2015, in Jean Decety and Thalia Wheatley (Eds.), The moral brain (pp.19-34). Cambridge, MA: MIT Press.

2 Adaptationist Approaches to Moral Psychology

Andrew W. Delton and Max M. Krasnow

Keep your word. Love your neighbor. Help those in need. Only marry the child of your father's sister. Stone those who have had sex out of wedlock. Kill outgroup members.

Human moral communities have developed a variety of moral rules, injunctions, prescriptions, and suggestions. How are we to understand the origins and nature of human morality, both at the level of universal building blocks and the level of cultural variation and elaboration? As this volume illustrates, understanding human moral psychology is a truly interdisciplinary endeavor, drawing important contributions from psychologists, neuroscientists, philosophers, biologists, legal scholars, and many others. Our goal in this chapter is to illustrate the utility of taking an adaptationist approach from evolutionary biology to understand universal aspects of moral psychology. We first describe what it means to take an adaptationist approach. We next give several examples of how an adaptationist approach has informed the study of certain aspects of moral psychology. We then briefly conclude with what we see as the value of this approach to the study of moral psychology broadly.

Adaptation

Charles Darwin's *The Origin of Species* is justly celebrated for revolutionizing biology. Although Darwin himself described the book as "one long argument," the book actually develops several important and logically separable hypotheses (see Mayr, 1982). For example, all life is descended from one or a small number of ancestral forms. Additionally, evolution (i.e., descent with modification) is a process of mainly small gradual changes. Most revolutionary, natural selection is *the* driving force causing organisms to become ever more adapted to their environments.

Why is natural selection the best known and most used of these hypotheses? This is because natural selection was the first, and is still the only, non-question-begging explanation for the origin of complex functional organization in animals and other organisms (Williams, 1966). Prior to Darwin, scientists had regularly invoked special creation to explain the existence of complex functionality. William Paley famously asked, if we were to find a working pocket watch on the heath, would we think that random forces conspired to create it or, more parsimoniously, that it was fashioned by the activities of a mind? Similarly, if one comes across a platypus or Venus flytrap, which is more likely, that their ability to survive and reproduce was produced by chance or by an intelligent designer? The theory of natural selection showed, for the first time, a way outside of this false dichotomy (Dawkins, 1986).

Modern evolutionary biologists distinguish between the process of natural selection and its products, *adaptations*. Wings, chemosensory systems, eyes, parental love: all of these are adaptations—features of organisms that came into existence through natural selection because they caused faster genetic replication. Of course, not all features of organisms are adaptations. Some are merely *by-products*, features pulled along incidentally with adaptations. Although bones are white and human blood is red, these traits are by-products, not design features: calcium phosphate helps make sturdy bones, and iron-rich hemoglobin usefully transfers oxygen; their respective features of whiteness and redness do not contribute to these functions. In addition to adaptations and by-products, organisms are also composed of random noise, genetic and developmental insults to an organism's system. For instance, there are a precise number of hairs on your head. This particular number is not a design feature nor is it a by-product shared by all humans like the whiteness of bones; instead, it is simply random noise.

Importantly, by-products and random noise do not lead to complex, functional systems. To the extent that psychologists, including moral psychologists, are interested in complex aspects of human moral nature, they are interested in psychological adaptations. Thus, the conceptual tools from evolutionary biology are additional sources of information that psychologists can use to generate testable hypotheses about the moral mind. Connecting moral psychology with evolutionary biology and adaptationism is one part of the increasingly interdisciplinary nature of moral psychology. We suspect that using the tools of adaptationism will contribute just as

much to moral psychology as have connections with neuroscience, the law, and philosophy.

So what kind of moral psychology might selection have designed? Although not explicitly couched in the language of adaptationism, theories of moral development like Piaget's and Kohlberg's implicitly assume that the adaptations that give rise to moral psychology are not specialized for moral thought. Instead, they reflect more general processes such as reasoning and induction. Other work, however, suggests that evolution has equipped the mind with quintessentially morality-specific concepts such as "good" and "bad" (Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Hamlin, Wynn, & Bloom, 2007; van Leeuwen, Park, & Penton-Voak, 2012). Moreover, evolution may have equipped us with skeletal domains of moral intuitions such as intuitions about community, harm, and fairness (Haidt & Bjorklund, 2008). Our goal in this chapter is to push this argument farther. Natural selection may have designed very specific moral concepts and abilities into the human cognitive architecture, ones more fine-grained than simple right and wrong. Here we present three examples of such fine-grained concepts.

Moral Concepts for Group Cooperation

What constitutes morality and moral psychology? This is a big question and one we have no wish to discuss exhaustively here. Some scholars treat morality as encompassing many aspects of human behavior—altruism, helping, generosity, virtue, and so forth. On this view, helping our kin and our friends is part of morality (e.g., Ridley, 1996). From an evolutionary perspective, psychological mechanisms for aiding kin and friends were likely created by selection pressures such as kin selection and selection for reciprocity. But what makes these selection pressures and the psychologies they create moral? Many other animal minds seem to have been shaped by them, and it is not clear that these animals have anything approaching human moral systems. Theorists taking a restrictive view propose instead that the core of morality are judgments of right and wrong, the moral praise or condemnation that follows, and attendant emotions such as anger, disgust, guilt, and shame (e.g., DeScioli & Kurzban, 2009, 2013).

We are agnostic on this issue. Indeed, we are not sure whether at the margins moral psychology can be cleanly separated from social psychology.

Nonetheless, to make our argument as relevant to the widest audience possible, here we focus mostly on examples that fall under a more narrow reading of the moral domain.

Our first example is drawn from human group cooperation. Humans have evolved a unique kind of cooperative foraging system, often called *resource pooling* or *communal sharing* (Fiske, 1992; Gurven, 2004). Foraging is risky business. On any given day a forager may be a victim of bad luck and return empty handed. Worse, injury and illness can keep a forager disabled for significant lengths of time (Sugiyama, 2004a). Other foragers in one's group, however, may have more food than they can usefully consume. By sharing, foragers can buffer against shortfalls and hardships. Over the long term everyone benefits: yes, you may be giving up some food now to help a person in need, but later, when you are eventually in need, your generosity will be repaid. Such sharing can have huge benefits: in a study of the Shiwiar of the Ecuadorian and Peruvian Amazon, approximately 65% of the sample would be dead had they not been part of a resource-pooling system (Sugiyama, 2004a, 2004b).

Although they provide huge benefits, resource-pooling systems are not trivial to maintain. One problem is that resource-pooling systems (and other types of group cooperation) are vulnerable to free riders, people who take the benefits of a resource-pooling system without contributing to it. Because they take collective benefits but do not pay the costs of creating them, free riders are materially better off than cooperators. This means that, in the absence of any countervailing mechanism, free riders are favored by natural selection over cooperators, and a design for free riding would come to predominate in a population. Of course, humans do engage in resource pooling so somehow the free-rider problem must be solved.

Evolutionary modeling work shows that the free-rider problem can be diminished or eliminated if free riders are either excluded from cooperation or are punished to induce their cooperation (Boyd & Richerson, 1992; Sasaki & Uchida, 2013). Thus, free riders fall under the narrow version of moral psychology—they are likely to be judged as doing wrong and to elicit anger and punitive sentiment. The question we have explored is whether selection has designed a specific mechanism—a free-rider concept—to identify and respond to free riders. In addressing this question we have contrasted the hypothesis that there is a dedicated free-rider concept with a number of alternatives, including economic rationality, statistical learning, and a

moral psychology that includes concepts of right and wrong, but nothing so specific as a dedicated free-rider concept.

Why would a dedicated free-rider concept be necessary? Identifying and responding to free riders is a difficult and nuanced problem. Evolutionary analysis reveals that free riders cannot be identified merely by their overt contributions to a communal sharing system. Resource-pooling systems exist to allow people to survive when they are most vulnerable, such as when they are unable to contribute to the common pool (Tooby & Cosmides, 1996). Using a metric of overt contribution to identify free riders would be counterproductive—people in legitimate need of aid would also be classified as free riders. Instead, evolutionary modeling shows that the mind needs a way to "see through" overt behavior and instead make inferences about a person's underlying motivations (Delton & Robertson, 2012).

In a series of experiments we have found that the human mind draws precisely this distinction (Delton, Cosmides, Guemo, Robertson, & Tooby, 2012). These studies used an implicit measure of social categorization. This method uses confusions in memory as a proxy for categorization (Taylor, Fiske, Etcoff, & Ruderman, 1978): thus, if persons A, B, and C are selectively confused with each other, and persons X, Y, and Z are selectively confused with each other, then the two sets are assumed to be represented as distinct categories. When presented with foragers who differ in their motivations—some desiring to exploit a sharing system, others well-intentioned cooperators—experimental subjects sharply distinguished between the two and viewed those desiring to exploit the system as free riders deserving of exclusion or punishment. Importantly, the experimental stimuli were arranged such that free riders and cooperators contributed identical amounts; although free riders purposefully withheld a resource, cooperators attempted to contribute but were prevented by reasons beyond their control. Despite identical contributions, free riders were categorized separately from cooperators based on their motivation and elicited a clearly moralistic response, such as being viewed as deserving exclusion or punishment.

These experiments show that motivations are used as a cue to free riding independent of overt contributions. But perhaps, as predicted by an economic rationality account, overt contributions are the cue of interest, and motivations are only secondarily informative. After all, over the long term, people with exploitive motives are likely to contribute less than people with cooperative motives. Thus, the mind might detect a correlation between

motivation and long-term contribution and learn to use motivation as a secondary cue to free riding. This alternative predicts that differences in overt contributions should be used as cue to free riding, perhaps as an even stronger cue than motivation. We tested and experimentally falsified this prediction: when motivation is held constant while overt contributions are manipulated, people did categorize based on this difference, but it did not lead to moralistic responses—no desire to punish or avoid. Instead, people who contribute less were viewed as less competent—an important, but orthogonal, dimension in person perception (Cuddy, Fiske, & Glick, 2008).

Let us consider a final alternative hypothesis. An adaptationist analysis of the logic of resource pooling suggests that the mind has a dedicated free-rider concept for identifying and responding to free riders. Although economic rationality and statistical learning do not fully explain the data, could a more general moral psychology do the job? Perhaps all that is needed is a system that categorizes people as good or bad based on intentional actions as well as other information (Cushman, 2013). Indeed, current evidence suggests that the mind contains reliably developing systems for making general judgments about immorality (Hamlin et al., 2007; van Leeuwen et al., 2012). The question, however, is not whether the mind has a superordinate system for making moral discriminations. The question is instead whether the mind also has finer-grained levels of categorization. These fine-grained levels would exist because different adaptive problems require different solutions. We have tested this issue by comparing free riders to other types of moral violators. In one study, we tested whether free riders were distinguished from what, intuitively, is a very distinct type of moral violation: intentional, unprovoked physical battery. Experimental subjects did indeed categorize free riders separately from batterers and responded more moralistically to batterers (Delton, Cosmides, Guemo, Robertson, & Tooby, 2012).

But still, a free-rider concept may not be needed. The mind may distinguish unprovoked battery from violations involving entitlements and obligations (some theories treat entitlement and obligation as psychologically identical) (Cheng & Holyoak, 1985). Free riders could be construed as violating these general deontological principles but nothing so specific as the rules of a sharing system. To test against this we contrasted free riders with an even more closely matched moral violation: a person who steals a resource communally owned by the group. Thus, both free riders and

thieves were in one sense expropriating a group resource for their own benefit. Nonetheless, the logic of resource pooling and the logic of ownership and theft are different adaptive problems. Does the mind categorize them separately? Yes: experimental subjects categorized the two separately and reacted to free riders with greater moralistic punishment.

Although more work needs to be done, the current evidence points to a very specialized piece of conceptual machinery—a specialized free-rider concept. This moral concept was uncovered by taking an adaptationist approach to the human mind, by looking at the kinds of problems humans have regularly solved and generating psychological theories about the kinds of computations required to solve these problems. In other work we have extended this approach to uncover other concepts related to group cooperation, at least some of which are clearly moral. For instance, the mind appears to have a specialized concept for understanding and responding to newcomers to a coalition (Cimino & Delton, 2010; Delton & Cimino, 2010). The mind also appears to contain a concept for identifying and selectively associating with highly valuable cooperation partners (Delton & Robertson, 2012). And the mind appears to have a concept of public goods: not all types of group cooperation produce goods exploitable by free riders; public goods are the subset vulnerable to free riding. Our recent research shows that the mind distinguishes public goods from other cooperatively produced goods and mobilizes relevant moral sentiments to prevent free riding on public goods (Delton, Nemirow, Robertson, Cimino, & Cosmides, 2013).

Moral Concepts for Social Exchange

Our second example is drawn from human dyadic cooperation, often organized as social exchange (Cosmides & Tooby, 2005). When two parties have goods or services that are more valuable to the other than they are to the bearer, then both are better off if they can trade. These items can be traded contemporaneously: if I have an abundance of peanut butter but no jelly, and you have an abundance of jelly but no peanut butter, we can exchange on the spot and both realize gains in trade. Or, exchange can happen over time, such as when I care for you while you are ill now and you care for me when I am ill in the future, allowing both of us to survive our misfortunes. Exchange is a culturally universal (Brown, 1991) major facet of human

social life, and the majority of the improvement in the human condition over the last 10,000 years of human history is attributable to the increasing prevalence of such exchange and the specialization and division of labor it enables (Ridley, 2010).

But, realizing gains in trade is a risky business: if you trust me and invest in an exchange (e.g., give me some of your jelly), there is always the chance that I will try to cheat you and not reciprocate (e.g., keeping all my peanut butter for myself). The possibility of cheating suggests that you cannot be uniformly trusting. But neither can you be uniformly *dis*trusting: you can never realize the benefits of exchange without sometimes taking the risk of trusting someone else. Some of this risk can be mitigated by effective partner choice; some partners will present better targets for cooperative investment than others, given their past history or reputation. And some of this risk can be mitigated by effective responses to cheating; either terminate the relationship and thus cut off future loses or deploy sanctions in an attempt to recalibrate the partner's future behavior and thus salvage the future potential for long-term gains in trade. These are adaptive problems, and their adaptive solutions are mechanisms for relationship cultivation, partner selection, and relationship maintenance (see also Cosmides & Tooby, 2005).

It is common wisdom that the outcomes of exchange decisions are moralized: upholding a bargain is morally right, and cheating is morally wrong and often worthy of punishment. But is that as specific as the system gets, merely sketching this moral dimension? Because there are long-enduring benefits, costs, and complexity to social exchange, an adaptationist approach predicts that there should be greater specificity and functional organization in the design of the psychological mechanisms that underpin it. For example, moral responses to exchange violations should be structured to improve long-term profitability for the responder. As such, moral responses to cheating should vary by their relevance to both partner selection and the prospect of improving a dissatisfactory relationship.

We tested these questions in a series of experiments (Krasnow, Cosmides, Pedersen, & Tooby, 2012). We first measured participants' cooperative dispositions. In one study cooperative dispositions were measured by a survey of scenarios in which participants could cheat a business, a friend, a roommate, a co-worker, and so forth, assuming that they could not be caught. A second study measured cooperative dispositions by having participants

play a series of one-shot prisoners' dilemma (PD) games, each game ostensibly played with a different partner. (All partners in these games and those described below were actually played by a computer.) In the PD games subjects were randomly paired, and each was given a resource that was more valuable to her or his partner than it was to her- or himself. Participants then each decided whether to keep the resource or give it to his or her partner. In these games the payoff-maximizing decision is to not give. However, if both participants give, they are made better off than if neither gave.

After these measures participants were paired with a partner, and, for some partners, participants were given information about the partner's decisions in the cooperative disposition task. They then played a two-round trust game with this partner. In the game the first mover was given a sum of money that could be evenly split, or he or she could risk trusting the second mover and transfer the money to that person, with the sum multiplying along the way. The second mover could then evenly split this multiplied sum (cooperate by reciprocating the first mover's trust) or could claim the lion's share for him- or herself (cheat). If the second mover cheated, the first mover could spend a small amount of money to impose a large cost on the second mover (punish) or do nothing. Participants first played one round of this game taking the role of first mover and then a second round of the game as the second mover.

The results of these experiments are telling:

Participants exhibited a default openness to exchange. In conditions without any information about a partner's past history of cooperation, the majority trusted their partner.

Participants were sensitive to reputation and used reputation as a cue to how the partner would treat them. There may be individuals who are lousy at exchange with everybody, and they can be avoided by paying attention to how they treat third parties. But every person represents a unique mix of interests, and someone might be a better cooperative partner for me than he or she is for you. If the architecture responsible for trust merely encodes instances of cheating for avoidance and the attribution of moral wrongness, it would miss out on this second class of actor. If instead the architecture used cues hierarchically in the order of their utility in predicting who will most likely cooperate with the self, then being cheated in the past should carry more weight in the regulation of trust than if that same person cheated someone

else. And the mind agrees: when given information about their partners' previous decisions to cheat third parties in the survey, participants were much less likely to trust partners with a history of cheating. But when given information about the partner's decisions in two PDs, one with the participant and one with a third party, only the participant-specific information now regulated trust while the third-party information did not. The moral architecture that interprets the violation of exchange norms appears to be designed to discriminate profitable exchange partners and not merely to attribute moral wrongness generally.

Participants deployed punishment to maintain an ongoing relationship. If the moral architecture responsible for attribution of wrongness and punitive sentiment responded to cheating per se, it would not differentially respond to cheaters the person intends to continue cooperating with. But expending punitive effort to reform a relationship you intend to abandon is a waste of resources. And the mind agrees: participants were much more likely to spend money to punish their cheating partner if they intended to cooperate with them in the next round. Participants who abandoned the relationship were much less likely to punish their partner for cheating. Interestingly, convergent evidence for this two-pronged strategy can be found in the domain of criminal justice, where "high-value" offenders are more likely to receive punishment in the form of rehabilitation than are "low-value" offenders (Petersen, Sell, Tooby, & Cosmides, 2012).

In short, the cognitive architecture that produces our moral responses to being wronged and learning that others are wronged does not appear designed for dispassionately evaluating transgressions and thus appears not to result from only a general system of moral right and wrong. Instead, the system appears to contain features that deploy moral approval and disapproval, trust and punitive sanctions, in order to manage the risks of engaging in social exchange in order to effectively profit from gains in trade.

Condemnation and Coordination

Moral communities the world over regularly condemn various classes of behavior. Among the Yanomamo of the Amazon, it is morally condemnable to marry a parallel cousin (i.e., mom's sister's child or dad's brother's child); instead, it is considered appropriate to marry a cross cousin (i.e.,

mom's brother's child or dad's sister's child) (Chagnon, 1996). Among the !Kung San of the Kalahari desert, not sharing your meat with others is condemnable (Cashdan, 1980), although they feel little compunction wearing clothing that would make pop stars blush. And many modern Americans morally condemn smoking tobacco (Rozin & Singh, 1999), an activity common in many other times and places.

Although moral condemnation is common, it is anything but simple. One complexity of condemnation is that the classes of condemned behavior vary across time and space. A related question, which we discuss here, is how multiple people can successfully coordinate their condemnation. At a very general level moral condemnation serves to change someone's behavior. If I have done something wrong, you might condemn me to get me to refrain from repeating the wrong in the future. But if you condemn me alone, I might simply retaliate against you rather than change myself. Given the potential costs of condemning alone, condemnation is often a social act: multiple people coordinate their condemnation on a specific target (Boehm, 1993; Boyd, Gintis, & Bowles, 2010).

Coordination, however, is difficult to accomplish. In a series of recent papers DeScioli and Kurzban have shown how the difficulties of coordination impact our moralistic judgments (DeScioli & Kurzban, 2009, 2013). They started with an enduring mystery in moral psychology and philosophy: the difference in culpability between sins of omission and sins of commission. Committing a wrong is usually judged more harshly than not preventing a wrong, even if the prevention would have cost nothing. For instance, imagine that Bill is alone in the control room for a train hub. Bill sees a train hurtling down track A and a person standing on track A. Bill could push a button to divert the train to track B, but he does not and the person is killed. How wrong was that? Now imagine that the train is still hurtling down track A, but the person is standing on track B. Bill presses the button and diverts the train to track B and person is killed. How wrong was that? If you are like most people, Bill was more morally wrong in the second scenario where he actively diverted the train.

Why? Logically the distinction appears meaningless. In both cases a death occurred, and Bill could have prevented it. Yet Bill's act of commission is judged more harshly. DeScioli, Bruening, and Kurzban (2011) argue that one of the functions of our moralistic reactions is to help coordinate multiple people in condemnation. Because of this, one factor that should

influence the strength of our moral judgments is the ease with which coordination can be achieved. If coordination cannot easily be achieved, it will be too costly to enact condemnation. When condemnation would be costly, there are not benefits in becoming morally outraged (and possibly costs if the retaliation is severe). But if coordination can be achieved, then condemnation will be relatively cheap and moral outrage worthwhile.

What do omission and commission have to do with this? DeScioli and Kurzban argue that sins of omission leave behind little in the way of public evidence that a wrong was committed. Without public evidence that can be used to form common knowledge about the wrong (cf. Pinker, Nowak, & Lee, 2008), sins of omission make coordination difficult. Sins of commission, conversely, do leave behind clear public evidence, making coordination much easier to achieve. DeScioli, Bruening, and Kurzban (2011) designed an ingenious series of experiments to test this hypothesis against the most prominent class of alternative theories. These alternative views propose that direct causal responsibility is the distinction driving the omission–commission effect: sins of commission are directly caused; sins of omission are not. To test against this, DeScioli, Bruening, and Kurzban created scenarios similar to the one described above. They also added a critical test condition, which to simplify and continue with the above examples worked as follows: as Bill watched the train on track A hurtle toward a person on track A, he pressed a button that simply recorded his lack of desire to move the train onto a different track. This was a sin of omission—he did not directly cause the train to kill the person. But it was a *public* sin of omission—there was clear evidence that he allowed the person to be killed. In these studies, when sins of omission leave behind clear public evidence, the omission–commission effect vanishes: sins of public omission elicit almost identical moralistic responses as sins of commission. Moreover, people appear to be sensitive to the fact that they will not elicit condemnation when public coordination is difficult, and thus they behave more self-interestedly in these conditions (DeScioli, Christner, & Kurzban, 2011).

These studies only scratch the surface of the complex phenomenon of moral condemnation. Nonetheless, we believe they illustrate the utility of taking an adaptationist approach. By thinking about the information-processing problems that arise when engaging in moral condemnation, DeScioli and Kurzban were able to generate novel, testable hypotheses.

Parting Thoughts: Why You Should Care about Adaptations for Moral Psychology

In this chapter we have sketched a brief outline of how the adaptationist toolkit has informed the study of human moral psychology in the domains of group cooperation, social exchange, and coordinated condemnation. Our argument, however, is not that adaptationist reasoning is limited to elucidating these few domains of moral thought. The only way for a complex biological structure to endure against the constant destructive forces of entropy and mutation is to reproduce its design faster than its competitors. Thus, to the extent that components of human moral judgment and reasoning are complex in structure, robust in their development, and universal in our species—the components of interest to many moral psychologists—they are very likely to be the product of psychological adaptations. Adaptations are the product of long-term selective regimes: some genetic variants lead to the development of neural systems that, in interaction with environmental regularities, produce fitness-enhancing behavior. These variants thus increase in frequency in the population. The adaptationist approach therefore examines the long-enduring problems that humans have faced to generate hypotheses about the kinds of psychological adaptations that would have evolved to solve these problems (see Delton, Krasnow, Cosmides, & Tooby, 2010, 2011; Krasnow, Delton, Tooby, & Cosmides, 2013). These hypotheses then guide empirical research designed to confirm or falsify them. As illustrated by this volume, insights from neuroscience, philosophy, and the law have begun to inform our understanding of moral psychology. We suggest that the addition of thinking as an adaptationist can improve this understanding.

References

Boehm, C. (1993). Egalitarian behavior and reverse dominance hierarchy. *Current Anthropology*, *34*(3), 227–254.

Boyd, R., Gintis, H., & Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science*, *328*, 617–620.

Boyd, R., & Richerson, P. J. (1992). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology*, *13*(3), 171–195.

Brown, D. E. (1991). *Human universals*. New York: McGraw Hill.

Cashdan, E. (1980). Egalitarianism among hunters and gatherers. *American Anthropologist*, *82*, 116–120.

Chagnon, N. (1996). *Yanomamo* (5th ed.). San Diego: Harcourt Brace.

Cheng, P., & Holyoak, K. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, *17*, 391–416.

Cimino, A., & Delton, A. W. (2010). On the perception of newcomers: Toward an evolved psychology of intergenerational coalitions. *Human Nature*, *21*(2), 186–202.

Cosmides, L., & Tooby, J. (2005). Neurocognitive adaptations designed for social exchange. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 584–627). Hoboken, NJ: Wiley.

Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in Experimental Social Psychology*, *40*, 61–149.

Cushman, F. (2013). Action, outcome, and value: A dual-system framework for morality. *Personality and Social Psychology Review*, *17*(3), 273–292.

Dawkins, R. (1986). *The blind watchmaker*. New York: W. W. Norton.

Delton, A. W., & Cimino, A. (2010). Exploring the evolved concept of newcomer: Experimental tests of a cognitive model. *Evolutionary Psychology*, *8*(2), 317–335.

Delton, A. W., Cosmides, L., Guemo, M., Robertson, T. E., & Tooby, J. (2012). The psychosemantics of free riding: Dissecting the architecture of a moral concept. *Journal of Personality and Social Psychology*, *102*(6), 1252–1270.

Delton, A. W., Krasnow, M. M., Cosmides, L., & Tooby, J. (2010). Evolution of fairness: Rereading the data. *Science*, *329*(5990), 389.

Delton, A. W., Krasnow, M. M., Cosmides, L., & Tooby, J. (2011). The evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. *Proceedings of the National Academy of Sciences USA*, *108*(32), 13335–13340.

Delton, A. W., Nemirow, J., Robertson, T. E., Cimino, A., & Cosmides, L. (2013). Merely opting out of a public good is moralized: An error management approach to cooperation. *Journal of Personality and Social Psychology*, *105*(4), 621–638.

Delton, A. W., & Robertson, T. E. (2012). The social cognition of social foraging. *Evolution and Human Behavior*, *33*, 715–725.

DeScioli, P., Bruening, R., & Kurzban, R. (2011). The omission effect in moral cognition: Toward a functional explanation. *Evolution and Human Behavior*, *32*(3), 204–215. doi:10.1016/j.evolhumbehav.2011.01.003.

DeScioli, P., Christner, J., & Kurzban, R. (2011). The omission strategy. *Psychological Science*, *22*(4), 442–446.

DeScioli, P., & Kurzban, R. (2009). Mysteries of morality. *Cognition*, *112*(2), 281–299.

DeScioli, P., & Kurzban, R. (2013). A solution to the mysteries of morality. *Psychological Bulletin*, *139*(2), 477–496.

Fiske, A. P. (1992). The four elementary forms of sociality: A framework for a unified theory of social relations. *Psychological Review*, *99*(4), 689–723.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, *293*(5537), 2105–2108.

Gurven, M. (2004). To give and to give not: The behavioral ecology of human food transfers. *Behavioral and Brain Sciences*, *27*(4), 543–583.

Haidt, J., & Bjorklund, F. (2008). Social intuitionists answer six questions about moral psychology. In W. Sinnott-Armstrong (Ed.), *Moral psychology: The cognitive science of morality* (pp. 181–254). Cambridge, MA: MIT Press.

Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, *450*, 557–559.

Krasnow, M. M., Cosmides, L., Pedersen, E. J., & Tooby, J. (2012). What are punishment and reputation for? *PLoS ONE*, *7*(9), e45662.

Krasnow, M. M., Delton, A. W., Tooby, J., & Cosmides, L. (2013). Meeting now suggests we will meet again: Implications for debates on the evolution of cooperation. *Nature Scientific Reports*, *3*, 1747. doi:10.1038/srep01747.

Mayr, E. (1982). *The growth of biological thought*. Cambridge, MA: Harvard University Press.

Petersen, M. B., Sell, A., Tooby, J., & Cosmides, L. (2012). To punish or repair? Evolutionary psychology and lay intuitions about modern criminal justice. *Evolution and Human Behavior*, *33*(6), 682–695.

Pinker, S., Nowak, M. A., & Lee, J. J. (2008). The logic of indirect speech. *Proceedings of the National Academy of Sciences USA*, *105*(3), 833–838.

Ridley, M. (1996). *The origins of virtue*. London: Viking, Penguin Books.

Ridley, M. (2010). *The rational optimist*. New York: HarperCollins.

Rozin, P., & Singh, L. (1999). The moralization of cigarette smoking in the United States. *Journal of Consumer Psychology*, *8*(3), 321–337.

Sasaki, T., & Uchida, S. (2013). The evolution of cooperation by social exclusion. *Proceedings of the Royal Society B: Biological Sciences*, *280*, 20122498.

Sugiyama, L. S. (2004a). Illness, injury, and disability among Shiwiar forager-horticulturists: Implications of health-risk buffering for the evolution of human life history. *American Journal of Physical Anthropology*, *123*(4), 371–389.

Sugiyama, L. S. (2004b). Patterns of Shiwiar health insults indicate that provisioning during health crises reduces juvenile mortality. In M. Alvard (Ed.), *Socioeconomic aspects of human behavioral ecology: Research in economic anthropology* (Vol. 23, pp. 379–402). Greenwich, CT: Elsevier.

Taylor, S. E., Fiske, S. T., Etcoff, N. L., & Ruderman, A. J. (1978). Categorical and contextual bases of person memory and stereotyping. *Journal of Personality and Social Psychology*, *36*(7), 778–793.

Tooby, J., & Cosmides, L. (1996). Friendship and the banker's paradox: Other pathways to the evolution of adaptations for altruism. In W. G. Runciman, J. Maynard Smith, & R. I. M. Dunbar (Eds.), Evolution of social behaviour patterns in primates and man. *Proceedings of the British Academy*, *88*, 119–143.

van Leeuwen, F., Park, J. H., & Penton-Voak, I. S. (2012). Another fundamental social category? Spontaneous categorization of people who uphold or violate moral norms. *Journal of Experimental Social Psychology*, *48*, 1385–1388.

Williams, G. C. (1966). *Adaptation and natural selection*. Princeton, NJ: Princeton University Press.