

Variation in Decision Making

Sasha R. X. Dall, Samuel D. Gosling, Gordon D. A. Brown,
Niels J. Dingemanse, Ido Erev, Martin Kocher,
Laura Schulz, Peter M. Todd, Franjo J. Weissing, and Max Wolf

Abstract

Variation in how organisms allocate their behavior over their lifetimes is key to determining Darwinian fitness, and thus the evolution of human and nonhuman decision making. This chapter explores how decision making varies across biologically and societally significant scales and what role such variation plays when trying to understand decision making from an evolutionary perspective. The importance of explicitly considering variation is highlighted, both when attempting to predict economically and socially important patterns of behavior and to obtain a deeper understanding of the fundamental biological processes involved. Key elements of a framework are identified for incorporating variation into a general theory of Darwinian decision making.

Why Is Variation in Decision Making Important?

Decision Making from a Darwinian Perspective

How individuals behave is of profound importance in both the human-centered and the biological sciences. In the human-centered sciences, as well as in attempting to understand the human condition and what its key influences are, it is often important to be able to predict individual behavior and what will happen to it as a result of specific interventions. From an evolutionary perspective, for any evolved entity, how this entity behaves over its lifetime is central to determining how it accumulates resources and allocates them to the crucial demands of surviving to influence how its genes persist into future generations. From such perspectives, decision making—how individuals choose to allocate their behavior—is likely to be a key focus. Here we explore how decision making varies, what role such variation is likely to play when trying to understand decision making from a Darwinian perspective, and how we can use such an understanding to promote more effective models of human behavior.

Variation in Decision Making Matters!

Why focus on variation in decision making? The simple answer is that there is abundant evidence from a range of perspectives that the way in which the allocation of behavior varies really matters for understanding and predicting the human condition and in broader biological contexts. To briefly illustrate:

Variation Is Needed to Predict What Actual People Do

Comparing the most popular process-based models of human behavior with the most successful applications of models of human behavior in the “real world” reveals an interesting distinction. The popular models (e.g., prospect theory; Kahneman and Tversky 1979) focus on capturing typical human behavior and pay little attention to individual differences. In contrast, many of the successful applications of behavioral models are based on the assumption that observed behavior reflects robust individual differences in decision making. The most important examples of these applications are recommender systems that are used in e-commerce sites, like amazon.com (see Resnick and Varian 1997). These systems assume that future behavior of a target individual (e.g., a person’s tendency to buy a specific product) can be predicted from the behavior of individuals that are similar to the target individual. In other words, these systems use past behavior to classify individuals to one of several classes, and use this classification to predict behavior. Another important set of examples involve psychometric and psychological tests. These tests measure individual differences and are effectively used to predict future behavior. For instance, performance in the SAT test correlates significantly (with coefficients between 0.36 and 0.65) with performance in college (Ramist et al. 1993).

Personality Variation Influences Life Outcomes

In the human domain, personality measures have been shown to predict consequential life outcomes, such as mortality, divorce, occupational choice, occupational attainment, health, community involvement, and criminal activity (Ozer and Benet-Martinez 2006; Roberts et al. 2007). For example, one 14-year prospective study found that the personality trait of trust predicted mortality with an effect size (r) of $-.22$ (Barefoot et al. 1998). Another study found that girls’ scores on a test of aggression predicted ($r = .30$) their likelihood of being divorced (vs. intact marriage) 28 years later (Kinnunen and Pulkkinen 2003). A recent meta-analysis of prospective longitudinal studies that controlled for background factors (e.g., existing health conditions, age, gender) found that the magnitude of effect sizes for personality on such important life outcomes were similar to those found for socioeconomic status (SES) variables and cognitive ability (Roberts et al. 2007). For example, in the domain of marital outcome, scores on conscientiousness, neuroticism, and agreeableness

predicted divorce with average effects (r) of $-.14$, $.18$, and $-.16$, respectively; in the same set of studies, SES predicted divorce with average effects of $.05$ (Roberts et al. 2007).

Variation Influences Society

Differences among human societies can result from a range of factors. These include the fact that small groups often consist of heterogeneous decision makers whose composition can differ substantially due to stochastic sampling effects, specific decision makers self-select into certain types of groups (e.g., assortative matching), specific voting or decision-making rules vary within groups, and preferences often shift (group polarization; for an early overview, see Pruitt 1971) as a direct consequence of group membership. An example of the latter, the “risky shift” (Stoner 1961), implies that groups consisting of, on average, more risk-averse group members come up with even more risk-averse unanimous decisions, while groups whose members are, on average, more risk-prone make more risky unanimous decisions. Thus, existing individual predispositions can be reinforced in the group decision-making process.

In general, it is unclear whether unitary groups—groups which do not have any internal conflict in terms of objectives and that have to come up with a joint group decision—are “better” decision makers than individuals. The relative advantage of groups over individual decision makers depends on the nature of the task, on the organization of the group, and the inclination to fall prey to group decision-making biases such as “groupthink” (Janis 1972). One important advantage of group decision making is the possibility of aggregating knowledge if information is distributed heterogeneously among group members. An apparent disadvantage of group decision making is the decision-making process, which is supposed to be more complicated and slower in groups than in individuals. Both effects are likely to be significant for individuals that self-select into groups. Across different simple economic tasks, a surprisingly uniform proportion of about two-thirds of human decision makers self-select into a group decision-making mode, whereas one-third prefers to decide alone (Kocher et al. 2006).

Cultural Variation: Practices, Freedom, and Wasteful Consumption

The study of dress codes reveals two important variation-related problems. First, conservative religious groups use strict dress codes to control their members and reduce their ability to leave the group (Arthur 1999). The strict codes reduce the members’ opportunities to interact with members of other groups and facilitate enforcement of the groups’ rules. Thus, the among-group variation, implied by strict dress codes, helps reduce within-group variability and can impair the freedom of their members. At the same time, however, strict dress codes have many positive effects. Without dress codes, rich group

members can benefit from signaling their high status with expensive clothes and/or with high within-person variability (e.g., changing and washing trousers every day by modern Israelis). It is easy to see that such outcomes can lead to wasteful cultural practices since “wastefulness” is often a reliable signal of wealth (e.g., modern Israelis use much more water than they otherwise would).

Variation Has Ecological Impacts

Behavioral variation is often correlated with resource-use and resource-exploitation strategies. Since competition is most intense among the most similar strategies, intraspecific competition may be reduced in the presence of variation. Variation may therefore be associated with a higher carrying capacity of the population. The same holds when different variants have a synergistic effect upon each other (such as in division of labor). For example, bold-shy pairs of great tits seem to have a higher reproductive output than either bold-bold or shy-shy pairs (Both et al. 2005). Indeed, such variation seems to be important for reproductive outputs even where such effects are more complex, such as in zebra finches where partner matching is only important for the most exploratory and aggressive individuals (Schuett et al. 2011).

Some recent theoretical and empirical work in the field of animal personality may be used to illustrate that variation may affect ecological processes (Cote et al. 2011; Fogarty et al. 2011). For example, Fogarty et al. (2011) modeled the spread of invasive species and concluded that the populations consisting of a mix of social and asocial individuals spread faster than populations consisting of either one. Furthermore, experimental work on invasive fish species showed that the average level of boldness in the population affected dispersal distance of individuals, implying that the social environment imposes selection on individual actions.

Variation Influences Evolution

It is well known to game theorists that the existence of even small degrees of random variation can have major implications for game theoretical predictions. The trembling hand approach (where optimal decisions are implemented with error) of Reinhard Selten (1975) is used in economic and evolutionary game theory to distinguish “unreasonable” equilibria (which only make sense in the absence of variation) from more reasonable ones. Indeed, such decision errors are crucial for solving the more computationally intensive game theoretic models such as state-dependent dynamic games, which are growing in popularity in evolutionary biology (Houston and McNamara 1999).

In addition to selecting among equilibria, new equilibria may become available if there is sufficient (stable) variation in behavior. Several modeling studies on the evolution of cooperation demonstrate, for example, that stable cooperation is achievable and stable once there is sufficient behavioral

variation in the population (McNamara and Leimar 2010). McNamara et al. (2004) consider a finitely repeated prisoner's dilemma game where the interaction between two players is terminated once any of the players' defects. In the absence of variation, "always defect" is the only equilibrium. If, however, nonadaptive alternative strategies arise repeatedly in the population (e.g., by mutation), then it may actually be adaptive to "exploit" this variation by adopting a slightly more cooperative strategy than always defect. This may initialize a positive feedback process of "variation begets variation" that eventually leads to the establishment of cooperation. There is a second reason why variation can stabilize cooperation. Several models have stressed that partner choice and partner inspection are generally favorable for the evolution of cooperation. However, choosiness and inspection will typically be costly. To persist stably in the population, these costs have to be off-balanced by some benefits. Such benefits can only accrue if there is variation (choosiness does not make sense, if there is no variation to choose among). Indeed, McNamara et al. (2008) showed that in a snowdrift game with partner choice, the degree of cooperation achieved was positively related to the degree of behavioral variation. In a different setup, Wolf et al. (2011) showed that small degrees of variation can select for social sensitivity which, in turn, stabilizes and enhances this variation. Again, the outcome is qualitatively different than in the absence of variation. Finally, McNamara et al. (2009) have shown how small amounts of variation, in how trustworthy individuals are, select for conditionally trusting strategies which, in turn, select for further individual differentiation in trustworthiness.

In addition to changing the outcome of evolution in a qualitative way, the existence of variation can have a quantitative effect, since it can speed up evolution enormously. Evolutionary biologists have always been amazed at how rapidly populations can adapt to novel circumstances. In particular, the rapid evolution of a highly integrated phenotype (requiring the evolution of correlations among a multitude of traits) is, at least to some evolutionary biologists, an unresolved mystery; this is sometimes referred to as Haldane's dilemma (Haldane 1957). The existence of a "standing stock" of variation may provide an explanation. First, evolution is much less mutation limited in such cases. If selection can "feed" on existing variation, it does not have to wait for the appearance of the right type of mutations. In individual-based simulations with realistic parameter settings, the mutation process typically imposes severe limitations on the speed of adaptation. This is particularly true in cases where selection has to "solve" a design problem in a multidimensional trait space. West-Eberhard (2003) suggested that a correlational structure as found in personalities might enhance the adaptive potential of a population. These correlations might gradually be shaped by selection in reaction to small-scale fluctuations of the environment. When the environment suddenly changes more dramatically (in the direction of the earlier small-scale fluctuations), the required correlational structures are already in place and do not have to evolve from scratch.

In addition to these direct effects on the course of evolution, the phenomenon of personality differences (stable behavioral differences among individuals) may force scientists to reconsider evolutionary processes in a fundamentally new way. Until now, the trend has been to atomize behavior into different functional contexts (e.g., finding food vs. finding mates vs. caring for young etc.), which are subsequently tackled in separate evolutionary models. This may be quite misleading. If, for example, boldness in response to novel and potentially dangerous environments is fundamentally linked to aggressiveness in intraspecific contests, it does not make sense to investigate the “boldness game” as being separate from the “aggressiveness game.” In fact, a behavior can appear maladaptive when viewed in isolation but make perfect sense (e.g., as a costly signal) when viewed from a more integrative perspective. There are several examples where this does indeed seem to be the case (e.g., Johnson and Sih 2005).

What Are the Scales at Which Decision Making Varies?

There is an interesting dichotomy between the human-centered sciences and biology in that there is a tendency in the former to assume stable individuality as the default. In contrast, students of nonhuman animal behavior who think about its evolution tend to presume that behavior is highly flexible and optimally tactical (from an evolutionary perspective) in every possible context. For the purposes of this chapter, it is important to provide a brief overview of the patterns of variation in behavior and decision making that are of interest to researchers of human and nonhuman behavior (i.e., that have significant economic, social, and ecological consequences).

Temporal patterns of variation can be distinguished by whether they operate on developmental versus intermediate versus moment-by-moment timescales. Evolutionarily significant units of variation range from within an individual, among individuals, to groups and populations. The latter are distinguished biologically by the fact that populations tend to be divided by gene flow, whereas groups are divided by behavior (interactions). In the human species, populations can be thought of as being synonymous with human cultures as they are both defined by limits to replicator spread (population: genes; culture: cultural replicators like memes). The key difference is that cultural barriers are very often permeable to highly adaptive (useful) variants (e.g., metal axes replacing stone axes), whereas this will not necessarily be the case for gene flow among populations.

Are all patterns of variation interesting? The term “animal personalities,” for example, refers to variation in highly structured behavioral types: individuals differ in their behavior and these differences are correlated over time and across different contexts (e.g., aggression, foraging, mating). To judge whether a given correlation is surprising, it seems important to have a “null model”

at hand; that is, a model of trait correlations (over time, across context) that serves as a standard of comparison. At present such a null model does not exist. This may be for good reasons: developing a nontrivial “general-purpose” null model in the context of behavioral types may be very difficult or even impossible. The trivial null model is that there are no correlations. This null model is not really interesting: rejecting it is like tearing down a straw man in which nobody really believes. The reason one does not believe in the trivial null model is that any “realistic” mechanism will create some correlations. However, the details of such correlations will depend very much on the mechanisms in question. As a consequence, a nontrivial null model will have to refer to a specific underlying mechanism. This is problematic, since we typically do not have a good idea of what the “real” mechanisms are and what kind of correlation structure they will generate. So, even if we ignore the fact that there is a multitude of such nontrivial null models, we are currently unable to characterize their statistical properties (specifically or in general). Therefore, for the time being we will have to continue using “biological intuition” when judging whether observed correlation structures are interesting or not.

Are there any general patterns of variation nature? Repeatability, the proportion of the observed variance in the population that can be explained by stable differences among individuals, is of key importance when studying labile phenotypic traits like behavior and physiology, and it has therefore been documented for a wide variety of behavioral traits across a wide range of species. Recent meta-analyses show that values of repeatability are on average 0.37, implying that a major proportion of the observed variation in behavior (e.g., up to 0.63) is within rather than between individuals of the same population (Bell et al. 2009). Some of the variation in behavior is evident at the population level, shaped by local adaptation to population-specific ecological conditions like predation risk. For example, work on three-spined sticklebacks, *Gasterosteus aculeatus* (Figure 15.1), revealed that 9.4% of phenotypic variation in predator-inspection behavior could be attributed to variation among 12 Welsh populations, 41.2% to variation among individuals within these populations, and 49.4% to within-individual variation (Dingemanse et al. 2007). To further illustrate how behavior and decision making varies at different scales, we now proceed by highlighting some interesting examples from the human-centered and biological literature.

Variation within Individuals

Standard economic theory relies on the assumptions that individual preferences are relatively stable over time (Stigler and Becker 1977) and that behavior is well-approximated by assuming a representative individual, especially if the decision-making situation has a unique optimum or equilibrium. The former assumption implies that decision makers choose the same option in the same decision-making situation (under the same conditions/states) over time. The

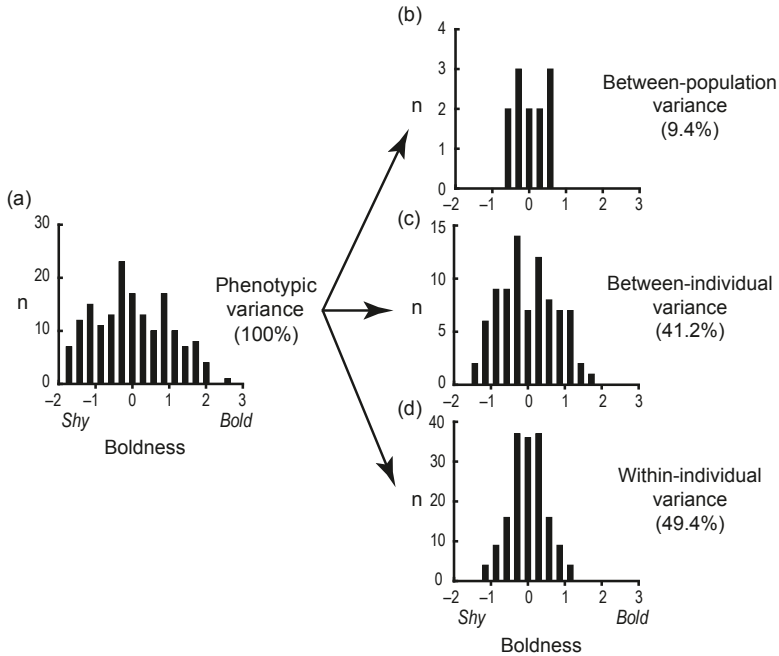


Figure 15.1 Variation in predator-inspection behavior (ranging from shy to bold) across hierarchical levels. Histograms show the variation in “boldness” at each level. (a) Distribution of the raw phenotypic data ($n = 168$ assays: 2 repeat assays for each of 7 individuals within each of 12 populations). (b) Distribution of population-average values ($n = 12$ populations). (c) Distribution of individual-mean values ($n = 84$ individuals) expressed in deviations from the population mean (i.e., we show here the distribution of individual mean values within populations). (d) Distribution of the deviations from individual-mean values ($n = 168$ deviations) (i.e., the within-individual variance). The data used in this example are from exploration of altered environments by three-spined sticklebacks (see Dingemanse et al. 2007).

latter assumption is not critical for the application of (game) theory, and more recent models often simply assume the existence of different types of decision makers with different optimal points (e.g., selfish types and reciprocal types) and specify their interactions formally. There is little dispute about the fact that the assumption of a representative individual is too restrictive in many situations. Different levels of other-regarding preference (e.g., Fehr and Schmidt 1999; Bolton and Ockenfels 2000; Charness and Rabin 2002; Fehr et al. 2007), different degrees of rationality or bounded rationality (e.g., Nagel 1995; Costa-Gomes and Crawford 2006; Crawford and Iriberri 2007) and different attitudes toward uncertainty (e.g., Abdellaoui et al. 2011) across decision makers lead to heterogeneity in final choices and behavior.

The intertemporal stability/plasticity of individual preferences in human decision makers is a much more controversial issue, and it affects the basic modeling methodology in standard microeconomics (e.g., crucial axioms

associated with preference revelation; Samuelson 1938). No one denies that preferences could and would change over long time horizons (e.g., in the context of trust/trustworthiness and age; Bellemare and Kröger 2007; Sutter and Kocher 2007). However, studies that look at changes in behavior over short time horizons by the same decision makers have recently been conducted, and the results are not fully conclusive. Some find quite high levels of robustness in behavior over time (Volk et al. 2012), others report very low levels of robustness for the same behavioral tendencies in slightly different contexts (Blanco et al. 2011). Imprecision is also noted as one important reason for intertemporal instability (Butler and Loomes 2007). An empirical problem is that providing decision makers with exactly the same tasks over time is confounded with learning, and providing decision makers with similar tasks over time is confounded with potential context or framing effects. Hence, intertemporal stability of individual preferences is hard to measure precisely in the controlled environment of an experiment.

Recent studies of learning reveal an interesting general pattern: large within-subject variability, which can sometimes swamp species differences. One indication of this pattern is presented by Erev and Haruvy (2012). Their review of the classical learning literature shows that many of the phenomena, originally documented in studies with different animals, can be reliably replicated in the study of human behavior in simple computerized experiments. One example is the partial reinforcement extinction effect (Humphreys 1939). Human and nonhuman subjects learn less in a noisy environment (under a partial reinforcement condition), but learning in this noisy setting is more robust to extinction. At the same time, the results of these studies reveal large within-subject variability (see Shafir et al. 2008). For example, when faced with a repeated choice between “3 with certainty” or “80% chance to win 4; 0 otherwise” the typical subject does not learn to prefer one of the two options. Rather, the typical subject selects the risky prospect in about 60% of the trials. High within-subject variability was documented in human subjects even after 400 trials with immediate feedback. Erev and Haruvy (2012) show that the main experimental results can be captured with simple models which share the assumption that subjects tend to rely on small sets (about four) of past experiences in similar situations. Different samples are used in different trials; as a result, these models imply a payoff variability effect (Myers 1960; Busemeyer and Townsend 1993): large payoff variability increases within-subject variability.

Variation among Individuals

Variation in decision making among individuals can occur at different levels. Individuals can differ in (a) the mechanisms that underlie behavior (e.g., physiological or cognitive systems underlying behavior), (b) their evaluation systems, and/or (c) states that affect decision making (e.g., information about features of the decision problem). A “decision-making mechanism” may be

viewed as a device that “chooses” an action (possibly in a nondeterministic way) for a given set of external stimuli and (external or internal) state variables. It makes sense to speak about “variation in decision making” if the underlying decision-making mechanisms vary (i.e., if in a systematic way different actions are chosen despite similar states and stimuli). Such variation can, for example, occur if individuals differ in the way they receive information (e.g., variation in receptors, as described for female sticklebacks), process that information, evaluate the processed information, translate this evaluation into action, and perform the corresponding action. One complication in the above description is the concept of “state.” When do we consider a “state” to be part of the decision-making mechanism, and when is it something external to it? For example, “memory” may be viewed as a state that is external to the decision-making process (e.g., the memory of previous outcomes is an input of conventions like “winner-loser” effects), but the use of memory may be a distinguishing feature of a decision-making mechanism. Even if state variables are clearly external to the mechanism, they may still create considerable confusion. Let us assume that males and females differ substantially in the way they make decisions. Technically speaking, we could still consider this as the outcome of one and the same decision-making mechanism that has “gender” as one of its input variables. On the other hand, we can always assume that there are subtle, unobserved differences in state, when individuals are seemingly taking different decisions. So we have to be pragmatic. If we are interested in systematic variation in behavior that cannot in any “obvious” way be explained by differences in environment or differences in state, we should not forget that we are also interested in variation that is patterned, stable in time, and consistent across contexts.

Ecologists and evolutionary biologists have provided convincing evidence for the existence of individual variation in decision making in a wide range of species. Empirical evidence for individual differences comes from studies in which the same set of individuals are assayed for the same behaviors repeatedly, such that the variation in the sample can be decomposed into between- and within-individual variation. Values of repeatability, the proportion of total variance that comes from individual differences, suggest that the majority of the variation in behavior is within rather than among individuals in animal populations (Bell et al. 2009). Nevertheless, repeatable variation implies that the average level of behavior differs among individuals; however, it does not imply that individuals are completely stable in their decision making (see Figure 15.2a). Fleeson (2001, 2004) discusses similar issues in a human personality context. While the existence of repeatable variation is well documented, there is growing awareness that individuals may differ in how they change their behavior in response to variations in environmental conditions or age—phenotypic plasticity (Dingemanse et al. 2010b; Nussey et al. 2007). Plastic individuals alter their behavior as a function of context, whereas nonplastic individuals do not (Figure 15.2b). Individual variation in plasticity,




Figure 15.2 Graphic depiction of how between- and within-individual variance components are separated by plotting seven measurements of a behavioral decision (Y-axis) for five individuals (numbered) whose behavior was assayed over an environmental gradient (X-axis). For example, the Y-axis might represent aggressiveness and the X-axis conspecific density. (a) Gray lines represent the average phenotypic value of each individual; the variance among lines represents the between-individual variance. The variance in within-individual deviation from individual means represents the within-individual variance. (b) Individuals differ in average phenotype as well as in phenotypic plasticity.

termed “individual by environment interactions” in the evolutionary literature, usually explains about 5% of the variation in decision making within animal populations (Nussey et al. 2007).

As an example of long-term calibration of behavioral phenotype by some other characteristic, in humans, men who are physically large tend to be more aggressive. However, it turns out that the key parameter is not adult size, but size relative to others at a certain age in adolescence (e.g., Pellegrini and Long 2002). Men who are larger than their peers at this age adopt a more aggressive phenotype, and this persists even though their peers may subsequently grow to be as large as they are. This could be an example of what is referred to in the human personality literature as reactive heritability. That is, size variation is heritable, and aggression is calibrated to this in development; thus aggression ends up showing effective heritability in behavioral genetic studies. A similar effect occurs with extraversion, where it has recently been argued that much of the heritable variation is in fact variation in physical attractiveness. Attractive people are reinforced for initiating social interaction and become more extroverted as a consequence, causing an effective heritability of what is in fact a developmentally calibrated trait (Lukaszewski and Roney 2011; Meier et al. 2010).

More generally, dominance hierarchies provide good examples for systematic individual differences in behavior: individuals at the top of the hierarchy behave very differently in many respects (depending on the species in

question) than individuals at the bottom. One might think that the position in the dominance hierarchy and, hence, behavior related to that position, just reflects individual differences in fighting ability (“resource holding potential”). In a modeling study, van Doorn et al. (2003) showed that the “winner-loser effects” on which dominance hierarchies seem to be based can evolve from scratch even if all individuals have the same fighting ability, irrespective of the outcome of previous fights. In other words, winner-loser effects (and, hence, dominance hierarchies) may not be more than a “social convention”; that is, a conditional strategy which allows individuals to coordinate their actions, thereby avoiding escalated contests. Of course, differences in fighting ability will often be reflected in an individual’s position in a dominance hierarchy. However, there are various indications that this is not the whole story. First, experiments with a variety of organisms have revealed that the position on a dominance hierarchy can actually be quite arbitrary (Hsu et al. 2006). In these experiments, groups with an established dominance hierarchy are taken apart and the individuals are reshuffled over new groups. After reshuffling, the position of individuals in the newly establishing dominance hierarchies is often only loosely related to their position before reshuffling. Second, dominance is often site specific. Territorial animals are good examples: within its territory, an animal is dominant, although it is subdominant in the territory of another individual. The relative dominance status of two individuals *A* and *B* on neighboring territories can thus be shifted from high to low by observing these individuals on the transect from the center of territory *A* to the center of territory *B*. Similarly, site-specific dominance is also observed in nonterritorial species, like oystercatchers in the winter (Ens and Cayford 1996) or white-throated sparrows (Piper and Wiley 1989). Here an individual *A* can be dominant over *B* at site *X*, while *B* is dominant over *A* at site *Y*; in contrast to territory ownership, the individuals are only loosely attached to the sites where they are dominant. In these cases, dominance is apparently a social convention that is not explained by “obvious” differences, such as differences in fighting ability.

Beyond differences in competitive behavior, individuals can differ systematically in other ecologically important ways. For instance, lions manifest two types of social organization. Some are residents, living in groups, called prides. The pride usually consists of five or six related females, their cubs of both sexes, and a coalition of males who mate with the adult females. The number of adult males in a coalition is usually two. Other individuals are called nomads, who range widely, often in pairs of related males. Males have to go through the nomad lifestyle before they can become residents in a pride, which is always different from the pride into which they were born. The resident males associated with a pride tend to stay on the fringes, patrolling their territory. Both males and females defend the pride against intruders. The males associated with the pride must defend their relationship to the pride from nomad males who attempt to take over their residency position. When a group of male nomads oust the previous males associated with a pride, the conquerors often kill

any existing young cubs, presumably because females do not become receptive until their cubs mature or die. A lioness will often attempt to defend her cubs fiercely from a usurping male, but such actions are rarely successful. In addition to these sex and social role differences, there are strong individual differences in the behavior of both male and female lions. The classical example concerning female behavior comes from Heinsohn and Packer (1995:1260):

Female lions (*Panthera leo*) showed persistent individual differences in the extent to which they participated in group-territorial conflict. When intergroup encounters were simulated by playback of aggressive co-vocalization, some individuals consistently led the approach to the recorded intruder, whereas others lagged behind and avoided the risks of fighting. The lead females recognized that certain companions were laggards but failed to punish them, which suggests that cooperation is not maintained by reciprocity. Modification of the “odds” in these encounters revealed that some females joined the group response when they were most needed, whereas other lagged even farther behind. The complexity of these responses emphasizes the great diversity of individual behavior in this species and the inadequacy of current theory to explain cooperation in large groups.

Several follow-up studies have been conducted and similar differences were found in experiments with males (Grinnell et al. 1995). Again, the males within a coalition strongly differ in their tendency to launch an attack on rivaling males.

Variation among Groups

There have been many reports of systematic between-group differences in lion behavior. For example, there are pronounced differences in how prides of similar size hunt down prey in cooperative hunting efforts. There are also pronounced differences in the ways groups of male nomads try to conquer a reproductive position in a pride. These differences, however, are difficult to quantify, and it is not clear whether such consistent differences among groups corresponds to something like a “group culture” or whether it is just a reflection of the different “personalities” of their constituent members. The differences alluded to above have mainly been reported in books (e.g., Schaller 1972), more general reviews (e.g., Packer 1986), or—not unimportantly—in nature documentaries (e.g., National Geographic: Super Pride). Nevertheless, work by Sapolsky (1990) indicates that there may also be stable differences in behavior among baboon groups that persist over generations. In one group, Forest troop, an outbreak of bovine encephalitis led to the death of several of the dominant males in the group. Compared to neighboring control groups, social relations in the Forest troop were peaceful. There was less aggression and more reciprocal grooming. Interestingly, this difference persisted over a ten-year period in which all of the original males died and were replaced by immigrant males.

Group differences that emerge from simple interactions among individuals are well documented. Group composition is often influenced by locomotion performance (Krause and Ruxton 2002) and can be seen in flocks of birds, schools of fish, and herds of mammals (Couzin and Krause 2003). For example, herds of African ungulates are often structured according to walking speed, with faster individuals at the front and slower individuals at the back of the group. In the extreme, this can result in social segregation with lactating and nonlactating females in zebras forming distinct herds (Gueron et al. 1996). Another mechanism that can structure ungulate herds is active preference. For example, Thomson's gazelles, *Gazella thomsoni*, actively associate with Grant's gazelles, *Gazella granti*, because the latter are taller and have a greater ability to detect an approaching predator, such as a cheetah, *Acinonyx jubatus* (FitzGibbon 1990). Grant's gazelles appear to benefit, in turn, from the presence of the Thomson's gazelles because the latter are the preferred prey of cheetah. Given an attack on the group, the per capita risk of a Grant's gazelle is therefore lowered by the greater group size, due to the presence of Thomson's gazelles.

Variation may also exist at higher levels in a hierarchy. In particular, populations might be comprised of social neighborhoods differing in decision making. Variation among social groups has, for example, been observed in shoaling fish, where shoals within the same population differ in average behavior (Magnhagen and Staffan 2005). Indeed, between-group variation can emerge at equilibrium when members of different groups are motivated to behave like members of their group, and are punished when they behave like members of other groups. This can be illustrated in a multiagent market game that is played by 20 players: 10 Red and 10 Blue players. Each player has to select a location: left or right. Agent j gains 1 point for every other member of j 's group who selects like her, and loses 1 point for each member of the outside group that does likewise. At equilibrium, all Red players will select one location and all Blue players will select the other.

Factors that motivate group members to behave similarly are well documented. Social contagion ("peer effects" in economics) is typically thought of as imitative behavior that is not merely a result of homophily or of the fact that neighboring social agents are likely to be subject to common influences. Specifically, there has been much recent interest in research surrounding suggestions that imitative behavior evolves in networks over time, such that smoking cessation and levels of alcohol consumption propagate readily through social networks (Christakis and Fowler 2007, 2008). It is, however, difficult to separate correlated influences (Ioannides and Topa 2010; Manski 1993), so interpretation remains controversial (e.g., Cohen-Cole and Fletcher 2008). Why does social contagion occur? Imitation can serve a social learning role (Chamley 2003), may reduce cognitive effort (Epstein 2001), or may reflect a concern with relative social position (e.g., Clark and Oswald 1998). The social contagion literature has remained largely silent (or has made claims

inconsistent with other literature) as to how exactly the relevant social comparisons occur mechanistically. Effects related to social contagion have been extensively studied in agent-based network models. Such approaches (e.g., Schelling 1978) have provided useful insights into areas such as collective behavior in ants and traders (Kirman 1993), swarming behavior (Reynolds 1987), crowd behavior (Dyer et al. 2009), population group size (Axtell et al. 2002), cultural dissemination (Axelrod 1984), and imitative voting (Bernardes et al. 2002). Most of these models incorporate an imitation parameter or mechanism of some kind, such that agents in a network are likely to change in the direction of or influence other (typically local) agents with whom they are connected. It is then typically shown that interesting emergent behavior or segregation arises as such networks evolve.

Variation among Populations/Cultures

Variation among populations has been given substantial attention in the animal behavior literature, particular with regard to the question of whether populations with different ecologies (e.g., that vary in predation risk, population density, food availability) differ in the average level of behavior expressed (Magurran 1998). For example, populations of fish which live in environments with a high predation risk are on average less bold compared to populations that experience a low predation risk (Huntingford et al. 1994).

Human personality research has implied that the structure of behavior, as captured by the Big 5 personality traits, is similar in different cultures: one might not always get exactly the same structure, but it is remarkable how often similar dimensions arise (e.g., Gosling and John 1999; Weinstein et al. 2008; McCrae and Allik 2002). In contrast, comparisons of behavioral syndrome structure (the direction and strength of correlations between behaviors) across animal populations do not generally confirm such ubiquity. For example, comparative work on three-spined sticklebacks mentioned above reveal that the ecological conditions of populations, particularly predation risk, predict syndrome structure. Specifically, populations of sticklebacks may generally be characterized by positive associations between aggressiveness, activity, and exploratory behaviors, but these associations are tighter in populations living sympatrically with predatory fish, both in European (Dingemanse et al. 2010a; Dingemanse et al. 2007) and North American populations (Bell 2005; Bell and Sih 2007). Therefore, it appears that local selection regimes shape trait associations, which may explain population differences in genetic correlation structure across North American populations (Bell 2005). At the same time, there is also evidence that certain genes are only expressed in the presence, and others only in the absence, of predators (Dingemanse et al. 2009), implying that any population differentiation at the phenotypic (observable) level is not necessarily underpinned by population-genetic variation.

In studies on humans, there is an enormous body of literature which assesses the robustness of results to cultural variation first found in simple lab-based economic decision-making situations in Western countries. Attitudes to uncertainty, negotiation behavior, cooperative behavior, reciprocity, trusting behavior, sanctioning behavior, and many more have been compared in controlled and incentivized studies with human decision makers across different cultures. Overall, differences between student populations in different countries are much smaller than expected a priori. Exceptions to this trend seem to be for punishment behavior (Herrmann et al. 2008) and perhaps conditional cooperation (Kocher et al. 2008); other differences have also been found, but they are not always robust. Sometimes the variation in behavior within a given culture is greater than the variation in behavior across cultures (Kocher et al. 2008), but defining the boundaries can be difficult. Numerous studies for the ultimatum game, for instance, have reported small or nonexistent differences in the behavior of student participants in experiments across different countries around the world (for an overview, see Camerer 2003).

In contrast, larger and persistent differences emerge when comparing standardized behavioral tests in small-scale societies. Henrich et al. (2004) find that the behavioral variation across 15 selected small-scale societies when playing standard economic games (ultimatum game, public goods game) is extremely high, but it is consistently related to two factors: the higher the degree of market integration and the higher the payoffs to cooperation, the greater is the level of cooperation in experimental games. Nevertheless, recent studies show that potential cross-cultural differences in norms are not simply changed in different cultural environments. While humans, for instance, tend to adjust their tipping behavior when they travel (Azar 2004), they do not make (full) adjustments in other domains. Fisman and Miguel (2006) show that UN diplomats from high corruption countries accumulate significantly more parking tickets in New York. Note that diplomatic immunity implies that there is almost zero legal enforcement for these violations. Hence, a standard model of decision making would not predict any differences, and the results indicate that norms are transported to different situations.

What Proximate Mechanisms Underpin Variation in Decision Making?

Patterns of human and nonhuman behavior have been the subject of substantial research for many years. It is evident from our brief tour of behavioral variation at different scales that such patterns can vary substantially. Therefore, it is important to review the range of mechanistic processes that can give rise to variation in decision making over the psychologically, economically, and ecologically significant scales that we have focused on in this chapter. Such processes

can be thought of as falling within the remit of the so-called “proximate” explanations (Tinbergen 1963) for the existence of variation in decision making.

Within Individuals

At the most fundamental level, the physiological mechanisms that control behavior (e.g., neural processes) are likely to be subject to stochasticity in the basic biochemical processes that underpin them (e.g., due to quantum fluctuations at the molecular level). The details of such “essential stochasticity” (e.g., synaptic noise) are likely to be important as they may generate substantial within-individual variation in behavioral outcomes.

Beyond neurophysiological noise, cognitive processes can also generate behavioral variability. The decision-by-sampling model (Boyce et al. 2010; Stewart et al. 2006)—a cognitive model of human judgment and decision making—assumes that attributes are judged in terms of their ranked position within a retrieved sample context (for further details, see Brown et al., this volume). Thus variability within individuals can result from different samples being retrieved from memory on different occasions, or from changes in choice context. Cognitive models have addressed the role of cognitive (and potentially variable) samples in drawing inferences and making judgments (see, e.g., Fiedler and Juslin 2005). For example, the MINERVA-DM model (Dougherty et al. 1999) has been used to explain judgments of likelihood within a memory model, and limitations of and individual differences in the capacity of general working memory have been referred to in accounting for bias in probability judgment (Dougherty and Sprenger 2006). Another tradition of research relevant to variability has suggested that participants represent distributions (e.g., of the location of an item in a sequence) but use samples from these distributions to make judgments, which can be repeated leading to improved average estimates such that successive estimates have uncorrelated errors (Vul et al. 2009; Vul and Pashler 2008).

Models of decision making via heuristics (e.g., Kahneman et al. 1982; Payne et al. 1993; Gigerenzer et al. 1999) posit that many choices and judgments are made using simple heuristics or rules of thumb, which often yield “good” decisions when used in appropriate contexts. These heuristics typically use little information and process it in a limited manner to allow rapid (or “fast and frugal”) decision making, making them an evolutionarily plausible alternative to traditional axiomatic approaches to rationality. Many heuristics have been explored that are algorithmically specified in terms of their exact information-processing steps. Less well-explored and understood are the means by which different heuristics are selected to apply to different tasks, which could be a considerable source of both within- and between-person variability in decision making. Studies have found evidence for appreciable variation in the use of heuristics across individuals (Bröder 2012; Gigerenzer et al. 1999), often because different heuristics can produce approximately the same quality

of decisions in particular tasks. Indeed, it is considered important to compare the ability of multiple heuristics to account for participant data, so that this between-person variability can be discovered (Brighton and Gigerenzer 2011); less is known about variability in each individual's use of particular heuristics over time or in different situations. Variability in the decisions made by an individual (as opposed to variation in which heuristics they use) can be accounted for by specific heuristics. Many heuristics are deterministic in their operation and thus will not produce this variation—one example is take-the-best, which chooses between two options (e.g., which of two desserts has more calories) by comparing one cue or feature at a time until the first cue is found that points to one or the other option (e.g., first considering how much butter each dessert contains, and then if butter content is equal, considering how much sugar each has, and so on). In this case, the cues are searched in a fixed order, which can vary between people based on their learning history (Gigerenzer and Goldstein 1996). Other related heuristics incorporate external environmental influences or stochasticity to create decision-to-decision variability. For instance, the take-the-last heuristic operates similarly to take-the-best but uses cues that were previously successful in producing choices (introducing variability through contingencies of learning), whereas the minimalist heuristic uses cues in a random order. Other heuristics base their decisions on aspirations determined by samples experienced over time, which will vary from one set of experiences to another (as in sequential mate search by Miller and Todd 1998; see also the sample-based heuristics in Pachur et al. 2011).

From the standard economic perspective, it is constructive to distinguish between two sources of variability in behavior. The first class involves variability that can be observed even if individuals are fully informed of the incentive structure and are fully rational (i.e., behavior is at a Nash equilibrium). Within-subject variability emerges at equilibrium in situations in which the agents do not want to be predictable. One set of examples of this comes from constant sum games with unique equilibria, like the matching penny game presented in Table 15.1.

A second source of variability includes variation that is likely to emerge when the agents do not know the incentive structure and have to rely on their past experiences (Erev and Haruvy 2012). One source of within-subject variation in this setting is exploration. Some level of exploration is necessary (and assumed by all learning models) to collect information. Another source of within-subject variability is assumed by sampling models of learning. These models imply that different choices are expected to be based on different

Table 15.1 A two-person matching penny game.

	A2	B2
A1	1, -1	-1, 1
B1	-1, 1	1, -1

samples of past experiences. A third source of individual variability involves behavioral changes that can be described as adjustments to changes in incentive structure. One set of examples comes from the effect of aging, which can lead to within- and between-subject variation in behavior. Indeed, there is an extensive literature on age changes in personality, but to link this to changes in “incentive structure” it is probably worth distinguishing the “internal” (e.g., puberty) and “external” (e.g., having children, getting a job) influences that can impact age changes (e.g., Roberts et al. 2006; Roberts and Wood 2006; Soto et al. 2008, 2011; Srivastava et al. 2003), and translate these into incentives.

Among Individuals

Continuing to think economically, between-subject variability can emerge at equilibrium in situations in which the payoff for each rational agent decreases with the number of other agents that behave like her. A simple subset of this set of situations is represented in the market game, presented in Table 15.2, in which each of two sellers has to select between two locations, and payoffs are maximized (for both players) when each seller select a different location. At the pure strategy equilibrium of the game, the sellers select different locations (i.e., there is variation in the choice of position).

More generally, there are a range of processes that can underpin variation among individuals in how behavior is expressed. On the one hand, it is possible for individuals to be employing exactly the same decision-making process (e.g., the same rule for allocating behavior as conditions vary) while experiencing different local conditions from one another (e.g., being more or less hungry due to idiosyncratic—chance—recent history of access to food). This can cause individuals to differ from each other in how they behave at any given moment in time. Indeed, one source of learning-induced between-subject variability is implied by the fact that different agents are likely to experience different outcomes. Alternatively, it could be that individuals follow distinct decision-making processes (e.g., due to genetic variation in condition-dependent behavioral rules), in which case differences in behavior among individuals will be evident even if individuals have had exactly the same experiences. For instance, some models of learning assume between-subject variability in learning parameters.

There can be substantial impacts of early developmental conditions on how individuals behave. Such ontogenetic programming (e.g., early environment and maternal effects) can amplify slight (e.g., stochastic) initial differences

Table 15.2 A two-person market game.

	A2	B2
A1	-1, -1	1, 1
B1	1, 1	-1, -1

in experience or nutrition and lead to profound differences in the way that individuals behave. For example, Bateson et al. (2004) have argued substantial variation in adult human eating behavior can be explained by early life physiological “programming” by maternal nutritional status during pregnancy. Mothers that are nutritionally stressed while pregnant induce fetuses to “expect” food to be limited and so switch on a “thrifty phenotype” in their progeny, who crave energetically dense foods and store fat whenever possible. This effect is argued to underpin a significant amount of the variation in patterns of obesity and Type II diabetes being documented in modern, food-rich developed societies.

Among Groups

As in between-individual variation, it is important to appreciate that variation among groups can emerge in two distinct ways. First, different groups may be formed by distinct types of individuals. For instance, in many species, males and females often have different basic (e.g., energetic) demands, and therefore they behave differently and form separate groups (Ruckstuhl 2007). Second, individuals may not necessarily vary in any fundamental sense among groups (i.e., they all follow the same decision rules) but the group-conditions may vary idiosyncratically (e.g., different groups just happen to be different sizes, which cause different levels of antipredatory vigilance). A good illustration of this: mechanisms for individual social learning—essentially, copying the behaviors of others—can result in behavioral differences between groups as a result of within-group social convergence (e.g., conformity) after different starting conditions. That is, if asocial learners in one group hit upon a particular behavior, and those in another group discover a different behavior, social learners in both groups may copy those asocial learners, as well as each other, to the point where the two groups diverge in their behavioral profiles. A number of social learning rules have been explored, which can be roughly broken down into those determining when copying should happen and others directing whom to copy (Laland 2004). “When” rules trigger the application of a social learning strategy and include copying when a current behavior is unproductive, copying when learning asocially (e.g., through trial-and-error) is too costly, and copying when the environment, and hence the appropriate behavior, is uncertain. “Who” rules specify the other model individuals whose behavior should be copied, and different rules can determine whether or not and how quickly a population will converge on a shared behavior. Such conformity is promoted by the widely studied copy-the-majority rule (Boyd and Richerson 1985), as well as by copying successful individuals, and possibly copying others if they are just doing better than oneself. Within-group convergence is unlikely if individuals use a copy-the-rare rule, adopting uncommon behaviors of others, though this can be individually advantageous if having a novel behavior (e.g., mating display) enhances local (or even group) competitiveness. Rules for

copying others in one's social network can lead to within-group homogeneity or heterogeneity depending on the network structure (Lieberman et al. 2005).

Language acquisition can be a pervasive social force for group differentiation in humans. Infants come into the world prepared to be part of any culture and language group. At birth, humans (like other primates) can distinguish any phoneme produced in any language around the world (Kuhl 1991; Kuhl et al. 1992). However, during the first year of life, perceptual tuning results in a loss of function: sounds that a newborn can distinguish (e.g., phonemic distinctions present only in Mandarin) are no longer distinguishable by, for example, an English-speaking one-year-old. Even brief exposure to a nonnative language can keep this perceptual window open. However, preservation of this ability depends on live social interaction. A 10-month-old from an English-speaking household who interacts with a Mandarin speaker for an hour a week will continue to distinguish the Mandarin phonemes at 12 months; a 10-month-old exposed to identical input from a video display will not (Kuhl et al. 2003). Case studies of hearing children of deaf adults also suggest that live human interaction is critical for language learning. A hearing child (1 years 8 months) who interacted only with his deaf mother but had abundant exposure to spoken English through television had no productive spoken language before intervention; his brother (3 years, 9 months), who briefly attended an English-speaking preschool, also exhibited severe delays (Sachs et al. 1981). These findings suggest that although widespread cultural transmission through media may reduce cross-cultural variability in human populations, specific aspects of human cognitive development may tend to maintain cultural differentiation. In particular, as long as there are distinct language groups among adults, those distinctions will be conserved insofar as children only learn language by interacting with the human beings in their immediate vicinity.

How Can Variation in Decision Making Evolve?

Given our focus on decision-making agents that are subject to Darwinian selection, it is appropriate to consider the evolution of variation in decision making, the so-called "ultimate" explanations (Tinbergen 1963) for its existence. Indeed, for many of the scales of variation we have discussed so far, there has been significant research effort devoted to understanding how variation in decision making can be favored by Darwinian selection.

Within Individuals

It is possible that individual behavior will vary in different contexts because there are distinct modules in the central nervous system that operate independently to control behavior for functionally (biologically) distinct types of problems (e.g., when foraging vs. choosing mates vs. caring for young, etc.).

Such modularity may be adaptive because the fundamental (e.g., statistical) properties of the problems posed in the different contexts may differ quite substantially, such that completely different control processes are required to maximize performance in each biological “domain.” It is also possible that modularity could act as a within-individual “hedging of bets” when there is substantial, irreducible uncertainty about the environment on a moment-by-moment timescale (Dall 2010), such that knowing what the appropriate set of responses is likely to be at any moment in time is too difficult (and requires too broad a range of responses) to allow a general-purpose decision rule to evolve. Nevertheless, the existence of module-general common resources (e.g., information, energy) or control traits (e.g., sensory systems) will erode modularity and could structure individual behavior across contexts (functional domains). Wolf and Weissing (2010) discuss how the latter can select for cross-context correlations in behavior.

Individual behavior often varies over time as a result of learning. The conditions under which we would expect learning to evolve are relatively well understood (Stephens 2007). For prior experiences to improve decision making (i.e., allow for a better fit between behavior and current conditions) there has to be some degree of patterning of ecologically important events over time. However, because learning can be costly in terms of effort, time, and the maintenance of complex decision-making “machinery,” if environments are too patterned in their states (i.e., do not change enough over generational time) then “hard-wired” behavior (e.g., that does best in the average environment) is likely to evolve. From such a perspective, Stephens and colleagues (reviewed in Stephens 2007) have developed a framework for predicting variation in the adaptive value of learning/behavioral plasticity by specifying uncertainty and reliability of experience on orthogonal axes. Their predictions were corroborated using experimental evolution in *Drosophila* and setting blue jays operant tasks in the lab. In a similar way, ontogenetic changes in reproductive value and ability (more generally: changes in fitness trade-offs with state) can drive longer-term variation in behavior as a function of the judicious use of information over a lifetime, when such phenotypic plasticity is adaptive.

One influential approach to analyzing the behavioral consequences of information use (including learning) from an evolutionarily adaptive perspective involves the application of statistical decision theory (Dall et al. 2005), which makes extensive use of inductive (Bayesian) inference. Most organisms, regardless of cognitive complexity, have to make guesses about the world: Is that a bright red berry in dim light or a dull red berry in bright light? Is that prickly looking animal going to eat me or can I eat him? Should I go to this restaurant or that one? The problem of inductive inference is the problem of trying to make a decision under uncertainty; many conclusions are possible given the data, and the organism must choose just one. Bayesian inference provides a computational-level, ideal (adapted) observer and analysis of how background knowledge should be integrated with statistical data to narrow

the hypothesis space. These models can be applied to problems of induction across species and cognitive domains. They have been used to describe foraging decisions in a wide range of species (Krebs et al. 1978): problems of visual perception (Yuille and Kersten 2006), decision making in sensorimotor control (Körding and Wolpert 2006), language acquisition (Chater and Manning 2006) and many aspects of abstract, higher-order cognition (Tenenbaum et al. 2011). Specifically, Bayes's law provides a rule for how an organism might evaluate a hypothesis, h , about the process that produced some data, D . Bayes's law states that the probability of the hypothesis given the data, $P(h|D)$, depends on both the prior probability of the hypothesis, $P(h)$, and the likelihood of the hypothesis, $P(D|h)$: the probability that the data would have been observed if the hypothesis were true. That is, $P(h|D) = P(h)P(D|h)$. We can borrow a simple example to illustrate (see Tenenbaum et al. 2011): Imagine you observe some data D —your child is coughing. You can consider many hypotheses, including, ($h1$) your child has a cold, ($h2$) your child has lung cancer, ($h3$) your child has the stomach flu. Colds and stomach flu in children are common; lung cancer is not. Thus prior probabilities favor $h1$ and $h3$. However, if your child has either a cold or lung cancer, it is very likely that she will cough; it is less likely that she will cough if she has the stomach flu. Thus the likelihood, $P(D|h)$, is higher for $h1$ and $h2$ than $h3$. Integrating both the prior and the likelihood suggests that you should decide that your child has a cold. From where do prior beliefs originate? In our example, they came from common cultural knowledge about the prevalence of childhood diseases. The transmission of such knowledge is critical to cultural learning in human cognition. However, constraints due to prior knowledge can come from many sources, including both the individual organism's own past experience and evolutionary adaptations (i.e., selective "experience" of one's reproductive lineage; Dall et al. 2005). In solving problems of visual perception, for instance, prior knowledge may include an evolved constraint to assume that illumination comes from a single, overhead source (i.e., the Sun). Bayesian inference, therefore, formalizes the claim that background knowledge is integrated with new data to affect judgment and decision making. This ability to integrate prior knowledge and statistical information supports rapid, accurate, inference across a range of problems and goals. Particularly relevant to decision making, formal analyses suggest that organisms whose decisions approximate the output of Bayesian inference models will out-compete organisms using other strategies when animals are contending for resources (Ramsey 1926/1980; de Finetti 1937/1980). Critically, however, Bayesian inference models provide computational-level accounts of cognition (Marr 1982): they describe optimal decision making given a particular set of constraints and data, but they do not prescribe the ways these computational outcomes are instantiated. Many different algorithms can approximate the output of Bayesian inference models, and these algorithms might, in turn, be implemented by many kinds of neural systems. Thus although evolution might favor organisms that approximate the output

of normative models, there might be substantial variability in the mechanisms underlying decision making across species.

Another major source of behavioral variability in many species comes from variation in “self-control” in different contexts. Nonhuman animals commonly prefer immediate rewards. This is paradoxical in the sense that they often prefer smaller, sooner benefits even when they could achieve a higher overall intake rate by choosing a more delayed but larger alternative. There is, of course, a large literature on this topic that goes under several headings: delay discounting, intertemporal choice, self-control, and failure to delay gratification. Superficially, at least, this so-called self-control problem resembles naturally occurring patch exploitation because when an animal chooses to stay longer in a food patch, it is “choosing” an alternative in which it takes longer and acquires more food. In the “self-control” literature, however, the time between trials typically does not affect preference, whereas in patch exploitation the so-called travel time is a consistently powerful variable: animals spend longer in patches when travel times are longer (Stephens and Krebs 1986). This problem has been studied by Stephens and colleagues by creating two types of choice situation (e.g., Stephens and Anderson 2001). One test situation involves a typical self-control situation: blue jays were offered a binary, mutually exclusive choice between a smaller-sooner and larger-later option. The second situation sought to create an economically similar situation that was more “patch like.” To achieve this, a situation is created in which the jays had to choose between leaving and staying. The “stay” option led to a “larger-later” option, whereas the “leave” option resulted in a smaller-sooner option. The key result is that jays in the self-control situation behaved in a typically impulsive way, and they performed relatively poorly; in contrast, jays in the patch situation achieved high rates of intake (nearly optimal). In short, in patch-like situations, jays achieve a high level of performance, whereas in the self-control situation they perform relatively poorly. The hypothesis put forth by Stephens et al. is that the choice rules used by the jays evolved to make decisions that are more like patch exploitation than binary mutually exclusive choice; thus the rules perform better in the situation that more closely resembles the problems they have faced in their evolutionary past. They remark that it does not necessarily follow that jays are less impulsive in the patch situation; instead, the patch situation could be a case where an “impulsive rule” performs well (for the mathematical reasoning to support this claim, see Stephens and Anderson 2001). This has been termed the ecological rationality hypothesis of impulsivity. Here, the phrase ecological rationality means, crudely speaking, that the jay’s decision rule “makes sense” in an ecological context. The phrase ecological rationality has come into use because it is associated with research on decision heuristics conducted by Gigerenzer and the ABC group (e.g., Todd et al. 2011b). However, some biologists (Kacelnik et al. 2006) have criticized Gigerenzer’s use of this term, because when applied to human heuristics it is used to describe a fit between a decision mechanism

and the current environment. From an evolutionary perspective, however, the environment that matters is the ancestral selective environment (sometimes called the adaptively relevant environment). To make this distinction, Kacelnik has suggested the term biological rationality to describe a rule that is rational in the adaptively relevant environment. Of course, if we are willing to assume stationarity (e.g., the statistical properties of patches have remained the same since the Pleistocene), then this distinction would be relatively unimportant. The distinction is clearly quite important for the evaluation of human decision heuristics because the environment of modern humans is often thought to be quite different from the adaptively relevant environment for humans (although recent molecular-genetic evidence suggests that humans have been subject to multiple bouts of substantial selection significantly more recently than the Pleistocene; e.g., Bustamante et al. 2005).

Patterning of behavior over time can often be selected for directly. Indeed, variability in behavior can be selected for when being predictable would facilitate exploitation by competitors and natural enemies. Such selection pressures can be particularly acute in antipredatory contexts, where exposure to certain types of predators (e.g., stalkers) can favor individuals that behave as randomly as possible while at risk (Bednekoff and Lima 1998).

Among Individuals

When thinking about individual differences in behavior from an evolutionarily adaptive perspective, we have to distinguish at least three questions (Wolf and Weissing 2010): First, what factors favor the evolution of variation in behavior, and/or the mechanisms underlying behavior, among individuals? Second, what factors favor the evolution of consistency in behavior within a functional context? And third, what factors favor the evolution of behavioral correlations across different functional contexts?

In terms of the first question, substantial effort in the evolutionary biology literature has been devoted to understanding the causes of adaptive variation. Evolutionary biologists typically recognize three major processes in this context:

1. Frequency-dependent selection: Competition for limiting resources, for example, can often select for mixtures of tactics to be expressed within populations, causing the increasingly common use of a tactic or strategy to render it less effective. Thus the fitness returns from expressing it depend negatively on its frequency within the population.
2. Spatiotemporal variation in the environment: Adaptive behavior can often vary within populations because different things are selected for at different times and in different places within the larger niche that a population occupies.

3. Nonequilibrium dynamics: The constant influx of variants associated with nonadaptive processes, such as random dispersal and mutation, can give rise to substantial standing variation in traits (including behaviors) within populations.

Concerning the second question (consistency within contexts), it is important to consider that, from an adaptive perspective, the optimal action that an animal should take typically depends on its current state (e.g., energy level, experience) and thus rarely remains constant from one moment to the next. When positive feedbacks between state and behavior occur, they can amplify stochastic variation among individuals and lock individuals into distinct regions of state space, causing individuals to differ consistently while such positive feedbacks persist (Dall et al. 2004). Indeed, learning can generate such positive feedback as initial differences in early experience can encourage individuals to develop different skills and behave differently (Tinker et al. 2009). Finally, recent work suggests that behaving consistently can be selected for directly in social situations when doing so can cause social partners to respond more favorably than they otherwise would. This can occur whenever there is coevolution between social responsiveness (social information use) and behavioral consistency, which may be the case during aggressive competition over resources (Dall et al. 2004; Wolf et al. 2011) and in some types of cooperative interactions (McNamara et al. 2009).

In terms of the third question (correlations across contexts), it is important to examine the factors that are likely to select for adaptive differences in underlying mechanisms that affect several traits at the same time: What should select for differences in metabolism or physiological commitments (e.g., neural tissue) to information processing? Wolf et al. (2007), for example, addressed the question of why, in many organisms (ranging from octopuses to chimpanzees) is boldness in novel and potentially dangerous situations often associated with aggressiveness in intraspecific interactions, whereas shyness is associated with less aggressive behavior. Based on the asset protection principle of life history theory (Clark 1994; Houston and McNamara 1988), they argue that individuals of high reproductive value (i.e., that have high future fitness expectations) should be risk-averse because they have much to lose, whereas individuals with low fitness expectations should be more risk-prone because they could hardly end up worse off than they already are. Since this basic principle applies to behavior in all kind of risky situations (Dall 2010), differences in future fitness expectations should give rise to correlated differences in all kinds of risk-related behaviors despite limited evidence for such domain-general attitudes to risk in humans (e.g., Weber et al. 2002). In the meantime, various empirical researchers have concluded that this principle can indeed explain a diversity of phenomena, ranging from the risk-prone “personalities” of wild guinea pigs that are born relatively late in the season and, hence, have lower fitness expectations (Groothuis and Trillmich 2011), to the risk-averse behavior of

oystercatchers breeding on high-quality territories (Goss-Custard 1996). More convincing than such indirect evidence is the outcome of a recent experiment on house mice, which was designed to test the a priori prediction that female house mice of genotype $t/+$ should be more risk-averse than their wild-type $+/+$ counterparts, while the opposite should be the case in males. Here, t refers to a certain gene locus (the t -complex) that does not directly affect the behavior of house mice (e.g., it is associated with sperm motility). The prediction was based on the observation that heterozygous $t/+$ females have higher expected fitness than homozygous $+/+$ females, while heterozygous males have lower fitness expectations than homozygous wild-type males. Auclair et al. (in preparation) performed several personality tests, with outcomes generally in line with the predictions: in comparison to wild-type individuals, heterozygote females are more risk-averse, while heterozygote males seem to be more risk-prone. Finally, Wolf et al. (2008) provided a theoretical explanation for the observation that in many populations some individuals readily react to changes in their environment, while other individuals exhibit more rigid, routine-like behavior that is much less affected by environmental cues. Moreover, Wolf et al. predicted that responsiveness and rigidity should be relatively stable over lifetime and consistent across contexts. Both features are consistent with various experiments, for example in ducks (Shettleworth 1998) and spice finches (Mottley and Giraldeau 2000), where environmental conditions were repeatedly changed by the experimenters.

Among Groups

As discussed above, group differences can be an important source of variation in decision making when there is selection for social information use and/or conformity. Social information use can be adaptive when evolutionary conflicts of interest are minimal within groups (e.g., where group members are highly related or when competition for resources is reduced), and an individual's uncertainty about how to allocate behavior can be significantly reduced by observing the behavior of group mates (Dall et al. 2005). Copying the behavior of social partners, or conformity, is a particular form of social information use; it is often the cheapest form of information use but carries substantial risks, stemming from the erroneous copying of inappropriate behaviors or "informational cascades" (Dall et al. 2005). Cultural learning (both within and across generations) facilitates the tracking of intermediate rates of environmental change ("red noise" environments) and will be often be selected for in species that are dependent on the populations of other species (e.g., predators), especially if they are marine (Whitehead 2007). This form of learning will maintain very substantial differences among groups, even without conformity, but conformity will exacerbate this type of variation.

Variation and a Theory of Darwinian Decision Making

We have discussed variation in decision making by individuals that have evolved by Darwinian selection from a range of perspectives: Does this really matter? What are the interesting scales of variation? What kind of mechanisms can generate it? How could it evolve? In the process, we hope to have highlighted the importance of explicitly considering variation both when attempting to predict economically and socially important patterns of behavior as well as to obtain a deeper understanding of the fundamental biological processes involved. Our approach so far has been rather piecemeal, with little attempt at any general understanding of the role that variation should play in an evolutionary account of decision making. Thus we finish by sketching the key elements that we feel should be included when incorporating variation into a general theory of Darwinian decision making:

- Game theoretic approaches. The underlying structure of the problems that evolved entities is essentially coevolutionary. This is because major components of the selective environments, to which individuals are subject, are biological in origin and will thus evolve. Therefore, specifying the outcomes of evolution naturally lends itself to applying evolutionary game theory. Moreover, variation is emerging as a key factor in determining outcomes predicted by evolutionary game theory. This has been illustrated a number of times in this chapter (e.g., how variation and socially acquired information use can coevolve to maintain individual variation in trust and trustworthiness; McNamara et al. 2009).
- Does one general, all-purpose mechanism exist? Or does it reflect diversity of forms and function? The specific set of problems faced by individuals is likely to be different for different species (set by their specific ecological circumstances or niche, which will include their social environments). Nevertheless, selection acts on existing variation so evolved systems are likely to share the basic components of a decision apparatus in proportion to their phylogenetic distance. Observable behavioral variation must therefore reflect the constraints on the system. Thus the evolutionary history of a lineage must be considered explicitly when incorporating variation in Darwinian accounts of decision making.
- Statistical decision theory (Bayesian inference). All evolved systems are contingent, yet such contingencies will represent the prior experience of ancestral and developmental environments (contexts) as a result of prior adaptation. For adaptive decision making, such genetically or developmentally induced priors will often need to be updated by more current experience. Specifying when or how such updates occur should be a major focus of an evolutionarily sensible theory of decision making, and statistical decision theory offers a formal framework that incorporates

this intuition very comfortably (Dall et al. 2005). Variation in behavioral allocation emerges naturally in Bayesian decision makers.

- The performance of alternative models of decision making should be assessed according to their fitness consequences, which may or may not include explicit analysis of potential persistence over evolutionary time. Because of this, such assessments must include explicit consideration of the fitness costs and benefits of the rules considered. Moreover, performance must be assessed using invasibility analysis, which means that the competitive environment must be specified, including the set of possible rules (as well as by assessing the evolutionary stability of rules). Thus standing variation, or the potential for variation, in decision making will be crucial when determining which rules are likely to evolve in particular systems.
- The statistical properties of the ecological (including social) context in which decision making evolved must be explicitly considered. What are the limits to potential specific models of decision making? (What is the ecological problem set?) While this is often difficult to specify in full detail for any given system, even identifying key statistical properties of typical ecological problems would help to identify the kinds of decision-making process that we should expect to see, and how they can vary in different contexts.
- For organisms subject to epigenetic development (i.e., develop from a totipotent cell), ontogeny will likely have a strong impact on individual variation (earlier environments will matter disproportionately). This suggests that attention should be focused on detailing such early influences and how persistent they are likely to be.
- The adaptive tinkering and contingency to which evolved systems are subject will tend to limit the dimensionality of control mechanisms (key state variables), which can account for imperfect optimization across contexts of higher dimensionality. The challenge will be to identify the key control mechanisms (e.g., sensory, cognitive processes) and how likely they are to vary both individually and phylogenetically.
- For humans (and some other animals), cultural evolution is likely to play an influential role in driving variation in decision making. This suggests that social contingency and/or historical influences are likely to be pervasive when attempting to predict specific outcomes. Furthermore, cultural variation will be bounded by genetic adaptation (e.g., many cultures have cleanliness-based practices that are likely to have evolved to mitigate the risk of infectious disease) and so studying the interplay between these factors will prove fruitful. Finally, basic psychological processes (e.g., salience) are likely to play a key role in determining how behavior varies over time and across cultural units (e.g., oddity effects, appeal to basic biological “drives,” etc.).

Conclusions

It is clear that variation is central to understanding how individuals should and do behave, from both biological and human-centered perspectives. Perhaps this should not be too surprising since evolution is at the heart of any scientific (i.e., natural) account of biological (including human) systems, and there is no evolution by Darwinian selection without variation. Moreover, how individuals allocate their behavior is key to determining the evolutionary success, and hence existence, of their specific (heritable) traits. This interplay between behavior, variation, and evolutionary outcomes was a major theme of this chapter. Indeed, we hope to have illustrated how this interplay is central to promoting models of human behavior that can successfully predict individual behavior and identify interventions that are going to be effective in determining specific outcomes. To this end, our recommendations for incorporating variation into a theory of Darwinian decision making will hopefully draw attention. To paraphrase Theodosius Dobzhansky: We feel that no behavior, human or otherwise, makes sense except in the light of variation and, therefore, evolution!