

6

Altruistic Punishment

Is there one word which may serve as a rule of practice all one's life?" The Master said, "Is not reciprocity such a word?"

Confucius The Analects. (1938[4th C.B.P.])

6.1 Introduction

It has been conventional since Thomas Hobbes' *Leviathan* to attribute the maintenance of social order to states. But, for at least 95 percent of the time that biologically modern humans have existed, our ancestors avoided the chaos of the Hobbesian state of nature without the assistance of governments. Moreover, much of the cooperation observed in modern societies includes self-policing not based on state coercion at all. Without this spontaneous and voluntary prosocial behavior, it is doubtful that social life could go on.

Essential to the maintenance of cooperation without state *fiat* are the two dispositions that make up the behavioral repertoire that we call strong reciprocity, namely, a concern for the well-being of fellow group members who adhere to one's ethical norms, and a willingness to punish those who violate social norms even at a cost to oneself. We will see that the second, altruistic punishment, plays a central role.

The key to the evolution of altruistic punishment, as we will see and was shown in Boyd et al. (2003), is that it differs in a little-noticed way from unconditional altruism. The fitness disadvantage of an unconditional altruist relative to an unconditional defector is independent of the population composition, so whatever the frequency of types in the population, the altruist suffers a fixed disadvantage. This is not the case for the altruistic punisher, who only suffers a fitness disadvantage with respect to those who do not punish when there are norm violators deserving punishment in the group. When a group is composed mostly of altruistic punishers and cooperators, the disadvantage of the punishers becomes relatively small. Thus while punishers may be greatly disadvantaged when their type is rare, making their initial emergence unlikely, once established they can be maintained in

a population as a result of a very weak form of group competition, as we will see presently. Moreover, if punishers refrain from incurring the costs of punishing when they face large numbers of shirkers in a group, as they do in our model, their disadvantage when rare is also very small, being limited to punishing through error, so punishing behavior can grow even when initially absent from the population, and introduced into the population only through a low (0.00005 per birth) rate of mutation.

We begin with an overview of the foundations of social order, surveying the various ways that altruistic punishment may implement a high level of norm adherence even in the absence of states. We then present our model and elucidate its dynamics insofar as this is possible analytically. In Section 6.4 we provide an agent-based version of the model, exploring its dynamic properties under a variety of parameter values. In section 6.5 we extend the model to take account of differing demographic patterns.

The argument to this point in the chapter has focused on the evolutionary dynamics of altruistic punishment, but said relatively little concerning how self-interest, unconditional altruism, and strongly reciprocal motives interact in a cooperative setting to determine the relationship between the personal characteristics of group members and the level of cooperation sustained by the group. We address this issue in section 6.7, where we show, among other things, that a strong reciprocator contributes when he considers that others are contributing adequately, and shirks when he considers that others are shirking excessively. This sections shows that the beliefs, preferences, and constrains approach to modeling choice allows us to draw powerful comparative static results even when preferences are non-standard.

We then ask why altruistic punishment is distinctively human. In the conclusion we point out that, unlike the repeated game models of the previous chapters, the model here conforms to our *desiderata* for an explanation of the evolution of cooperation, namely empirical relevance, plausible information requirements, and evolutionary stability. We also note that a population in which altruistic punishment is common constitutes an evolutionary environment likely to favor the evolution of individual social emotions such as shame and guilt, as well as group-level institutions of socialization supporting general adherence to group-beneficial norms. We return to these two topics in subsequent chapters.

6.2 Foundations of Social Order

Because the models we will presently offer represent radical simplifications of observed social practices, we begin with the ethnographic evidence concerning the maintenance of order in stateless societies. Balikci (1970) reports that among the Netsilik, an isolated hunting people on the Arctic coast,

...there is a general rule...according to which all able bodied men should contribute to hunting, and the returns of the hunt should be shared according to established custom. Any activity in exception to this rule was bound to provoke criticism, various forms of conflict, and frequently social ostracism. (176)...lazy hunters were barely tolerated by the community. They were the objects of back biting and ostracism...until the opportunity came for an open quarrel. Stingy men who shared in a niggardly manner were treated similarly. (177)

In like fashion, Richard Lee (1979):458 describes the moral world of the !Kung in Southern Africa:

The most serious accusations one !Kung can level against another are the charge of stinginess and the charge of arrogance. To be stingy, or far-hearted, is to hoard one's goods jealously and secretively, guarding them "like a hyena." The corrective for this is to make the hoarder give "till it hurts"; that is to make him give generously and without stint until everyone can see that he is truly cleaned out. In order to ensure compliance with this cardinal rule the !Kung browbeat each other constantly to be more generous and not to hoard.

Most of the lethal violence documented by Lee concerns such problems as adultery, not stinginess. By contrast to the reports of Lee and Balikci, Endicott (1988):118 reports horror expressed by a Batek informant at the thought of exiling a member whose laziness had caused some resentment.

Punishment often goes considerably beyond browbeating. Using data from forty-eight small scale societies, Boehm (1993):228 found evidence that arrogant members of the group are constrained by public opinion, criticism and ridicule, disobedience, and extreme sanction:

...assassination is reported in 11 out of the 48 [surviving small-scale societies surveyed]...behaviors that terminated relations with an overly assertive individual or removed him from a leadership role involved 38 of the 48 societies, while in an additional 28 instances the person was manipulated by social pressure...the great majority of these misbehaviors involve dominance or self-assertion. (231)

among simple foragers, ...group execution of overassertive persons seems to be rather frequent. (239)

While these passages are typical of the ethnographic record, the only quantitative investigation of social order in pre-state society of which we are aware is remarkable study by Wiessner (2005). During 1974 and 1996-7, she recorded 308 three- to four-hour time blocks of conversations among the Ju/'hoansi (!Kung) of the Dobe-/Kae/kae area in Northwest Botswana. In the 1970s, the Ju/'hoansi obtained virtually all of their subsistence through hunting and gathering. By the 1990s their society had been transformed, with government rations, wages, the sale of crafts, and old age pensions occupying a substantial role in their subsistence. Of the conversations studied, 56% included criticism of one or more members of the group including harsh criticism in two fifths of the cases, and ridicule and mild forms of expression of displeasure in a fifth of the cases. The remaining two fifths of the cases were explicit but not harsh criticism. Only 7 percent of the conversations included praise. Frequently, targeted behaviors included stinginess, greed or laziness, including reclusive behavior associated with failure to share, big-shot behavior, failure to honor kin obligations, and drunkenness. The intent of criticism was behavior modification. In sixty-nine percent of the cases the target was either present or within earshot. In 21 percent of the cases the target was absent but his close relatives were present. The remaining cases involved coalition building for punishment at a later date.

Wiessner recounts one of these episodes: "In 1974 a highly respected leader and one of the few cattle owners was accused of big-shot behavior and failure to share meat. After suffering extensive criticism in silence, he noted that one of his cows was behaving dangerously, slaughtered it, and distributed the meat widely." Another "big-shot" was initially subjected to ridicule, but when this proved ineffective, harsh criticism followed along with a refusal to share meat with the offender, who eventually left the group for a period. A disruptive woman was ostracized from the group and subsequently

died, after which her family was welcomed back into good standing. In one case criticism aimed at expulsion of a segment of the group led its relocation at a new camp 9km distant after which friendly relations were resumed.

The content and targeting of these conversations did not differ between the 1970s and 1990s. Men and women initiated criticism about equally, but men were overwhelmingly the targets of criticism because they more frequently engaged in big-shot or disruptive behaviors. High status (“strong”) group members engaged in more criticism than others, and were as likely to be targeted as “weak” members. The most respected figure in the camp (the *n!ore kxao*) tended to refrain from criticism, thereby facilitating reconciliation so as to avoid the loss of a group member.

Four aspects of social order among the Ju/'hoansi are important in what follows. First, though Wiessner stresses the many ways that the Ju/'hoansi reduce the costs of maintaining order, those engaging in norm enforcement bore costs in the form of disrupted patterns of sharing. In 3% of the cases violent brawls resulted. Second, most of the criticisms were carried out by coalitions (three or more persons) rather than individuals, and this was true in virtually all of the cases where the target had threatened community stability or harmony. As a result, the costs of engaging in punishment of norm transgressors was relatively low as long as the number of altruistic punishers was sufficiently great. Third, the punishment of those who refrained from punishing was entirely absent: “I have not observed any second order punishment in all of my years working with the Ju/'hoansi.” Fourth, Wiessner finds no evidence that punishing norm transgressors is a signal of otherwise unobservable quality as a mate, coalition partner, or competitor. Single reproductive-age individuals are the least likely to punish, contrary to the predictions of the signaling model. Moreover, while mediation ability is highly valued, those who punished frequently or harshly were not sought out as coalition partners but were considered to be *tchi n!ai*, “angry, sharp or biting things,” and were told to desist.

The other study of the maintenance of order in a pre-state society on which we have based our model concerns ostracism among the Pathan Hill tribes an acephalous, egalitarian kinship-based people in Afghanistan. Niloufer Mahdi (1986) writes:

The most striking use of ostracism among the Pathans is the rejection by the tribe or clan of one of its members whose behavior might lead to a feud...[O]stracism functions...to deter behavior that violates customary legal norms, to punish spe-

cific acts that are culturally defined as improper, and to unify the primary reference group on which individuals depend for protection and economic support. (p. 295)

The Pathans, “do not recognize permanently established or permanently functioning authority either in the central government or at any level of the tribal structure.” (p. 296) Rather, social order is sustained by peer punishment. Ostracism is commonly the fate of those who have violated the *Pukhtunwali* (Code of the Pathans).

The obligation to the Code is not secured in the first instance by any coercive force. Ask any Pathan why he upholds *Pukhtunwali* and the answer will be because of *izzat* (honor). Obedience to *Pukhtunwali* is freely given. 297

Similarly, ostracizing those who have violated the Code “becomes the obligation of every Pathan, acting individually or as part of a relevant tribal segment.” (p. 297) Punishment of those who violate social norms is itself a norm.

Ostracism, writes Mahdi, may include “avoidance, exclusion from social participation, shunning or derision.” (295). If a man has to be avoided, he may not be considered for a matrimonial alliance or invited to participate in a *Jirga* (Council of Elders).”(p. 302) Notably, the aim of ostracism seems to be punishment, not behavior modification:

the unique expression of overt verbal disapproval that is practiced is *paighor* (insult or taunt). Its main purpose is to shame the person. It is most commonly applied in those cases where a man has been perceived to have failed in valor or in the discharge of an obligation of honor. *Paighor* is not given with the conscious aim of prompting its object to action, nor is it calculated to reform. However very often it has precisely that effect. (p. 303)

As in Boehm’s survey, punishment may take the lethal forms, such a *badal* (revenge killing) or physical expulsion, the consequences of which are often fatal:

when a Pathan is exiled from his group his situation becomes untenable. He is vulnerable not only to his *dushman* (the one who has targeted him), but can be victimized by anyone without fear of *badal*. (p. 301)

The fact that *badal* applies equally to all members of one's kin group provides strong incentives for discipline to be applied within families, therefore reducing the frequency of more costly between-kin group feuding (as in the "community responsibility" model of Greif, 2002).

As among the !Kung, the egalitarianism of the Pathan peoples is essential to the process of peer punishment, as is the on-going nature of social interaction:

...neither economic or social status deflects its course...if a man seeking *badal* is weak *vis a vis* his *dushman* he will pass on his obligation to his sons and they in turn to their sons. And if a man is well protected enough to escape *badal* himself, it is extremely doubtful that the protection can be extended to his kin or to successive generations, who would constitute legitimate targets of *badal*. (p. 298)

Like the Ju/'hoansi, the Pathans attempted to minimize the costs of punishment. To mitigate the disadvantages of *badal*, *Pukhtunwali* makes provision for mediation through a *Jirga*...when it appears that honor has been satisfied and it is an opportune time to halt the cycle of violence....Certain conditions such as monetary compensation may be imposed to achieve a balance in the redress of grievances. (p. 299)

We base our model on the structure of interaction among members of the mobile hunter-gatherer bands. Modern accounts of these societies record considerable variety in social organization and livelihood (Kelly 1995, Binford 2001). But widespread participation in joint projects such as hunting, raiding, and common defense as well as the sharing of food, valuable information, and other necessities of survival among many of these societies in the modern world is well established. A good case can be made that these cooperative projects were at least as important among our Late Pleistocene ancestors as they are among modern mobile foraging bands (Boehm 2002,2007).

Our model therefore reflects the following empirical considerations appropriate to a single band of foragers. First, groups are sufficiently small that members directly observe and interact with one another, yet sufficiently large that the problem of shirking in contributing to public goods is present. On the basis of Marlowe's assessment of those foraging groups in the historical or ethnographic record that are most likely to approximate our ancestral foraging bands we take groups to be of about 27 individuals, or nine members per generation (Cavalli-Sforza and Bodmer 1970, Marlowe 2005). Second,

there is no centralized structure of governance (state, judicial system, Big Man, or other) so the enforcement of norms depends on the participation of peers. Third, average relatedness of group members is not sufficient high that altruism could be explained entirely by kin selection mechanisms of the type reviewed in Chapter 4. On the basis of genetic evidence presented in the next chapter, we think it likely that average relatedness in foraging groups was considerably less than that among cousins, so the kin altruism model would support the evolution of an altruism only if the benefit exceeded the cost by a factor of at least eight. Fourth, status differences are limited, especially by comparison to agricultural and later industrial societies, which justifies our assumption that anyone can punish anyone, and our treatment of individuals as similar other than by behavioral type and by the group to which they belong. Fifth, the sharing on which our model is based, either of food individually acquired or of the common work of acquiring food, for example, is characteristic of these societies. For a recent survey see Kaplan and Gurven (2005). Sixth, the individuals in our model do not store food or accumulate resources. This, too, is a characteristic of at least those hunter-gather bands based on what Woodburn (1982) calls an “immediate return” system of production. Seventh, consistent with genetic data surveyed in Chapter 11 and Bowles (2007), and with direct observation (Woodburn 1982), membership in our groups is quite fluid, with considerable between-group migration due to exogamy, among other reasons. Finally, our model supports behavioral heterogeneity in the long-run, a property that corresponds to the ethnographic record of foraging bands, as well as from the experimental evidence on both hunter gatherers (Henrich et al. 2004) and modern market based societies (Loewenstein, Thompson and Bazerman 1989, Andreoni and Miller 2002).

6.3 Working, Shirking and Punishing

Consider a population in which, in each period, by cooperating an individual produces an amount b that is shared by all group members, at personal cost c . Thus, if all group members cooperate, each has a net payoff $b - c > 0$ per period. We assume behavior is the expression of genetic inheritance and individuals are haploid (an individual has one copy of each gene), but reproduction is diploid (an individual inherits each gene with equal probability from one of two parents).

Individuals have two genetic loci. One locus determines if an individual is a Selfish or a Cooperator type. A Cooperator cooperates unconditionally, producing b for the group at personal cost c , while a Selfish type cooperates only if the cost of being punished, s , times the probability of being punished p , is greater than the cost c of cooperating. However, with a small probability ϵ (1.5% in our simulations), which we call the *execution error rate*, Selfish types miscalculate and defect even when this is more costly than cooperating, or Punishers punish shirking even when there are too few Punishers in a group to deter shirking. Individuals have a second genetic locus that determines whether they are Punishers or Nonpunishers. Punishers impose a cost s on each member who is detected shirking, where the probability of detection is proportional to the fraction of Punishers in the group. The cost of punishing a shirker, c_p , is shared by all Punishers in the group. Nonpunishers never punish. If Punishers are sufficiently outnumbered by Selfish types in a group, they refrain from punishing. We implement this by specifying the maximum number n_{\max} of punishments a single Punisher will mete out, set to $n_{\max} = 1.7$ in the simulations. Thus, if the ratio of Selfish types to Punishers is greater than n_{\max} , Punishers do not punish, except accidentally, which occurs with probability ϵ . Note that Punishers have a $\epsilon(1+n_{\max})c_p = 3.7\epsilon c_p$ disadvantage as compared with Nonpunishers even when there is no intentional shirking on the part of Selfish types.

There are therefore four types of agents: Cooperator-Punishers (CP), Selfish Punishers (SP), Cooperator-Nonpunishers (CN), and Selfish-Nonpunishers (SN). We also tested a model with unconditional defectors, but such individuals are quickly eliminated from the population because they defect even when they would achieve higher payoffs by cooperating. Mutation occurs, so that a newborn inherits one of his parents' Selfish/Cooperator genes, but with probability $\mu/2$ mutates from Selfish to Cooperator or *vice-versa*, and inherits one of his parents' Punisher/Nonpunisher genes, but with probability $\mu/2$ mutates from Punisher to Nonpunisher or *vice-versa*. We call μ the *rate of mutation*.

The spatial structure of our model is as follows. Each group is represented by a node on a torus (a two-dimensional surface without boundaries), so that each group has exactly eight neighbors. When individuals migrate, they move to a neighboring group, and when they mate, they choose a mate from a neighboring group. When a group becomes unsustainably small (size less than n_{\min}), it recruits new members from neighboring groups.¹

At the end of each production period, after punishment, if any, is inflicted, a fraction r of each group reproduces. An individual A is chosen randomly from the population and is permitted to reproduce with a probability proportional to his payoff, in which case a mate B for A is chosen in the same manner from among A's eight neighboring groups. A and B then have a single offspring and a random member of the population is killed, so population size remains constant. Each offspring is randomly assigned to either A's group or B's group. We assume full recombination so that, in the absence of mutation, the offspring of a Selfish-Nonpunisher and a Cooperator-Punisher has equal probabilities of being each of the four possible types. Note that on average each individual is chosen as a candidate for reproduction once as the A party (randomly from the population) and once as the B party (a random member of the neighborhood of an A party).

6.4 The Evolution of Strong Reciprocity

Groups change size through reproduction, death, migration, and repopulation of unsustainably small groups. We assume that if a group falls below some n_{\min} in size, its size is brought up to initial group size by randomly moving members from the largest group in the neighborhood to the under-sized group. This procedure ensures that extremes in group size are avoided.

For our baseline simulation, we set up 100 groups, each starting out with nine members of a single generation. We set the initial frequency of Selfish individuals at 50%, and the initial frequency of Cooperators at 50%. These levels allow the simulation to achieve its stationary distribution rapidly. As we explain below, even starting out with no Punishers, the model eventually, through the mutation process, produces a few Punishers, and when these Punishers happen to be in the same group, the group does sufficiently well that the frequency of Punishers in the population increases. In the long run, the system achieves a mean level of shirking independent from the initial frequency of Punishers.

We assume 25 production periods occur per generation, so the periods are roughly years. Each individual mates and reproduces on average once per generation, so the per-period reproduction rate is $r = 0.04$. We assume also

¹All of our results obtain if we replace migration and mating based on physical propinquity by random migration and random mating, although in the latter case the number of periods before cooperation is established is generally larger than in the local migration and mating case.

Variable	Value	Description
b	2	Output per Individual, No Shirking
c	1	Cost of Production, No Shirking
c_p	1	Cost of Punishing
s	2	Cost of Being Punished
r	0.04	Reproduction Rate per Period
μ	0.0001	Mutation Rate
ϵ	0.015	Execution Error Rate
m	0.02	Migration Rate per Period
n	100	Number of Groups
pf	0-1	Resource Constraint
n_{\max}	1.7	Selfish to Punisher Ratio, below which Punishers punish
	9	Initial Size of one generation/group
n_{\min}	4	Minimum Group Size
	0.50	Initial Fraction of Punishers
	0.50	Initial Fraction of Selfish

Table 6.1. Baseline Parameters. These parameters are used in all simulations, unless otherwise noted.

that half of all group members migrate per generation, so the migration rate is $m = 0.02$ per period. The baseline parameters are listed in Table 6.4. Note that the total migration rate is considerably larger than this because half of a parent's offspring locate in groups other than the parent's, and there is considerable repopulation of under-sized groups by their neighbors.

Using these parameters, Figure 6.1 shows the evolution of the distribution of individual types and the average shirking rate for the whole population over 15,000 periods, which we think of as 600 generations. The results shown from period 100 on are backward-moving averages over the previous 100 periods. To study the behavior of the model, we ran the simulation 10 times with the baseline parameters for 30,000 periods, and calculated the average fraction of each type and the average shirking rate, averaged over the last 1000 periods. These averages and their standard deviations are reported in the figure legend.

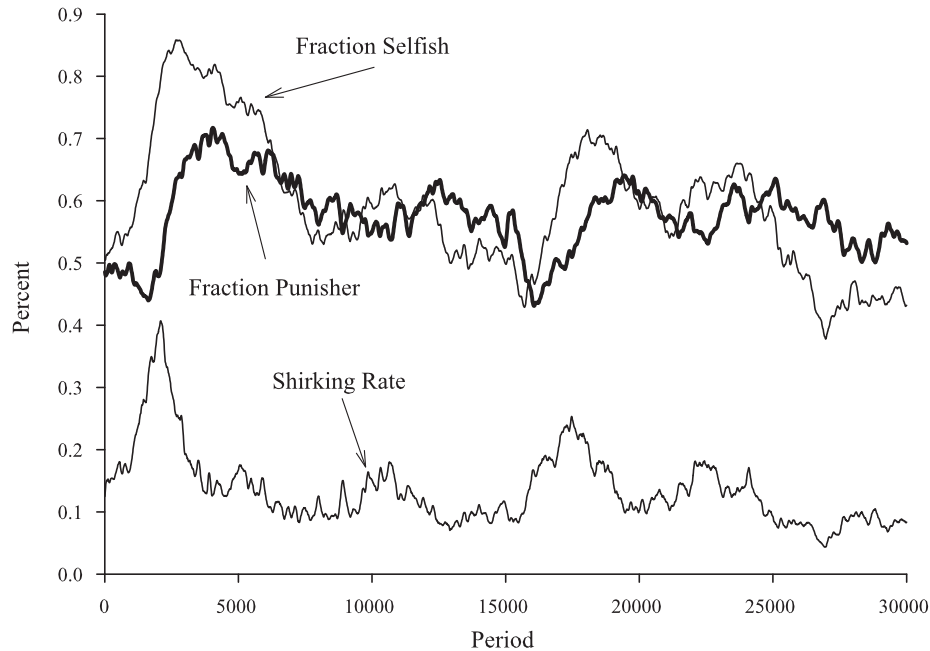


Figure 6.1. The Evolution of Strong Reciprocity. The statistics are 100 period running averages. The baseline parameters are as in Table 6.4. The means and (standard deviations) over the final 1000 periods (40 generations), were 0.89(0.11) for Fraction Punisher, 0.50(0.30) for Fraction Selfish, 0.46(0.20) for Fraction Selfish Punisher, 0.43(0.19) for Fraction Cooperative Punisher, 0.05(0.02) for Mean Shirking Rate.

Punishers evolve in this model despite that fact that, in every group in which they punish, they have lower fitness than Nonpunisher group members. The capacity of Punishers to evolve depends on the fact that members of groups with many Punishers have higher fitness than groups with few Punishers, and most Punishers are in groups with many Punishers. To see that this is the case, we calculate two numbers that capture respectively the payoff advantages of the punishers resulting from the fact that most are located in highly cooperative groups, just mentioned, and their within-group payoff disadvantages, arising from the fact that they bear the costs of punishing shirkers, which the Cooperators avoid, and the costs of contributing, which the shirkers avoid when there are few Punishers in the group.

The first we call the across-group fitness f_a of Punishers, defined as the expected value of average fitness of group members, weighted by the fraction of Punishers in the population that are members of the group, minus the average fitness of the whole population. This is the advantage enjoyed by punishers (compared to the population average) resulting from their location in high payoff groups. The second number is the within-group fitness of Punishers f_w , calculated as the average across groups of the fitness of Punishers within a group minus the average fitness of members of the same group. Figure 6.2 shows the averages of f_a and f_w over the previous 100 periods. The figure shows that while Punishers are altruistic in the sense of sacrificing ($f_w < 0$) in a manner that increases the fitness of group members, they thereby increase their fitness relative to the population as a whole ($f_a > 0$). Thus, groups composed entirely of Cooperators and Punishers will be very successful, but if there are a few shirkers requiring punishment, Nonpunishers will do better than the Punishers, eventually taking over and opening the door to a takeover of Selfish types who in the absence of Punishers will shirk, leading the group to have low average payoffs.

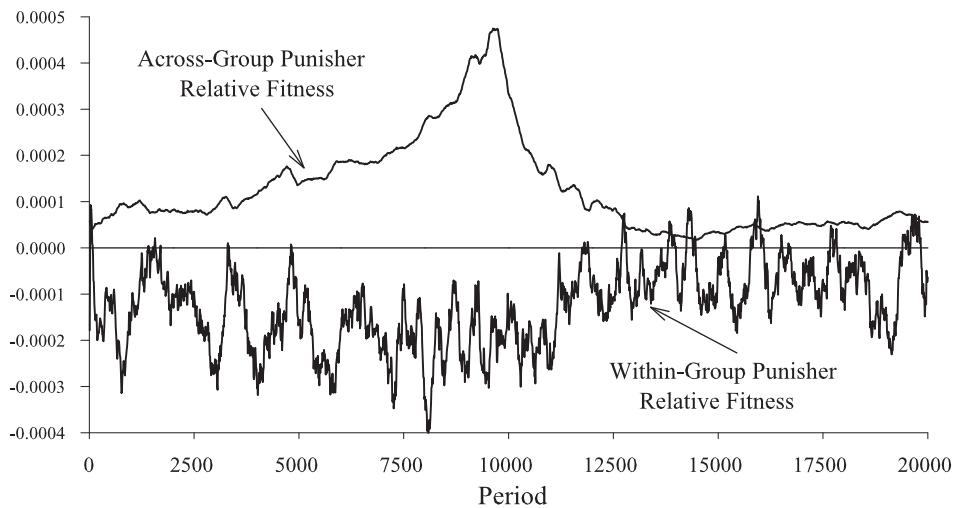


Figure 6.2. The Dynamics of the Evolution of Punishing

Figure 6.3 show that punishers can evolve even when very rare. In this case, there are zero punishers at the start of the simulation. This occurs because eventually one group will have enough Punishers to induce Selfish

types to cooperate, thus implementing a high level of cooperation with very low costs of punishing, since the only defections in this situation are due to rare execution errors by Selfish types. When Punishers predominate in a group, all members have high fitness, and the group grows in size, and then seeds other groups by migration and repopulates the sites of undersized groups.

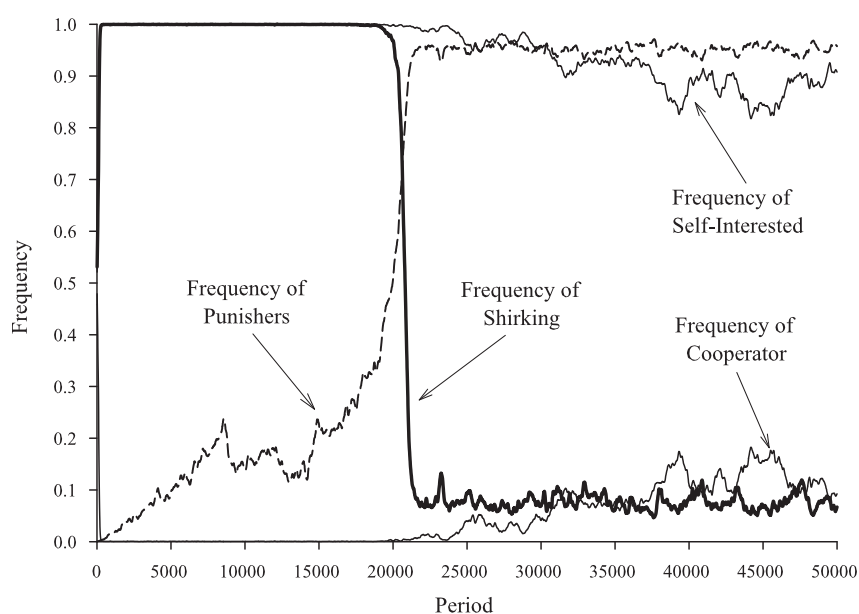


Figure 6.3. The Emergence of Punishers from a Nonpunishing Initial State. This simulation uses the same parameters as before, but the initial fraction of Punishers is zero, so Punishers appear only through mutation. For the first 11,000 periods the fraction of Punishers remains extremely small and the shirking rate is unity, since Selfish Nonpunishers drive out Cooperator Nonpunishers. Thereafter, the rapid increase in Punishers drives the shirking rate to between 10% and 20%.

Our model, and that of Boyd, Gintis, Bowles and Richerson (2003), illustrate a dynamic observed in many of the models to follow: long term persistence of a particular distribution of behavioral types in a population

is consistent with on-going change in the distribution of types within subgroups of the population. The same dynamic occurs in Chapter 7: altruists are always selected against within groups and hence their within-group frequencies decline, but differential group extinction and out-migration from large groups, which favor groups with higher fractions of altruists, may balance these within-group selection effects.

6.5 Variations and Extensions

Our model is also readily extended to other aspects of the enforcement of norms that have a public goods character. To see this, consider the commonly observed norm in foraging societies that restricts copulations to monogamous pairs, which when violated leads to strife within a group or lessens its effectiveness in acquiring food, insuring against stochastic events, sharing information, or defending itself, all of which reduce fitness levels of group members. Those who violate the norm, however, like our selfish individuals, enhance their fitness unless punished. Cooperator-Punishers never violate the norm and punish those who do. Cooperator-Nonpunishers adhere to the monogamy norm but never punish. Selfish individuals violate the norm when the expected cost of punishment is sufficiently low. This model could easily be cast in the form that we have just simulated. We therefore believe the model is applicable to a wide range of concrete problems of norm adherence likely to arise in small stateless groups.

In further simulations, we explore a large parameter space with two objectives. First, we check that the simulation results respond in plausible ways to parameter shifts, thus confirming our interpretation of the causal structure of the model. Second, we show that the model works for values approximating human ancestral environments. See Figure 6.4. In these simulations, we created 1000 groups and followed them through 60,000 periods (2400 generations), the recorded figures being averages over the final 10,000 periods. We also initialized the population with zero Punishers, so we are here investigating the conditions under which cooperation through strong reciprocity can evolve, rather than the conditions under which it can persist once established. The conclusions differ little when the persistence question rather than the emergence question is posed, except that migration tends not to disrupt cooperation once Punisher become prevalent, since most migrants are then punishers, who do not disrupt cooperation within groups.

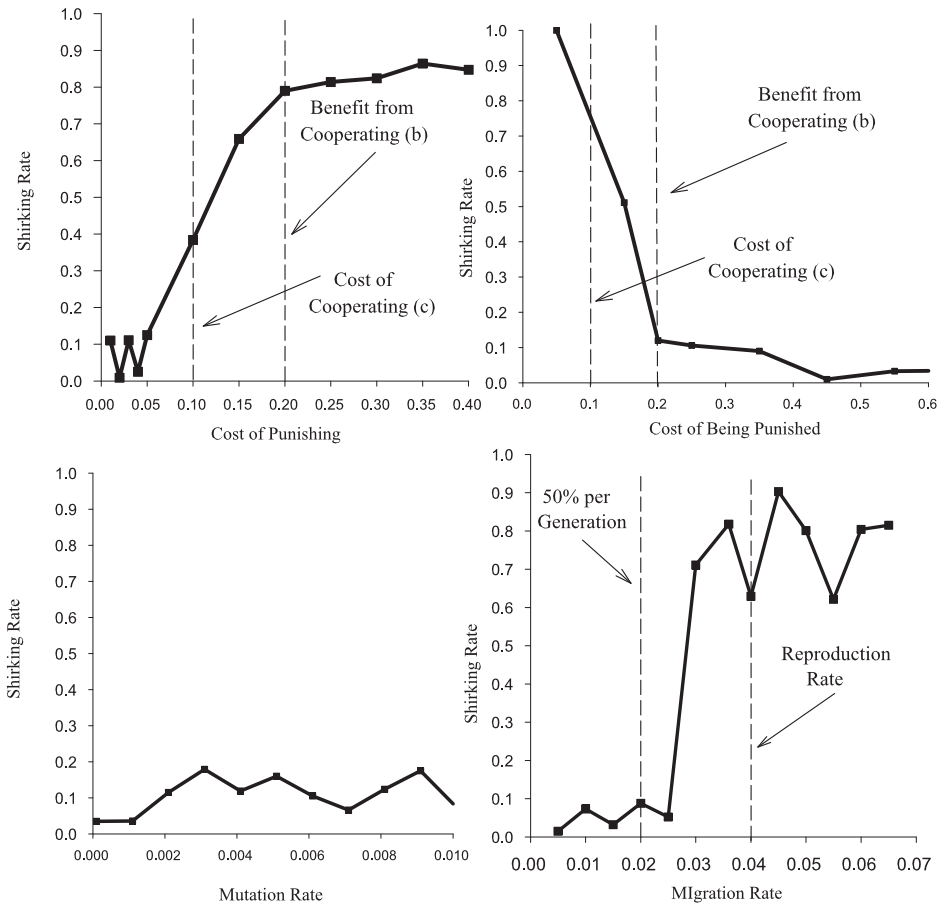


Figure 6.4. Sensitivity to Variation of Model Parameters

The upper left graph in Figure 6.4 shows that when the cost of punishing a single shirker, which is shared by all punishers, is less than about the cost of cooperating (c), a high level of cooperation can be maintained, as the average shirking level in the population is about 10%. However, when the cost of punishing is greater than $c/2$, the amount of shirking rises rapidly. As we might have expected, a low punishing cost is a prerequisite to effective cooperation through strong reciprocity. The top right graph shows that the cost of being punished is also an important determinant of the level of cooperation. The lower left panel shows that over a wide range, the mutation rate has no systematic effect on the level of cooperation. This is important to know, because the effects of excessively high or low mutation rates in some models can be erroneously attributed to the major causal variables investigated in the model. The graph on the lower right shows that the model responds in the expected manner to changes in the migration rate. When migration rates are very low, positive assortment of Punishers allows cooperation flourish, but this positive assortment is undermined when the migration rate is very high. Since a 50% migration rate, when added to the other reasons for between group movement in the model per generation is considered high, we can assert that cooperation can evolve using plausible migration rates.

Punishers and Cooperators benefit from between-group differences in their prevalence because most are in groups with above average numbers of individuals with like traits. But, for the same reason, Punishers and Cooperators are also more likely to compete over group specific resources (Taylor 1992, Wilson, Pollock and Dugatkin 1992). To this point we have ignored the ecological limits of sites and the resulting density-dependent selection process by assuming that groups are able to grow without imposing fitness costs on their members. Is the evolutionary process we have just described robust in the presence of resource constraints?

Technically, the ‘no resource constraints’ condition is implemented by our assumption that the fitness of each individual is proportional to his share of payoffs in the population as a whole. Were groups at their maximum size so that resource constraints were absolute, however, the selection process would make fitness proportional to an individual’s share of the group’s payoffs, reflecting the fact that members of a group are competing with other group members, not other population members, for reproductive success. We know that if we implement absolute resource constraints, neither Cooperators nor Punishers can evolve, because their evolutionary advantage is their tendency

to be in groups that produce more offspring than other groups, and if this is precluded so that they compete only against fellow group members, then their within group payoff disadvantages will result in their elimination.

To explore the consequences of resource constraints, we let the number of offspring of a couple be a weighted average of their fraction of the groups payoffs multiplied by group size and their fraction of the entire population's payoffs multiplied by total population size. The weight in this average is the degree of resource constraint ranging from 0, the model above, to 1, an absolute binding limit on group size. Figure 6.5 presents the results of this exploration. As expected the resource constraints penalize Punishers and favor shirking; but a substantial level of Punishers are sustained in the population even for quite stringent resource constraints (up to 0.6).

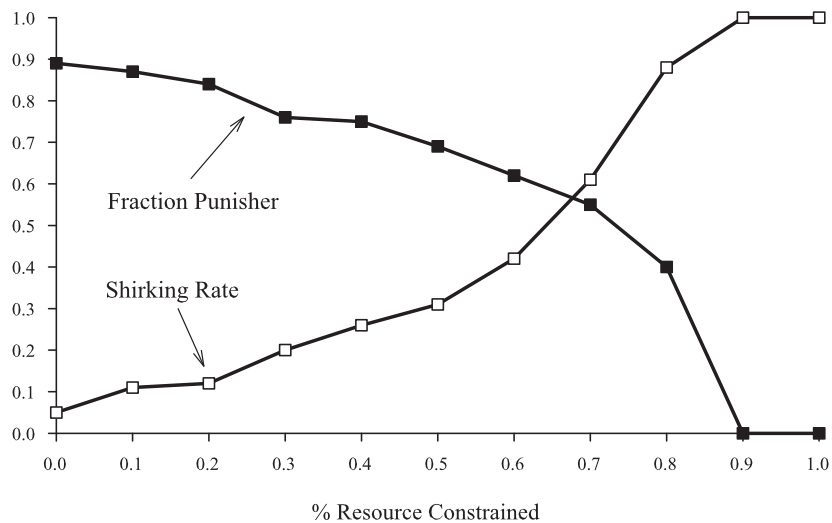


Figure 6.5. Sensitivity to Local Resource Constraints

But how could strong reciprocators come into existence *de novo*? We speculate that they could have emerged through a rather trivial modification of fitness-enhancing behaviors. For example, strong reciprocity among kin might have emerged and proliferated as a form of kin-based altruism, and then be generalized to unrelated individuals. Another possibility is that strong reciprocity arose through a modification of the individual fitness

enhancing strategy of reciprocal altruism. In this case the modification is trivial: simply ignore the future payoffs to current behavior. Like the extension of kin-based strong reciprocity to non-kin, the mutation or mutations to convert reciprocal altruist strategy to a strong reciprocal one involves a reduction in complex discrimination rather than an increase in complexity. These strategy conversions thus might occur with high probability.

Because offspring of group members remain in the group of one parent until they either migrate or die, our model is also likely to have some kin selection involved in both strengthening the altruism of Punishers and damping the free-riding of Cooperators. Since individuals on average have only one offspring who remains in the group, and since the migration rate is 50% per generation, the amount of kin selection is very small. We confirmed this by verifying that the behavior model does not change when we locate offspring in a group other than that of either parent.

Our simulations also address the problem of second order free riding, one of the most persistent criticisms of models of altruistic punishment (see also Boyd, Gintis, Bowles and Richerson 2003). Our Nonpunishers never punish, yet they never drive out Punishers from the population, and in equilibrium, even when the shirking rate is 10% or lower, the frequency of Nonpunishers in the population in the long run varies from 10% to 50%, depending on chance and the parameters of the model. The reason for the persistence of Nonpunishers is that when a Nonpunisher replaces a Punisher in a group, the expected level of shirking in the group increases, so Nonpunishers “pollute their own nests.”

It has been suggested that models in which strong reciprocity evolves depend on tightly linking the altruistic cooperation and altruistic punishment aspects of strong reciprocity, and that were they permitted to evolved independently, altruistic punishment would disappear (Lehmann and Keller 2007, Lehmann, Rousset, Roze and Keller 2007). Our analysis here shows that this concern is unfounded. In our model, punishing vs. non-punishing and altruistic cooperation vs. self-interested behavior are not linked in any way. These traits evolve independently, and are constantly broken up by recombination in the reproduction process. By assuming 100% exogamy and the absence of positive assortment in reproductive matching, we surely exaggerate the forces tending to separate rather than to link cooperation and punishment. Nonetheless, a high level of cooperation evolves with a high frequency of Punishers.

6.6 Why is Strong Reciprocity Uniquely Human?

Our model is extremely simple and the group sizes, migration rates, and other parameters range over which it exhibits a high level of altruistic punishment and hence effective cooperation should apply to many primate species and even social mammals, such as wolves, lions, and other cooperative hunters. Yet, altruistic punishment is rare if not entirely absent in other species, excepting the eusocial insects. We suggest that the following factors account for this unique position of *Homo sapiens*.

First, humans are capable of inflicting punishment upon transgressors at very low cost to the punishers. While size, strength, and vigor generally determine the outcome of animal disputes, victory often involving great cost even to the winner, in human societies even a small number of attackers can defeat the most formidable single enemy at very low fitness cost to the attackers through the use of coordination, stealth and deadly weapons. Bingham (1999) has correctly stressed the importance of the superior abilities of humans in clubbing and throwing projectiles as compared with other primates, citing Goodall (1964), Plooij (1978) on the relative advantage of humans, and Darlington (1975), Fifer (1987), and Isaac (1987) on the importance of these traits in human evolution. Calvin (1983) argues that humans are unique in possessing the same neural machinery for rapid manual-brachial movements that allow for precision stone-throwing.

A second factor in the success of strong reciprocity is the capacity of punishers to coordinate their activities. Punishers do not punish unless their frequency in the group is sufficiently high that the cost of punishing to an individual Punisher is not excessive. This requires truthful communication among group members as to which members have shirked and which have not, as well as how many Punishers will share the costs of punishing. Thus, the effectiveness of altruistic punishment is predicated upon a broader social structure of mutual cooperation based on truthful communication and the widespread use of social sanctions.

If our speculation that strong reciprocity emerged through a modification of reciprocal altruist behaviors is correct, this provides another reason why strong reciprocity might be uniquely human, given that reciprocal altruism appears to be rare in other species (Stephens et al. 2002, Hammerstein 2003).

We have modeled a process of genetic transmission. But were the altruistic punishment trait subject to cultural transmission, as it certainly is in humans, then additional reasons for human uniqueness arise. The process of decay of a

state in which virtually all group members are Punishers leading to a takeover of the group by Selfish types (described above) could be slowed or even halted if cultural updating responded not only to payoff differences but also to the frequency of types in the group, so that more common behaviors were preferentially copied. Guzman, Sickert and Rowthorn (2007) demonstrate that conformist cultural transmission of this type supports the evolution of altruistic punishment even for groups of enormous size (over 1,000 persons). They also demonstrate that a genetic predisposition to engage conformist updating could co-evolve with altruistic punishment in this environment.

6.7 From Evolutionary to Proximate Explanations of Strong Reciprocity

Our evolutionary model shows how a predisposition to punish selfish behavior could have evolved. For this to have occurred people must have come to have proximate motives inducing them to act in this way—whether anger at transgressors of norms, elevation at performing one’s civic duty, or other. It remains, therefore to study the preferences that could induce people to behave as strong reciprocators. This will allow a more nuanced treatment of behavior in which individuals may continuously vary the amounts they contribute to a common project, and how much punishment they mete out to others. We will also replace the four genotypes of the evolutionary model with a continuum of types.

In this section we use a simple representation of the utility function of individuals with other-regarding preferences, and show that with this specification, individuals with strongly reciprocal motives do indeed behave in this manner.

To model strategic interaction with social preferences in a standard social dilemma, the n -person public goods game, we say that individual i ’s utility depends on his own material payoff π_i and the payoff π_j to other individuals $j \neq i$ according to:

$$u_i = \pi_i + \sum_j (\alpha_i + \lambda_i \beta_{ij}) \pi_j, \quad (6.1)$$

where $\alpha_i, \beta_{ij} \in \mathbf{R}$ and $\lambda_i \geq 0$ (Rabin 1993, Levine 1998). The parameter α_i is i ’s level of unconditional altruism if $\alpha_i > 0$ and unconditional spite if $\alpha_i < 0$, β_{ij} , which may be of either sign, is i ’s assessment of j ’s type, and λ_i

is i 's strength of reciprocity motive, valuing j 's payoffs more highly if j 's type conforms to i 's concept of good behavior, and conversely. If λ_i and α_i are both positive, the individual is termed a strong reciprocator, motivated to behave generously towards individuals about whom he knows nothing, but willing to reduce the payoffs of an individual who reveals a bad type even at a cost to himself. Table 6.2 shows the correspondence between the genotypes of our evolutionary model and the parameters of the preference function (6.1).

	P: $\lambda > 0$	N: $\lambda = 0$
C: $\alpha > 0$	PC: Strong Reciprocator	NC: Unconditional Altruist
S: $\alpha = 0$	PS	NS: Homo economicus

Table 6.2. Social and Self-regarding Preferences Associated with the Four Genotypes

Consider a group of size $n > 1$, where member i supplies an amount of effort $1 - a_i \in [0, 1]$. We call a_i the *level of shirking* of member i , and write $\bar{a} = \sum_{j=1}^n a_j/n$ for the average level of shirking. We assume working ($a_i = 0$) adds $\chi > 1$ dollars to the output of the common project, while the cost of working is one dollar. We call χ the *productivity of cooperation*. We assume the members share their output equally, so member i 's payoff is given by

$$\pi_i = \chi(1 - \bar{a}) - (1 - a_i). \quad (6.2)$$

The payoff loss to the group from one member shirking is χ , of which the shirker's share is χ/n , so the shirker's net gain from shirking is

$$\frac{\partial \pi_i}{\partial a_i} = 1 - \frac{\chi}{n}, \quad (6.3)$$

which we assume is strictly positive.

We assume member i can impose a monetary cost μ_{ij} on $j \neq i$ at cost $c_i(\mu_{ij})$ to himself. We treat μ_{ij} as the monetary equivalent of such forms of punishment as public criticism, shunning, ostracism, physical harm and the like. We assume that acts of punishment, like work effort, are non-verifiable

and hence not subject to contract. We also assume $c_i(0) = c'_i(0) = c''_i(0) = 0$ and $c_i(\mu_{ij})$ is increasing and strictly convex for all i, j with $\mu_{ij} > 0$.

Member i can judge member j 's type only on the basis of j 's level of shirking, which we assume i observes correctly, and the harm that inflicts on i , which we assume is public knowledge. Therefore, we specify

$$\beta_{ij} = \chi(1 - 2a_j)/n \quad (6.4)$$

in (6.1), so $\beta_{ij} = -\chi/n$ if i observes j completely shirking, $\beta_{ij} = \chi/n$ if i observes j completely working, and $\beta_{ij} = 0$ if i does not observe j 's behavior. This means that $a_j = 1/2$ is the point at which i evaluates j 's type as neither good nor bad. This point could be shifted to any value between 0 and 1 by treating β_{ij} as a quadratic in a_j , but the added generality is not illuminating. We then define

$$\rho_{ij}(\mu_{ij}, a_j) = (\lambda_i q(2a_j - 1)/n - \alpha_i)\mu_{ij} - c_i(\mu_{ij}), \quad (6.5)$$

which is i 's *subjective gain* from punishing j when j shirks at level a_j and i chooses a level μ_{ij} of punishment. Note that i 's propensity to punish $j \neq i$, $\rho_{ij}(\mu_{ij}, a_j)$, is not affected by the propensities to punish or the observed rates of punishing of other members. This means that punishing is ‘‘warm glow’’ rather than instrumental towards affecting j 's behavior, or towards inflicting the ‘right’ amount of punishment on j . Of course these other motives for punishment may be important, but the evidence for the existence of ‘warm glow’ punishment is convincing, and the case worth exploring (Casari and Luini 2007).

Member i will then choose $\mu_{ij}^*(a_j)$ to maximize utility in (6.5), which is equivalent to maximizing u_i in (6.1), giving rise to the first order condition (assuming an interior solution)

$$c'_i(\mu_{ij}^*) = \lambda_i q(2a_j - 1)/n - \alpha_i, \quad (6.6)$$

or the marginal cost of punishing is equal to the marginal benefit of reducing j 's payoffs given i 's assessment of j 's type. When

$$a_j \leq a_i^0 = \frac{1}{2} \left[\frac{n\alpha_i}{\lambda_i q} + 1 \right], \quad (6.7)$$

the maximization problem has a corner solution in which i does not punish. For $a_j > a_i^0$, denoting the right hand side of (6.6) by ϕ , and differentiating

(6.6) totally with respect to any parameter x , we get

$$\frac{d\mu_{ij}^*}{dx} = \frac{\partial\phi}{\partial x} \frac{1}{c_i''(\mu_{ij}^*)}. \quad (6.8)$$

In particular, setting $x = a_j$, $x = \lambda_i$, $x = \chi/n$ and $x = n$ in turn in (6.8), we see that for $a_j > a_i^0$, the level of punishment by i μ_{ij}^* , is (a) increasing in the level a_j of j 's shirking; (b) increasing in the level λ_i of i 's reciprocity motive; (c) increasing in the harm χ/n that j inflicts i by shirking; (d) decreasing in group size n ; and (e) decreasing in i 's unconditional altruism α_i .

For notational convenience we write $\rho_{ij}(a_j) \equiv \rho_{ij}(\mu_{ij}^*(a_j), a_j)$. Since $c_i''(0) = 0$ and $\mu_{ij}^*(a_j^0) = 0$, $\mu_{ij}(a_j)$ is differentiable.

Let $\mu_j(a_j)$ be the expected punishment inflicted by all $i \neq j$ on j if j shirks. We have

$$\mu_j(a_j) = \sum_{i \neq j} \mu_{ij}^*(a_j), \quad (6.9)$$

which is then differentiable and increasing in a_j . Since it is plausible that agents will never punish another agent who does not shirk, we assume $\alpha_i/\lambda_i + 1 > 0$ for all i , so $\mu_i(0) = \mu_i'(0)$ for all i .

From the first order condition on a_i from (6.1), we must have

$$1 - \frac{\chi}{n} = \frac{\chi}{n} \sum_{j \neq i} (\alpha_i + \lambda_i \beta_{ij}) + \mu_i'(a_i), \quad (6.10)$$

so the member shirks up to such point as the net benefits of shirking (the left hand side) equal i 's valuation of the cost imposed on others by his shirking (the first term on the right hand side) plus the marginal cost of shirking entailed by the increased level of punishment that the shirker may expect. This defines i 's optimal shirking level a_i^* for all i , and hence closes the model. Differentiating the first order condition $\partial u_i / \partial a_i = 0$ totally with respect to a parameter x gives

$$\frac{\partial^2 u_i}{\partial a_i^2} \frac{\partial a_i}{\partial x} + \frac{\partial^2 u_i}{\partial a_i \partial x} = 0. \quad (6.11)$$

We say that i 's partners *shirk on balance* if

$$\sum_{j \neq i} \left(a_j - \frac{1}{2} \right) \geq 0. \quad (6.12)$$

From the fact that $\partial^2 u_i / \partial a_i^2 < 0$ by the second order condition, it is straightforward to show that (a) an increase in unconditional altruism α_i leads i to shirk less; (b) the stronger i 's strength of reciprocity motive λ_i , the more i shirks when i 's partners shirk on balance, and the less i shirks otherwise; (c) the larger group size n , the more i shirks, provided his partners shirk on balance; and (d) the greater the productivity of cooperation χ , the more i works, provided his partners do not shirk on balance.

6.8 Conclusion

Our reading of the ethnographic and paleoanthropological evidence is that our model may capture the salient social and ecological conditions of the late Pleistocene. Thus, we think that our model, suitably extended, can capture the environments that may have supported high levels of cooperation among our ancestors living in mobile foraging bands during the late Pleistocene. We do not know that a human predisposition to strong reciprocity evolved as we have described. But our model and simulations suggest that it could have.

If cooperation did evolve by this route, it is not surprising that, as Cosmides and Tooby (1992) and others have shown, cheater detection is an advanced cognitive capacity of humans. The model may also resolve one of the big puzzles raised by behavioral experiments: the fact that while contributing to a public good and punishing those who fail to do so are both altruistic behaviors, subjects are considerably more avid about the latter than the former. As we have seen in the Introduction of this chapter and in Chapter 3, the experiments of Ernst Fehr and his collaborators show that inflicting punishment on norm violators is a strong human motive. The fact that experimental subjects actively take pleasure in punishing transgressors, as indicated by their own accounts and by behavioral neuroscience experiments, could thus be the result of the evolutionary processes we have modeled.

Our model presents a sharp contrast to the repeated game approaches surveyed in the previous chapter. First, it is consistent both with the empirical facts about forager societies outlined at the beginning of this chapter. Second, cooperation persists over thousands of generations despite the fact that there are no cooperative asymptotically stable Nash equilibria in this model. Finally, unlike the fictive strategies invoked by the Folk Theorem, the main causal mechanism in our model, the altruistic punishment of norm violators, is widely observed, as we have seen in Chapter 3.

The behavior of the Selfish types in our model remains to be explored in greater depth, however. We have assumed that they contribute when it is fitness maximizing to do so. But, were we to drop the fiction that fitness is explicitly maximized and introduce proximate motives, it is something of a mystery how this could occur. The benefits of shirking occur now, while the punishment should it occur, happens later. Humans are more impatient than would maximize their fitness and thus would not correctly weigh the future consequences of shirking even if they were correctly known. We think that the evolution of the social emotions like shame may be explained in part by their ability enhance the present motivational salience of future punishments and thus to offset what would otherwise be a fitness reducing myopia of would be shirker's. In the presence of altruistic punishment by one's fellow group members, individuals with levels of impatience that lead them to shirk when it is not fitness maximizing to do so would improve their fitness were they to be motivated in part by feelings of shame and guilt. Boehm (2007) calls this process 'sanctioning selection.' We will return to it when we address the social emotions in Chapter 10.

A further implication to be taken up in Chapter 9 is that individuals who developed the capacity to internalize group beneficial norms, and groups that devote their socialization practices to this end would be evolutionarily favored.

In this chapter we have opted to clarify the evolution of altruistic punishment in the simplest possible manner, postponing our investigation of the essential roles played in human cooperation by group level social institutions and between-group conflict. To these we now turn.