

Decisions from experience: Competitive search and choice in kind and wicked environments

Renato Frey*

Abstract

Information search is key to making decision from experience: exploration permits people to learn about the statistical properties of choice options and thus to become aware of rare but potentially momentous decision consequences. This registered report investigates whether and how people differ when making decisions from experience *in isolation* versus *under competitive pressure*, which may have important implications for choice performance in different types of choice environments: in “kind” environments without any rare and extreme events, frugal search is sufficient to identify advantageous options. Conversely, in “wicked” environments with skewed outcome distributions, rare but important events will tend to be missed in frugal search. One theoretical view is that competitive pressure encourages efficiency and may thereby boost adaptive search in different environments. An alternative and more pessimistic view is that competitive pressure triggers agency-related concerns, leading to minimal search irrespective of the choice environment, and hence to inferior choice performance. Using a sampling game, the present study ($N = 277$) found that solitary search was not adaptive to different choice environments ($M = 14$ samples), leading to a high choice performance in a kind and in a moderately wicked environment, but somewhat lower performance in an extremely wicked environment. Competitive pressure substantially reduced search irrespective of the choice environment ($M = 4$ samples), thus negatively affecting overall choice performance. Yet, from the perspective of a cost-benefit framework, frugal search may be *efficient* under competitive pressure. In sum, this report extends research on decisions from experience by adopting an ecological perspective (i.e., systematically varying different choice environments) and by introducing a cost-benefit framework to evaluate *solitary* and *competitive* search — with the latter constituting a challenging problem for people in an increasingly connected world.

Keywords: decisions from experience, choice environments, exploration, competition, competitive search

1 Introduction

When people make decisions, they often do not know the consequences of choosing one of the available options *a priori*. Instead, people typically have to first explore the possible outcomes (e.g., which outcomes can be obtained, and how likely do these occur?) and ultimately make a decision — implicitly or explicitly — based on “statistical probabilities” (Knight, 1921). For example, people engage in pre-decisional search when comparing different offers while looking for a hotel room on an online platform, or when dating potential partners to explore the local mating market (Miller & Todd, 1998).

Past research on such *decisions from experience* (Hertwig et al., 2004; Hertwig, 2015) has focused on how peo-

ple search and choose *in isolation*, investigating how pre-decisional search is influenced by various task characteristics such as the magnitude of incentives (Hau et al., 2008), the role of gains versus losses (Lejarraga et al., 2012), the variability of payoffs (Lejarraga et al., 2012; Mehlhorn et al., 2014), or choice set size (Hills et al., 2013; Frey et al., 2015a). Similarly, past research has also investigated the role of various person characteristics such as emotional states (Frey et al., 2014b), cognitive abilities (Rakow et al., 2010; Frey et al., 2015a), or age (Frey et al., 2015a; Spaniol & Wegier, 2012).

Yet, in an increasingly connected world people rarely make decisions from experience in isolation but often in the physical or virtual presence of others. That is, others might simultaneously aim to identify and choose the best out of the same limited set of choice options (Phillips et al., 2014; Schulze et al., 2015; Schulze & Newell, 2015). The presence of others may lead to challenging trade-offs between pre-decisional search (i.e., exploration) and choice (i.e., exploitation; Hills et al., 2014; Markant et al., 2019).

To study the role of competitive pressure in people’s decisions from experience, this registered report pits two theoretical accounts against each other: an optimistic view presumes that competitive pressure boosts efficiency and

This paper is a registered report.

I am grateful to the members of the Center for Cognitive and Decision Sciences (University of Basel), Gilles Dutilh, Tomás Lejaraga, and Robin Hogarth for helpful comments to a first draft of this paper, and to Laura Wiles for editing the manuscript. This work was supported by a grant of the Swiss National Science Foundation provided to R.F. (PZ00P1_174042).

Copyright: © 2020. The authors license this article under the terms of the [Creative Commons Attribution 3.0 License](https://creativecommons.org/licenses/by/3.0/).

*Behavioral Science for Policy Lab, Princeton University, USA, and Department of Psychology, University of Basel, Switzerland. Email: rfrey@princeton.edu. ORCID: 0000-0002-3190-3523

thus triggers adaptive search in different choice environments. Conversely, a more pessimistic view presumes that competitive pressure triggers minimal search irrespective of the choice environment because competition may induce agency-related concerns — that is, people might be worried about not being able to make an active choice. Consequently, competitive pressure may hamper choice performance in environments that would require ample exploration.

1.1 An ecological perspective: The critical role of pre-decisional search in different environments

Human behavior can often be evaluated meaningfully only in light of the choice environments in which people make decisions (Simon, 1956), which is why this article adopts an ecological perspective by taking into account choice environments that systematically vary regarding fundamental statistical properties. Specifically, when people explore choice options by sequentially sampling possible outcomes, search can be viewed as an inference task with the goal of learning about the available options' underlying statistical properties. People may pursue several possible goals during this process, but one frequent assumption in research on decisions from experience has been that people aim to learn about each option's average reward that can be expected in the long run (i.e., the option's expected value, EV; Mehlhorn et al., 2014; Ostwald et al., 2015; Phillips et al., 2014). Indeed, tests of the natural-mean heuristic ("the psychologically plausible pendant to the estimation of an option's EV"; Hertwig & Pleskac, 2008; Ostwald et al., 2015) and evidence from cognitive modeling analyses (Erev et al., 2010; Frey et al., 2015b,a) suggest that people tend to rely on the experienced sample means to identify and choose their preferred option.

In so doing and all else being equal, more search promises to yield more precise estimates of an option's long-run consequences, because rare but potentially momentous outcomes will be observed more likely (Poisson, 1837). Yet, whether large samples truly pay off — or, conversely, whether frugal search leads to biased representations of choice options — depends on the properties of the choice environment. A basic distinction has previously been proposed between two paradigmatic types of choice environments, bearing fundamental implications for search:

"In *kind* environments, people receive accurate and complete feedback that correctly represents the situation they face and thereby enables appropriate learning. Thus, observing outcomes in a kind environment typically leads people to reach unbiased estimates of characteristics of the process. In contrast, feedback in *wicked* environments is incomplete or missing, or systematically biased,

and does not enable the learner to acquire an accurate representation." (Hogarth & Soyer, 2011, p. 435)¹

To illustrate, Figure 1 depicts exemplary decision problems of a kind, a moderately wicked, and an extremely wicked environment. Each decision problem involves two choice options (i.e., a payoff distribution with higher and lower EV). The difference between the modes of the two choice options is identical in all environments (i.e., 2), such that a person who sequentially draws outcomes with replacement will experience a difference of about 2 between the options *most of the time*. However, in the decision problems of the wicked environments, one of the two choice options has a marked bimodal distribution, with a subset of outcomes being rare negative outliers (i.e., these distributions have a global and a local maximum, or a "mode" and an "anti-mode"; see Figure A1 for all decision problems used here, including reversed cases with rare positive outliers).

As a result, the two options of a decision problem in the wicked environments differ substantially in terms of their EVs (i.e., a difference of 40). These EV-differences are 20 times larger than the relatively small differences between the two distributions' modes. In short, whenever the differences between two options are relatively small most of the time, but dramatically different in rare occasions — as in the decision problems of wicked environments — it might be particularly profitable to search extensively and learn about the options' long-run consequences.

Figure 2 depicts this relationship, namely, how likely a higher-EV option will be chosen in a kind, a moderately wicked, and an extremely wicked environment, depending on different sample sizes ranging from 2 to 20, and as a function of different choice sensitivities (i.e., whether the option with the higher experienced sample mean is chosen with $p=1$, $p=.9$, or $p=.8$). This simulation analysis illustrates that different search efforts do not substantially influence the likelihood of choosing the higher-EV option in kind environments, whereas sample size matters substantially in this respect in the wicked environments. Strikingly, in the wicked environments very small samples will systematically misguide decision makers and (wrongly) suggest that lower-EV options are the advantageous options.

In real life, wicked environments with J-shaped distributions as depicted in Figure 1 abound, ranging from domains

¹The implications of the distinction between kind and wicked environments are naturally not limited to information search. Yet, this distinction is helpful to examine the role of search in paradigmatically different environments, and is conceptually most closely related to situations B and C depicted in Figure 1 of Hogarth et al. (2015): in kind environments, the elements of the learning setting (i.e., a small sample from a payoff distribution) are approximately the same as the elements of the target setting (i.e., the full underlying distribution). In wicked environments, in contrast, the elements of the learning setting systematically exclude elements of the target setting, namely, the rare events of the distribution that will usually be under-represented in small samples.

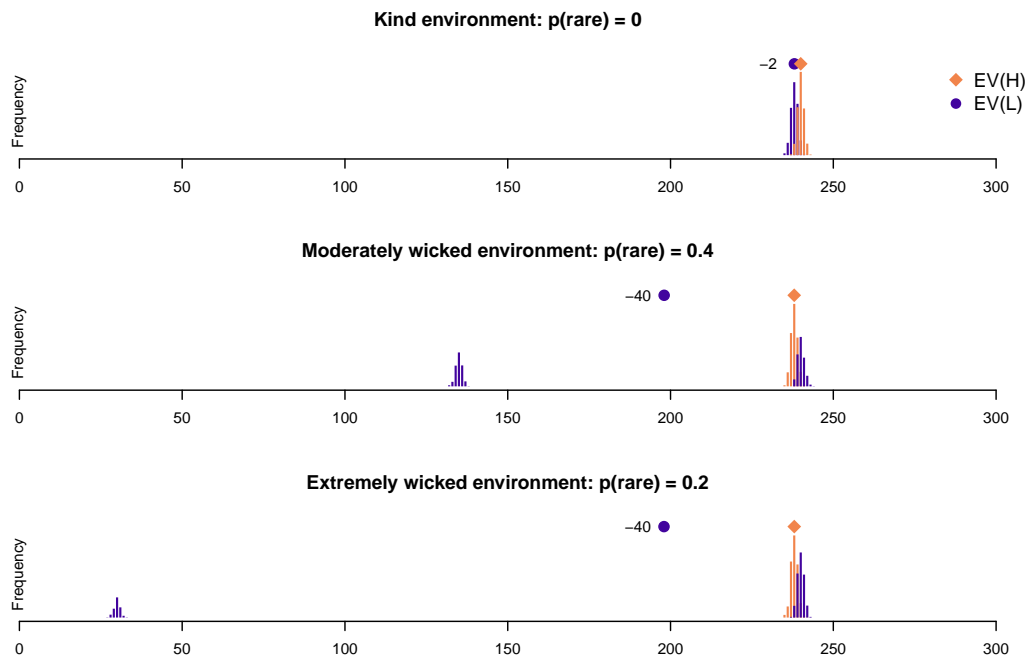


FIGURE 1: Exemplary decision problems for each of the three implemented choice environments. The diamonds and circles above the distributions depict the expected values (EV) of the higher-EV (H) and lower-EV (L) options, respectively. The numbers indicate the differences between the distributions' EVs. The full set of decision problems is depicted in Figure A1.

such as academic citation counts to product ratings: citation counts have a bimodal and heavily skewed distribution, with good publications naturally getting most citations, but exceedingly bad publications being cited more often than mediocre ones (Bornstein, 1991; Nicolaisen, 2002). Similarly, the distributions of online product reviews typically tend to be J-shaped, with many highly positive and only few very negative ratings, and almost no ratings in the middle-range of the scale (Hu et al., 2009; Wulff et al., 2014). A systematic implementation of kind and wicked environments thus promises to be a good framework for studying how adaptive people's solitary and competitive search is when making decisions from experience.

1.2 How adaptive is search without competitive pressure?

To study people's pre-decisional search, a simple sampling paradigm is often used in research on decisions from experience (Hertwig et al., 2004). In this paradigm, participants explore two payoff distributions by sampling outcomes with replacement, and only the final choice counts towards their payment. A recent meta-analysis reported that people draw, on average, 20 samples before making a final choice in this sampling paradigm, and more search has been observed when both choice options were risky (i.e., had variability in their outcomes, as in the wicked environments introduced

above) compared to when one of the two choice options was safe (i.e., had no variability in the outcomes; Wulff et al., 2017). Do people thus indeed search adaptively and adjust pre-decisional search, contingent on the choice environment?

Although no final answer has yet been provided to this question, there exist two explanations for why people might do so (Ostwald et al., 2015; Lejarraga et al., 2012; Mehlhorn et al., 2014). One possibility is that people have specific prior beliefs about a choice environment, which they update after experiencing a final outcome. A second and possibly complementary mechanism entails that people adjust search on the fly, in response to their short-term experiences during the actual sampling process. That is, the experience of variability may trigger increased search, which permits learning how frequently outliers occur. One study reported such a positive association between experienced variance and sample size — yet without providing a conclusive answer as to whether it is indeed the experience of variance that triggers more search, or whether more search leads to an increased experience of variance (Lejarraga et al., 2012) — and more recently another study has challenged the causality of this effect (Mehlhorn et al., 2014).

Taken together, still relatively little is known regarding how strongly people adjust search to paradigmatically different choice environments when making decisions from experience. Assuming that people have long-run aspirations (Wulff et al., 2015; Phillips et al., 2014) and enough oppor-

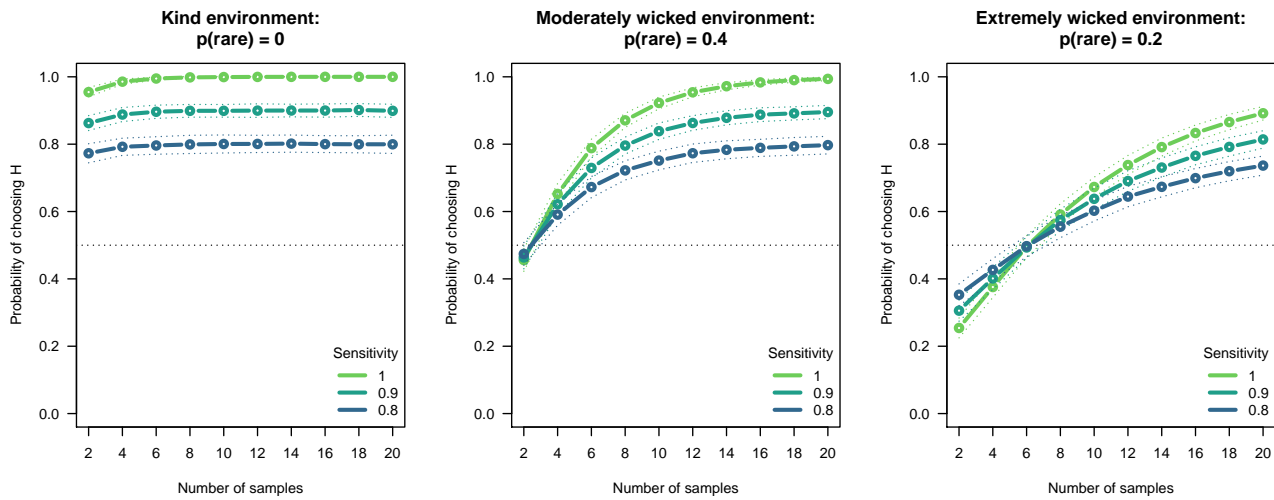


FIGURE 2: Simulation analysis for the three implemented choice environments. The curves show the predicted choice proportions of the options with the higher EV (H), based on three different sensitivities for choosing the option with the higher experienced sample mean (H_{exp}), and for different sample sizes ranging from 2 to 20. The simulation analysis was run for 1,000 experiments (each involving 30 participants), and aggregated across all eight decision problems in each choice environment (see Figure A1). Solid lines depict the average choice proportions and the dotted lines the mean proportions ± 1 SD across all 1,000 simulation runs.

tunity for exploration, one would expect that people search sufficiently to identify and choose advantageous options with growing experience. The first contribution of this article is thus to quantify the extent to which people adapt search in decisions from experience with increasing experience, as a function of different choice environments.

1.3 Does competitive pressure boost or hamper adaptive search?

The effects of competition have previously been studied extensively at the population level, and competition has adopted a foundational role in major theories of various disciplines: for example, in biology competition has been regarded a key driving force behind natural selection and thus to be a crucial element for evolution (Darwin, 1867). Specifically, according to the competitive exclusion principle, organisms less suited to competition should either adapt or die out (Hardin, 1960). Alike, in classic economic theory, competition is considered a key component of the “invisible hand” and thus the hallmark of liberal trading, based on the assumption that competitive pressure might encourage different forms of efficiency (Smith, 1776).

But how does the presence of competitors influence the search and choice behaviors of *individual persons* in decisions from experience? Using a “competitive sampling game”, Phillips et al. (2014) examined how pairs of participants explore choice options simultaneously. The first player to stop pre-decisional search became the “chooser” and could

freely pick one of the available options, whereas the other player became the “receiver” and had to accept the remaining option. Competitive pressure led to a dramatic reduction of pre-decisional search, namely from a median sample size of 18 (control condition with solitary search) to a median sample size of 1 (condition with competitive search; Phillips et al., 2014). Despite their minimal search effort choosers obtained the advantageous options (i.e., higher-EV options) in 58% of trials and thus clearly above chance level.

Although Phillips et al. (2014) employed a simulation analysis to study different “social environments” (i.e., whether competitors choose quickly or slowly), they did not investigate the role of choice environments with different statistical properties. Therefore, the mechanisms underlying the observed effects remain largely unknown: did competitive pressure make people highly *efficient* in the sense that their search adapted strongly to the statistical properties of the (kind) choice environment? Or were people simply lucky to be making decisions in a kind environment, but would have been led astray in a wicked environment?

1.3.1 The optimistic view: Competition as a driver of efficiency

The former possibility is rooted in a key assumption of standard economic theory, according to which competitive pressure is expected to lead to different types of efficiency (Smith, 1776). For example, the notion of “productive efficiency” implies that institutions will operate at the lowest point of their average-cost curve when facing competitive pressure:

that is, they will invest exactly the amount of costs that promises to lead to the best ratio of costs per unit of return. Conversely, without competitive pressure too much will be invested, a phenomena that has been labeled “x-inefficiency” (Leibenstein, 1966). Does a similar mechanism potentially also operate at the level of individual decision makers, in terms of how they explore choice options in different environments? This theoretical prediction can be tested on two different levels.

First, in most research on decisions from experience using the sampling paradigm, it is implicitly assumed that opportunity costs such as the time invested for exploration are negligible. Under this — potentially not quite realistic — assumption, competitive pressure might be the only driver for why people adjust search in different choice environments, making search costly in the sense that a competitor could choose the advantageous option first. According to this rationale, under competitive pressure people should search minimally in kind environments (Phillips et al., 2014) but increasingly more in wicked environments. Conversely, without competitive pressure people may over-sample (particularly in kind environments) as search does not entail any costs. Empirically these predictions can be tested by comparing participants’ actual search efforts with the optimal levels of search, as derived separately for the different choice environments in the simulation analysis shown in Figure 2.

Second and arguably more realistically, efficiency might be evaluated by taking into account opportunity costs such as the time allocated for pre-decisional search (which likely correlate with other potential costs, e.g., cognitive effort, and will thus be used as a proxy in the remainder of this article). As outlined above, efficiency may imply that people operate at the lowest point of their average-cost curves. These curves reflect the ratio between expected rewards and a person’s average cost of search for different sample sizes. Specifically, expected returns result directly from the simulation analysis shown in Figure 2, by multiplying the likelihoods of choosing the higher-EV option (which depend on the sample size, choice sensitivity, and crucially, the choice environment) with the relative payoff of choosing the higher- over the lower-EV option (which is +2 in the kind environment, and +40 in the wicked environments). Average costs can be measured empirically, by computing the average time (in seconds) it takes a participant to sample another outcome. Thus, as search costs may differ between individuals and search modes (e.g., solitary vs. competitive search), *efficiency* reflects an idiosyncratic measure, defined as the distance between the lowest point of a participant’s average-cost curve and the empirically observed sample size.

Only few studies have started to investigate the role of competitive pressure in people’s decisions from experience while simultaneously taking into account additional aspects of the choice ecology (Schulze et al., 2015; Schulze & Newell, 2015; Markant et al., 2019). The present approach is the first

to systematically vary the statistical properties of choice environments, thus specifically permitting to test the prediction that competitive pressure is a driver of efficiency. According to this optimistic view, competitive pressure should result in smaller differences between people’s empirical search efforts and the sample sizes implied by (the lowest point of) their idiosyncratic average-cost curves, relative to these differences in the conditions without competitive pressure.

1.3.2 The pessimistic view: Competition as a threat of agency

There also exists a more pessimistic view on the role of competitive pressure in people’s decisions from experience, which provides an alternative explanation for the observation made by Phillips et al. (2014): specifically, competitive pressure may trigger agency-related concerns because it implies a threat to people’s choice autonomy (Bandura, 2006; Moore, 2016; Leotti et al., 2010). As a consequence, people may search minimally to retain the possibility of an active choice — irrespective of the choice environment — which should drastically affect choice performance in wicked environments (see Figure 2).

This prediction is plausible given the empirical evidence from various domains suggesting that people strongly cherish agency and choice autonomy: in medicine, for example, some physicians are reluctant to use clinical decision support systems — despite these systems often increasing diagnostic accuracy (Dawes et al., 1989) — because they tend to be perceived as a threat to professional autonomy (Walter & Lopez, 2008). Similarly, relatives of incapacitated patients value having a voice and making a surrogate decision within the family, rather than delegating such a difficult decision to a physician or to a statistical prediction rule (Frey et al., 2014a, 2018). These findings also resonate with recent observations of algorithm aversion in the context of human versus statistical forecasting (Dietvorst et al., 2015).

Taken together, according to this rather pessimistic view competitive pressure will trigger minimal search across the board, whereas people may search and choose adaptively in the absence of any competitive pressure. Competitive pressure should therefore not affect choice performance negatively in kind environments (Phillips et al., 2014), yet drastically so in wicked environments.

1.4 Overview and research questions

This registered report adopts an ecological perspective to study whether and how people’s decisions from experience differ during solitary and competitive search. To this end, an adapted version of a competitive sampling game (Phillips et al., 2014) was employed, with different choice environments that systematically vary in terms of whether they are *kind* or *wicked* (in different degrees; see Figs. A1 and A2,

and Table A1). In each decision problem, participants were tasked with exploring two choice options before making a final choice between them. After each draw, participants had to indicate whether they prefer to continue exploration (i.e., draw another sample) or to make a final choice. In the solitary condition, participants could draw samples from the two payoff distributions for as long as they liked, prior to making a final choice between them. In the competitive condition, players explored the available options individually yet simultaneously in pairs of two (i.e., at the same rate), and the first player to stop information search became the “chooser”, with the opportunity to freely make a final choice between the two payoff distributions. The other player (who opted to continue information search) became the “receiver”, and was forced to accept the remaining option — much like gathering information about two unfamiliar products too extensively, and then having to accept the only option that is left available. If both participants indicated readiness to make a final choice at the same step, the two options were allocated randomly between them. Finally, to account for the fact that people often compete against anonymous and novel competitors when exploring choice options (e.g., searching for a hotel room on an online platform), participants were paired with new, anonymous competitors after each decision problem.

Research question 1. The first research question pertains to how adaptive “solitary search” (i.e., search in the absence of any competitive pressure) is in different choice environments. To this end, one third of the participants played solitary trials only (eight trials). At an absolute level and based on the assumption of adaptive search, these participants should search more in the wicked environments (particularly so in the extremely wicked environment) as compared to in the kind environment, in which frugal search is sufficient to choose higher-EV options due to the lack of outliers (Prediction 1a). As a corollary and assuming at least somewhat adaptive search in the absence of competitive pressure, choice performance should not vary substantially across the different choice environments (Prediction 1b).

Research question 2. The second and main research question pertains to whether and how strongly competition leads to increased “efficiency” or, alternatively, whether competitive pressure may trump the potential adaptive effect of solitary search. According to the former view, competitive pressure should lead to more efficient search relative to solitary search, implying that the differences between the optimal levels of search (either not taking vs. taking search costs into account) and participants’ actual search efforts are smaller under competitive pressure — as opposed to when participants search in isolation, where they may tend to oversample (Prediction 2a). As a corollary, choice performance

should not differ substantially between the two search modes (Prediction 2b).

Alternatively, however, if the past observations of minimal search (Phillips et al., 2014) were in fact a backfiring effect of competitive pressure rather than a manifestation of efficiency, competitive search will be minimal in all choice environments, whereas solitary search may decrease across trials in kind environments, but increase in wicked environments with participants’ growing experience (Prediction 3a). As a consequence, under competitive pressure choice performance should be substantially inferior in the wicked environments relative to the choice performance in the solitary condition (Prediction 3b).

2 Methods

The main study of this registered report was designed based on the insights gained from a set of pilot studies, which were conducted in advance. The stage-I registration of this registered report² including the theoretical rationale and introduction, full methods section, prospective design analysis, and results from the pilot studies) can be retrieved from <https://osf.io/5vs83/>.

2.1 Participants and inclusion criteria

Participants were recruited from Amazon mTurk ($N = 277$; see Table A2 for sociodemographic information).³ Only participants with 500 or more completed HITs (human intelligence tasks) and at least 99% approval rating were selected for the study. Moreover, participants who a) did not successfully complete two instructional manipulation checks or b) reported that they were not strongly focused during the study (i.e., a rating of 25 or lower on a scale from 0 to 100) were removed from the dataset. Recruitment continued until reaching the aspired sample size (see stage-I registration for the exact sampling plan and an overview of the experimental design).

2.2 Experimental design

Participants were randomly assigned to one of three between-subjects conditions of the choice environment: “kind”,

²The registered report format has been advocated by the Center for Open Science and others (pro, 2017; Munafò et al., 2017; Chambers, 2013; Center for Open Science, 2017), because it promises to help in avoiding hindsight bias, p-hacking (Simmons et al., 2011), and HARKing (Kerr, 1998). Furthermore, preregistration may foster an improved use of theory and stronger research methods, as well as a decline in false-positive publications (Gonzales & Cunningham, 2015).

³Recent work has demonstrated that mTurk yields data comparable in quality to that of traditional lab data (Crump et al., 2013; Paolacci, Gabriele & Chandler, Jesse, 2014), and mTurk participants have been found to pay even more attention to a study’s instructions than participants of traditional subject pools (Hauser & Schwarz, 2016).

“moderately wicked”, or “extremely wicked” (see section “Choice environments” in the Appendix). Furthermore, participants were randomly assigned to either the “solitary” or the “competitive” condition — in the competitive condition, participants were randomly paired with another participant after each decision. problem — and played eight solitary or eight competitive trials. In each time slot of the study (see “Procedure” below), the decision problems were presented in a new randomized order.

2.3 Incentive structure

Participants received a fixed compensation of 1 USD. In addition, they earned a performance-contingent bonus payment: to motivate choices of the higher-EV options (i.e., to trigger long-run aspirations), the average of 100 randomly drawn outcomes from the chosen option counted as participants’ score in each trial. This incentive scheme follows the procedure of Wulff et al. (2015) and is also similar to that used by Phillips et al. (2014). At the end of the study, two of the eight trials were randomly selected, and the sum of the respective scores constituted participants’ bonus payment.

2.4 Procedure

Upon accepting a HIT on mTurk, participants could freely choose one of the available time slots. The study started at exactly the same time for all participants of a time slot. Before beginning with the actual study, participants were informed about the general procedures (e.g., that they cannot interrupt and restart the study; that they can only participate once, which was enforced by verifying mTurk worker IDs) and that they must not be distracted during the study. Also, participants learned that they can earn an additional bonus payment between 0 and 6 USD, depending on their choice performance.

If participants accepted these conditions and provided informed consent, they read the instructions of the study (see Appendix) and played one practice trial. Moreover participants in the competitive condition were informed that they will be paired with one of the other players after each trial. After having completed the eight trials, participants provided sociodemographic information, reported how focused they were during the study, and responded to the following four questions (participants in the solitary condition only answered the first two of these questions) using a continuous slider that yielded values from 0 to 100: “During the decisions that you made in this study... i) how important was it to you to choose the option with the highest average outcome? ii) how important was it to you to choose the option with the maximum outcome? iii) how important was it to you to be able to choose an option ahead of the other player? iv) if there was a trade-off, would it be more important to you to choose the “better” option, or to make a choice ahead

of the other player?” Finally, participants were shown the outcomes of their choices in the eight trials, out of which two were selected and displayed as participants’ final bonus.

2.5 Simulation analyses

Prior to conducting the study, a simulation analysis was run to assess how different search efforts should affect choice performance in the different choice environments (see Appendix for a detailed description of how the decision problems of the different environments were generated). Specifically, this analysis simulated 1,000 experiments for each of the three choice environments, each involving 30 players (i.e., the aspired sample size per condition). These players were simulated to sample between 2 (i.e., once per option) and 20 times, in steps of two, and from all eight decision problems. Based on the “experienced” samples, the players were simulated to choose the option with the higher experienced sample mean (H_{exp}) with three different choice sensitivities, namely, 100%, 90%, and 80% — all plausible values according to the results observed in the pilot studies. Finally, the analysis determined the probability with which these simulated players would choose the option with the higher EV (H); that is, the criterion participants were incentivized for.

As Figure 2 shows, sample size does not affect the probability of choosing H in the kind environment, and the different choice sensitivities are directly reflected in the resulting probabilities after only about 4 samples. In the wicked environments, in contrast, the likelihood of choosing H-options strongly depends on sample size, and the more extreme the wicked environment is, the larger the sample sizes required to achieve a probability of greater than 50% for successfully choosing H-options. Thus in the wicked environments (particularly so in the extremely wicked environment), frugal search may make a critical difference for choice performance.

2.6 Prospective design analysis

An extensive prospective design analysis (i.e., “Bayesian power analysis”) has been conducted to make sure that the aspired sample size will provide conclusive evidence given the proposed experimental design. The details of this analysis are reported in the stage-I registration and can be retrieved from <https://osf.io/5vs83/>.

2.7 Analysis plan

2.7.1 Main analyses

To address the main research questions, four separate Bayesian mixed-effects regression models were estimated: sample size (i.e., “search effort” ignoring opportunity costs; not to be confused with N , the number of participants) was the dependent variable (DV) in the first model (as sample size represents count data, a Poisson distribution with an

identity link function was used). The second model was analogous yet used a different DV to test the efficiency of search — specifically, the differences between the observed search efforts and the lowest point on participants' average-cost curves (which were determined individually for each participant, see introduction). In the third model, the DV was whether the option with the higher experienced sample mean (Hexp-choice) was chosen (to evaluate “choice sensitivity”). Finally, in the fourth model the DV was whether the option with the higher EV (H-choice) was chosen (to evaluate “choice performance”). Hexp-choice and H-choice are binary variables, thus a logistic distribution with a logit link function was used.

The fixed effects were “search mode” (with the reference level “solitary search” and the effect level “competitive search”), “choice environment” (with the reference level “kind environment” and the effect levels “moderately wicked environment” and “extremely wicked environment”), and “trial index” (1–8; to account for potential sequence effects). Due to the repeated-measurement design, random effects across participants were implemented (i.e., random intercepts and random slopes for trial index, to keep the model “maximal”; Barr et al., 2013). One of the several advantages of the mixed-effects models is that all effects can be estimated robustly, despite that the number of trials in the analysis may vary between participants (in the “competitive trials” the data of “choosers” and “receivers” are inversely redundant, therefore only the data of “choosers” were analyzed). Two-way interactions between search mode and the environment were estimated to examine differences in search and choice as a function of the different conditions. Moreover, for the DV “sample size” three-way interactions with trial index were implemented, to examine a potential effect of adaptive search with increasing experience (i.e., whether search unfolds differentially with increasing experience in the different environments⁴).

2.7.2 Complementary analyses

To evaluate the potential motivations underlying different search strategies, the posterior distributions of the responses to questions i–iv (described in “Procedure” above) were modeled using four Bayesian linear models with a Gaussian distribution. For questions i) and ii) the posterior means were compared across the two conditions (solitary vs. competitive search). For questions iii) and iv) the distributions were modeled merely using an intercept, as these questions were only provided to participants in the competitive condition.

A final analysis in the competitive condition tested whether previous “receivers” may reduce search effort in

⁴Note that the prospective design analysis included trial number as a covariate, but as it focused on the main effects of search mode and the environment and their interactions, no two- or three-way interactions with trial number were implemented.

subsequent trials, in order to increase the chance of becoming “choosers” themselves. To this end, a separate Bayesian mixed-effects model was implemented with the DV “chooser” (binary: yes, no), the fixed effects “was receiver in the previous trial” (no, yes) and “trial index” (1–8), and random intercepts for participants (as the DV is a binary variable, a logistic distribution with a logit link function was used).

All models used the weakly informative default priors as implemented in the R-package *rstanarm* (Stan Development Team, 2016); namely, $\mathcal{N}(0, 10)$ for the intercept and $\mathcal{N}(0, 2.5)$ for the predictors. Weakly informative priors provide some statistical regularization and thus guard against overfitting the data. Three chains with 2,000 iterations were run per model. The medians of the posterior distributions are reported as a measure of central tendency, along with the 95% highest-density intervals (HDI) of the posterior distributions.

2.8 Open data and open code

The entire dataset and the analysis scripts are available from <https://osf.io/5vs83/>.

3 Results

3.1 Search

3.1.1 Sample size

In the reference condition (solitary search in the kind environment), participants on average sampled 13.7 times (HDI: 12.0 – 15.2) before making a final choice (Figure 3 and Table A3). As can be seen in Figure 3, there were no indications for adaptive search in the solitary mode, as participants did not sample credibly more in the moderately wicked (14.8 samples [HDI: 13.1 – 16.5]) and in the extremely wicked (14.3 samples [HDI: 12.7 – 16.2]) environments.

Yet, there was a marked effect of competitive pressure on sample size: in the kind environment, there was a credible reduction of –9.3 samples (HDI: –11.2 – –7.4), leading to an average sample size of 4.3 [HDI: 3.2 – 5.7]. As in the solitary mode, the sample sizes in the moderately and extremely wicked environments were not credibly different from the sample sizes observed in the kind environment (Table A3). Yet, compared to the sample sizes observed for solitary search (i.e., reference condition), the reductions remained highly credible.

Finally, there was a weak but credible effect of increasing experience, with a reduction of –.5 samples [HDI: –.6 – –.3] in each additional trial. There were no credible two- or three-way interactions with trial, suggesting that with increasing experience participants did not differentially adjust search as a function of the different choice environments or




FIGURE 3: Sample size as a function of search mode (solitary vs. competitive), and separately for the three different choice environments. Histograms (i.e., the horizontal bars) depict the average sample sizes per participant (i.e., aggregated across the eight trials). The horizontal lines depict the posterior means for each combination of environment and search mode (resulting from the Bayesian mixed-effects models), with the red shaded areas representing 95% highest density intervals (HDI).

search modes. These interactions were thus omitted in the subsequent models.

3.1.2 Search (in)-efficiency

The second analysis concerned participants’ search efficiency, defined as the distance between a participant’s actual sample size in a given trial and the optimal (i.e., lowest) point of a participant’s idiosyncratic average-cost curve (Figure 4). As such, larger values reflect increasing search *inefficiency*. Average-cost curves were determined by i) computing the expected rewards of different sample sizes (from 2 to 20) in the participant’s choice environment,⁵ and ii) computing the average costs (in terms of seconds required for each sample) per expected reward, separately for the different sample sizes.

The average-cost curves of each participant are depicted in Figure 4, along with the mean curves across participants in the various conditions (depicted in black, with black circles indicating the different sample sizes). As can be seen in this figure, in the competitive condition the costs for search were generally higher (i.e., higher elevation of the curves in all three environments) because exploration tended to be slower due to the synchronization of search between pairs of participants. In the kind environment, the maximum expected reward (+2 when choosing H over L) could be realized after only 2 samples, and more search merely resulted in additional costs. For this reason, the average-cost curves formed

vertical lines in the kind environment (depicted in green). In the moderately wicked (orange) and extremely wicked (red) environments, larger sample sizes monotonically increased the expected rewards (up to the maximum expected reward of +40 when choosing H over L).⁶ As to be expected, the average costs for obtaining the same expected rewards were relatively higher in the extremely as opposed to in the moderately wicked environment (higher elevation of curves). Crucially, in the solitary mode the average-cost curves declined with increasing sample size up to a specific sample size — indicating that up to a certain point, increasing search paid off in terms of larger expected rewards. Conversely, in the competitive mode there was virtually no such decline in the average-cost curves. Instead, the curves were essentially flat in the range of small sample sizes, implying that sampling more than twice hardly pays off in these environments — due to the relatively higher search costs in the competitive condition.

Table A3 denotes the changes in search inefficiency as a function of the different conditions. In the reference condition (solitary search in the kind environment), participants on average over-sampled 12.5 times (HDI: 10.8 – 14.3), relative to the lowest point of their idiosyncratic average-cost curves. In the solitary condition, participants’ search only started to pay off in the extremely wicked environment, where search inefficiency credibly decreased by –3.8 (HDI: –6.3 – –1.4). Conversely, search inefficiency was substantially smaller in

⁵The vast majority of participants had a choice sensitivity close to 1 (see Figure A4), which is why for reasons of comparability this analysis was conducted with a sensitivity of 1 for all participants.

⁶Note that for reasons of consistency, this analysis only covered the range of up to 20 samples. Therefore, in the extremely wicked environment the curves do not span up to the maximum reward of +40, as even larger sample sizes would have been required to reliably reach this level.

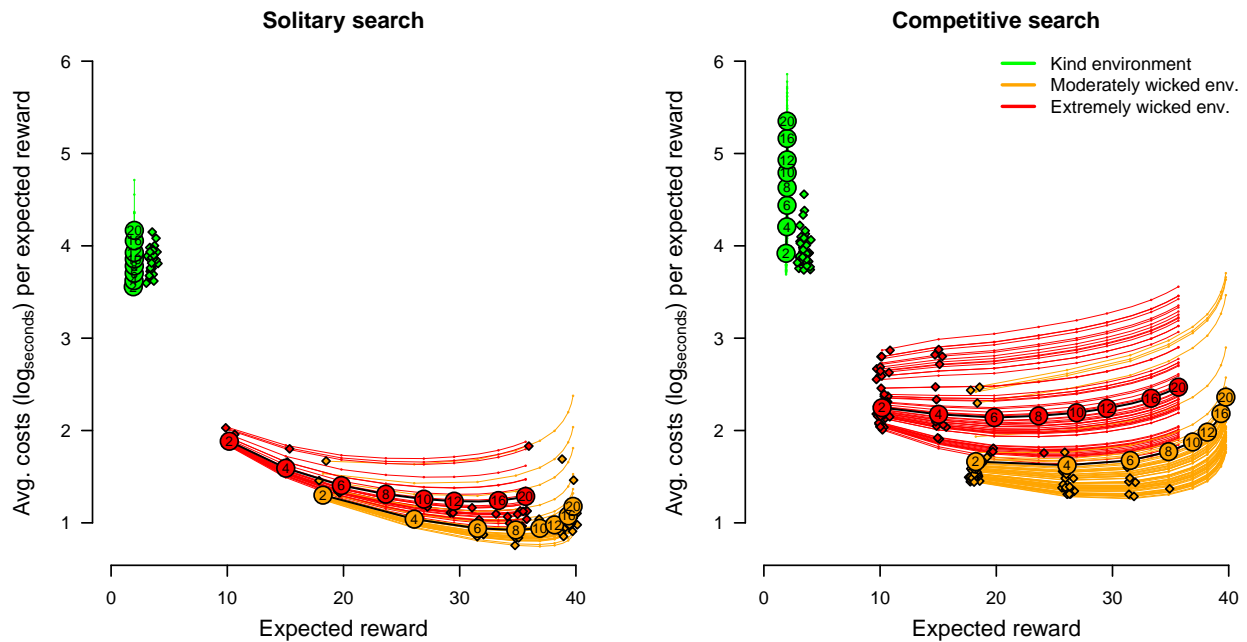


FIGURE 4: Average-cost curves of all participants, separately for the three environments and solitary and competitive search. The expected rewards for different sample sizes (up to 20 samples) were determined based on the simulation analysis displayed in Figure 2. Average costs per expected reward were computed based on participants' actual costs in terms of the time needed for each sample. The black circles with numbers depict different sample sizes, and the black lines the mean average-cost curve of participants in the respective experimental condition. In the kind environment (green lines), the curves form vertical lines because increasing search only resulted in additional costs, but did not increase expected rewards. Conversely, in the moderately (orange) and extremely (red) wicked environments, increasing search led to higher expected rewards. In the solitary condition, these curves initially declined for small sample sizes, indicating that the ratio between search costs and expected rewards improved up to a certain point (i.e., lowest point on each participant's curve). In contrast, in the competitive condition there was virtually no such decline, indicating that the relatively higher search costs for extensive search were not outweighed by the increases in expected rewards. Diamonds represent participants' positions on their idiosyncratic average-cost curves (small jitter added, particularly so for the kind environment).

the competitive mode: relative to the reference condition, search inefficiency decreased by -9.5 (HDI: $-11.6 - -7.2$) in the kind environment, by -10.5 (HDI: $-12.8 - -8.4$) in the moderately wicked environment, and by -11.3 (HDI: $-13.4 - -9.1$) in the extremely wicked environment. These patterns are also reflected in Figure 4, where each participant is depicted by a small diamond. In the solitary condition, most participants did not cluster at the lowest points of their idiosyncratic curves in all three environments, whereas most participants in the competitive condition did so. To illustrate the former, in the kind environment most participants in the solitary condition (green diamonds in the left panel) searched substantially more than the lowest point of their curve (i.e., 2 samples) would imply (“x-inefficiency”). Finally, the slight reduction in search with increasing experience (i.e., higher trial number; see previous section) also resulted in a slight reduction in search inefficiency of -0.4 (HDI: $-0.5 - -0.2$) in each additional trial.

3.1.3 Motivation for different search strategies

The third analysis of search explored participants' potential motivations for searching either extensively or frugally; and in particular the role of agency-related concerns participants might have during competitive search (i.e., lack of possibility to make an active choice). As Figure A3 illustrates, overall participants reported that it was highly important to them to choose the option with the higher average reward (i.e., the criterion used to determine their final payoff). Yet, participants in the solitary condition provided a credibly higher mean rating (94.4 [HDI: $91.8 - 97.0$]) than participants in the competitive condition (89.7 [HDI: $87.7 - 91.5$]), suggesting that the latter may have pursued at least some additional goals other than maximizing their payoffs. Similarly, participants in the solitary condition rated it to be more important (91.0 [HDI: $86.9 - 94.6$]) than participants in the competitive condition (87.7 [HDI: $85.1 - 90.3$]) to choose the option with the highest maximal outcome.

In the competitive condition, participants on average considered it fairly important to be able to choose ahead of the other player (70.1 [HDI: 66.1 – 74.6]). However, when having to make a trade-off between being able to choose the “better” option (low values in Figure A3) and being able to choose ahead of the player (high values in Figure A3), participants on average considered it more important to choose the more advantageous option (30.5 [HDI: 25.4 – 35.3]).

3.2 Choice

To analyze participants’ “choice sensitivity” and “choice performance” in the competitive mode, only the data of the “choosers” were included. The choices of the “receivers” as well as all random allocations were excluded from these analyses, because otherwise all choice proportions would converge to 50% by definition, as the data of choosers and receivers are inversely redundant.

3.2.1 Choice sensitivity

The first analysis examined participants’ choice sensitivity, that is, how frequently participants chose the option with the higher experienced sample mean (i.e., Hexp-options). In the reference condition (solitary search in the kind environment), participants chose the Hexp-options in 91% of cases (HDI: .87 – .95; Figure A4 and Table A3). As can be seen in Figure A4, for participants in the solitary condition the level of choice sensitivity did not credibly differ in the moderately and in the extremely wicked environments.

Under competitive pressure, choice sensitivity declined somewhat to 89% (HDI: 83% – 94%) in the kind environment, yet this difference was not credible (Table A3). Alike, in the moderately wicked environment there was no credible decline in choice sensitivity relative to the reference condition (-4% [HDI: -10% – 3%]), but in the extremely wicked environment choice sensitivity credibly declined to 78% (HDI: 70% – 86%). Finally, with increasing experience (i.e., higher trial number) choice sensitivity neither increased nor decreased (Table A3).

3.2.2 Choice performance

The second analysis examined choice performance, in terms of how frequently participants chose the option with the higher expected value (i.e., H-options; this is the criterion that mattered for participants’ final bonus payment, see “Methods” section above). In the reference condition (solitary search in the kind environment), participants chose the H-options in 90% of cases (HDI: 85% – 95%; Figure 5 and Table A3). In the solitary condition, there was no credible difference in choice performance in the moderately wicked environment (89% [84% – 94%]), yet choice performance credibly declined to 74% (HDI: 67% – 80%) in the extremely wicked environment.

In the competitive condition, choice performance dropped by 7 percentage points (HDI: -14 – -1) to 83% (HDI: 77% – 89%) in the kind environment. In the moderately wicked environment, the decline relative to the reference condition consisted of -16 percentage points (HDI: -24 – -9), resulting in a choice performance of 74% (HDI: 67% – 80%); and in the extremely wicked environment, this decline consisted of -28 percentage points (HDI: -37 – -21), resulting in a choice performance of 62% (HDI: 55% – 70%). Finally, increasing experience (i.e., higher trial number) had no credible effect on choice performance (Table A3).

3.2.3 Interdependencies between competitors’ behaviors

The third and final analysis on choice examined whether participants were responsive to the behavior of their competitor in the previous trial — potentially reducing search in subsequent trials to become choosers themselves. Indeed, the probability of becoming a “chooser” increased by 13 percentage points (HDI: 5 – 26) if a participant was not the chooser in the previous trial. This result thus corroborates the conclusions obtained from participants’ self-reports of their (search and choice) motivations reported above, namely, that they perceived some benefit of being able to make an active choice.

4 Discussion

This registered report employed a sampling game to study the role of competitive pressure in people’s decisions from experience. To this end, the statistical properties of three choice environments were systematically varied, in line with an ecological perspective presuming that human behavior can only be evaluated meaningfully in light of the respective choice ecology. Beyond replicating previous findings (e.g., Phillips et al., 2014), this study made a series of novel contributions concerning how people search and choose when making decisions from experience — with and without competitive pressure.

4.1 How adaptive is search without competitive pressure?

A first and basic question of this article concerned whether solitary search is adaptive to different choice environments. That is, to what extent is pre-decisional search more extensive in environments in which this truly pays off? To date, this question has rarely been addressed in research on decisions from experience. By implementing a kind, a moderately wicked, and an extremely wicked choice environment, the present study has found no indications that people




FIGURE 5: Proportion of choices of the options with higher EV (H), as a function of search mode (solitary vs. competitive), and separately for the three different choice environments. Histograms (i.e., the horizontal bars) depict the proportion of H-choices aggregated across trials (i.e., one value per participant). The horizontal lines depict the posterior means for each combination of environment and search mode (resulting from the Bayesian mixed-effects models), with the red shaded areas representing 95% highest density intervals (HDI).

adapt their search to the statistical properties of the decision problems that they encounter (Figure 3). Specifically, participants’ search effort of about 14 samples per decision problem did not systematically differ across the three choice environments.

This observation is at odds with the hypothesis that people adjust search based on the degree of variance that they experience during exploration. For example, Lejarraga et al. (2012) have tested whether the experience of variance triggers increased search, and reported evidence supporting this idea. Yet, this analysis only distinguished between “variance experienced” and “no variance experienced”, and was based on decision problems that were not implemented to systematically differ concerning their variance. In contrast, the present study has experimentally varied three choice environments, as a result of which participants experienced substantially different degrees of variance (Figure A5). Yet, the correlation between experienced variance and sample size was small ($r = .15$) in the solitary condition, and virtually non-existent in the competitive condition ($r = .06$). Moreover, the degree of experienced variance evidently did not trigger different search efforts in the three choice environments. In sum, there was no evidence for the hypothesis that solitary search adapts to (the statistical properties of) different choice environments (prediction 1a).

One interpretation of this finding is that participants tended to err on the side of caution, aiming to explore the choice environment thoroughly irrespective of the potential search costs, at least up to a certain point (see also next section). In fact, given participants’ high level of choice sen-

sitivity, they over-sampled relative to what would have been required to identify the higher-EV options in the kind environment (Figure 2) — resulting in a high choice performance in the kind and in the moderately wicked environment. Yet, choice performance dropped in the extremely wicked environment, where participants tended to under-sample. In sum, largely due to the lack of adaptive solitary search, choice performance did vary across the different environments, thus not supporting prediction 1b (i.e., no substantial differences in choice performance between the different choice environments during solitary search).

4.2 Does competitive pressure boost or hamper adaptive search?

The lack of adaptive solitary search may have partly resulted because there was no driving force rendering participants’ search *efficient*. That is, as there is no risk of a competitor making a choice first, search might have been inefficient at an absolute level (i.e., ignoring any search costs), and at least in the kind environment exceeded the number of samples required to identify the advantageous options (Figure 2). Moreover, solitary search might have been inefficient also in the sense that participants invested too much search cost (e.g., the time required to sample outcomes; Figure 4) relative to the marginal increase in expected reward associated with more search (“x-inefficiency”; Leibenstein, 1966).

Did competitive pressure make people’s pre-decisional search more efficient, as assumed by the optimistic view (prediction 2a)? At an absolute level, this would imply that

participants draw no more samples as are required to identify the higher-EV options. This did not turn out to be the case; instead, search effort did not vary across the three choice environments (as in the solitary condition), and except for the kind environment, was too low to reliably identify the advantageous options (Figure 2). However, when taking individual search costs into account, there was a different picture. Specifically, as figure 4 shows, participants in the competitive condition tended to be much closer to the lowest points of their idiosyncratic average-cost curves, as opposed to participants in the solitary condition — thus supporting prediction 2a. To illustrate, in the kind environment there was clear indication for a reduction in search inefficiency under competitive pressure: most participants only sampled twice — that is, the lowest point on the curves for both the solitary and the competitive condition — whereas participants in the solitary condition over-sampled substantially relative to this point. Similarly, in the moderately wicked environment many participants sampled only 2 or 4 times under competitive pressure — which again tended to be the lowest points of their average-cost curves. The small samples implied by these optimal points indicate that the marginal increase in expected rewards beyond this search effort was too small, given the respective costs for additional search.

In light of the relatively high search costs for the participants in the competitive condition, and the average-cost curves that consequently emerged in this study, the observation of minimal search under competitive pressure may be interpreted as a sign of high efficiency (see previous paragraph). Yet, the fact that search was minimal across all three environments is also in line with the more pessimistic view, predicting competitive pressure to lead to minimal search irrespective of the choice environment — such as because of agency-related concerns (i.e., prediction 3a). Indeed, the minimal search effort as observed under competitive pressure resulted in a substantially lower choice performance compared to that in the solitary condition, thus invalidating prediction 2b and instead supporting prediction 3b (i.e., that competitive pressure degrades choice performance in decisions from experience across the board).

Finally, participants' self-reports concerning their motivations and goals when performing the task have provided some additional support for the more pessimistic view on the effects of competitive pressure. Although participants reported that they considered it highly important to identify and choose the option with the higher average payoff (i.e., the criterion used to determine their final bonus payment), in the competitive condition a substantial number of participants also considered it fairly important to make a choice ahead of the other player — which may particularly backfire after frugal search in the wicked environments. This finding is in line with earlier research demonstrating that people cherish choice autonomy (Bandura, 2006; Moore, 2016; Leotti et al., 2010), and a lack thereof (e.g., due to competitive

pressure) may be perceived as aversive — irrespective of whether making a final choice after frugal search consists of an advantage (kind environment) or a disadvantage (wicked environments).

4.3 Limitations and further research

This study has introduced an ecological perspective to studying decisions from experience (i.e., by evaluating search and choice in paradigmatically different choice environments), as well as a cost-benefit framework taking into account the costs of pre-decisional search. Although these innovations have resulted in several important insights, the study naturally had some limitations, which should be addressed in future research.

First, the study focused on gains only (i.e., as when people research products online before making a buying decision, thus in principle hoping to obtain a “positive outcome”). As past research (e.g., Wulff et al., 2017) has found that people tend to search more extensively in the loss domain, future research should thus test whether competitive pressure also reduces search effort to a similar degree in contexts of losses.

Second, participants played the task only within one search mode (i.e., either solitary or competitive search). In some decision contexts, people may be able to engage in both search modes alternately. That is, solitary search (e.g., searching for hotel room well in advance) and the associated insights concerning a choice environment may systematically bolster people for subsequent search under competitive pressure (i.e., when the demand for hotel rooms increases). Thus, it would be worthwhile studying to what extent people may be able to transfer their experience about choice environments from solitary to competitive decisions from experience.

Third, in the present study the costs of competitive search were higher as compared to the costs of solitary search. This is automatically taken into account in participants' idiosyncratic average-cost curves, and also resembles many real-life situations (e.g., sequentially exploring options may be more arduous when having to wait for one's competitors to make decisions, as opposed to when being able to explore choice options in isolation). Nevertheless, testing a setting in which search costs do not differ between solitary and competitive search promises to lead to interesting predictions in the context of the proposed cost-benefit framework.

Fourth and finally, the game in the present study employed only two options and (in the competitive mode) two players, implying that not making an active choice ahead of the other player forces the “receiver” to accept the only option that is left available. As in many real-life settings, it may be exactly this combination of a limited choice set (i.e., “only one room left at this price”) and the presence of competitors that triggers agency-related concerns. Future research could further examine how competitive search unfolds in

other configurations, such as when more choice options are available Markant et al. (2019) as well as when more competitors are present — which would, however, substantially complicate the analysis of search and choice in the context of a cost-benefit framework.

4.4 Conclusions

People make many decisions that require a prior exploration of the possible outcomes that can be obtained from different choice options, as well as an (implicit or explicit) estimation of how frequently different outcomes occur. Past research on such *decisions from experience* has mostly focused on solitary search and typically did not systematically take into account essential aspects of the choice ecology, such as its variability (e.g., whether people explore decision problems in kind or wicked environments).

The current article thus contributes to this literature by studying solitary *and* competitive search and by adopting an ecological perspective, which may inspire future research on decisions from experience to evaluate search and choice in more nuanced ways (e.g., by means of the proposed cost-benefit framework). Taken together, this registered report has resulted in the following four main findings.

First, solitary search was not adaptive to different choice environments, implying that participants did not explore more extensively in decision problems that would have required so for participants to make advantageous decisions. Second and relatedly, although participants' search effort was sufficient to make advantageous choices in kind and moderately wicked environments, choice performance decreased in an extremely wicked environment — characterized by decision problems with rare but high-impact consequences. Third, across all choice environments competitive pressure substantially reduced pre-decisional search to very small sample sizes. Fourth and finally, although frugal search under competitive pressure may be *efficient* from the perspective of a cost-benefit framework, it led to substantially inferior choice performance as compared to the choice performance found for solitary search. This observation suggests that under competitive pressure people may at least in part pursue other goals than simply maximizing their monetary payoffs, such as maintaining choice autonomy and thus retaining the possibility of making an active choice.

References

- Bandura, A. (2006). Toward a psychology of human agency. *Perspectives on Psychological Science*, 1(2), 164–180, <https://doi.org/10.1111/j.1745-6916.2006.00011.x>.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278, <https://doi.org/10.1016/j.jml.2012.11.001>.
- Bornstein, R. F. (1991). The predictive validity of peer review: A Neglected issue. *Behavioral and Brain Sciences*, 14(1), 138–139, <https://doi.org/10.1017/S0140525X00065717>.
- Center for Open Science (2017). Registered Reports <https://cos.io/rr/>.
- Chambers, C. D. (2013). Registered Reports: A new publishing initiative at Cortex. *Cortex*, 49(3), 609–610, <https://doi.org/10.1016/j.cortex.2012.12.016>.
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLOS ONE*, 8(3), e57410, <https://doi.org/10.1371/journal.pone.0057410>.
- Darwin, C. (1867). *The Origin of Species: By Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*. Cambridge University Press, 6 edition.
- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, 243(4899), 1668–1674, <https://doi.org/10.1126/science.2648573>.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126, <https://doi.org/10.1037/xge0000033>.
- Erev, I., et al. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23, 15–47, <https://doi.org/10.1002/bdm.683>.
- Frey, R., Hertwig, R., & Herzog, S. M. (2014a). Surrogate decision making: Do we have to trade off accuracy and procedural satisfaction? *Medical Decision Making*, 34(2), 258–69, <https://doi.org/10.1177/0272989X12471729>.
- Frey, R., Hertwig, R., & Rieskamp, J. (2014b). Fear shapes information acquisition in decisions from experience. *Cognition*, 132(1), 90–99, <https://doi.org/10.1016/j.cognition.2014.03.009>.
- Frey, R., Herzog, S. M., & Hertwig, R. (2018). Deciding on behalf of others: A population survey on procedural preferences for surrogate decision-making. *BMJ Open*, 8(7), e022289, <https://doi.org/10.1136/bmjopen-2018-022289>.
- Frey, R., Mata, R., & Hertwig, R. (2015a). The role of cognitive abilities in decisions from experience: Age differences emerge as a function of choice set size. *Cognition*, 142, 60–80, <https://doi.org/10.1016/j.cognition.2015.05.004>.
- Frey, R., Rieskamp, J., & Hertwig, R. (2015b). Sell in May and go away? Learning and risk taking in nonmonotonic decision problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(1), 193–208, <https://doi.org/10.1037/a0038118>.

- Gonzales, J. E. & Cunningham, C. A. (2015). The promise of pre-registration in psychological research <http://www.apa.org/science/about/psa/2015/08/pre-registration.aspx>.
- Hardin, G. (1960). The competitive exclusion principle. *Science*, 131(3409), 1292–1297, <https://doi.org/10.1126/science.131.3409.1292>.
- 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, <https://doi.org/10.1002/bdm.598>.
- Hauser, D. J. & Schwarz, N. (2016). Attentive Turkers: MTurk Participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, 48(1), 400–407, <https://doi.org/10.3758/s13428-015-0578-z>.
- Hertwig, R. (2015). Decisions from experience. In G. Keren & G. Wu (Eds.), *Blackwell Handbook of Judgment and Decision Making*.
- 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, <https://doi.org/10.1111/j.0956-7976.2004.00715.x>.
- Hertwig, R. & Pleskac, T. J. (2008). The game of life: How small samples render choice simple. In N. Chater & M. Oaksford (Eds.), *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*. Oxford University Press.
- Hills, T. T., Noguchi, T., & Gibbert, M. (2013). Information overload or search-amplified risk? Set size and order effects on decisions from experience. *Psychonomic Bulletin & Review*, <https://doi.org/10.3758/s13423-013-0422-3>.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2014). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, <https://doi.org/10.1016/j.tics.2014.10.004>.
- Hogarth, R. M., Lejarraga, T., & Soyer, E. (2015). The two settings of kind and wicked learning environments. *Current Directions in Psychological Science*, 24(5), 379–385, <https://doi.org/10.1177/0963721415591878>.
- Hogarth, R. M. & Soyer, E. (2011). Sequentially simulated outcomes: Kind experience versus nontransparent description. *Journal of Experimental Psychology: General*, 140(3), 434–463, <https://doi.org/10.1037/a0023265>.
- Hu, N., Zhang, J., & Pavlou, P. A. (2009). Overcoming the J-Shaped distribution of product reviews. *Communications of the ACM*, 52(10), 144, <https://doi.org/10.1145/1562764.1562800>.
- Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2(3), 196–217, https://doi.org/10.1207/s15327957pspr0203_4.
- Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Houghton Mifflin Company.
- Leibenstein, H. (1966). Allocative efficiency vs. "X-efficiency". *The American Economic Review*, 56(3), 392–415 <https://www.jstor.org/stable/1823775>.
- Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, 124, 334–342, <https://doi.org/10.1016/j.cognition.2012.06.002>.
- Leotti, L. A., Iyengar, S. S., & Ochsner, K. N. (2010). Born to choose: The origins and value of the need for control. *Trends in Cognitive Sciences*, 14(10), 457–463, <https://doi.org/10.1016/j.tics.2010.08.001>.
- Markant, D. B., Phillips, N., Kareev, Y., Avrahami, J., & Hertwig, R. (2019). To act fast or to bide time? Adaptive exploration under competitive pressure. *PsyArXiv Preprints* <https://osf.io/3jwjtq>.
- Mehlhorn, K., Ben-Asher, N., Dutt, V., & Gonzalez, C. (2014). Observed variability and values matter: Toward a better understanding of information search and decisions from experience. *Journal of Behavioral Decision Making*, 27(4), 328–339, <https://doi.org/10.1002/bdm.1809>.
- Miller, G. F. & Todd, P. M. (1998). Mate choice turns cognitive. *Trends in Cognitive Sciences*, 2(5), 190–198.
- Moore, J. W. (2016). What Is the sense of agency and why does it matter? *Frontiers in Psychology*, 7, <https://doi.org/10.3389/fpsyg.2016.01272>.
- Munafò, M. R., et al. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1(1), 0021, <https://doi.org/10.1038/s41562-016-0021>.
- Nicolaisen, J. (2002). The J-shaped distribution of citedness. *Journal of Documentation*, 58(4), 383–395, <https://doi.org/10.1108/00220410210431118>.
- Ostwald, D., Starke, L., & Hertwig, R. (2015). A normative inference approach for optimal sample sizes in decisions from experience. *Decision Neuroscience*, (pp. 1342), <https://doi.org/10.3389/fpsyg.2015.01342>.
- Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science*, 23(3), 184–188, <https://doi.org/10.1177/0963721414531598>.
- Phillips, N. D., Hertwig, R., Kareev, Y., & Avrahami, J. (2014). Rivals in the dark: How competition influences search in decisions under uncertainty. *Cognition*, 133(1), 104–119, <https://doi.org/10.1016/j.cognition.2014.06.006>.
- Pleskac, T. J. & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, 143(5), 2000, <https://doi.org/10.1037/xge0000013>.
- Poisson, S.-D. (1837). *Recherches sur la probabilité des jugements en matière criminelle et en matière civile: précédées des règles générales du calcul des probabilités*. Bachelier.
- Promoting reproducibility with registered reports. *Nature*

- Human Behaviour*, 1(1), 0034, <https://doi.org/10.1038/s41562-016-0034>.
- Rakow, T., Newell, B. R., & Zougkou, K. (2010). The role of working memory in information acquisition and decision making: Lessons from the binary prediction task. *The Quarterly Journal of Experimental Psychology*, 63, 1335–1360, <https://doi.org/10.1080/17470210903357945>.
- Schulze, C. & Newell, B. R. (2015). Compete, coordinate, and cooperate: How to exploit uncertain environments with social interaction. *Journal of experimental psychology. General*, 144(5), 967–981, <https://doi.org/10.1037/xge0000096>.
- Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology*, 78, 78–98, <https://doi.org/10.1016/j.cogpsych.2015.03.002>.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359–1366, <https://doi.org/10.1177/0956797611417632>.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129.
- Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. BiblioLife.
- Spaniol, J. & Wegier, P. (2012). Decisions from experience: Adaptive information search and choice in younger and older adults. *Frontiers in Decision Neuroscience*, 6, <https://doi.org/10.3389/fnins.2012.00036>.
- Stan Development Team (2016). Rstanarm: Bayesian applied regression modeling via Stan <https://mc-stan.org>.
- Walter, Z. & Lopez, M. S. (2008). Physician acceptance of information technologies: Role of perceived threat to professional autonomy. *Decision Support Systems*, 46(1), 206–215, <https://doi.org/10.1016/j.dss.2008.06.004>.
- Wulff, D. U., Hills, T. T., & Hertwig, R. (2014). Online product reviews and the description-experience gap. *Journal of Behavioral Decision Making*, <https://doi.org/10.1002/bdm.1841>.
- Wulff, D. U., Hills, T. T., & Hertwig, R. (2015). How short- and long-run aspirations impact search and choice in decisions from experience. *Cognition*, 144, 29–37, <https://doi.org/10.1016/j.cognition.2015.07.006>.
- Wulff, D. U., Mergenthaler Canesco, M., & Hertwig, R. (2017). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*.

Appendix

Choice environments

To maintain maximum control over the choice options' distributions (e.g., the differences between the choice options' modes, their EVs, their variances, the number of unique discrete outcomes, etc.), 1,000 outcomes were fixed prior to the experiment for each option (see Table A1 for a summary). Specifically, the eight decision problems (DPs) of each of the three choice environments were implemented as follows.

For the kind environment, eight choice options (A) were created by sampling 1,000 values from eight normal distributions with the means 10, 35, 60, 85, 215, 240, 265, 290 and standard deviation 1 in a first step. In a second step eight associated options (B) were generated equivalently, except shifting the eight normal distributions by +2 (DPs one to four) or by –2 (DPs five to eight). To obtain discrete outcomes, all sampled values were rounded to integers. The resulting eight decision problems are depicted in the first column of Figure A1.

For the two wicked environments, 600 values (moderately wicked environment) or 800 values (extremely wicked environment) were sampled from the same normal distributions as used in the first step for the kind environment. The remaining 400 values (moderately wicked environment; $p(\text{rare}) = .4$) or 200 values (extremely wicked environment; $p(\text{rare}) = .2$) — that is, the outliers in the bimodal distributions — were sampled from a second set of normal distributions. The means of these distributions constitute the anti-modes (i.e., local maxima) in the bimodal distribution and resulted from shifting the original distributions by +105 (DPs one to four) or by –105 (DPs five to eight) in the moderately wicked environment, and by +210 (DPs one to four) or by –210 (DPs five to eight) in the extremely wicked environment. As in the kind environment, all sampled values were rounded to integers to obtain discrete outcomes. The resulting eight bimodal distributions per environment were paired with the distributions generated in the second step of the kind environment, and are depicted in the middle and right column of Figure A1.

All in all, this systematic approach of generating the decision problems resulted in the following properties (see also Figure A2 and Table A1): i) The absolute difference between the two distributions' modes is 2 in all decision problems and all environments (the difference between the modes of H- and L-options is +2 in the kind environment and –2 in the wicked environments). ii) The difference between the EVs of H-options and L-options is 2 in all decision problems of the kind environment, and 40 in all decision problems of the wicked environments (i.e., 20 times larger EV-differences in the wicked environments as opposed to in the kind environment). iii) In half of the decision problems of the wicked environments, the H-options are the bimodal

distributions with larger variance, and vice versa for the other half of the decision problems (to be able to control for “variance-aversion”, cf. Figure A2). iv) Due to the previous point, despite a substantial amount of variance there exists no correlation between EV and variance *across all* H- and L-options in the wicked environments. That is, unlike in environments in which risks (i.e., variance) and rewards are positively correlated with each other (Pleskac & Hertwig, 2014), a crucial feature of the wicked environments is that one cannot use the shortcut of learning about only one of the two statistical properties (e.g., variance) to eventually infer the other (e.g., EV). v) Finally, all unimodal distributions involve about seven discrete outcomes, whereas the bimodal distributions involve about fourteen discrete outcomes.

Participants saw all values in units of USD, that is, divided by 100 (e.g., an outcome of 265 is displayed as “\$ 2.65”). The R-script to generate the choice environments and the full set of resulting decision problems can be downloaded from <https://osf.io/5vs83/>.

Instructions

Participants read the following instructions for the task, apportioned on multiple screens, and with a practice trial interspersed:

Instructions 1 / 3. In the main part of the study you will play a choice game consisting of 8 independent trials. After this game you will only have to answer a few survey questions, which will take no longer than 2-3 minutes. So let’s get started with the main part! In each trial of the choice game you will see two blue boxes as shown below:

[Depiction of unlabelled choice options]

Both boxes contain multiple and different payoffs in US dollars. The boxes may contain high or low payoffs, and the payoffs in a box might be relatively constant or quite variable.

Instructions 2 / 3. Each trial consists of two stages: in the first stage you can preview payoffs from both boxes (we will explain this on the next page). In the second stage, you have to make a final choice between the two boxes.

Once you make a final choice, 100 payoffs will be automatically drawn from the chosen box. The average of these 100 payoffs will then be saved as your score of the current trial. At the end of the study, 2 of your 8 scores will be randomly selected and you will receive these two amounts as an additional bonus payment on Amazon mTurk.

The total bonus payment can range up to \$6, so it might really be worthwhile to explore the boxes

thoroughly to identify and choose the box that yields the higher average payoff “in the long run”!

Instructions 3 / 3. To preview payoffs before making a final choice, simply click on a box. One of the existing payoffs will then be randomly drawn from that box and shown to you. After a short period of time, the previewed payoff disappears again and will be put back into the box. The boxes are shuffled after every draw, therefore the payoffs do not have a specific sequential order.

After you have previewed a payoff, you have to indicate whether you wish to preview another payoff from one of the two boxes or, alternatively, whether you feel like you have explored the boxes enough and would like to make a final choice.

[Instructional manipulation check] On the next page, there is also a small textbox on the left side. Please type the number [random number for each participant] into that box to demonstrate that you have read and understood the instructions.

Let’s have a look at an example and try this out!

[Practice trial]

[The following paragraph will only be displayed in the competitive condition.]

Second player. There is one last piece of important information: you are always going to play together with a second live player (you will be matched with a new second player after each trial). The second player will explore the identical two options simultaneously. The player who first stops the “exploration rounds” and opts to make a final choice can freely choose one of the two options. The other player will then be allocated the remaining option. If both players opt to make a final choice at the same time, there will be a random allocation of the two options.

The system will now pair you with a “second player” for the first trial. Afterwards, the game starts.

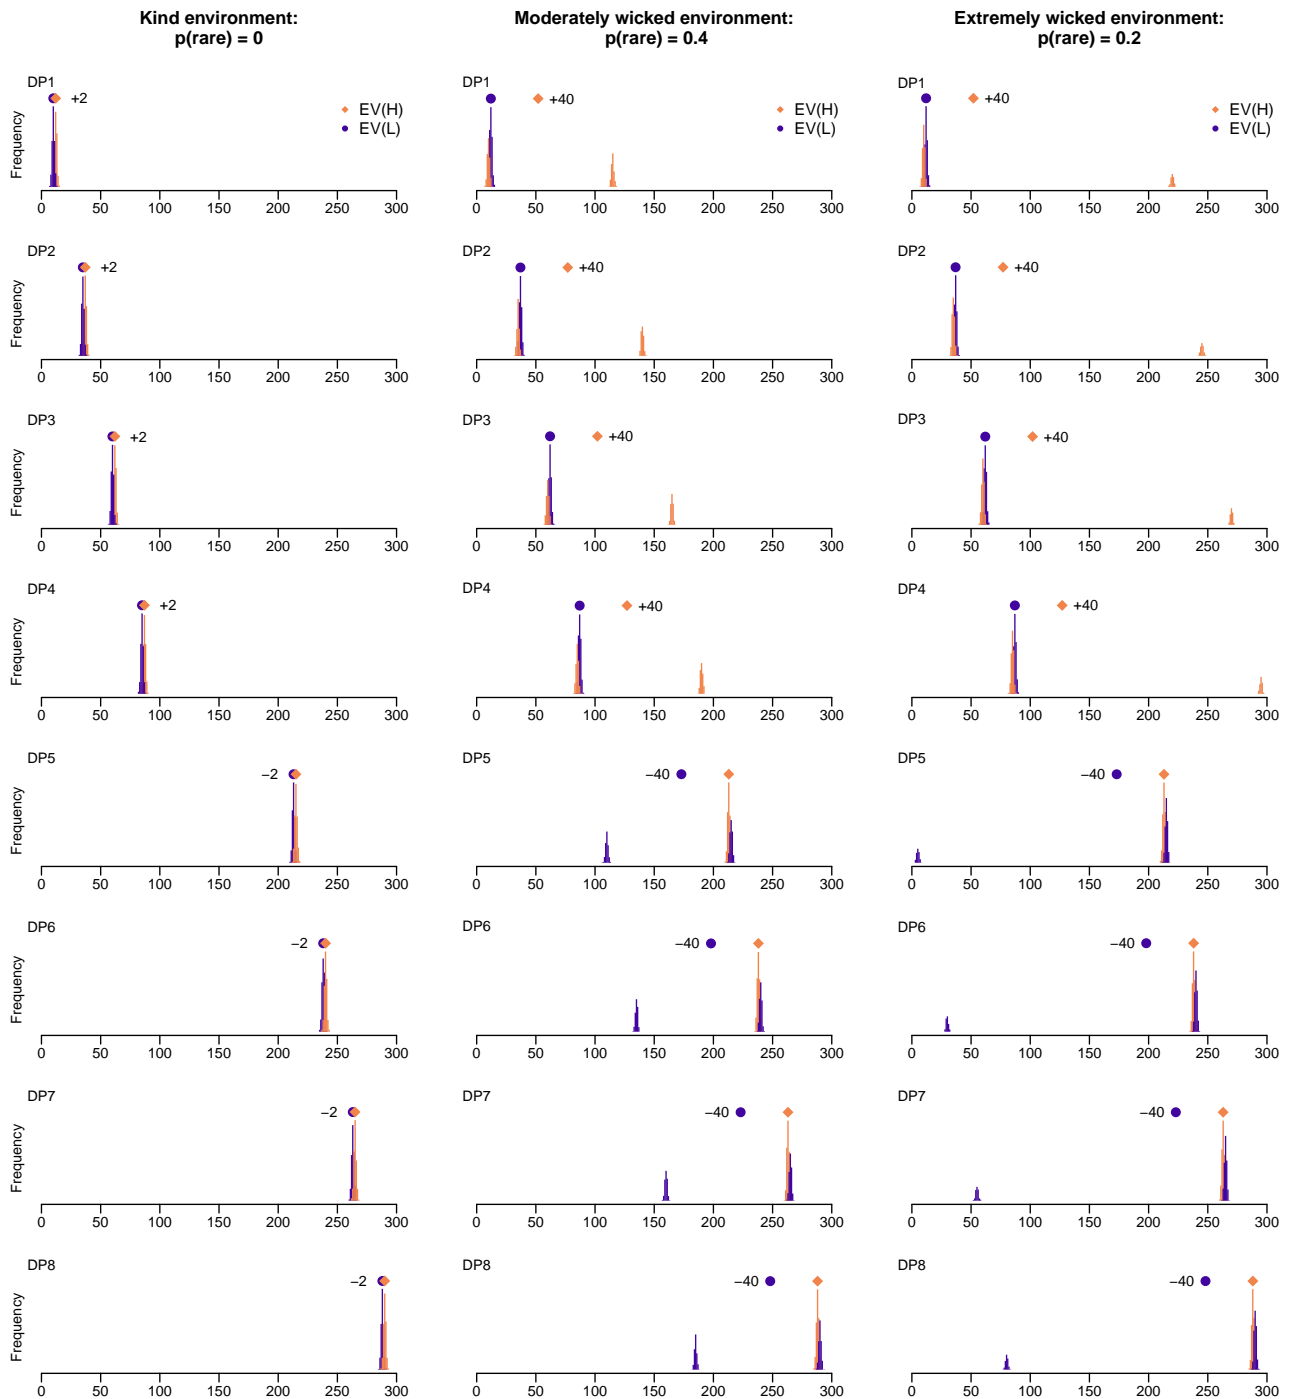


FIGURE A1: Decision problems of all choice environments. The diamonds and circles above the distributions depict the expected values (EV) of the higher-EV (H) and lower-EV (L) options, respectively. The numbers indicate the differences between the distributions' EVs. The H-options of the wicked environments have a larger variance than the L-options in half of the decision problems and vice versa in the other half (see Figure A2).




FIGURE A2: Expected value (EV) and standard deviation of all choice options in the three different choice environments. The corresponding higher- (H) and lower- (L) EV options of a decision problem are depicted next to each other (kind environment) or connected by a line (wicked environments).




FIGURE A3: Ratings of the importance of four strategies that may be pursued in the sampling game. The red horizontal lines depict mean ratings per condition.

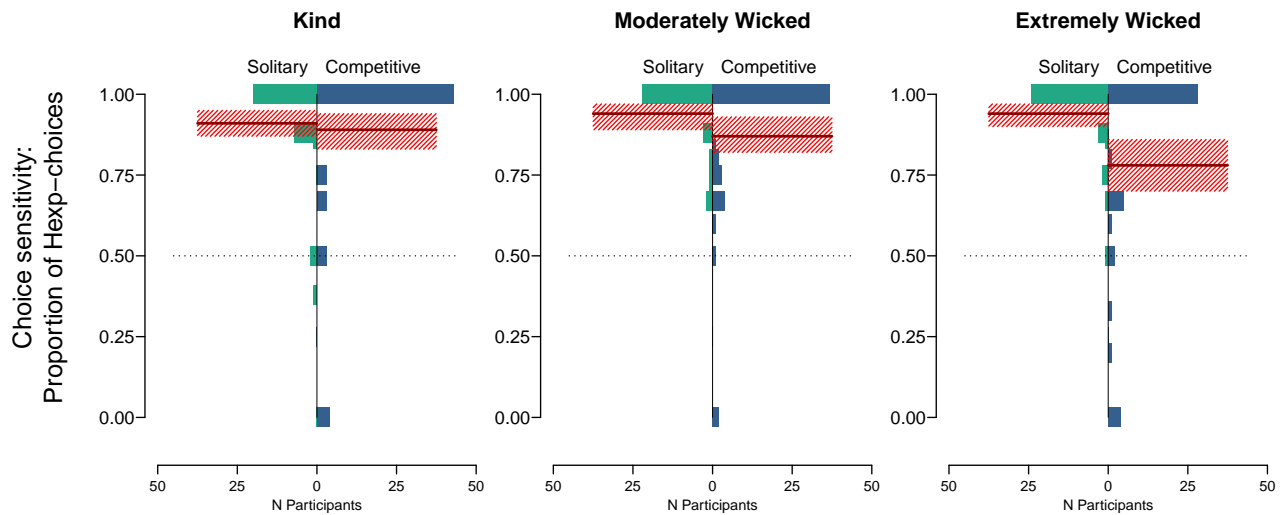


FIGURE A4: Proportion of choices of the options with higher experienced sample mean (Hexp), as a function of search mode (solitary vs. competitive), and separately for the three different choice environments. Histograms (i.e., the horizontal bars) depict the proportion of Hexp-choices aggregated across trials (i.e., one value per participant). The horizontal lines depict the posterior means for each combination of environment and search mode (resulting from the Bayesian mixed-effects models), with the red shaded areas representing 95% highest density intervals (HDI).

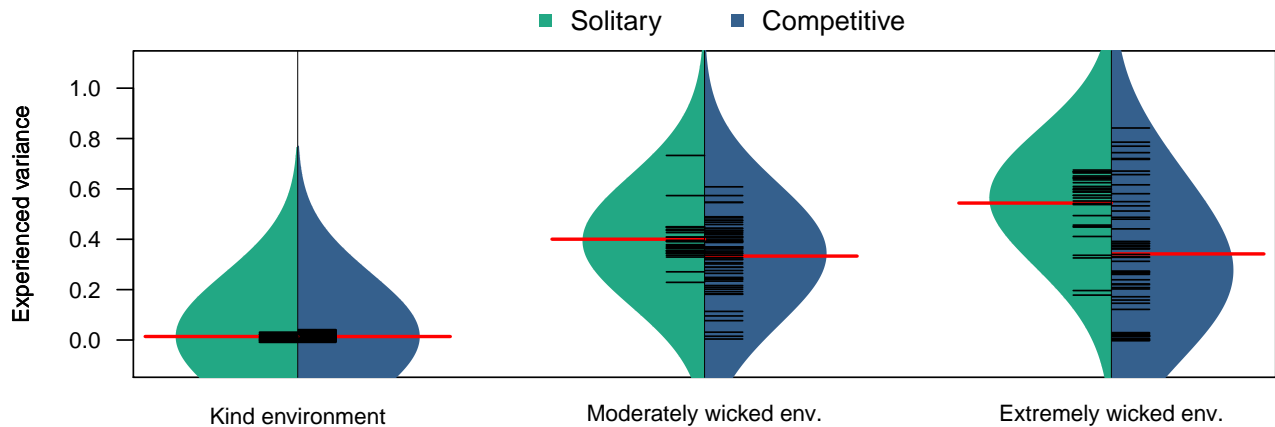


FIGURE A5: Experienced variance during pre-decisional search, separately for the two search modes and the three conditions. Each horizontal line depicts the mean variance a participant experienced across the eight decision problems. Note that in the kind environment, the experienced variance was close but not equal to 0 (i.e., even in the kind environment there were no safe options with fixed outcomes.)

TABLE A1: Summary of decision problems.

Environment	p(rare)	DP	M(H)	M(L)	SD(H)	SD(L)	Diff(EVs)	Diff(Modes)	bimodal
Kind	0.0	1	12	10	1	1	2	2	none
Kind	0.0	2	37	35	1	1	2	2	none
Kind	0.0	3	62	60	1	1	2	2	none
Kind	0.0	4	87	85	1	1	2	2	none
Kind	0.0	5	215	213	1	1	2	2	none
Kind	0.0	6	240	238	1	1	2	2	none
Kind	0.0	7	265	263	1	1	2	2	none
Kind	0.0	8	290	288	1	1	2	2	none
Moder. wicked	0.4	1	52	12	51	1	40	2	H
Moder. wicked	0.4	2	77	37	51	1	40	2	H
Moder. wicked	0.4	3	102	62	51	1	40	2	H
Moder. wicked	0.4	4	127	87	51	1	40	2	H
Moder. wicked	0.4	5	213	173	1	51	40	2	L
Moder. wicked	0.4	6	238	198	1	51	40	2	L
Moder. wicked	0.4	7	263	223	1	51	40	2	L
Moder. wicked	0.4	8	288	248	1	52	40	2	L
Extrem. wicked	0.2	1	52	12	84	1	40	2	H
Extrem. wicked	0.2	2	77	37	84	1	40	2	H
Extrem. wicked	0.2	3	102	62	84	1	40	2	H
Extrem. wicked	0.2	4	127	87	84	1	40	2	H
Extrem. wicked	0.2	5	213	173	1	84	40	2	L
Extrem. wicked	0.2	6	238	198	1	84	40	2	L
Extrem. wicked	0.2	7	263	223	1	84	40	2	L
Extrem. wicked	0.2	8	288	248	1	84	40	2	L

TABLE A2: Sociodemographic information.

N		277
Sex	Female	129 (46.6%)
	Male	147 (53.1%)
	NA	1 (0.4%)
Age		M=38.3 (SD=11.3)
Occupation	Employed	193 (69.7%)
	Retired	4 (1.4%)
	Self-employed	46 (16.6%)
	Student	7 (2.5%)
	Unemployed	24 (8.7%)
	NA	3 (1.1%)
Monthly net income	0 - 999 USD	64 (23.1%)
	1,000 - 1,999 USD	70 (25.3%)
	2,000 - 2,999 USD	56 (20.2%)
	3,000 - 3,999 USD	34 (12.3%)
	4,000 - 4,999 USD	22 (7.9%)
	> 5,000 USD	29 (10.5%)
	NA	2 (0.7%)
Years of education		M=15.1 (SD=2.4)

TABLE A3: Bayesian regression analyses.

	Sample size	Search inefficiency	Choice of Hexp	Choice of H
Intercept	13.7 [12 – 15.2]	12.5 [10.8 – 14.3]	0.91 [0.87 – 0.95]	0.9 [0.85 – 0.95]
Environment: Moderately wicked	1.2 [-0.8 – 3.2]	-1.5 [-4.1 – 0.8]	0.02 [-0.02 – 0.07]	-0.01 [-0.06 – 0.04]
Environment: Extremely wicked	0.7 [-1.5 – 2.6]	-3.8 [-6.1 – -1.4]	0.02 [-0.02 – 0.07]	-0.16 [-0.21 – -0.09]
Search mode: Competitive	-9.3 [-11.2 – -7.4]	-9.5 [-11.6 – -7.2]	-0.02 [-0.09 – 0.03]	-0.07 [-0.13 – -0.01]
Competitive x Moderately wicked	-1.1 [-3.6 – 1.5]	0.6 [-2.6 – 3.7]	-0.04 [-0.11 – 0.05]	-0.09 [-0.18 – 0.01]
Competitive x Extremely wicked	-1.1 [-3.4 – 1.5]	2.1 [-1 – 5.1]	-0.14 [-0.24 – -0.03]	-0.08 [-0.21 – 0.06]
Trial	-0.5 [-0.6 – -0.3]	-0.4 [-0.5 – -0.2]	0 [0 – 0.01]	0 [-0.01 – 0.01]
Trial x Competitive	0.2 [0 – 0.4]			
Trial x Moderately wicked	0 [-0.2 – 0.2]			
Trial x Extremely wicked	0.2 [0 – 0.5]			

Note. Credible coefficients (with highest density intervals excluding 0) are printed in bold. Intercept depicts the reference level “trial 1 in the solitary mode of the kind environment”. The three-way interactions for sample size were not credible and are not shown in the table.