

Joint Action and Hominin Cognitive Evolution

Brian McLoone
Thesis Committee:
Patrick Forber (Chair)
Stephen Bailey

Prologue

Put broadly, this paper is about the role that joint, coordinated action played in hominin cognitive evolution. Typically, the theoretical methods used to explain cooperation among hominins reference the role that kinship plays in such interactions, the effect that iterated interactions have on the long-term pay-offs of Prisoner's Dilemma games, and the role that cultural norms play in suppressing subversion within groups and promoting pro-social behavior (see, for instance, Cartwright 2008; Richerson and Boyd 2005). There has been too little attention paid to an alternative model, recently advocated by Skyrms (2004), in which the emergent benefits of joint action are configured as the impetus behind some forms of cooperation among humans. The idea, put simply, is illustrated by the following scenario, known as a Stag Hunt game: We can imagine that two hunters can each snag three hares over the course of a day. But, if they team up to hunt together, they can hunt much larger game, like stag. The crucial point is that a stag is worth more than the sum of both hare hunts. The behavioral potential of the group is emergent; it does not cleanly reduce to the inputs of the actors involved. So if the yield of such joint action is distributed at least somewhat fairly, then cooperating agents are poised to benefit.

In Chapter I, I'll discuss how joint action fits into ethological accounts of cooperation among non-kin. After doing so, I'll draw attention to psychological research which suggests that humans have an early-emerging capacity to engage with others in joint activities (Tomasello, Carpenter, Call, et al. 2005; Warneken, Chen and Tomasello 2006), and I'll argue that the emergent benefits of joint action – as depicted through the Stag Hunt game – help us to understand how these capacities evolved. In Chapter II, I'll provide a re-shaping of the *Social Intelligence Hypothesis* in the manner of Sterelny (particularly 2007, but also 2003, 2008b) and will argue that joint action played a crucial role in the cognitive evolution of hominins. In its standard guise, the Social Intelligence Hypothesis configures

hominin cooperative behavior as following the payoff-structure of a Prisoner's Dilemma game, and the pressures of hominin social life as principally involving the need to detect cheaters within your group, to avoid detection when you cheat, and to keep track of group mates and their relationships to each other. While this account is good and right, it is not complete. There was also selection for conceptual capacities which allowed hominins to team and take part in joint action – in carcass defense, anti-predation defense, group raids, collective hunting, and social organization. These points need to be a part of the Social Intelligence Hypothesis if we're after a complete account of the selective pressures which guided hominin cognitive evolution, and so Chapter II is, in essence, an effort to sketch just such an account.

Finally, in Chapter III, I'll consider interesting conceptual issues which emerge when we consider how to describe a group of agents engaged in joint action, an issue which becomes tricky when we acknowledge (1) that the behavioral potential of the group is emergent and (2) that groups themselves can be units of selection. I'll end up defending the view that groups engaged in joint action are distinct entities in the world, and need not be decomposed into component parts. To a large extent, then, in this project I am attempting to apply somewhat recent conceptual developments in our understanding of cooperation and natural selection to issues which have to do with our species' evolutionary history.

I want to extend my gratitude to Patrick Forber, who saw this project from the start to the finish, and who over the course of that time provided thorough comments, criticism, and reading suggestions for each of these chapters, all of which he read multiple times. This paper is much better than it would have been had he not been involved, and so I want to express my thankfulness for his help. I also want to thank Stephen Bailey, who provided useful recommendations on these topics, and whose seminar on human evolution I attended last fall helped me greatly in writing Chapter II. I would also like to thank him because his Physical Anthropology class, which I took as a sophomore, was one the reasons I became interested in human evolution and human behavior to begin with.

Table of Contents

Chapter I: The Benefits of Joint Action	5
I.1 Explaining Cooperation among non-Kin	8
I.1.1 Cooperation in Interactions where $DC > CC$	10
I.1.2 By-product Mutualisms and Stag Hunts, where $CC > DC$	13
I.1.3 Game Theory: Too Vague for Anyone's Good?	22
I.2 Multi-Agent Differentiated Cooperation	23
I.2.1 MADC and Comparative Ethology	25
I.2.2 The Benefits of Joint Action and Differentiation	28
Chapter II: Joint Action and Its Role in Hominin Cognitive Evolution	31
II.1 The Social Intelligence Hypothesis	34
II.2 Integrating Joint Action and the Social Intelligence Hypothesis	39
Chapter III: Groups and Emergence, Reduction, and Selection	48
III.1 Joint Action and Emergent Behavioral Potential	50
III.2 Groups are Units of Selection	53
III.2.1 Substrate Neutrality and Multi-Level Selection	53
III.2.2 The Ways Groups are Units of Selection	55
Conclusion	59
Bibliography	61

Chapter I

The Benefits of Joint Action

Among primates, humans have the rare capacity to coordinate with others to achieve a shared goal, and to utilize differentiated sub-tasks to do so. What's more, this capacity emerges early ontogenetically – at about 14 months – and plays a key role in human social and cultural interactions (Tomasello, Carpenter, Call, et al. 2005; Warneken, Chen, and Tomasello 2006). In this chapter, I'll look at how this capacity evolved. Put briefly, my argument draws attention to the point that when two or more agents team up, they are often able to achieve more as a collective than the sum of what each individual could achieve were he or she to work alone. Each agent therefore stands to benefit from joint activity, assuming the yield of the activity is distributed in a somewhat equitable fashion.

I'll explicitly tie these points to the Stag Hunt game, which illustrates this phenomenon quite lucidly. The game describes the following scenario: Two hare hunters might each be able to catch three hares over the course of the day. But, they could also team up with each other to hunt larger game, like stag. What's crucial to note is that a stag is not worth six hares, the sum of the two hare hunts. It is worth far more than that. Thus, assuming that the goods of the collective action are shared somewhat fairly, this means that both agents have an incentive to partake in joint action (Skyrms 2004). This basic theoretical depiction can, of course, be extended to other forms of joint activity, and to activities which involve more than two agents.

In Section I.1, I'll better explicate these points and configure the Stag Hunt model within contemporary ethology. To do this, I'll identify two broad types of cooperative scenarios ethologists are typically dealing with when observing cooperation among non-kin. The first describes animal interactions in which defecting with a cooperator yields the highest return for an individual in a one-time interaction. Over the years, game theoreticians and ethologists have developed a theoretical

framework to explain why animals cooperate in such situations; in Section I.1.1, I'll detail what this framework amounts to. But put briefly for now, defection might close off the possibility of future interactions, so that in the long run, the benefits of cooperation might come to outweigh the benefits of one-time defection. Then, in Section I.1.2, I'll look at another class of animal interactions, those in which *cooperating* with a cooperator yields the highest return for an individual in a one-time interaction. Cooperation in these circumstances is not particularly vexing, since it is the best strategy for all involved. But conceptual snags arise when one considers how an agent comes to feel confident that (a) her partner will *actually* cooperate, not defect, and (b) that her partner will share the resources of the joint action such that it is worthwhile to team up in the first place.

With this ethological background in place, in Section I.2 I'll describe what I call *multi-agent differentiated cooperation* (henceforth, MADDC). As I have defined it, MADDC occurs when:

- (1) Two or more agents team up with each other to pursue a common goal
- (2) These agents use differentiated sub-tasks to achieve this goal
- (3) The payoff-structure of the interaction is such that joint cooperation yields a higher return than defecting with a cooperator

To be clear, not all cooperation in a Stag Hunt game is a form of MADDC, since agents can participate in a Stag Hunt game *without* utilizing (2) – i.e., differentiated sub-tasks – when they do so. Also, a Stag Hunt game has a highly-specific payoff-structure, and not all instances of MADDC mirror this payoff-structure perfectly. But the payoff-structures of Stag Hunts and MADDC *are* similar in that, in both scenarios, joint cooperation yields a higher return for an individual than does defecting with a cooperator. I have chosen to draw particular attention to MADDC because the capacity to engage in this behavior is quite rare among primates. Put very briefly for now, laboratory and field studies indicate that humans and chimpanzees are the only primates which cooperate with each other in this way, and chimpanzees do so

for the most part in highly specific contexts – namely, group hunting. From a comparative approach, therefore, MADC is a decidedly salient phenomenon, one worth considering in some detail. Giving a label to the behavior makes this task easier.

There is, in fact, an underlying theoretical shape to this chapter, which perhaps it is best to explicate from the get go: Namely, that we *need* a selective account of humans' ability to cooperate in joint action. Why is this so? The first reason is that the capacity to cooperate with others in this way is highly complex. As with wings, eyeballs, sonar, and so on, we should expect that this capacity didn't evolve because of a lucky evolutionary jump, but, rather, from a process of incremental improvement motivated by incremental selective advantages (Dawkins 1986). My argument is that the basics of the Stag Hunt quantify these advantages, and therefore provide a framework to understand the reasons a capacity of this sort has come into existence. The second reason is that this capacity is necessarily cooperative. By choosing to hunt stag (or whatever) with a partner, an agent is not only helping herself, but also her partner. We need an account of why agents don't defect in such scenarios. Again, I'll argue that the Stag Hunt illustrates the answer: Because joint action yields returns greater than the sum of what both individuals could do alone, both stand to benefit from cooperating.

Before moving on, it is also important, for the sake of clarity, to be maximally explicit about the main explanatory goal of this chapter. I am interested in how the *capacity* to coordinate joint action with others evolved in our lineage. Thus, I am interested in the evolution of an individual-level trait, since a psychological capacity is something in an individual's brain. Alternatively, joint action the *behavior*, at least in cases where tasks are differentiated, is a group-level behavior. In this chapter, I am not so interested in this topic, but will take it up in Chapter III. Again, here I am interested in providing an individual-level account of how our capacity to coordinate action in this way evolved. I'll deal with the interesting group-level feature of joint action – namely, the behavior itself – only at the end of this paper.

I.1 Explaining cooperation among non-kin

Before getting too far along, it's very much mandatory that I dwell on what I mean by *cooperation*, because unclear use of the word – or its cousin, *altruism* – can create unnecessary problems. In this paper, I will use the word cooperation in a broad sense, and will apply it to an animal's behavior which increases the fitness of one or many others. I will also use this word to describe behavior which an animal performs which decreases its own fitness, but increases the fitness of its close relatives. I count behavior which evolved through *kin selection*, in other words, as cooperative. In contrast, I'll use the word altruism to describe a behavior which an animal performs which increases the fitness of another, unrelated animal, and which decreases the performer's own fitness. Now, it should be noted that an altruist, as I am using the notion, might decrease its fitness *within* a group, but enhance its fitness *across* groups, since altruistic groups might showcase greater reproductive output than non-altruistic groups. That is to say, I am using a *multi-level interpretation of altruism* (in the sense of Kerr, Godfrey-Smith, and Feldman 2004). Employing these definitions, MADC is indeed cooperation, even though some writers are apparently uneasy with the idea that cooperation can generate shared goods (see Dugatkin 1997, pp. 32-33; Rosenberg 1992; and D.S Wilson and Dugatkin 1992, for discussion of this). Ultimately, I'm interested in the biological interactions themselves, not what we call them. For that reason, I am not too concerned with terminological issues, and attempt to avoid them simply by using the terms consistently.

Throughout the animal world, many instances of cooperative behavior can be accounted for through appeal to a theoretical construct typically called kin selection or *inclusive fitness* (Hamilton 1964; Dugatkin 1997). When an animal helps another, closely-related animal, it is in some sense helping itself, since the two share so much genetic material. Brown hyenas in the Kalahari Desert showcase an interesting social structure which bears on this issue. A particular clan contains males and females of various ages, but during any given breeding season, it is common for only one female to bear all of the

clan's offspring. When these new pups reach the age of three months or so, they are moved from a private den to a communal den with the rest of the clan, where the older hyenas provision food for them. What's important is that around 80% of this provisioning is typically carried out by hyenas who aren't the pups' mother, though those who provision *are* likely either aunts or cousins to the pups. So in large part, we can understand this behavior as deriving from the fact that the hyenas are closely related. Even though provisioning for the pups takes time and energy, the costs are likely less than the benefits of helping an organism which shares many of the same genes, such that the gene which codes for this provisioning trait – helping pups other than your own – is poised to evolve (Dugatkin 1997, p. 104).

Yet not all cooperative behavior occurs among kin. At the extreme, there are a variety of scenarios where members of different *species* form partnerships, in which case there is no non-trivial genetic similarity which could drive the behavior. For the rest of this section, I'll look at how ethologists have dealt with cooperation of this sort. Specifically, I'll identify two *types* of interactions non-kin cooperate in, and detail the theoretical frameworks ethologists use to understand cooperation in each. In the first class of interactions, the subject of Section I.1.1, the highest individual reward in a one-time interaction goes to the agent who defects with a cooperator, such that $DC > CC$, where *D* stands in for *defect*, *C* stands in for *cooperate*, *DC* means *defect with a cooperator*, and *CC* means *cooperate with a cooperator*. I briefly noted at the start of this chapter how cooperation in such cases can turn out to be a profitable strategy, and I'll expand on those points below. In Section I.1.2, I'll look at another class of animal interactions, those in which the highest individual level payoff in a one-time interaction goes to the agent which *cooperates* with a cooperator, such that $CC > DC$. The issue here, as I also briefly noted, is not so much why it makes sense for animals to cooperate in such scenarios, but how an animal comes to trust that its partner will reciprocate, and how it can ensure that the yields of joint action are distributed so that that it makes sense to cooperate in the first place.

I.1.1 Cooperation in interactions where $DC > CC$

Mitani has documented that chimpanzees at the Kibale National Park in Ngogo, Uganda, exchange meat for grooming. So let's imagine that chimpanzee A comes back from a hunt, and gives chimpanzee B some meat, expecting that chimpanzee B will repay the favor with a grooming session. Chimpanzee B, one would think, would be tempted to take the meat and then just walk off. That way, chimpanzee B gets the benefits of the interaction, and doesn't have to pay any cost. But this rarely happens. Much more often, chimpanzee B *does* repay the favor (Mitani 2006, p. 115).

How are we supposed to account for interactions like these? It's helpful if we understand this interaction using the framework provided by the Prisoner's Dilemma game. There are a number of ways to illustrate the take home message of this game, but one story is particularly popular. Imagine that the police have just caught two thieves, and put them in different cells, and the cells are far enough apart that the thieves can't see or hear each other. Each prisoner has one of two options: He can defect from his partner by confessing the crime, or cooperate with his partner by keeping silent. If one prisoner defects on his partner, he can get a reduced sentence of, say, one year, while his partner will get punished for committing the crime and not confessing. If they both keep silent, however, then there is less evidence to charge them, and each will serve time, but each will serve a reduced sentence of, let's say, three years. And if they both defect, then they'll both be charged harshly, but less harshly than the worst situation (cooperating with a defector), since they were both at least forthcoming (Vincent and Brown 2005. p. 67). Here's are the payoffs in a Prisoner's Dilemma game:

	Cooperate	Defect
Cooperate	4,4	0,7
Defect	7,0	2,2

Looking at this payoff-matrix, the best option for an individual is to defect with a cooperator, while the worst possible situation is that you end up cooperating with a defector.¹ Cooperating with a cooperator isn't as bad as cooperating with a defector, but it's not as good as defecting with a cooperator – hence, $DC > CC$.² (As an aside, something worth noting is that the best *group-level* strategy is to cooperate with a cooperator, since that results in the lowest collective amount of jail time, but this is at odds with the best individual-level strategy. In the next section, we'll see that in Stag Hunt games this tension between individual-level and group-level fitness is erased.)

But an interesting transformation occurs to the payoff-structure of Prisoner's Dilemma games when they are run over and over again between the same players. The benefits of cooperating in such cases has been understood since the 1950's, but Trivers (1971) explicitly connected this work to ethology. Consider the following example, which he used in his seminal paper on the topic. Wrasses are a small fish, and they will often swim into the mouths of groupers, to clean off some of the ectoparasites which have accumulated there. The wrasse gets protection and a meal out of the interaction, and the grouper gets a cleaning. But consider the interaction from the point of view of a grouper. The grouper has the opportunity to get a cleaning from the wrasse and then to eat the wrasse afterward. This would amount to the grouper defecting on the cooperating wrasse, which would yield the highest return for the grouper. Yet the grouper doesn't defect. Why? Trivers pointed out that when interactions are iterated, the payoffs can change. Specifically, he argued that repeated interactions can make "apparently detrimental" behaviors between unrelated organisms adaptive, a phenomenon he called *reciprocal altruism*. So consider again the interaction between groupers and wrasses. Under Trivers's model, we can understand this behavior because the grouper benefits more from the wrasse's future work than it does from an extra meal now (pp. 40-43).

¹ The better strategy in this payoff-matrix is indicated by having the *higher* number; thus, a higher number indicates less time in prison.


² Unless otherwise noted, payoff-structures will refer to the relative yields which a particular strategy will have for an individual in a one-time interaction.

We can apply this model to the interaction between chimpanzees A and B, mentioned above. It's true that chimpanzee B stands to gain a lot for this *one* interaction if he doesn't give anything in return for the meat chimpanzee A gave to him. But over the long-term, it might be best for him to actually reciprocate and groom his group mate. Because if chimpanzee B doesn't do so, chimpanzee A might not want to do business with him in the future. Thus if chimpanzee B has reason to believe interactions like these might take place in the future (and, that he benefits at least partially from honest trades), then he might be best off cooperating, since, over the long-term, the modest gains he gets from honest trades will come to outweigh the one-time, big gain he gets when he defects.

Not long after Trivers's work, Robert Axelrod and W.D. Hamilton made an important and related point. Their contribution came about following the results of two computer tournaments Axelrod had held, for which various academics each submitted a program, one which competed against all other programs in a tournament based upon the payoff-matrix of a Prisoner's Dilemma game. A highly successful program in both tournaments was Rapaport's, which employed a "tit-for-tat" strategy. That is, the program began an interaction cooperatively, then repeated whatever move its partner had just made. Its success demonstrated that cooperation in iterated games can be both stable and profitable, even in an environment of many different strategies (Axelrod 1984, p. 21). This work was important, because while in principle cooperation in iterated Prisoner's Dilemma games might make sense, such a strategy would seem liable to *subversion from within*, as selfish strategies would have their backs scratched, as it were, but never scratch back, such that cooperation would fall apart within a population. What Axelrod and Hamilton showed is that by repeating the move your partner just made, one can avoid this threat, leading the authors to write that Trivers's work and their own, along with Hamilton's explication of *kin selection*, were evolutionary theory's "recently acquired two kinds of extension," by which researchers could "account for the manifest existence of cooperation and related group behavior, such as altruism and restraint in competition" (Axelrod and Hamilton 1984, p. 89).

1.1.2 By-product Mutualisms and Stag Hunts, where $CC > DC$

Curiously, there has been decidedly less attention paid to an alternative payoff-structure in cooperative behavior, one which illustrates the point that joint cooperation can, in certain circumstances, yield the highest returns for individuals in a one-time interaction, such that $CC > DC$. For instance, yellowtail hunt mackerel by pushing the mackerel up against the shore, so that they have nowhere to escape, at which point the yellowtail send in an attacker to disperse the mackerel into the mouths of the other yellowtail (Dugatkin 1997, p. 57).



The yellowtail benefit by coordinating action to hunt the mackerel. Image taken from Dugatkin 1997, p. 57.

By cooperating with others, each yellowtail is able to access rewards which would have been difficult to get were an individual yellowtail to act alone. Thus, so long as each yellowtail gets its share, cooperation in these situations is the best individual-level strategy. CC is the most fitness-enhancing strategy because the behavioral potential of the group exceeds what a lone yellowtail could accomplish.

In a similar manner, chimpanzees in the Tai forest work together to hunt red colobus monkeys, and their success rates go up as more chimpanzees engage in the joint action, so that joint cooperation yields the highest returns (Boesch 2006). Like with the yellowtail, each individual chimpanzee benefits

from cooperating jointly with others, because the collective makes possible what a lone individual simply could not do, or would have a very slim chance of being able to do, were it to act alone. Through his field research of these hunts, Boesch has concluded that “Cooperation in hunting among Tai chimpanzees is not readily explained by kin selection, but rather is the result of a mutualistic process in which all participants gain more than if they were acting alone”(Boesch 2006, p. 146-148). It should be noted that interactions of this sort are typically what Tomasello, Carpenter, Call et al. (2005) call *triadic*: They involve at least two agents and a goal which is external to those agents. This contrasts with most interactions for which $DC > CD$, because in such scenarios the interaction typically involves just the two cooperating (or defecting agents), and their interactions with each other.

Often, when an interaction is such that $CC > DC$, it is called an instance of *by-product mutualism* (Dugatkin 1997, p. 31). Recently, particularly in behavioral ecology, interactions with this payoff-matrix are called *mutualistic* (see, for example, Cartwright 2008; Kappeler and van Schaik 2006, p. 11; Boesch 2006). Unfortunately, both of these terms are problematic: The word by-product makes such behavior sound misleadingly accidental, which, as we have just seen, it often is not. Chimpanzees coordinate their hunts quite purposefully, for instance, and the strategies they use to hunt, as we’ll see below, only work because other chimpanzees are pursuing complementary strategies. So not all scenarios where $CC > DC$ are accidental. And the word mutualism is problematic too, since it has traditionally been used to refer to inter-specific interactions. Yet I don’t want to confuse the discussion by adding another term, so I’ll refer to behaviors in which $CC > DC$ as being examples of mutualisms, or examples of joint action.

A model for understanding certain forms of this type of cooperative behavior is the Stag Hunt game. Though it dates back to Rousseau, the game has been re-emphasized, with a more Darwinian twist, by Skyrms (2004). In short form, the Stag Hunt is a game theoretic model which quantifies the benefits of coordinated action, as the payoff-matrix below illustrates.


	Cooperate	Defect
Cooperate	9,9	0,5
Defect	5,0	3,3

While defecting with a cooperator yields the highest yield in a Prisoner’s Dilemma, cooperating with a cooperator yields the highest yield in the Stag Hunt. (It is important to note, for clarity’s sake, that cooperation in Stag Hunts is mutualistic, but not all mutualistic cooperation follows the entire payoff-structure of a Stag Hunt game. For in some mutualistic encounters, we can imagine that, for example, $DD = CD$, or $DC = DD$, and, strictly speaking, Stag Hunts are interactions where $CC > DC > DD > CD$.)

Unlike the examples from the last section, when scenarios follow the Stag Hunt payoff-structure, or the more general $CC > DC$ payoff-structure of mutualisms, we don’t need to include some *additional* variable (say, relatedness, or the prospect of future encounters) to understand why it makes sense for animals to cooperate. They cooperate because the collective is able to achieve what each individual within the collective would have a hard time achieving alone; indeed, in some cases, it might be *impossible* for lone individuals to achieve what the collective can. Carcass defense, Sterelny points out, is a good example of a Stag Hunt game, and, hence, is a good example of a mutualistic encounter. A lone defender of a carcass has very little chance of fending off prey animals. But as a few more agents get involved, defense becomes possible, because agents will then be able to surround the kill (Sterelny 2007). No one individual can defend a carcass, but when a few individuals team up, they have a shot.

Maynard-Smith and Szathmary make an important distinction which is related to this discussion when they describe *skulling games* versus *rowing games*. In a skulling game, one of the individuals can sit back and let the other do all of the work, since both people have a pair of oars. But in a rowing game, each person has only one oar, so the boat can’t move in the right direction unless both are putting in

effort (Maynard-Smith and Szathmary 1995). Yellowtail hunting mackerel, or chimpanzees hunting red colobus, or bi-peds defending carcasses, mirror the rowing scenario: It only works, or works consistently, when all the agents are *actually* putting in their fair share.




Defecting with a cooperator is the best strategy when skulling, but when rowing, cooperating gets both you and your partner the highest returns. Image taken from Maynard-Smith and Szathmary (1995, p. 262).

Of course, an agent could just sit back even in a rowing game and the interaction would never get off the ground. So what might be the reason she cooperates? For one, the rewards of successfully completing the shared goal might outweigh the costs of cooperating to achieve that goal. Or, likewise, there might be situations in which cooperation is mandatory, such as anti-predator defense, where not cooperating would be disastrous for your own fitness. Finally, agents with whom one is cooperating might monitor the input everyone is contributing. Boesch has documented that Tai chimpanzees, for example, provision captured monkeys on the basis of participation in the hunt, and exclude provisioning to those who did not play legitimate roles in snagging the monkey (Boesch 2001).

Stag Hunt games in particular, and mutualisms generally, illustrate that in some situations, you expand your behavioral range by teaming up with another – either because some tasks become easier, or some become possible – and you are thereby poised to access resources which are in many cases

non-linearly related to the input of the agents. For instance, if two agents can collectively hunt six hares over the course of the day, we might expect whatever they hunt together to be equal to six hares. But this is not the case. There is a big jump in behavioral potential when agents coordinate action, as a stag (or whatever prize joint action yields) is worth more than the sum of what the two (or more) agents would have accessed had they worked alone. So long as the accessed goods are distributed somewhat fairly, a point I'll return to shortly, each agent therefore stands to benefit from the interaction. The graph below illustrates this non-linear increase in returns.



Let the x-axis describe the number of hunters, and the y-axis, the yield of a hunt. Hunting alone, an agent can bag three hares (at one unit each), but by hunting with another, she can down bigger game, like stag (represented as thirty units).

Stag Hunts, to be clear, needn't involve actual hunting. The game theoretic model is meant to describe all sorts of interactions where collective action results in expanded behavioral potential, such that participating agents stand to benefit even if they share the goods. That is, there is an interesting emergent property of many types of such behavior where the behavior of the collective is non-linearly related to the sum of the behavioral potential of the collective's parts. It just so happens, however, that the Stag Hunt game is often best utilized as a description of hunting – not only because it describes so well the benefits of coordinated action, but because coordinated action among animals typically occurs when they are hunting.

What we have, then, is a framework which depicts the potential benefits of joint action, and field evidence that joint action actually occurs in the animal world along these lines, in that it follows the payoff-structure of $CC > DC$. With these points in mind, it is interesting to return to an assumption which runs through the work of Trivers, Axelrod and Hamilton: that cooperation is, *ceteris paribus*, a fitness-reducing option. This is why these researchers searched for a factor outside of the interaction itself to account for cooperative behavior – and as we saw, they settled on the influence of genetic relatedness or the effects of iterated interactions. This assumption follows from the fact that all three authors implicitly employed the payoff-structure of a Prisoner's Dilemma game in their studies, in which cooperation in a one-time interaction is *not* the best strategy. But, as we have seen, this payoff-structure does not exhaust the range of the types of interactions out in the world.

Despite the field evidence that there are interactions in which joint cooperation is the best strategy, and actually *is* the strategy some animals use, Binmore, for one, is skeptical that cooperation of the sort could get off the ground. He writes that “Alice and Bob may agree to cooperate in hunting a stag, but when they separate to put their plan into action, each may be tempted to abandon the joint enterprise by the prospect of bagging a hare for themselves” (2007, p. 68). This would only be true to the extent that hunting a hare is more likely successful than hunting a stag, however, which is an empirical point, not a conceptual one. The real threat to cooperating in Stag Hunts is that you don't know if your partner is going to reciprocate, or is going to leave you hanging, because doing so might put you at a relative fitness disadvantage. I don't configure this threat as particularly significant, though, since your partner has an incentive to participate in Stag Hunts because she'll benefit from the enhanced returns they can yield, too.

A more legitimate problem with the evolvability of mutualistic cooperation – particularly mutualistic hunting – is that we could imagine that Smith and Jones, who each catch two hares over the course of the day, team up to catch stag, but Jones then won't let Smith in on the yield of the collective

kill. If this happened, then Smith has no incentive to cooperate. We just saw that Binmore was skeptical of Stag Hunts because you can't be sure that your would-be partner will not defect. In fact, the more pertinent worry is that your partner will cooperate, but then expropriate the prize.

How might we deal with this issue? First, it is important to note that, because there is this non-linear increase in the behavioral potential of the collective relative to its component parts, Jones could take more than half of the hunt's resultant goods and Smith could still come out of the interaction doing better than if he had stuck to hunting hare. Additionally, in the same way that it might be best for an agent to cooperate in a Prisoner's Dilemma if there could potentially be future encounters from which she could benefit, it might make more sense for Jones to make the hunt worth Smith's while, so that Smith is willing to team up in the future for another hunt.

Important, too, is that if cheating were to occur after a joint action, it would typically be difficult to mask. Unlike with various forms of reciprocity, the memory demands of figuring out who has taken part in a collective hunt, for instance, are much simpler, because dividing the kill occurs immediately and in plain view. Making this point, Sterelny writes that joint action can avoid "the cognitive costs of policing reciprocation. If joint action is successful, the profit is shared jointly and simultaneously (and shared automatically, in the case of successful collective defence). Dividing the spoils after a successful buffalo hunt is not difficult. It is not delayed in time. The buffalo can be dismembered on the spot in circumstances in which each is monitored by all" – in other words, tacit manipulation on the part of exploitative hunters is difficult to achieve and easy to look-out for (Sterelny 2007, p. 724). This would seem to apply particularly well in many cases of collective hunting, where often there are more than just two agents involved, so that a dominant agent couldn't simply bully a subordinate – instead, the dominant would have to bully a whole cadre of individuals, which is a far less feasible task.

Another way to conceptualize the task of dividing the yields of joint action is to quantify it as a Dictator's game. A Dictator's game is quite simple: One participant gets to decide how a good is

distributed, and the other agent can do nothing but accept what she is given. This contrasts with an Ultimatum game, in which one participant decides how to divide the goods, but if the offer is not to the liking of the second participant, then neither gets anything. Dividing the yield of, for instance, a cooperative hunt is obviously far more like a Dictator's game. If the "dictator" in this interaction does not dole out enough of the prize, then the subordinate has no real incentive to cooperate in future tasks. Indeed, when chimpanzees were able to access a food prize through a cooperative task, and the dominant expropriated the entire prize, the subordinate was less likely to participate in future encounters (Melis, Hare, and Tomasello 2006).

Despite the above points, cooperation in cases where $CC > DC$ is often simply not recognized as an option by those writing on animal cooperation. Commenting on the evolution of language, for example, Pinker has written that the cooperative nature of speech is not especially odd. After all, information sharing is just another manifestation of the "standard problem of the evolution of any form of cooperation or altruism," and one for which the "solution" is the same. Thus, "By sharing information with our kin, we help copies of our genes inside those kin" and "if we inform only those people who are likely to return the favor, both of us can gain the benefits of trade" (Pinker 2003, p. 28). The quotation assumes the point that helping another agent means you necessarily hurt yourself, unless some other factor – relatedness or future interactions – justifies the putative sacrifice. A workable framework in that discussion, but one which is rarely mentioned, is that communication allowed the behavioral potential of those communicating with each other to expand, for mutual benefit.

So why have mutualistic explanations of cooperation, and the Stag Hunt game in particular, been overlooked? Commenting on this, Boesch writes, "Mutualism has long been recognized as one of the mechanisms leading to the evolution of cooperation, but because of its obvious nature, 'if both gain more, then they should cooperate', theoreticians have neglected it and concentrated on less obvious mechanisms, such as delayed reciprocity and altruism" (Boesch 2006, p. 149). This is a distinct

possibility. But another factor might have to do with the types of animals ethology has typically focused on. Going back to Trivers's example of a wrasse and a grouper, there is no single, one-time behavior the two can partake in which would allow each to achieve something that would have been otherwise impossible to achieve. Wrasses and groupers can't build houses, and their dietary needs give them no reason to hunt together – though that would be quite a spectacle. So perhaps ethology's traditional focus on interactions among relatively "simple" organisms has biased the field against recognizing the notable benefits of coordinated action among some other types of animals.

What we have, in sum, are two broad theoretical categories which we can use to describe two classes of interactions – $DC > CC$ and $CC > DC$ – and we have looked at how we can explain cooperation in each. The first class of interactions are situations in which defecting with a cooperator yields the highest reward, because an agent gets the rewards of the interaction and doesn't have to pay any cost. Cooperation in such cases can make sense if agents are closely related, or if interactions are iterated. The second class of interactions describes situations in which joint cooperation yields the highest rewards. Such interactions are typically triadic, and joint cooperation is the best option because working with others opens up new behavioral possibilities, which allow the agent to access resources which would otherwise be out of reach. Often, for these behaviors to work, each agent has to actually contribute, like in a rowing game, and I identified a few mechanisms by which legitimate participation could be motivated or enforced, and some reasons agents have to provision the yield of joint action in a somewhat fair way.

1.1.3 Game Theory: Too vague for anyone's good?

One possible criticism of the above discussion is that game theory is too simple to accurately depict what is going on in the animal world. Animal behavior, and particularly animal psychology, are acutely complex. One criticism of game theoretic approaches to the evolution of behavior is that they

are too simplified to tell us anything accurate or interesting about what is really going on in the world these games are used to represent. I think I am now in a better position to respond to this criticism. In the above discussion, I've only used game theory in a very restricted way, and this was done intentionally. Typically, a given payoff matrix in game theory will involve four pairs of behaviors. As we've seen, the payoff-structure of a Prisoner's Dilemma game, for instance, is $DC > CC > DD > CD$, while for Stag Hunts, it is $CC > DC > DD > CD$. And though I've explicated the payoff-structure of these games, I've used them only to showcase the point that some behaviors are such that $DC > CC$, while others are such that $CC > DC$, and I've kept to this deflated payoff-structure to avoid making the discussion unnecessarily complicated by including the other potential behavioral options, and their relative fitness returns.

As I see it, game theory provides useful models to quantify fitness-enhancing and fitness-reducing behaviors that we can actually observe out in the animal world. That is, there really are interactions where defection is the best option for an agent in a one-time interaction, and there really are interactions where cooperation is the best option for an agent in a one-time interaction. We looked at Mitani's documentation of chimpanzees in Uganda who exchange meat for grooming (Mitani 2006, p. 115). In a one-time interaction, getting the meat but not grooming would yield the highest reward. That is a basic empirical fact, and so we can say of the situation $DC > CC$. Alternatively, we just saw that participating in joint hunts with other chimpanzees increases the amount of food a chimpanzee would receive if he hunted alone or hunted not at all. Again, an empirical point, one which allows us to say that in such a situation, $CC > DC$.

To be clear, DC and CC are not themselves precise replicas of what is going on, they are constructs which we use to give value to various sorts of behavior or potential behavior, and to determine how certain behaviors make sense from the standpoint of natural selection. By keeping to this basic pair of descriptions, I have been able to make nice, general statements about various sorts of

animal interactions. Indeed, to say that game theoretic explanations are too simple in fact points out a feature of such explanations: They enable us to make broad characterizations of how cooperation works, then utilize those frameworks to explain particular token instances of cooperative behavior. So long as you acknowledge that the two examples just discussed do indeed display different payoff-structures – and it doesn't seem clear to me how one couldn't acknowledge that – then in essence you are acknowledging that the game theoretic models capture *something* about interactions in the world. And as we saw already, these theoretical constructs have quite clear explanatory power. This next section will showcase this point even more, when we use the benefits of joint action and divvying up labor to reverse-engineer humans' capacity to engage in MADC.

Section I.2 Multi-Agent Differentiated Cooperation: What is it, and why did it evolve?

As I mentioned at the beginning of this chapter, laboratory work has demonstrated that humans are able to coordinate with others to accomplish a shared goal, using differentiated sub-tasks to do so, from quite an early age (Tomasello, Carpenter, Call, et al. 2005; Warneken, Chen, and Tomasello 2006). I have chosen to label the capacity to engage in this sort of behavior as the capacity to engage in MADC, and MADC occurs when:

- (1)** Two or more agents team up with each other to pursue a common goal
- (2)** These agents use differentiated sub-tasks to do so
- (3)** The payoff-structure in the interaction is such that $CC > DC$

This capacity is cognitively complex, typically involves forming partnerships with non-kin, and unfolds with very little explicit instruction, since, on the one hand, children are able to cooperate in this way long before they are able to comprehend and produce language competently, and also, all normal children achieve this capacity at around the same time. In short, this capacity cries out for an account of its evolution. Albeit, we might say this about other humans' pro-social conceptual competencies as well.

And this is fair. But I have chosen to hone-in on MADC because it picks out one of the most sophisticated conceptual capabilities humans possess which relates to sociality, and one which is rare among primates. I have given it a name so as to make discussion of the capacity itself easier.

So then, why do humans, from so early in life, have the ability to coordinate with others in this way? My argument in this section is that we can understand how this capacity evolved by utilizing the theoretical points of Section 1.2, which showcased that joint action can be a fitness-enhancing strategy in many scenarios, because often, the returns of joint action are greater than the sum of what both (or more) individuals would be able to access, were they to work alone. But to be clear, this might explain the evolution of the capacity to engage in mutualistic behavior *generally*. MADC, keep in mind, is a sophisticated type of mutualistic interaction, because it involves agents using *differentiated sub-tasks* to achieve a shared goal, which is not a requirement of mutualistic cooperation on the whole.



So in this section, I'll devote some discussion to why the capacity to divvy up labor while pursuing a shared task evolved. Specifically, I'll look at Calcott's (2008) work in this area, which illustrates how dividing up tasks in this way is a good strategy for "generating benefit" most effectively. Put roughly, my argument is that the capacity to engaged in MADC evolved, first, because of the behavioral fact that when two or more agents team up, what they can achieve as a collective often exceeds the sum of what both could achieve were they to act alone, and second, because the most fitness-enhancing way to engage in joint action is to divide labor as you do so. MADC exploits physical aspects of the world – specifically, that there can be a non-linear increase in returns when one teams up with others, and that divvying up labor enhances these returns even more. To some extent, then, we might treat the capacity to engage in MADC as an evolutionary target toward which our ancestors' minds incrementally evolved. So I want to start off this section by exploring MADC as a psychological phenomenon, paying particular attention to its ontogeny and whether the capacity for it exists in other

primates. Then, I'll expand on the theoretical points I just made and describe how we can reverse-engineer MADC and come up with an account of why the capacity evolved.

1.2.1 MADC and comparative ethology

It has been well demonstrated that humans have the ability to engage in joint activities from quite an early age. At around their second birthday, children are able to take part in tasks that require two people to perform, and re-engage participants in such tasks when they slack off. For instance, in an experiment conducted by Warneken, Chen, and Tomasello, 18-24 month-olds would hold on to one side of a small trampoline, and make a block on the tarp of the trampoline jump up and down by moving the trampoline in unison with the experimenter. The children knew that they could not simply move the trampoline up and down randomly, but that they had to coordinate their actions with their partner. In similar scenarios, when the experimenter would stop participating in a joint activity, children typically attempted to re-engage her with a gesture – for instance, pointing to the part of the apparatus she was supposed to be operating – or verbally (Warneken, Chen, and Tomasello 2006).

To what extent do we see this behavior among non-human primates? It seems clear that at least chimpanzees can take part in joint cooperation in which sub-tasks are *not* differentiated. In an experiment conducted by Melis, Hare, and Tomasello (2006), chimpanzees were placed in a barred room. In front of the chimpanzee, on the other side of the bars, was a platform with food on it. To access the food, the chimpanzee would have to pull in the ropes.



Notice that the rope will not pull in the plank if just one end of the rope is pulled. Image taken from Melis, Hare, and Tomasello (2006)

But in one version of the experiment, the ropes were placed too far apart for a lone chimpanzee to pull them in by herself. And if she were to pull just one rope, the rope would come off of the platform, and the food would stay out of reach. There was a door to the adjacent room that could be opened easily once a peg (which served as the key) was taken out of the frame. When the task was such that chimpanzees *needed* a partner to get access to the food – i.e., when the ropes were too far apart for one chimpanzee to pull alone – it was much more likely that the chimpanzee would open the door to let the partner in. When the ropes were arranged so that one chimpanzee could pull the food in without any help, those chimpanzees were far less likely to allow the partner to come in.

Clearly, chimpanzees understand that some tasks require two. They struggle, however, at least in some circumstances, with tasks which require joint cooperation and differentiated sub-tasks. Humans, we know, are able to engage in tasks which necessitate this division of labor quite early in life. For instance, in a cooperative problem-solving activity, Tomasello and his colleagues tested the ability of 14-month-old children to access a prize that was encased in a tricky apparatus. The apparatus required coordination between a child and an adult, who was a confederate. One of the two had to push up a lift which put the prize in reaching distance of the other. But the person who pushed up the lift couldn't also reach around to grab the prize, and the person who could grab the prize couldn't also push up the lift. So the tasks were complementary, and also quite different from each other. When the lab

technician stopped participating in the activity, the child would prompt her verbally and in some cases gesturally to finish her task. Quite interesting, when a parallel experiment was performed with chimpanzees, the chimps failed to impel the experimenters to cooperate again, and would attempt instead to perform their tasks with greater intensity. The chimps failed to recognize that their actions alone – no matter how intensely applied – would not allow them to access the prize; put another way, that the prize was unreachable without the help of the other (Warniken, Chen, and Tomasello 2006). This is particularly interesting, because we just saw with the plank experiment that chimpanzees *could* conceptualize that some tasks required two to be achieved. Perhaps the difference is that, when the sub-tasks are differentiated, the chimpanzees are unable to grasp the dynamics of the scenario.

Boesch's account of chimpanzees' hunting strategy in the Tai forest, however, should forestall any broad conclusions that the above experiments might motivate. Describing such hunts, he writes that a "*driver* initiates the hunt by slowly pushing the arboreal prey in a constant direction, *blockers* climb trees to prevent the prey from dispersing in different directions, the *chaser* may climb under the prey and by rapidly running after them try a capture, and the *ambusher* may silently climb in front of the escape movement of the prey to block their flight and close a trap around the prey" (Boesch 2005, p. 692, emphasis in original, see also Boesch 2001, 2006). There is also a degree of flexibility that the chimpanzees showcase in such hunts, as the monkeys try different tricks to avoid capture, showcasing that the chimpanzees are able to recognize their role in the hunt in relationship to the role of others. That chimpanzees understand the dynamics of differentiated labor is buttressed by the point that the chimpanzees with the toughest jobs in the hunt typically get the most meat when the monkey carcass is provisioned (Boesch 2001).

So the claim that humans only are capable of taking part in MADC is very probably incorrect. What is interesting, however, is that chimpanzees have so far been unable to cooperate in this way in situations other than collective hunting. So what we might conclude is that what's unique about the

human capacity to engage in MADC is that we can do it – indeed, children can do it – in a whole host of settings, whereas chimpanzees can do it in a specific setting. This difference is important also because perhaps chimpanzees need to *learn* how to behave in this way far more than humans do, which could account for why chimpanzees are unable to showcase MADC in such artificial laboratory experiments. This coheres nicely with the fact that, among Tai chimpanzees, whose hunting skills are the most highly sophisticated of any group catalogued, males don't begin to hunt until the age of ten, after a long process of learning (Boesch 2001).

1.2.2 *The benefits of joint action and differentiation*

In this section, Section 1.2.2, I will “reverse engineer” (in the sense of Dennett 1995) the capacity humans have to engage in MADC, and argue that we can understand the capacity's existence by (1) utilizing the broad message of the Stag Hunt model from game theory – “broad” in the sense that I only care that $CC > DC$, not that $CC > DC > DD > CD$ – and (2) acknowledging the important theoretical point that dividing up labor is an effective way to further enhance the benefits of joint action.

What we have then, as a starting point, is observational data which shows that present-day human infants are able to take part in cooperative action which involves a shared goal and requires differentiated sub-tasks to achieve that goal. What we need, however, is a theoretical framework to understand how the capacity to engage in this sort of behavior evolved. The Stag Hunt model provides just such a framework. As an illustration: Smith and Jones might be able to individually catch three hares over the course of a day. But this does not mean their joint hunting will yield a payoff equal to just six hares. For some activities, when you team up with another, you don't just get double the goods of what you would have gotten by yourself. It is an aspect of our world that some joint interactions showcase an emergent property, whereby the behavioral potential of the collective is not linearly

related to the inputs of the collective's parts. MADC exploits this aspect of the world, and that is why the capacity to engage in MADC evolved.

As I have mentioned, however, we need not only an account of why the capacity to engage in joint action evolved, but why the capacity to engage in joint action and to use differentiated sub-tasks to do so evolved. To this end, it is useful to explore some recent work by Calcott (2008). Calcott has an interest not only in the way organisms come to cooperate with each other, but how they come to *maximize the benefits of cooperation* when they do. He identifies "internal differentiation and the division of labor" as a factor which enhances the benefits of group action, writing that "Division of labor provides efficiency gains, and the ability to perform novel tasks, and is thus about the generation of benefit" (Calcott 2008, pp. 180-1). He illustrates this phenomenon perfectly when he describes the scenario of a standard Hollywood heist movie: Different agents – the muscle, the computer nerd, the explosives expert – perform different tasks to complete a shared goal with shared payoffs. It wouldn't be possible to nab the diamonds – or whatever – if all of the thieves were computer experts, or strongmen, or they all like to blow things up (pp. 179-80). Rather, because they can each specialize in a sub-task, the collective takes on a greater behavioral capability. The point is that dividing up labor in this way allows individual sub-units to specialize on a certain goal, such that the collective begins to operate, to some non-trivial extent, like a superorganism.

Important to note is that the efficiency and efficacy of dividing tasks is not specific to human or non-human primate cooperation. Calcott, for instance, explicitly links his discussion to work on the major evolutionary transitions. In their work on this exact topic, Maynard Smith and Szathmary note that the evolution of multifunctional enzymes into those which are used specifically for a certain reaction occurred because multifunctional enzymes were not able to specialize in particular reactions in which they were involved (Maynard Smith and Szathmary 1995, p. 110) – multifunctional enzymes would know a little about a lot, but it is more efficient, from a systems-level perspective, to have

specialized enzymes which know a lot about a little. Tying this back to the discussion, my point is that the capacity to engage in MADC evolved because joint action can create greater rewards for all, and because using differentiated sub-tasks to achieve shared goals is the most efficient way to do so.

To rehearse the central point of this section, it is a fact about our world that sometimes when agents team up to achieve a shared task, their behavioral potential can be non-linearly related to their individual inputs. This is why different species have discovered joint action as a behavioral strategy. It is also a fact about our world that when the sub-components of a system specialize by using differentiated tasks, it makes joint action more effective for the individuals agents who specialize, *and*, as we'll see in Chapter III, for the system itself. So, in sum, as I have configured it, MADC is an adaptive target. It is a fact about our world that dorsal appendages help guide objects through water, and this fact motivated the analogous evolution of dorsal fins in fishes and sharks, whose last common ancestor didn't have any sort of dorsal fin (Dennett 1995).

To some extent, we might understand this last section to have been a *distal* explanation about why MADC evolved. That is, I identified very abstract properties about the world which fueled the evolution of MADC. This is right. In the next chapter, I'll look at more *proximate* mechanisms which led hominin lineages to evolve a conceptual apparatus which allowed them to better take part in joint action. I'll do this by re-modeling the *Social Intelligence Hypothesis* so that it takes into consideration the importance that mutualistic cooperation played in hominin cognitive evolution.

Chapter II

Joint Action and Its Role in Hominin Cognitive Evolution

In broad form, the *Social Intelligence Hypothesis* (henceforth, SIH) posits that the primary driver of hominin cognitive evolution was the complexity of ancestral hominin social life.³ The basics of the idea have been around for quite some time, but they took on their modern form most notably with the work of Jolly (1966) and Humphrey (1976). The ideas became more highly codified (and more widely spread) with the publication of Barkow, Cosmides, and Tooby's (1992) anthology on evolutionary psychology, which endorsed the SIH in various ways, and with the publication of two anthologies put together by Byrne and Whiten, both of which were largely devoted to illustrating the importance that "social chess" and "Machiavellian intelligence" played in hominin cognitive evolution (Byrne and Whiten 1988, 1997).

To a rough approximation, the SIH goes like this: Group life has its advantages. The sheer size of a large group of primates reduces predation; also, the more members a group has, the more eyes to spot predators when they do come close (Cheney and Seyfarth 2007). Group life also opens up opportunities to engage in reciprocal relationships, such as trading meat for grooming, which can be beneficial to all parties involved, as we saw in the last chapter (Mitani 2006). But group life is also cognitively demanding. To make the best decisions, one needs to know her place within the group: Her relationships with others, and *others'* relationships with others, both socially (that is, in terms of dominance hierarchies) and genetically. Crucially, one needs to be able to keep track of past encounters to determine who makes a good or naive trading partner, both to avoid getting cheated in reciprocal exchanges and, also, to cheat when it's safe to do so. Those which were able to navigate such a social

³ Throughout, I will use hominin to mean those bi-peds who are our ancestors and who diverged from the lineage which led to modern chimpanzees.

world were at a fitness advantage, and there was, hence, selection for the cognitive mechanisms which would allow for such navigation – cheater-detection, detection-avoidance, and various other instantiations of behavioral and group monitoring cognitive processes. This is why Byrne and Whiten described such group life as having promoted the evolution of Machiavellian intelligence, a term which is now widespread (Byrne and Whiten 1988, 1997). Importantly, the SIH holds that the co-evolution of cheater and cheater-detection is particularly illuminating if we want to understand how our brains got to be so big. As cheater-detection mechanisms improved, better detection-avoidance mechanisms were needed, which required better detection mechanisms, and so on. This dynamic is often referred to as a *cognitive arms race*; hominin cognitive evolution, the model proposes, underwent runaway selection.

I agree with the broad message of the SIH, but argue that it is incomplete. The social lives of our ancestors do seem to have been the predominant force shaping the evolution of their brains. But the SIH puts forward a largely under-developed model of what types of social capacities were selected for. Sterelny (particularly 2007, but also 2003 and 2008a) provides a re-structuring of the SIH, in which niche construction, and the non-genetic transmission of knowledge about one's environment from one generation to the next, play central roles. But Sterelny's model has an important, additional feature: It acknowledges the importance that collective action and cooperation played in our ancestors' evolutionary past *before* hominin cultural complexity exceeded much beyond what we observe in other primates today, a point which is also made by Moll and Tomasello (2007). I'm agnostic about the extent to which niche construction and cultural transmission has shaped hominin cognition. But I think this second part of Sterelny's approach is absolutely on the right track, and in this chapter, I'll expand on his model, utilizing the paleontological record to buttress its claims. Simply, there appears to have been significant selection for hominins which possessed the capacity to coordinate and cooperate with others on joint tasks; there was not selection *just* for those hominins who were able to monitor expanding social networks and exploit them.

In rough form, my argument is as follows. The capacities which allowed particular hominins to navigate strategically within a group setting would have been selected for. And the type of cognitive suite which was selected for would have been that described by models which endorse the SIH – mechanisms which allowed cheater detection, detection avoidance, and various other forms of group monitoring. But hominins also teamed up with each other – to scavenge, defend carcasses, repel predators, raid other hominin bands, and hunt big game. Those hominins which were better able to engage in such joint action were at an advantage. Not only could they repel predators (including other hominins) from attacking, they could coordinate more effectively to raid other hominin bands, and they could access greater resources, through collective hunting, than would be possible through hunting alone. Those hominins with the capacity to coordinate joint action would have, *ceteris paribus*, been at an advantage, and so there was selection for *this* suite of cognitive adaptations, too.

The resultant framework that I'm putting forward is, like Sterelny's, one which expands "social intelligence" to include discussion of how hominins teamed up to defend themselves and navigate their ecological surroundings. Perhaps, therefore, one might argue that I'm no longer talking about *social* intelligence. But I disagree, because these behaviors are only possible in a social scenario, and the fitness advantages of such behaviors, I'll argue, explain to a significant degree how humans came to possess the suite of social competencies we possess today. So I treat the model I'm putting forward as an extension of the SIH, not a replacement of it.

To a large extent, then, the supplements to the SIH which I'll explicate here involve integrating the lessons of the last chapter into a standard, mainstream account of hominin cognitive evolution. As group life is typically configured in most strands of the SIH, interactions basically follow the payoff-structure of a Prisoner's Dilemma game, where there is the opportunity to come out ahead, but only if you can make sure you're not getting cheated, or if you can cheat and get away with it. My aim here is to point out that some important hominin interactions can be described using a Stag Hunt game and

other basically mutualistic models, since joint action yielded enhanced benefits for participating agents in many of the activities cited in the above paragraph. Thus, a goal of this chapter is to re-play the lessons of Chapter 1, and argue that the benefits of coordinated action played an important role in shaping our minds.

In Section II.1, I'll provide an overview of the central theoretical commitments of the SIH, and provide examples of the theoretical and empirical reasons we have to believe that the model is largely accurate. In Section II.2, I'll show why the SIH, while not wrong, is nevertheless quite an under-developed picture of (1) how human brain expansion occurred and (2) the nature of the selective pressures which drove this expansion. I'll argue that the benefits of coordinated action and between-group competition created selective pressure for ancestral hominins to evolve cognitive mechanisms which allowed them to work together cooperatively and to achieve shared goals.

II.1 *The Social Intelligence Hypothesis*

As I mentioned at the beginning of the chapter, the SIH is committed to the idea that the primary driver of human cognitive evolution was the social pressure which group life created for our ancestors. *Australopithecines*, fossils of which date from around about 3.5 mya to 1.5 mya, have brains that are not significantly larger than what we see among modern chimpanzees, about 500 cc. It is not until *H. erectus*, when brain volume tops 1000 cc, some 1.2 mya, that we begin to come close to the brain size of anatomically modern humans, which comes in at around 1350 cc (Lewin and Foley 2004; Cartwright 2008; Larsen, Matter, and Gebo 1988). In short, the *explanandum* of much research in human cognitive evolution generally, and of the SIH in particular, is to provide a self-consistent and empirically-substantiated explanation as to what mechanisms drove the evolution of the massive brain expansion which occurred between our divergence with chimpanzees – often, with the assumption that such an account will illuminate broader questions about the evolution of human cognition, as well. The


SIH posits that ancestral hominins underwent selection for capacities which allowed them to navigate effectively in social settings – to monitor conspecifics, secure resources, find mating partners, exploit naïve group mates, and so on – and that the resultant suite of cognitive adaptations we possess are the outcomes of this process.

The SIH stands in contrast to models which, for example, argue that the complexity and variability of our ancestors' *ecological surroundings* drove hominin brain expansion (see, for instance, Holekamp 2007 for an up-to-date version of this view). This proposal seems unlikely, however, since many animals, including many primates, live in complex ecological environments, but do not show the sort of brain expansion we humans possess (Sterelny 2007, p. 720). Also, certain primates appear to be more adept in their social lives than at navigating their ecological niches. Cheney and Seyfarth, as one example, document that vervet monkeys show quite stunning capabilities to recognize kin, remember who has crossed them in the past, recall how group mates are related to each other, and to act in accord with the dominance hierarchies of the group. But they are unable to, for instance, recognize that a dead gazelle hanging from a tree means that a leopard is nearby, or that python tracks – the only animals which leave such tracks – signify that a python was, in fact, in the vicinity (Cheney and Seyfarth 2008, pp. 127-131). No doubt ecological knowledge played an important role in hominin cognitive evolution, but the role it played was likely not primary, and no such model which posits its primacy can serve as a good framework to understand the degree of brain expansion we showcase.

In this way, the SIH is a needed framework to understand human cognitive evolution. In rough form, the theoretical thrust of the SIH is that group life opens up avenues of opportunity, since groups serve as buffers against predation and give rise to reciprocal interactions. But group life is informationally rich and complex, which makes the right behavioral strategies and predictive heuristics less and less clear. To make smart decisions, one needs to know whom to interact with, using knowledge about one's own social standing, one's own genetic relationships, and also the ways in which


others are related to each other, socially and genetically. The larger the group, and the more information available, the trickier the relationships become (Dunbar 1992, 1997). More sophisticated cognition will be selected for, because making smart decisions means making fitness-enhancing decisions. But this selection process raises the bar ever-higher, so that there is continual selection for *ever more* enhanced cognitive abilities. Traditionally, the SIH posits that a particular manifestation of this dynamic was the co-evolution of cheater-detection and detection-avoidance mechanisms. As cheaters became better at avoiding detection, cheater-detection mechanisms became a more important capacity, and were therefore selected for, which then raised the bar on detection-avoidance mechanisms, and so on. This co-evolutionary trajectory is one instantiation of the general cognitive arms race that the SIH argues was definitive of ancestral hominin group life and was, hence, central to hominin cognitive evolution.

So these are the broad theoretical commitments of the SIH. What evidence is there to support these claims? Consider humans' capacity to detect cheaters. There is nowadays a virtual cottage industry dedicated to what is called the *Wason Selection Task* and its bearing on these issues. Let me illustrate. Below are four cards. Imagine that after eating a spicy meal, someone had beer afterward. Which cards would you need to turn over to determine whether this is true?



The answer can be a bit tricky, certainly not automatic. The cards you would need to turn over would be "Water" and "Chili Pepper". You wouldn't need to turn over the "Beer" card, since it would be fine if someone had a beer but didn't have a spicy meal beforehand. Now consider a variation of this problem

below. If you're a bartender, and you need to make sure everyone who is drinking is 21, then what cards do you need to flip over?



Rather quickly, we can tell that we need to flip the second and fourth cards. This experiment has been repeated many times, across cultures, and it seems that people really struggle with the first type of problem, but not the second. Cosmides and Tooby use these results to argue that we have a specific, fast and unconscious ability to discriminate cheaters. That is, we have a quick, basically automatic ability to figure out who's breaking the rules in a social situation. We know this is "for" cheater detection, because when the basic structure of the problem is presented in a different context (one that doesn't involve cheating), we cannot produce a judgment as accurately or as quickly (example adapted from Baum 2004; Cosmides and Tooby 1992). So this is one item which proponents of the SIH often bring out to support their model: Hominin social life must have played a prominent role in our brain's evolution because we have a highly sophisticated capacity to recognize cheaters, and cheaters are only a perennial threat in group settings.

Pursuing this message from a different angle, Dunbar (Dunbar 1993; Barton and Dunbar 1997; Barrett, Dunbar, and Lycett 2002) has argued that there is significant physiological evidence among primates that social complexity drives the evolution of their brains. Dunbar found that "a neo-cortex ratio (neo-cortex volume divided by the volume of the rest of the brain) was positively correlated with average group sizes of primate genera," with group size "taken to be a proxy for social complexity" (Barton and Dunbar 1997, p. 246). The idea, then, is that as groups become larger, they become

exponentially more complex, and that members of the primate species in question need to devote more energy and more brain matter to organizing all the information such groups exhibit. To make group life worthwhile, in short, one must be able to keep track and monitor the behavior of others and others' relatives and friends, and to remember past interactions and how they played out. Those which were able to utilize the benefits of group life without paying (too many) costs will be those which were selected for.

Central to this basic model is that *predicting* one's behavior becomes increasingly difficult as information – particularly, information about agents' relationships – becomes richer and richer, owing to the expanding network of social relationships within the group which results from this expansion. As Cheney and Seyfarth have documented, baboons are quite good at knowing where within the social and genetic landscape they fit: who outranks them and whom they outrank, and who their close relatives are. But many primates, baboons included, know not only their own social rank and relations, but also how *others* are ranked compared to others, and how *others* are related to each other. What's particularly interesting to note is that rank is often orthogonal to genetic relatedness. Thus many primates have to take into account these over-lapping social networks in order to make good decisions about how to interact with group mates.

Nuanced, complicated networks like these, Sterelny argues, motivate populations to develop the capacity to use *decoupled representations* (Sterelny 2003, pp. 51-8). When agents' lives are intricately tied up with those of others, the best strategy might be to develop a behavioral regime by which particular cognitive representations are *not* tied to particular behaviors - hence, they are decoupled from them. This allows action to be mediated by more than one environmental cue, and allows an agent to utilize all sorts of information – memory of past interactions, knowledge of an agent's relationships, etc. – before committing to a decision. So if a group mate ranks lower than you and is near a food source, the initial, pre-programmed cue might be to bully him and take the food for

yourself. But if this conspecific is closely related to baboons which are bigger and outrank you, then the safest strategy might be to leave him alone. Being able to make socially smart decisions like this, the SIH holds, is what determined which genotypes were represented most in future generations. Since the conceptual apparatus which would have allowed for better strategizing is partially a result of pre-programmed cognitive biases, the SIH posits, each new generation is poised to be a bit more socially intelligent, on average, than the last.

11.2 Integrating joint action and the Social Intelligence Hypothesis

There is clearly much to recommend the basic structure of the SIH. Primate group life is complex, and the larger the group, the more complex it becomes. What's more, group life is always under siege from cheaters, who undermine the reciprocal advantages of cooperative interactions and threaten the stability and advantages of living in a group setting. The runaway process of cheater-detection and detection-avoidance, along with other stakes-raising cognitive advancements among hominins, might very well explain a big part of why our brains are so big, so metabolically expensive, and so hard to birth.

But most strands of the SIH come up short. Simply, they do not pay adequate attention to the evidence we have which suggests that there was significant selective pressure on our ancestors to evolve a cognitive apparatus which would allow them to coordinate with conspecifics effectively. Being able to detect cheaters, remember and recognize group members, and behave strategically with respect to one's relationship with others and their relationships to each other, are cognitively demanding and impressive. But there is evidence that the selective advantages of coordination and joint action played a central role in our ancestors' cognitive evolution as well. Coordination of this sort allows groups to defend carcasses which they come upon while scavenging, team up to hunt big game, and to defend against and overtake competing hominin bands. What's more, individuals who made up groups which

were able to cooperate in this way would have produced offspring at higher rates than those who could not, for such individuals would have been better defended, and would have had access to resources (through collective hunting) which only become possible when agents coordinate behavior. In short, my message is that the SIH is right to stress the importance that our ancestors' social lives played in shaping our present-day brains, and it is true that cheater-detection, detection-avoidance, group monitoring, and social knowledge were big parts of this evolution. But there was also selection for the capacity to coordinate action with others. For the remainder of this section, I'll look at the evidence we have for believing this is so.

In Chapter 1, I drew attention to ethological and laboratory data which suggests humans are rare among primates in their ability to coordinate with others to accomplish shared goals, and to use differentiated sub-tasks when they do (Tomasello, Carpenter, Call et al. 2005; Warneken Chen, and Tomasello 2006). Here, I'll argue that the complexity of this behavior, and the capacity which underlies it, suggests that it evolved because of the selective benefits it conferred on those who possessed it. Let me detail a bit more the phenomenon I'm interested in, then I'll explain why its existence suggests coordinated action played an important role in hominin cognitive evolution.

To take part in joint action, an agent must not only understand how her sub-task will result in the completion of the collective task, but how others' tasks come together to fulfill the ultimate goal of the collective, too. Gilbert has illustrated this phenomenon through her analysis of what goes into two people walking together (Gilbert 1990, 1992). When two people go on a stroll down the street, each is not simply walking by herself. Rather, the behavior is coordinated – “oh, I'm going a bit too fast” or “I need to pick up the pace” – so that the two are engaged in the behavior *jointly*. As Gilbert (1992) puts it, “In order for X and Y to be doing A together, some condition of behavioral appropriateness must presumably hold, but both X and Y must have a special conception of what is going on. Crucially, each

must intentionally act...as a constituent of a plural subject of a certain goal, a subject whose other constituent is the other person in question” (p. 166).

Jackendoff (2007) addresses this phenomenon as well. When two individuals move a couch, the behavior is not one individual lifting one end of the couch, while the other simply happens to lift the other end. The behavior only makes sense insofar as both actors know that *the other* knows the goal, too. As he puts it, such joint action is not “*I* am doing such-and-such...and *you* are doing such-and-such’...but ‘*We* are doing such-and-such out of a joint intention, and my role in it is such-and-such and your role in it is such-and-such” (Jackendoff 2007, p. 176). Recall the experiment run by Warneken et al. (2006) discussed in the last chapter, in which children were paired up with an experimenter who would stop participating in the middle of experiment. The children prompted the experimenter to continue – through gestures and vocalizations – because they realized that the goal could only be accomplished if both tasks were completed. Alternatively, the chimps pursued their own tasks with greater enthusiasm, but had little inclination to motivate the experimenter to continue with her sub-task. It was as though the chimpanzees didn’t recognize that it was a *joint task*, and *required* two people for it to work.

Why have I mentioned this sort of behavior? Because it is exceedingly complex. Coordinating action with other animate objects is non-trivial, and the conceptual capacities which underpin this behavior are more complex still. They allow us to represent our individual actions within a larger framework of a joint action, and to configure the role that another’s (or others’) actions play in achieving the collective goal. My argument is that, like most complex traits of this sort, it is fairly safe to posit that this capacity emerged because it was selected for. What’s more, complex traits like these – and eyes, and ears, and camouflage – are typically the result of drawn-out processes of selection in which incremental improvements yield incremental advantages, as opposed to short *macro-mutational* events in which a few lucky substitutions create hugely adaptive traits (Dawkins 1986). That is, it is far more likely that the capacity to engage in joint action was built up and refined slowly, over time,

because gradual increases in our ancestors' ability to coordinate with others in this way garnered incremental fitness advantages for them. This gradualist formulation of the evolution of such capacities seems all the more plausible when we consider that nascent forms of this behavior exist in chimpanzees, suggesting an incremental process of the capacity to act in this way could have, at least in principle, been the result of a gradual process of evolution, one re-enforced by the selective advantages of this sort of joint action. In sum, the first reason we have to suppose that coordinated action was an important part of our species' cognitive evolution is that we have salient, complex capabilities in this area that are early to emerge and that we don't share with other primates. There was apparently selective pressure to evolve the capacity to cooperate in this way, or else, such a complex capacity wouldn't have evolved.

But there is other evidence which suggests coordinated action played an important role in hominin cognitive evolution. This comes about when we consider the foraging, hunting, and inter-group competition these populations engaged in. There are cut marks on bones, probably the result of *Australopithecines*, which date to about 2.5 mya (Lewin and Foley 2004), and there is evidence that such scavenging practices continued among some groups through 1.5 mya (Leakey 1994, p. 71). It is doubtful these are the result of hominins teaming up to down large game, however, then cutting up the remains. Rather, these *Australopithecines* were probably scavenging the kill-sites of big cats. This behavior is non-trivial, though, for it is poised to create a scenario which motivates joint, coordinated action. If a group of *Australopithecines* has come across a carcass, how is it to fend off other scavengers, like hyenas? Notice that one bi-ped defending a carcass doesn't have much of a chance. But if there are two or more involved, defense becomes much easier, since all can surround the kill, giving rise to an interesting instantiation of the Stag Hunt game (Sterelny 2007). The need for effective coordination in such scenarios is all the more necessary when we consider that fossil evidence indicates that the cats

which lived in Africa during this time were in fact much larger than they are now (Kruuk 2002, pp. 103-112).

But the advantages of joint, cooperative *hunting*, more so than collective defense, seem to have played an absolutely essential role in hominin cognitive evolution. The benefits of teaming up with others to hunt in this way are quite salient, as we saw in Chapter I. To re-use an example from that chapter, Smith and Jones might each be able to hunt three hares over the course of the day, but if they team up, they can start hunting larger game, like stag, which is worth more than the collective yield of those would-be hare hunts. Chimpanzees hunt in this way, as we saw, but typically they hunt rather small game, like monkeys. Even though they benefit from such hunts, particularly since such protein-rich meat is a rare aspect of the chimpanzee diet, the returns of collective hunting for early hominins were far more striking. Of course, it is difficult to place precisely when hominins began to forego scavenging for collective hunting. But it is not entirely closed-off to analysis. In 1979, bones with concussive fractures and scrapes from tools were found at Olduvai gorge, a site associated with *H. habilis*, with the bones dating to about 1.5 mya. In Germany, researchers found spears which dated to around 400,000 years. So we can be sure that at some point between the *Australopithecine* bone scrapers from 2.5 mya, and these spears in Germany, humans began to down large game (Lewin and Foley 2004).

In either case, the dates are not particularly important in this present context; what matters is that we have evidence that collective hunting was a part of the hominin behavioral regime, for this is precisely the sort of evidence which supports my argument that coordinated action played a role in hominin cognitive evolution. How so? In a chimpanzee hunt, those which are better able to hunt will get more of the kill, while those which don't hunt are actively excluded from the meat provisioning process (Boesch 2001). So on the one hand, there would have been an individual-level benefit to possessing a cognitive apparatus which would allow you to coordinate action with others.

But there is another reason such a capacity would have been selected for. Namely, groups of agents who were able to cooperate in this way would have out-competed those groups which didn't, in two different ways.⁴ First, groups which were better able to coordinate joint action with each other and partner-up to achieve shared goals would have repelled raids by other hominin tribes and would have been better and more successful at raiding other hominin tribes. Commenting on this phenomenon, Cartwright writes, "It is possible in fact that the greatest threat to our ancestors was the existence of other hominin groups tightly bound in a system of mutualistic cooperation and willing to instigate violent conflicts" (Cartwright 2008, p. 205). So that is one sense in which coordinated action would have been selected for in ancestral hominin environments: Those hominin bands which were able to repel attacks better and launch attacks in more sophisticated, coordinated ways would have been at an advantage over others, and the cognitive capacities which would have allowed for such behavior would have been selected for.

But between-group conflict doesn't have to exist in the upfront sense I've just described. It can also be indirect. Even in such indirect competition, those individuals which are better able to coordinate action and partake in joint, cooperative tasks will be those that, *ceteris paribus*, are selected for and therefore come to dominate the population. Let's return to the Stag Hunt game to see how this works. If one group of early hominins hadn't developed the conceptual capacity to coordinate action effectively enough to hunt larger game, the individuals in that group would not have had access to the relative wealth of resources which a group of individuals which *had* developed that capacity would have had. It is a big difference if one set of groups continues to forage, while another is engaging in pro-active kills. The latter group will not only be more poised to repel predation by hominins and big cats, it will also simply be accessing more calorically rich resources. Individuals in such groups will live relatively longer and be able to provide for more offspring, for they will have access to more food and be amongst

⁴ These two ways are analogous, though not perfectly, to *multi-level selection 1 and multi-level selection 2*, which I'll address in detail in Chapter III.

conspecifics who are capable defenders against predation. Such groups are keenly poised to produce offspring at a greater rate than those who wouldn't be able to access those types of resources, because they had not come out of the "hunt hare" strategy, for whatever reason, to the "hunt stag" strategy which yields higher returns.

This coheres nicely with a model which attempts to explain why Neandertals went extinct. Within the past sixty-thousand years, Neandertals and anatomically modern humans lived in Europe at the same time. But around forty-thousand years ago, the Neandertals start to disappear. The latest fossils anyone has found are from a cave in Spain around Gibraltar, and date to about thirty-thousand years ago (Finlayson, Pacheco, Rodriguez-Vidal 2006). One prominent writer, Trinkaus, has argued that humans and Neandertals inter-bred, and Neandertal populations disappeared because they were integrated into those of anatomical moderns (see, for a recent appraisal, Soficaru, Petrea, Dobos, and Trinkaus 2007). But recent anatomical and genetic analysis have run up against this claim (see Weaver and Roseman 2006; Callaway 2008). Others, acknowledging this, suggest the Neandertals disappeared because anatomical moderns killed them all.

An alternative model, proposed by Kuhn and Stiner (2006), is that different reproductive rates could be a big part of the story. Specifically, they claim that Neandertals went extinct because they didn't showcase a sophisticated division of labor, while their human neighbors did. Such a division of labor allowed anatomical moderns to hunt and gather more efficiently, enter into more sophisticated reciprocal relationships, and allowed men and women to specialize in sub-tasks, such as caring for developing young. Dividing labor in this way, as we saw in Chapter I, allows agents to more effectively "generate benefit" in a manner that is potentially advantageous for all. Kuhn and Stiner argue that this sort of division of labor would have meant that individuals in such groups were yielding more resources and reproducing at a higher rate than their Neandertal contemporaries. This reproductive advantage alone, they argue, could account for the Neandertal's extinction.

This claim is buttressed by computer simulations Zubrow ran, in which he sought to determine what type of reproductive advantage was needed for one population to over-take the other in a meta-population of two groups. Surprisingly, he found that, over a millennium, a two-percent advantage was enough to do so (discussed in Leakey 1994, p. 98). To be clear, as I've described it, I *do* imagine our hominin ancestors engaged in direct conflict. But the take home message is that direct competition is not the *only* sense in which groups can compete with each other, and not the only way in which advanced coordination and cooperation would have been selected for among ancestral hominins. Indirect competition, too, likely played a key role in this process.

To sum up, then, what I've proposed in the last few pages is a modified sketch of the SIH, one in which I agree that the primary vehicle of human brain expansion in particular, and the evolution of human mentality generally, came about because of the complex social worlds in which our ancestors found themselves. But I've argued that most formulations of the SIH put too much emphasis on the role that so-called Machiavellian intelligence played in this process. My criticism of the SIH, to a large extent, is that the sort of intelligence it focuses on, this Machiavellian sort, is in fact quite phylogenetically old, and mischaracterizes or (more accurately) simply ignores the truly important selective pressures our ancestors faced *after* they broke off from the lineage which has led to modern chimpanzees.

Moll and Tomasello make this point when they write that "A reasonable proposal is therefore that primate cognition in general was driven mainly by social competition, but beyond that the unique aspects of human cognition— the cognitive skills needed to create complex technologies, cultural institutions and systems of symbols, for example—were driven by, or even constituted by, social cooperation" (Moll and Tomasello 2007, p. 1). What came about after we split from the lineage which led to chimpanzees was a different class of social problems to address. The pressures of living in a group were still important, but there was selection for the capacity to coordinate with other hominins to access resources, repel predators, divide labor, and over-take other hominin bands. These might have

been issues that primates dealt with before the split, but not to the extent that they were issues a million or so years after that split occurred. If we want to explain why hominin brain expansion is unique, and how it happened, we should integrate into our models of this explanation an account of the unique cognitive capacities which underpin human social life – for instance, the capacity for joint action – and how the evolution of these capacities occurred.

Chapter III

Groups and Emergence, Reduction, and Selection

The goal of the first chapter of this paper was to establish in what ways joint action was beneficial to individual organisms, to address possible conceptual snags which arise when thinking about the evolution of joint action, and to argue that the benefits of joint action generally (and joint action with differentiated sub-tasks in particular) help us to understand why humans have the early-emerging capacity to cooperate with others in the manner carved out by the definitional parameters of MADC. The goal of the second chapter was to take these lessons and apply them more concretely to an analysis of hominin cognitive evolution. I did so by explicating the SIH, endorsing its basic message, but illustrating, along the lines of Sterelny, that it needs to integrate the evidence we have that there was selection among hominin populations for a capacity to engage effectively with others to achieve shared goals.

In this chapter, I am interested in addressing a question which arises when we talk about agents teaming up in a joint task in such a way that the *collective* showcases a behavioral potential which is non-linearly related to the sum total of its constituent parts. To re-use a previous example: Smith and Jones can each snag two hares over the course of the day. But when they team up, they can down larger game, like stag. But neither has much of a chance of downing stag alone, and the yield of a successful stag hunt is more than the sum total of the hare hunts. So what sort of thing is a group engaged in joint action? Should we pay attention to the group *itself*, or understand the group as an ignorable by-product of *individual* interactions?

My argument is that we have good reason to treat such a group as a new, distinct entity in the world, one which need not be decomposed by putting forward a description which merely references

the individuals which constitute it. Such a group becomes a “plural subject,” and it is therefore difficult to explain what is happening at the group level in terms of individual-level inputs. So I will argue that there is no *a priori* reason to reduce such a group; instead, we should reserve the right to treat it as a distinct entity. Now, in some cases, it might be heuristically useful to describe a group engaged in joint action as a collection of individual actors. (Say, for instance, when describing why food is distributed in this or that way after a hunt.) There is nothing wrong with this. My thesis in this chapter is that in some *other cases* (say, when discussing between-group competition among hominins), it is perfectly legitimate to treat the group as a distinct entity itself.

Why? There are three reasons this is so, and each stands on its own. First, as I’ll detail in Section III.1, the behavioral potential of a group of this sort is non-linearly related to the input of its constituent parts; to reduce the group behavior to individual behaviors is to erase reference to the important emergent property which occurs during joint action, and this property is essential to understanding the evolution of the behavior and the capacities which underlie it. Second, it is quite unclear to me why a reduced depiction of the phenomenon is preferable in the first place – or why, for that matter, reduction is treated as the default procedure in these circumstances. Nevertheless, I’ll address this “fine-grain preference,” to use terminology from Jackson and Pettit (1992), and suggest why it is wrong.

Third, and finally, a group engaged in joint action is a distinct unit of selection, and therefore we can “pick out” such groups in our descriptions as independent entities. This third argument comes in two parts. In Section III.2.1, I’ll explicate the conditions for natural selection, and point out that just as individuals can evolve because of differential reproduction and inheritance, so can other items. What matters is *that* the conditions of natural selection are met, not *what* meets the conditions. With this view in hand, I’ll posit that certain groups – not *just* the individual members of these groups – undergo between-group competition and can be selected for, and evolve in part because of selective forces

operating at the group-level. This configuration of selection, typically called *multi-level selection*, posits that selection occurs throughout the biological hierarchy, and therefore acts on genes, individuals, and groups simultaneously. This makes me a *pluralist* with regard to the levels of selection, though I should mention that I will endorse (and utilize) a *realist* approach to the causal processes of this multi-level selection framework. That is, selection is having a real causal effect on what evolves at each level of the hierarchy (see Sober 2008).

Having established that groups themselves *could* be selected for (at least in principle), in Section III.2.2 I'll detail the ways in which this actually occurs when we're talking about groups engaged in joint action. Specifically, the two ways groups of this sort are objects of selection involve processes which are often demarcated as *multi-level selection 1* and *multi-level selection 2* (henceforth, MLS1 and MLS2), a construct which goes back to Damuth and Heisler (1988). Put briefly, the individuals within a group in which joint action is a behavioral strategy will be at a relative fitness advantage over individuals within groups which don't utilize this behavior. In this way, groups compete with each other in the relative offspring each group contributes to the total population – this is MLS1. MLS2 occurs when groups compete with each other by producing daughter *groups* at various rates. This process describes the selection of groups which utilize joint action, too, as I'll detail later on. If we were to automatically reduce joint action into individual-level behaviors, we would ignore this between-group competition. This is part of the reason I argue here that, in our descriptions of their evolution and behavior, we should reserve the right not to decompose such groups into smaller parts.

Section III.1 *Joint action and emergent behavioral potential*

As I have been stressing for the past two chapters, humans and some other animals are capable of engaging in a type of cooperative behavior which results in an exponential increase in returns. We saw this initially in Chapter I, as we looked at yellowtail and chimpanzee group hunting. I tied this

behavior to Skyrms' work on the Stag Hunt game, and used Calcott's extrapolation of these ideas to make sense of MADDC (Skyrms 2004; Calcott 2008). I argued that we can make sense of this type of cooperation because each individual is poised to benefit from the returns of the joint action.

Here, I want to push this analysis of joint action further and argue that when a group of agents are engaged in it, they in fact constitute a new entity in the world. Above, I briefly highlighted the three reasons I am making this claim. In this section, I will focus on the first of those three. Namely, that we can treat such a group as a new, distinct entity because the behavioral potential of the group is not equal to the sum of the behavioral potential of its parts. For instance, to bring to the forefront again an example I have been using for the past two chapters, two individuals can each hunt hare by themselves, or they can join forces and thereby make possible hunting larger game, like stag. What the collective is able to accomplish exceeds the sum of what both individuals within the collective could accomplish on their own. For that reason, we're forced to resist the temptation to describe a group of agents engaged in joint action as equivalent to a description of the behaviors of each of the group's members. For when the behavior of an individual is placed within the context of other individuals, the collective behavior takes on qualities which are not instantiated on the individual-level. To be sure, this is not magic, and I hasten to add that I am using "emergent" in a strictly materialist sense.

Another reason we have for treating a group engaged in joint action as a distinct entity is that there is simply no good, *a priori* reason to reduce it. Why is the default position a reductive one? Indeed, the ramifications of this reductive explanatory methodology are pointed out lucidly in Jackson and Pettit's (1992) piece on this topic. They write that "given that scientific levels take us to more and more micro or small-grained structures, it should be clear that the small-grained preferences will have a general implication of a dramatic and controversial kind. It will lead us to prefer a lower level of scientific explanation to any higher level, and to prefer micro-physical explanation to all else" (p. 9). That is, if an individual-level account of group behavior is somehow more accurate than an account at

the level of the group, then we would be obliged to reduce this individual-level account down further still.

What is important, as I see it, is whether the level at which one is describing a phenomenon provides a useful picture of the phenomenon. For instance, Pinker (2002) rehearses a popular example when he points out that we *could* explain the start of World War I in the language of particle physics, but it is not clear that such an explanation would do anyone much good. Likewise, we could, in principle, describe group-level behavior in terms of individual-level behavior – like I said, there’s nothing magical happening with the emergent property – but if we do so, we’re drawing attention away from the non-linearity between a collective’s behavioral potential and the sum of the behavioral potential of its constituent parts, which is *precisely* the aspect of the phenomenon which explains why the individual-level behaviors evolved and is, hence, the phenomenon we’re most interested in. Along these lines, Jackson and Pettit endorse what they call *explanatory ecumenicalism* because “relatively coarse-grained explanations...can provide information on causal structure that more fine-grained stories suppress” (Jackson and Pettit 1992, p. 19). So if a description at one level helps us organize our concepts better – which it clearly does in the case of joint action, because the emergent potential of the group explains why agents team up to engage in it – then we’re perfectly justified in staying at that level when we describe the behavior of the entity. For one, there is no reason that a description at a lower level is more accurate, and a lower-level description might “suppress” aspects of the phenomenon which are especially important.

Along these lines, if one is tempted to question my allegiance to a group-level description of behavior in cases of joint action, then I’d point out that a gene-level description commits the same error. For we could describe genic evolution in terms of sugar molecules, phosphates, hydrogen bonds, and base pairs, but we choose not to, because such descriptions are conceptually hard to grasp, and they don’t present a truer story of what is occurring in any case. This is precisely what I am saying about joint

action. It is easier to understand the phenomenon from a group-level, because it is at that level that behavioral potential takes off, so that it can benefit the individuals which constitute the group. In some sense, then, my claim is that the heuristic value of *not* dropping down a level in an account of group behavior would merit group-level description. For not only is there no *a priori* reason to suppose that a reduced explanatory story is somehow more right than one at a higher level, it is clear that such a reduced story in fact glosses over the important points that certain forms of collective action showcase emergent properties.

Section III.2 *Groups are units of selection*

III.2.1 *Substrate neutrality and multi-level selection*

The second line of argumentation I will pursue here is that we should be able to treat a group engaged in joint action as a distinct entity in the world because it is a distinct unit of selection. That is, a group engaged in joint action follows an evolutionary trajectory that does not necessarily correspond to the evolutionary trajectory of its individual members. Groups propagate and radiate, and so too do the individuals who make up these groups. But these processes are different processes, and often what is good for the group will not be good for the individual, and vice versa (Sober and Wilson 1998). Even *when* individual-level fitness interests correspond to the fitness interests of the group, it does not follow that this means there are not distinct causal forces of selection operating at both levels simultaneously.

How is this so? This claim hinges on an important point, that natural selection is a substrate neutral process (in the sense of Lewontin 1970; Dennett 1995). This means that any sort of media – alive, dead, inorganic, etc. – is liable to evolve in part because of selective pressure. Consider a standard definitional explication of natural selection: Selection occurs when (1) individuals are differentiated, (2) these differences affect reproductive success, (3) and these differences are inherited (Sober 2000, p. 9). Looking at feature (1) of the above formula, the word “individuals” can refer to individual FOXP2 genes,

individual marmosets, individual religious rituals, and so on. In other words, given that *any* objects meet the above three conditions, selection is poised to occur. Put simply, what matters is *that* the conditions are met, not how.

As an illustration of this, consider an example from Sober (2000, pp. 214-215). Why are businesses profit-maximizers? There are at least two available explanations: one, business leaders, being “rational actors,” are sufficiently plastic so as to adjust to market demands. Or two, profit-maximizing firms dominate the market because those that *weren't* profit-maximizing were unable to survive, and hence our current representation of business strategies is biased toward representing profit-maximization. Notice that the firms’ strategies have nothing to do with genetic survival. Indeed, these strategies might be housed within a computer program, in which case the “leaders” of the businesses have no genes at all. A similar non-biological process of selection is posited to explain the existence and retention of certain scientific beliefs, in what is typically called *evolutionary epistemology*. According to this model, scientific ideas survive and get passed along according due to their practical and conceptual superiority. This is the sense in which selection is substrate neutral, and the sense in which selection is liable to act on items other than those typically associated with evolutionary biology.

I mention these examples because this theoretical framework is important when we are considering group selection. Just as all sorts of media can undergo processes of selection, so long as the conditions of selection are met, so too can groups of individual organisms evolve in a manner according to natural selection. Sober and Wilson (1998), as one example, have pursued this idea, and have argued that selection operates on all three of the traditional levels of the biological hierarchy – genic, organismal, and group – simultaneously. This framework, the multi-level selection framework, simply follows from acknowledging that natural selection is substrate neutral in the sense I have described; so long as each of the levels of the hierarchy meets the conditions for selection, selection will take place. As Sober and Wilson sum up their multi-level selection approach, “Natural selection occurs when genes

differentially survive and reproduce within single individuals, when individuals differentially survive and reproduce within single groups, and when groups differentially survive and reproduce within a global population” (p. 89). To relate these points back to the message of this chapter, my argument is that we should treat a group engaged in joint action as a distinct entity not only because it exhibits behavior which exceeds the collective input of its members, but because the group itself is an entity upon which selective forces operate. In the Section III.2.2, I’ll detail how this occurs.

III.2.2 The ways groups are units of selection

Now, I have just spent the last few paragraphs arguing why groups could be the target of selection *in principle*; namely, because selection can operate on a variety of media, so long as they instantiate the conditions of selection. Here, I want to explain how this actually occurs in the case of groups engaged in joint action. A useful theoretical division which is helpful for this task is that between MLS1 and MLS2 (Damuth and Heisler 1988; Okasha 2006, pp. 53-9). These constructs are not particularly complicated, but do require a fair amount of explication. Let me begin with an example of MLS1. We might imagine a population of ten different groups, which are reproductively isolated from each other. If some of these groups are composed largely of altruists, while others are composed largely of selfish individuals, the reproductive output of altruistic groups will be greater than the reproductive output of selfish groups (though the selfish individuals within the altruistic groups will increase in proportion to altruists). Now, if after reproduction, the individuals of these groups merge into a meta-population, the altruists will outnumber the selfish individuals because the individuals within altruistic groups reproduced at higher rates than those within selfish groups. So, the number of altruists can increase in the meta-population even though the proportion of altruists is decreasing within each particular group (Sober and Wilson 1998). If we understand this as a process of group competition

and group selection, it is group selection in the sense of MLS1: Competition is cashed out in terms of how many individuals each group contributes to the entire population.

MLS2 differs from this in an important way. With MLS2, groups compete with each other not in the number of offspring each group produces, but in how many daughter *groups* each group produces. At least in principle, therefore, MLS2 could occur intra-generationally, whereas for MLS1 to take place, it would be requisite that at least one reproductive cycle had occurred. In short form, the measure of “success” among groups within the framework of MLS2 is not how many individuals from each group have now entered the meta-population, but how many new groups of a certain sort have entered the population of other groups. Okasha depicts this distinction clearly when he writes that “First, a collective’s fitness could be defined as the average or total fitness of its constituent particles; so the fittest collective is the one that contributes most offspring *particles* to future generations. Secondly, a collective’s fitness could be defined as the number of offspring *collectives* it leaves; so the fittest collective is the one that contributes the most offspring collectives to future generations of collectives” (Okasha, p. 54, emphasis in original). The former type describes fitness in the sense of MLS1; the latter, in the sense of MLS2. With MLS1, we might treat the groups as the *interactors*, while the members of the group are the *replicators*; with MLS2, the group themselves are the replicators (Lloyd 1992).

In a Stag Hunt game, unlike in a Prisoner’s Dilemma game, there is no tension between individual-level and group-level interests. So consider MLS1. If a group of humans, let’s specify it as Group A, engages in joint action, and Group A is housed within a meta-population of hominins which do not, which we’ll call Groups B, C, and D, then the reproductive output of the group which engages in joint action will be higher, *ceteris paribus*, than the reproductive output of those groups which do not. For group A will access resources which are simply unavailable to Groups B, C, and D, because joint action allows for behavioral opportunities (for example, “hunt stag”) that greatly outweigh the fitness benefits of isolated action (for example, “hunt hare”). Thus, from the viewpoint of the meta-population,

Group A is contributing proportionally more offspring to future generations. So this is one sense in which certain groups are selected for: Within a meta-population, individuals are contributing variously to the genetic pool of future generations based on the type of group they are in.

It should be noted, however, that success in the sense of MLS1 can result in success in MLS2 when discussing joint action, a point to which I will now turn. Again, imagine a meta-population composed of a Group A which can engage in joint action and a bunch of other groups of hominins, Groups B, C, and D, which do not. We just saw that Group A produces more relative offspring than the other groups. But Group A also reproduces itself. That is, it gives rise to daughter groups which showcase joint action. How so? The process, like above, is inter-generational: The capacity to partake in token instances of joint action relies in part on a particular genetic makeup. Not just any animal can learn how to cooperate with others in this way, as we saw in earlier chapters. The offspring which enter into the population from Group A will be pre-disposed to engage in instances of joint action to a greater extent than those offspring which came from groups which were incapable of cooperating in this way. If we use “the capacity to engage in joint action” as the marker of relatedness between parent and daughter groups, then joint action is poised to evolve in a manner described by MLS2. So a group engaged in joint action competes against the reproductive output of other groups (MLS1), and also spawns new groups of the same sort (MLS2).

In rough form, the argument of this section is that we should be able to treat a group engaged in joint action as a distinct entity because there is selection which occurs at the group level, not just selection which occurs at the level of the individuals within such a group. What’s more, my arguments stand in contrast to those of, for instance, Sterelny and Kitcher (1988), because I have shown that these selective processes have real causal influence on the type of entities which evolve, and it is not a matter of preference how we talk about evolution occurring at different levels of the biological hierarchy. I first substantiated this claim in principle by illustrating the reasons why selective forces can operate on

groups - namely, because selection is substrate neutral. From here, I identified two mechanisms which actually describe how selective forces would operate on a group engaged in joint action – namely, MLS1 and MLS2.

These points relate quite nicely to some of those mentioned in the last section. Specifically, reduction in the case of a group of agents engaged in joint action is problematic for the additional reason that the selective process occurring at the group-level is different from that occurring at the individual- or gene-level. We cannot easily reduce a group of agents engaged in joint action to a story about those individual agents, because the group itself is an entity in the world which is being selected for and evolving. If we were to reduce MLS1 and MLS2 to a story which appealed solely to individual-level selection, we would, in fact, be committing an empirical error, since we would thereby be describing a *different* evolutionary process.

In sum, for three independent reasons, we should treat a group engaged in joint action as a distinct entity, one that should not be reduced to its smaller parts, and one that merits attention itself. The first reason was that the behavioral potential of a group engaged in joint action does not neatly reduce to the inputs of its members, so that a reduction does a disservice to an important emergent property of the collective's behavior. The second reason is that there is no good, *a priori* reason to prefer a reduced explanation in the first place. The third is that a group engaged in joint action undergoes selective pressures itself, and we saw the mechanisms by which this occurs. Thus, describing these processes at a lower level, such as an individual level, would pick out a different process of selection, for evolution occurs at various levels of the hierarchy.

One can fairly ask why the three reasons I have picked out are *sufficient* reasons to reserve the right to treat a group engaged in joint action as a distinct entity in the world. To this, it is difficult to provide a definitive response, since we mortals lack a clear-cut list of which entities go into which ontological categories. But since we are here concerned with evolutionary biology, we can see how two

of these reasons are at least *good* reasons, if not necessarily definitive ones. For one, we want to explain animal behavior out in the world, and so in our model of cooperative hunting, it is important to recognize that group behavior exceeds the behavioral potential of the individual inputs which produce it. We lose this when we just look at individual action, and so we lose our account of why this behavior occurs. Also, since groups engaged in joint action are units of selection, something also would be missed if we acted as though the individuals within these groups are the only units of selection. An additional reason I gave – that there is no *a priori* reason to reduce entities of this sort to begin with – is far more epistemological, and to justify why it explains why groups engaged in joint action are distinct entities would be to re-state the argument I have already put forward.

Conclusion

To sum up, in this paper I have tried (1) to establish that the benefits of joint action can help us understand human's rare capacity to take part in joint action with others, that (2) joint action played an important role in hominin cognitive evolution, such that, to be complete, the SIH needs to incorporate that fact, and (3) a group of agents which is engaged in a joint task showcases an interesting property whereby the individuals are no longer two (or more) *individuals* working independently; rather, they become a plural subject, one which needn't necessarily be reduced to a lower level of description.

As I mentioned in the *Prologue*, the goal of my paper was to apply an abstract conceptual idea – that there is an emergent property of agents working together – to very concrete issues about the nature of hominin evolutionary history. By doing so, I hope I was able to show that the individual-level view of selection that predominates in the SIH, in which hominins are depicted as conniving champions of social chess, always watching their backs for cheaters, and cheating when they get the chance to, is only part of the story. There was also selection for hominins to coordinate behavior and take part in joint action, and this occurred because our world is such that, sometimes, when agents coordinate with

each other, there is an emergent behavioral potential, so that each agents stands to benefit from the interaction. To be sure, I do not think this view of our ancestors is any more appealing from a moral standpoint, since it, as I argued, probably resulted at least in part from hominin tribes massacring each other. But it's more appealing from a theoretical standpoint because it helps us to explain how humans have the ability to so easily coordinate action with each other.

Bibliography

- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Bailey, S. (2008). "Human Evolution," a course taught at Tufts by Professor Stephen Bailey, Fall 2008.
- Barkow, J.H., Cosmides, L. and Tooby, J. (1992). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press.
- Barrett, R., Dunbar, R. and Lycett, J. (2002). *Human Evolutionary Psychology*. Princeton University Press.
- Barton, R., and Dunbar, R. (1997). Evolution of the Social Brain, in *Machiavellian Intelligence: Extensions and Evaluations*. Oxford University Press.
- Baum (2004). *What is Thought?* MIT Press.
- Binmore, K. (2007). *Game Theory: A Very Short Introduction*. Oxford University Press.
- Boesch, C. (2001). Cooperative hunting roles among Tai Chimpanzees, in *Human Nature*, 13:1, pp. 27- 46.
- _____. (2005). Joint cooperative hunting among wild chimpanzees: Taking natural observations seriously, as Peer Commentary to Carpenter, M., Call, J., Behne, T., and H. Moll. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sci.* 28: 675-691.
- Boesch, C., Boesch, H., and Vigilant, L. (2006). Cooperative hunting in chimpanzees: kinship or mutualism? in Kappeler, P. and van Schaik, C.P. (eds.) *Cooperation in Primates and Humans: Mechanisms and Evolution*. Germany: Springer.
- Buller, D. (2006). *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, MA: MIT Press.
- Byrne, R., and Whiten, A., eds. (1988). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Oxford University Press.
- _____. (1997). *Machiavellian Intelligence II: Extensions and Evaluations*. Oxford University Press.
- Brown, J.L. (1988). Cooperation – a biologist's dilemma. In *Advances in the Study of Behavior* (ed. J.S. Rosenblatt), pp. 1-37. Academic Press: New York.
- Callaway, E., (2008). Neandertal genome already giving up its secrets. *New Scientist*, 10 December 2008.

- Cartwright, J. (2008). *Evolution and Human Behavior: Darwinian Perspectives on Human Nature*. Cambridge, MA: MIT Press.
- Calcott, B. (2008). The other cooperation problem: generating benefit. *Biology and Philosophy*, 23: 179- 203.
- Cosmides, L., and Tooby, J. (1992). Cognitive adaptations for social exchange, in Barkow, J.H., Cosmides, L. and Tooby, J. (eds.) *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press.
- Cheney, D. and Seyfarth, R. (2007). *Baboon Metaphysics: The Evolution of a Social Mind*. Chicago: University of Chicago Press.
- Damuth, J., and Heisler, L. (1988). Alternative formulations of multilevel selection. *Biology and Philosophy*, 3: 407-30.
- Dawkins, R. (1986). *The Blind Watchmaker*. London: Longman.
- Dennett, D.C. (1995). *Darwin's Dangerous Idea*. New York: Simon and Schuster.
- Dugatkin, L.A. (1997). *Cooperation among Animals: An Evolutionary Perspective*. Oxford University Press.
- Dunbar, R. (1992). Neocortex size as a constraint on group size in primates, in *Journal of Human Evolution*, 20: 469-93.
- _____. (1993). Coevolution of neocortical Size, group size and language in humans. *Behavioral and Brain Sciences* 16 (4): 681-735.
- _____. (2000). On the origin of the human mind in Carruthers, P., and Chamberlain, A. (eds.) *Evolution and the Human Mind: Modularity, Language, and Meta-Cognition*. Cambridge University Press.
- Finlayson, C. Giles Pacheco, F., Rodríguez-Vidal, J. (2006). Late survival of Neanderthals at the southernmost extreme of Europe" *Nature* 443: 850-853.
- Gardenfors, P. (2003). *How Homo became Sapiens: On the Evolution of Thinking*. Oxford University Press.
- Gilbert, M. (1990). Walking together: a paradigmatic social phenomenon, *Midwest Studies in Philosophy*, Vol. 15, pp. 1-14.
- _____. (1992). *On Social Facts*. Princeton University Press.
- Godfrey-Smith, P. (2007). Conditions for evolution by natural selection. In the *Journal of Philosophy*, 104:489-516.

- Hamilton, W.D. (1964a). The genetical evolution of social behavior, I., *Journal of Theoretical Biology*, 7, 1-16.
- _____. (1964b). The genetical evolution of social behavior, II., *Journal of Theoretical Biology*, 7, 17-52
- Holekamp, K. E. (2007). Questioning the social intelligence hypothesis, in *Trends in Cognitive Science*. 11: 65-69.
- Humphrey, N.K. (1976). The social function of intellect, in Bateson, P., and Hinde, R.A. (eds.) *Growing Points in Ethology*. Cambridge University Press.
- Jackendoff, R. (2007). *Language, Consciousness, Culture: Essays on Mental Structure*. MIT Press.
- Jackson, F., and Pettit, P. (1992). In defence of explanatory ecumenicalism. *Economics and Philosophy*, 8: 1-21.
- Jolly, A. (1966). Lemur social behaviour and primate intelligence, in *Science*, 153: 501-6.
- Kappeler, P.M., and van Schaik, C.P. (2006). *Cooperation in Primates and Humans: Mechanisms and Evolution*. Germany: Springer.
- Kerr, B., Godfrey-Smith, P., and Feldman, M. (2004). What is altruism?, in *TRENDS in Ecology and Evolution* Vol.19 No.3 March 2004.
- Kruuk, H. (2002). *Hunter and Hunted: Relationships Between Carnivores and People*. Cambridge University Press.
- Kuhn, S. and M. Stiner (2006). What's a mother to do? A hypothesis about the division of labor among Neandertals and modern humans in Eurasia, in *Current Anthropology* 47(6):953-980.
- Larsen, C.S., Matter, R., Gebo, D. (1998). *Human Origins: The Fossil Record (3rd Edition)*. Waveland Press, Inc.
- Leakey, R. (1994). *The Origin of Humankind*. Basic Books.
- Lewin, R., Foley, R. (2004). *Principles of Human Evolution: Second Edition*. Blackwell Publishing.
- Lloyd, E.A. (1992). Unit of selection, in Keller, E.F. and Lloyd, E.A. (eds.) *Keywords in Evolutionary Biology*. Cambridge, MA: Harvard University Press.
- Lewontin, R. (1970). The units of selection, in *Annual Review of Ecology and Systematic*, 1: 1-14.

- Maynard Smith, J., and Szathmary, E. (1995). *The Major Transitions in Evolution*. W.H. Freeman Spektrum Press.
- Melis, A., Hare, B., & Tomasello, M. (2006). Chimpanzees recruit the best collaborators. *Science*, 31, 1297 - 1300.
- Mellars P. (2006). A new radiocarbon revolution and the dispersal of modern humans in Eurasia, in *Nature*, 439:931-935.
- Mitani, J. (2006). Reciprocal exchange in chimpanzees and other primates, in Kappeler, P.M., and van Schaik, C.P. (eds.) *Cooperation in Primates and Humans*. Germany: Springer.
- Moll, H., and Tomasello, M. (2007). Cooperation and human cognition: The Vygotskian Intelligence Hypothesis, *Philosophical Transactions of the Royal Society*, 362: 639-348.
- Pinker, S. (2002). *The Blank Slate: The Modern Denial of Human Nature*. Viking.
- _____. (2003). Language as an adaptation to the cognitive niche, in Christiansen, M.H., and Kirby, S. (eds.) *Language Evolution*. Oxford University Press.
- Richerson, P., and Boyd, R. (2005). *The Origin and Evolution of Cultures*. Oxford University Press.
- _____. (2006). *Not by Genes Alone: How Culture Transformed Human Evolution*. The University of Chicago Press.
- Roebreks, W. (2008). Chronology of the Middle-Upper Paleolithic transition in Eurasia. *Journal of Human Evolution*, Vol. 55, Issue 5: 918-926.
- Rosenberg, A. (1992). Altruism: theoretical contexts, in Keller E.F. and Lloyd E.A. (eds.) *Keywords in Evolutionary Biology*. Harvard University Press.
- Searle (1995). *The Construction of Social Reality*. Free Press.
- Shanahan, T. (1997). Pluralism, antirealism, and the units of selection. *Acta Biotheoretica*, 45: 117-126, 1997.
- Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.
- Sober, E. (2000). *Philosophy of Biology: Second edition*. Westview Press.
- _____. (2008). Realism, conventionalism, and causal decomposition in units of selection: Reflections on Samir Okasha's *Evolution and the Levels of Selection*, in *Philosophy and Phenomenological Research*, forthcoming.

- Sober, E. and Wilson, D.S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Soficaru, A., Petrea, C., Doboş, A., Trinkaus, E. (2007). The human cranium from the Peştera Cioclovina Uscată, Romania: Context, age, taphonomy, morphology, and paleopathology *Current Anthropology*. Volume 48, Issue 4, Page 611–619, Aug.
- Sterelny, K. (2003). *Thought in a Hostile World: The Evolution of Human Cognition*. Blackwell Publishing.
- _____. (2007). Social intelligence, Human Intelligence and Niche Construction. *Philosophical Transactions of the Royal Society, Series B*. Volume 362
- _____. (2008a). Plasticity, Learnability, Differentiation: An Alternative Model of Evolutionary Psychology, *Jean Nicod Lectures, Session I*.
- _____. (2008b). The model in action: Revisiting the problem of cooperation, *Jean Nicod Lectures, Session IV*.
- Sterelny, K., and Griffiths, P. (1999). *Sex and Death: An Introduction to Philosophy of Biology*. The University of Chicago Press.
- Sterelny, K., and P. Kitcher (1988). The return of the gene, in *Journal of Philosophy* 85: 339-361.
- Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., and H. Moll (2005). Understanding and sharing intentions: The origins of cultural cognition, in *Behavioral and Brain Sci.* 28: 675-691.
- Trivers, R.L. (1971). The evolution of reciprocal altruism, in *Quarterly Review of Biology* 46: 35-57.
- Van Schaik, C., and Kappeler, P.M. (2006). Cooperation in primates and humans: Closing the gap, in *Cooperation in Primates and Humans: Mechanisms and evolution*, van Schaik, C., and Kappeler, P.M. (eds.). Germany: Springer.
- Vincent, T. and Brown, J. (2005). *Evolutionary Game Theory, Natural Selection, and Darwinian Dynamics*. Cambridge University Press.
- Warneken, F., Chen, F., & Tomasello, M. (2006). Cooperative activities in young children and chimpanzees, in *Child Development*, 3, 640-663.
- Weaver, T., Roseman C. (2005). Ancient DNA, late Neandertal survival, and modern human–Neandertal genetic admixture, in *Current Anthropology* Volume 46, Number 4, August/October 2005.
- Wilson, D.S. Levels of Selection: An alternative to individualism in biology and the human

sciences, in Sober, E. (2006). *Conceptual Issues in Evolutionary Biology: Third edition*. MIT Press.

Wilson, D.S., and Dugatkin, L.A. (1992) Altruism: Contemporary debates, in Keller E.F. and Lloyd E.A. (eds.) *Keywords in Evolutionary Biology*. Harvard University Press.

