The ecology of competition: A theory of risk–reward environments in adaptive decision making

Timothy J. Pleskac

University of Kansas Max Planck Institute for Human Development

Larissa Conradt, Christina Leuker, & Ralph Hertwig

Max Planck Institute for Human Development

Author Note

Abstract

In many choice environments it seems that risks and rewards or probabilities and payoffs are tightly coupled such that high payoffs only occur with low probabilities. This structure may afford an adaptive mind an opportunity to exploit it—for instance, by inferring the probabilities of different outcomes in the face of uncertainty. However, a mind can only adapt to and exploit an environmental structure if it is frequent and recurrent. Here, we show that the ecology of competition makes low probabilities of high payoffs ubiquitous. They are a consequence of an ecological principle known as the ideal free distribution where the number of competitors in a resource patch is proportional to the gross total amount of resources in the patch. This principle implies a predictable inverse relationship between probabilities and payoffs, the risk–reward structure. Moreover, the ecological theory identifies important boundary conditions for the structure. For instance, heterogeneity of resources in a given patch, computational limits among competitors, and scarcity of resources, systematically distort the risk–reward relationship. Finally, we show that people's representations of the risk–reward structure are consistent with these distortions. Grounding people's inferences in the competitive ecology theory of the risk–reward structure makes it possible to predict when adaptive decision makers should use the risk–reward structure. Such an ability demonstrates that the coupling of a theory of the mind with a theory of the environment is pivotal for understanding and predicting the behavior of a system of bounded rationality.

Keywords: risk–reward, uncertainty, ecological rationality, environment, ideal free distribution

The ecology of competition: A theory of risk–reward environments in adaptive decision

making

> For—believe me—the secret for harvesting from existence the greatest
>
> fruitfulness and the greatest enjoyment is— *to live dangerously*! Build your
>
> cities on the slopes of Vesuvius! Send your ships into unchartered seas! Live at
>
> war with your peers and yourselves! Be robbers and conquerors as long as you
>
> cannot be rulers and possessors, you seekers of knowledge!

—Friedrich Nietzsche, *The Gay Science*

People have many tools they can draw on to traverse the uncertainties they face in the world. One such tool is probability theory. This tool, which Pierre-Simon Laplace (1814/1902) described as "common sense reduced to calculus" (p. 196-197), provides a means to quantify and tabulate the uncertainties people face by assigning a probability to the event thereby turning the uncertainty into a risk (Keynes, 1921; Knight, 1921; Luce & Raiffa, 1957). Yet,because of the computational limitations of the user or limitations of the theory itself, probability theory is not a catchall tool. As an alternative people have available a simple rules to navigate uncertainty. One such rule is the *principle of indifference* whereby the possible events under consideration are treated as equally likely to occur (Keynes, 1921; Laplace, 1902/1814). Another strategy is to use the similarity or representativeness between a possible event and a large class of events as a proxy for the likelihood of it occurring (Silver, 2012; Tversky & Kahneman, 1974). Here we examine a different rule, a rule that draws on the sage advice that great rewards can only be harvested in the face of great risk. This life lesson is proffered not only by philosophers like Nietzsche, but also by financial advisers, gamblers, athletes and artists alike. The ubiquity of this advice suggests that an inverse relationship between risks and rewards, henceforth the *risk–reward structure*, may be a useful structure for the human mind to use to navigate uncertainty.

Indeed this regularity appears to be enlisted by the mind is in terms of a heuristic that Pleskac and Hertwig (2014) called the *risk–reward heuristic*. This cognitive tool helps the mind to reckon with uncertainty—that is, situations in which the probabilities of an event occurring are not given or not known (Keynes, 1921; Knight, 1921; Luce & Raiffa, 1957). Consider a monetary gamble that costs €2 to play and offers €100 if a specific event occurs, but otherwise nothing. The risk–reward heuristic can be used to infer the unknown probability of winning a payoff p from the ratio of the entry costs of playing the gamble c to the total amount of possible winnings as $p = c/(c + g)$ where g is the possible payoff (€100). Thus, if the gamble costs €2 to play, the probability of winning should be approximately 2%.

There is some evidence that people use the risk–reward heuristic to infer missing probabilities (Pleskac & Hertwig, 2014; Leuker et al., 2018a). The heuristic also appears to help inform people's prior beliefs as they learn from experience (Hoffart et al., 2018). Enlisting an environmental structure like the risk–reward structure—be it in a heuristic or within a larger framework like Bayesian inference(e.g., Chater et al., 2010; Gershman et al., 2015; Griffiths et al., 2010) or reinforcement learning(e.g., Daw et al., 2005; Niv, 2009)—is appealing because it makes use of environmental regularities instead of forcing the decision maker to seek further information on the likelihood of an event time and again. The challenge, however, is that cognition can adapt to the environment only if the structure in question has been frequent and recurrent (Barkow et al., 1995). That is, the risk–reward heuristic or any other cognitive process that enlists the risk–reward structure can only operate reliably if the environmental regularity is robust. The central questions we tackle in this paper are is the risk–reward structure a frequent and recurrent structure in choice environments, and can the emergence of the risk–reward structure be predicted *a priori*?

The risk–reward structure does seem common in human choice environments. For instance, Figure 1 illustrates the structure that exists in the game of roulette. It plots, for each unique possible payoff (relative to a $1 bet), the corresponding probability of winning
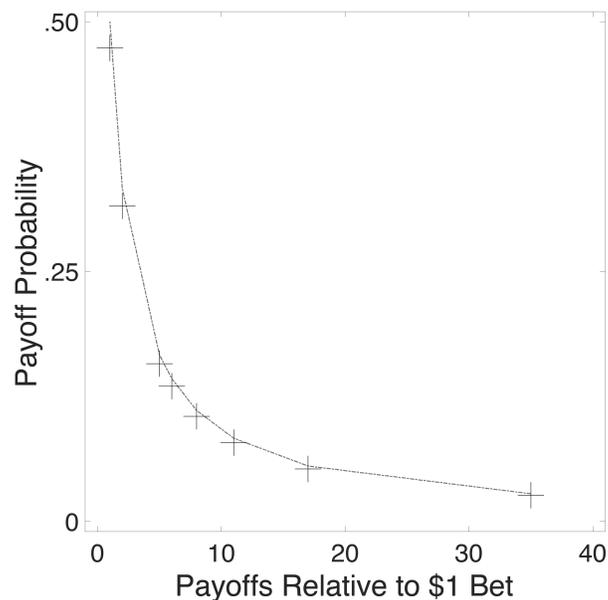
Figure 1. The risk–reward structure in the game of roulette. The dotted line is the estimated probability assuming a fair bet.

it. This risk–reward structure takes a precise form—namely, a hyperbolic function—where the chances of winning are an inverse function of the payoffs. As a result, the probability of winning gets smaller as payoffs increase. A survey of many everyday choice environments showed that the same or a similar regularity exists for bets people take at the horse track, the lottery tickets people buy, the journals scientists submit their article to (trading acceptance rate for impact factor), and the semen dairy farmers purchase to inseminate their cows (trading the semen's conception rate for profitability; Pleskac & Hertwig, 2014).

In these modern environments, the risk–reward structure can be traced back to the forces of the marketplace. Buyers desire high payoffs with the highest of probabilities at the lowest cost. Sellers want the opposite. These opposing forces push lotteries toward a fair bet (see also Samuelson, 1965). Consequently, any increase in a potential payoff would have to be offset with an increase in cost in order to make the gamble attractive for both parties (a Pareto-efficient state). More formally, in this state expected gains correspond to

expected losses such that

$$p \times g = (1 - p) \times c. \tag{1}$$

A gamble represents a fair bet *if and only if* the probability of winning is

$$p = \frac{c}{c + g}. \tag{2}$$

Thus, the probability of a payoff will be inversely related to its magnitude because the expected payoffs and costs of gambles are pushed toward their fair price. The dashed line in Figure 1 plots this fair-bet probability for roulette. In roulette, as they say, the house always wins. This means that the actual probabilities are systematically lower than those predicted by the assumption of fair bets (with a mean absolute deviation of .005). Nevertheless, Figure 1 shows a close correspondence between the actual and fair bet probabilities.

While this kind of gamble may be common, there seems to be a much larger class of situations in which people face competition for resources, such as harvesting food, searching for a mate, or even finding a new home. Even in these situations, which lack a clear institutionalized market with buyers and sellers, it seems that people face a risk–reward structure. In this paper we develop a competitive ecology theory of the risk–reward structure to begin to understand under what conditions a risk–reward structure is expected in choice environments and by extension how frequent and recurrent the structure is. We could attempt to arrive at such a framework inductively, as Pleskac and Hertwig (2014) did, by surveying choice environments and quantifying the frequency of different risk–reward structures. This inductive form of an ecological analysis can be productive and is perhaps the most common course for such an analysis (e.g., Anderson & Schooler, 1991; Brunswik, 1944; Czerlinski et al., 1999; Gibson, 1979; Hertwig et al., 2008; Şimşek & Buckmann, 2015; Stewart et al., 2006; Todd et al., 2012; Ungemach et al., 2011). However, such inductive analyses are piecemeal and make it difficult to predict when a specific environment has a particular structure a priori. These limitations of inductive

ecological analysis also impact the ecological rationality of strategies—that is, conclusions about how adaptive choice strategies are in a given choice environment (Marewski & Schooler, 2011; Rieskamp & Otto, 2006; Todd et al., 2012). For instance, a person must use the risk–reward heuristic in a suitable structure (i.e., a risk–reward structure) in order to be acting in an ecologically rational manner. However, without a principled understanding of when and where a particular environmental structure occurs, there is no normative benchmark to which to compare a person's behavior. As a result, conclusions regarding ecological rationality are in jeopardy of being post hoc rationalizations.

We therefore take a deductive approach in order to identify some basic conditions that can give rise to a risk–reward structure. We focus on competition over limited resources—a common dynamic in many choice environments—and ask under what conditions a risk–reward structure arises. In the classes of situations with which we are concerned, limited resources are distributed across different patches in a resource landscape. We examine what conditions cause a coupling of the size of the resource or the reward with the probability of successfully obtaining it. Moreover, we establish important boundary conditions for this risk–reward structure. Finally, we show how people's beliefs about the risk–reward structure are consistent with some of these theoretically derived properties, and we develop a computational model that captures them. We conclude by discussing how a complete theory of adaptive behavior must go beyond merely describing the environment to which the mind is adapting; it must also contain a principled theoretical account of that environment. It is on the basis of such theoretical accounts of the environment and its structure that adaptive behavior can be predicted.

Limited resources and the distribution of competitors

Our starting point is the observation that often, a critical factor in determining the probability of obtaining a resource is the number of individuals competing for it. The more competitors there are, the less likely it is that any given individual will obtain that

resource. It follows that in seeking to understand the risk–reward structure one must first establish the distribution of competitors an individual will encounter. To do so, we draw on a well-known principle from behavioral ecology known as the ideal free distribution principle (Fretwell & Lucas, 1970). According to this principle, competitors distribute themselves to match the total gross amount of resources within patches (see Davies et al., 2012, Chapter 5).

To get an intuition of the ideal free distribution, consider fishing vessels heading out to catch their limit. Each vessel wants to minimize their effort in catching their limit. There are many locations vessels could search for their fish. It also means vessels have to identify the location they believe has the best chance of meeting this goal and compete with other vessels in getting there first. As vessels do this a dynamic equilibrium forms with effectively the same catch-per-unit effort in each fishing location (Gillis et al., 1993; van der Lee et al., 2013). This relationship also means that the number of fishing vessels in each location is proportional to the amount of fish in each location, thereby creating an ideal free distribution of fishing vessels .

Formally, an ideal free distribution is present when the number of competitors n_y (e.g., vessels) in each patch y (e.g., fishery) is proportional to the amount of resources r_y in each patch (e.g., the gross total amount of fish),

$$n_y \propto r_y. \tag{3}$$

Animals including mallard ducks foraging for bread crumbs in a pond (Kennedy & Gray, 1993), pike searching for food in natural lakes (Haugen et al., 2006), and bumblebees seeking nectar among flowers (Dreisig, 1995) have been shown to distribute themselves in a way that is consistent with ideal free distribution. This relationship between the amount of resources and the number of competitors also extends to other types of resources including mates (Parker, 1974), breeding sites (Alatalo et al., 1992), spawning sites (Smith et al., 2000; Valladares & Lawton, 1991), and nesting sites (Nicolai et al., 2014). An ideal free distribution also appears to occur when resource patches are organized conceptually (e.g.,

flower types for bees) rather than spatially (Chittka et al., 1999; Waser, 1986). And, as we have seen, ideal free distribution also occurs with humans. Besides fishing, people also distribute themselves in a way consistent with ideal free distribution when they look for rewards in behavioral studies (Goldstone & Ashpole, 2004; Goldstone et al., 2005; Sokolowski et al., 1999), and when they apply for jobs (Krueger, 1988; Holzer et al., 1991).

An ideal free distribution of competitors is a group level phenomenon that occurs as animals or humans seek to maximize their individual resource consumption.[1] The ideal free distribution principle rests on two key assumptions: that individuals are *ideal* and that they are *free*. The first assumption means that individuals can detect the patches of resources that permit the highest rate of consumption. The second assumption means that individuals can move between patches with no restrictions or costs. Both assumptions are idealized descriptions of individuals and the conditions they face, which makes it important to study how violations of these assumptions impact the predictions of the ideal free distribution principle (e.g., Kennedy & Gray, 1993; Sutherland et al., 1988; Tregenza, 1995). Before we turn to those violations, we show how an ideal free distribution of competitors can imply a risk–reward structure.

The ideal free distribution and the risk–reward structure

To determine what the ideal free distribution principle implies about the relationship between risk and reward we need to establish how the size of the resource s_y (i.e., reward) and the *success probability* of obtaining the resource p (i.e., risk) covary. To do so, we need to break down the total amount of resources r_y available in a given patch into the number

---

[1] The ideal free distribution principle bears some resemblance to Herrnstein's (1970; 1974) matching law, where an individual tends to equate the rate of a response to the rate of reinforcement. The crucial difference is that the matching law concerns behavior at the individual level, while ideal free distribution refers to group level behavior (see also Houston & McNamara, 1988).

of resources m_y multiplied by the size of the resources,

$$r_y = m_y \times s_y. \tag{4}$$

Thus, formally the ideal free distribution principle can be rewritten, substituting Equation 4 into 3, as

$$n_y \propto m_y \times s_y. \tag{5}$$

The success probability p_y will, holding all else constant, increase with the number of resources m_y in the patch. Thus, within a limited time window the success probability in a patch is proportional to the number of resources in it, $p_y \propto m_y$. Furthermore, the success probability will, holding all else constant, decrease with the total number of competitors in the same patch, $p_y \propto 1/n_y$. Taken together, these assumptions imply that the success probability is proportional to the ratio of the number of resources to competitors in a patch

$$p_y \propto \frac{m_y}{n_y}. \tag{6}$$

The link between success probability and resource size is established by substituting Equation 5 into Equation 6. This reveals that the success probability is inversely proportional to the size of the resource in a patch,

$$p_y \propto \frac{1}{s_y}. \tag{7}$$

This means that if there is variation in the size of resources between patches but not within patches (e.g., because different patches have different sizes of fish), then competitors like our fishing vessels face a choice between patches that trade resource size for success probability. Thus, a basic ideal free distribution of competitors predicts that the success probability is inversely proportional to resource size across a landscape: a risk–reward structure.

Three points are worth emphasizing. First, the fact that an ideal free distribution can imply a risk–reward structure is not limited to markets with clear buyers and sellers.

Rather it is a structure that one can expect to emerge whenever competitors vie for a finite resource that is spread across a landscape. Second, demonstrating that an ideal free distribution can imply a risk–reward structure moves beyond merely describing an environment agents are adapting to—a common stipulation of adaptive theories of cognition (Simon, 1990). Instead it establishes a set of sufficient conditions that can give rise to a risk–reward structure and help establish how frequent and recurrent the structure is. Third, by systematically changing the assumptions, the theory makes it possible to examine when and how the risk–reward structure will change and even break down as a result of changes in the ecology.

The Risk–Reward Structure and Changes in the Ecology of Competition

A core aspect of adaptive behavior is that an individual makes use of an ecological structure only when it is present in the environment. Therefore, in order to predict adaptive behavior one must first establish the conditions for when the ecological structure will be present and when it will be weakened or disappear. The competitive ecology theory helped identify how changing eight different properties of the ecology affect the predicted risk–reward structure.

Two conditions are relatively straightforward. One relates to the temporal dynamics of an ideal free distribution. An ideal free distribution of competitors is expected to be present when the average consumption of resources is equal across patches. If the system has not yet, or cannot, reach such an equilibrium, then the distribution of competitors should not conform to an ideal free distribution and a reliable inverse relationship between the success probability and resource size is not expected (see Table 1, assumption 1).

The second relatively straightforward condition is the availability of resources. An ideal free distribution of competitors is only expected when resources are limited. With unlimited resources (e.g., air, sunlight) competitors do not have to maximize their consumption of resources and no ideal free distribution develops. This also means that a

risk–reward structure—which is caused by competition over a resource—should not be expected for unlimited resources. The reason is that for unlimited resources the success probability will be essentially one for all resource sizes (see Table 1, assumption 2). Both of these conditions—the system at equilibrium and availability of resources—thus establish boundary conditions for the emergence of a risk–reward structure. Table 1 lists six additional boundary conditions. Next, we analyze each of these boundary conditions and quantify how the risk–reward structure systematically changes.

Table 1

Boundary conditions of a risk–reward structure in the ecology of competition over limited resources.

| Ideal Free Distribution Assumption	Violation of the Assumption	Does Risk–Reward Structure Still Hold?
(1) System is in equilibrium.	Rate of resource consumption is not equal across patches.	When the system is not in equilibrium, an ideal free distribution does not hold and no reliable risk–reward structure is expected.
(2) Limited availability of resources.	Resources are unlimited.	When resources are unlimited an ideal free distribution does not hold and no risk–reward structure is expected.
(3) Homogeneous resource sizes within patches.	Heterogeneity in resource size within patches.	As within-patch variability in resource size increases, the success probability equalizes across all sizes, reducing the risk–reward structure.
(4) Landscape saturation of competitors and/or resources.	Competitors and/or resources are sparsely distributed across the landscape.	As landscape saturation decreases the risk–reward structure weakens and becomes more disperse at low resource sizes.
(5) Competitors are ideal.	Competitors' ability to discriminate between the quality of the patches is limited.	As the ability to discriminate decreases the risk–reward structure weakens and becomes more disperse at low resource sizes.
(6) Competitors are equally competitive.	Competitors have different competitive abilities.	The risk–reward structure holds within categories of abilities, but not across categories.
(7) Competitors do not interfere with each other in the acquisition of resources.	Interference is possible in obtaining resources.	The risk–reward structure holds, but will take longer to emerge.
(8) Competitors only require a single resource type.	Competitors require multiple different limited resources.	The more limited and limiting the resource, the more reliably the risk–reward structure can be expected to emerge.

Heterogeneity in resource sizes

Although the assumption that resource sizes are constant within patches is a plausible approximation in some cases (e.g., fruits and vegetables), there are also situations where resource sizes can vary within patches. Indeed, Appendix A proves that for landscapes made up of two resource sizes (large and small), when the patches are more heterogeneous in terms of resource size the success probability degrades such that at maximal heterogeneity of resource size there will be no risk–reward structure.

Using simulations, we analyzed how heterogeneity in resource size impacts the risk–reward structure, using more than two resource sizes. The results of the simulations are plotted in Figure 2 (for details see Appendix A). At minimal within-patch heterogeneity of resource size, the success probability is inversely related to resource size (panel a). At maximal heterogeneity, the success probability is identical for all resource sizes (panel d). For intermediate levels of heterogeneity, the risk–reward structure falls within these two extremes (panels b and c). An important observation is that in all of these cases, even as within-patch heterogeneity increases, an ideal free distribution is still present. Thus, an ideal free distribution and a risk–reward structure in the wake of competition over finite resources are not synonymous.

Landscape saturation

Another simplifying assumption of the ideal free distribution principle is that the landscape is sufficiently saturated in terms of both resources and competitors. If the number of competitors or the total number of resources is small relative to the number of patches, there can be patches that attract relatively few or even no competitors. This change in the saturation of the landscape—via either the number of resources or the number of competitors—impacts success probabilities. To find out how it impacts the risk–reward structure, we simulated four different landscapes with different levels of saturation of competitors and/or resources. The results of the simulation are shown in
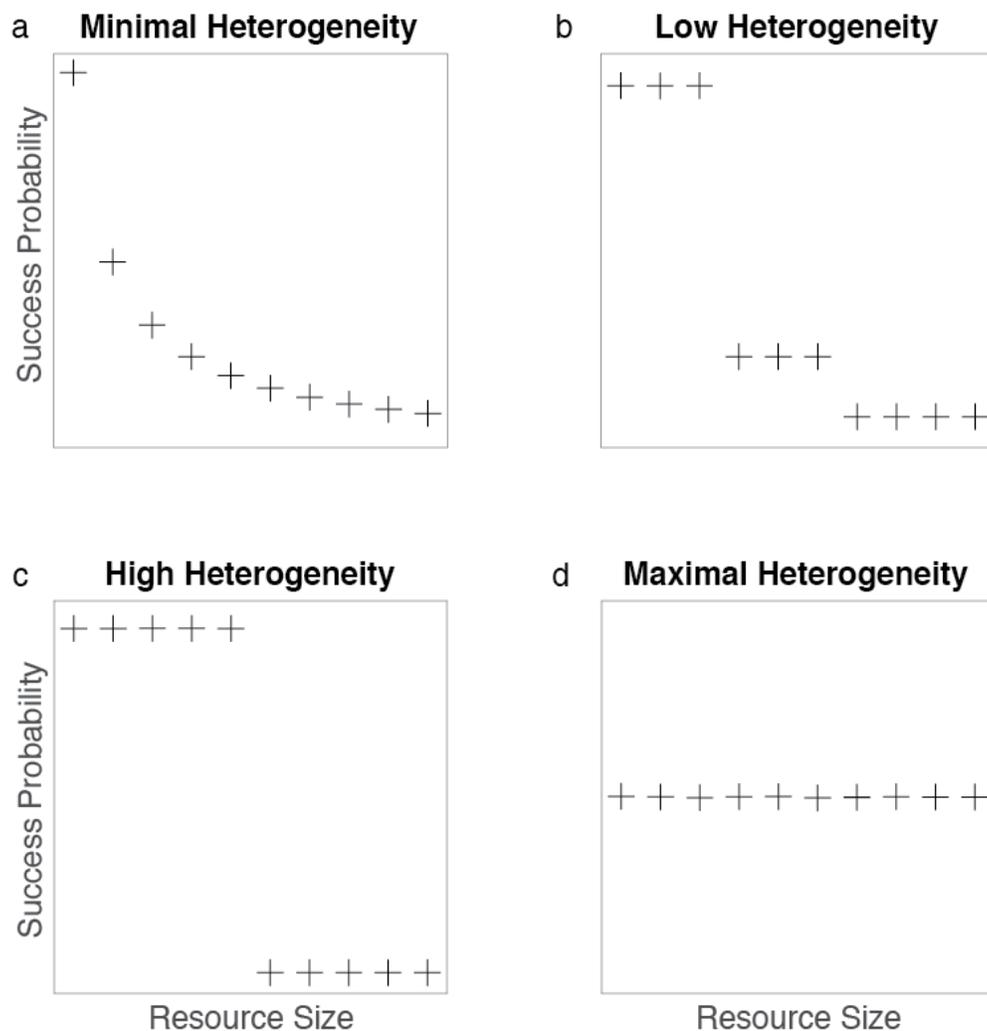
Figure 2. Success probability as a function of resource size in landscapes with different levels of resource size heterogeneity within patches. Within-patch heterogeneity of resource size increases from panels a to d. For details on how the probabilities were estimated see Appendix A.

Figure 3. In general, the risk–reward structure is robust. Figure 3, however, shows that there are some quantitative changes to the structure. Comparing the top two panels (with low-competitor saturation in the landscape: 100 total competitors across 200 patches) to the bottom two panels (1,000 total competitors across 200 patches), shows that with fewer

competitors and thus lower saturation the slope of the function is shallower. Moreover, the variation in the success probability for small resource sizes is higher when either the total number of competitors is small, or resources are sparse. Thus, in less saturated landscapes, there is no longer a guarantee that small resources will always be associated with a high success probability.

The increasing variability in success probability as resource size gets smaller means resource size becomes increasingly unreliable as a predictor of success probability. Our simulations show that this problem does not arise to any noteworthy degree with respect to estimating the success probability for large resources: Even in sparse landscapes, the variability in success probability is low for large resources (Figure 3). For instance, in the simulation with low competitor saturation and low resource saturation the range of success probabilities for the smallest resources is 3.6 times larger than that for the largest resources. As we see next, a similar regularity emerges when the cognitive constraints of competitors are taken into account.

Non-ideal individuals

Perhaps ideal free distribution principle's boldest assumption is that competitors are ideal—that is, that they are sensitive to and able to move to patches that maximize their consumption rate. Like the proposed behavior of homo economicus, this optimal behavior requires a complete representation of the environment and unlimited cognitive resources (Gigerenzer & Selten, 2002; Simon, 1955). What are the consequences of a more realistic portrayal of competitors? Abrahams (1986) demonstrated that if competitors have limited ability to detect differences in the quality of a patch, relative to landscapes with ideal competitors, real competitors will overuse poor patches and underuse rich patches. This pattern suggests the following reformulation of the ideal free distribution principle from Equation 3 to

$$n_y \propto r_y^{\alpha}. \tag{8}$$
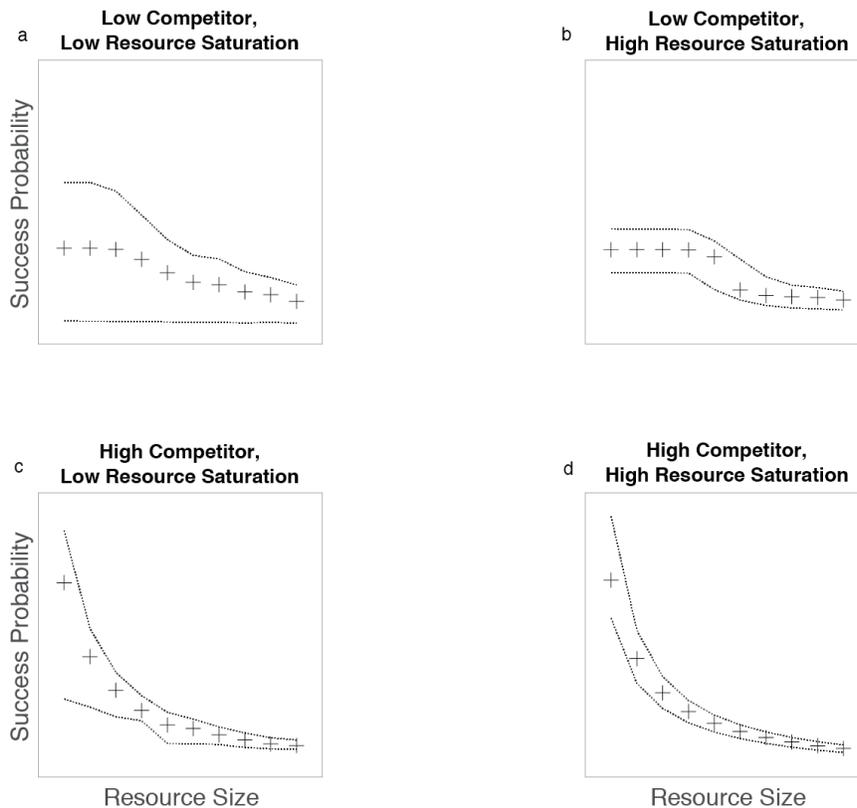
Figure 3. Success probability as a function of resource size in an ideal free distribution

model with different levels of landscape saturation. The median success probability ($+$)

decreases with resource size. The fewer the number of competitors within the landscape,

the shallower the decrease. The dashed lines plot the 2.5% and 97.5% quantiles of success

probabilities, showing that the fewer competitors and resources in the landscape, the

greater the variability in success probabilities—particularly at smaller resource size. In

each simulation there were 200 patches with either 100 (low) or 1,000 (high) competitors

and either 100 (low) or 1,000 (high) resources per size. The results are based on 2,000

simulations per parameter combination. See Appendix B for further details.

The parameter α $(0 < \alpha < 1)$ captures the degree to which competitors, in the aggregate,

mismatch with respect to the resource amount. As sensitivity to patch quality at the

individual level increases, the distribution of competitors increasingly conforms to an ideal free distribution, reflected in an increase in α. Kennedy and Gray (1993) found that across 52 animal studies, an estimate of $\alpha = 0.66$ gave the best account of the data.

Diminished sensitivity to the quality of the patches changes the distribution of competitors, thereby impacting the risk–reward structure. To see how, recall that the amount of resources in a patch is $r_y = m_y \times s_y$ (Equation 4). As a result, the predicted distribution of competitors will be

$$n_y \propto (m_y \times s_y)^{\alpha}. \tag{9}$$

Using this more general version of the ideal free distribution principle, we can determine how nonideal individuals will impact the success probability. This is done by incorporating Equation 9 into the assumption where $p_y \propto m_y/n_y$ (Equation 6). Doing so shows that the success probability with nonideal individuals is

$$p_y \propto \frac{m_y^{1-\alpha}}{s_y^{\alpha}}. \tag{10}$$

As Equation 10 shows, once the limited sensitivity of competitors is acknowledged, success probability p_y is proportional to both the size of the resource s_y and the number of resources m_y. If there is no variability in the number of resources across patches, then m_y is a constant and thus in terms of proportionality will be factored out of the equation. Consequently, in landscapes where each patch has the same number of resources but competitors are nonideal, success probability is less than inversely proportional to resource size, $p_y \propto 1/s_y^{\alpha}$. This means that relative to the risk–reward structure with ideal competitors, nonideal competitors produce success probabilities that will be less extreme for small resources and more extreme for large resources. In other words, a landscape with nonideal competitors produces a flatter risk–reward structure than one with ideal competitors (Figure 4).
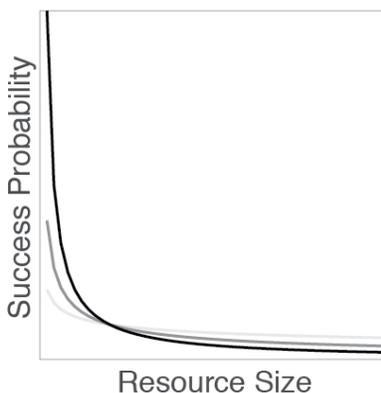
Figure 4. Success probability as function of resource size for different levels of sensitivity to patch quality α ($\alpha = 0.3, .6,$ and $.8$). As sensitivity to patch quality lessens and competitors become less ideal (shown above with increasingly lighter lines), success probabilities become less extreme for small resources and more extreme for large resources.

We can also examine what happens if the number of resources per patch varies between patches. If the median number of items per patch is similar for patches with different resource sizes, the median success probability still decreases with resource size s_y, just as before, but it will be shallower than what is predicted for a landscape with ideal individuals (Equation 10). However, because there is variability in the number of resources there will be variability in the success probability for a given resource size. To see the impact of the variability, assume, for example, that the 95% range of number of resource items per patch ranges from $\dot{m}$ to $\ddot{m}$. In landscapes with nonideal individuals, Equation 10 implies that the 95% range of the predicted success probabilities will range from values proportional to $\dot{m}^{1-\alpha}/s_y^\alpha$ to $\ddot{m}^{1-\alpha}/s_y^\alpha$. Figure 5 shows how nonideal individuals in these landscapes change the range in possible success probabilities as a function of resource size. The pattern is similar to what we saw in sparse landscapes (Figure 3). The success probability associated with a small reward item now varies widely, so that in a large number of cases the success probability will be low. In contrast, the success probability for a large resource item is less variable and nearly always low.
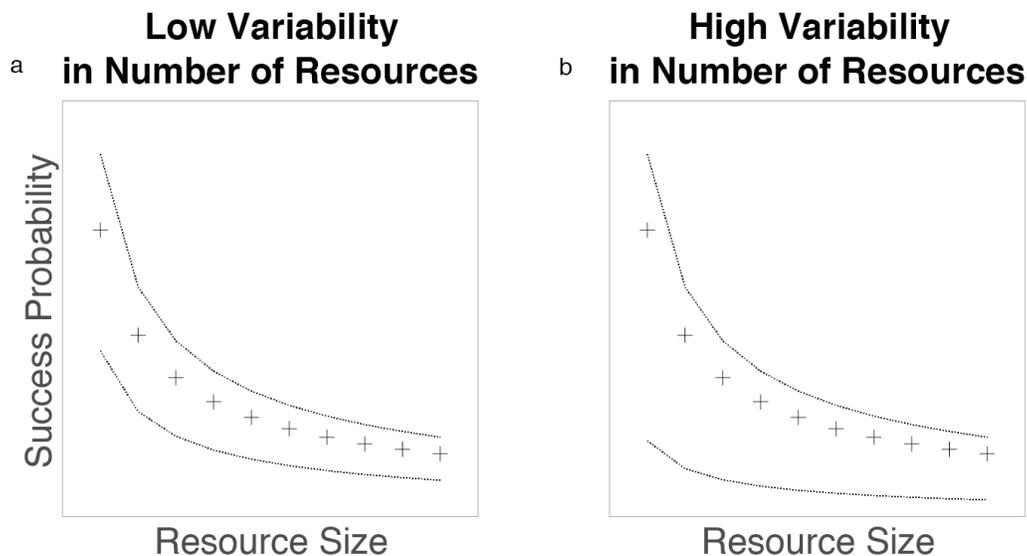
Figure 5. Success probability as function of resource size when constraints lead to competitors undermatching patch resources. The median success probability (+) decreases with resource size. The 95% ranges of success probabilities at each resource size (dashed line) depends on the variation in the number of resource per patch (left panel: $\dot{m} = 100$ to $\ddot{m} = 1000$; right panel: $\dot{m} = 10$ to $\ddot{m} = 1000$).

In sum, limited ability to discriminate between patches of different qualities results in a risk–reward structure that is less extreme for small resources and more extreme for larger resources than one that occurs in landscapes with an ideal free distribution. Furthermore, if there is variability in resource sizes across patches, then these cognitive constraints result in an asymmetry in the relationship between risks and rewards. In particular, echoing a similar pattern with low landscape saturation, large resource sizes will reliably have a low

success probability while small resource sizes will no longer reliably guarantee a high success probability.

Unequal abilities between competitors

The basic ideal free distribution principle assumes that all individuals are equal in their ability to obtain a resource. This is obviously a gross simplification; in reality competitors often differ (Houston & McNamara, 1988; Parker & Sutherland, 1986; Sutherland et al., 1988). We show in detail in Appendix C that if individuals differ in their competitive ability, an ideal free distribution of competitors results in a risk–reward structure *within* each class of competitors with the same abilities, but not between classes. The intuition behind this result is that competitors with greater ability are more likely to obtain a resource, regardless of the resource size. However, within the same class of competitors the success probability is inversely proportional to number of competitors, which results in a risk–reward structure. A corollary of this result that the risk–reward structure will appear weaker in these environments when ignoring or collapsing across abilities or classes.

Interference between competitors

The basic ideal free distribution principle assumes that competitors do not interfere with each other's ability to acquire a resource. However, in many situations individuals do interfere with each other's resource acquisition (Sutherland, 1983), including for reasons other than competition (e.g., predators may frighten off prey). Nevertheless, interference does not alter the risk–reward structure. For a formal proof see Appendix D. The intuition behind the result is that interference impacts both the distribution of competitors in a patch and the success probability of any given competitor. Consequently, interference between competitors cancels itself out, thus retaining the risk–reward structure.

Multiple resource types

Finally, what happens to the risk–reward structure when competitors in a landscape require more than one type of resource (e.g., water, food, mates)? Multiple resource types may render it impossible to achieve equilibrium across all types of resources simultaneously. For example, if most of the food is located in specific patches and most of the mating opportunities are in other patches, then the distribution of competitors between those patches can reflect either the goal of consuming food or the goal of finding of mates or some compromise thereof. How does this affect the risk–reward structure? A possible answer comes from applying what is known as Liebig's law in behavioral ecology (Bloom et al., 1985; Danger et al., 2008; Farrior et al., 2013; Gorban et al., 2011; Harpole et al., 2011; Rosenheim et al., 2010). According to Liebig's law, growth is not determined by the total amount of resources, but by the most limiting one. The more limiting (i.e., essential and scarce) a resource type is, the more likely it is that its distribution dictates the distribution of competitors. For example, in a drought, water might be a more limiting resource, whereas in a wet season, mates or food might be more limiting. Consequently, the more limiting a resource is, the stronger the risk–reward structure for that resource.

In summary, the competitive ecology theory establishes how competition over limited resources results in the emergence of a risk–reward structure. Building on the concept of ideal free distribution from behavioral ecology, its first contribution is to establish that an inverse relationship between the size of a resource and the probability of obtaining it is to be expected whenever there is competition over finite resources. The second contribution of our theory is to establish how robust the risk–reward structure is against changes in the ecology of competition. Specifically, we showed how changes in resources across the landscape and in competitors systematically affect the risk–reward structure. This structure proves to be relatively robust: The strength of the relationship weakened, but never reversed. If the risk–reward structure is weakened, the success probabilities of smaller rewards are more affected than those of larger rewards. We now examine the extent to

which people's beliefs in the risk–reward structure reflect some of the predicted properties.

Do people's beliefs reflect the ecological properties of a risk-reward structure?

The competitive ecology theory is a theory about an environmental structure. An adaptive approach to cognition like the one we take suggests that the mind reflects key properties of the environment. Competition over finite resources is frequent and recurrent, and so is the risk–reward structure resulting from this competition. Moreover, our theory establishes that two properties can be commonly expected in the risk–reward structure as a result of a competitive ecology with realistic assumptions. First, the relationship between risks and rewards can be expected to be less extreme for small resources and more extreme for larger resources. As we have shown, this property arises both due to heterogeneity in resource sizes and when competitors have diminished sensitivity to patch quality. Second, the resulting risk–reward structure in a landscape will often be one where the variability in success probabilities increases as resources gets smaller. This heteroscedasticity in estimates arises as saturation in the number of competitors or number resources decreases, or when competitors are not perfecting discriminate between resource sizes. We examined whether these two properties were present in people's beliefs in an estimation task first used by Pleskac and Hertwig (2014). There are subtle variations the task, but all the tasks asks participants to consider a bet like the following:

Imagine you have been asked to play the following lottery. The lottery offers the opportunity to win € x, but it costs you € 2 to play. If you choose to play you would pay the € 2 and, without looking, draw a ball from a basket. In the basket there are 1,000 balls. The balls are either black or red. If the ball is red you will win € x; otherwise, if the ball is black you will receive nothing. Thus, the number of red balls in the basket determines the probability that you will win. You are less likely to win the € 2 the lower the number of red balls in the basket.
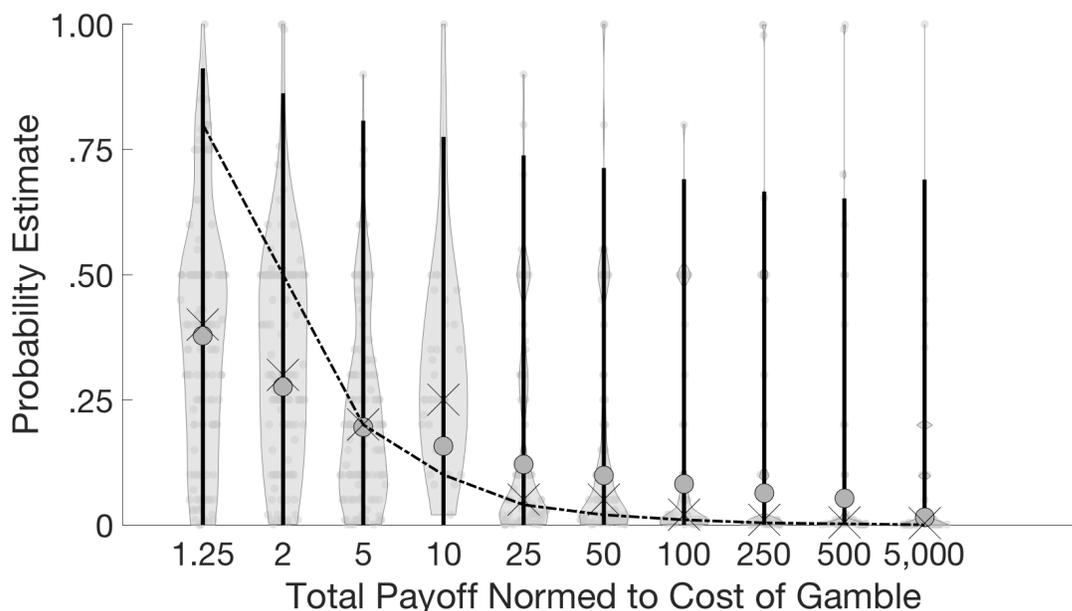
Figure 6. Probability estimates by reward magnitudes from the lottery task. The median observed estimates are denoted by ×. The light gray dots denote individual participant estimates. The median of the posterior predicted distribution from risk–reward estimation model is denoted by the large dot and the 95% HDI of the posterior predicted distribution is denoted by the solid black line. The estimates are shallower than an inverse risk–reward relationship (dashed line) and the variability in the estimates reduces as the reward magnitudes increase. Both of these properties are consistent with the properties derived from our framework.

Different participants saw different total payoffs x (e.g., 2.5, 4, 10, 20), and were asked two questions: (1) How many balls do you think are in the basket? and (2) Would you pay € 2 to play?

Here we focus on how participants revealed their beliefs in the first question. We draw on responses across three studies. The first two come from Pleskac and Hertwig (2014): In the first study, participants ($N = 138$) could win and lose real money with the bets (in a behavioral laboratory) and the bets were about an urn with 100 balls. In the second study participants ($N = 196$) answered questions about a hypothetical bet (online)

Table 2

Interquartile range (IQR) of the probability estimates at different payoff levels.

Payoff	IQR (estimates)
1.25	.30
2	.35
5	.20
10	.30
25	.29
50	.19
100	.10
250	.10
500	.10
5000	.02

about an urn with 1000 balls. The third study is a new online study where participants ($N = 455$) also answered questions about a hypothetical bet with 1000 balls (see Supplementary Materials for full details on the third study).

Figure 6 plots the estimates collapsed across the three studies. Participants' estimates (in light gray) reflect an inverse risk–reward structure. Consistent with the ecological theory's prediction that the risk–reward structure is flatter with nonideal competitors, participants' estimates were flatter than predicted by a perfect inverse risk–reward structure (dashed line). Moreover, the estimates also became more variable as the reward size decreased, as indicated by the interquartile range of the probability estimates listed in Table 2. To formally test the degree to which the properties identified in our theory are reflected in people's beliefs, we developed the *risk–reward estimation model*.

Risk–reward estimation model

The risk–reward estimation model predicts the estimates participants provide in a pay-to-play gamble like the scenario described above. The model consists of four core assumptions: that people use the risk–reward heuristic, that ecological distortion of probability estimates occurs, that people employ the principle of indifference, and that estimates vary.

Risk–reward heuristic. The first assumption is that people estimate the probabilities using the risk–reward heuristic,

$$p(x) = c/x. \tag{11}$$

The variable c is the cost of playing (i.e., \$1 or € 1) and x is the total payoff (typically including the cost of the gamble c plus the gain g).

Ecological distortion of probability estimates. Second, as we have established, the risk–reward structure people experience may not be a perfect inverse relationship. For instance, as Figure 4 shows when competitors have diminished sensitivity to patch quality the success probabilities are less extreme for small resources and more extreme for large resources. To capture these differences in people's beliefs, the risk–reward estimation model allows the estimated probabilities from the risk–reward heuristic $p(x)$ to be distorted so that the distorted estimate $p'(x)$ is

$$p'(x) = \frac{\delta p(x)^\alpha}{\delta p(x)^\alpha + [1 - p(x)]^\alpha}. \tag{12}$$

The α parameter measures the degree to which the risk–reward relationship is flatter than a perfect inverse risk–reward relationship. As α decreases, the risk–reward relationship becomes flatter. That is, the probabilities associated with low payoffs are less extreme and those associated with high payoffs are more extreme. The δ parameter measures an overall bias in the estimates. A decrease in δ results in an overall downward shift of the

probabilities.[2]

The principle of indifference.   Third, across several studies with the risk–reward estimation task, we have noted that there is a small but noticeable cluster of estimates at .5. This is consistent with people using the principle of indifference, whereby if people are uncertain about the probability of an event they estimate it to be equally likely to occur or not (Keynes, 1948; Laplace, 1814/1902). There is good evidence that people sometimes rely on this rule when judging the probability of an event (Fox & Clemen, 2005; Fox & Levav, 2004; Fox & Rottenstreich, 2003). With this in mind, we equipped the risk–reward estimation model to capture the rate at which people use the principle of indifference rather than relying on the risk–reward heuristic.[3] According to the model, with probability $0 \leq \lambda \leq 1$ people rely on the principle of indifference; otherwise they are assumed to use the risk–reward heuristic

$$p''(x) = \lambda \times .5 + (1 - \lambda) \times p'(x). \tag{13}$$

Variability in estimates.   Finally, we sought to capture the variability in people's estimates. Across a range of conditions as resource size decreased the variability in the success probabilities increased (see Figures 3 and 5). To test whether this regularity is also reflected in people's estimates, and to account for people's baseline noise in their estimates (Erev et al., 1994), we assume that the observed estimates are distributed according to a beta distribution, with a mean of

$$\mu = p''(x). \tag{14}$$

———

[2] Bias has been examined in the context of the ideal free distribution principle, but no systematic bias was found (Parker & Sutherland, 1986).

[3] From an ecological perspective, we might expect an increased rate of the use of the principle of indifference when the environmental conditions are not expected to create a risk–reward structure (e.g., when the system is out of equilibrium or when the resource is unlimited).

The beta distribution has a second parameter ϕ called the dispersion parameter.[4] It determines the spread of the distribution or how noisy people are in reporting their estimates. We parameterized ϕ so that with increasing reward magnitudes the variability in estimated probabilities can decrease. Formally, we set

$$\phi = \exp(d_0 + d_1 x). \tag{15}$$

The parameter d_0 is a free parameter that captures baseline variability in the estimates. The parameter d_1 captures the degree to which the dispersion or spread of the estimates changes across payoff levels with $d_1 = 0$ indicating no change in the dispersion across the payoff levels. When $d_1 < 0$ then there is an decrease in dispersion across payoff levels and when $d_1 > 0$ there is an increase in dispersion.

Capturing ecological properties with the risk–reward estimation model

We fit the risk–reward estimation model to the data using Bayesian estimation techniques (see Supplementary Materials for details and the OSF website for code). Briefly, the estimation model provides a better fit than a standard beta regression model that models the probability estimates as a linear function of payoffs (see Supplementary Material for a model comparison).

Risk–reward heuristic.   Generally, Figure 6 shows that there is a close correspondence between the posterior predicted distributions of the model and the observed estimates: Participants' estimates are consistent with the use of the risk–reward heuristic. The posterior predictions also show that the risk the model appears to capture some of the deviations from the perfect inverse relationship between payoffs and probabilities. But the parameters are perhaps more revealing in terms of the remaining predictions of the competitive ecology theory, as we now show.

---

[4] Typically the beta distribution is characterized by the two independent parameters a and b where $a = \mu\phi$ and $b = (1 - \mu)\phi$. The variance of the beta distribution is $\sigma^2 = \mu(1 - \mu)/[\phi + 1]$. Thus, the larger ϕ is, the lower the variance.

Table 3

Estimated parameters for the risk–reward estimation model.

Parameter	M	HDI
Sensitivity α	0.30	[0.18, 0.44]
Bias δ	0.42	[0.33, 0.52]
Rate of principle of indifference use λ	.20	[.06, .33]
Baseline variability d_0	0.63	[0.73, 0.54]
Effect of payoff magnitude on variability d_1	-0.10	[-0.19, -0.01]

HDI denotes the 95% highest density interval of the distribution. The values

for d_0 and d_1 are in terms of standardized payoff values.

Distortion in the risk–reward relationship.   Table 3 summarizes the posterior estimates of the risk–reward estimation model. In terms of the distortion prediction, the parameter estimates for α are below 1, which is consistent with a flatter risk–reward relationship than a perfect inverse risk–reward relationship (dashed line in Figure 6).[5,6]

Variability in the risk–reward relationship.   The d_1 parameter indicates a small but credible increase in the dispersion of the estimates as the payoff magnitude decreased (Table 3). According to the model, the dispersion for the highest payoff level (5,000) ($M = 0.92$; HDI $= [0.34, 1.61]$) was on average half as large as the dispersion for the lowest payoff level (1.25) ($M = 1.94$; HDI $= [1.75, 2.13]$) ($M = 0.48$; HDI $= [0.17, 0.84]$).

Use of the principle of indifference.   The risk–reward estimation model also reveals that not all participants relied on a risk–reward structure to estimate the

―――――

[5] A special case of the risk–reward estimation model where α and δ were set to 1 provided a substantially worse fit to the data (see Supplementary Material).

[6] The values for δ also show that the estimates were biased away from the perfect inverse relationship in a manner consistent with individuals being pessimistic in their estimates. Lacking an a priori prediction for this result we do not interpret it further.

probability of winning. An estimated λ parameter of .20 indicates that the principle of indifference had some weight in the estimates. This weight is notable in Figure 6 where there is a bump at .5 across the different payoff levels. Given the between-subjects nature of the studies it is impossible to identify whether this effect is due to differential use of the principle of indifference or the risk–reward heuristic, or if the effect is more the result of an integration between the two. Regardless, the model shows that the risk–reward structure carries substantially higher weight in participants' estimates.

Summary

In summary, we have shown that not only do people's beliefs in the probability of obtaining different payoffs adhere to the risk–reward structure, they also reflect two core predictions of our ecological theory of the risk–reward structure. First, beliefs about the risk–reward relationship were, relative to a perfect inverse risk–reward relationship, was less extreme for low payoffs and more extreme for high payoffs. This property is consistent with the competitive ecology theory of the risk–reward structure, according to which this pattern arises in landscapes with heterogeneity in resource sizes or due to nonideal competitors. Second, participants' estimates were increasingly variable with decreasing payoffs. This, according to our theory, is a common property due, for example, to sparsity of resources or competitors, or to nonideal competitors. We now explore the behavioral consequences of both of these aspects in the General Discussion.

General Discussion

Herbert Simon's (1990) principle of bounded rationality states,

Since we can rarely solve our problems exactly, the optimizing strategy suggested by rational analysis is seldom available. We must find techniques for solving our problems approximately, and we arrive at different solutions depending on what approximations we hit upon. Hence, to describe, predict

and explain the behavior of a system of bounded rationality, we must both construct a theory of the system's processes and describe the environments to which it is adapting." (p. 7)

Simon's principle has led to extensive work on how attention, learning, memory, affective processes, and other factors shape people's judgments and decisions (Weber & Johnson, 2009). At the same time, the principle has promoted an extensive cataloging of many of the choice environments in which these judgments and decisions are made (Hertwig et al., 2013; Hogarth & Karelaia, 2007; Todd et al., 2012). A weakness of this cataloging approach, however, is that one may end up with a different process for each discernible environment or environmental structure, resulting in a multitude of descriptions of environment-by-process associations. In order to avoid such "description inflation," we suggest reframing Simon's goal as follows: In order to describe, predict, and explain the behavior of a system of bounded rationality, one must construct theories of the system's processes as well as theories of the mechanisms behind the environmental structures to which the system is adapting.

The competitive ecology theory of the risk–reward structure is a step in this direction. It shows that the ecology of competition is a sufficient condition for linking risks with rewards, and therefore—because competition for limited resources is a ubiquitous condition—a widespread and recurrent environmental regularity. Importantly, our theory also identifies how differences in the ecology, whether in environmental conditions (e.g., limitedness of resources) or in the competitors' abilities, alter the coupling of risk and reward. For instance, the more limiting a resource is, the more reliably one can expect a coupling. Furthermore, the relationship is stronger with more competitors and with competitors who are better able to discriminate between different rewards. The relationship is also stronger the more clumped the resources or rewards are (i.e., divided into distinct patches or categories), and the more homogeneous in size the resources are within a patch relative to between patches. Finally, while in most cases a large reward

corresponds to a low probability, a smaller reward is not reliably tied to a high probability. This implies that people should be wary of banking on a large probability when presented with a small reward.

How an ecological theory can help make sense of human behavior

Identifying the specific forms the risk–reward structure takes and the causes for variation helps, in turn, to better understand human behavior. Empirically, we found that people's representations of the risk–reward structure display properties that are predicted by the ecological theory. Not only is the reported risk–reward relationship flatter, it also exhibits greater dispersion in estimated probabilities as the payoff magnitudes decrease (see Figure 6). This is consistent with the predictions of the ecological theory that suggests various reasons for the relationship to be flatter than in an ideal world.

How do people come to reflect these ecological properties? We do know that people are able to learn risk–reward structures incidentally as they make decisions about risky and uncertain prospects without feedback and when learning is not the central goal (Leuker et al., 2018a). This may suggest that the mind is ready to learn this relationship and perhaps even expects it. But learning is only one mechanism through which people adapt to the risk–reward structure. Social transmission may be another (Canini et al., 2014; Moussaïd et al., 2015). Indeed, the cross-cultural occurrence of proverbs like "A bird in the hand is worth two in the bush" is an indicator that the risk–reward structure has entered our collective mind. The German version of this proverb has it that a sparrow in the fist is better than a pigeon on the roof, but the gist remains the same: It is better to hold on to something than to risk losing it by trying to attain something better. Expressed in the probabilistic terminology of the risk–reward structure: One is better off contenting oneself with a smaller and safer reward than striving for a larger but less likely reward. Understanding how people pick up recurrent and stable regularities in their environment via individual, social, and cultural learning (Henrich & McElreath, 2003; Mesoudi, 2011;

Richerson & Boyd, 2005) will be important next steps in understanding how the risk–reward structure and other structures impact human behavior.

The competitive ecology theory can be informative for other uses of the risk–reward structure as well. For example, since the risk–reward relationship is bidirectional, it can also be exploited to make inferences from known risks to the magnitude of rewards (Skylark & Prabhu-Naik, 2018). Like it does for the risk–reward heuristic, the theory identifies conditions where inferences about rewards based on the risks are impacted by the ecology. For instance, the systematic change in the variability in success probabilities affects an individual's ability to infer a reward from a given risk. In particular, small probabilities of success do not guarantee a large reward size, but large probabilities of success will be more likely to be linked to small resources (see Figures 3 and 5). In other words, an object that is easy to obtain is probably of low value but an objects that is harder to obtain does not promise high value.

A conjecture of ecological rationality requires an ecological theory

The concept of ecological rationality (Todd & Gigerenzer, 2007) highlights the fit between a heuristic or, as Hertwig, Pleskac, Pachur, and the Center for Adaptive Rationality (in press) propose, of any decision-making tool and an environment. To reach a comprehensive understanding of this fit, however, an ecological theory like the one presented here is required. Consider the initial work on the risk–reward heuristic (Pleskac & Hertwig, 2014). Across the different ecologies surveyed, there was a fair amount of variation in the exact shape of the risk–reward structure. Moreover, when they evaluated how well people's beliefs tracked the risk–reward relationship again there was substantial variation. Although estimates obeyed the predictive qualitative pattern, people's estimates did not perfectly map onto the prediction of the risk–reward heuristic. Our ecological theory, however, helps explain this variability in the world and in people's estimates. By drawing on the ideal free distribution framework, we can identify how key properties of the

ecology (e.g., resource size, limits on the resource) as well as properties of the organisms'
cognitive systems (e.g., discriminatory abilities) shape the risk–reward structure in a
specific environment.

This theoretical approach to understanding choice environments thus offers the
opportunity to predict when the mind is expected to assume and exploit the risk–reward
relationship and when it is not, as well as when it would be maladaptive to do so. For
instance, the competitive ecology predicts the existence of a risk–reward relationship only
when an organisms–environment system has reached its ecological equilibrium. In
environments that have not reached equilibrium, ecologically rational actors facing
uncertainty should not rely on the risk–reward structure. Of course, a person may not
know that the system is newly forming and not yet in a state of equilibrium. Therefore, the
risk–reward structure may still anchor a person's cognition, for instance via a Bayesian
prior. The question of the extent to which people overgeneralize the risk–reward
relationship to unsuitable environments is an empirical one. Nevertheless, an ecological
theory can be used to map environments in which the risk–reward heuristic is better than
other strategies.

A normative theory about domain-specific cognitive tools

Being able to predict when the mind should and should not expect a risk–reward
structure can also inform normative theories of decision making. The overarching
assumption in normative decision theory is that one should maximize utility when making
decisions (Keeney & Raiffa, 1993; Savage, 1954; von Neumann & Morgenstern, 1947). This
assumption has several consequences. First, it is applicable across environments, resulting
in a domain-general normative theory. Second, it is largely divorced from the properties of
the agent. As Simon (1990) stated, "the economist who wishes to predict behavior studies
the environment in which the behavior takes place, for the rational economic actor will
behave in whatever way is appropriate to maximize utility in that environment" (p. 6). In

this case, all the agent has to do to act in a rational manner is to act in accordance with the rules or axioms that permit utility maximization. Human cognition, however, is limited and conformity to these axioms is not an easy feat (Gigerenzer & Selten, 2002; Simon, 1955; Weber & Johnson, 2009).

In contrast, an adaptive approach to cognition takes a different path and assumes that decision makers exploit both core psychological capacities and ecological structures to make decisions (Anderson, 1990; Brunswik, 1943; Simon, 1956; Fiedler, 2000; Kareev, 1995; Perkovic & Orquin, 2017; Stewart et al., 2006). Following this logic, the human mind is assumed to be composed of a toolbox of environment-specific decision strategies (Gigerenzer et al., 1999; Hertwig et al., in press; Payne et al., 1993). A great challenge to this view is the lack of a theory about the tool that should be selected in a given situation. Our ecological theory illustrates one way to address this thorny normative issue in the future. The competitive ecology theory of the risk–reward structure can be understood as an explicit, normative theory to evaluate where a particular tool—the risk–reward heuristic—should be applied or not. In addition, we can also record deviations from the heuristic and evaluate whether they were justified, appropriate, or adaptive, rather than dismissing them as "irrational." Taking such an ecological approach to normativity also means that the theory is not limited to some ideal environment, as is often the case with utility theory. Instead, it can suggest how ecologically rational actors should decide when conditions are less than ideal or when ecological anomalies happen like opportunities arising that are too good to be true (see Kaunitz et al., 2017; Leuker et al., 2018b).

Conclusion

Tversky and Kahneman (1974) famously argued that "people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors" (p. 1124). The

largely neglected issue in their work and the work it has inspired has been the question of when these heuristic principles are useful. Ecological theories like the one presented here will help to provide an answer to this question that goes beyond simply cataloging environments in which heuristics succeed or fail. In the present case, the toolkit of behavioral ecology has helped to reveal when a risk–reward structure can reliably be expected. This theoretical grounding predicts when the risk–reward heuristic (or other cognitive strategies that take advantage of this environmental structure) is appropriate and when it is not. In sum, this work shows that a system of bounded rationality can only be understood if a theory of the mind is coupled with a theory of the environment.

References

Abrahams, M. V. (1986). Patch choice under perceptual constraints: A cause for departures from an ideal free distribution. *Behavioral Ecology and Sociobiology*, *19*(6), 409–415. doi: 10.1007/BF00300543

Alatalo, R. V., Höglund, J., Lundberg, A., & Sutherland, W. J. (1992). Evolution of black grouse leks: Female preferences benefit males in larger leks. *Behavioral Ecology*, *3*(1), 53–59. doi: 10.1093/beheco/3.1.53

Anderson, J. R. (1990). *The adaptive character of thought.* New York, NY: Psychology Press. doi: 10.4324/9780203771730

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*(6), 396–408. doi: 10.1111/j.1467-9280.1991.tb00174.x

Barkow, J. H., Cosmides, L., & Tooby, J. (1995). *The adapted mind: Evolutionary psychology and the generation of culture.* New York, NY: Oxford University Press.

Bloom, A. J., Chapin III, F. S., & Mooney, H. A. (1985). Resource limitation in plants—an economic analogy. *Annual review of Ecology and Systematics*, *16*(1), 363–392. doi: 10.1146/annurev.es.16.110185.002051

Brunswik, E. (1943). Organismic achievement and environmental probability. *Psychological Review*, *50*(3), 255–272. doi: 10.1037/h0060889

Brunswik, E. (1944). Distal focussing of perception: Size constancy in a representative sample of situations. *Psychological Monographs*, *56*(254), 1–49. doi: 10.1037/h0093505

Canini, K. R., Griffiths, T. L., Vanpaemel, W., & Kalish, M. L. (2014). Revealing human inductive biases for category learning by simulating cultural transmission. *Psychonomic Bulletin & Review*, *21*(3), 785–793. doi: 10.3758/s13423-013-0556-3

Chater, N., Oaksford, M., Hahn, U., & Heit, E. (2010). Bayesian models of cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(6), 811–823. doi: 10.1002/wcs.79

Chittka, L., Thomson, J. D., & Waser, N. M. (1999). Flower constancy, insect psychology, and plant evolution. *Naturwissenschaften*, *86*(8), 361–377. doi: 10.1007/s001140050636

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97–118). New York, NY: Oxford University Press.

Danger, M., Daufresne, T., Lucas, F., Pissard, S., & Lacroix, G. (2008). Does Liebig's law of the minimum scale up from species to communities? *Oikos*, *117*(11), 1741–1751. doi: 10.1111/j.1600-0706.2008.16793.x

Davies, N. B., Krebs, J. R., & A., W. S. (2012). *An introduction to behavioral ecology. 2nd ed.* New York, NY: Wiley-Blackwell.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–11. doi: 10.1038/nn1560

Dreisig, H. (1995). Ideal free distributions of nectar foraging bumblebees. *Oikos*, *72*(2), 161–172. doi: 10.2307/3546218

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, *101*(3), 519–527. doi: 10.1037/0033-295X.101.3.519

Farrior, C. E., Tilman, D., Dybzinski, R., Reich, P. B., Levin, S. A., & Pacala, S. W. (2013). Resource limitation in a competitive context determines complex plant responses to experimental resource additions. *Ecology*, *94*(11), 2505–2517. doi: 10.1890/12-1548.1

Fiedler, K. (2000). Beware of samples! a cognitive-ecological sampling approach to judgment biases. *Psychological Review*, *107*(4), 659–676. doi: 10.1037//0033-295X.107.4.695

Fox, C. R., & Clemen, R. T. (2005). Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior. *Management Science*, *51*(9), 1417–1432. doi: 10.1287/mnsc.1050.0409

Fox, C. R., & Levav, J. (2004). Partition-edit-count: Naive extensional reasoning in judgment of conditional probability. *Journal of Experimental Psychololgy: General*, *133*(4), 626–42. doi: 10.1037/0096-3445.133.4.626

Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, *14*(3), 195–200. doi: 10.1111/1467-9280.02431

Fretwell, S. D., & Lucas, H. L. J. (1970). On territorial behavior and other factors influencing habitat distribuion in birds. I. Theoretical development. *Acta Biotheoretica*, *14*, 16–36. doi: 10.1007/BF01601953

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278. doi: 10.1126/science.aac6076

Gibson, J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Gigerenzer, G., & Selten, R. (2002). *Bounded rationality: The adaptive toolbox.* Cambridge, MA: MIT Press.

Gigerenzer, G., Todd, P. M., & ABC Research Group. (1999). *Simple heuristics that make us smart.* Oxford, United Kingdom: Oxford University Press.

Gillis, D. M., Peterman, R. M., & Tyler, A. V. (1993). Movement dynamics in a fishery: Application of the ideal free distribution to spatial allocation of effort. *Canadian Journal of Fisheries and Aquatic Sciences*, *50*(2), 323–333. doi: 10.1139/f93-038

Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, *11*(3), 508–514. doi: 10.3758/BF03196603

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, *12*(1), 81–87. doi: 10.3758/BF03196350

Gorban, A. N., Pokidysheva, L. I., Smirnova, E. V., & Tyukina, T. A. (2011). Law of the minimum paradoxes. *Bulletin of Mathematical Biology*, *73*(9), 2013–2044. doi: 10.1007/s11538-010-9597-1

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, *14*, 357–364. doi: 10.1016/j.tics.2010.05.004

Harpole, W. S., Ngai, J. T., Cleland, E. E., Seabloom, E. W., Borer, E. T., Bracken, M. E., . . . Smith, J. E. (2011). Nutrient co-limitation of primary producer communities. *Ecology Letters*, *14*(9), 852–862. doi: 10.1111/j.1461-0248.2011.01651.x

Haugen, T. O., Winfield, I. J., Vøllestad, L. A., Fletcher, J. M., James, J. B., & Stenseth, N. C. (2006). The ideal free pike: 50 years of fitness-maximizing dispersal in Windermere. *Proceedings of the Royal Society B: Biological Sciences*, *273*(1604), 2917–2924. doi: 10.1098/rspb.2006.3659

Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews*, *12*(3), 123–135. doi: 10.1002/evan.10110

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, *13*(2), 243–266. doi: 10.1901/jeab.1970.13-243

Herrnstein, R. J. (1974). Formal properties of the matching law. *Journal of the Experimental Analysis of Behavior*, *21*(I), 159–164. doi: 10.1901/jeab.1974.21-159

Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(5), 1191. doi: 10.1037/a0013025

Hertwig, R., Hoffrage, U., & ABC Research Group. (2013). *Simple heuristics in a social world.* New York, NY: Oxford University Press.

Hertwig, R., Pleskac, T. J., Pachur, T., & the Center for Adaptive Rationality. (in press). *Taming uncertainty.* Cambridge, MA: MIT Press.

Hoffart, J. C., Rieskamp, J., & Dutilh, G. (2018). How environmental regularities affect people's information search in probability judgments from experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition.* doi: 10.1037/xlm0000572

Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review*, *114*(3), 733–758. doi: 10.1037/0033-295X.114.3.733

Holzer, H. J., Katz, L. F., & Krueger, A. B. (1991, 8). Job queues and wages. *The Quarterly Journal of Economics*, *106*(3), 739–768. doi: 10.2307/2937925

Houston, A. I., & McNamara, J. M. (1988). The ideal free distribution when competitive abilities differ: An approach based on statistical mechanics. *Animal Behaviour*, *36*(1), 166–174. doi: 10.1016/S0003-3472(88)80260-4

Kareev, Y. (1995). Through a narrow window: Working memory capacity and the
    detection of covariation. *Cognition*, *56*(3), 263–269. doi: 10.1016/0010-0277(95)92814-G

Kaunitz, L., Zhong, S., & Kreiner, J. (2017). Beating the bookies with their own
    numbers-and how the online sports betting market is rigged. *arXiv preprint
    arXiv:1710.02824*.

Keeney, R. L., & Raiffa, H. (1993). *Decisions with multiple objectives: preferences and
    value trade-offs.* Cambridge University Press.

Kennedy, M., & Gray, R. D. (1993). Can ecological theory predict the distribution of
    foraging animals? A critical analysis of experiments on the ideal free distribution. *Oikos*,
    68(1), 158–166. doi: 10.2307/3545322

Keynes, J. M. (1921). *A treatise on probability.* London, United Kingdom: Macmillan and
    Co.

Knight, F. H. (1921). *Risk, uncertainty, and profit.* New York, NY: Sentry Press.

Krueger, A. B. (1988). The determinants of queues for federal jobs. *Industrial and Labor
    Relations Review*, *41*(4), 567–581. doi: 10.3386/w2499

Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*
    (2nd ed.). New York, NY: Academic Press/Elsevier.

Laplace, P. S. (1902). *A philosophical essay on probabilities.* Wiley and Sons. (Original
    work published 1814)

Leuker, C., Pachur, T., Hertwig, R., & Pleskac, T. J. (2018a). Exploiting risk–reward
    structures in decision making under uncertainty. *Cognition*, *175*, 186–200. doi:
    doi:10.1016/j.cognition.2018.02.019

Leuker, C., Pachur, T., Hertwig, R., & Pleskac, T. J. (2018b). Too good to be true?
    Psychological responses to uncommon options in risk–reward environments (Manuscript

submitted for publication). Retrieved from psyarxiv.com/z4knb. doi: 10.17605/OSF.IO/XAUTN

Luce, R. D., & Raiffa, H. (1957). *Games and decisions.* New York, NY: John Wiley & Sons, Inc.

Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological Review*, *118*(3), 393–437. doi: 10.1037/a0024143

Mesoudi, A. (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences.* Chicago, IL: University of Chicago Press.

Moussaïd, M., Brighton, H., & Gaissmaier, W. (2015). The amplification of risk in experimental diffusion chains. *Proceedings of the National Academy of Sciences*, *112*(18), 5631–5636. doi: 10.1073/pnas.1421883112

Nicolai, C. A., Sedinger, J. S., Ward, D. H., & Boyd, W. S. (2014). Spatial variation in life-history trade-offs results in an ideal free distribution in Black Brant Geese. *Ecology*, *95*(5), 1323–1331. doi: 10.1890/13-0860.1

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, *53*(3), 139-154. doi: 10.1016/j.jmp.2008.12.005

Parker, G. A. (1974). The reproductive behaviour and the nature of sexual selection in Scatophaga stercoraria L. (Diptera: Scatophagidae). IX. spatial distribution of fertilization rates and evolution of male search strategy within the reproductive area. *Evolution*, *28*(1), 93–108. doi: 10.1111/j.1558-5646.1974.tb00730.x

Parker, G. A., & Sutherland, W. J. (1986). Ideal free distributions when individuals differ in competitive ability: Phenotype-limited ideal free models. *Animal Behaviour*, *34*(4), 1222—1242. doi: 10.1016/S0003-3472(86)80182-8

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker.* New York, NY: Cambridge University Press.

Perkovic, S., & Orquin, J. L. (2017). Implicit statistical learning in real-world environments leads to ecologically rational decision making. *Psychological Science*, *29*(1), 34–44. doi: 10.1177/0956797617733831

Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, *143*(5), 2000–2019. doi: 10.1037/xge0000013

Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution.* Chicago, IL: University of Chicago Press.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207–236. doi: 10.1093/acprof:oso/9780199744282.003.0011

Rosenheim, J. A., Alon, U., & Shinar, G. (2010). Evolutionary balancing of fitness-limiting factors. *The American Naturalist*, *175*(6), 662–674. doi: 10.1086/652468

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, *6*, 41–49.

Savage, L. J. (1954). *The foundations of statistics.* New York, NY: John Wiley & Sons.

Silver, N. (2012). *The signal and the noise: why so many predictions fail–but some don't.* The Penguin Press.

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118. doi: 10.2307/1884852

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. doi: 10.1037/h0042769

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19. doi: 10.1146/annurev.psych.41.1.1

Şimşek, Ö., & Buckmann, M. (2015). Learning from small samples: An analysis of simple decision heuristics. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, & R. Garnett (Eds.), *The 29th annual conference on neural information processing systems* (pp. 3159–3167). Red Hook, NY: Curran Associates.

Skylark, W. J., & Prabhu-Naik, S. (2018). A new test of the risk-reward heuristic. *Judgment and Decision Making*, *13*(1), 73.

Smith, C., Reynolds, J. D., Sutherland, W. J., & Jurajda, P. (2000, 6). Adaptive host choice and avoidance of superparasitism in the spawning decisions of bitterling (Rhodeus sericeus). *Behavioral Ecology and Sociobiology*, *48*(1), 29–35. doi: 10.1007/s002650000212

Sokolowski, M. B., Tonneau, F., & Freixa i Baqué, E. (1999). The ideal free distribution in humans: An experimental test. *Psychonomic Bulletin & Review*, *6*(1), 157–161. doi: 10.3758/bf03210824

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. doi: 10.1016/J.Cogpsych.2005.10.003

Sutherland, W. J. (1983, 10). Aggregation and the 'ideal free' distribution. *The Journal of Animal Ecology*, *52*(3), 821–828. doi: 10.2307/4456

Sutherland, W. J., Townsend, C. R., & Patmore, J. M. (1988). A test of the ideal free distribution with unequal competitors. *Behavioral Ecology and Sociobiology*, *23*(1), 51–53. doi: 10.1007/BF00303058

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, *16*(3), 167–171. doi: 10.1111/j.1467-8721.2007.00497.x

Todd, P. M., Gigerenzer, G., & ABC Research Group. (2012). *Ecological rationality: Intelligence in the world.* New York, NY: Oxford University Press.

Tregenza, T. (1995). Building on the ideal free distribution. *Advances in Ecological Research*, *26*, 253–307. doi: 10.1016/S0065-2504(08)60067-7

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. doi: 10.1126/science.185.4157.1124

Ungemach, C., Stewart, N., & Reimers, S. (2011). How incidental values from the environment affect decisions about money, risk, and delay. *Psychological Science*, *22*(2), 253–260. doi: 10.1177/0956797610396225

Valladares, G., & Lawton, J. H. (1991). Host-plant selection in the holly leaf-miner: Does mother know best? *Journal of Animal Ecology*, *60*(1), 227–240. doi: 10.2307/5456

van der Lee, A., Gillis, D., & Comeau, P. (2013). Comparative analysis of the spatial distribution of fishing effort contrasting ecological isodars and discrete choice models. *Canadian Journal of Fisheries and Aquatic Sciences*, *71*(1), 141–150. doi: 10.1139/cjfas-2012-0511

von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior.* Princeton, NJ: Princeton University Press.

Waser, N. M. (1986). Flower constancy: Definition, cause, and measurement. *The American Naturalist*, *127*(5), 593–603. doi: 10.1086/284507

Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, *60*, 53-85. doi: 10.1146/Annurev.Psych.60.110707.163633

Appendix A

Risk–Reward and Heterogeneous Patches

Here we examine what happens to the predicted risk–reward relationship when we relax the assumption that resource sizes are constant within patches.

Two resource types

As a first step, we consider a relatively simple environment that has two resource sizes: small and large, with the constraint that $s_{large} > s_{small} = 1$. We further assume that across the landscape there are two types of patches: rich and poor. A rich patch contains more large resources than small ones, $m_{large} > m_{small}$. Similarly, a poor patch contains more small resources than large ones, $m_{small} > m_{large}$.

It is useful to further specify the relationship between resource sizes in the rich and poor patches. To do so, we set the number of large resources in the rich patch and the number of small resources in the poor patch to

$$m' = m_{rich,\ large} = m_{poor,\ small}. \tag{A1}$$

We also set the number of small resources in the rich patch and the number of large resources in the poor patch to

$$m_{rich,\ large} = m_{poor,\ small} = m'/\mu. \tag{A2}$$

The parameter $\mu \geq 1$ determines the similarity of the patches across the landscapes. As $\mu \to 1$ the rich and poor patches become more and more similar to each other, but also more heterogeneous in terms of resource size. That is, both rich and poor patches grow to contain approximately m' large and small resources and as a result the within-patch heterogeneity of resources approaches between-patch heterogeneity. Analogously, as μ increases, rich and poor patches become more dissimilar to each other such that as $\mu \to \infty$ rich patches will contain only large resources and poor patches only small resources (i.e.,

patches become homogeneous with respect to resource size). Hence, as μ increases, the between-patch heterogeneity increases and the within-patch heterogeneity decreases.

Recall that the amount of resources available in a patch is a multiplicative function of the number and size of the number of resources, $r = m \times s$ (Equation 4). Thus, the total amount of resources r available in rich patches is

$$r_{rich} = s_{large} \times m' + s_{small} \times \frac{m'}{\mu}. \tag{A3}$$

In poor patches, the total amount of resources is

$$r_{poor} = s_{large} \times \frac{m'}{\mu} + s_{small} \times m'. \tag{A4}$$

According to the ideal free distribution principle (Equation 3; and remembering that $s_{small} = 1$), the number of competitors in the rich patch is

$$n_{rich} \propto m'(s_{large} + \frac{1}{\mu}). \tag{A5}$$

The number of competitors in the poor patch is

$$n_{poor} \propto m'(\frac{s_{large}}{\mu} + 1). \tag{A6}$$

Generally, the risk–reward relationship is measured in terms of the size of the resource. To establish how the risk–reward relationship works in heterogeneous patches with different resource sizes, it is useful to find the average number of competitors per resource of a specific size. Assuming an equal number of rich and poor patches, the average number of competitors per large resource across the landscape of a total of Y patches is

$$\begin{aligned}
\bar{n}_{large} &= \frac{\sum_{y \in Y} m_{y,\ large} \times n_y}{\sum_{y \in Y} m_{y,\ large}} \\
&= \frac{m' \times c \times m'(s_{large} + \frac{1}{\mu}) + \frac{m'}{\mu} \times c \times m'(\frac{s_{large}}{\mu} + 1)}{m' + \frac{m'}{\mu}} \\
&= cm'(\frac{s_{large} + \mu^2 s_{large} + 2\mu}{\mu^2 + \mu})
\end{aligned} \tag{A7}$$

where c is a proportionality constant. Analogously, the average number of competitors per small resource is

$$n_{small} = cm'(\frac{1 + \mu^2 + 2s_{large}\mu}{\mu^2 + \mu}). \tag{A8}$$

Recall that the success probability p of a patch is inversely proportional to the total number of competitors per resource, $p_y \propto m_y/n_y$ (Equation 6). Thus, from Equation A7 the success probability for a large resource is

$$p_{large} = \frac{\mu^2 + \mu}{m'c(s_{large} + \mu^2 s_{large} + 2\mu)}, \tag{A9}$$

and for a small resource is

$$p_{small} = \frac{\mu^2 + \mu}{m'c(1 + \mu^2 + 2s_{large}\mu)}. \tag{A10}$$

With the number of competitors and the success probability established for rich and poor patches, we can use the parameter μ to establish how the risk–reward relationship changes as the rich and poor patches become more similar and, as a result, as the heterogeneity between patches matches the heterogeneity within patches. To do so, we will examine the ratio of the probability of success for small resources relative to that of large resources,

$$\frac{p_{small}}{p_{large}} = \frac{s_{large} + s_{large}\mu^2 + 2\mu}{1 + \mu^2 + 2s_{large}\mu}. \tag{A11}$$

When poor and rich patches are equivalent and there is maximal within-path heterogeneity ($\mu = 1$), then according to Equation A11 the ratio of success probabilities for small and large resources will be 1, $p_{small}/p_{large} = 1$. In words, when the number of large and small resources is equal in both rich and poor patches (maximal within-patch heterogeneity) then success probability is on average equivalent for small and large resources and there is no risk–reward relationship.

Now consider what happens when the rich and poor patches in the landscape become maximally dissimilar and there is minimal within-patch heterogeneity in resource size (i.e., rich patches only have large resources and poor patches only have small resources). This

occurs when $\mu \to \infty$. In this case, taking the limit of Equation A11 shows that

$lim_{\mu \to \infty} \frac{s_{large} + s_{large}\mu^2 + 2\mu}{1 + \mu^2 + 2s_{large}\mu} = lim_{\mu \to \infty} \frac{2\mu s_{large} + 2}{2\mu + 2s_{large}} = s_{large}$. That is, as rich and poor patches

become maximially dissimilar the ratio of success probabilities for small to large resources

(p_{small}/p_{large}) approaches s_{large} and as a result the success probability grows to be inversely

related to resource size (a risk–reward relationship).

To address what happens for intermediate levels of μ we need to establish how

Equation A11 changes as a function μ. To do so, we take the partial derivative of the

function with respect too μ. Doing so shows that

$$d(\frac{s_{large} + s_{large}\mu^2 + 2\mu}{1 + \mu^2 + 2s_{large}\mu})/d\mu \propto (s_{large}^2 - 1)(\mu^2 - 1) > 0. \tag{A12}$$

Therefore, for intermediate levels of μ when rich and poor patches are transitioning from

maximally similar to maximally dissimilar, the ratio of success probabilties for small and

large resources lies between 1 and s_{large}. That is, the success probability is smaller for large

resources than for small resources, but the relationship is shallower than predicted when

patches are maximially dissimilar (i.e., $p_{small}/p_{large} < s_{large}$). In particular, as rich and

poor patches become more dissimilar (i.e., as μ increases), the relationship between the

success probability and resource size approaches an inverse-proportional relationship and

resource size becomes a more reliable predictor of the success probability.

Multiple resource types

In order to generalize to more complex situations with larger numbers of different

reward sizes, we used numerical simulations. Please see the accompanying website at the

Open Science Framework for the code. In the simulation, we assumed 10 resource sizes

(ranging from $s = 1, 2, ..., 10$). There were $\sum m_s = 1,000$ total resources with $m_s = 100$

resources of each size class s. These resources were distributed among $Y = 10$ patches

following one of four procedures:

1. Each patch received 100 resources of one size class; different patches received
   resources of different sizes ("minimal within-patch heterogeneity");

2. Patches were divided into three different classes, each receiving 100 resources as follows:

- Three "very poor" patches received resources drawn randomly from sizes 1 to 3;

- Three "intermediate" patches received resources drawn randomly from sizes 4 to 6;

- Four "very rich" patches received resources drawn randomly from sizes 7 to 10 ("low within-patch heterogeneity");

3. Patches were divided into two different classes, each receiving 100 resources as follows:

- Five "poor" patches received resources drawn randomly from sizes 1 to 5;

- Five "rich" patches received resources drawn randomly from sizes 6 to 10 ("high within-patch heterogeneity");

4. All patches received 100 resources drawn randomly from sizes 1 to 10 ("maximal within-patch heterogeneity").

For each patch, we calculated the total amount of resources using Equation 4. Next, $\sum_{y \in Y} n_y = 1,000$ competitors were distributed among the patches according to the ideal free distribution principle so that the number of competitors was inversely proportional to the total amount of resources in a patch. The relative probability of success for each resource size was calculated by determining the inverse of the average number of competitors per resource of this size class across all patches,

$$p_s = \frac{\sum_{y \in Y} m_{s,y}}{\sum_{y \in Y} n_y \cdot m_{s,y}}. \tag{A13}$$

Figure 2 plots the average probability of success across 100 iterations of each landscape. They confirm that the analytical results for two resource sizes generalizes to multiple resource sizes: As within-patch heterogeneity increases, the relationship between

probability of success and resource size becomes increasingly shallower. Finally, when within-patch heterogeneity is maximal, probability of success is the same for all resource sizes (Figure 2d).

Appendix B

Landscape Saturation

Assume a landscape in which the number of patches is large, but the number of competitors is small and resource density is low, so that resources are spread out thinly. Several patches will contain no competitors, either because these patches contain such small amounts of resources that they attract no competitors, or because there are simply not enough competitors for each patch. Do we still expect a negative relationship between reward size and probability of success in such a non-saturated landscape? We address this question with simulations. Please see the accompanying website at the Open Science Framework for the code.

Simulation procedure

We assumed a landscape with $Y = 200$ patches, and resources of 10 different sizes ranging from $s = 1, 2, ...10$. The exact number of patches is of little relevance for the present purpose, as long as it is large relative to the number of competitors and/or number of resources. For simplicity, the number of resources of each size was presumed to be equal ($m_s = 20$). We further assumed that individual patches were homogeneous with respect to resource size, and that there were equal number of patches with resources of each size class so that 20 patches contained resources of size $s = 1$, another 20 patches contained resources of size $s = 2$, and so on.

In each simulation, the resources of each size were randomly distributed among the 20 patches with that resource size according to a uniform distribution. Then competitors were distributed one by one in the landscape according to the maximal per capita consumption rate at the time. This ensured that competitors were distributed in an "ideal–free" manner. More specifically, when a new competitor entered the landscape the amount of resources r_y was calculated for each of the 200 patches (Equation 4). Then the per capita consumption rate u_y for each patch which this competitor would encounter at

that time was calculated such that

$$u_y = \frac{r_y}{n_y + 1}.$$ (B1)

The new competitor then joined the patch with the highest per capita consumption rate at that time. If multiple patches had the same maximum uptake then the competitor was randomly placed in one of those patches. This procedure continued until all competitors were distributed.

At the end of each simulation, the success probability for a newly arriving focal competitor was calculated using the assumption that the success probability in a given patch is proportional to the ratio of the number of resources to the number of competitors in that patch (Equation 6). In particular, the success probability for patch y was calculated as

$$p_y = \frac{\frac{m_y}{n_y+1}}{\sum_{y \in Y} \frac{m_y}{n_y+1}}.$$ (B2)

Note that to make the estimated success probabilities comparable across different landscapes, we normalized the ratio of the number of resources to the number of competitors to the value of the sum of this ratio across the patches.

Since patches were homogeneous in terms of resource size, the success probability in a given patch was also the probability of success for the relevant resource size, p_s. Thus, in each simulation we obtained 20 estimates of the success probability for any given resource size.

We investigated low and high levels of competitor saturation of the landscape by distributing either $n = 100$ total competitors across the landscape or $n = 1,000$ competitors. These different population sizes of competitors either entered a landscape with low resource density with $m_s = 100$ resources of each of 10 resource sizes (i.e., 1,000 resources in total), or the competitors entered a landscape with high resource density with $m_s 1,000$ resources of each of 10 resource sizes (i.e., 10,000 resources in total). For each of

the resulting four parameter combinations, we ran 1,000 simulations.

Figure 3 summarizes the results of the simulations. We found that the median (and mean) success probability still decreases with resource size, but the more shallow the decrease, the less saturated the landscape is with respect to the number of competitors. Further, the probability of success is more variable the lower the saturation of the landscape with either competitors or resources, particularly for small resource sizes.

Appendix C

Unequal Competitors

The ideal free distribution principle assumes not only that competitors are ideal, but also that they are equally competitive in the landscape. That is, they have an equal ability to obtain a resource. To examine the consequences of this assumption we relax it, allowing for two classes of competitors: bullies and wimps. Bullies are twice as good at acquiring resources as are wimps. Thus, the presence of a bully impacts a wimp's resource consumption twice as strongly as the presence of another wimp does, while any wimp impacts on a bully's intake only half as strongly as another bully does.

With these constraints, based on an ideal free distribution, we set the per capita resource consumption rate u in a given patch y for bullies as

$$u_{y,\text{bullies}} = \frac{r_y}{n_{y,\text{bullies}} + \frac{1}{2}n_{y,\text{wimps}}}. \tag{C1}$$

The per capita resource consumption rate u in a given patch y for wimps is

$$u_{y,\text{wimps}} = \frac{r_y}{2n_{y,\text{bullies}} + n_{y,\text{wimps}}}. \tag{C2}$$

It follows from Equations C1 and C2 that the per capita consumption rate for wimps is half that for bullies in each patch with both classes of competitors. An ideal free distribution of unequal competitors across patches is possible: If both bullies and wimps move between patches so as to maximize their own resource consumption, then the per capita resource consumption rate for bullies at equilibrium is the same in all patches, and that for wimps is also the same in all patches. However, the constant per capita resource consumption rate for wimps is only half as large as that for bullies,

$$u_{\text{bullies}} = 2u_{\text{wimps}} = \kappa. \tag{C3}$$

It follows from Equations C1 and C3 that the distribution of competitors in each patch is

$$n_{y,\text{bullies}} + \frac{1}{2} \cdot n_{y,\text{wimps}} = \frac{r_y}{\kappa} \tag{C4}$$

Recalling the amount of resources is determined by the number and size of resources,

$r_y = m_y \times s_y$ (Equation 4), this implies

$$n_{y,\text{bullies}} + \frac{1}{2} \cdot n_{y,\text{wimps}} = \frac{r_y}{\kappa} = \frac{m_y \times s_y}{\kappa}. \tag{C5}$$

We are interested in a bully's (or wimp's) success probability within a patch. Based

on the resource to competitor assumption (Equation 6), the success probability for bullies is

$$p_{y,\text{bullies}} \propto \frac{m_y}{n_{y,\text{bullies}} + \frac{1}{2} \cdot n_{y,\text{wimps}}}. \tag{C6}$$

Recall that each wimp only impacts half as much as another bully on a bully's resource

uptake, and, thus, on a bully's probability of successfully gaining a resource within a given

time period. Similarly, the success probability for wimps is

$$p_{y,\text{wimps}} \propto \frac{m_y}{2 \cdot n_{y,\text{bullies}} + n_{y,\text{wimps}}} \tag{C7}$$

Substituting Equation C4 into Equations C6 and C7, we can determine the expected

relationship between success probability and resource size at equilibrium. For bullies, this

results in

$$p_{y,\text{bullies}} \propto \frac{m_y}{\frac{m_y \times s_y}{\kappa}}. \tag{C8}$$

Because κ is a constant this expression simplifies to

$$p_{y,\text{bullies}} \propto \frac{1}{s_y}. \tag{C9}$$

For wimps, their reduced competitiveness results in the success probability being

proportional to

$$p_{y,\text{wimps}} \propto \frac{m_y}{\frac{m_y \times s_y}{\frac{1}{2}\kappa}}. \tag{C10}$$

This expression simplifies to

$$p_{y,\text{bullies}} \propto \frac{1}{2s_y}. \tag{C11}$$

That is, for bullies and for wimps the success probabilities are different: In our example the

success probability of bullies is always twice as high as that of wimps for all resource sizes.

This reflects the fact that, as competitors, bullies are twice as strong as wimps. However, within each class of competitors, the success probability is still inversely proportional to resource size. Hence, just as for equal competitors, the success probability for an individual competitor of a given class can be predicted by resource size. The conclusions are similar when one considers more than two different classes of competitors.

Appendix D

Interference Competition

If competitors interfere with each other's ability to acquire resources (e.g., because they frighten each other's prey off, or engage in antagonistic interactions), the impact of competitors on the resource consumption rate is more than just proportional (Parker & Sutherland, 1986),

$$u \propto \frac{r_y}{n_y^\theta}, \tag{D1}$$

where $\theta > 1$ and captures the degree of interference. Since the per capita resource consumption rate at equilibrium is the same in all patches, the ideal free distribution principle (Equation 3) changes to

$$n_y \propto r_y^{\frac{1}{\theta}}. \tag{D2}$$

Because $r_y = m_y \times s_y$ (Equation 4), it follows that

$$n_y \propto (m_y \times s_y)^{\frac{1}{\theta}}. \tag{D3}$$

The success probability is also impacted by interference such that it is no longer inversely proportional simply to the number of competitors (Equation 6). Instead we assume that the interference impacts the success probability to the same degree it impacts the per capita consumption rate so that

$$p_y \propto \frac{m_y}{n_y^\theta}. \tag{D4}$$

Substituting Equation D2 into Equation D4 shows that

$$p_y \propto \frac{m_y}{n_y^\theta} = \frac{m_y}{[(m_y \times s_y)^{\frac{1}{\theta}}]^\theta}. \tag{D5}$$

Thus, with interference among competitors the success probability is still inversely proportional to the size of the resources,

$$p_y \propto \frac{1}{s_y}. \tag{D6}$$

That is, while interference of competitors changes the distribution of competitors (see Equation D2), it has no effect on risk–reward relationship.

Supplementary Material for
"The ecology of competition: A theory of risk–reward
environments in adaptive decision making"

Timothy J. Pleskac[1,2], Larissa Conradt[2], Christina Leuker[2] and Ralph Hertwig[2]

[1]University of Kansas
[2]Max Planck Institute for Human Development, Berlin, Germany

December 6, 2018

Contents

1

1 Risk–Reward Estimation Study

Our goal was to replicate the risk–reward estimation studies in Pleskac & Hertwig (2014) with a German sample. We used the scenario (in German) described in Study 2. Moreover, we manipulated the order in which participants were asked to estimate how many balls they thought were in the basket and if they wanted to pay to play the gamble. Half of the participants completed the estimation question before the choice question, while the other half completed it in the other order. We also extended the range of payoffs (payoff levels, between participants: €2.5, €4, €10, €50, €100, €200, €500, €1,000, and €10,000).

1.1 Methods

1.1.1 Participants

A total of 455 participants were recruited using a web panel hosted at the Max Planck Institute for Human Development in Berlin, resulting in approximately 50 participants per payoff × order condition (between-participants). Each participant earned €2.0 for participation. The study was approved by the IRB at the Max Planck Institute for Human Development in Berlin.

1.2 Results

1.2.1 Estimation

	Coefficient	HDI (lower)	HDI (higher)
Intercept (Mean)	-1.45	-1.60	-1.29
Reward (Mean)	-1.48	-2.12	-0.79
Order:Choice first (Mean)	-0.19	-0.41	0.03
Order:Choice first × Reward (Mean)	1.43	0.56	2.28
Intercept (Dispersion)	0.34	0.14	0.53
Reward (Dispersion)	1.35	0.50	2.10
Order:Choice first (Dispersion)	0.48	0.20	0.76
Order:Choice first × Reward (Dispersion)	-2.10	-3.11	-1.05

Table S1: Bayesian beta regression table using normalized rewards (0–1), estimates (0–1), and question order (baseline: estimates, then choice) as predictors for both the mean and the dispersion of the estimates.

We modeled the effect of payoff on the mean and dispersion (variance) in the estimates using a beta regression with Bayesian estimation techniques (Kruschke, 2014), using Stan in R for regression analyses with the rstanarm package (*RStanArm Version 2.9.0-4*, 2016). In the regression, we used the reward value (normalized to be between 0 and 1), question order (estimate first vs. choice first), and the interaction between reward value and order, as predictors via a logit link of the estimates (normalized to be between 0 and 1). In addition, we used the same variables to simultaneously predict the dispersion (a proxy for the variance) of the estimates (Smithson & Verkuilen, 2006) via a log link.

Results of the beta regression are summarized in Table S1. Replicating past results, and consistent with participants using the risk–reward heuristic, as rewards increased the probability estimates decreased. There was also a credible order effect on the estimates as well as an interaction between

order and reward level: Participants who chose whether or not to play first (and then estimated the chances of winning) generally showed a diminished effect of reward on their estimates. This also meant that participants who chose first appeared to estimate their chances of winning average, higher.

In terms of the dispersion, there was a credible decrease in the dispersion of the estimates as payoff increased (due to a negative log-link function). As with the mean estimates, there was a credible order effect on the dispersion of the estimates as well as an interaction between order and reward level. For participants who chose first (and then estimated the chances of winning), there was more dispersion in the estimates. Moreover, the effect of reward value changed from showing a credible decrease when estimating the probability first ($b = 1.38$, CI $= [0.52, 2.13]$) to a credible increase when making a choice first ($b = -0.76$, CI $= [-1.41, -0, 14]$; dispersion coefficient from a beta regression for estimate first condition).

It is difficult to draw strong conclusions about the order effect as the study was not designed to examine this effect. Nevertheless, the risk–reward estimation model (described in detail in section 2) can help understand these results. The risk–reward estimation model suggests that the order effects are due to an increase in the rate of participants using the principle of indifference (i.e. estimating a probability of .5) when choosing whether to play first compared to when estimating before choosing: Analyzing the data from Study 3 alone showed that the rate of the principle of indifference was $M = .28$ (HDI $= [.12, .41]$) for participants who chose before estimating, and $M = .24$ (HDI $= [.07, .37]$) for participants who estimated before choosing (difference: $M_{choose>estimate} = .04$, HDI $= [-.03, .12]$).

1.2.2 Choice

	Coefficient	HDI (lower)	HDI (higher)
Intercept	-0.52	-0.89	-0.16
Reward	1.25	0.35	2.19
Estimate	1.32	0.15	2.51
Order:Choice first	0.69	0.18	1.20
Order:Choice first $\times$ Reward	0.62	-0.90	2.29
Order:Choice first $\times$ Estimate	-1.15	-2.88	0.58

Table S2: Logistic regression table using normalized rewards (0–1), estimates (0–1), and question order (baseline: estimates, then choice) as predictors for both the mean and the variance estimates.

We also examined the choices participants made in terms of whether or not to pay to play. We did so using a logistic regression with Bayesian estimation techniques (Kruschke, 2014). The logistic regression was run using Stan in R for regression analyses with the rstanarm package (*RStanArm Version 2.9.0-4*, 2016). In the regression, we entered the normalized reward value, the normalized estimates, whether participants entered an estimate first (1) or choice first (0), and the interaction between order and the reward value and the interaction between the order and the estimate as predictors of the choice. The results are summarized in Table S2. They largely replicate the studies reported in Pleskac & Hertwig (2014). On average, participants were more likely to choose to pay €2.0 to play the gamble as the reward value increased. Consistent with participants using the risk–reward heuristic to help decide whether to pay to play, their estimates were also credible predictors of their choice. Finally, note there was a credible order effect such that participants that chose first were more likely to pay to play.

2 Risk–reward estimation model

2.1 Model estimation

We estimated the posterior distributions over the parameters of the hierarchical models using Markov Chain Monte Carlo (MCMC) methods. Our model estimation procedure was implemented with JAGS 4.3 using Matlab via matjags to interface with JAGS. The general JAGS code used is in the OSF. The estimation of the model used three parallel chains. Each chain consisted of 1,000 burn-in steps (unrecorded samples to allow the chain to reach the reasonable parameter space) and 20,000 samples for a total of 60,000 samples.

The chains were evaluated for representativeness and accuracy using the procedures outlined by Kruschke (2014). Representativeness was evaluated using visual inspection of trace plots of the chains and density plots. All the chains at the group level were inspected visually with random samples of chains from the individual level. Representativeness was also evaluated numerically using the Gelman-Rubin statistic with the conventional heuristic that values of the Gelman-Rubin statistic above 1.1 were worrisome. All the chains met these standards suggesting representativeness of the posterior distributions. Accuracy was evaluated by examining the autocorrelation and the effective sample size. The effective sample size estimates the sample size of the chain after accounting for the autocorrelation present in the samples. As a rough standard we sought to have approximately an effective sample size of approximately 10,000. Our main focus in this paper was on comparing the group level mean estimates of the parameters.

2.2 Model comparisons

As a means for testing for the importance of different components of the RREM, we also fit several variants. We compared the fits of the different models with the Deviance Information Criterion (DIC; Spiegelhalter et al., 2002) where smaller values are better. Table S3 lists the DICs for the different variants of the risk–reward estimation model where different models were included (1) or not (0). The comparisons show that the best model has all three components, which is the model we report in the paper. However, DICs do show that the distortion and ignorance prior components play a larger role than the change in dispersion.

It is also informative that a beta regression on the estimates where the mean and variance are allowed to vary as a function of the payoff level has a DIC of -1,1139.2 That is, the Risk–Reward Estimation Model provides a better fit to the data than a statistical model, demonstrating the importance of accounting for the cognitive strategies used in making estimates.

Distortion	Dispersion	Ignorance Prior	DIC
1	1	1	-1,249.4
1	1	0	-1,247.4
1	0	1	-1,247.3
0	1	1	-1,241.6
1	0	0	-1,247.6
0	1	0	925.1
0	0	1	-1,228.9
0	0	0	936.7

Table S3: DICs for different variants of the Risk–Reward Estimation Model.

References

Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). New York, NY: Academic Press/Elsevier.

Pleskac, T. J., & Hertwig, R. (2014). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*, *143*(5), 2000–2019. doi: 10.1037/xge0000013

RStanArm version 2.9.0-4. (2016). `cran.r-project.org/web/packages/rstanarm/`.

Smithson, M., & Verkuilen, J. (2006). A better lemon squeezer? maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods*, *11*(1), 54–71. doi: 10.1037/1082-989X.11.1.54

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583–639.