1 Introduction

Whether signals are reliable or deceptive has been a central question in the study of animal communication in recent years. The crux of the issue is whether animal signals are honest, in the sense of conveying reliable information from signaler to receiver, or deceitful, in the sense of conveying unreliable information, the falsity of which somehow benefits the signaler. This issue arises in a variety of contexts. When a male courts a female, do his signals honestly convey his quality relative to other males? Or does he exaggerate his quality in order to win over females that would otherwise choose some other male? When one animal signals aggressively in a contest over a resource, does the signaler honestly convey its likelihood of attack? Or does the signaler exaggerate that likelihood in order to intimidate competitors that would otherwise defeat him? The question of reliability versus deceit arises even in interactions that, on the face of things, seem to be predominantly cooperative. When an offspring begs for food from its parents, does it honestly convey its level of need? Or does the offspring exaggerate its need in order to get more food than the parents would otherwise provide?

The issue of reliability and deceit in animal communication resonates with human observers for a variety of reasons. One is that the occurrence of deceit is fraught with moral implications. In the view of many, human communication is permeated with deceit. Do humans stand apart in this regard, or are other animals as bad or worse? The answer might have considerable effect on how we view ourselves, as well as on how we view other animals. A second reason for interest in this issue is that the occurrence of deceit, if deceit is defined appropriately, can have considerable implications for our understanding of animal cognition. Some definitions of deceit are framed so as to require cognitive processes of considerable sophistication, such as the ability to form intentions and beliefs and to attribute beliefs to other individuals. If we employ such a definition, and if we can then determine that nonhuman animals deceive each other according to this definition (a big "if"), then we have provided support for a greater level of cognitive capacity than many earlier views of animal behavior have allowed.

Our own interest in reliability and deceit revolves around neither morality nor cognition, but instead derives from the evolutionary implications of the issue. The way one expects animal communication systems to function in terms of reliability and deceit depends on how one views the operation of natural selection. Early students of animal behavior often assumed implicitly that selection operates at the level of groups, so that behavior evolves toward

what is best for the population or species as a whole, leading to the view that animal communication consists primarily of the cooperative exchange of reliable information. The predominant view nowadays, however, is that selection acts largely at the level of the individual, so that behavior evolves toward what is best for the individual performing the behavior, and not toward what is best for the group. If behavior is commonly selfish, in this sense, then it is not always obvious why animals should exchange information cooperatively. Instead, one might expect many instances in which signalers would attempt to profit individually by conveying dishonest information. But because individual selection works on the receiver as well as the signaler, receivers ought to respond to signals only if doing so is to their advantage, on average. Therefore, if dishonesty is common, it also is not obvious why receivers should respond to signals.

Taking the argument one step further, if receivers fail to respond to signals, it is not obvious how signaling systems can exist at all. Thus if one accepts the view that selection acts predominantly at the level of the individual, as we do, and if one at the same time accepts the idea that animals do communicate with each other, as seems obvious, then one is left with a series of evolutionary puzzles. Are animal signals in reality reliable or unreliable? If animal signals are reliable, what mechanisms maintain reliability despite the tempting advantages of dishonesty? If animal signals are deceitful, do receivers respond to them anyway, and, if so, why? Our principal purpose in this book is to work through possible answers to evolutionary puzzles such as these.

Definitions

Before we get to these puzzles, we need to define some terms. First, we need to define what we mean by "signal," in order to delimit the set of traits whose honesty and dishonesty we will examine. In one of the first rigorous evolutionary analyses of communication, Otte (1974, p. 385) defined "signals" as "behavioral, physiological, or morphological characteristics fashioned or maintained by natural selection because they convey information to other organisms." Otte explicitly rejected group-selectionist explanations for the evolution of traits, so in his view the transmission of information had to confer some reasonable advantage on the signaler itself in order to satisfy the definition. Thus Otte excluded as signals those traits that convey information to predators or parasites without any benefit to their possessors; he cited the chemicals in human sweat that attract disease-carrying mosquitoes as a possible example. Otte also rejected as signals those traits, such as body size, that may be used by other individuals of the species to assess their possessors but did not evolve for that function. Clearly included under Otte's definition would be vocalizations, color patterns, and body movements that have evolved be-

cause they transmit information in a way that benefits the individual that exhibits those traits. More ambiguous are traits, such as the form of a bird's tail, that originally evolved for some other function but have been modified by selection for information transmittal. We will regard such traits, or more precisely their modified properties, as signals; thus the bird's tail itself is not a signal but the tail's length is, if that length has been exaggerated beyond its aerodynamic optimum in order to influence receivers.

This brings us to our definitions of reliability and deceit. In everyday English, "reliable" means that "in which reliance or confidence may be put; trustworthy, safe, sure" (Little et al. 1964). An animal signal, then, would be reliable if one could have confidence in its veracity, or truthfulness—if, that is, one could trust the signal to convey whatever it is supposed to convey. The difficulty with this formulation is in ascertaining what the signal is "supposed to" convey. "Supposed to" in this context must be interpreted from the viewpoint of the receiver rather than the signaler; what matters is whether the signal conveys something that the receiver would benefit from knowing. If we are certain what it is that the receiver benefits from knowing, such as some attribute of the signaler or its environment, then we can ascertain the reliability of the signal by measuring the correlation between the signal and the attribute of interest.

Suppose, for example, that we think that female frogs are interested in the size of conspecific males, and we find that calls communicate information on male size by a negative correlation between call frequency and caller size (males with deeper croaks are larger). We can then determine the reliability of this information by measuring the correlation between call frequency and caller size. The trouble is that we can never really be certain that caller size is what the females "want" or "need" to know. Even if we can show that call frequency is well correlated with caller size, and that the females show a behavioral preference for calls of lower frequency, we cannot be sure that their true interests are not in some other characteristic—perhaps, in this example, male age. The best we can do is to measure as carefully as we can the benefits that the receivers obtain from different types of information. If we can show that female frogs benefit from mating with larger males but not from mating with older ones, we at least can have some confidence that size is what matters to the receivers, and then evaluate reliability of call frequency in terms of its correlation with signaler size.

To formalize this definition, we suggest that an animal signal is reliable if:

1. Some characteristic of the signal (including, perhaps, its presence/absence) is consistently correlated with some attribute of the signaler or its environment; and
2. Receivers benefit from having information about this attribute.

A remaining problem is how to specify what we mean by "consistently correlated." We can never expect a perfect correlation between signal characteristic

and the attribute being signaled. Even if signalers are striving for perfect honesty, errors must be expected in the production of the signal and in our measurements of it, either of which would prevent our observing perfect reliability. How good, then, does the correlation have to be for us to conclude that the signal is on the whole reliable? One answer is provided by the concept of "honest on average" (Johnstone and Grafen 1993, Kokko 1997). A signal can be considered honest on average if it contains enough information, sufficiently often, that the receiver on average is better off assessing the signal than ignoring it. Consider again the example of male frogs communicating their size to females via the frequency of their call. The correlation between male size and call frequency can never be expected to be perfect, and in reality is often rather low (see chapter 4). The male's call can be considered honest on average if the correlation between male size and call frequency is good enough that the female benefits on average from using the call to assess male size, instead of ignoring this signal feature. In practice, it will be difficult to determine whether this criterion is being met, but at least it provides a theoretical standard against which reliability can be judged.

A simple way to define "deceptive" would be as the opposite of reliable, but for many the concept of deception carries more baggage, and consequently requires a more complex definition. A relatively simple definition of deception is provided by Mitchell (1986, p. 20), who suggested that deception occurs when:

1. A receiver registers something Y from a signaler;
2. The receiver responds in a way that is appropriate if Y means X; and
3. It is not true here that X is the case.

Note that the definition requires specifying what the signal (Y) means to the receiver. The meaning of Y to the receiver is judged by the response of the receiver to Y together with an observed correlation between Y and X, across many such signals. In other words, we infer that Y means X to the receiver because signalers usually produce Y in association with X, and because the receiver responds to Y in a way that is appropriate if X is true. To make this more concrete, let Y be an alarm call given by the signaler. The alarm call is usually produced when a predator (X) is present, and the receiver typically responds to the alarm call by fleeing, an appropriate (i.e., beneficial) response if a predator is indeed nearby. Deception occurs if the signaler produces the alarm and the receiver reacts by fleeing when in fact no predator is present.

A difficulty with Mitchell's (1986) definition, which he himself points out, is that deception so defined cannot be distinguished from error on the part of the signaler. If the signaler has produced an alarm in error, would we want to call such an action deceptive? This problem can be solved if the definition of deception further stipulates that the signaler benefits from the receiver's response to the signal. Mitchell (1986) himself is uncomfortable with the notion

of benefit, remarking that the "idea of benefit is taken from human affairs" and when applied to nonhuman animals typically refers to what a human observer "believes is good for them." For an evolutionary biologist, however, "benefit" has a straightforward meaning—an individual benefits from an action if that action increases the individual's fitness, in the sense of the representation of the individual's genes in subsequent generations. Benefit in this sense is not an anthropocentric idea, but one that applies equally well to all organisms. With the added stipulation about a benefit to the signaler, we will define deception as occurring when:

1. A receiver registers something Y from a signaler;
2. The receiver responds in a way that
   a. benefits the signaler and
   b. is appropriate if Y means X; and
3. It is not true here that X is the case.

Deception defined in this way has sometimes been termed "functional deception" (Hauser 1996), meaning that the behavior has the effects of deception without necessarily having the cognitive underpinnings that we would require of deception in humans.

Other definitions specify that deception must have more complex cognitive underpinnings, that is, that the signaler has an "intention" to cause the receiver to form a false "belief" about the true situation (Russow 1986, Miller and Stiff 1993). Deception defined in this way has been termed "intentional deception" (Hauser 1996). "Intentions" and "beliefs" are mental states, and as such are difficult to measure in nonhuman animals, to say the least. Whether animals possess such mental states, and whether they can ascribe them to others, is of great interest to philosophers (Dennett 1988) and cognitive ethologists (Cheney and Seyfarth 1990, Seyfarth and Cheney 2003, Byrne and Whiten 1992), as well as to the general public. A major goal of some researchers studying deception in nonhuman animals is to use this type of interaction as a window onto the mental states of those animals, in an effort to determine whether they do indeed form intentions, beliefs, and so forth. Although we applaud such efforts, we repeat that our own interests lie elsewhere, in the analysis of reliability and deceit from a functional, evolutionary viewpoint. Another way of saying this is that we are interested in how natural selection shapes animal communication to be either honest or dishonest. From this viewpoint, the question of mental states is largely irrelevant; the costs and benefits to the signaler of giving a false alarm, and to the receiver of responding, ought to be the same whether or not the signaler is able to form an intention and the receiver to form a belief.

Another issue in defining deception is whether to include the withholding of signals. Some authors have argued in favor of this inclusion, suggesting that under certain circumstances, a failure to signal can be considered just as

deceptive as producing a dishonest signal (Cheney and Seyfarth 1990, Hauser and Marler 1993a, Hauser 1996). Hauser (1996), for example, states that if an animal fails to produce a signal in a certain context in which that signal is typically produced, and if the animal benefits from failing to signal, that failure constitutes functional deception. This idea seems to us to have little application to a large majority of signaling contexts, such as those involving aggression or mate choice, in which cooperation is not expected from the interactants. In practice, the idea that withholding information is deceptive has most often been applied to cooperative interactions, most notably to interactions in which an animal signals the discovery of a food source to others of the same species (Hauser and Marler 1993a,b). Even here, the concept seems to us to be problematic. Say, for example, that a signaler follows the convention of calling when it finds a large amount of food, more than it can eat itself, and not calling when it finds a smaller amount. The signal then is consistently correlated with an aspect of the environment that receivers benefit from knowing, and so meets our criteria for reliability. Of course the receivers would be even better served by knowing more (i.e., from hearing about the small amount of food as well), but the signaler has not broken its convention in denying them this information.

Before we move on, let us reiterate in less formal terms the definitions of reliability and deceit we plan to use. Reliability requires that there be a correlation between some characteristic of the signal and some attribute of the signaler or its environment that the receiver benefits from knowing about, and that the correlation be good enough that the receiver on average benefits from assessing the signal rather than ignoring it. Deceit requires not only that the correlation between signal characteristic and external attribute be broken at times, but that the signaler benefits from this breakdown. Therefore, if a breakdown occurs in the correlation between signal characteristic and external attribute from which the signaler does not benefit, this would constitute unreliability but not deceit. A breakdown of this type we would describe as "error."

Some History

Opinions about the prevalence of reliability and deceit in animal communication have swung back and forth in recent decades. A convenient place to enter this history is with a seminal paper published by Richard Dawkins and John Krebs in 1978 titled "Animal signals: Information or manipulation?" In writing this paper, Dawkins and Krebs were reacting to what they labeled as the "classical ethological" view of animal communication, which in their opinion treated communication as a cooperative interaction between signaler and receiver. The ethological view assumed that receivers (reactors) were "selected to behave as if predicting the future behaviour" of signalers, while the signalers were "selected to 'inform' reactors of their internal state, to make it easy for reactors to

predict their behaviour" (Dawkins and Krebs 1978, p. 289). Thus the classical ethological view held that "it is to the advantage of both parties that signals should be efficient, unambiguous and informative" (Dawkins and Krebs 1978, p. 289). Dawkins and Krebs objected to this Panglossian picture of communication on the grounds that it is not what one would expect to evolve under natural selection. Natural selection favors behavior that enhances the actor's own survival and reproduction, rather than anyone else's, so that "cooperation, if it occurs, should be regarded as something surprising, demanding special explanation, rather than as something automatically to be expected" (Dawkins and Krebs 1978, p. 289). Dawkins and Krebs proposed replacing the cooperative view of communication with one that interprets signaling as an attempt on the part of a signaler to manipulate the behavior of the receiver to the signaler's advantage. Under this alternative, the signaler communicates not in order to tell the receiver what the receiver wants to know, but to induce the receiver to do something that will benefit the signaler. "If information is shared at all it is likely to be false information, but it is probably better to abandon the concept of information altogether" (Dawkins and Krebs 1978, p. 309).

The manipulative interpretation of communication proposed by Dawkins and Krebs reflected the growing consensus among animal behaviorists that individual selection, rather than group selection, plays the preeminent role in shaping the evolution of behavior. Group selection is selection stemming from the births and deaths of groups (such as populations and species) and favoring traits that benefit groups, whereas individual selection is selection stemming from the births and deaths of individuals and favoring traits that benefit individuals. The consensus in favor of individual selection arose in large part in reaction to Wynne-Edwards' (1962) overtly group-selectionist ideas, which brought the distinction between group and individual selection into focus. Group selection as articulated by Wynne-Edwards was sharply criticized, and individual selection championed, by influential evolutionary biologists such as Hamilton (1963), Lack (1966), Williams (1966), and Maynard Smith (1976a). Although argument over these issues has not entirely died away, from the 1970's onward researchers investigating animal behavior have interpreted their results almost exclusively in terms of individual selection.

In describing the "classical ethological" view of communication, Dawkins and Krebs (1978) gave a series of quotations from earlier papers, some of which had a decidedly group-selectionist ring. Tinbergen (1964, p. 206), for example, was quoted as saying "One party—the actor—emits a signal, to which the other party—the reactor—responds in such a way that the welfare of the species is promoted." Ethologists objected that the quotations chosen by Dawkins and Krebs did not correctly represent the central ideas of ethology with respect to communication. Hinde (1981, p. 535), for example, claimed that Dawkins and Krebs had erected a "straw man" that "neither accurately nor adequately conveys the main stream of ethological studies." Hinde (1981)

argued that the ethologists, rather than assuming that "signals carry precise information of what the actor will do next," had actually emphasized the use of signals in "conflict" situations, where the animal was torn between "incompatible tendencies" such as attack and retreat. An animal caught in a conflict situation in this way cannot itself predict what it will do next, let alone inform others. And in fact, if one reads the Tinbergen article cited above (Tinbergen 1964), one finds little discussion of either the information transmitted by signals or the selective benefits of signaling to signaler or receiver. Instead, Tinbergen's principal interests were in the evolutionary origin of displays, in the sense of the movements from which signals were originally derived, and in the proximate causation of display, in the sense of what "motivates" the animal to signal. In discussing motivation, Tinbergen (1964) indeed emphasized the "conflict hypothesis," but although this hypothesis may imply that displays have low information content, Tinbergen himself did not draw this inference. As for the level of selection question, Hinde (1981) claimed that the forthrightly group-selectionist statement quoted by Dawkins and Krebs was "not representative" of Tinbergen's writings in general. Hinde (1981) has a point here, in that Tinbergen at times discussed the evolution of behavior in terms of individual as well as group advantage (e.g., Tinbergen 1951), and the same can be said of other "classical" ethologists as well. In truth, Tinbergen's interpretation of the group selection/individual selection distinction was rather different from a contemporary one; for example, he tended to attribute the evolution of any behavior that furthered the reproduction of individuals (rather than their survival) to group advantage rather than individual advantage (Tinbergen 1951).

Whether or not the cooperative information-exchange view of animal communication truly represented the main trend of ethological thinking, this viewpoint was thoroughly discredited by Dawkins and Krebs' analysis. The manipulative view of communication that Dawkins and Krebs suggested in its place had its own problems, however. According to this view, the signaler communicates in order to maneuver the receiver into performing some action that will benefit the signaler, and if the signal can be said to convey any information, that information is at least as likely to be false as true. The critical flaw with this reasoning is that it does not explain why the receiver would be selected to respond to the signal at all. If there is, on average, no information of benefit to the receiver in a signal, then receivers should evolve to ignore that signal. If receivers ignore the signal, then signaling no longer has any benefit to the signaler, and the whole communication system should disappear. In short, it was fairly easy to construct an individual-selectionist argument showing why signaling should not be honest, but when this argument is followed to its logical conclusion, it is not obvious why signaling would occur at all.

A partial solution to this dilemma had already been proposed by Zahavi (1975), with reference to signals used in mate choice. Mating signals provide

an excellent example of the honest-signaling dilemma. Females (if they are the ones exercising choice) will benefit from choosing males of superior quality. If superior males can give a signal that identifies them as being superior, then they will benefit from the ability of females to identify and choose them. Poor-quality males, however, would also benefit from being chosen, so they will also be selected to give the signal. If all males, regardless of their quality, give the signal, then the signal contains no information on male quality. Females then should be selected to ignore the signal, and males should cease to give it. The solution proposed by Zahavi was that a mating signal must confer a "handicap" on the survival of the signaler. In Zahavi's words, the handicap serves "as a kind of test imposed on the individual." A male with a highly developed handicap "is an individual which has survived a test" and therefore has demonstrated he is of superior quality. As possible examples of such handicap traits, Zahavi (1975) cited the exaggerated train of the peacock and singing in exposed positions by warblers, both of which he thought would expose a male to predators. "Since good quality birds can take larger risks it is not surprising that sexual displays in many cases evolved to proclaim quality by showing the amount of risk the bird can take and still survive" (Zahavi 1975, p. 211).

Zahavi's handicap idea initially met with skepticism. Maynard Smith (1976b), Bell (1978), and others formulated genetic models to analyze whether a handicap trait would evolve under Zahavi's assumptions. Typically, these models assumed that single genes controlled both the handicap trait in males and the preference for the handicap in females. A third gene controlled viability in both sexes, where viability was defined as fitness exclusive of mating success (Maynard Smith 1985). Maynard Smith (1985) reviewed the results of these modeling efforts, and concluded that with realistic parameter values neither the handicap trait nor the female preference for it would increase from an initially low frequency. A simple quantitative-genetics model by Pomiankowski and Iwasa (1998) recently reinforced these conclusions. The problem is that females benefit from preferring the handicap only to the extent that the preference gene comes to covary with viability, and this covariance arises very indirectly, from the association of the preference with the handicap and the association of the handicap with viability. The resulting weak benefit is counterbalanced by the cost to females of having their sons inherit the handicap, along with its deleterious effect on survival. This cost is sufficient to prevent both the handicap trait from increasing among males and the preference for the handicap from increasing among females (Maynard Smith 1985, Pomiankowski and Iwasa 1998).

The negative conclusions of these genetic models did not lead to the demise of the handicap idea; instead, handicaps were rescued by some new ideas about how the link between a handicap and viability could come about. The original models by Maynard Smith (1976b) and others assumed that the handicap becomes associated with viability because only males of high viability can sur-

vive with the handicap. A handicap linked to viability in this way is sometimes referred to as a "pure epistasis handicap" (Maynard Smith 1985), but we prefer the term "Zahavi handicap" (Pomiankowski and Iwasa 1998), both because this label is less cumbersome and because the way this handicap is defined accords well with our reading of Zahavi's original formulation (Zahavi 1975). Following the criticism of his original idea, Zahavi (1977a, p. 603) suggested a new kind of handicap, in which "the phenotypic manifestation of the handicap is adjusted to correlate to the phenotypic quality of the individual." Such an adjustment could be made in a couple of different ways. One is for the handicap trait to be expressed only if a male has both the gene for the handicap trait and the gene or genes for high viability; a handicap of this sort has been termed a "conditional" or "condition-dependent" handicap (West-Eberhard 1979, Andersson 1994). A second way is for the handicap to be expressed in all males that have the handicap gene, but with its size or conspicuousness made to correlate with the viability of its possessor; this has been termed a "revealing" handicap (Maynard Smith 1985, 1991b).

Analysis of genetic models indicated that condition-dependent and revealing handicaps, together with female preferences for them, are much more likely to evolve than are Zahavi handicaps (Andersson 1986, Iwasa et al. 1991). One reason is that whereas all males with the handicap genes pay the full cost of a Zahavi handicap, only a subset of males—those with genes for high viability—pay the full cost of condition-dependent and revealing handicaps (Andersson 1994). A second reason is that the link between handicap and viability is more direct for condition-dependent and revealing handicaps, making them better, more informative signals of viability, and increasing the average benefit of choosing a male with such a handicap (Pomiankowski and Iwasa 1998). Models of condition-dependent and revealing handicaps thus indicated that, given certain assumptions, reliable signals of mate quality might evolve in the context of male courtship of females.

Meanwhile, the reliability of signals given in aggressive contexts was also being debated. In a conflict between two animals, two kinds of information might well be valuable to either contestant: the other animal's "quality," in the sense of its fighting ability or "resource holding potential" (RHP), and its "intentions," in the sense of whether it was likely to attack or to retreat. The near-universal occurrence of display in animal conflicts implied that either or both types of information were commonly being conveyed, but why this was the case was not obvious. In a series of influential papers published in the 1970's, Maynard Smith and colleagues used game theory to analyze this problem (Maynard Smith and Price 1973, Maynard Smith 1974, 1979; Maynard Smith and Parker 1976). We will look at the details of some of these models later (see chapter 4), but the gist of the argument can be presented without mathematics. Suppose that aggressive contests occur between individuals that are well matched for fighting ability, in which case the winner is likely to be

the one willing to escalate to a higher level of aggression. If the population had a set of signals that reliably communicated the signaler's aggressive intentions, then it would pay one animal to give way if the other signaled a higher level of aggression than its own. Such a population could be invaded, however, by a cheating strategy whereby the cheater signaled the highest level of aggressiveness, regardless of its true intentions. Cheaters would win many contests without having to fight, and the cheater strategy would increase in frequency. Once cheating became sufficiently common, however, we would expect receivers to evolve to ignore the signal (Maynard Smith and Parker 1976, Maynard Smith 1979). A similar argument could be given for why signals of fighting ability should be dishonest, at least for signals that are not directly constrained by the signaler's phenotype. These theoretical arguments against reliability were widely accepted, and researchers shifted their attention to demonstrating empirically that aggressive displays, on the one hand, did not predict aggression and, on the other, were largely ignored by receivers (Caryl 1979).

Of course if aggressive signals are largely useless in terms of their information content and consequently are ignored by receivers, then it makes little sense for signalers to go on signaling—and yet in practice signaling is nowhere more common than in aggressive conflicts. Here, too, Zahavi's handicap idea, that signals can be honest if they are costly, provided a possible way out of the paradox. Enquist (1985) used game theory to show that reliable signaling of either aggressive intentions or fighting ability could be evolutionarily stable, provided that signaling was costly and that either the cost of signaling or the benefit of winning varied between individuals (see chapter 4). Enquist (1985) acknowledged the importance of Zahavi's (1975, 1977a,b) work in directing attention to the role of signal cost in maintaining signal reliability.

The work of Enquist (1985) on aggressive signals and of Andersson (1986), Pomiankowski (1987), and others on mating signals initiated a swing in opinion back to the expectation that animal signals generally are reliable. This swing was completed by the publication in 1990 of two papers by Alan Grafen. In the first of these, Grafen (1990a) presented a population-genetics model of the evolution of female choice for a male handicap trait. The model assumes that male advertising is costly and that the cost is higher for low-quality than for high-quality males. Given these assumptions, Grafen (1990a) showed that an evolutionary equilibrium exists at which the level of male signaling is a strictly increasing function of male quality—meaning that the signal is reliable—and at which females prefer males with higher levels of signaling—meaning that receivers respond to the trait. In the second paper, Grafen (1990b) presented a game-theory model of honest signaling, which he extended to aggressive signaling as well as mate choice. To obtain an evolutionarily stable strategy, Grafen (1990b) had to assume: (1) that signaling is costly, in the sense that signaler fitness declines as the level of signaling increases; (2) that the receiver's assessment of the signaler's quality increases as the signaler's sig-
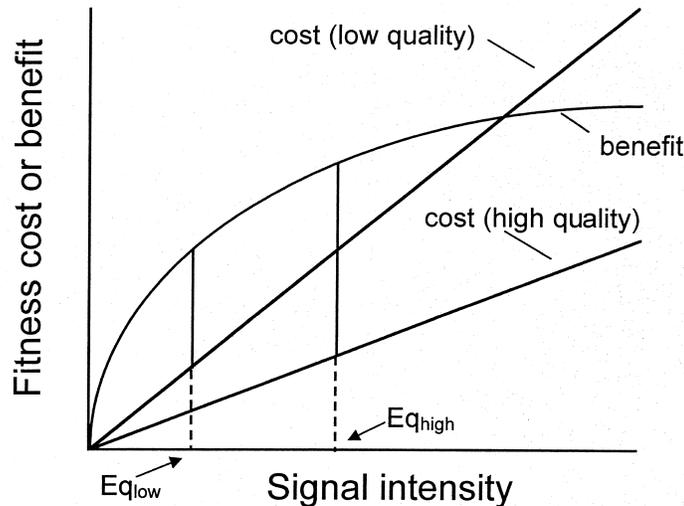
FIGURE 1.1. Johnstone's (1997) first graphical signaling model, in which the signal conveys signaler quality. Two cost lines are drawn, one for a signaler of high quality and a second for a signaler of low quality. The relationship between signal benefit and signal intensity is assumed to be the same for all signalers. The equilibrium signaling level is found as the signal intensity at which the difference between benefit and cost is greatest. The equilibrium for the high-quality signaler (Eq_{high}) is greater than the equilibrium for the low-quality signaler (Eq_{low}).

naling level increases; and (3) that the signaler benefits from being given a higher assessment. In addition, it was necessary to assume (4) that the ratio of the marginal cost of signaling (taken as a positive term) to the marginal benefit of a higher level of assessment is a decreasing function of signaler quality. The latter condition is satisfied if the benefits of improved assessment are the same regardless of quality and the cost of signaling is greater for signalers of poor quality than for those of high quality. Grafen (1990b) also turned the logic around, to argue that the existence of stable signaling systems implies that signals are reliable and costly in a way that meets the above conditions.

We review Grafen's models, as well as other signaling models, in more detail later in the book. For now, we will use graphical models developed by Rufus Johnstone (1997) to aid understanding of the honest-signaling argument. Figure 1.1 shows a version of Johnstone's model that is appropriate for the kind of situation envisioned by Grafen, in which the signal conveys the quality of the signaler, and the signaler benefits from receiving a higher assessment. This benefit increases monotonically with increasing signal intensity, that is, the higher the intensity of the signal, the more effective is the signal in terms of receiver response and the higher is the benefit to the signaler. This version

of the model assumes that the benefit of a given signal intensity is the same for all receivers regardless of their intrinsic quality. Signal costs also increase monotonically with signal intensity, but the costs rise more rapidly for a signaler of poor quality than for a signaler of high quality. The optimal signaling level for any signaler occurs at the signal intensity where the difference between benefit and cost is maximized. In order to generate a simple solution, costs are assumed to increase linearly with increasing signal intensity, while benefits increase to an asymptote. Under these assumptions, the optimal signaling intensity is higher for a signaler of high quality than for a signaler of low quality. Thus by assuming that each individual is following its own evolutionary interests, the model generates a signaling system that is reliable, in the sense that levels of signal intensity accurately convey levels of signaler quality.

In this first version of the model, Grafen's fourth assumption concerning the ratio of marginal cost to marginal benefit is satisfied in what has come to be viewed as the standard way—by making the costs of signaling dependent on signaler quality and the benefits independent of quality. Johnstone (1997) uses a second version of the graphical model to show that reliable signaling will occur if these assumptions are reversed. This second version of the model is appropriate for situations in which the signal conveys level of need rather than level of quality; it might be applied, for example, to the case of nestling birds begging for food from their parents. In figure 1.2 the relationship between signal cost and signal intensity is assumed to be the same for all signalers, whereas the benefit of signaling rises more rapidly for a signaler of high need than for a signaler of low need. Again, optimal signaling levels are found where the difference between benefit and cost is maximized for a given signaler, and again the result is reliable signaling, in this case because the signaler with the higher level of need signals at a higher intensity than does the signaler with the lower level of need.

Categories of Signal Costs

Grafen's and Enquist's models convinced many researchers that signal costs can be important in stabilizing signaling systems. Attention then turned to determining whether signals actually do have costs, and how those costs might come about. Many categories of costs have been described, but it is important to note that all must be reducible to a fitness cost if they are to be effective in enforcing signal reliability. This does not necessarily mean that a signaler's fitness is lower the higher its level of signaling, because (among other considerations) it might receive fitness benefits from the responses that receivers make to its signals. A precise definition of cost was provided by Grafen (1990b), who considered a signal to be costly if the partial derivative of fitness with respect to signaling level was negative, holding receiver assessment and sig-
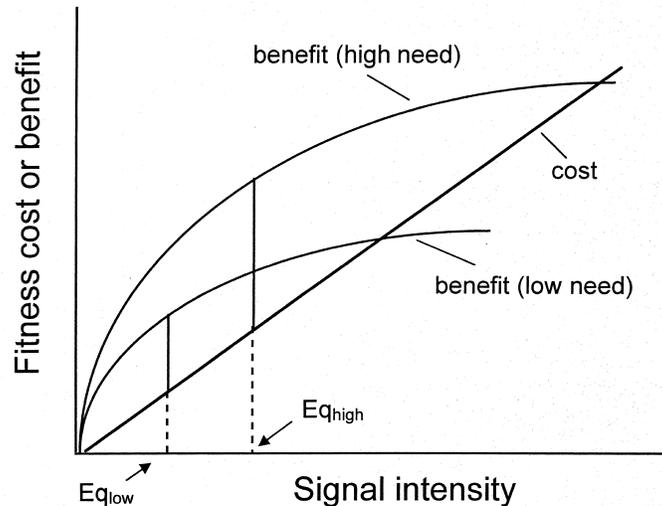
FIGURE 1.2. A second version of Johnstone's (1997) graphical model, in which the signal conveys signaler need. Two benefit curves are drawn, such that the signaler with higher need receives a greater benefit than the signaler with lower need, at any given signaling level. The relationship between signal cost and signal intensity is assumed to be the same for all signalers. The equilibrium signaling levels (compare figure 1.1) are again found as the signal intensities at which the difference between signal benefit and signal cost is greatest. The signals are predicted to be reliable, in the sense that the signaler with the higher need signals at higher intensity than does the signaler with lower need.

naler quality constant. In other words, the concept of signal costs requires that signaler fitness go down as signaling goes up when any fitness effects of receiver response and of signaler quality are held constant.

In categorizing costs, a primary division can be made between so-called "receiver-dependent costs" and "receiver-independent costs"—costs that stem from some response of receivers to a signal versus costs that are imposed regardless of whether or how receivers respond (Guilford and Dawkins 1995, Vehrencamp 2000). Vehrencamp (2000) suggests further dividing the receiver-dependent category into "vulnerability costs" and "receiver retaliation rules." A vulnerability cost occurs because the action of producing the signal opens the signaler to an increased chance of injury, if a receiver chooses to attack. Zahavi (1987) gave putative examples of signals in this category, such as postures and vocalizations that require relaxation and are given as aggressive displays. Zahavi argued that relaxation in the proximity of an aggressive rival was dangerous, and that the risk would be less for an individual of strong fighting ability than for one of poor fighting ability. The basis of a receiver-retaliation rule is that receivers are more likely to attack, or punish in some

other way, those signalers that give one kind of signal rather than another. Enquist (1985) used a receiver-retaliation rule as the cost of an aggressive signal in one of his original signaling models: one of two signals is more effective than the other, in terms of helping the signaler to win the contest, but is also more likely to provoke the opponent to fight. Receiver-retaliation rules are most likely to apply to aggressive signals, but they can apply to mating signals, if rivals of the same sex are more likely to attack a signaler if it gives a more effective mating signal than if it gives a less effective one.

We consider receiver-independent costs to include three categories: production costs, developmental costs, and maintenance costs. Production costs are costs that are paid at the time the signal is exhibited to the receiver. Included in this category would be the considerable energy consumed by calling in frogs or roaring in red deer, the time taken away from foraging and other activities by singing in a songbird, and the increased risk of predation a male stickleback may experience when it exposes its red coloration. Developmental costs are costs paid at the time a signal develops, well before the signal is displayed. The concept of developmental costs is usually applied to display structures whose growth requires considerable investment, such as the antlers of deer (Andersson 1986). We have argued that certain display behaviors also have developmental costs, especially complex behaviors that are supported by specialized neural systems (Nowicki et al. 1998, 2002a). Maintenance costs are ones that are a consequence of having to bear a display structure once it has been developed, and which are paid regardless of whether the display is actually given. A prime example of this category is the cost paid by birds for an elongated tail. Elongation of tail feathers beyond a certain point makes flight more clumsy and more expensive (Evans and Thomas 1992), so a male with an elongated tail is likely to expend more energy in flying, to be more vulnerable to predation, and to have decreased foraging success—all of which are detrimental to fitness.

Some displays must have multiple costs. The extravagant train of the peacock provides a familiar example. Growing these greatly elongated feathers must require a considerable investment in energy and nutrients, a clear instance of developmental costs. Fanning the train to display it must require some energy expenditure, and may also expose the signaler to an increased risk of predation, both of which are production costs. Displaying the train might have some receiver-dependent costs, if a large display tends to elicit attacks from rival males. And as the peacock's train must have an aerodynamic impact, maintenance costs due to decreased flight performance certainly apply. Measuring the summed effects of all these different costs, with their different units (energy versus risk) and timing, would be very difficult indeed.

In the real world, signals often attenuate and/or degrade while propagating between signaler and receiver, making it difficult for receivers to discriminate signals from irrelevant energy (Wiley 1994). Given this problem, and assuming

that a response in the absence of a signal is costly, it follows that receivers may be selected to set a high threshold for response in order to avoid "false alarms" (Wiley 1994, Johnstone 1998). If receiver thresholds are high, signalers will be selected to produce intense and hence costly signals in order to ensure detection (Johnstone 1998). By this argument, signals may be costly for reasons of "efficacy" rather than reliability. To support the handicap principle, one needs to show that a signal has a "strategic cost" over and above its "efficacy cost" (Maynard Smith and Harper 2003); however, in real-world signaling systems the boundary between strategic and efficacy costs may be difficult or impossible to delineate.

Alternative Explanations for Reliability

The handicap principle—the idea that signals can be reliable if they are costly in an appropriate way—is not the only viable explanation for honesty in animal signals. We have already introduced a second explanation, that embodied in Johnstone's (1997) second graphical mode (figure 1.2). In that model, it is the relationship between signal intensity and benefit that chiefly acts to make signaling honest, rather than the relationship between signal intensity and cost. If the benefit of signaling is sufficiently different for various categories of signalers, then optimal signaling levels will be quite different for different signalers, even if signaling costs are minimal.

A third explanation for reliability is the lack of a conflict of interest: if signaler and receiver agree on the rank order of possible outcomes of their interaction, then signals can be reliable without being costly. We introduce this idea at greater length in chapter 2. A fourth explanation, based on a model by Silk et al. (2000), is that deceit can be disadvantageous if receivers remember acts of deception by particular signalers and discriminate against signals from those individuals in the future. We call this mechanism "individually directed skepticism." Again, this hypothesis is explained more fully in chapter 2.

A fifth explanation for reliability is that some signals are constrained to be honest because of the mechanisms by which they are produced. Maynard Smith and Harper (2003) term a signal that is constrained to reliability in this way an "index," which they define as "a signal whose intensity is causally related to the quality being signalled, and which cannot be faked." To give a concrete example, the fundamental frequency of a vocalization might be considered to be an index of body size. The argument is that the fundamental frequency is determined primarily by the size of the vocal-production apparatus, for example by the length of the vocal folds (or vocal "cords") in many vertebrates. Longer cords produce lower frequencies, vocal-cord length is correlated with body size, and therefore small animals are constrained to produce higher frequencies than large animals.

Note that this argument requires two assumptions: first, that there is an inherent relationship between a signal property and the structure that produces the signal, and second, that there is some necessary relationship between the structure that produces the signal and an attribute of the signaler of interest to the receiver. We consider the validity of these assumptions with respect to vocal signals and body size in chapter 4. For now, we want to point out that an argument based on these sorts of constraints can often be recast in terms of developmental costs. In the example just described, the size of the vocal apparatus, and thus the length of the vocal cords, is not absolutely determined by body size; rather, the size of the vocal apparatus can vary independently of overall size to some degree. It is possible, then, for an individual to develop a vocal apparatus of larger size than that of other individuals of identical body size, but it might pay various developmental costs for doing so. Thus the boundary between handicap and index signals is not always clear.

Deception Redux

Grafen's (1990a, b) models in particular were tremendously influential in convincing researchers that signal reliability is not only possible but probable. Suddenly, the major theoretical puzzle was not how signals could possibly be reliable, but how they could ever be deceptive. Grafen (1990b) himself was moved to ask "What happened to cheating?" As an example of a scenario that might allow deception to occur, Grafen (1990b) suggested that two groups of signalers might exist, for one of which signaling was cheaper than the other, holding signaler quality constant. Those for whom signaling was less costly would signal at a higher level than expected and would benefit from the discrepancy, fulfilling our criteria for deception. Another possibility is that some signalers (for whatever reason) receive greater benefits from signaling than others do and signal more intensely than expected for that reason. More complicated scenarios are also possible, in which both the costs and benefits of signaling differ across signalers. Costs and benefits might differ with respect to signaler age, history, physiological state, and so forth.

In subsequent chapters we will review models in which the balance between costs and benefits makes deception an evolutionarily stable strategy, or ESS, within stable signaling systems. In such "ESS models" (Grafen and Johnstone 1993), a signaling system with some admixture of deception can only be stable if receiver response to the signal is adaptive on average. This condition does not necessarily require that deceptive signals be rare, or even in the minority; deception can be the rule rather than the exception if the benefit of responding to an honest signal is high enough and the cost of responding to a false signal is sufficiently low.
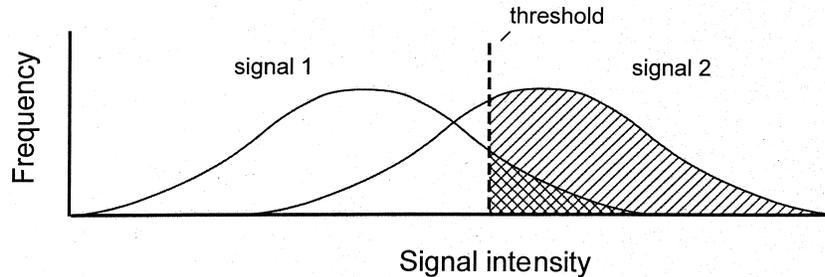
FIGURE 1.3. The problem of discriminating between two signals that vary along a single dimension, in this case signal intensity (adapted from Wiley 1994). Signal 1 is given by signalers of low quality; signal 2 is given by signalers of high quality. The receiver sets a threshold and responds only to signals that exceed that threshold. Those signals above the threshold and under the signal 2 curve represent correct detections (hatching with positive slope); those above the threshold under the signal 1 curve represent false alarms (hatching with negative slope). Sliding the threshold to the right decreases the number of false alarms but at the cost of decreasing the number of correct detections as well. Conversely, sliding the threshold to the left increases the number of correct detections but at the cost of increasing the number of false alarms as well. It is impossible to set the threshold at a value that simultaneously minimizes false alarms and maximizes correct detections.

Another way to view deception is as a consequence of failure on the part of receivers to discriminate between classes of signals; this viewpoint puts deception into the realm of signal detection theory (Wiley 1994, Getty 1995, 1996). Suppose we have two classes of signalers differing in some attribute important to receivers, such as quality. The two classes produce signals that differ along a single dimension, such as intensity (figure 1.3). Signals vary within a class, and there is some overlap between the two classes. A receiver that benefits from discriminating in favor of the signalers of high quality can set a threshold below which it will not respond and above which it will. A receiver that sets a high threshold minimizes its chances of responding to signals from low-quality individuals (i.e., it minimizes false alarms), but at the cost of failing to respond to signals from some high-quality individuals (i.e., of missing correct detections). Thus, no threshold can simultaneously minimize false alarms and maximize correct detections. The optimal receiver strategy depends on the frequency of the two signaler classes and the payoffs of the various possible outcomes (false alarms, correct detections, etc.) (Wiley 1994). If the cost of a false alarm is low and the benefit of a correct detection high, then "adaptive gullibility" may be favored, with the threshold set low so that a receiver often responds to incorrect signals (Wiley 1994). Those incorrect signals that are above threshold can be considered to deceive the receiver.

The signal-detection and ESS approaches to deception are not antithetical. To simplify somewhat, the signal-detection approach assumes that the correct

and incorrect signals cannot be completely separated by a receiver, and receiver response is maintained because the benefits of responding to the subset of correct signals above the optimal threshold outweigh the cost of responding to the subset of incorrect signals above that threshold. The ESS approach assumes that the correct and incorrect signals are not separable at all, and receiver response is maintained because the benefits of responding to the complete set of correct signals outweigh the costs of responding to the complete set of incorrect signals. The two approaches can be combined, by introducing signal variation into ESS models (Johnstone 1998), but most ESS models do not make this step. For many of the actual signaling systems that we will examine, honest and deceptive signals appear to be identical, so that a signal-detection approach does not apply. One could still argue, however, that the honest and dishonest signals might be separable on the basis of context.

Dawkins and Guilford (1991) have argued that less-reliable conventional signals will replace honest-handicap signals because of the costs that handicap signals impose on receivers. Costs to receivers are particularly likely in systems, such as roaring contests between red deer (Clutton-Brock and Albon 1979), where one individual can induce another to signal maximally only by signaling maximally himself. A more general cost to receivers is the time investment that in many cases is necessary to assess displays; for example, a female songbird might have to listen to a male for a considerable period in order to estimate the size of his song repertoire. Time spent attending to a display imposes an opportunity cost, in the sense that the receiver might be investing its time in something else of value, and may also impose a survival cost, if proximity to a signaler increases risk of predation. Dawkins and Guilford (1991) argue that, given these and other possible costs of receiving elaborate signals, it would often be advantageous to both receiver and signaler for the signaler to use cheap conventional signals instead of handicaps.

The term "conventional signal" has been used with various meanings in the animal-communication literature (e.g., Maynard Smith and Harper 1988, Dawkins and Guilford 1991), but the clearest definition in our opinion is that provided by Guilford and Dawkins in a later paper (Guilford and Dawkins 1995). For them, a conventional signal is one in which there is "a degree of arbitrariness in the relationship between signal design and signal message" and therefore a need for an agreement (a convention) on what the signal means (Guilford and Dawkins 1995, p. 1692). A nonarbitrary signal might be one in which the signal has a cost that is related to the message, as when the "signal 'uses up' some of the quality being signalled about" (Guilford and Dawkins 1995). For a conventional signal there is no such relationship between signal design, signal cost, and signal message. Instead, the costs of conventional signals are inherently receiver-dependent, and the convention of what the signal means is maintained solely by the response of the receivers.

Implicit in Dawkins and Guilford's (1991) argument, then, is the assumption that receiver-dependent costs are less effective in maintaining signal reliability than are other types of costs. As we discuss later (chapter 4), both reliability and deceit can emerge from signaling models in which the only signal costs are receiver-dependent. The same is true, however, of models in which signals have production, developmental, or maintenance costs (chapter 3); that is, these models too can generate reliable signals but also support some level of dishonesty. Thus it is not obviously true that receiver-dependent costs are more likely to allow dishonesty in signaling systems than are other categories of cost. At the same time, it does seem to be true that theory allows some level of dishonesty in signaling systems of most types.

Evolutionary Interests of Signalers and Receivers

Honest signaling in the absence of signal costs would be expected if the signaler and receiver have identical interests in an evolutionary sense, meaning that a fitness gain experienced by one individual produces an equal fitness benefit for the other. Communication between two such individuals would be akin to communication between two cells or two organs within an individual, and one in general would not find reliability puzzling for signaling systems that operate within individuals (but see Zahavi and Zahavi 1997, who suggest that costs are important in ensuring the reliability of signaling between cells within single individuals). Evolutionary interests also ought to be identical for separate individuals if they are genetically identical.

More commonly, communication occurs between genetically distinct individuals, and here we can distinguish three likely cases with respect to evolutionary interests. In the first, the interests of signaler and receiver are overlapping though not identical. This description is most likely to apply when signaler and receiver are genetic relatives. For related individuals, a fitness benefit experienced by one is necessarily experienced by the other as well, though to a reduced degree. Conflicts of interest are still possible, whereby the optimal outcome of an interaction differs for the two interactants. In chapter 2, we discuss theoretical models that have been proposed for cost-free, reliable signaling between individuals with overlapping interests. We then examine what is known empirically about two of the best-studied examples of communication between relatives: begging, in which offspring solicit food or some other resource from their parents, and alarm signaling, in which one individual warns another of the approach of a predator. Alarm signaling can be interpreted in ways other than as signaling between relatives, for example as signaling between prey and predator. We consider these possible interpretations before examining aspects of reliability and deceit in this type of system. Finally, in chapter 2 we introduce another explanation for reliability, which can be

thought of either as existing outside of the costly signaling paradigm or as simply positing a novel kind of cost. Here reliability is advantageous to a signaler because it interacts multiple times with a receiver able to recognize signalers as individuals and to remember their past record of reliability (Silk et al. 2000). In this situation, even if the signaler gains a benefit from deceiving a receiver on one interaction, that benefit may be outweighed by the cost of having the receiver fail to respond to its signals during subsequent interactions.

A second possibility is for the evolutionary interests of signaler and receiver to be separate but not necessarily opposing. We describe this situation as "divergent interests," and consider it to apply to mate attraction and mate choice. Suppose a male acts as the signaler, trying to attract a female, who acts as the receiver. Signaler and receiver in this context typically are genetically unrelated, so genetic relatedness does not provide a tie between their respective fitnesses. A large body of both empirical evidence and theory suggests that females will benefit from assessing prospective mates on some aspect or aspects of "quality." Females therefore will be selected to attend to signals of male quality, as long as those signals are reliable. Males, however, might benefit from exaggerating their quality, other things being equal, thus pushing the signals toward unreliability. This kind of system was the principal focus of Grafen's (1990a,b) signaling models. We review those models in greater detail in chapter 3, together with more recent models that attempt to show how deception might coexist with reliability in mate-attraction signals. We then review some of the empirical results on reliability and deceit in mating signals. We do not attempt a comprehensive review of the vast literature on male attributes that affect female choice of mates (Andersson 1994). Instead, we focus on just three categories of mating signals, chosen because they illustrate particularly well the issues of reliability and deceit. These systems are carotenoid signals in fish and birds, song in songbirds, and tail length in birds.

A third possibility is for the evolutionary interests of signaler and receiver to be in direct opposition. This description applies most generally to cases in which two animals engage in an aggressive contest for possession of some resource, such as food, territory, or mates. Here the interests of the two interactants are necessarily opposed, in the sense that if one animal wins the resource the other loses it. Signaling models, starting with that of Enquist (1985), have suggested that reliable signaling can occur in aggressive interactions despite the opposing interests of signaler and receiver. We review those models in chapter 4, and then examine some actual aggressive signaling systems. Again, we focus on just a subset of the systems that have been studied, chosen because they illustrate the issues of interest. The aggressive signaling systems we have chosen to review are postural displays and badges of status in birds, weapon displays in crustaceans, and calling in frogs and toads.

In many signaling interactions, there are individuals other than the primary interactants that benefit from acquiring whatever information is exchanged.

When individuals that are not directly involved in a signaling interaction nevertheless gather information from it, their behavior is termed "eavesdropping" (McGregor 1993, McGregor and Dabelsteen 1996). Eavesdropping may have important implications for the reliability of a signal, in the sense that unintended receivers can impose additional costs on a signaler (or, in theory at least, may provide additional benefits), and thus a signaler may be selected to modify the reliability of signals that are subject to eavesdropping. Signaling very commonly occurs in "networks" of signalers, rather than in closed dyadic interactions, suggesting that eavesdropping may often be a key factor in the evolution of signaling systems. In chapter 5, we discuss evidence that eavesdropping occurs, and we explore the implications of eavesdropping for signal reliability.

In examining each of the signaling systems that we review in chapters 2, 3, and 4, we will take a standard approach, one that is based on the logic of the reliable signaling problem. One way of stating this logic is that the existence of a signaling system, in which a signaler signals and a receiver responds, implies that the signal has some appreciable level of reliability. The reliability of the signal in turn implies, according to theory, that the signal has some appreciable cost. Therefore, in addressing each system we begin by reviewing evidence on whether receivers actually respond to the signal in question. Demonstrating that receivers respond is the crucial step in establishing the existence of a signaling system. Once a signaling system has been shown to exist, the next step is to examine whether the signal is reliable, as predicted by the logic sketched out above. Reliability can be established by assessing the correlation between attributes of the signal and whatever it is that the receiver benefits from knowing. If the signal is indeed reliable, we next assess the signal's costs. Unfortunately, theory is not terribly clear on the magnitude of the costs needed to maintain reliability; nevertheless, we can make some headway in determining whether the signal is more costly than seems needed simply for transmission. Finally, for each signaling system we review, we will discuss any evidence adduced for the deceptive use of the signal. Certain of the systems that we will focus on have been chosen because they do provide convincing evidence of deception.

Throughout we focus on natural signaling systems, rather than on signaling systems imposed on animals by humans. This focus is in keeping with our evolutionary interests; we want to see what signaling systems natural selection has come up with, not what humans can induce animals to do. We also will concentrate, although not quite exclusively, on communication within species. We have made this choice because it is in the within-species context that there is a clear contrast between the older view of communication as essentially cooperative and the newer view of communication as the product of selection for behavior that furthers each individual's own interests. This decision causes us to exclude some classic cases of deception, such as Lloyd's (1965, 1986)

demonstration that females of the predatory firefly genus *Photuris* mimic the flash patterns of female fireflies of the genus *Photinus*, and in this way lure male *Photinus* to their deaths. Although fascinating, in our view such examples of interspecific deception do not pose the same kind of evolutionary puzzles as intraspecific deception, because no one would expect communication across species to be cooperative. Where we do discuss interspecific communication is in cases where there is debate about whether a signal has evolved for a within-species or between-species function.

Another decision we have made is not to make use of one-time observations of the behavior of individual animals. Such observations have long been used as evidence for the occurrence of deception in nonhuman animals, dating back at least to the work of Romanes (1883). Romanes (1883) gathered examples of deception from lay observers as part of his attempt to establish the occurrence of intelligent behavior in animals, and then interpreted those examples rather liberally ("Another of my correspondents, after giving several examples of the display of hypocrisy of a King Charles spaniel . . ."). In more modern analyses, one-time observations of apparent deception, made by scientifically trained observers, have been systematically collected and categorized for primates by Whiten and Byrne (1988, Byrne and Whiten 1992), again with an eye chiefly to the implications of this type of evidence for animal intelligence. As our interests are not in the cognitive aspects of deception, we can avoid anecdotal evidence without passing any general judgment on its scientific usefulness (Burghardt 1988, Byrne and Whiten 1988). We confine our attention to signaling interactions that occur regularly enough that they can be statistically analyzed using data from single studies. We emphasize experimental evidence whenever possible, but note that experimental methods are more applicable to some of the questions we address (such as the response of receivers) than to others. In particular, the reliability of signals often can be addressed only by measuring correlations between the attributes of signal and signaler.