royal Society Interface*, 5(27), 1193–1202.
- Hertwig, R. (in press). Decision from experience. In G. Keren & G. Wu (Eds.), *Blackwell handbook of judgment and 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.
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517–523.
- Hertwig, R., Hoffrage, U., & the ABC Research Group (2013). *Simple heuristics in a social world*. New York, NY: Oxford University Press.
- Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, 115(2), 225–237.
- Ho, T. H., Camerer, C., & Weigelt, K. (1998). Iterated dominance and iterated best response in experimental “p-beauty contests”. *The American Economic Review*, 88(4), 947–969.
- Jacobs, Andrew (2011, April 11). *For many Chinese men, no deed means no dates*. <<http://www.nytimes.com/2011/04/15/world/asia/15bachelors.html>>.
- Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford, UK: Oxford University Press.
- 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.
- Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, 124(3), 334–342.
- Nowak, M., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364(6432), 56–58.
- Park, A., & Smith, L. (2008). Caller number five and related timing games. *Theoretical Economics*, 3(2), 231–256.
- Real, L. A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science*, 253(5023), 980–986.
- Rotjan, R. D., Chabot, J. R., & Lewis, S. M. (2010). Social context of shell acquisition in *Coenobita clypeatus* hermit crabs. *Behavioral Ecology*, 21(3), 639–646.
- Savage, L. J. (1954). *The foundations of statistics* (2nd rev. ed.). New York, NY: Dover.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Taleb, N. N. (2007). *The black swan: The impact of the highly improbable*. New York, NY: Random House.
- Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, 111(2), 430–445.
- Wulff, D. U., Hills, T. T., & Hertwig, R. (2014). The impact of long- and short-run frames on search and choice in decisions from experience. Submitted for publication.