# Taming Uncertainty

Ralph Hertwig, Timothy J. Pleskac, Thorsten Pachur,
and The Center for Adaptive Rationality

The MIT Press
Cambridge, Massachusetts
London, England

# 12 Rivals in the Dark: Trading Off Strategic and Environmental Uncertainty

**Doug Markant and Ralph Hertwig**

## 12.1 From Environmental to Strategic Uncertainty

Choosing a home is one of the most important decisions an animal can face—it warrants careful inspection of any potential habitat. At the same time, the desire to make a well-informed choice is often eclipsed by the need to act before competitors. Unsurprisingly, many species have evolved habitat selection strategies that adapt to competitive pressure. Take hermit crabs, which must repeatedly find new shells as their bodies grow. Upon discovering an empty shell, a lone hermit crab may diligently inspect it, decide it is a good fit, vacate its previous shell, and inhabit the new one. But empty shells of the right size are not always easy to come by. In fact, the surest way to secure the right shell before competitors is simply to clamber in the moment it is discarded by another crab. So groups of house-hunting hermit crabs sometimes self-organize into queues ordered by size, with each crab perched on the back of a slightly larger (but still occupied) shell. When the largest crab at the head of the queue moves into a new shell, it sets off a chain reaction: each crab quickly occupies the shell vacated by the next-largest neighbor before any competitors can muscle their way in (Rotjan, Chabot, & Lewis, 2010).

Honey bees provide another example of adapting the search for a new home when facing competition. In contrast to the unilateral selection made by each hermit crab, honey bees collectively decide where to build a new hive by relying on scouts that explore and "vote" on potential nest sites (Seeley, Visscher, & Passino, 2006), a process that would seem to prioritize consensus-building over speed. Yet honey bees are also sensitive to the cost of losing a high-quality location to competitors. If an attractive site is under inspection by members of a rival colony, scouts will fight to prevent it from being explored until their own colony has reached a decision (Rangel, Griffin, & Seeley, 2010).

Most humans no longer resort to physical conflict over real estate, but they are nonetheless keenly attuned to the need to act before competitors. The mere suggestion that others are considering the same house can push people to make an offer more quickly than they would otherwise—a fact that is well known not just to real estate agents but to all salespeople. Why is this sales technique so effective? It is not simply due to people being impulsive or suggestible. Assuming the threat of competition is credible, acting quickly to claim a resource—even one of uncertain value—can be a rational response. Decision strategies that might be rewarding in a solitary world, such as a meticulous search for the perfect home, are less effective in a social world where there is a risk that others will swoop in to claim attractive options for themselves.

Indeed, an inherent challenge of competition is that one rarely knows for certain how competitors will behave, as they are typically loath to reveal their intentions before they act. People thus commonly face two sources of uncertainty when making decisions under competition: *environmental uncertainty* about the value of choice options and *strategic uncertainty* about how others will act (Brandenburger, 1996). Environmental uncertainty is present when the payoff from choosing an option (e.g., a potential home) is not fully known to the decision maker (see chapters 2, 3, 7, and 8). Strategic uncertainty arises in any interactive setting with other agents who might adopt one of many decision-making strategies (see chapters 5 and 14). Nonsocial decisions, which can be conceptualized as "games against nature" (Hertwig, Hoffrage, & the ABC Research Group, 2013), can involve environmental uncertainty but are typically free from strategic uncertainty. The risks involved in playing roulette depend only on the structure of the roulette wheel, which does not change to anticipate or outwit gamblers. In social settings, however, other people's behavior has consequences for one's own decisions, and uncertainty about how they will act makes it difficult to identify the best strategy for oneself.

### 12.1.1 The Need to Know vs. the Need for Speed

Searching for a home is just one example of a common conflict between environmental and strategic uncertainty: people must balance the benefits of exploration—an important means for resolving environmental uncertainty—against the potential costs of being preempted by rivals. This conflict exists in many competitive situations, from the momentous to the mundane, but is perhaps most salient when uncertainty is high. Consider the relatively routine task of booking a hotel room in an unfamiliar tourist destination (an example not unlike that of house-hunting, but with considerably lower stakes). How much time should be spent researching different hotels given that other tourists are exploring the same options? How long are those competitors likely to bide their time before claiming the more attractive options for themselves? The same questions could be asked of any resource that is in demand but of uncertain value. In this chapter, we examine how individuals navigate the fundamental trade-off between exploration and the drive to claim resources before competitors do so.

One motivation of this research is to understand whether this trade-off is distinct from other *exploration–exploitation trade-offs* seen in solitary contexts (see chapter 7). Is the cost of losing out to a competitor interchangeable with other costs, such as opportunity costs or costs of acquiring information? Is strategic uncertainty substantively different from environmental uncertainty, or is it just a label for a class of environments that involve social agents? We think there are at least two reasons to expect that strategic uncertainty is more than a label, both of which we explore further here. First, people can mitigate strategic uncertainty to some extent by drawing on social knowledge to predict others' behavior. Even when competing with strangers, knowledge about social norms and past experiences in similar social situations can provide good guesses as to what other people will do. Second, competing with other agents may invoke strategies that are matched to the structure of social interactions and are robust to uncertainty about others' goals or strategies (Hertwig et al., 2013). For example, reasoning about the intentions of others can involve complex chains of perspective-taking ("I think that they think that I think that they will..."), but this process may quickly outstrip the capacity or time available to a human decision maker. Adapting to strategic uncertainty may depend on simpler strategies which exploit regularities that are unique to social, competitive environments (see also chapter 5).

In the following, we use two threads of research to draw out the tension between environmental and strategic uncertainty. We begin with an example from behavioral game theory of how people respond to strategic uncertainty when the costs of being preempted by competitors are fully transparent (i.e., in the absence of environmental uncertainty). We then step beyond the standard confines of game theory by considering the intersection of strategic and environmental uncertainty, where the imperative to act before competitors is in conflict with the drive to learn about choice options. Finally, we argue for an ecological understanding of how people adapt their decision strategies in a heterogenous, competitive world in which the "rules of the game" may themselves be uncertain.

## 12.2 Strategic Uncertainty: How Will Competitors Act?

How *should* people act when faced with strategic uncertainty? The most prominent theoretical tool for answering this question is *game theory* (see also chapter 5). In game theory, individuals are conceptualized as "players" in a game that reflects the structure of a real-world interaction and which affords multiple courses of action, referred to as "strategies." Game theory aims to explain how the value of a strategy depends on the strategies adopted by other players. As a normative framework founded on the principles of expected utility theory (Luce & Raiffa, 1957; von Neumann & Morgenstern, 1944/2007), it builds on the assumptions that first, all players act to maximize their own utility, and second, they believe their opponents will do the same. A common goal of game theoretic analysis is to identify *equilibria* for a given interaction. An equilibrium occurs when each player selects a best strategy (i.e., the course of action that maximizes their expected utility) given uncertainty about their opponents' strategies. When a set of competitors are at an equilibrium, no individual has an incentive to unilaterally change their strategy. Equilibria are useful normative guideposts for how people (or more generally, companies, countries, or any other kind of agent) should act in strategic interactions, but in practice they often provide a poor description of behavior. Throughout the history of game theory, behavioral researchers have sought to understand why the gap between predicted equilibria and actual behavior exists and what it reveals about the way that people respond to strategic uncertainty (Camerer, 2003; Gächter, 2004).

### 12.2.1 The Centipede Game: Competing for Known Rewards

Consider a classic problem involving competition for limited resources referred to as the *centipede game* (see figure 12.1; McKelvey & Palfrey, 1992; Rapoport, 2003; Rosenthal, 1981). Imagine that two players are seated

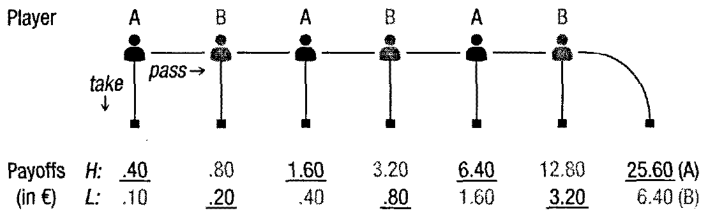


**Figure 12.1**
Depiction of the 6-move centipede game from McKelvey and Palfrey (1992). Player A and player B take turns deciding whether to take the higher value option H (vertical lines) or pass to the other player (horizontal lines). Each pass causes both payoffs to double. The payoffs earned by player A following each take decision are underlined. If neither player decides to take at any point, the game terminates with player A receiving the H option (€25.60) and player B receiving the L option (€6.40).

across a table, with two piles of money lying between them (a high-value pile *H* and a low-value pile *L*). The players are informed that the amount of money in each pile will double with each passing round of the game, up to a known maximum number of rounds. For example, the game might start with pile *H* containing €0.40 and pile *L* containing €0.10. After six turns the amounts in each pile will have increased to $H$=€12.80 and $L$=€3.20. The gameplay is simple: players take turns deciding whether to take the *H* pile, thereby ending the game and forcing the other player to receive the *L* pile, or to pass and give the other player the chance to decide.

The centipede game is free of environmental uncertainty about the choice options. Players know the precise payoffs they will receive should they choose to stop and take an option, and they know how those payoffs will change over future rounds. It also provides perfect information about the game state: players observe the choices made by opponents and know that their opponents have the same knowledge of the game structure. This simple, transparent game nevertheless generates a powerful conflict of interests. On the one hand, the increase in payoffs across rounds incentivizes players to allow the game to continue as long as possible, as both piles will double in value after every round. On the other hand, the disparity between the *H* and *L* options produces a strong preemption motive, in that there is a clear cost to *not* being the first person to stop and take the *H* option. As figure 12.1 shows, whenever a player's opponent decides to take *H*, the player's own payoff is half of what it would have been if they had stopped on the previous round.

This conflict is fueled by strategic uncertainty. If an opponent's strategy was known in advance (e.g., that the opponent plans to stop on the third round), the best response would be to stop one round earlier, thereby securing the largest possible payoff. What should a player do, however, if they do not know how long an opponent will wait before claiming the *H* option? Imagine that the game progresses to player B's final choice of whether to *take* or *pass* (see figure 12.1). The payoff for passing (€6.40) is half that of taking (€12.80), so a payoff-maximizing player would be expected to take. Now consider player A's immediately preceding choice. Predicting that player B will take at their next opportunity, player A's payoff from passing (€3.20) is less than that of taking (€6.40), similarly leading to the decision to take. The same logic applies backwards up to the first round of the game; any decision to pass leads to a worse payoff under the assumption that the opponent will take at their next opportunity. This reasoning process, known as *backwards induction*, creates a "race to the bottom" that converges on the first round, such that the equilibrium strategy in this game is to pounce on the *H* pile at the very first opportunity (Rapoport, 2003; Rosenthal, 1981).

Several behavioral studies on the centipede game have shown that people act quite differently from this equilibrium strategy (Bornstein, Kugler, & Ziegelmeyer, 2004; Fey, McKelvey, & Palfrey, 1996; Krockow, Pulford, & Colman, 2015; McKelvey & Palfrey, 1992; Nagel & Tang, 1998; Parco, Rapoport, & Stein, 2002; Rapoport, Stein, Parco, & Nicholas, 2003). People neither allow the game to continue for the maximum number of rounds, which might be expected if they only paid attention to the overall magnitude of the payoffs, nor do they typically stop on the first round. The most common stopping points fall somewhere in between; the modal choice is to claim the *H* option when approximately half the number of possible rounds have passed (Krockow et al., 2015; McKelvey & Palfrey, 1992; Nagel & Tang, 1998). Subtle changes in the payoff structure of the game can lead to earlier stopping points, including when larger incentives are at stake (Parco et al., 2002; Rapoport et al., 2003), when the magnitude of the payoffs does not increase across rounds (i.e., constant-sum centipede games; Fey et al., 1996), or when both players receive nothing at the final node (Krockow et al., 2015; Rapoport et al., 2003). On the whole, however, the equilibrium strategy of taking the *H* option at the first opportunity is relatively uncommon in experimental studies of the centipede game.

What explains this stark difference between the game-theoretic prediction and human behavior? In the following, we describe behavioral studies which suggest that people make choices based on their beliefs about how opponents will act—beliefs which may depart in significant ways from the standard assumptions of game theory. In short, strategic uncertainty may be moderated by prior knowledge about other people's competence or preferences, and it may be reduced as a result of learning through experience.

**12.2.1.1 Prior beliefs about competitors** People may not view opponents as the self-interested, rational actors presumed by game theory. For example, players may be more willing to pass in the centipede game if they believe that competitors are somewhat altruistic (i.e., care about the payoffs received by the other player; McKelvey & Palfrey, 1992) or are likely to make errors (Fey et al., 1996). Such beliefs imply that nonequilibrium strategies are more likely among opponents, leading to a different best response. According to this account, players are more willing to bide their time before claiming the *H* option because they perceive less competitive pressure than assumed by the game-theoretic equilibrium.

In an experimental test of this idea, Palacios-Huerta and Volij (2009) argued that players were more likely to adopt the equilibrium strategy when they had reason to believe their opponent would. The authors examined the behavior of both college students and expert chess players in the centipede game. In games between two chess players, the most common strategy was to stop on the first round in accordance with the equilibrium strategy (and in games between the most skilled grandmaster players, the equilibrium strategy was used 100% of the time). In games between students, behavior was similar to that observed in previous studies, with the majority of pairs reaching half the number of possible rounds. However, when students were informed that they were playing against chess players, they were more likely to stop on early rounds. Thus, given a hint that opponents may be strategically adept, naive players appeared to adjust their own strategy to preempt them.

**12.2.1.2 Learning through repeated interaction** A number of studies suggest that this kind of adaptation can also emerge as a result of experience in the centipede game (Fey et al., 1996; Rapoport et al., 2003), consistent with evidence from other strategic games that learning supports convergence toward equilibria (Fudenberg & Levine, 2016; Roth & Erev, 1995). Although Palacios-Huerta and Volij (2009) focused on players' knowledge of their opponents' abilities, their results also suggest that naive players converged toward the equilibrium more quickly across multiple games when they competed against chess players than when they faced other students. Being told that an opponent is strategically sophisticated may not be as powerful as actually experiencing competition with them and being forced to adapt.

An even more striking example of this convergence is provided by Rapoport et al. (2003). This study used a three-person variant of the centipede game in which college students played a series of 60 games against randomly matched opponents. The experiment involved considerably higher stakes than in previous studies, with the payoffs received after nine turns equivalent to [$1,280, $128, $128] (versus [$5, $0.50, $0.50] on the first round). Despite this strong incentive to prolong the games, most ended within the first three rounds. Moreover, over the course of repeated competition, players strongly converged toward the equilibrium of stopping on the first turn. This adaptation was best described by a simple learning model in which players' beliefs about when the next game would terminate were updated following each game. On the whole, these students' behavior did not resemble the backwards induction strategy (in contrast to that of the expert chess players from Palacios-Huerta and Volij, 2009), but rather relied on incremental learning from competitive experience.

To conclude, the centipede game provides an elegant example of adaptation to strategic uncertainty in a competitive environment, and illustrates the divide between human behavior and the normative strategies predicted by game theory. It is not the case that people are not motivated to earn larger payoffs for themselves or that they do not understand the need to claim better options before an opponent. Rather, their choices appear to be mediated by their uncertainty about how others will act—uncertainty that can be lessened by additional knowledge about opponents or direct experience. This adaptation takes the form of faster decisions to secure a higher payoff before competitors swoop in, even if it means sacrificing greater rewards in future rounds. Of course, being offered an ever-increasing stack of money is a rare occurrence—usually people compete for resources that are themselves uncertain. What if, in addition to strategic uncertainty imposed by competition, people lack perfect knowledge about the values of available options? How does their response to competitive pressure change when the stakes can only be gleaned through exploration? Next, we examine this interplay between strategic and environmental uncertainty.

## 12.3 Combining Strategic and Environmental Uncertainty

Unlike the transparent, guaranteed returns of the centipede game, competitors must often contend with environmental uncertainty about available options in addition to strategic uncertainty. To return to the real estate example: Can house hunters be sure a property is really their dream home, or are there lurking risks like a leaky roof or encroaching development? If they are in a hot market, can they devote time to learn more about a home before someone else makes an offer? Exploration is key to making effective choices in the face of environmental uncertainty, whether people are selecting a home, an investment, or an employee. Research on the *sampling paradigm* (Hertwig, Barron, Weber, & Erev, 2004; see chapter 7) has examined how people explore uncertain options prior to making a choice between them. In the sampling paradigm, an individual is faced with a set of choice options of unknown value, each associated with an underlying distribution of outcomes with different probabilities. The individual learns about each option through exploration, generating new experiences one by one, until ready to make a final choice that leads to a consequential payoff. In contrast to the payoff structure of the centipede game, in which players know that options' values increase by a fixed factor over time, what changes in the sampling paradigm is not options' underlying values but the decision maker's uncertainty about them. With increasing experience, people can more accurately choose the option with a greater value; the sampling paradigm thus entails an incentive to explore in order to reduce environmental uncertainty.

As we have seen, however, competition may not afford the luxury of such well-informed decisions. Spending time learning about available options carries the risk that others will swoop in first and claim the best ones for themselves. The value of reducing environmental uncertainty through exploration must therefore be weighed against the possibility that others will act first. In highly competitive settings, it may be necessary to act fast based on relatively little information in order to choose before competitors. Are people able to manage this trade-off effectively?

### 12.3.1 Rivals in the Dark: Competing for Uncertain Rewards

Phillips, Hertwig, Kareev, and Avrahami (2014) investigated the effects of competitive pressure on exploration using a variant of the sampling paradigm referred to as the *rivals-in-the-dark game* (or "rivals" game). In this game, two players choose from a number of options of uncertain value—in this case, urns on a computer screen (see figure 12.2). The urns are filled with virtual balls, each labeled with a number representing its value. The urn's value is the average of the values of the balls it contains. As in the sampling paradigm, each option is associated with a set of outcomes—here, the virtual balls—that occur with different probabilities. In each round, players simultaneously draw from an urn of their choice, seeing only the

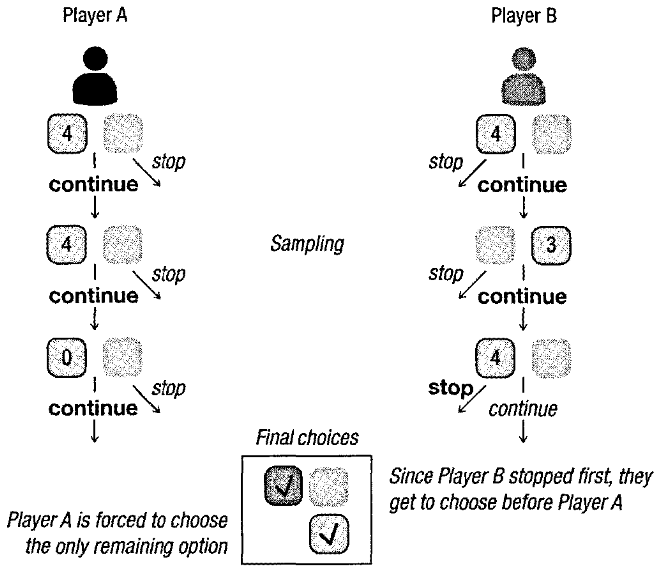


**Figure 12.2**
In this example of the rivals-in-the-dark game, two players are faced with the same set of two choice options. As in the sampling paradigm, the players can learn about each option through exploration. For example, player A experiences three outcomes from the option on the left (4, 4, and 0), corresponding to potential payoffs if that option is chosen at the end. After every simultaneous draw, each player decides whether to continue learning or to stop and choose one of the two options. In this example, player B is the first person to stop. As a result, sampling ends for both players and player B is able to make the first choice. The other player is then forced to take whichever option remains. If both players simultaneously decide to stop exploring, the first person to choose an option is determined with a coin flip.

results of their own selection. They then each decide whether to continue learning (i.e., voting to allow the game to proceed to the next round) or to stop and claim an option that will produce a consequential payoff. In contrast to the typical, solitary gameplay of the sampling paradigm, in the rivals game one player's decision to choose an option removes it from the pool of available options for the other player. As a result, deciding to continue exploring the options through repeated sampling is associated with the risk that the other player will stop and choose the better option first. You can play the rivals game yourself by accessing interactive element 12.1 (at https://taming-uncertainty.mpib-berlin.mpg.de/).

Phillips et al. (2014) found that—unlike solitary players, who typically experience 10–20 outcomes in the sampling paradigm—competitors in the rivals task drastically curtailed exploration prior to making a choice: their median sample size was just one draw, the minimum allowed. In a large proportion of games, players made decisions without observing at least a single outcome from both options. Is this meager exploration an overreaction to competitive pressure, or is it an effective strategy? Although the rivals game differs from the centipede game in a number of respects,[1] the same game-theoretic analysis based on backwards induction can be applied here. If an opponent's strategy was known beforehand (e.g., that they will stop and choose after 10 draws), the best response would be to stop one round earlier and claim the option that appears to have a higher value based on the observed outcomes up to that point. A rational opponent would be expected to respond similarly. As a result, applying the same logic backwards over decreasing sample sizes, the equilibrium strategy in the rivals game—much like that of the centipede game—is to stop exploring after the first draw. Observing even a single outcome conveys information about the relative value of the options, giving players a slight edge if they are able to make the first choice.

Minimal exploration, such as that observed by Phillips et al. (2014), may therefore be an effective response to competitive pressure in settings that

---

1. The centipede game is a *sequential move* game with *perfect information* (i.e., all players have full knowledge of the game structure and all events that occur during the game). In contrast, the rivals game described here is a *simultaneous move* game (players make decisions concurrently on each round) involving *private information* (e.g., the outcomes experienced by one player are not known to the other player).

are high in environmental uncertainty. It amounts to a bet on the value of choosing first while ignoring any potential downsides of choosing based on little information. To the extent that small samples of experience are correlated with the value of an option, people can benefit by pouncing on an option that appears favorable based on their initial experiences. The results of Phillips et al. suggest that players in the rivals game relied on such a strategy. Experiencing an attractive (positive) outcome on the first draw frequently caused them to immediately stop exploring and to choose that option; experiencing an unattractive (negative) outcome was typically associated with continued exploration. Although they often made decisions based on only very small samples of experience, players who stopped first were nevertheless more likely to obtain the *H* option than their opponents. Like other examples of simple heuristics that curtail information search by exploiting environmental structures (Gigerenzer, Hertwig, & Pachur, 2011; Todd & Gigerenzer, 2012), acting on minimal information can be advantageous in competitive settings, even when the level of uncertainty about the options' relative quality is high.

### 12.3.2 Adapting Exploration to Competitive Pressure

Should people always make such fast decisions when competing for uncertain options? A cost–benefit analysis would suggest that, when deciding whether to continue exploring, people should weigh the perceived competitive pressure (i.e., the costs of not choosing first) against the benefits of additional experience in the environment in question. Following such a strategy is particularly important when the risks of competition vary from situation to situation. The intensity of competitive pressure can change considerably as a function of resource type (e.g., plentiful vs. scarce goods), social structures (e.g., dominance hierarchies), or even time (e.g., seasonal demand). Although prioritizing fast decisions increases the chance of choosing before competitors, such a strategy may forgo rewards when competitive pressure is actually low (e.g., when there are few competitors relative to the number of available options). Using computer simulations, Phillips et al. (2014) examined how the payoffs from different sample sizes in the rivals game depend on competitive pressure. When the ratio of competitors to options is high (e.g., two players competing for two options, as in their behavioral study) and competitors are expected to make fast decisions, the minimal exploration strategy is most effective. If there are more options than competitors or competitors are expected to make slow decisions, however, larger sample sizes lead to higher payoffs.

Just as the costs of competition may change from situation to situation, the value of exploration can vary based on the structure of a given choice environment, including the distribution of potential outcomes and their respective probabilities. To take an extreme case, if the decision maker knows that the first draw from an option is a perfectly valid indicator of its value (a "sure thing" that always generates the same outcome), no information is gained from sampling it more than once. In contrast, if small samples are misleading with respect to option value (e.g., when the variance in outcomes is high or when the distribution of outcomes is strongly skewed, with infrequent extreme outcomes), exploration is more worthwhile even under intense competition. Of course, the critical question is whether people *know* that continued exploration is beneficial. Although it is difficult to make direct comparisons between choices in the rivals game and the centipede game, one reason that naive participants in the Phillips et al. (2014) study may have come closer to following the equilibrium strategy is their uncertainty about the value of further exploration (unlike the certain increase in payoffs that is common knowledge in the centipede game). After all, if the gains from exploration are unknown, why not focus on beating the other player to the punch?

In another set of experiments, Markant, Phillips, Kareev, Avrahami, and Hertwig (2018) tested whether people adapt their search based on these situational trade-offs between competitive pressure and the value of exploration. The choice environment involved options that tended to have skewed outcome distributions and in which the importance of learning about rare outcomes through repeated sampling was transparent. Each outcome was associated with a low-magnitude common outcome ($p$=.8) and a potentially high-magnitude rare outcome ($p$=.2) that could have a large impact on the quality of an option (see figure 12.3a). Participants were informed about this structure, including the number of possible outcomes and their probabilities, but had to learn the values of each outcome through exploration. In this environment, exploring until at least one of the rare outcomes is experienced improves the chances of identifying the $H$ option by more than 20% on average.

Competitors in this version of the game consistently explored far less than solitary counterparts, just as they did in the environment in the

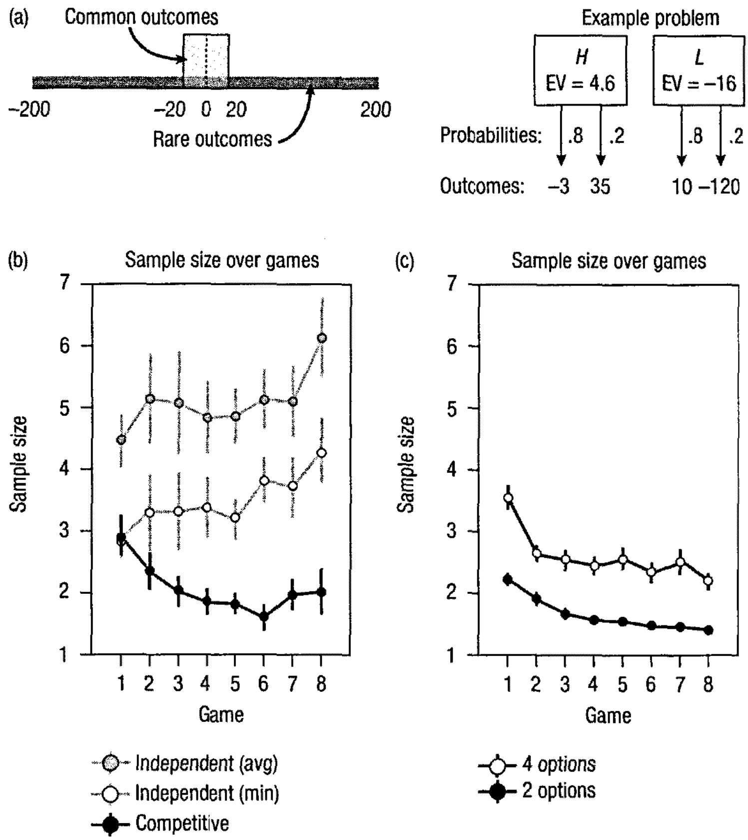


**Figure 12.3**

Design of choice environment in Markant et al. (2018). (a) Each option was associated with a common outcome that occurred with probability p=.8 and was drawn from a uniform distribution over the interval [–20, 20], and a rare outcome that occurred with probability p=.2 and was drawn from a much wider distribution over the interval [–200, 200]. (b) In this version of the task, sample sizes were higher among independent (solitary) pairs of players than among competitive pairs, and sample sizes changed in different directions as a function of experience. (c) Although competitive pairs tended to sample relatively little, sample sizes were higher when more options were available (4 vs. 2).

Phillips et al. (2014) study (see figure 12.3b). In addition, players stopped exploring prior to experiencing any rare outcomes in more than half of all games. This truncated search is especially notable to the extent that players were aware of the potential for rare, extreme outcomes that could have an outsize impact on the value of the options. However, truncated search was only part of the story. Players also adjusted their exploration in ways consistent with the behavior and simulation results of Phillips et al. When two players were competing for four options rather than two, they collected roughly double the number of samples (see figure 12.3c). In addition, players were more likely to explore when options with negative value (losses) were more common in the environment than when they were rare. Finally, even when the intensity of competition was high (e.g., when two players competed over two options), sample size declined over the course of repeated play, indicating a short-term convergence toward the equilibrium strategy that was driven by repeated experience in the game. Taken together, these results suggest that people adapt their exploration to the current choice context, including the perceived intensity of competition and the potential value of further sampling.

How do people manage this trade-off between exploring options and beating their competitor to the punch? Markant et al. (2018) developed a model in which the expected benefits of further exploration (given previously observed outcomes) are weighed against the expected cost of losing out to an opponent if they decide to stop. If a player believes that the opponent is likely to stop on the current round, the expected benefit of exploring further is outweighed by the cost of being the second person to choose an option. On the other hand, if the risk of an opponent stopping seems to be low, it may be better to continue exploring in order to reduce uncertainty about the options' relative values. This model provided a good account of behavior in the task, including the link between players' beliefs about opponents (i.e., the probability that they will stop) and the reduction in exploration over repeated play. As players gained more experience playing against the same opponent, they learned how likely the opponent was to stop on every trial and adapted their own strategy accordingly. These findings echo the behavior seen in the repeated centipede game (Rapoport et al., 2003), in that naive players appear to incrementally adapt to their competitors rather than relying on more complex forms of reasoning like backwards induction.

### 12.3.3 Know Thine Enemy: Social Norms and Learning

To date, studies involving the rivals game have focused on an extreme form of competition, with relative anonymity, minimal interaction, and little reason to care about the payoffs received by other players. Although this may resemble some especially cutthroat environments (e.g., investment trading), many competitions take place in the context of richer social interactions that change the way people perceive and respond to uncertainty. We briefly consider two examples here.

Many interactions invoke social norms that may conflict with the self-interested focus seen in extreme competition. Those norms may reduce strategic uncertainty by providing background knowledge about how people will act. For example, Fleischhut, Artinger, Olschewski, Volz, and Hertwig (2018) examined exploration in a version of the sampling paradigm where individuals had to decide how, or whether, to share a payoff. In this *mini-ultimatum game*, a player could suggest one of two ways of sharing a sum of money with their partner. The offer could then be either accepted by the partner, in which case both players received the designated amounts, or rejected, in which case both earned nothing. The choice of offers (e.g., 50:50 or 80:20 split) varied with each new game. Players were able to learn about the kinds of offers that would be accepted or rejected by other players through repeated sampling. Relative to a solitary setting with the same environmental uncertainty, people explored much less in the social framing of the mini-ultimatum game. They appeared to rely on knowledge of social norms to infer the probability of different offers being accepted, reducing the need to gain experience through exploration. For example, lopsided offers that appeared greedy were more likely to be rejected than others due to norms of fairness. Social norms, including aversion to inequitable outcomes, may operate in what seem to be purely competitive settings, especially when there is an opportunity for players to punish unfair play (Fehr & Fischbacher, 2004).

In addition to knowledge of social norms, real-world competition often affords people more opportunity to reduce strategic uncertainty by learning about competitors' goals or preferences. Whereas players in the rivals game were given little to no information about opponents' identities and observed very little of their behavior (only their final choices), observing how competitors explore might provide a clearer picture of the true competitive pressure. Taking note of how others explore can also facilitate learning about the environment. Research on behavioral ecology has examined the use of *public information,* defined as observations of competitors' choices that are used to assess the quality of a resource (Danchin, Giraldeau, Valone, & Wagner, 2004). For instance, one advantage of foraging in a group (rather than alone) is that the individual can learn about the distribution of resources by observing other group members' sampling behavior. Although patches of resources are depleted more quickly due to exploitation by a larger number of competitors, groups of foragers can use public information about others' successes to better discern when a patch is exhausted and it is time to explore further afield (Valone & Templeton, 2002). The foraging behavior of competitors can signal the quality of chosen options even when the actual outcomes are not observed (Goldstone, Ashpole, & Roberts, 2005). Extending this idea to the rivals game, if an opponent samples an option once and immediately switches to a different option, one might infer that they experienced an unambiguously negative outcome and are unlikely to claim that option. This opportunity to glean information from competitors' exploration leads to further reductions in both strategic and environmental uncertainty that would not be possible under minimal exploration. We examine the benefits of information sharing in noncompetitive social settings in further detail in chapter 13.

## 12.4 Learning to Compete: Toward an Ecological Perspective

Strategic uncertainty poses a qualitatively distinct challenge from environmental uncertainty. For all its vagaries and dynamism, nature is, like a roulette wheel, oblivious to the plans of its inhabitants and does not intentionally act to thwart them. In contrast, competitors are often aware of their opponents and can attempt to reason through how interactions will unfold. Yet whereas sophisticated strategic reasoning is often considered a pinnacle of human cognition (as exemplified by master chess or Go players), the experimental findings discussed above (see sections 12.2.1.2 and 12.3.2) suggest that people rely on incremental, experiential learning to compete in new interactive settings. In both the centipede game and the rivals game, players appear to incrementally update their beliefs about how competitors will act and adjust their strategies accordingly with repeated experience. In both tasks, learning to compete entails acting more or less quickly in response to the competitive pressure experienced.

Of course, as any chess grandmaster knows, competing effectively may at times require more sophisticated or complex strategies in order to outplay opponents. Some researchers have argued that strategic interaction is an important evolutionary driver of cognitive ability (Dunbar, 1998; Whiten & Byrne, 1997). As we have illustrated in section 12.3.1, however, competition itself does not necessitate a sophisticated response (Hertwig & Hoffrage, 2013), particularly when it is a stable property of the social environment. An ecological perspective would suggest that people adopt more complex strategies for dealing with strategic and environmental uncertainty only when compelled to do so by the trade-offs in their natural environment.

Using evolutionary simulations, Hintze, Phillips, and Hertwig (2015) have provided an example of this process in environments similar to the rivals game, each with a different level of competitive pressure, and with a focus on evolutionary adaptation (see chapter 15). Under *direct* competition, two agents could explore to learn about the value of a common option. They could then either choose that option or a private reserve option of known value (which could not be taken by the opponent). In comparison, an *extreme* competition condition, similar to the game used by Phillips et al. (2014), involved two agents who learned about two options and could claim either one, with no private options to fall back on. The strategy that evolved under direct competition was one of adaptive exploration: agents frequently sampled more than once, and the likelihood of continued exploration increased with the variance of the options' outcomes (i.e., when environmental uncertainty was higher). In contrast, the strategy that evolved under extreme competition was to choose after a single draw, regardless of the degree of environmental uncertainty. Thus, a minimal exploration strategy (i.e., one that is insensitive to environmental uncertainty) emerges when extreme competition is a stable property of the social environment, whereas adaptive exploration persists only under less stringent competitive conditions.

Understanding the range of competitive pressures faced by individuals may help to explain how they respond to strategic uncertainty. Does their willingness to explore depend on the intensity of competition they have experienced? Do different life histories influence the ability to adapt to new strategic circumstances? Recent work has suggested that other examples of seemingly impulsive decision making may in fact reflect adaptation to stressful or unreliable social environments (Frankenhuis, Panchanathan, &

Nettle, 2016; Griskevicius et al., 2013; Kidd, Palmeri, & Aslin, 2013). Just as "Depression babies," who have endured periods of financial hardship, may be less likely to take financial risks (Malmendier & Nagel, 2011; see also chapter 10), people who have experienced intense competition for resources (e.g., due to their socioeconomic background or experience in a highly competitive industry) might tend to explore less, even when competitive pressure is low. Experiencing a wide range of strategic situations, on the other hand, may foster the ability to balance the value of exploration against the current risks of competition (as in Markant et al., 2018) or to seek out social information to reduce strategic uncertainty about opponents. If the "rules of the game" are themselves uncertain, competing successfully may require that people adapt quickly based on their early experiences in a new domain.

## 12.5 Conclusion

Traditional approaches to studying decision making under uncertainty often cast the decision maker as a lone explorer on uncharted seas, buffeted by forces beyond their influence or affinity. The rivals-in-the-dark game highlights that it is necessary not only to examine "games against nature," but also to understand how behavior changes when people's landscape is populated with others. The consequences of a choice often depend on the social context, including how other people anticipate, judge, and respond to one's actions. We have examined how strategic uncertainty arises in competitive settings and interacts with environmental uncertainty in driving decision making. We have also stressed the need for an ecological approach that considers the kinds of competition experienced in different environments, including the information channels through which people can learn about others.