3-2011

On the Production and Ramification of Cooperation: The Cooperation Afforder With Framing Hypothesis

Steven. O. Kimbrough
University of Pennsylvania

Recommended Citation
Kimbrough, S. O. (2011). On the Production and Ramification of Cooperation: The Cooperation Afforder With Framing Hypothesis. *Philosophy of the Social Sciences, 41* (1), 111-136. http://dx.doi.org/10.1177/0048393110387883

On the Production and Ramification of Cooperation: The Cooperation Afforder With Framing Hypothesis

Abstract

This article presents a new proposal for understanding the establishment and maintenance of cooperation: the cooperation afforder with framing hypothesis, producing what can be called cooperation from *afforder-framing*. Three key moves are present. First, a special variety of the Stag Hunt game, the Cooperation Afforder game, will reliably produce mutualistic cooperation through an evolutionary process. Second, cognitive framing is a credible candidate mechanism to meet the special conditions and requirements of the Cooperation Afforder game. Third, once mutualistic cooperation is established in this way, it will plausibly lead to broader forms of cooperation, even to limited forms of altruism.

Keywords

cooperation, altruism, evolution, framing, game theory, strategic interaction, simulation

Disciplines

Other Social and Behavioral Sciences | Theory, Knowledge and Science

On the Production and Ramification of Cooperation: The Cooperation Afforder with Framing Hypothesis

Steven O. Kimbrough*

November 5, 2010

Abstract

This paper presents a new proposal for understanding the establishment and maintenance of cooperation: the cooperation afforder with framing hypothesis, producing what can be called cooperation from *afforder-framing*. Three key moves are present. First, a special variety of the Stag Hunt game, the Cooperation Afforder game, will reliably produce mutualistic cooperation through an evolutionary process. Second, cognitive framing is a credible candidate mechanism to meet the special conditions and requirements of the Cooperation Afforder game. Third, once mutualistic cooperation is established in this way, it will plausibly lead to broader forms of cooperation, even to limited forms of altruism.

Keywords

cooperation, altruism, evolution, framing, game theory, strategic interaction, simulation

1 Introduction

The problem of cooperation is not one thing; it is many. Broadly speaking the problem is to explain and predict cooperative behavior, when it occurs, and failure of cooperation when that occurs. Further, we wish to support intervention in order to promote or discourage cooperation. How should we design and manage institutions in order to achieve optimal levels of cooperation? Add the philosophers' perennial issue of characterizing the subject: Just what is cooperation? How does altruism, surely a form of cooperation, differ from other kinds of pro-social behavior? and so on. Each of these problems spawns multiple sub-problems and varieties.

What makes the *subject* of cooperation so forcefully problematic is the fact that there seems to be so much of it, which fact is also without fully satisfactory explanation. Tipping in some circumstances, helping strangers, and so on are some of many well-known, accepted behaviors that appear to be highly cooperative, even altruistic. Testimony may be had from many sources as to the prevalence and puzzlingness of such behavior. To pick just one, here is Michael Tomasello on altruism:

> In the contemporary study of human behavioral evolution, the central problem is altruism, specifically, how it came to be. [SOK: And we might add, how come it has stayed

*University of Pennsylvania, Operations and Information Managment, 3730 Walnut Street, Philadelphia, PA 19104 USA. Email: kimbrough@wharton.upenn.edu

> around?] There is no widely accepted solution to that question, but there is no shortage of proposals either. (Tomasello, 2009, page 51)

(The edited volume by Hammerstein (2003) is rich in biological examples.)

In what follows, I wish to add to the list of proposals for the origin and maintenance of cooperation, including altruism and other forms of pro-social, or helping, behavior. It is not my aim to offer a settled explanation of any particular cooperative behavior. Instead, I wish to add a new, or at least revised, account to the extant list. I want to add a conceptual tool to the resources of empirical science. Uses and tests of the tool must perforce follow later, although I shall in passing present some evidence.

The argument I want to make, or the story I want to tell, comes in several steps, beginning in the next section. In preparation I'll address two issues quickly.

1.1 Definitions

First, what counts as cooperation, altruism, pro-social behavior, helping, and so on? There are problems with trying to answer this question fully. The fact is that the relevant terms are commonly used with varying senses in the literature. To note just two examples, Bronstein (2003, pages 186–7) writes that

> Diverse terms are currently in use to define interspecific cooperative interactions. Some of these terms have well-accepted alternative meanings, however. Given such inconsistencies, the terminology we choose to employ is more than a strictly semantic matter: it can determine whether we are all trying to explain the same phenomena. . . .
>
> I use the term *mutualism* to refer to all mutually beneficial, interspecific interactions, regardless of their specificity, intimacy, or evolutionary history. . . .
>
> Whereas mutualism [as Bronstein will use the term] denotes two-species beneficial interactions, *cooperation* has usually been used somewhat more vaguely to denote benefits in a within-species context. Some researchers have used the terms mutualism and cooperation interchangeably . . . [and she goes on to note a wide variety of conflicting uses and senses]

In a footnote to the first occurrence of the term cooperation, Richerson et al. (2003, page 358) write this:

> "Cooperation" has a broad and a narrow definition. The broad definition includes all forms of mutually beneficial joint action by two or more individuals. The narrow definition is restricted to situations in which joint action poses a dilemma for at least one individual such that, at least in the short run, that individual would be better off not cooperating. . . . The "cooperate" vs. "defect" strategies in the Prisoner's Dilemma and Commons games anchor our concept of cooperation, making it more or less equivalent to the term "altruism" in evolutionary biology. Thus, we distinguish "coordination" (joint interactions that are "self-policing" because payoffs are highest if everyone does the same thing) and division of labor (joint action in which payoffs are highest if individuals do different things) from cooperation.

Even a more ordinary word such as exploitation can be problematic because it has (at least) the senses of "putting to good use" and "taking unfair or immoral advantage of". For example, Bronstein (see above) writes of "reciprocal exploitation (i.e., mutualism)" (Bronstein, 2003, page 191) which she contrasts with "unilateral exploitation," which itself may or may not be particularly harmful to the exploited party.

A further impediment to any attempt to finalize the relevant definitions is that the list of relevant terms is growing, with no apparent natural end point in view. At a minimum we need also to consider *mutualism, reciprocity, pseudoreciprocity, indirect reciprocity, by-product mutualism, kin selection,* and *reciprocal altruism.*

The notion has been raised that perhaps our language is inadequate to the phenomena, and that we would be well-advised to focus primarily on describing kinds of phenomena that seem to be relevant and worry later about what to call them (e.g., Leimar and Connor (2003)). I have considerable sympathy for this advice and intend to follow it here. Nevertheless, some rough characterization of (some) terms will be useful. Here it is.

Altruism is commonly described as "selfless concern for the welfare of others" and I think that's fine as far as it goes. I'd prefer to keep examples, or exemplars, to mind. At the low end, helping strangers by giving directions and giving to charity I'll take as (normally) counting. The agent freely provides—transfers—something of value without expectation of any immediate or very specific reward. I'll call this *altruism in the small.* At the high end, of course, are taking big risks to save someone else or even sacrificing one's life for another. I'll call this *altruism in the large.* Of course, altruism may not itself be one thing (I'm inclined to think it isn't), and its nature is itself contested. Some theorize that apparent altruism is often (always?) discretely selfish. I want to move on and leave these questions for some other time. Roughly, I'll take as altruistic any behavior that benefits someone else without any immediate prospect in sight of commensurate benefit for the giver.

I'll use *cooperation* approximately in the broad sense of Richerson et al. (2003, page 358, see above), which I take to be essentially the same as mutualism in Bronstein's sense Bronstein (2003, pages 186, see above), but with a relaxation to allow more than two individuals and to allow intraspecific interactions. And I'll use *pro-social* and *helping* interchangeably to include both cooperation and altruism, as well as anything else that is arguably similar but not similar enough to count as one or the other, such as strong reciprocity (Fehr and Henrich, 2003).

1.2 Extant Proposals

The second housekeeping issue before I begin the argument directly is the list of extant proposals for explaining pro-social behavior. A number of kinds of explanations for the establishment and maintenance of cooperation or altruism have been put forward. Sometimes called *evolutionary pathways to cooperation* (Henrich and Henrich, 2007, page 40), among others, they include:

- kin selection

  Cooperation, even altruism, among relatives is favored by selection because they share the same genes ((Wilson, 2000) or any standard text on evolution).

- reciprocal altruism

  The cooperation and altruism that is selected for is conditioned on reciprocity by the individual receiving the benefits of the cooperative or altruistic behavior ((Wilson, 2000) or any standard text on evolution).

- group selection

While is it individually deleterious be be altruistic, the trait is preserved in populations because populations in which the trait is prevalent out-compete populations in which it is not (Sober and Wilson, 1998).

- culture

  Norms of cooperation are promoted and enforced by cultural institutions (Richerson and Boyd, 2005).

- bounded rationality

  Cooperation, even altruism, is due to the side-effects of imperfectly-rational heuristics which on balance are individually favorable and favored by selection (Gigerenzer and Selten, 2001).

- sexual selection

  Males (typically) compete for mating access to females, and may be selected because of their cooperative behaviors ((Wilson, 2000) or any standard text on evolution).

- social selection

  A generalization of sexual selection in which competition is for contribution to the entire reproductive process, not merely mating (Roughgarden, 2009).

- selective investment theory

  Cooperative behavior is an investment, which if applied effectively will be advantageous on straightforward darwinian grounds (Brown and Brown, 2006).

- pseudoreciprocity

  A helps B because B's flourishing is itself helpful to A (Dugatkin, 1997, circa page 49).

- indirect reciprocity

  A helps B because it will lead to someone else, a C, helping A (Nowak and Sigmund, 2005).

- by-product mutualism

  Mutualism maintained by selfish behavior that incidentally benefits others (Mann et al., 2000, circa page 202).

These pathways, causes of cooperation, are not mutually exclusive. Determination of the actual mix present is an empirical question and likely very much depends upon the particulars of the case at hand.

A final comment for now. We can distinguish between *proximal* and *distal* explanations of pro-social behavior. Some of the pathways to helping behavior are best understood as one more than the other. For example, reputation (a form of indirect reciprocity), social norms, ethnicity (Henrich and Henrich, 2007, page 32) and "identity economics" (Akerlof and Kranton, 2000, 2005, 2010) are proximal explanations of cooperative behavior. What explains why we rely on reputation effects to modify our behavior or why we value "identities" and rely on them to support cooperative behavior? These are questions that call for distal, presumably evolutionary and/or cultural explanations.

Now to the story to hand. The reader may judge the extent to which it differs from existing explanations.

2 The Cooperation Afforder Game

It will help to set the scene. Imagine a tribe of hunter-gathers in a primitive state of development, compared to modern humans. In particular there is little or no supporting apparatus, technology, or even cultural practices useful to the individuals *jointly* as they go out and hunt for food. There are in short little in the way of *cooperation afforders* for the hunt, things that make possible and support cooperation in the hunt.[1] In consequence, there is no joint gain to be made from pairwise cooperation. Everyone hunts hare, as it were, and there are no available alternatives. Into this situation comes a new afforder with the property that it affords cooperation between two individuals if they *both* have and use this afforder. Think of it perhaps as a net that must be held by two people for it to work in catching a stag, but the net must be assembled from two components, one brought to the field by each of the two hunters. Further, let us assume that the cost of bringing a single component to the field is small but real; the afforder is not free. More on all of this later. For concreteness, let the strategic situation be as in Figure 1.

	C	D
C	(2,2)	(0.99, 1)
D	(1, 0.99)	(1, 1)

Figure 1: Cooperation Afforder game, standard payoffs

Think of the payoffs as being given in days of food harvested. If an agent has the afforder (and uses it), we call this the C or cooperative strategy; if not then this is the D or defecting strategy. Under our standard payoffs (Figure 1), two cooperating agents will harvest each two days of food. Two defecting agents will each harvest one day of food. When a cooperating agent plays with a defecting agent, the cooperator harvests a day of food minus a small cost of bringing the afforder to the field, while the defecting agent gets the day of food without the extra cost.

The upshot is that we have a kind of Stag Hunt game. Brian Skyrms has, very commendably, brought the general Stag Hunt game forcefully to our attention. He has explored general Stag Hunt games and documented the problems of achieving sustained cooperation in them (Skyrms, 2001, 2004).

The so-called *replicator dynamics* is a standard model for investigating the behavior of games such as Stag Hunt under evolution by natural selection (Taylor and Jonker, 1978; Hofbauer and Sigmund, 1998). Skyrms, with others, uses this model. In its discrete form (generations are discrete) the frequency of an element (here, a strategy, either C or D) in the next generation, $p(s_i)'$ (probability of strategy s_i in the next generation) equals its frequency in the current generation, $p(s_i)$, plus the difference between its average performance, or return, r, and that of the entire population $((\overline{r}(s_i) - \overline{\overline{r}(s_i)})$, weighted by is current frequency. That is:

$$p(s_i)' = p(s_i) + p(s_i) \cdot (\overline{r}(s_i) - \overline{\overline{r}(s_i)}) \tag{1}$$

Thus, if the average return of strategy s_i is higher than the average of the averages of all the individual strategies (all the s_is in the population), then the frequency of s_i increases in the next generation; and if not, not.

[1] I choose the term afforder with allusion to Donald Norman's use of the term affordance (Norman, 2002). A design, in Norman's usage, provides certain affordances, things we can have or do because of the design. By the term afforder, I simply mean that which affords something, so a design is an afforder and what it affords is an affordance.

1. Preparatory steps:
   (a) Determine *populationSize*, the population size to be used.
   (b) Determine the game, G, with its payoffs, and the available strategies, $\mathcal{S}$.
   (c) Determine *fInitial*, the initial frequency distribution of strategies from $\mathcal{S}$.
   (d) Determine *sampleSize*, the number of pairwise games to be play each generation.
   (e) Create the initial population of *populationSize* strategies drawn from *fInitial*.
   (f) Create and set to 0 reward accumulators, $r(s_i)$, for each strategy $s_i \in \mathcal{S}$.
2. Process generations until done:
   (a) Do *sampleSize* times:
      i. Draw two members of the population at random, noting their types (which strategy they are).
      ii. Play the two members against each other in the game G; allocate the outcome rewards to their respective strategies, $r(s_i)$.
   (b) Determine *fNext*, the frequency distribution of strategies for the next generation, as

$$fNext(s_i) = \frac{r(s_i)}{\sum_i r(s_i)} \tag{2}$$

   (c) Create the next generation population of *populationSize* strategies drawn from *fNext*.
   (d) Set to 0 the reward accumulators, $r(s_i)$, for each strategy $s_i \in \mathcal{S}$.

Figure 2: Pseudocode for the discrete generation replicator dynamics with a finite population

Figure 2 presents pseudocode for calculating the distribution of strategies, as it changes over discrete generations, with a finite population. It is this process that I shall use, or that has been used, to generate the results to be described. It is slightly different from the equational form, expression (1), but the results are not materially different.

	C	D
C	(3, 3)	(0, 2)
D	(2, 0)	(1, 1)

Figure 3: Stag Hunt, standard payoffs

With the payoff regimes Skyrms has typically handled, e.g., Figure 3, it is indeed quite possible for evolution in the form of the replicator dynamics to drive a population to all D, depending on initial conditions. For example, for the game of Figure 3 and a start of Cs and Ds in equal numbers, the normal result is conquest by Ds. Herein lies a puzzle for cooperation, as Skyrms has emphasized. How is it that cooperation can get started in the first place?

	C	D
C	(R, R)	$(N - c, N)$
D	$(N, N - c)$	(N, N)

Figure 4: Cooperation Afforder game, abstracted. R: reward for cooperation. N: payoff for no afforder. c: (small) cost of the afforder. $R > N > c > 0$.

The Cooperation Afforder game, abstracted in Figure 4, is (technically) a Stag Hunt (or Assurance) game because it is symmetric and $R > N > N - c$. According to classical game theory, only the ordering of R, N and $N - c$ payoffs matter, not their magnitudes. But not for us. Our concern is with evolution and other causal processes, not with what happens under some a-causal theory of ideal rationality (as in classical game theory). So payoff magnitudes do matter. Under the replicator dynamics for the game in Figure 1, so long as the frequency of Cs is at or above about 2–3%, the population will reliably evolve to all Cs. With balking, multiple rounds of play and reactive strategies, etc., starting frequencies of less than 2% for Cs can reliably result in conquest by cooperators.These are empirical claims, arrived at using ReplicatorDynamicsWithBalking.nlogo, but they are robust. Figure 5 shows ReplicatorDynamicsWithBalking.nlogo set up to run the Cooperation Afforder game with initially 3% cooperators, shown as black small squares. The defectors are white small squares. After a number of generations, on the order of 75, the defectors (represented by white squares) will be eliminated. The entire population will consist of cooperators.[2]

These tests were with large populations, of about 95,000 individuals. With small populations it is quite plausible that strategy C could appear and "drift" to a large enough frequency that its selective advantage would come into play and drive the population to all C, and then drive the population to expand at the expense of populations without the afforder.

There are three key properties of the Cooperation Afforder game that afford the evolution of cooperation.

1. There are substantial benefits to be had from cooperation.

[2]To run the ReplicatorDynamicsWithBalking.nlogo program, point your Java-enabled browser to `http://opim.wharton.upenn.edu/~sok/agebook/applications/nlogo/41/` and open up ReplicatorDynamicsWithBalking.html.

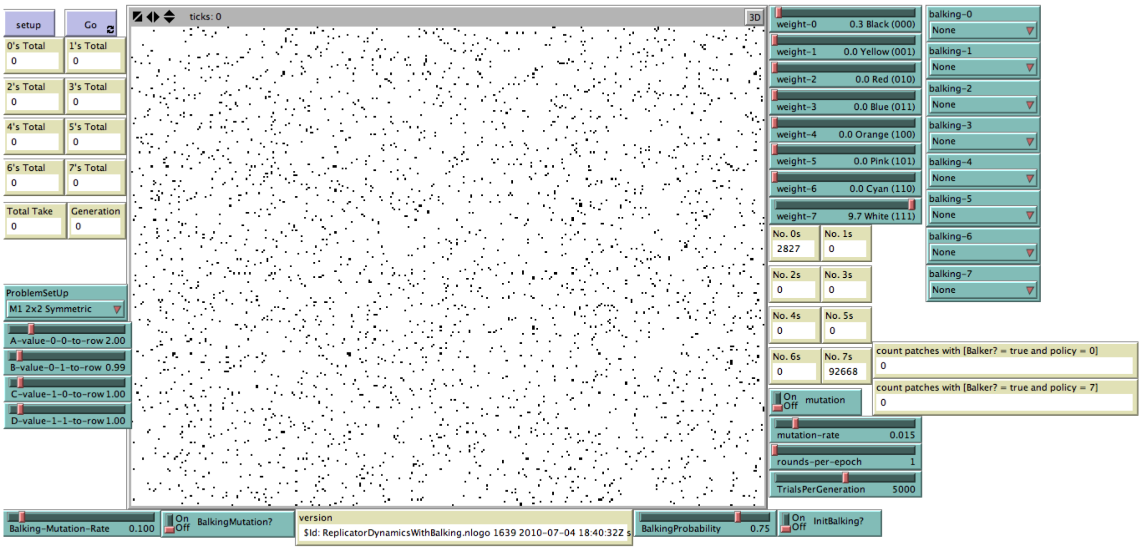

Figure 5: ReplicatorDynamicsWithBalking.nlogo set up for a run of the Cooperation Afforder game, starting with 3% cooperators and balking and mutation off.

For example, in the Cooperation Afforder game with standard payoffs, Figure 1, two days of food from mutual cooperation is substantially better than one day of food when not cooperating.

2. Unrequited cooperation is not very costly.

   It is not free, but neither is it immediately life-threatening to offer cooperation and not have it reciprocated. In terms of our example, the fitness loss is low enough, compared to the benefits of cooperation, that even relatively rare encounters with fellow cooperators will pay for the cost of the carrying the afforder.

3. Unrequited cooperation is not exploitable.

   When one agent cooperates and the other defects, the defecting agent gains no positive benefit of any size. There is a fitness loss by the cooperating agent, but the fitness gain (relatively) is divided among all of the non-cooperators. The non-cooperating counterpart in question has no idiosyncratic gain. The replicator dynamics calculations (simulations) with the Cooperation Afforder game are simply a concrete, specific way to demonstrate this.

Under these general conditions, which are not specific to the Cooperation Afforder game, we should expect cooperation (often) to be favored by evolution and/or by learning. Put otherwise, when these conditions obtain, we are able to answer the question of how cooperation can get started and then be maintained. Broadly speaking, these conditions suffice for cooperation.

3 Name that Technology

The upshot of the previous section is that there is no fundamental impediment to understanding how a cooperation afforder could conquer a population and a species via something like the Cooperation Afforder game. This naturally raises the questions of what afforder or afforders could, or even did, play this rôle. Does evolution have to find anew every form of cooperation or is there a more general cooperation afforder? And so on. To address these questions and to motivate my general answer, I want to take a brief digression.

3.1 Hi-Lo

Consider for concreteness a simple, no-conflict, even trivial game represented in strategic form:

	H	L
H	2, 2	0, 0
L	0, 0	1, 1

(a) Standard form

	H	L
H	A, A	0, 0
L	0, 0	B, B

(b) Generic form

Figure 6: Hi-Lo, after Bacharach (2006). Assumed: $A > B$.

The game is called Hi-Lo, for obvious reasons, and is an example of an *inessential* game, a type normally passed over with dismissal because it is so trivial (Rapoport et al., 1976, page 22). Notably and commendably, Bacharach (2006) disagrees and discusses it at length. I agree with much of his analysis, but would depart considerably from elements of it. These matters, however, do not matter for the present.

Any simple equilibrium account of Hi-Lo (or indeed Stag Hunt), such as classical game theory, will fail for definiteness in the presence of multiple equilibria, and the classical account fails for that reason here.[3]

> Nash equilibrium fails to confer determinacy anyway, so there is no need to dispute the principle in order to show that standard theory [that is, classical game theory] fails for Hi-Lo. (Bacharach, 2006, page 61)

But things get worse. The theoretical problem of multiple equilibria is exacerbated in the case of Hi-Lo because the practical problem is utterly unproblematic. People's behavior and all our intuitions agree that playing H is the rational thing to do. The theoretical challenge is to explain why—descriptively, in practice—it *is* unproblematic to play high, H. Bacharach puts it nicely.

> You are to play Hi-Lo, and it is common knowledge that you and your coplayer are intelligent people. It seems quite obvious that you should choose A [H in our Figure 6]. However, the question why it seems obvious, and the related question of why people almost always do choose A [H], have turned out to be anything but easy to answer. (Bacharach, 2006, page 35)

[3] Bacharach is aware of (Harsanyi and Selten, 1988) and related work on equilibrium refinement, which neither he nor I think succeeds in eliminating the paradox. See (Bacharach, 2006) for a discussion, which would take us too far afield.

This is the paradox, the Hi-Lo paradox of classical game theory. Game theory should, one would think, easily be able to offer a satisfying account of why players will coordinate on Hi in Hi-Lo, but it does not. Game theory should, one would think, also easily be able to offer a satisfying prescriptive account of why players should play Hi in Hi-Lo, but again it does not. We now turn to the essentials of Bacharach's alternative account.

In a nutshell, Bacharach observes that classical game theory presumes individual thinking, in which each of the playing agents asks of itself What is it best for me to do in this situation? Unfortunately for classical game theory, there is no determinate answer to this question via the theory. That is the point of the Hi-Lo paradox. As a way out of the paradox Bacharach proposes a different framing: team thinking. A team thinking agent, confronted by a strategic decision problem asks instead What is it best for my team to do and what is my rôle in it? Having answered these questions, the agent follows through by enacting its rôle in what on balance it judges is best for the team to be doing.[4]

Credible examples of team reasoning come easily to mind in the case of team sports. Two players perceive an opportunity for the team and perceive that the other shares the perception. They then spontaneously coordinate their actions. Simultaneously, one throws the ball to a spot and the other runs towards the spot. This happens, has to happen, too quickly for the players to signal each other and agree. They act autonomously on what they see and very often will achieve effective coordination. Ordinary life is rich in activities that will often be best understood as evidencing team reasoning by their members. These include group hunting, rowing, dancing, singing, military activity, and various forms of work (McNeill, 1995). On the last, McNeill thinks we see this not only in gangs of laborers, but in the efforts of modern corporations to instill company spirit through propaganda and (in Japan) calisthenics.

Bacharach's view is that coordination in Hi-Lo works similarly, but of course in quotidian interactions where it hardly requires the skills of professional sports players. The story is that the agents see the game (perceive the situation), are cued for team thinking, reason—trivially—that what is good for the team is for each player to play H, and end by . . . playing H.

If the players are in fact engaging in team reasoning (or something like it), it is a virtue of this account that we can easily explain and predict what should be obvious, coordination on H, but is unexplained entirely by individual thinking and classical game theory. End of digression.

3.2 Pro-Social Framing

The thesis, the explanatory path to cooperation, I wish to put on the table employs a version of, or something like, Bacharach's account of team reasoning. There are several parts to the thesis. Here are the first few.

1. **Causation.** Pro-social frames, or framing, once present in a population could, by the argument above with regard to the Cooperation Afforder game, evolve to domination or fixation in a population.[5]

[4] There are many complications possible which for present purposes need not detain us. For example, what if the agent thinks that some of the other team members cannot discern or enact their rôles? See Bacharach (2006) for discussion of some of the related nuances.

[5] Frames for Bacharach and in the other literature I discuss are alternate ways of organizing experience. One might speak of them as collections of presuppositions, which are used to direct action, arrive at new beliefs, and so on. So we might think of frames as collections of beliefs and a framing effect as something believed or done in consequence of a particular collection of beliefs. But this is not necessary. Frames may organize experience without recourse to

2. **Behavior.** There exist frames of thinking, among them, prominently, a selfish, self-centered frame and a pro-social frame. Individuals using their pro-social frames will tend to behave cooperatively, even altruistically. In a team thinking frame, individuals identify with the team and predicate their actions on their estimates of what they should do for the good of the team.

In short, *pro-social framing is an afforder whose proliferation may be explained straightforwardly on evolutionary grounds.* I submit this as a candidate for explaining, or contributing to the explanation of, many cooperative phenomena. Call it the *cooperation afforder hypothesis with framing* for explaining cooperative phenomena. Cooperation that comes about and is maintained in virtue of the process hypothesized may be called *afforder-framing cooperation.*

The existence point—that there are both selfish and pro-social frames of thinking—is not a new idea. It goes by different names from different authors, who may or may not cite each other.[6] The idea has been in the air. Here is a recent representative passage.

> In shared cooperative activities, we have a joint goal that creates an interdependence among us—indeed, creates an "us." If we are carrying a table to the bedroom, I cannot simply drop it and run off without hurting us and our goal. In shared cooperative activities, my individual rationality—I want to transport the table to the bedroom so I should do X—is transformed into a social rationality of interdependence: *we* want to transport the table to the bedroom, so I should do X and you should do Y. (Tomasello, 2009, page 41)

The case I am making is not very specific regarding the details of the pro-social afforder and associated framing. Intendedly so; I'm making a basic point. There is much remaining to do by way of modeling and empirical investigation. I note in passing, however, several features commonly identified with pro-social framing (we-mode thinking, team thinking, etc.):

1. **Shared goals.** The agent directs its activities towards fulfilling the goals of the team, group, society.

2. **Imitation.** The agent emulates and conforms to, insofar as feasible, characteristic behaviors of members of the group (society).

3. **Enforcement.** The agent undertakes to enforce norms of conformance. Defectors from the group are to be punished.

I wish to suggest, either alternatively or in addition (let the facts reign), another afforder as a key player, one that complements and abets cooperative framing and team reasoning. The afforder is communication and again Brian Skyrms has broken important ground in exploring the subject (Skyrms, 2010). The ability to communicate, even in a rudimentary way, will plausibly often have scope to instantiate the three key properties of the Cooperation Afforder game:

anything directly propositional, for example by using a procedure to interpret sensory signals. The account I am giving here, and the accounts given by others employing framing, are mostly neutral with regard to how frames are built and to how framing effects result from them. This said, it is unlikely that frames are fundamentally propositional or doxastic. In the end, the issue is an empirical one whose resolution is of peripheral consequences for the case I am making.

[6] Tomasello (2009) cites philosophers: Nagel (1979), **?**, Gilbert (1992), Searle (1997) and Tuomela (2007) who earlier introduced the terms I-mode thinking and we-mode thinking, corresponding roughly to individual and team thinking. Tomasello does not, however, cite economists and game theorists, e.g., Bacharach (2006), Sugden as in Mehta et al. (1994); Sugden (2008), and Colman as in Colman et al. (2008a). Nor do they cite the philosophers or the biologists.

1. There are substantial benefits to be had from cooperation.

   The potential benefits of cooperation have to be provided, exogenously as it were, by a combination of the environment and the agents' needs and other capabilities. If these are such that with cooperation abetted by communication would in fact pay well, then cooperation may well ensue.

2. Unrequited cooperation is not too costly.

   We are not safe in assuming that the act of communication is free, especially considering the requisite apparatus for being able to communicate. Given the apparatus, however, the energetic costs of signaling are surely low, or often can be.

3. Unrequited cooperation is not exploitable.

   If the target of the communicative act does not reciprocate with cooperation, there is, or can be, little strategic cost to the sender. This may happen because the target fails to understand or in some other way to have uptake of the message. The target might have perfectly adequate uptake, but choose to ignore the message. What is being excluded, for the Cooperation Afforder game to apply, is that the target of the communication is able to use the message to its advantage and to the disadvantage of the sender, say by signaling a third party and effectively denying the original sender access to the resource. Of course it can happen that unrequited cooperation is exploitable and is exploited. The suggestion here is that this can be a rare enough eventuality that the conditions for the Cooperation Afforder game do in fact generally apply.

Note in this regard Grice on communication and the important rôle of the cooperation principle (Grice, 1967). Communication in Grice's analysis presumes cooperation and to violate the rule of cooperation is to abuse the institution. Could it be that a propensity for pro-social framing afforded the development of communication and even reinforced the propensity?

4 Evidence

I will be brief, given the intended scope of this paper, which is to present an evolutionary pathway to cooperation, as a candidate for explaining phenomena. It is beyond the scope of this paper to establish this pathway—the cooperation afforder hypothesis—as *the* explanation of anything.

The cooperation afforder hypothesis can explain, or help explain, the presence of something on the order of team thinking, we-mode thinking, shared cooperative activity thinking. Whether it is uniquely qualified to do so is beyond the scope of this paper. Little or nothing really rides on whether that is the case. The more urgent question, I suggest, is whether there is in fact evidence for something on the order of team thinking, we-mode thinking, shared cooperative activity thinking. Do we indeed have the two distinct modes of thinking, one pro-individual, one pro-social?

We do. I'll cite some of the indirect and direct evidence. On the indirect side, the presence of framing and framing effects in human psychology is very well established. Richard Gregory (1997, 2009) has made us well aware of visual illusions. What we see depends on how see, how we frame things. The effect occurs outside of vision and even perception. We see it in our conceptualizing of the world. This is, I take it, a main point Wittgenstein was after in his duck-rabbit discussions. The observation is an ancient one. Is Antigone's opposition to Creon a defiance of law and duty to her fellow citizens or is it a noble display of fidelity to and honoring of her dead brother? The frames, the presuppositions, of Antigone and Creon conflict with tragic upshot. The effect is often present in moral dilemmas, but has been established in many contexts and is well accepted in the psychology literature (Tversky and Kahneman, 1981). So, the indirect point is that we see framing effects all over and should not be surprised to find one more.

Direct experimental evidence for team reasoning in Hi-Lo and related games is also available, although only a few controlled experiments have been reported. Team reasoning in Hi-Lo is especially difficult to explore experimentally because there is no variation to work with: everyone always picks H. Similar games can be designed, however, that do not elicit successful team coordination universally. Subjects may then be cued with team-promoting or individual-promoting messages and we may expect, if the team reasoning hypothesis is correct, that team-cued subjects would achieve coordination more often than individual-cued subjects. In experiments of this sort Colman et al. (2008a,b) convincingly do find support for the team reasoning hypothesis.

More or less direct evidence of another sort is available from the literature on friendship, recently reviewed by Silk (2003). To summarize, her money quote is: "Friendship is not mutualism" (Silk, 2003, page 47). Silk's concept of mutualism is, unlike Bronstein's, tied to near-term reciprocity, which requires close bookkeeping that is absent in friendships. She puts her point as follows.

> Taken together these experiments provide empirical support for the distinction between exchange and communal relationships. More importantly, they support the hypothesis that communal relationships are not based on strict Tit-for-Tat reciprocity. People use Tit-for-Tat reciprocity as a diagnostic criteria for the existence of close friendships; when benefits are balanced directly, relationships are assumed to casual and ephemeral [i.e., exchange rather than communal]. People seem to make concerted efforts to obscure the accounting of costs and benefits among their friends—in joint tasks, they hide their own contributions and avoiding [sic] monitoring their friends' contributions. (Silk, 2003, page 46)

Continuing, she writes,

> ...the exchange-communal distinction implies that the process that preserves the balance in these two different kind [sic] of relationships differs. In exchange relationships, help is given with the explicit expectation that it will be reciprocated. In communal relationships, help is given because it is needed or desired; wehn both partners have the same communal orientation, benefits will flow back and forth, but they will not be strictly contingent on expectations of future benefits. (Silk, 2003, page 46)

The suggestion I want to make is that communal thinking is much the same as team thinking, perhaps a generalization of it. In any case, it is a form of framing, which can be and is escaped by experiences, just not by comparatively small deviations from strict reciprocity. Team thinking and communal thinking make us more patient.

We should regard the existence of framing and effects thereof as reasonably well established. Further, we should see the existence of something on the order of team reasoning, we-mode thinking, communal thinking, et cetera as a credible hypothesis, one to be investigated further.

Before beginning to generalize the argument, in the following section, it may be helpful to rehearse in compressed form the argument, or perhaps narrative is a better term, so far.

1. The Cooperation Afforder game is a *model* that may be useful in explaining the origin and maintenance of cooperation in some cases. Calculations with the model demonstrate how a

low-cost, reasonably high-benefit afforder of mutual cooperation could evolve to fixation in a population, once it appeared even in very small numbers.

2. Given the model and its interesting property of explaining the fixation of cooperation, we would like to know how accurately it tracks relevant aspects of the real world. That is a very large question, one that I am not attempting to dispose of in this paper. Weaker, but still very important questions are
   (a) Are there credible candidates for the role of cooperation afforder?
   (b) If there are credible candidates, what does this tell us about further consequences of the narrative?
3. I take up question (2b) in §5, question (2a) immediately below.
4. Regarding question (2a), candidates that are very specific to particular circumstances are less interesting and less credible as sources for the high level of cooperativeness we find about us. The Cooperation Afforder game may explain the fixation of several afforders, but it would be a stretch to ask it to explain hundreds of them (e.g., one for hunting animals, another for hunting fish, another for farming, and so on).
5. Uncooperative behavior, the propensity to act in terms of one's perceived selfish interests, without regard to the interests of one's group, is a behavior that has not been seen to be problematic from an evolutionary perspective.
6. The Cooperation Afforder model teaches us that a propensity to act in terms of the interests of one's group (of 2 so far) is, under certain circumstances (not too costly, etc.), also not problematic from an evolutionary perspective.
7. So there are (at least) two possible propensities, one for selfish behavior and one for cooperative behavior.
   (a) Do we need both to explain the phenomena?
   (b) If so, how do they work and what are their properties?
8. Regarding (7a), the question of whether we need both propensities, that we do has prima facie support from the evident fact, very broadly accepted, that there is a lot of selfish behavior and a lot of cooperative behavior. We can of course try to eliminate the need for one or the other, and this is what some of the paths to cooperation (above, §1.2) try to do. The exercise here is, in part, one of offering a kind of explanation that would serve to buttress the need for both selfish and cooperative propensities. Recall as well Bacharach's argument, discussed above in §3.1, that the behavior universally found for the Hi-Lo game *cannot* be explained with the assumed selfish propensity of contemporary game theory.
9. Regarding (7b), the question of how these propensities work, the suggestion to hand is that they are, or operate like, the cognitive frames we are familiar with and that have received considerable evidential support quite independently of the issues immediately before us. If not quite evidence, this is an attractive property of the suggestion.

10. Cognitive frames are considered to be cued by experiences and once in place the agent operates with them effectively as presuppositions. Additional experience may cause a frame in use to be replaced by another. Sometimes, as with certain optical illusions, frames are easily displaced. In other cases, such as with friendships, frames not so easily displaced; they may be said to be sticky.

Given the evolutionary argument for the fixation and maintenance of a mutualistic propensity for cooperative behavior, and the hypothesis that the cognitive basis for this relies on reasonably sticky framing, what else can we explain? I take up this question in the next section.

5 Ramifications: On Beyond Cooperation?

The Cooperation Afforder game models the capacity for two-player mutualistic cooperation spreading and becoming established in a population. How do we get from this to an explanation of altruistic behavior, of contributing to public goods, and in general of defaulting to cooperative, even altruistic behavior? It is these behaviors that are indeed most in need of explaining.

If what evolves or is learned that supports cooperation is specific to context, goals, and so on, it is indeed hard to see how this would generalize to other circumstances. If what we acquire is a particular ability to cooperate in hunting stag, why would it apply elsewhere? There is little reason to expect altruistic donation to the Red Cross, in consequence of cooperation in Stag Hunt.

The first thing to notice is that there is nothing in the Cooperation Afforder game account that requires cooperation be limited to *two* agents. If teams of agents, of communicating, communal-reasoning agents, find opportunities for cooperative behavior in which the costs of offering cooperation a low and cannot be exploited, then we can expect groups of cooperators to thrive.

Further, I want to add another part of the cooperation afforder with framing hypothesis. I call it the principle of imperfect specificity.

1. **Causation.** Pro-social frames, or framing, once present in a population could, by the argument above with regard to the Cooperation Afforder game, evolve to domination or fixation in a population.

2. **Behavior.** There exist frames of thinking, among them, prominently, a selfish, self-centered frame and a pro-social frame. Individuals using their pro-social frames will tend to behave cooperatively, even altruistically.

3. **Imperfect specificity.** Pro-social frames will compete within (and among) agents with pro-individual frames. They may, in many cases, prove to be of superior value to pro-individual frames.

Consider the explanation for play in Hi-Lo. The player sees and understands the game, goes into team reasoning mode, reasons that its job is to play H, and plays H. The pattern is:

1. Perceive a situation calling for action.

2. Be cued into pro-social thinking in some form.

3. Arrive at an action from the presumed framework and assumptions of pro-social thinking.

The suggestion here is that given a frame, ordinary thinking is all that is required for action, but different frames will lead to different actions because different assumptions are being made. Once a frame is available that supports collaborative activity, cooperation, et cetera, the agent may acquire cues for that frame that are different from those originally calling it up. Further, because cues are always imperfect, there will be a tendency to apply the frame in new situations. In short, a pro-social frame, if available, will compete to be cued with a pro-individual frame. And if the rewards are superior it will tend to be used.

And what might make pro-social frames out-compete pro-individual frames in contexts beyond simple mutualism? Sexual selection is an obvious source (hypothetically).[7] It is not much of a stretch to imagine that, very relevantly in our evolutionary history, a male conspicuously exhibiting pro-social behavior would be especially attractive to females because of the prospect of contributing to raising the children. (Roughgarden (2009) makes a similar point about sexual selection in birds. I don't believe any generalized pro-social framing is envisioned, however; it is not necessary for contributing to brood support.)

Punishment is a second possible source. If a pro-social frame includes a propensity to imitate others in the group, to conform to the group's behavior, and to punish those who do not, we should not be surprised if this frame pervades a given group. If in addition the cooperation so induced is productive, the group may be expected to grow and to out-compete groups without the frame. Support for this kind of story may be found in (Boyd and Richerson, 1992) originally; see (Henrich and Henrich, 2007) for an updated treatment.

I'll mention one other source; there are others: transaction costs. The reason there is imperfect specificity (at least in part) is that it is too expensive (or even impossible) to maintain an apparatus that is highly accurate in discriminating when to and when not to engage in communal thinking and communication, that is to engage one's cooperation afforder. In consequence we may often engage in seemingly altruistic behavior—tipping in places not to be visited again, helping strangers, and so on—as a default, which is not overridden by careful scrutiny and calculation. If the opportunity cost is low (Just how much richer would you be if you always refused to give directions when asked by a stranger?) and you live in a society in which others mostly have made similar cost-benefit tradeoffs, then neither learning nor natural selection may disfavor it. The alternative is simply too costly for the available benefit.

This kind of argument may go far to account for altruism in the small. As for altruism in the large, we should expect less of it. Even so, kin selection plus imperfect specificity may account for some cases, as surely must catastrophic errors that occur so rarely that, again, imperfect specificity rules because the alternatives are themselves too expensive.

6 Conclusion

To summarize, here is the mode of explanation I want to offer for explaining some forms of pro-social, helping behavior.

1. The Cooperation Afforder game (in conjunction with the replicator dynamics) teaches us that a low-cost afforder that supports cooperation and avoids exploitation could, once it appeared even in small numbers, evolve to dominance and even fixation in a population.

[7] Or social selection in Roughgarden's sense (Roughgarden, 2009).

2. If the afforder acquired from such a process were a generalized ability to engage in pro-social thinking (and perhaps communication), this afforder would inevitably be tried in circumstances beyond those originally occasioning it.

3. There are plausible circumstances in which pro-social thinking would out-compete pro-individual thinking, and in consequence be reinforced.

4. A cooperation afforder (including a generalized ability to engage in pro-social thinking and communication) could also grow and be maintained in a population due to its uses in situations that are (somewhat) exploitable.

   Although agents might in principle have incentive to defect in many circumstances from pro-social thinking, the monitoring and discrimination costs may well be prohibitive and/or the frequency of exploitation by counterparts may be low.

5. This process may explain much altruistic behavior. In the small, agents are simply relying on an imperfect heuristic—pro-social thinking—that is generally, usually successful in the circumstances as understood. In the large, catastrophic mistakes in very unusual circumstances are also subject to the logic of cost-benefit tradeoffs for monitoring and discrimination.

The picture, then, is this. Pro-social framing is (with communication) a cooperation afforder that in many circumstances is directly favored by selection and/or learning. It is a heuristic that has been supported by evolution and learning. In consequence, humans pervasively adopt it as the default stance in social encounters, although we can, and do, learn to recognize exceptions. Often, what looks like (or simply is?) altruism is an unexcepted application of a pro-social framing heuristic.

This picture fits nicely with several well-established bodies of experimental findings.

1. Human infants, beginning as early as 14 months, spontaneously exhibit helping behavior. The effects are quite remarkable and have been documented by Tomasello and his group (see (Tomasello, 2009) for an overview).[8]

2. Humans as adults robustly display *strong reciprocity,* that is "people willingly repay gifts and punish the violation of cooperation and fairness norms even in anonymous one-shot encounters with genetically unrelated strangers" (see Fehr and Henrich (2003, page 55), who also provide a useful review of the data).

3. In the well-known Dictator game, one of two agents is given the power to decide on a split of an amount of money. This is played as a one-shot game between anonymous counterparts. Typically, the dictator allocates a non-trivial amount of money to the counterpart, who must accept the dictator's decision. Under the assumptions of classical game theory and economics, the dictator should always appropriate the entire amount, leaving nothing for the counterpart. Interpreted in light of the cooperation afforder hypothesis, allocating a non-trivial amount of money to the counterpart indicates the presence by default of a pro-social framing heuristic and measures its strength or resistance to being excepted.

4. In the well-known Ultimatum game, one of two agents (the proposer) is given the power to propose a split of an amount of money and the other agent (the disposer) is given the power

[8] Perhaps the reader will be struck, as I have been, with the incisive insight revealed in the lyrics of the song "You've Got To Be Carefully Taught" from the musical *South Pacific.* It would not be the first time that a work of art has done this.

to accept or reject the proposal. If the disposer accepts the proposal it is implemented and the game is over; if the disposer rejects the proposal, the game is over and neither player receives any portion of the money. Ultimatum is also played as a one-shot game between anonymous counterparts. Typically, proposers propose a split that allocates a substantial portion (50% or even more, but more usually in the 20-40% range) of the available fund. Also typically, disposers reject portions amounting to less than, say, 20% of the fund.

Under the assumptions of classical game theory and economics, the proposer should always appropriate almost the entire amount, leaving next to nothing for the counterpart, and the counterpart should always accept such a proposal. Interpreted in light of the cooperation afforder hypothesis, the disposer's rejection of ungenerous proposals indicates the presence by default of a pro-social framing heuristic and measures its strength or resistance to being excepted. The proposer's allocating a substantial portion of the fund to the disposer both indicates the presence by default of a pro-social framing heuristic and measures its strength or resistance to being excepted, and indicates the proposer's understanding of the disposer's likely frame of thought.[9]

Now, with the cooperation afforder hypothesis on the table as an evolutionary pathway to cooperation, we may hope that systematic empirical and modeling investigations will explore it and its ramifications, including those pertaining to social policy and design.

It is appropriate at this point to remind the reader that cooperation is normatively ambiguous.[10] If a finite common resource requires cooperation to exploit it, the account here explains how cooperation could be supported by evolution (and more generally an evolutionary dynamics) with the resulting cooperation leading to a tragedy of the commons (Hardin, 1968). Concretely, exploitation of fossil fuels requires extensive cooperative action, yet it may well lead to destruction of the earth's ability to support our species.[11]

[9]The literature on both the Dictator game and the Ultimatum game is voluminous. Camerer (2003) provides a good overview and summary of these games, and of much else that is relevant to the cooperation afforder hypothesis.

[10]I thank an anonymous reviewer for reminding me to remind the reader.

[11]An earlier version of this paper was read at the 2010 Philosophy of Social Science Roundtable, meeting in St. Louis. I am grateful for the ensuing discussion and comments. I also want to acknowledge and thank three anonymous reviewers for their comments, which were materially useful in improving the clarity of the presentation. Finally, this research was funded in part by the Wharton-SMU Research Center at Singapore Management University. The felicitous environment at SMU also aided materially in bringing this work to culmination and I am most grateful for that.

References

Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, 115(3):715–53. doi:10.1162/003355300554881.

Akerlof, G. A. and Kranton, R. E. (2005). Identity and the economics of organizations. *Journal of Economic Perspectives*, 19(1):9–32. doi:10.1257/0895330053147930.

Akerlof, G. A. and Kranton, R. E. (2010). *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being*. Princeton University Press, Princeton, NJ.

Bacharach, M. (2006). *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton University Press, Princeton, NJ. Edited by Natalie Gold and Robert Sugden.

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

Bronstein, J. L. (2003). The scope for exploitation with mutualistic interactions. In Hammerstein, P., editor, *Genetic and Cultural Evolution of Cooperation*, pages 185–202. The MIT Press, Cambridge, MA.

Brown, S. L. and Brown, R. M. (2006). Selective investment theory: Recasting the functional significance of close relationships. *Psychological Inquiry*, 17(1):1–29.

Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Russell Sage Foundation and Princeton University Press, New York, NY and Princeton, NJ.

Colman, A. M., Pulford, B. D., and Rose, J. (2008a). Collective rationality in interactive decisions: Evidence for team reasoning. *Acta Psychologica*, 128:387–397.

Colman, A. M., Pulford, B. D., and Rose, J. (2008b). Team reasoning and collective rationality: Piercing the veil of obviousness. *Acta Psychologica*, 128:419–412.

Dugatkin, L. A. (1997). *Cooperation Among Animals: An Evolutionary Perspective*. Oxford University Press, New York NY.

Fehr, E. and Henrich, J. (2003). Is strong reciprocity a maladaptation? on the evolutionary foundations of human altruism. In Hammerstein, P., editor, *Genetic and Cultural Evolution of Cooperation*, pages 55–82. The MIT Press, Cambridge, MA.

Gigerenzer, G. and Selten, R. (2001). Rethinking rationality. In Gigerenzer, G. and Selten, R., editors, *Bounded Rationality: The Adaptive Toolbox*, pages 1–12. MIT Press, Cambridge, MA.

Gilbert, M. (1992). *On Social Facts*. Princeton University Press, Princeton, NJ.

Gregory, R. L. (1997). *Eye and Brain*. Princeton University Press, Princeton, NJ, fifth edition.

Gregory, R. L. (2009). *Seeing through Illusions*. Oxford University Press, Oxford, UK.

Grice, P. (1989 (originally, 1967)). *Studies in the Way of Words*, chapter Logic and Conversation, pages 22–40. Harvard University Press, Cambridge, MA. ISBN: 0-674-85270-2.

Hammerstein, P., editor (2003). *Genetic and Cultural Evolution of Cooperation*. The MIT Press, Cambridge, MA.

Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859):1243–1248.

Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. The MIT Press, Cambridge, MA.

Henrich, N. and Henrich, J. (2007). *Why Humans Cooperate: A Cultural and Evolutionary Explanation*. Oxford University Press, Oxford, UK.

Hofbauer, J. and Sigmund, K. (1998). *Evolutinary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.

Leimar, O. and Connor, R. C. (2003). By-product benefits, reciprocity, and pseudoreciprocity in mutualism. In Hammerstein, P., editor, *Genetic and Cultural Evolution of Cooperation*, pages 203–222. The MIT Press, Cambridge, MA.

Mann, J., Connor, R. C., Tyack, P. L., and Whitehead, H., editors (2000). *Cetacean Societies: Field Studies of Dolphins and Whales*. University of Chicago Press, Chicago, IL. ISBN-10: 0226503410 ISBN-13: 978-0226503417.

McNeill, W. H. (1995). *Keeping Together in Time: Dance and Drill in Human History*. Harvard University Press, Cambridge, MA.

Mehta, J., Starmer, C., and Sugden, R. (1994). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84:658–73.

Nagel, T. (1979). *The Possibility of Altruism*. Princeton University Press, Princeton, NJ.

Norman, D. A. (2002). *The Design of Everyday Things*. Basic Books. ISBN-10: 0465067107. ISBN-13: 978-0465067107.

Nowak, M. A. and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437:1291–1298. doi:10.1038/nature04131.

Rapoport, A., Guyer, M. J., and Gordon, D. G. (1976). *The 2×2 Game*. The University of Michigan Press, Ann Arbor, MI.

Richerson, P. J. and Boyd, R. (2005). *Not by Genes Alone: How Culture Transformed Human Evolution*. The University of Chicago Press, Chicago, IL.

Richerson, P. J., Boyd, R. T., and Henrich, J. (2003). Cultural evolution of human cooperation. In Hammerstein, P., editor, *Genetic and Cultural Evolution of Cooperation*, pages 357–388. The MIT Press, Cambridge, MA.

Roughgarden, J. (2009). *The Genial Gene: Deconstructing Darwinian Selfishness*. The University of California Press, Berkeley, CA.

Searle, J. R. (1997). *The Construction of Social Reality*. Free Press.

Silk, J. B. (2003). Cooperation without counting. In Hammerstein, P., editor, *Genetic and Cultural Evolution of Cooperation*, pages 37–54. The MIT Press, Cambridge, MA.

Skyrms, B. (2001). The stag hunt. World Wide Web. Proceedings and Addresses of the American Philosophical Association. `http://www.lps.uci.edu/home/fac-staff/faculty/skyrms/StagHunt.pdf`. Accessed December 2009.

Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Cambridge, UK.

Skyrms, B. (2010). *Signals: Evolution, Learning & Information*. Oxford University Press, Oxford, UK.

Sober, E. and Wilson, D. S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Harvard University Press, Cambridge, MA.

Sugden, R. (2008). Nash equilibrium, team reasoning and cognitive hierarchy theory. *Acta Psychologica*, 128:402–404.

Taylor, P. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.

Tomasello, M. (2009). *Why We Cooperate*. The MIT Press, Cambridge, MA.

Tuomela, R. (2007). *The Philosophy of Sociality: The Shared Point of View*. Oxford University Press, New York, NY.

Tversky, A. and Kahneman, D. (1981). The framing of decision and the psychology of choice. *Science*, 211:453–458.

Wilson, E. O. (2000). *Sociobiology: The New Synthesis, Twenty-fifth Anniversary Edition*. Belknap Press of Harvard University Press, Cambridge, MA.

Bio

Steven O. Kimbrough is Professor in the Operations and Information Management Department at the University of Pennsylvania. Trained in philosophy, he works as a decision scientist and focuses on the science of strategic interaction and on model-based deliberation.

`$Id: mutualism-ramifications-paper.tex 1812 2010-10-24 22:17:06Z sok $`