

Methods, Models, and the Evolution of Moral Psychology

Cailin O'Connor

Abstract

Why are we good? Why are we bad? Questions regarding the evolution of morality have spurred an astoundingly large interdisciplinary literature. Some significant subset of this body of work addresses questions regarding our moral psychology: how did humans evolve the psychological properties which underpin our systems of ethics and morality? Here I do three things. First, I discuss some methodological issues, and defend particularly effective methods for addressing many research questions in this area. Second, I give an in-depth example, describing how an explanation can be given for the evolution of guilt—one of the core moral emotions—using the methods advocated here. Last, I lay out which sorts of strategic scenarios generally are the ones that our moral psychology evolved to ‘solve’, and thus which models are the most useful in further exploring this evolution.

1. Introduction

Why are we good? Why are we bad? Questions regarding the evolution of morality have spurred an astoundingly large interdisciplinary literature with contributions from (at least) biology, sociology, anthropology, psychology, philosophy, and computer science. Some significant subset of this body of work addresses questions regarding our moral psychology: how did humans evolve the *psychological* properties which underpin our systems of ethics and morality?

Debates in this literature are ongoing, and often heated. One reason for this is that there are a number of methodological problems researchers must tackle in addressing the evolution of moral psychology. Human psychology evolved in the deep past, over the course of millions of years, and as a result of countless interactions between individuals. The data available to us about these events are often sparse. Further, it is often not even clear exactly what we are trying to account for when we attempt to explain the evolution of human moral psychology. There are several interrelated issues here. One issue arises from disagreements about what counts as our moral psychology in the first place. Are we worried about traits we share with many animals, like prosocial feelings? Or more high level traits like the ability to make moral judgments? The second has to do with determining which aspects of human moral behavior are in fact the result of a biologically evolved psychology, and which are culturally learned.

Email address: cailino@uci.edu (Cailin O'Connor)

Draft of February 20, 2019

PLEASE CITE THE PUBLISHED VERSION IF AVAILABLE!

Despite these, and other, issues, headway can be made. But one must be careful. The goal of this paper will not be to survey everything that has been done on the evolution of moral psychology—that would require a (large) book. (Those who wish to learn more might look at De Waal (1996); Sober and Wilson (1999); Joyce (2007); Churchland (2011) or else James (2010) for a short overview.) Here I will do three things. First, I will discuss the methodological issues at hand at greater length, and defend what I think are particularly effective methods for addressing many research questions in this area. In particular, I will argue that researchers need to use input from empirical literature and also evolutionary modeling techniques to triangulate well-supported theses regarding the evolution of various moral psychological traits.

Second, I will give an in-depth example of this proposed methodology using one of the best modeling frameworks for this sort of exploration—game theory and evolutionary game theory. I will describe how a partial explanation can be given for the evolution of guilt—one of the core moral emotions—using the methods suggested here. Last, I will carefully lay out more broadly which sorts of strategic scenarios are the ones that our moral psychology evolved to ‘solve’, and thus which models are the most useful in further exploring this evolution. Along the way, I will briefly discuss how work on each of these games can inform the evolution of moral psychology.

The paper will proceed as just described. In section, 2, I discuss methodological issues in the study of the evolution of moral psychology. In section 3, I discuss the evolution of guilt, demonstrating how empirical work and evolutionary modeling complement each other in such an exploration. And in 4, I describe several types of games—including the famous prisoner’s dilemma, the stag hunt, bargaining games, and coordination games—to make clear what challenges and opportunities of social life led to the emergence of moral behavior and moral psychology.

2. Methods

As mentioned in the introduction, there are a number of methodological issues plaguing the study of the evolution of moral psychology. Here I outline them in more detail. As we will see, these issues are sometimes intertwined.

Many of the historical sciences face problems related to sparse data. In paleontology, scientists may have access to a single bone, and from that evidence alone attempt to infer features of an entire species.¹ In the case of human behavioral evolution there are various types of historical data available: bone and tool remains, footprints, cave drawings etc. Scientists also draw on data from modern primates, hunter-gatherer societies, and human psychology in developing coherent evolutionary narratives regarding our psychology.

As Longino and Doell (1983) painstakingly outline, sometimes historical scientists can draw fairly dependable inferences based on their data. Other times, especially when it comes to the evolution of human psychology, these inferences will have to depend on premises that

¹For more on methodology in the historical sciences see, for example, Chapman and Wylie (2016); Currie (2016, 2018).

are themselves shaky or unconfirmed.² One central worry about much work in evolutionary psychology—a discipline aimed specifically at explaining human behavior from an evolutionary standpoint—is that the ‘gaps’ in these inferences can leave room for current cultural beliefs and biases to creep in.³

A way to further constrain these evolutionary narratives, and thus avoid some of these worries, is to use evolutionary modeling techniques. One branch of work along these lines comes from biology where researchers have used population genetic models to assess, for example, the conditions under which altruistic behaviors can evolve.⁴ Another related branch of work using evolutionary game theory starts with static representations of strategic social interactions and adds dynamical aspects to the model to see how behaviors in these scenarios evolve. More on this in the next section.

These sorts of models have been very successful, as we shall see, in explaining the evolution of moral behaviors such as altruism, cooperation, and apology. For example, they have been used to disconfirm evolutionary narratives about these behaviors that seemed coherent, thus improving on what we can do with sparse data and human reasoning alone.⁵ For this reason, I advocate here for the importance of modeling in the study of the evolution of moral psychology. But the use of these models raises another methodological issue. As mentioned, they address the evolution of behavior, not psychology. The reason is that, when it comes to evolution, behavior is where the rubber meets the road. The internal psychological organization of an organism only matters to its fitness inasmuch as it influences how the organism behaves. Of course, psychology and behavior are tightly connected since psychology evolves for the purpose of shaping effective behavior. But if we start with a model for a successful, fitness enhancing behavior, it takes an extra step to argue that such a model explains the evolution of a psychological trait.⁶

To see why, let’s consider an example. Suppose we look at a model of the prisoner’s dilemma and observe that under some set of conditions altruistic behavior is selected for. If we want to connect this observation to human psychology, we can argue that under these same conditions psychological traits that promote altruism will likewise be selected for. But the model does not tell us what those psychological traits are—anything that causes altruistic behavior might do. To make the connection to human psychology, then, we need to know what the candidate psychological traits responsible for altruism are ahead of time. This is the sense in which empirical data and models must be combined in studying the

²Another way to diagnose the issue here is that in some cases the historical evidence more severely underdetermines what hypotheses are drawn (Stanford, 2017).

³For an example, see the debate between evolutionary psychologists (Tooby and Cosmides, 1992; Buss, 1995) and social structural theorists (Eagly and Wood, 1999) on the origins of gender differences in psychology.

⁴The most influential work along these lines begins with Hamilton (1964) on kin selection. See Okasha (2013) for an overview.

⁵A famous example comes from John Maynard-Smith’s modeling work (Smith, 1964) which played a central role in refuting work on the group selection of altruism (Wynne-Edwards, 1962).

⁶Another traditional worry about the use of evolutionary models to understand behavior has to do with employing simplified representations to understand a complex world. Philosophers of modeling have worked extensively on this topic. (Weisberg (2012) surveys much of this work.) I do not address this worry here.

evolution of human psychology. One has to know empirically what psychological traits humans have, and what behaviors associate with those traits, in order to use models to explain them. Note that this makes it quite difficult to predict what sorts of psychological traits *will* evolve given a selective environment, even if we can generate good evidence about why certain psychological traits *did* evolve.

So, to use evolutionary models to study the evolution of human moral psychology, we need to know what moral psychology we are trying to explain. But this is not always simple. A first issue along these lines has to do with figuring out which aspects of moral psychology are biologically evolved. This can be tricky because much of human moral behavior is culturally shaped. Moral systems vary greatly across cultures, so we know there is no robustly hardwired biology that fully determines our moral psychology. These systems emerge on a cultural evolutionary timescale, and individual psychology is shaped by these cultural systems in multiple ways during the course of a lifetime.

On the other hand, it is also clear that much of our moral behavior is shaped by evolved biological tendencies. Some moral ‘non-nativists’ have argued that this is not so. One view is that human morality emerged as a side-effect of human-level rationality (Ayala, 1987). More common is the view that morality is entirely, or almost entirely, learned, not evolved (Prinz, 2008). Empirically, these views seem untenable. Joyce (2007) surveys literatures showing that, for example, 1) children show early, strong empathetic tendencies, 2) children develop a remarkably early understanding of norms, including the difference between conventional and moral norms, 3) primates and other animals display psychological traits connected to human morality, like prosocial pleasure and prosocial distress, 4) it is well established that emotions play a key role in moral decision making, and 5) sociopaths display perfectly good reasoning and learning skills, but nonetheless fail to engage in normal moral behavior as a result of emotional deficits. So to sum up, the evidence is that moral psychology is produced by three interacting processes on different timescales, 1) biological evolution, 2) cultural evolution, and 3) individual learning.⁷

How do we pick apart just those aspects that are well-explained by biological evolution? There is no general answer. Researchers have to argue it out on a case-to-case basis using the empirical evidence available. In many cases a full explanation of some moral psychological feature will have to draw on all three processes. For an emotion like guilt, for instance, we know that its production is highly culturally dependent, i.e., responsive to culturally evolved norms. On the other hand, it has clear biological components since sociopaths seem to lack normal capacities to feel guilt (Blair, 1995; Blair et al., 2005). Furthermore, various aspects of moral education play a role in the prevalence and strength of guilty feelings (Benedict, 2005). And, to add one more complication, in this case cultural evolutionary processes and biological evolutionary processes likely feedback on each other in a gene-culture co-evolutionary process (Gintis, 2003). For instance, culturally evolved moral norms for punishment create a selective scenario where guilt is more likely to emerge (O’Connor, 2016). At the same time, biological tendencies towards guilt shape which moral systems are culturally viable. The Baldwin effect from biology occurs when a learned behavior eventually

⁷For more on some of the ways these processes might interact see Nichols (2004); Sripada (2008).

becomes more and more innate.⁸ There may well have been various moral psychological traits that emerged on a cultural timescale, but have subsequently been stabilized by selection in this way.

In other words, it is complicated. These particular complications suggest another step for the methodology proposed above. One should see what behaviors are produced by a moral psychological trait, and use modeling to understand what scenarios support selection of those behaviors. Once this is done, though, a further question to ask, which must necessarily depend on empirical data is: are those selection pressures culturally created or not? Is the solution to those pressures a cultural system that shapes our psychology or a biological one? (Of course, this need not be a clean distinction, and, in fact, will probably not be.) Answering these questions may help researchers refine their models to produce more satisfactory explanations of moral psychological traits. (In addition, as we will see in section 3, returning to empirical data with plausible modeling results can often lead to refinements of other sorts as well.)

Another methodological issue, again regarding which psychological traits are candidates for evolutionary explanation, has to do with determining which traits are really the ‘moral’ ones. There are two sub-problems here. The first has to do with general problems regarding the identification and delineation of psychological traits. Which are really distinct from the others? For instance, emotions researchers mostly agree about the existence of the ‘big five’ emotions, but there is huge disagreement about how many other emotions there are, what these are, and how they are related.⁹ These debates get muddied by cultural influences on the development of moral psychological traits. For example, debates about the definitions of guilt and shame have been confused by the significant cultural variation in the production of these emotions (Wong and Tsai, 2007).¹⁰

The second sub-problem, which has mostly arisen in philosophy, has to do with identifying which behaviors genuinely count as ‘moral’. Should we be focusing mostly on prosocial instincts? Emotions? Or higher level abilities unique to humans? Theory of mind? The ability to make moral judgments? I follow Churchland (2011) in being generally uninterested in trying to disambiguate truly moral behavior from behavior in the animal kingdom with moral character. The most compelling, naturalistic accounts admit that our moral psychology involves many features shared with other animals, and also a number of features, like theory of mind and language use, that few species (or just humans) exhibit. As such, there are a large variety of diverse moral psychological traits one might want to make sense of in the light of evolution.

Before completing this section, I think it will be helpful to develop an (incomplete) list of aspects of human psychology that contribute to moral behavior and thus might be appropriate targets of evolutionary explanations of moral psychology. Note that this list will include concepts that overlap, and others that are ill-defined. Many of the things on this

⁸This was first proposed by Baldwin (1896).

⁹Compare, for example, the eight emotions identified in Plutchik (1991), and organized as opposites, with the 27 emotions identified by Cowen and Keltner (2017).

¹⁰Benedict (2005) gave an influential, early analysis of these cultural differences.

list can be tackled to some degree of success using the kind of methodology I advocate here, though. (Indeed, many of them already have been addressed from an evolutionary point of view already, but it is beyond the scope of this paper to fully survey this literature.)

There are positive prosocial emotions, instincts, and feelings including social trust, love, empathy, sympathy, compassion, and gratitude. There are also negative prosocial feelings and emotions such as guilt, shame, and embarrassment. Empathy can also lead to prosocial distress, for instance if one sees a loved one in pain. In addition, we have a set of emotions that contribute to punishment and norm guidance: moral anger, urges for retribution, contempt, and disgust.

Besides the emotional side of things, Joyce (2007) argues that the ability to judge something as morally wrong is key to human morality. In other words, if we did not have the capacity to think thoughts like, ‘Jill ought not have done that’ or ‘Brian deserves punishment for breaking rule X’, we simply would not have full human morality. As Joyce points out, language is a prerequisite for this sort of judgment. He further argues that this judgment is a trait that has been biologically selected for, rather than emerging from other cognitive abilities. (On his account the benefits of this trait come from the fact that moral judgments solidify prosocial behaviors.)

Another such trait is moral projection. We often feel that moral norms like ‘do not kill innocent people’ are not just wrong for us, or wrong in our culture, but are simply wrong in a freestanding, objective way.¹¹ Stanford (2018) develops an evolutionary narrative where this sort of trait benefits individuals by improving moral coordination in that individuals who judge X right are also incentivized to perform X. For Joyce (2007) the evolution of moral projection was a key step in our ability to make genuinely moral judgments.

Churchland (2011) and others have argued that theory of mind—the ability to understand others as having minds like our own—is a key feature of our moral psychology. Without it, we cannot judge intent, and thus cannot make standard moral judgments. In addition, theory of mind plays a role in empathy by allowing us to conceptualize feelings in others that we do not feel ourselves.

Various scholars have noted that humans are fascinated by gossip. Emler (1990) argues that humans spend 60-70% of conversations on gossip and ‘reputation management’. This trait is important to our moral systems, where individuals use knowledge about past moral behavior to choose interactive partners (Joyce, 2007). It also plays a role in punishment and reciprocation.¹²

Some have argued that humans, and maybe even non-human primates have an innate aversion to inequity, or a ‘taste for fairness’ (Fehr and Schmidt, 1999). If so, this is likely connected to the prevalence of norms for fairness and justice across human groups. In a reflection of the methodological worries raised above, there has been debate in economics

¹¹For instance, in one study Amish school children were asked ‘if God said it was OK to work on Sunday would it be?’ They agreed. When stealing was substituted for ‘working on Sunday’, they largely did not agree (Nucci, 1985, 2001).

¹²Some have even argued that gossip is so important to humans, it drove the evolution of human language (Aiello and Dunbar, 1993).

about whether human behavior is better explained by inequity aversion (which presumably is at least partially a biological trait) or by social norms for equity (which emerge on the cultural timescale) (Fehr and Schmidt, 1999; Binmore and Shaked, 2010; Fehr and Schmidt, 2010). Most seem to agree that there is at least some innate inequity aversion of this sort.

Norms are rules in human groups which may vary cross-culturally, and which determine what behaviors ‘ought’ or ‘ought not’ be performed. While not all norms have a moral character, there seem to be psychological traits dedicated to norm governed behavior which play a strong role in human moral systems. As Sripada and Stich (2005) outline, one of these traits involves internalizing norms, which is a process whereby individuals come to believe they should follow them regardless of expectations for sanction. This psychological trait plays a key role in moral systems (Joyce, 2007; Stanford, 2018).

Associated with norm governance and norm psychology are learning behaviors that allow groups of humans to adopt group-specific norms and conventions. First is the simple fact that humans have evolved to be extraordinary social imitators. Beyond this, humans show biases towards learning common behaviors, as well as behaviors displayed by particularly prominent or successful group members (Boyd and Richerson, 1988). These learning abilities are key to the stability of our moral norms.

I will mention one more psychological trait which is certainly important in the production of our moral behavior, even if we might not want it to be. Humans show strong psychological tendencies towards in-group bias, i.e., treating in-group members better than out-group members.¹³ This tendency likely evolved to regulate human prosocial behavior.¹⁴

3. Modeling Guilt

The methodology I am advocating for, as noted above, involves 1) choosing some aspect of human moral psychology, 2) carefully investigating which behaviors it causes under what conditions, 3) using evolutionary models to explain the psychological trait by explaining the evolution of the behaviors it causes, and 4) returning to empirical data to refine the model, and further develop the narrative one can draw from it. At this point, I move on to the second part of the paper, which involves giving an in-depth example of how this methodology can work. In particular, I describe how empirical and modeling work complement each other in an investigation of the evolution of guilt. Before doing so, though, I introduce basics from the mathematical frameworks used in this investigation—game theory and evolutionary game theory.

¹³Evidence for this bias comes from experiments in ‘minimal group paradigm’. These experiments involve the formation of often arbitrary groups, i.e., by coin flip, and find that participants nonetheless use group structure to determine prosocial behavior (Tajfel, 1970).

¹⁴I will not describe the evolutionary models most germane to understanding this aspect of human behavior, but see, for example Boyd and Richerson (2005); O’Connor (2019).

3.1. Game theory and Evolutionary Game Theory

Game theory is a branch of mathematics used to model *strategic* behavior—behavior that involves multiple actors where each individual cares about what the others do.¹⁵ A *game* is a model of some such situation involving three elements: *players*, or who is involved, *strategies*, what actions they can take, and *payoffs*, or what each player gets for each set of strategies played.¹⁶

Traditional approaches to game theory model various strategic scenarios and use these models to predict or explain human behavior via assumptions about rationality. For example, by showing that taking behavior X will always yield a lower payoff than behavior Y in some game, a theorist can generate the prediction that humans will not take behavior X in analogous scenarios. Why use game theory here? Moral behavior occurs in the social realm, i.e., it evolves as a solution to social problems. Games are just those scenarios where actors interact socially in ways that matter to their fitness, and so just those scenarios morality emerged to regulate.

Evolutionary game theory, first developed in biology, and then imported into the social sciences, asks: what sorts of strategic behaviors are likely to evolve among group of agents playing games?¹⁷ These models take a population in some strategic scenario and add what are called *dynamics*, or rules for how the population will evolve. The most widely used class of dynamics in evolutionary game theory assumes that strategies that yield higher payoffs will expand in a population, while those that yield lower payoffs will contract. It should be clear why evolutionary game theory is a useful framework here. It is intended to model just those sorts of scenarios most likely to inform moral psychology: situations where individuals evolve social behavior.

3.2. Guilt

Guilt is a puzzling emotion in many ways. Some emotions, like fear, have clear fitness benefits (avoiding danger). Guilt is associated with a number of behaviors that seem straightforwardly detrimental, such as altruistic behavior, acceptance or seeking of punishment, making costly reparations, and even self punishment. For this reason, and because of its relevance to ethical behavior, a number of philosophers have gotten interested in the evolution of guilt (Joyce, 2007; Deem and Ramsey, 2016b,a; O'Connor, 2016; Rosenstock and O'Connor, 2018). The goal here will be to demonstrate how the method described earlier in the chapter can help explain some aspects of the evolution of guilt. I will draw on previous work from O'Connor (2016); Rosenstock and O'Connor (2018).

Let's start, as suggested, by looking at empirical work to identify what sorts of strategic behaviors are associated with feelings of guilt. These can be grouped into three rough categories. First, guilt seems to prevent social transgression and is thus associated with

¹⁵This framework originated with the work of Morgenstern and Von Neumann (1953).

¹⁶Usually games also define a fourth element, *information*, or what each player knows about the game. I ignore this element here since it is not usually relevant to evolutionary analyses.

¹⁷This field originated with Smith and Price (1973), though precursor work occurred in economics.

prosocial behaviors such as cooperation and altruism (Tangney et al., 1996; Regan et al., 1972; Ketelaar, 2006). Second, guilt leads to a suite of behaviors after social transgression, such as apology, acceptance of punishment, self punishment, and costly reparation (Silfver, 2007; Ohtsubo and Watanabe, 2009; Nelissen and Zeelenberg, 2009). Last, expressions of guilt seems to influence punishing and judging responses by other group members. In particular, individuals who express feelings of guilt and remorse are more likely to be judged guilty (in the sense that they did the deed), but also more likely to be forgiven and punished less harshly (Eisenberg et al., 1997; Gold and Weiner, 2000; Fischbacher and Utikal, 2013).

Let us focus, for the purposes of this chapter, on the second suite of behaviors associated with guilt—those that occur after transgression. These, as pointed out, lead to immediate fitness costs and so seem particularly in need of an evolutionary explanation. It is clear that these sorts of reparative behaviors play a role in avoiding ostracism after transgression (Joyce, 2007). In order to better understand how this might work, let us now turn to an evolutionary model.

In the next section, I will describe at much greater length the *prisoner's dilemma* and discuss how it can be used in understanding the evolution of moral emotions generally. For now, it will serve to know that this is a model where actors have two options—to behave altruistically by playing ‘cooperate’ or to behave selfishly by playing ‘defect’. There is always a payoff incentive to choose defection, which, as we will see, raises the question of why altruism has so often evolved. In the *iterated* version of the prisoner’s dilemma actors play again and again over the course of many round. This makes it a particularly good model for exploring behaviors associated with reparation and apology—actors engage in a long enough interaction to have time to potentially rend and repair their relationship. (It also is a good model of early human interaction, since the group structures of early humans led to repeated interactions with group members.)

In particular, let’s consider a version of this game where actors sometimes make mistakes. An actor might usually behave altruistically by playing cooperate, but sometimes defect either accidentally, or due to exigent circumstances. This might represent the sort of situation where a usually prosocial individual transgresses against ethical norms. Let’s suppose further that actors tend to use reciprocating strategies. We will discuss these further in the next section, but the basic idea is to meet cooperative behavior with cooperation and selfish behavior with defection. Accidental defection of the sort just described causes problems for these sorts of reciprocating strategies. In particular, reciprocators can get stuck in cycles of retribution where they continue to defect on each other, despite have underlying altruistic tendencies.¹⁸ Two altruistically-inclined individuals can lose the benefits of mutual altruism as a result of an error plus reciprocation.

One way to get around this issue is through apology. Models have shown that strategies where actors apologize after defection, and accept the apologies of others, can evolve since they avoid the payoff loss associated with mutual negative reciprocation.¹⁹ In showing how

¹⁸There are various ways to solve this problem (Axelrod and Hamilton, 1981; Nowak and Sigmund, 1993). As we will see, apology is one of these.

¹⁹See Okamoto and Matsumura (2000); Ohtsubo and Watanabe (2009); Ho (2012); Han et al. (2013);

apology and reparation can evolve, these models can also show how guilt might be selected for in order to mediate apology and reparation.

One of the most important take-aways from these models is that this sort of apology can only work if it is hard to fake, costly, or both. The reason is that unless there is some mechanism guaranteeing the apology, fake apologizers can take advantage of trusting group members. If some individuals forgive and forget after being defected on, fakers can keep on saying ‘I’m sorry’ and defecting in the next round. This prevents apology from evolving to promote altruism.

Let us talk through these ways of stabilizing apology, and discuss how they might connect up with the evolution of guilt. Frank (1988), in very influential work, suggests that moral emotions evolved for a signaling purpose. Guilt leads to altruistic behavior, and is also hard to fake (because it is an emotion). Thus guilt can allow individuals to choose altruistic partners and avoid defectors. Something similar might work for apology too. If we can tell which individuals are offering sincere apologies, because we can simply see that they feel guilty, then we can only forgive those who have real, altruistic tendencies. This is the sense in which apology can evolve if it is hard to fake (O’Connor, 2016).

If we return to the empirical literature, though, we see that this model does not quite fit the bill with respect to explaining the evolution of guilt. In particular, unlike some emotions (like anger and fear), there are no stereotypical body and facial postures associated with guilt. This means that others cannot identify a guilty individual just by looking. In other words, guilt does not display the features we would expect from an emotion that evolved for the purposes of signaling altruistic intent (Deem and Ramsey, 2016b).

What about costly apology? In models of apology and reparations, costs stymie fakers by making apologies not worth their while. This is because apology yields different benefits for those who intend to defect and those who intend to re-enter a long, cooperative engagement. Imagine you are the latter type of individual. You are being shunned from a group for a social transgression. If you are willing to issue a costly apology by accepting punishment, punishing yourself, or paying the one you transgressed against, your group members will start cooperating with you again. It should be well worth your while to pay even a large cost to re-enter the fold. Doing so earns you a lifetime of receiving the benefits of mutual altruism. Now imagine you are a defector and plan to transgress immediately after apologizing. Paying a cost to apologize will generally not be worth your while. As soon as you defect, you’ll be on the outs again, and have to pay another cost to find a cooperative partner.

The success of costs in facilitating the evolution of apology seems to tell us something important about the evolution of guilt. There is a reason that guilt leads to a number of behaviors that are individually costly, which is that without these costs it would not be an effective emotion for promoting apology and reparation.

We can again turn to the empirical literature to tune this model further. Remember that expressions of guilt seem to decrease punishing behaviors by group members. Although it

Pereira et al. (2016); O’Connor (2016); Rosenstock and O’Connor (2018). In Pereira et al. (2016); O’Connor (2016); Rosenstock and O’Connor (2018) authors connect these models to the evolution of guilt. In addition, Pereira et al. (2017a,b) take a different tack in showing how guilt can evolve by causing self-punishment.

does not make sense to posit that guilt evolved solely for a signaling purpose, this literature tells us that it does seem to play some signaling role in apology. In Rosenstock and O'Connor (2018) we explore in more detail why this might be. In particular, we go back to the models and show that the high costs necessary to guarantee apology have a downside. As costs, they directly decrease the fitness of those who feel guilt and issue costly apologies. This means that even when guilt and apology are evolutionarily viable as a result of these costs, they may be unlikely to evolve.²⁰ However, if expressions of guilt are even a bit trustworthy, this significantly decreases the necessary costs, and increases the likelihood that apology and guilt can evolve.²¹ In other words, these slightly more complicated models show how guilt and expressions of guilt might emerge to regulate costly apology, while reducing the necessary costs as much as possible.

This work provides only a partial explanation of guilt. Notice that the models invoked are evolutionary models, but they are not gene-culture coevolutionary models. And, as mentioned, guilt likely is the result of gene-culture coevolution. Furthermore, they do not explicitly represent or account for the role of culture in the production of guilt, or for the cultural variability in what sorts of transgression lead to apology and reparation. (It is not always failures of altruism in the real world, as we know.) As suggested, a good next step is to return to the empirical literature to hone the partial explanation developed here.

4. Models

We have now seen how empirical investigation and evolutionary modeling can be used in concert to elucidate the evolution of a particular moral psychological trait. The goal of this section is to sketch out the directions in which such a methodology might be further applied. I overview of the main sets of models used to explain the evolution of moral behavior. In other words, I lay out here the models that have the greatest potential to tell us something about the evolution of the panoply of moral psychological traits listed in section 2.

4.1. Altruism and the Prisoner's Dilemma

The prisoner's dilemma, mentioned above, is probably the most famous, and certainly the most widely studied, game in game theory. We have seen how it might play a role in modeling the evolution of guilt. Let's now fill in the details and try to understand more generally how it can be used to model the evolution of moral psychology.

The motivating story is that two prisoners are being asked to rat each other out. If they both stay silent they each get a short jail sentence. If they both rat, they each get a long one. If only one rats on the other, who stays silent, the rat gets to go free, while the other one serves an even longer sentence.

²⁰We show that the basins of attraction for guilt-prone strategies that are generally cooperative, retributive, apologetic, and trusting of apologies are relatively small when costs are high.

²¹This combination of costs and honest signals was inspired by Huttegger et al. (2015) who show how biological signals that are somewhat hard to fake can reduce the costs necessary to ensure they are trustworthy.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	2,2	0,3
	Defect	3,0	1,1

Figure 1: A payoff table of the prisoner’s dilemma. There are two players, each of whom choose to cooperate or defect. Payoffs are listed with player 1 first.

To turn this scenario into a game we must define the formal elements listed in section 3.1. The players here are the two prisoners. Their strategies, to rat or stay silent, are usually labeled ‘defect’ and ‘cooperate’, as noted earlier. Their payoffs in terms of *utility* gained for each outcome described, are listed in figure 1. Each entry in the table shows the payoffs for some combination of strategies with player 1 listed first. Utility here is meant to track preference or benefit for each player. The numbers are chosen arbitrarily, but the game will be a prisoner’s dilemma as long as they retain their ordering.

We can now see why this game is referred to as a ‘dilemma’, and why it has generated such a significant literature. Both players prefer mutual cooperation over mutual defection, because the first yields payoffs of 2 and the second 1. But regardless of what the other player is doing, each player prefers to defect rather than cooperate. If your opponent cooperates, you receive a 3 for defection and only a 2 for cooperation. If your opponent cooperates you receive a 1 for defection and only 0 for cooperation. This means that defect-defect is the only *Nash equilibrium* of the game—the only pairing of strategies where neither player wants to change what they are doing. (This is a central predictive concept in game theory and evolutionary game theory. Because no one wants to deviate, Nash equilibria tend to be stable, and also tend to emerge in evolutionary models.) Importantly, this payoff structure also means that the ‘cooperate’ strategy can represent altruistic behavior. Taking this strategy always involves lowering your own payoff while simultaneously increasing the payoff of your partner. Thus the prisoner’s dilemma serves as a model for a wide swath of human moral behavior, including those we discussed in the last section on guilt.

At first glance, the prisoner’s dilemma model seems to predict non-altruistic behavior. It is always better to defect than to cooperate. But, of course, humans engage in altruistic behavior all the time. (To give one cute example, I just watched an internet video where a group of men spent over an hour cutting a humpback whale free from a fishing net.) How do we explain this discrepancy? Evolutionary game theory provides one avenue. The suggestion is that a population of biological agents might evolve altruistic traits, even though on the face of it altruism seems irrational. And, regarding moral psychology, such an agent might thus come under selection pressure for psychological traits that lead to altruistic behavior.

The key to evolving altruism in a scenario where actors face a prisoner’s dilemma is *correlated interaction* between the strategies. Whenever cooperators meet cooperators and defectors meet defectors often enough, cooperation can win out. We can think of correlation as changing the payoff table in figure 1 to one where only the top left and bottom right

entries are available. And in such a case, cooperation does strictly better.

So what are the mechanisms that lead to correlated interaction of this sort? There are a number that have been identified.²² I will focus here on kin selection, direct and indirect reciprocity, and punishment as the ones most pertinent to the evolution of moral psychology.

Kin selection explains the emergence of altruistic behavior across the animal kingdom.²³ Kin, of course, are more likely to share genetics. If altruists interact with their own kin, then the benefits they give tend to fall on other altruists. This can make altruistic genes more successful than selfish genes. This also helps explain why throughout the animal kingdom, organisms are most likely to give care to their own young, siblings, and family groups. Most theorists think that the first origins of moral psychology—positive, caring feelings directed towards kin, and distress at pain or separation from kin—were selected for via kin selection in our ancestors (Joyce, 2007; James, 2010; Churchland, 2011).

Reciprocal altruism is common in human moral systems, and also present in a few species of non-human animal.²⁴ The basic idea behind reciprocity is that players tend to cooperate with those who have cooperated in the past, and defect against those who have defected. The net outcome is that cooperators end up fitter, despite their altruistic behavior, because they are treated altruistically. As Trivers (1971) first showed in biology, this means that reciprocally altruistic behavior can evolve.

Direct reciprocity occurs when individuals choose their strategies based on how a partner has treated them in the past. This sort of behavior is often modeled using the *iterated prisoner's dilemma* mentioned in the last section, where two actors play the prisoner's dilemma again and again. One example of a reciprocating strategy made famous by Axelrod and Hamilton (1981) is tit-for-tat (TFT), where each player starts cooperative and then takes the strategy that their opponent did in the last round. This strategy does well against defectors, cooperators, and itself. And (though things are a bit complicated), direct reciprocity in general, and the tit-for-tat strategy itself can evolve (Trivers, 1971; Axelrod and Hamilton, 1981; Nowak, 2006a). Another such strategy is the grim trigger, where actors cooperated until a partner defects, and then defect forever.²⁵ Notice that these evolutionary models lend themselves to explaining the emergence of moral psychology surrounding retribution, contempt, and moral anger, and the connection between this psychology and altruism in human groups.

Indirect reciprocity involves reciprocating on the basis of transmitted information. If John defects against me today, and I defect against him tomorrow, this is direct reciprocity. If John defects against me today, and I report it to you, and you defect against him tomorrow, this is indirect reciprocity. Reputation tracking of this sort can help promote the emergence of altruistic strategies (Alexander, 2017; Nowak and Sigmund, 1998). This sort of model,

²²See Nowak (2006b) for a description of five major mechanisms that have played significant roles in the literature on the evolution of altruism: kin selection, group selection, direct reciprocity, indirect reciprocity, and network reciprocity.

²³The theory of kin selection was introduced by Hamilton (1964), and has been developed extensively since then, as in Dawkins (2006); Michod (1982); Grafen (1984).

²⁴Famously, in the vampire bat (Wilkinson, 1984).

²⁵See Sachs et al. (2004) for a good overview of the evolution of reciprocal altruism.

notice, lends itself to explaining the moral psychology behind gossip and reputation tracking.

The last thing to discuss here is punishment, though this route to the evolution of altruism gets a bit complicated. In some sense reciprocation is a type of punishment, but we might also consider actors who actively lower the payoffs of defectors, rather than just defecting on them in the future. The idea here is that if defectors are regularly punished for defecting, altruists can get a leg up. We do, of course, see moralistic punishment, including for altruism failures, across human societies (Fehr and Gächter, 2002; Henrich et al., 2006). And we know that such punishment can help maintain altruistic behavior (Yamagishi, 1986; Ostrom et al., 1992). But the evolution of punishing behavior is complicated by the fact that punishing others is always at least a little costly, leading to what is termed the ‘second order free rider problem’. Why would I sacrifice my fitness to punish a violator? A number of solutions have been proposed, though discussing them is beyond the scope of this paper.²⁶ Note that these models may serve to help us further understand the moral psychology behind punishment—moral anger, indignation, etc.

To sum up, the prisoner’s dilemma helps us frame one of the most significant social problems facing early humans—why be altruistic when selfishness pays off? It also helps us answer one of the most pressing puzzles in the evolution of moral psychology—why are humans so psychologically inclined towards altruism? As we saw, modeling work on the evolution of altruism is actually relevant to many aspects of moral psychology including prosocial instincts and emotions, as well as the psychology of retribution and punishment.

Public Goods, Trust, and Punishment

Before continuing, I would like to briefly mention another game that shares some character with the prisoner’s dilemma. In *public goods games*, actors have the option to contribute some amount of resource to a public good. This could represent a group of individuals clearing a meadow for public grazing, or building a town hall. The total amount contributed is multiplied by some factor, and then equally divided among the contributors. This captures the idea that there is a benefit to joint production—a group of people together can produce marginally more than they could alone. It also captures the idea that in any such situation there is a temptation to shirk by not contributing, but still enjoying the fruits of others’ labor. The Nash equilibrium of this game is that each player contributes \$0, even though they all would do much better by each contributing the maximum (or anything more than 0). For this reason, public goods games present another sort of social dilemma, where giving represents altruism in that any increase of personal contribution increases the payoffs of group members at the expense of one’s own payoff.

This game has been widely studied experimentally, and punishment and reward have been shown to improve contributions (Fehr and Gächter, 2000). From an evolutionary perspective, punishment, reward, and reputation tracking can lead to the evolution of altruistic behavior in public goods games (Hauert, 2010).²⁷ Like various evolutionary models of the prisoner’s

²⁶For example, see Panchanathan and Boyd (2004). See Frank (2003) as well.

²⁷See Santos et al. (2008) for more on how diversity of population traits influences these processes.

		Player 2	
		Stag	Hare
Player 1	Stag	3,3	0,2
	Hare	2,0	2,2

Figure 2: A payoff table of the stag hunt. There are two players, each of whom choose to hunt stag or hare. Payoffs are listed with player 1 first.

dilemma, these models may be able to tell us something about moral anger, retribution, gratitude, and the moral psychology behind altruism. In addition, this model has been taken as informative of the evolution of social trust (Churchland, 2011).

4.2. Cooperation, Coordination, and the Stag Hunt

The prisoner’s dilemma gets a lot of attention, but altruism is not the only prosocial behavior that has emerged in human groups. We now investigate another type of strategic scenario often faced by humans—where cooperation or joint action is mutually beneficial, but risky.

This sort of scenario is typically modeled by a game called the *stag hunt*. The motivating story is that two hunters have the option to either hunt stag or hare. If they hunt hare, they each gather a dependable, small amount of meat. If they hunt stag, they have the opportunity to gather much more by cooperating—half a stag each. But they will only catch the stag if both work together. In other words, there is a risk to cooperating because one’s partner might decide to work alone.²⁸ Figure 2 shows this game. The payoffs are 3 for joint stag hunting, 2 for hunting hare, and 0 for hunting stag when your partner hunts hare. In influential work, Skyrms (2004) uses this game as a simple representation of the social contract—hare hunting represents a state of nature, and stag hunting a social organization. In general, the game can capture many scenarios where by working together humans can generate surplus resources, but doing so opens them up to risk.

Unlike the prisoner’s dilemma, this game has two Nash equilibria: the strategy sets where both actors hunt stag, or both hunt hare. In other words, either a cooperative, mutually-beneficial outcome, or an uncooperative outcome is possible from an evolutionary standpoint. From a naive point of view it might seem obvious which should evolve—stag hunting, right? It yields a greater payoff.

Things are not so straightforward. This is because stag hunting, while clearly better from a fitness standpoint, is also more risky. Should your partner hunt hare, you end up with nothing. The result is that in evolutionary scenarios, it is often the case that populations are more likely to move from stag hunting to hare hunting than the reverse (Skyrms, 2004).

As with the prisoner’s dilemma, though, mechanisms that correlate interaction among stag hunters can lead to the emergence and stabilization of prosocial behavior. There are

²⁸This motivating story is from Rousseau (1984).

a few ways this can happen. If individuals can choose their partners, stag hunters will tend to pick each other, leading to high payoffs for stag hunting. If individuals are able to communicate with each other, they can use even very simple signals to coordinate on stag hunting. And in populations with a network structure, so that individuals tend to keep meeting the same neighbors for interaction, stag hunting can emerge spontaneously (Skyrms, 2004; Alexander, 2009).

What can the stag hunt tell us about moral psychology? This particular model seems especially useful in thinking about social trust. An individual in a cooperative group will do poorly if they are too worried about social risks (and thus hunt for hare while the rest hunt stag). The stag hunt may also tell us something about emotions like guilt or shame. Individuals may be temporarily tempted to shirk social duties for their own benefit. If this leads to poor payoffs in the end, emotions like guilt that decrease the chances of such shirking are directly beneficial (O'Connor, 2016).

4.3. Justice and Bargaining Games

The two models thus far examine scenarios where actors can choose prosocial or antisocial types of behavior. Let us now consider a set of models that are less commonly drawn upon in thinking about moral psychology—bargaining games. These are games that capture scenarios where humans must divide some resource among themselves. Note that these scenarios are ubiquitous. Whenever social groups obtain resources, including food, tools, building materials, they must decide how to divide them. In addition, joint production involves another ubiquitous sort of bargaining—individuals must decide who will do how much of the work involved.

There are several games used to represent such cases. I will mention two, each of which correspond to different assumptions about control of the resource. In the *Nash bargaining game*, two actors make demands for a portion of a resource. If these demands are compatible, they each get what they wanted. If they are too aggressive and over-demand the resource, though, each gets a poor payoff.²⁹ For simplicity sake, assume a resource of size 10, where each individual can make demands of 3, 5, or 7. Figure 3 shows the payoff table for this game. As is evident, when the payoffs sum to 10 or less, each actor gets what they demanded, otherwise they get 0.

There are three Nash equilibria of this game—the three strategy pairings where the resources is exactly divided (3 v 7, 5 v 5, and 7 v 3). One of the central evolutionary findings in this model is that the ‘fair’ demand, where each player demands exactly half the resource, is more likely to evolve than populations with the other two demands (Young, 1993; Skyrms, 2004).³⁰ This has been taken to help explain justice—both the cultural practice, and inequity aversion in humans (Skyrms, 2004; Alexander, 2009).

The ultimatum game is a variation on the Nash bargaining game that yields very different predictions. In this game, one actor controls the resource, and choose how much of it to

²⁹This game was introduced by Nash (1950). It is sometimes called ‘divide the dollar’, ‘divide the pie’, or just the bargaining game.

³⁰Though, see O'Connor (2019); D'Arms et al. (1998).

		Player 2		
		Low	Med	High
Player 1	Low	3,3	3,5	3,7
	Med	5,3	5,5	0,0
	High	7,3	0,0	0,0

Figure 3: A payoff table of the Nash demand game. There are two players, each of whom choose one of three bargaining demands. Payoffs are listed with player 1 first.

offer to a partner. The partner’s choices are then to either accept what they are offered or to reject. If they reject, neither partner gets anything.³¹ The rational choice prediction is that the first player should offer as little as possible, on the expectation that the second player will prefer any offer to nothing, and thus accept. In experimental set-ups, in fact, individuals make fairly high offers, and reject offers that are too low, though the details vary cross culturally (Güth et al., 1982; Henrich et al., 2006). This amounts to a kind of costly punishment—the second player in the game is willing to lower her payoff in order to lower the payoff of the first player.

Evolutionary models of the ultimatum bargaining game can help explain this seemingly irrational behavior. Populations can evolve where actors make high offers, and reject low ones (Gale et al., 1995; Harms, 1997; Skyrms, 2014). These models may be informative in understanding the moral psychology of fairness, as well as retribution.

4.4. Coordination and Norms

Ethical systems tend to display in-group similarity and between-group variability. For instance, as just mentioned, there is cross-cultural variation in how individuals play the ultimatum game, but relatively less variation within each culture (Henrich et al., 2006). It is even the case that ethical behaviors one culture finds abhorrent—infanticide, out-group homicide, honor killings—will be the norm elsewhere (Joyce, 2007).

I will describe one last sort of model in this section—coordination games. These are games where individuals choose one of two actions, and where their ultimate goal is to coordinate. A classic example involves choosing which side of the road to drive on. Drivers can choose the left or the right. They generally care much less about which choice they make, and much more about making the same choice as other drivers. Figure 4 shows the simplest possible coordination game, where actors get a payoff of 1 for choosing the same acts and 0 for choosing different ones. The two Nash equilibria of this game are (unsurprisingly) the strategies where both actors do the same thing.

The need to coordinate behavior is ubiquitous in human groups. In the moral sphere, there are serious problems for groups who do not share expectations about which behaviors are required and which forbidden (Stanford, 2018). The coordination game may be a helpful

³¹This model was introduced by Güth et al. (1982).

		Player 2	
		A	B
Player 1	A	1,1	0,0
	B	0,0	1,1

Figure 4: A payoff table of a simple coordination game. There are two players, each of whom chooses A or B. Payoffs are listed with player 1 first.

model in thinking about the aspects of norm psychology that get groups of humans behaving in coordinated ways, including aspects of social learning like conformity bias.

Before finishing this section, I would like to flag for the reader that the strategic situations outlined here are surely not the only ones that have been relevant to human social lives. There are other models that might be used to help inform the evolution of moral psychology. In addition, there are aspects of our moral psychology that have emerged in response to problems and opportunities that are not well modeled by games. To give an example, most cultures have incest taboos. These are culturally shaped, but humans also have a psychological tendency—dubbed the Westermarck effect—to avoid those we grew up with as romantic partners (Shepher, 1983). This tendency likely evolved because incest leads to issues with deleterious gene mutations. In other words, this moral psychological trait solves a social problem, but not a strategic one.

5. Conclusion

Section 3 gave an example of how the sort of methodology I advocate here can work. It starts by carefully identifying a moral psychological trait that might benefit from evolutionary explanation. It next proceeds by drawing upon empirical literature to determine which behaviors are associated with this trait. One can then use evolutionary models, perhaps drawing on the ones described in section 4, to develop a clear understanding of what leads to the selection of such behaviors, and thus what environments might lead to the selection of the psychological trait associated with that behavior. Next one can use empirical literature to assess the success of the potential explanation developed. In the case of guilt, as we saw, this sort of validation led us to downplay the importance of the Frank (1988) explanation of guilt, and put more weight on models of costly apology.

Of course, this sort of process will not work for all moral psychological traits. These traits, as we saw in section 2, are highly diverse. The evolution of the psychology behind social imitation for example requires a very different sort of explanation than the evolution of guilt. Furthermore, as noted, some moral psychology did not evolve to regulate strategic scenarios, and so may require a very different sort of methodology. The goal here is not to be dogmatic. Rather the goal is to give some good guidelines for further work on the evolution of moral psychology that makes use of the best epistemic tools available.

Acknowledgments

Removed for review

References

- Aiello, Leslie C and Robin IM Dunbar (1993). "Neocortex size, group size, and the evolution of language." *Current anthropology*, 34(2), 184–193.
- Alexander, J McKenzie (2009). "The structural evolution of morality."
- Alexander, Richard (2017). *The biology of moral systems*. Routledge.
- Axelrod, Robert and William Donald Hamilton (1981). "The evolution of cooperation." *science*, 211(4489), 1390–1396.
- Ayala, Francisco J (1987). "The biological roots of morality." *Biology and Philosophy*, 2(3), 235–252.
- Baldwin, J Mark (1896). "A new factor in evolution." *The american naturalist*, 30(354), 441–451.
- Benedict, Ruth (2005). *The chrysanthemum and the sword: Patterns of Japanese culture*. Houghton Mifflin Harcourt.
- Binmore, Kenneth and Avner Shaked (2010). "Experimental Economics: Where Next? Rejoinder." *Journal of economic behavior & organization*, 73(1), 120–121.
- Blair, James, Derek Mitchell, and Karina Blair (2005). *The psychopath: Emotion and the brain*. Blackwell Publishing.
- Blair, Robert James Richard (1995). "A cognitive developmental approach to morality: Investigating the psychopath." *Cognition*, 57(1), 1–29.
- Boyd, Robert and Peter J Richerson (1988). *Culture and the evolutionary process*. University of Chicago press.
- Boyd, Robert and Peter J Richerson (2005). *The origin and evolution of cultures*. Oxford University Press.
- Buss, David M (1995). "Evolutionary psychology: A new paradigm for psychological science." *Psychological inquiry*, 6(1), 1–30.
- Chapman, Robert and Alison Wylie (2016). *Evidential reasoning in archaeology*. Bloomsbury Publishing.
- Churchland, Patricia S (2011). *Braintrust: What neuroscience tells us about morality*. Princeton University Press.
- Cowen, Alan S. and Dacher Keltner (2017). "Self-report captures 27 distinct categories of emotion bridged by continuous gradients." *Proceedings of the National Academy of Sciences*.
- Currie, Adrian (2016). "Hot-Blooded Gluttons: Dependency, Coherence, and Method in the Historical Sciences." *The British Journal for the Philosophy of Science*, 68(4), 929–952.
- Currie, Adrian (2018). *Rock, Bone, and Ruin: An Optimist's Guide to the Historical Sciences*. MIT Press.
- D'Arms, Justin, Robert Batterman, and Krzysztof Górný (1998). "Game theoretic explanations and the evolution of justice." *Philosophy of Science*, 65(1), 76–102.
- Dawkins, Richard (2006). "The selfish gene: with a new introduction by the author." *UK: Oxford University Press. (Originally published in 1976)*.
- De Waal, Frans (1996). *Good natured*. Harvard University Press.
- Deem, Michael and Grant Ramsey (2016a). "The Evolutionary Puzzle of Guilt: Individual or Group Selection?." *Emotion Researcher: ISRE's Sourcebook for Research on Emotion and Affect*. Ed. Andrea Scarantino.
- Deem, Michael J and Grant Ramsey (2016b). "Guilt by association?." *Philosophical Psychology*, 29(4), 570–585.
- Eagly, Alice H and Wendy Wood (1999). "The origins of sex differences in human behavior: Evolved dispositions versus social roles.." *American psychologist*, 54(6), 408.
- Eisenberg, Theodore, Stephen P Garvey, and Martin T Wells (1997). "But was he sorry? The role of remorse in capital sentencing." *Cornell L. Rev.*, 83, 1599–1637.
- Emler, Nicholas (1990). "A social psychology of reputation." *European review of social psychology*, 1(1), 171–193.

- Fehr, Ernst and Simon Gächter (2000). “Cooperation and punishment in public goods experiments.” *American Economic Review*, 90(4), 980–994.
- Fehr, Ernst and Simon Gächter (2002). “Altruistic punishment in humans.” *Nature*, 415(6868), 137.
- Fehr, Ernst and Klaus M Schmidt (1999). “A theory of fairness, competition, and cooperation.” *The quarterly journal of economics*, 114(3), 817–868.
- Fehr, Ernst and Klaus M Schmidt (2010). “On inequity aversion: A reply to Binmore and Shaked.” *Journal of economic behavior & organization*, 73(1), 101–108.
- Fischbacher, Urs and Verena Utikal (2013). “On the acceptance of apologies.” *Games and Economic Behavior*, 82, 592–608.
- Frank, Robert H (1988). *Passions within reason: the strategic role of the emotions*. WW Norton & Co.
- Frank, Steven A (2003). “Repression of competition and the evolution of cooperation.” *Evolution*, 57(4), 693–705.
- Gale, John, Kenneth G Binmore, and Larry Samuelson (1995). “Learning to be imperfect: The ultimatum game.” *Games and economic behavior*, 8(1), 56–90.
- Gintis, Herbert (2003). “The hitchhiker’s guide to altruism: Gene-culture coevolution, and the internalization of norms.” *Journal of theoretical biology*, 220(4), 407–418.
- Gold, Gregg J and Bernard Weiner (2000). “Remorse, confession, group identity, and expectancies about repeating a transgression.” *Basic and Applied Social Psychology*, 22(4), 291–300.
- Grafen, Alan (1984). “Natural selection, kin selection and group selection.” *Behavioural ecology: An evolutionary approach*, 2, 62–84.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze (1982). “An experimental analysis of ultimatum bargaining.” *Journal of economic behavior & organization*, 3(4), 367–388.
- Hamilton, William D (1964). “The genetical evolution of social behaviour. II.” *Journal of theoretical biology*, 7(1), 17–52.
- Han, The Anh, Luís Pereira, Francisco C. Santos, and Tom Lenaerts (2013). “Why is it so hard to say sorry? evolution of apology with commitments in the iterated Prisoner’s Dilemma.” *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. Ed. Francesca Rossi 177–183.
- Harms, William (1997). “Evolution and ultimatum bargaining.” *Theory and decision*, 42(2), 147–175.
- Hauert, Christoph (2010). “Replicator dynamics of reward & reputation in public goods games.” *Journal of Theoretical Biology*, 267(1), 22–28.
- Henrich, Joseph, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwina Gwako, Natalie Henrich, et al. (2006). “Costly punishment across human societies.” *Science*, 312(5781), 1767–1770.
- Ho, Benjamin (2012). “Apologies as signals: with evidence from a trust game.” *Management Science*, 58(1), 141–158.
- Huttegger, Simon M, Justin P Bruner, and Kevin JS Zollman (2015). “The handicap principle is an artifact.” *Philosophy of Science*, 82(5), 997–1009.
- James, Scott M (2010). *An introduction to evolutionary ethics*. John Wiley & Sons.
- Joyce, Richard (2007). *The evolution of morality*. MIT press.
- Ketelaar, Timothy (2006). “The role of moral sentiments in economic decision making.” *Social psychology and economics*, 97–116.
- Longino, Helen and Ruth Doell (1983). “Body, bias, and behavior: A comparative analysis of reasoning in two areas of biological science.” *Signs: Journal of Women in Culture and Society*, 9(2), 206–227.
- Michod, Richard E (1982). “The theory of kin selection.” *Annual Review of Ecology and Systematics*, 13(1), 23–55.
- Morgenstern, Oskar and John Von Neumann (1953). *Theory of games and economic behavior*. Princeton university press.
- Nash, John F (1950). “The bargaining problem.” *Econometrica: Journal of the Econometric Society*, 155–162.
- Nelissen, Rob and Marcel Zeelenberg (2009). “When guilt evokes self-punishment: evidence for the existence of a Dobby Effect..” *Emotion*, 9(1), 118–122.

- Nichols, Shaun (2004). *Sentimental rules: On the natural foundations of moral judgment*. Oxford University Press.
- Nowak, Martin and Karl Sigmund (1993). “A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game.” *Nature*, 364(6432), 56–58.
- Nowak, Martin A (2006a). *Evolutionary dynamics*. Harvard University Press.
- Nowak, Martin A (2006b). “Five rules for the evolution of cooperation.” *science*, 314(5805), 1560–1563.
- Nowak, Martin A and Karl Sigmund (1998). “Evolution of indirect reciprocity by image scoring.” *Nature*, 393(6685), 573.
- Nucci, Larry P (1985). “Children’s conceptions of morality, societal convention, and religious prescription.” *Moral dilemmas: Philosophical and psychological issues in the development of moral reasoning*, 137–174.
- Nucci, Larry P (2001). *Education in the moral domain*. Cambridge University Press.
- O’Connor, Cailin (2019). *The Evolution of Inequity*. Oxford University Press.
- Ohtsubo, Yohsuke and Esuka Watanabe (2009). “Do sincere apologies need to be costly? Test of a costly signaling model of apology.” *Evolution and Human Behavior*, 30(2), 114–123.
- Okamoto, Kyoko and Shuichi Matsumura (2000). “The evolution of punishment and apology: an iterated prisoner’s dilemma model.” *Evolutionary Ecology*, 14(8), 703–720.
- Okasha, Samir (2013). “Biological Altruism.” *The Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta. Fall 2013 edition. Metaphysics Research Lab, Stanford University.
- Ostrom, Elinor, James Walker, and Roy Gardner (1992). “Covenants with and without a sword: Self-governance is possible.” *American political science Review*, 86(2), 404–417.
- O’Connor, Cailin (2016). “The evolution of guilt: a model-based approach.” *Philosophy of Science*, 83(5), 897–908.
- Panchanathan, Karthik and Robert Boyd (2004). “Indirect reciprocity can stabilize cooperation without the second-order free rider problem.” *Nature*, 432(7016), 499.
- Pereira, Luís Moniz, TA Han, LA Martinez-Vaquero, and Tom Lenaerts (2016). “Guilt for Non-Humans.” *Proceedings of the Thirtieth AAAI conference on Artificial Intelligence*. Ed. Dale Schuurmans and Dale Wellman.
- Pereira, Luís Moniz, Tom Lenaerts, LA Martinez-Vaquero, and TA Han (2017a). “Evolutionary Game Theory Modelling of Guilt.”
- Pereira, Luís Moniz, Tom Lenaerts, LA Martinez-Vaquero, and TA Han (2017b). “Social Manifestation of Guilt Leads to Stable Cooperation in Multi-Agent Systems.” *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*. Ed. Kate Larson and Michael Winikoff 1422–1430.
- Plutchik, Robert (1991). *The emotions*. University Press of America.
- Prinz, Jesse (2008). “Is morality innate.” *Moral psychology*, 1, 367–406.
- Regan, Dennis T, Margo Williams, and Sondra Sparling (1972). “Voluntary expiation of guilt: A field experiment..” *Journal of Personality and Social Psychology*, 24(1), 42.
- Rosenstock, Sarita and Cailin O’Connor (2018). “When It’s Good to Feel Bad: An Evolutionary Model of Guilt and Apology.” *Frontiers in Robotics and AI*, 5, 9.
- Rousseau, Jean-Jacques (1984). *A discourse on inequality*. Penguin.
- Sachs, Joel L, Ulrich G Mueller, Thomas P Wilcox, and James J Bull (2004). “The evolution of cooperation.” *The Quarterly review of biology*, 79(2), 135–160.
- Santos, Francisco C, Marta D Santos, and Jorge M Pacheco (2008). “Social diversity promotes the emergence of cooperation in public goods games.” *Nature*, 454(7201), 213.
- Shepher, Joseph (1983). *Incest: A biosocial view*. Academic Press.
- Silfver, Mia (2007). “Coping with guilt and shame: A narrative approach.” *Journal of Moral Education*, 36(2), 169–183.
- Skyrms, Brian (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press.
- Skyrms, Brian (2014). *Evolution of the social contract*. Cambridge University Press.
- Smith, J Maynard (1964). “Group selection and kin selection.” *Nature*, 201(4924), 1145.
- Smith, J Maynard and George R Price (1973). “The logic of animal conflict.” *Nature*, 246(5427), 15.

- Sober, Elliott and David Sloan Wilson (1999). *Unto others: The evolution and psychology of unselfish behavior*. Harvard University Press.
- Sripada, Chandra Sekhar (2008). “Nativism and moral psychology: Three models of the innate structure that shapes the contents of moral norms.” *Moral psychology*, 1, 319–343.
- Sripada, Chandra Sekhar and Stephen Stich (2005). “A framework for the psychology of norms.” *Morality, Norms, and Religion*. .
- Stanford, Kyle (2017). “Underdetermination of Scientific Theory.” *The Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta. Winter 2017 edition. Metaphysics Research Lab, Stanford University.
- Stanford, P Kyle (2018). “The difference between ice cream and Nazis: Moral externalization and the evolution of human cooperation.” *Behavioral and Brain Sciences*, 41.
- Tajfel, Henri (1970). “Experiments in intergroup discrimination.” *Scientific American*, 223(5), 96–102.
- Tangney, June Price, Rowland S Miller, Laura Flicker, and Deborah Hill Barlow (1996). “Are shame, guilt, and embarrassment distinct emotions?.” *Journal of personality and social psychology*, 70(6), 1256–1269.
- Tooby, John and Leda Cosmides (1992). “The psychological foundations of culture.” *The adapted mind: Evolutionary psychology and the generation of culture*, 19.
- Trivers, Robert L (1971). “The evolution of reciprocal altruism.” *The Quarterly review of biology*, 46(1), 35–57.
- Weisberg, Michael (2012). *Simulation and similarity: Using models to understand the world*. Oxford University Press.
- Wilkinson, Gerald S (1984). “Reciprocal food sharing in the vampire bat.” *Nature*, 308(5955), 181.
- Wong, Ying and Jeanne Tsai (2007). “Cultural models of shame and guilt.” *The self-conscious emotions: Theory and research*, 209–223.
- Wynne-Edwards, Vero Copner. Animal dispersion: in relation to social behaviour. Technical report, (1962).
- Yamagishi, Toshio (1986). “The provision of a sanctioning system as a public good..” *Journal of Personality and social Psychology*, 51(1), 110.
- Young, H Peyton (1993). “An evolutionary model of bargaining.” *Journal of Economic Theory*, 59(1), 145–168.