

Foraging, Exploration, or Search?

On the (lack of) Convergent Validity Between Three Behavioral Paradigms

December 1 , 2017

Bettina von Helversen^{a,*}, Rui Mata^b, Gregory R. Samanez-Larkin^c, Andreas Wilke^d

^a Department of Psychology, University of Zurich, Switzerland

^b Department of Psychology, University of Basel, Switzerland

^c Department of Psychology and Neuroscience, Duke University, USA

^d Department of Psychology, Clarkson University, USA

* Corresponding author: University of Zurich, Department of Psychology,

Binzmühlestr. 14, box 19, CH-8050 Zurich, Switzerland.

Tel.: +41 44 6357181; E-mail address: b.vonhelversen@psychologie.uzh.ch.

Acknowledgements

We thank Amanda Sherman and the members of the Evolution and Cognition Lab for their help in data collection. This work was supported by a grant from the Swiss National Research Foundation (SNF) to the first author (grant no. 157432). Gregory Samanez-Larkin was supported by a Pathway to Independence Award from the US National Institute on Aging (R00-AG042596).

Abstract

Recently it has been suggested that individual humans and other animals possess different levels of a general tendency to explore or exploit that may influence behavior in different contexts. In the present work, we investigated whether individual differences in this general tendency to explore (exploit) can be captured across three behavioral paradigms that involve exploration–exploitation trade-offs: A foraging task involving sequential search for fish in several ponds, a multi-armed bandit task involving repeatedly choosing from a set of options, and a sequential choice task involving choosing a candidate from a pool of applicants. Two hundred and sixty-one participants completed two versions of each of the three tasks. Structural equation modeling revealed that there was no single, general factor underlying exploration behavior in all tasks, even though individual differences in exploration were stable across the two versions of the same task. The results suggest that task-specific factors influence individual levels of exploration. This finding causes difficulties in the enterprise of measuring general exploration tendencies using single behavioral paradigms and suggests that more work is needed to understand how general exploration tendencies and task-specific characteristics translate into exploratory behavior in different contexts.

Keywords

Decision-making; domain-general; domain-specificity; exploitation; exploration.

Public Significance Statement

Many tasks require balancing the need to exploit known options and exploring new ones. The present study investigated whether individual tendencies to explore

or exploit generalize across three different tasks and found that exploration was highly task specific: Someone who explored a lot in one task did not necessarily explore much in another. More work is needed to understand the task features that drive exploration.

Introduction

The need for search is ubiquitous. Humans, like most animals, have needed to search for food, mates, and shelter in order to survive and reproduce throughout evolutionary time. Search is costly, however, because finding new options involves opportunity costs and prevents exploiting already known ones, thus requiring to trade off exploration with exploitation. There has been considerable interest in understanding how humans deal with such exploration–exploitation trade-offs, involving the use of a number of behavioral paradigms that mimic requirements of searching for food, mates, and other resources (Mehlhorn et al., 2015; Wilke, in press), such as foraging for fish in ponds (Hutchinson, Wilke, & Todd, 2008), exploring options from an array (Cohen, McClure, & Yu, 2007), or searching for a suitable candidate from a pool (von Helversen, Wilke, Johnson, Schmid, & Klapp, 2011). But do individuals' exploration–exploitation tendencies generalize across such varied problems?

Some have argued that individuals' exploration tendencies may be a trait, akin to personality, in humans (Somerville et al., 2017) and other animals (Bell, Hankison, & Laskowski, 2009). Importantly, behavioral ecologists suggest that domain-general personality traits can arise whenever the environment cues the development of basic physiological or cognitive makeup in early life, with implications for the full life span (Dall, Houston, & McNamara, 2004; Sih & Del Giudice, 2012). For example, it has been suggested that stressful early life experiences cue future environmental uncertainty and lead to general explorative tendencies because these are adaptive in such harsh environments (Frankenhuis & Weerth, 2013; Sih & Del Giudice, 2012).

Another source of evidence for the generality of exploration tendencies comes from research suggesting that the control of exploration and exploitation is driven by

a domain-general mechanism that can affect exploration across different domains and tasks, such as internal (i.e., in memory) and external (i.e., in the environment) search (Hills, Todd, Lazer, Redish, & Couzin, 2015; Pirolli, 2007). This hypothesis is supported by research showing that different kinds of exploration–exploitation tasks, such as foraging for food or searching for words in memory, may rely on similar cognitive strategies (e.g., Hutchinson et al., 2008; Wilke, Hutchinson, Todd, & Czienskowski, 2009), have similar neurobiological roots (Hills, 2006; Hills, & Dukas, 2012; Hills et al., 2015), and can be behaviorally primed by each other (Hills, Todd, & Goldstone, 2008).

Despite these lines of evidence that suggest domain-generality of individual exploration–exploitation tendencies, little is known about the empirical convergence of individual differences in exploration–exploitation tendencies across different tasks. Yet, assessing whether exploration tendencies generalize across tasks could help determine the structure of a general control mechanism for exploration–exploitation; in particular, any lack of generalizability across specific situations could indicate potential boundaries of this general control mechanism or, alternatively, that the same mechanism interacts with task characteristics to produce different levels of exploration.

In the following, we briefly introduce three paradigms that have been used to understand how humans deal with exploration–exploitation trade-offs, and then report a convergent validity test of individual exploration–exploitation tendencies across these paradigms.

Three Paradigms Involving Exploration–exploitation Trade-offs

Foraging

In foraging tasks, people gather rewards by foraging in patches that differ in the number of rewards they contain (e.g., Bell, 1991; Winterhalder & Smith, 1981). Accordingly, people can decide to exploit a specific patch or to explore to find a new patch with more rewards. For instance, when collecting berries one will first harvest from a specific blackberry bush. After some time, most blackberries will be harvested and finding those remaining on that particular bush will be increasingly time intensive. At this point, the harvester has to make the decision whether to continue searching for blackberries on this bush or to look elsewhere—even if getting there might take additional travel time and involves uncertainty about the quality of the next bush. Exploration–exploitation trade-offs in foraging have been investigated in many animals (Blanchard, Wilke, & Hayden, 2014; Cuthill, Haccou, & Kacelnik, 1994; Louâpre, Alphen, & Pierre, 2010), including humans (Hutchinson et al., 2008; Wilke & Barrett, 2009; Fu & Gray, 2006; Hills, Todd, & Goldstone, 2010; Wilke et al., 2009). Recently, computerized game-like tasks have been developed to help understand the cognitive mechanisms underlying humans' exploration–exploitation decisions in such foraging scenarios, such as foraging for fish across ponds (Hutchinson et al., 2008; Mata, Wilke & Czienskowski, 2009, 2013).

Repeated Choice

Another class of paradigms frequently used to investigate how people trade off exploration and exploitation is the repeated choice or bandit task—a task that has raised considerable interest in understanding the cognitive and neural basis of such decisions in the past decade (Cohen et al., 2007; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Navarro, Newell, & Schulze, 2016; Steyvers, Lee, & Wagenmakers,

2009). Bandit tasks typically ask people to repeatedly choose from a set of two or more options that differ in their expected reward while receiving probabilistic feedback from their choices. This requires trading off exploration with exploitation: on the one hand, once people know which option provides the larger expected value they should exploit this option to maximize their payoff. On the other, exploiting too early involves the danger of choosing the wrong option. Furthermore, the quality of the options may change over time, which can be missed if only one option is exploited. Although most bandit tasks are structurally different from the foraging tasks discussed above, for example, in that they ask decision makers to repeatedly choose between the same options, bandit tasks can also be couched as foraging tasks in which individuals are faced with the problem of learning about whether two or more patches differ in their mean quality (cf. Rutledge, Lazzaro, Lau, Myers, Gluck, & Glimcher, 2009).

Sequential Choice

Sequential choice paradigms also involve exploration–exploitation trade-offs and have been studied extensively (Bearden, Rapoport, & Murphy, 2006; Burns, Lee & Vickers, 2006; Ferguson, 1989; Gilbert & Mosteller, 1966; Miller & Todd, 1998; Seale & Rapoport, 1997, 2000; von Helversen et al., 2011; von Helversen & Mata, 2012). When making decisions, it is not always possible to evaluate options simultaneously because patches or options can be distributed in space and time, forcing the decision maker to sequentially search through each one until an acceptable one is found. For instance, a person harvesting fruit may need to consider at each fruit tree whether it is worth the effort of trying to reach the fruit, or to continue searching for a more accessible tree. Furthermore, it is often not possible to return to a patch or option that has previously been rejected because, for instance, it has been exploited by

others. Exploration is necessary to reduce the risk of settling for an inferior option because usually no (or only imprecise) information exists about the quality of the options that can be expected. Furthermore, the probability of at least encountering the best option increases with the amount of search. However, too much exploration can be dangerous because to continue exploration may mean passing by (and hence rejecting) the best option.

The Present Study

To our knowledge, there has been no previous attempt to assess whether individual differences in exploration generalize across different tasks involving exploration–exploitation trade-offs. We filled this gap by asking participants to complete three different tasks: A foraging task involving foraging for fish (Hutchinson et al., 2008), a repeated choice task in the form of a multi-armed bandit problem (e.g., Daw et al., 2006), and a sequential choice task as a version of the secretary task (Seale & Rapoport, 1997). Furthermore, participants completed two versions of each task to allow us to test the consistency of exploration preferences both across and within tasks, with the latter being a precondition for stability across tasks.

Methods

Design and Procedure

Participants completed all three exploration–exploitation tasks.¹ Participants completed the task either on one day or within one week (due to different class

¹ In addition, we measured individual differences in cognitive abilities (processing speed, working memory, and attention switching), risk-taking propensity, and a number of personality characteristics related to decision making to be used in an exploratory fashion. For details see the online supplementary materials.

schedules). The order of the task was kept constant, starting with the fishing task, followed by the bandit task, and the secretary task (see Figure 1). If participants completed the tasks in two sessions they would typically complete the fishing task in the first session and the other tasks in the second session. To incentivize them, participants received course credit towards the research requirement of their PY151 Introductory Psychology class as well as the chance to win a \$20 lottery for each of the three tasks. At the end of each semester a winner for the three lotteries was chosen.

Participants

We planned to acquire data from at least 200 participants to obtain enough power to conduct factorial analysis testing the correlations between three latent variables corresponding to each of the three behavioral measures of exploration (Kline, 2011). To achieve our desired N , our sampling strategy consisted of recruiting participants across two full academic years (2010/2011, 2011/2012) at Clarkson University, NY. Overall, a total of 303 participants were recruited to take part in the study. The study took on average 2.5 hr. Participants received course credit and entered a monetary lottery for each task. We excluded 42 of the 303 participants because they failed to complete all measures or their data could not be matched between tasks. This resulted in a final sample of 261 participants (124 men, mean age = 19 years, $SD = 1.69$). In the fishing task, we additionally excluded 36 participants: 28 because they experienced a different foraging environment due to an experimenter mistake when setting up parameter files, and 8 because they missed a large number of

fish (> 2 *SDs*) in either one of the two task versions, indicating that they were not paying attention. The study was approved by Clarkson University's Institutional Review Board (IRB #10-26).

All data are accessible on the Open Science Framework (Link: https://osf.io/4wf8r/?view_only=53a7b98e41c1460e837ddc2bc3d144d5).

Materials

Foraging: The fishing task.

The foraging task was a computerized fishing task in which participants had to maximize the number of fish they could catch within 45 min (Hutchinson et al., 2008; Mata et al., 2009, 2013). In the game, participants waited at a pond until a fish appeared at the surface (for a screen shot see Figure 1a). Once a fish appeared they could catch it by clicking on it with the mouse cursor within 2 s. If they did not click on it, the fish disappeared again. The rate at which fish appeared at the surface of a pond was a stochastic function of the number of fish in the pond ($0.005 \times (\text{number of fish in pond}) \text{ s}^{-1}$) such that the fewer fish were left in a pond, the longer it took until the next fish appeared. At any point in time the participant could decide to leave the pond they were currently at and go to a new pond, but traveling to a new pond required time. The new pond was chosen randomly and its quality was unknown to the participant. Participants played two versions of the foraging task, a short and a long one. Whether participants started with the short or the long version was counterbalanced between participants. In the short version the traveling time between ponds was 15 s, in the long version the traveling time was 35 s. The number of fish in a new pond was drawn from a negative binomial distribution (see Hutchinson et al., 2008 for details). This distribution is characterized by a relatively low mean but a long tail, representing an environment in which most patches will have few fish but a

few have a large number. Specifically, 50% of the ponds experienced by the participants contained 10 fish or fewer, but 5% contained 25 fish or more. Participants received \$0.10 for each fish.

Repeated choice: The bandit task.

Each participant made choices from a five-armed bandit (e.g., Daw et al., 2006). Participants were presented with five abstract cues representing the different choice options and had to choose one cue in each trial (see Figure 1b). Each arm had a stationary (time-independent) Bernoulli reward distribution function. All five arms were independent of each other. Choices were self-paced by participants. After each choice participants received feedback regarding the outcome of the chosen cue. Possible outcomes for each choice were \$1 or \$0 and the outcome was displayed for the selected cue only. The probability of earning \$1 varied across cues but remained stationary within a block, which was told to participants prior to playing the task. For the entire duration of the task, the number of remaining trials was shown on the screen in order to make time salient. Participants were told to maximize reward in the task, and that they could win 10% of their total earnings.

Participants completed four versions of the bandit task based on a 2 x 2 factorial design. Each of the four versions was repeated twice resulting in a total of eight blocks of trials. For each version, we manipulated the number of trials (long vs. short) and the variance of the expected payoff (high vs. low). Short versions consisted of 25 trials, and long versions consisted of 75 trials. The expected payoffs for each of the arms were farther apart (\$0.30, \$0.40, \$0.50, \$0.60, and \$0.70) in the high variance versions and closer together (\$0.40, \$0.45, \$0.50, \$0.55, and \$0.60) in the low variance versions. All participants played the eight blocks in the same order (long

high, short high, long low, short low, long high, short high, long low, short low). A different set of cues was used for each of the eight blocks.

Sequential choice: The secretary task.

As a sequential decision-making task participants completed a computerized secretary game (Seale & Rapoport, 1997). In the game participants had to select the best secretary from a pool of applicants. Applicants were presented sequentially in a random sequence. Participants did not receive direct information about the quality of an applicant, but only the relative rank of an applicant. The relative rank denotes how good an applicant is compared with the applicants the participant had seen so far in this trial. For each applicant participants had to decide whether they wanted to accept or reject the applicant (for a screenshot see Figure 1c). If they rejected an applicant the next one would appear. If they accepted an applicant the trial was over and they received feedback about whether they had selected the best applicant or not. If participants did not accept an applicant before they reached the last one in the sequence, they had to select the last applicant. A rejected applicant could not be recalled at a later time. If participants selected the best applicant within one trial they received 10 points, otherwise they received 0 points in that trial. As before, participant played a short and a long version of the secretary game. As in the fishing task, it was counterbalanced across participants whether participants started with the short or the long version of the task. In the short version the pool of applicants consisted of 20 applicants. In the long version the pool consisted of 40 applicants. Participants played 30 games of each version.

Results

To investigate whether there was a domain-general exploration factor we first report measures of exploration and their stability within the three tasks (see Table 1) and then a structural equation modeling (SEM) approach to measure whether exploration shared variance across the tasks. SEM is a multivariate statistical technique that can be used to estimate and test models regarding latent constructs and the relations to each other and manifest variables (for an introduction to SEM see Ullman & Bentler, 2012).

We followed a behavioral definition of exploration as behavior that indicates a change in the option that is considered (see also Mehlhorn et al., 2015). Accordingly, in the fishing task we measured exploration with the number of ponds visited. In the multi-armed bandit task we used the percentage of choices that involved a switch between options as a measurement for exploration averaged across blocks.² In the secretary task we measured exploration by the mean percentage of options that participants looked at before accepting an option.

In general, within each of the three tasks, the amount of exploration in the long and the short version of the tasks was highly related, suggesting that the exploration measures reflect stable individual differences in how much people explore in the three tasks (see Table 1). If the degree of exploration is driven by a task-general factor exploration should also share variance across tasks. To explore this, we compared three different measurement models using AMOS (version 20), an SPSS software module for structural equation modeling, assuming that in each task exploration–exploitation measures of the short and the long version constitute a latent factor. An

² Because the variance manipulation in the bandit task did not significantly affect behavior, we aggregated exploration measures over the high and the low variance versions of the task.

overview is presented in Table 2. The first measurement model would correspond to a completely task-specific model where all variance in the trade-off between exploration and exploitation depends on the task. In contrast, the third measurement model corresponds to a completely task-general model where exploration is completely determined by the same mechanism across all tasks. The second model presents a middle version that allows measuring the magnitude of variance that is shared between the tasks. Thus, high inter-correlations between the tasks would speak for a general exploration factor, whereas no or low inter-correlations would suggest that exploration is largely determined by task-specific characteristics.

We tested for deviations from normality and log-transformed the number of ponds visited in the fishing task to achieve more normally distributed values. In addition, we removed seven outliers with the largest Mahalanobis distance to reach multivariate normality (descriptives are reported in Table S2 in the online appendix).³ Models were fitted using maximum likelihood estimation, estimating means and intercepts for the missing values in the fishing task. For each model we report four common and recommended measures of goodness of fit, the chi-square statistic, the root mean square error of approximation (RMSEA), the comparative fit index (CFI) to the independence model, and the Tucker Lewis index (TLI), see Table 2. Usually a good fit can be assumed if $RMSEA < .06$, and CFI and $TLI > .95$ (Ullman & Bentler, 2012; Schreiber, Nora, Stage, Barlow, & King, 2006).

³ Outliers identified with the Mahalanobis distance largely overlapped with outliers identified via visual inspection and the anomaly detection procedure implemented in SPSS. Including or excluding outliers does not change much the estimates for the standardized coefficients for the paths between exploration in the bandit and the fishing task and the bandit and the secretary task (estimated coefficients including all cases are -.13 and -.11 respectively) but it influences whether they reach significance. Including all cases, all relations between tasks are non-significant. Excluding further potential outliers the paths between the bandit and the fishing task and the secretary task stay significant.

Overall, the second model (see Figure 2) assuming a three-factor latent variable model with correlated latent factors for each exploration task performed best. The model fit indices indicated a good fit to the data and the standardized residuals were all within acceptable limits (i.e., < 2). It clearly performed better than model 3, which assumed that exploration in all tasks was stemming from a single factor. It also significantly increased model fit compared to model 1, which assumed that the three latent factors are not correlated, as indicated by a chi-square difference test, $\Delta\chi^2(3) = 8.43, p = .038$. In line with this result model 2 indicated a significant negative relationship between exploration in the bandit task and exploration in the fishing task, $r = -.15, b = -.015, SE = .007, C.R. = -2.11, p = .035$, and between exploration in the bandit and the secretary task, $r = -.14, b = -.005, SE = .002, C.R. = -2.04, p = .041$. But there was no relation between exploration in the secretary and the fishing task, $r = .01, b < .001, SE = .003, C.R. = 0.14, p = .892$.

Additional analyses focusing on exploration at the beginning and the end of the task as well as using computational models to extract parameters measuring a participant's level of exploration show a similar picture of small negative or non-significant relationships between the tasks. We also investigated whether individual difference measures of risk preference were related to exploration behavior, because previous research has suggested a relation between risk taking and exploration (e.g. Schunk & Winter, 2009). We found small but significant correlations between risk preference and exploration in the secretary task and the fishing task, but taking these into account did not change the relation of exploration between the tasks. A more detailed report of all additional exploratory analyses can be found in the online appendix (S1-S4) on OSF (Link: https://osf.io/4wf8r/?view_only=53a7b98e41c1460e837ddc2bc3d144d5).

Discussion

Humans and other animals encounter numerous trade-offs between exploration and exploitation in a range of different decision tasks. Recent research has suggested that exploration is guided by general tendencies (e.g., Sih & Del Giudice, 2012) and controlled by a domain-general mechanism (e.g., Hills et al., 2010; 2015), which imply that consistent individual differences in exploration tendencies may be observed across various paradigms involving exploration–exploitation trade-offs. Here, we investigated individual differences in exploration behavior across three research paradigms: a foraging task, a repeated decision task, and a sequential decision-making task. Although exploration behavior was rather stable within versions of each task, a psychometric analysis showed no common factor underlying exploration across tasks. If anything, the analysis suggested that *more exploration* in the bandit task may be related to *less exploration* in the fishing task and *less exploration* in the secretary task. However, one should be cautious in the interpretation of these negative correlations because we found in exploratory analyses that they depend on excluding outliers (not excluding outliers leads to similar but non-significant negative correlations between the two tasks). Exploration in the secretary task was not related to exploration in the fishing task. These results show little evidence for a general exploration factor and suggest that behavioral measures of exploration do not generalize across the three tasks we examined here.

There are a number of possible explanations for our findings. First, it could be that domain-general preferences for exploration–exploitation do not exist or are trumped by strong domain-specific expectations about the underlying statistic regularities of particular environments. For example, expectations concerning

exploration when foraging for food may be considerably different than those for monetary rewards or job markets. Future work could likely profit from understanding the statistical structure of past and present foraging environments and how different exploration–exploitation paradigms (and cover stories) elicit such representations in decision makers today.

Second, domain-general preferences may be further masked by individual differences in response to task-specific characteristics as a function of their demands on specific abilities such as memory, time perception, and numerosity or numeracy. For example, the tasks we used differed in a number of structural aspects including whether individuals chose between the same (bandit) or different options over time (fishing, secretary), or faced limits concerning time (fishing) or number of choices (secretary, bandit). Similarly, the monetary payoffs of exploration versus exploitation differed between tasks and it is unclear how to equate them (cf. Mehlhorn et al., 2015). Future work that varies such characteristics systematically in multiple paradigms could provide additional insights into the role of such characteristics and, more broadly, the possibility of capturing general individual tendencies across different situations.

Third, and perhaps more interestingly, it could be that general exploration tendencies operate at the level of goal disengagement from default or higher-level goals that are differentially relevant or prominent in the different tasks (Hills et al., 2010). For example, in the fishing task, the default behavior may be exploitation (foraging within a pond), and the decision maker needs to switch to another behavior (exploring a new pond). In contrast, in the bandit and secretary tasks the default behavior may be exploration, with the decision maker having to make an active decision to stop exploring (see also Mehlhorn et al., 2015). Individual differences in

goal setting or cognitive control abilities in engaging/disengaging from default goals, rather than exploration per se, could perhaps help explain the negative relation found between exploration the bandit task and the fishing task (e.g., Hills, 2011; Mata & von Helversen, 2015). One possible avenue for future work that evaluates the role of goals could involve manipulating goals directly, for example, through task instructions (cf. Mata et al., 2009).

In sum, our research found no general exploration tendency across three different tasks commonly used to study exploration–exploitation trade-offs. Our findings point out the need to understand how particular domains, task characteristics, or goals lead to differential exploration in various foraging-like situations.

References

- Bearden, J. N., Rapoport, A., & Murphy, R. O. (2006). Experimental studies of sequential selection and assignment with relative ranks. *Journal of Behavioral Decision Making, 19*, 229–250.
- Bell, A. M., Hankison, S. J., & Laskowski, K. L. (2009). The repeatability of behaviour: a meta-analysis. *Animal Behaviour, 77*, 771–783.
- Bell, W. J. (1991). *Searching behaviour: The behavioural ecology of finding resources*. New York: Chapman and Hall.
- Blanchard, T. C., Wilke, A., & Hayden, B. Y. (2014). Hot-hand bias in Rhesus monkeys. *Journal of Experimental Psychology: Animal Learning and Cognition, 40*, 280–286.
- Burns, N., Lee, M., & Vickers, D. (2006). Are individual differences in performance on perceptual and cognitive optimization problems determined by general intelligence? *Journal of Problem Solving, 1*, 5–19.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007) Should I stay or should I go? Exploration versus exploitation. *Philosophical Transactions of the Royal Society B: Biological Sciences, 362*, 933–942.
- Cuthill, I. C., Haccou, P., & Kacelnik, A. (1994). Starlings (*Sturnus vulgaris*) exploiting patches: Response to long-term changes in travel time. *Behavioral Ecology, 5*, 81–90.
- Dall, S. R. X., Houston, A. I., & McNamara, J. M. (2004). The behavioural ecology of personality: consistent individual differences from an adaptive perspective. *Ecology Letters, 7*, 734e739.
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature, 441*, 876-879.

- Ferguson, T. (1989). Who solved the secretary problem? *Statistical Science*, 4, 282–289.
- Frankenhuis, W. E., & de Weerth, C. (2013). Does early-life exposure to stress shape or impair cognition? *Current Directions in Psychological Science*, 22, 407–412.
- Fu, W.-T., & Gray, W. D. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology*, 52, 195–242.
- Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *American Statistical Association Journal*, 61, 35–73.
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30, 3–41.
- Hills, T. T. (2011). The evolutionary origins of cognitive control. *Topics in Cognitive Science*, 3, 231–237.
- Hills, T. T. & Dukas, R. (2012). The evolution of cognitive search. In P. M. Todd, T. T. Hills, & T. Robbins (Eds.), *Cognitive search: evolution, algorithms and the brain* (pp. 11–24). Cambridge, MA: MIT Press.
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2008). Search in internal and external spaces: Evidence for generalized cognitive search processes. *Psychological Science*, 19, 676–682.
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2010). The central executive as a search process: exploration and exploitation in generalized cognitive search processes. *Journal of Experimental Psychology: General*, 139, 590–637.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19, 46–54.

- Hutchinson, J. M. C., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: Can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour*, 75, 1331–1349.
- Kline, R. B. (2011). *Principles and practice of structural equation modeling* (3rd ed.). New York: Guilford Press.
- Louâpre P, van Alphen J. J. M., & Pierre J.-S. (2010). Humans and insects decide in similar ways. *PLoS ONE*, 5(12), e14251.
- Mata, R., & von Helversen, B. (2015). Search and the aging mind: The promise and limits of the cognitive control hypothesis of age differences in search. *Topics in Cognitive Science*, 7, 416–427.
- Mata, R., Wilke, A., & Czienskowski, U. (2009). Cognitive aging and adaptive foraging behavior. *Journal of Gerontology: Psychological Sciences*, 64, 474–481.
- Mata, R., Wilke, A., & Czienskowski, U. (2013). Foraging across the lifespan: Is there a reduction in exploration with aging? *Frontiers in Neuroscience*, 7, 53.
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M., Morgan, K., Braithwaite, V., ... Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2, 191–215.
- Miller, G. F., & Todd, P. M. (1998). Mate choice turns cognitive. *Trends in Cognitive Sciences*, 5, 190–198.
- Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world : An investigation of the explore-exploit dilemma in static and dynamic environments. *Cognitive Psychology*, 85, 43–77.
- Pirolli, P. L. T. (2007). *Information foraging theory: Adaptive interaction with information*. Oxford: Oxford University Press.

- Rutledge, R. B., Lazzaro, S. C., Lau, B., Myers, C. E., Gluck, M.A., & Glimcher, P. W. (2009). Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *Journal of Neuroscience* 29, 15104-15114.
- Schreiber, J. B., Nora, A., Stage, F. K., Barlow, E. A., & King, J. (2006). Reporting structural equation modeling and confirmatory factor analysis results: A review. *The Journal of Educational Research*, 99(6), 323–338.
Doi:10.3200/JOER.99.6.323-338
- Schunk, D., & Winter, J. (2009). The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior and Organization*, 71, 347-360.
- Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the secretary problem. *Organizational Behavior & Human Decision Processes*, 69, 221–236.
- Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, 13, 280–300.
- Sih, A., & Del Giudice, M. (2012). Linking behavioural syndromes and cognition: A behavioural ecology perspective. *Philosophical Transactions of the Royal Society of London B*, 367, 2762-2772.
- Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology: General*, 146, 155–164.

- Steyvers, M., Lee, M. D., & Wagenmakers, E.-J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53, 168–179.
- Ullman, J. B., Bentler, P. M. (2012). Structural equation modeling. In Weiner, I. (Ed.), *Handbook of psychology* (2nd Edition, pp. 661-690, 2:IV:23). Hoboken, NJ: John Wiley.
- von Helversen, B., & Mata, R. (2012). Losing a dime with a satisfied mind: positive affect predicts less search in sequential decision making. *Psychology and Aging*, 27, 825–39.
- von Helversen, B., Wilke, A., Johnson, T., Schmid, G., & Klapp, B. (2011). Performance benefits of depression: Sequential decision making in a healthy sample and a clinically depressed sample. *Journal of Abnormal Psychology*, 120, 962–968.
- Wilke, A. (in press). The adaptive problem of exploiting resources: Human foraging behavior in patchy environments. Forthcoming book chapter.
- Wilke, A., & Barrett, H. C. (2009). The hot hand phenomenon as a cognitive adaptation to clumped resources. *Evolution and Human Behavior*, 30, 161–169.
- Wilke, A., Hutchinson, J. M. C., Todd, P. M., & Czienskowski, U. (2009). Fishing for the right words: Decision rules for human foraging behavior in internal search tasks. *Cognitive Science*, 33, 497–529.
- Winterhalder, B., & Smith, E. A. (1981). *Hunter-gatherer foraging strategies: Ethnographic and archaeological analyses*. Chicago: University of Chicago Press.

Figure captions

Figure 1: Overview of the study with screenshots of the three computerized tasks.
A) Fishing task; B) Bandit task; C) Secretary task.

Figure 2. Best measurement model for exploration/exploitation behavior. Circles represent latent constructs and squares represent manifest variables. Error variances (e1 to e6) are drawn as double-headed arrows into manifest variables. The numbers next to single-headed arrows between manifest variables and latent constructs represent standardized factor loadings, which can be interpreted as correlations between the manifest variables and the respective latent construct. The numbers next to the double-headed arrows indicate the correlation between the latent constructs. All loadings and correlations are standardized; * $p < .05$.

Table 1.
Mean (SD) exploration in the three tasks and within-task stability

Measures	Short task			Long task			Corr.
	<i>Mean</i>	<i>SD</i>	<i>Median</i>	<i>Mean</i>	<i>SD</i>	<i>Median</i>	<i>r</i>
Fishing task (No. ponds visited)	10.9	4.77	10	8.79	3.40	8	.82
Bandit task (prop. switching)	0.67	0.28	0.72	0.67	0.28	0.74	.95
Secretary task (prop. options searched)	0.63	0.14	0.64	0.62	0.16	0.64	.62

Note. $N = 261$ in bandit and secretary task and 225 in the fishing task. Corr. reports the correlation of participants' exploration in the short and the long version of a task. In the fishing task short and long refer to the traveling time between ponds (15 s vs. 35 s), in the bandit task to the task horizon (25 vs. 75 choices), and in the secretary task to the number of options (20 vs. 40). Prop. = proportion. All reported correlations are significant at $p < .001$

Table 2


Overview of the SEM models

SEM models	Description	Hypothesis	Model fit					
			χ^2	df	p -value	$RMSEA$	CFI	TLI
Model 1	latent factors are uncorrelated	exploration is completely task specific	18.48	12	0.10	0.05	0.99	0.99
Model 2	latent factors are correlated	exploration is task specific but there is common variance shared across tasks	10.05	9	0.35	0.02	1	1
Model 3	latent factors are identical, i.e. exploration in all three tasks can be described by one latent factor	exploration is completely task general	834.37	9	< .001	0.60	0.16	-0.97

Note: In model 1 and 2, we set the unstandardized factor loadings for all exploration indicators to 1, assuming that the short and the long task equally contribute to the latent factor.

Figure1

A


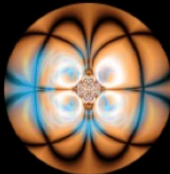


B

At the beginning of each trial, you will be told how many trials are remaining in that round.

34 trials remaining

Press 1, 2, 3, 4, or 5 to choose the corresponding picture.



1

2

3

4

5

Choose

C


Earned Points

20

Number of applicants left:

38

Game Nr. 5 of 30



Applicant Nr: 2

Relative rank: 1

Accept

Reject

