


This article was downloaded by: 10.2.97.136

On: 28 Mar 2023

Access details: *subscription number*

Publisher: *Routledge*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK



Handbook of Contemporary Behavioral Economics Foundations and Developments

Morris Altman

Group Selection and Behavioral Economics

Publication details

<https://test.routledgehandbooks.com/doi/10.4324/9781315703879.ch8>

Alexander J. Field

Published online on: 15 Jul 2006

How to cite :- Alexander J. Field. 15 Jul 2006, *Group Selection and Behavioral Economics from: Handbook of Contemporary Behavioral Economics, Foundations and Developments* Routledge
Accessed on: 28 Mar 2023

<https://test.routledgehandbooks.com/doi/10.4324/9781315703879.ch8>

PLEASE SCROLL DOWN FOR DOCUMENT

Full terms and conditions of use: <https://test.routledgehandbooks.com/legal-notices/terms>

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

GROUP SELECTION AND BEHAVIORAL ECONOMICS

ALEXANDER J. FIELD

Economists reading Stephen Pinker's 2002 book, *How the Mind Works*, might be forgiven some surprise at finding their discipline lumped together with sociology and anthropology and accused of a common error. Pinker's assault on the tendency to treat the human mind as a blank slate at birth extended the attack on the standard social science model begun by John Tooby and Leda Cosmides (1992). To these scholars, social science is largely undifferentiated by discipline—a shock to economists, who have traditionally seen themselves standing to one side of a divide separating those approaching human behavior through the study of rational choice and others emphasizing the influence of such variables as culture and social structure. To be sure, institutionalists, economic historians, and development economists have straddled these two camps uneasily, but many rational choice theorists will find it astonishing to be accused of the same error as traditional sociologists and anthropologists. Surely something must be wrong here.

And something is wrong in the sense that Pinker and others, in their efforts to blunt the continuing influence of behaviorist psychology, have papered over important differences within the social sciences. One major difference is adherence to methodological individualism, a principle embraced by rational choice theorists and rejected by traditional sociologists and anthropologists. Another distinction missed is that although the latter disciplines have been generally hostile to Darwinian explanations of human behavior that emphasize the legacy of natural selection on human biology and action inclinations, economists have been more open to them.

The downside of this openness, however, has been that with a few important exceptions, economists' understanding of Darwinian evolutionary theory has tended to be superficial. Such theory is frequently adduced as support for the assumption of a narrow version of human selfishness: preferences or goals are individualistic in the sense that they are assumed to guide action aimed at efficiently advancing the material self-interest of the actor. This assumption has a long pedigree and is often treated as axiomatic. Over a century ago, Francis Edgeworth wrote that "the first principle of economics is that every agent is activated only by self interest" (1881, 16). The pursuit of self-interest *might* mean that since I am happier when you consume more, acting to benefit you would be selfish. As an empirical matter, however, that's not what it typically *does* mean, and the hardheaded realism of the behavioral assumptions often underlying "economic reasoning" have been part of what has attracted many students to the field.

What is the basis, however, for this first principle? When pressed for justification, the heavy guns of what is presumed to be Darwinism are often wheeled in for support: of course we're selfish, because our parents, our grandparents, and our great-grandparents were, and if they hadn't

been, we wouldn't be here. Why? Had they been less selfish, the argument goes, they would have been at an evolutionary disadvantage and thus less likely to procreate.

This style of argument is one reason traditional sociologists and anthropologists, who tend to doubt that selfishness prevails in all realms of human action, want little to do with Darwinism, rational choice models, sociobiology, evolutionary psychology, or any type of thinking that even hints that genetics or biology might influence human behavior or cognition.

This chapter addresses controversial issues within the social and natural sciences. These controversies, which have implications for both major social scientific traditions, need to be considered if we are to make progress in understanding human behavior. I defend a version of methodological individualism against the emphasis on emergent properties of traditional sociology/anthropology, but at the same time I question the conclusions that many rational choice theorists have drawn about the implications of Darwinian thinking for human action inclinations.

ASSUMPTION OF SELFISHNESS

First and foremost, it is necessary to be clear about what rational choice theory does and does not claim. It is useful to think of such models as operating at one of three levels.¹ The first, level 1, simply asserts that people act in satisfaction of their desires. Analysis based on this approach may have a place in casual conversation but can have no claim to scientific standing, since there is no set of observations on human behavior that could contradict it.

Level 2 models, which represent the middle ground of economic theorizing, assume that preferences are stable and transitive. Without some stability over at least the short term, it is hard to conceive of humans engaging in sustained goal-oriented activity. Both observation and intuition suggest that we do so engage. Transitivity is justified by appeal to an evolutionary/money-pump argument. If you did in fact prefer A to B, B to C, and C to A, then, having given you C, I could get you to pay me successively to swap B for C, A for B, and finally C for A—thus inducing you to pay me three times for what I had originally given you for free. Presumably, it is argued, individuals with such preference structures would have been disfavored over time by evolutionary forces.

At least with respect to decisions of particular individuals, level 2 models are falsifiable. By presenting an individual with a sequence of binary choices, it is clearly possible for a person to choose B over C, A over B, and then C over A. If such a violation does occur, we must conclude that preferences lack either stability or transitivity. Level 2 models, however, do not require that preferences be self-interested in the sense in which Edgeworth intended it.

Level 3 models place the most restrictions on preferences. They assume, in addition to stability and transitivity, that the long history of natural selection has rendered humans, like other animals, behaviorally predisposed toward a particular version of selfishness: inclined toward actions that efficiently advance the material self-interest of the organism undertaking them.

Obviously, in many respects and many spheres of action, the assumption reflected in level 3 analysis—and the historical evolutionary explanation of why it is valid—have merit. Organisms that did not seek water when thirsty, food when hungry, or warmth when cold would have been more likely to perish before procreating, leaving those more inclined to such responses to contribute genes to future generations. In some circumstances, however, this assumption may imply behaviors different from those to which an organism predisposed to act in ways that fostered its inclusive reproductive fitness would be inclined.

In polite conversation, economists pursuing level 3 analysis will sometimes grant that “cultural influences” have here and there restrained such selfishness. Allowing restraint on selfishness to be explained in this manner represents an accommodation to the divide that the sociologist

Talcott Parsons (1937) and others tried to engineer. Parsons wished rational choice theory to focus on the choice of appropriate means to achieve given ends; sociology was to be about how human goals were socially determined or constructed. In less accommodationist moments, however, economists unambiguously reject the traditional sociological/anthropological approach, with its suggestion that culture and norms occupy a supra-individual realm, to be understood independently of the cognitions and behaviors of individuals whose activity they influence. In doing so, economists effectively reject as unscientific the concepts of culture and norms as used by sociologists and anthropologists because they violate the principles of methodological individualism (Field 1984).

This expression of disrespect for core disciplinary concepts is returned with accusations that economists are guilty of the sin of reductionism. Communication breaks down, and we are left pretty much where we have been for the last half century, if not longer. Traditional sociologists and anthropologists jealously guard such key concepts as culture and social norms, while the barbarian rational choice theorists take potshots from behind protected redoubts.

Culture is in fact a significant influence on human behavior, although critics have been right to insist that it not be treated as a superorganic force. Rational choice theorists, however, while rejecting superorganic interpretations of culture (and sometimes, wrongly, the importance of the concept itself), have nevertheless struggled to establish convincing explanations of the origins of human sociality from within a perspective based entirely on self-interested action.

There is in fact a large empirical literature documenting experimental and observational evidence at variance with this assumption. Whereas much of economists' interest in anomalies has focused on risky choices or the ways in which beliefs about the world are formed, my interest here is in simple choice situations where, from a wealth maximization standpoint, humans systematically choose less over more. I refer specifically, for example, to results in the ultimatum game, the trust game, the one-shot prisoner's dilemma, or voluntary contributions to public goods experiments (Kagel and Roth 1995; Fehr and Gächter 2000; Field 2001).

Sociologists and anthropologists have, as noted, tended to explain such behavior by appealing to norms, culture, or social structure, concepts that rational choice theorists have often viewed with suspicion. In considering culture, we need to reject both the idea that it is superorganic and the view that society is just the individual writ large. If we accept the merits of trying to understand culture—both its functioning and its origin—using a sophisticated version of methodological individualism, however, there remains the challenge of explaining human sociality. If we are to locate its origins in the characteristics of individuals, what should we assume about them? More specifically, if our social proclivities are ultimately rooted in a genetically mediated human nature, what are its characteristics? Economists have often insisted on trying to solve the problem using level 3 models. We need to understand why such models aren't up to the job.

WHY CAN'T LEVEL 3 MODELS WORK?

The basic problem of social organization has commonly and appropriately been modeled as a prisoner's dilemma. If we make level 3 behavioral assumptions, what we know from game-theoretic analysis can be easily summarized by considering the likelihood of nondefection outcomes in the indefinitely repeated game, the fixed- and known-duration (finitely repeated) game, and the one-shot game.

In the indefinitely repeated game, it is possible for mutual nondefection outcomes to emerge, although outcomes in which everyone defects are just as plausible: there are multiple equilibria, and we can't easily specify why one rather than another is selected. Once established, nondefection

outcomes can be self-policing—sustained by self-interest alone—but there is no guarantee they will emerge in the first place. Indefinite repetition thus may be necessary but is not sufficient to guarantee a nondefection rational choice equilibrium outcome (Fudenberg and Maskin 1986).

In the fixed- and known-duration game, defection right from the start is the only rationalizable course for either party unless one believes that one's counterpart might be irrational.² Thus, using level 3 behavioral assumptions (wealth maximization), the only way one can rationalize a strategy other than continuous defection is by positing the absence of rational choice in one's counterpart—hardly a satisfying foundation on which to build a rational choice explanation of cooperation. In the finitely repeated two-player game, continuous defection by both parties is the only Nash equilibrium, although, in contrast with the one-shot case, the strategy is not strictly dominant.

In the one-shot game, defection is the only rationalizable strategy. Irrespective of the play of the counterpart, choosing rationally requires that one defect, because defection is strictly dominant. There is simply no way that cooperation in a one-shot prisoner's dilemma game can be philosophically justified, assuming the payoffs accurately reflect the benefits to the individuals associated with the different strategy profiles.

Within the context of a rational choice–game-theoretic framework, individuals must somehow transition from one-shot to indefinitely repeated games in order to have a chance at moving from an asocial state to social interaction. Yet it is clear that assumptions reflected in level 3 models would guarantee that this could not happen. To accommodate evidence of human sociality, then, we must either assume that the state of nature has “always” been one of indefinitely repeated interaction, and then try to explain why sociality equilibria have been selected (Binmore 1994), or modify the behavioral assumptions in level 3 models. In trying to explain sociality, rational choice theorists have been stuck between assuming indefinite repetition at the outset, which is unrealistic, and which in any event gets us only part of the way toward the intended objective because of the equilibrium selection problem, and an aversion to modifying behavioral assumptions because of the evolutionary argument that any behavior that systematically disadvantaged the organism practicing it would have been disfavored by natural selection.

Many people—particularly those wedded strongly to level 3 rational choice models—seem to feel that it is somehow unfair to raise the issue of origin, that it is akin to insisting that we explain what happened before the Big Bang. But the issues are quite different. Based on our knowledge of the fossil record and the biology of surviving organisms, we can be fairly certain that there was a time when there were no social animals, just as there was a time when there were no multicellular organisms, and before that a time when there were no nucleated cells. Real historical transitions occurred, and exploring the mechanisms of natural selection that enabled them, as well as the legacy of the most recent transitional experience on human predispositions, are legitimate endeavors.

The most plausible explanation of human sociality is that as a legacy of that transition, humans possess a genetically mediated behavioral inclination that predisposes us to “solve” one-shot prisoner's dilemmas, even when such action is at variance with that consistent with level 3 behavioral assumptions. This solution module, or primary sociality algorithm, enables pairings of humans some of the time to achieve a cooperate-cooperate profile at initial encounter. Such achievements, in turn, allow the interaction to go into extra innings.

Once continued interaction has been established, a set of secondary sociality algorithms inclines us to spend inordinate time scrutinizing third parties, discriminating among those who do and do not violate group norms (Cosmides and Tooby 1992), and engaging in costly punishment of those who do (Fehr and Gächter 2000). In order to understand how any of these inclinations might possibly have been favored by the operation of natural selection, we need to have a more sophisticated understanding of evolutionary theory.

THE GENE'S-EYE VIEW AND GROUP SELECTION

The preponderance of biological opinion today is that the ultimate locus of selection pressure is at the level of genes, not the cells, organs, or organisms that may contain them. To be sure, in many instances there is a mutuality of interest between the gene and the organism: if the organism does not survive, procreation may not take place, and thus in many cases genetically mediated behavioral predispositions that threaten the survival of the organism will be disfavored. Thus, again, it is no accident that organisms inclined to drink when thirsty are more likely to survive and procreate than those for some reason lacking this inclination.

But to say that this logic holds in most instances does not necessarily mean that it applies in all. Return to the economist's retort to those who might question level 3 behavioral assumptions. It is also true that if our great-grandparents hadn't altruistically sacrificed their material welfare to raise our grandparents, and our grandparents to raise our parents, we wouldn't be here either. Well, yes, comes the answer, but what we meant by selfishness was selfishness toward nonkin, not necessarily within the family. This allowance, however, grants a pretty important exception to the principle that natural selection makes us universally selfish. Let us further consider its logic.

Most economists accept the evidence (including that of our own behavior) that humans, along with many other animals, have a tendency altruistically to sacrifice their material welfare for their offspring. Since the late William Hamilton's work (1964) we have understood that genes so predisposing might be favored in spite of their threats to organism survival because children share half the genetic endowment of each parent.³ If parental sacrifice, even sacrifice leading to premature death, enables several children to survive and procreate, then genes predisposing to such sacrifice might be favored. The logic and explanatory power of the inclusive-fitness calculus pioneered by Hamilton is now widely accepted by biologists and economists who have considered it.

Since humans may prefer their children's welfare to their own, we have established here a clear violation of the universal purview of level 3 behavioral assumptions, a violation based on a biologically and genetically mediated altruistic tendency that trumps the inclinations predisposing toward wealth maximization. What is happening here is that genes are "deviously" predisposing organisms (us) to self-sacrifice, because these behaviors have in the past resulted in an increase in the frequency of genes so predisposing. This example is an indication that the straitjacket natural selection places on assumptions about human behavioral inclinations is looser than a casual reading of Darwinism might suggest.

While the existence and explanation of altruistic tendencies toward kin is generally accepted, the question remains whether we might be biologically inclined, albeit at weaker levels, toward altruistic behavior toward nonkin. Is this possibility consistent with known mechanisms of natural selection? The straightforward and perhaps surprising answer is yes. As in the case of altruism toward kin, such a circumstance would require "devious" genes to incline an organism toward action potentially detrimental to the organism's material welfare, but nevertheless resulting in increasing frequencies of such genes in future generations. The rise from low to higher frequency of such genes would require a specific set of demographic conditions enabling the operation of selection above the level of the individual organism. This has historically been known as "group selection," although the term "multilevel selection" is often preferred today, as a means of emphasizing that current analysis is in no way inconsistent with a gene's-eye perspective.

It is important to understand both the logic and some of the intellectual history of this analysis in order to appreciate its implications for behavioral science. We begin by noting that the debate about the role of group selection in accounting for human sociality is one instance of a

set of closely related debates surrounding several key evolutionary transitions. These include the origins of chromosomes, the transitions from prokaryotes to eukaryotes (cellular organisms) and subsequently to multicellular organisms, and the origin of sexual reproduction (Maynard Smith and Szathmáry 1995; Keller 1999). In each instance, appeal to how selection at a higher level overcomes selection pressures at a lower level is necessary in order to provide a coherent account of transition.

For example, with respect to sexual reproduction, parthenogenetically reproducing organisms would appear to have a twofold genetic advantage over those engaged in sexual reproduction, since the parent in the former case transmits 100 percent rather than 50 percent of genetic endowment to offspring. On the other hand, it is commonly accepted today that a system of sexual reproduction gives a species an advantage based on access to a larger pool of genetic variability in the face of rapidly evolving pathogens. But that advantage accrues to the group and could not have favored the origin of sexual reproduction if selection occurred no higher than the organism level.

The flashpoint has come when scholars both begin and end an explanation of a characteristic with reference to its function. Modern evolutionary theory insists that we specify how this equilibrium originated. Given that the ultimate locus of selection is at the level of the gene, not the organism, let alone the species, this means one must explore the operation of natural selection at multiple levels. It is uncontroversial that selection can take place at the gene level, the organism level, or levels intermediate. Virtually all biologists also admit the theoretical possibility of selection above the level of the organism. Because of a series of contentious intellectual debates that began four decades ago, however, the issue of the empirical importance of selection at this level remains the third rail of biological discourse. Fifteen years ago, even broaching the subject practically guaranteed electrocution. Today, the topic is again sometimes admitted into polite conversation, although many biologists engage in semantic contortions to avoid calling higher-level selection group selection.

Appeals to higher-level selection were common in the 1940s and 1950s among biologists such as Sewall Wright (1945), ethologists such as Konrad Lorenz (1966), and some social scientists who analogized social structures to organs of a body and “explained” them with reference to their function. But the empirical importance of group-level selection came under withering attack in the 1960s, in particular by George Williams (1966). In light of these attacks, the dominant position among biologists from the late 1960s on was that these processes were empirically unimportant in the evolution of humans and other species. Those who continued to consider or investigate its operation were marginalized, consigned in the disciplinary pecking order to a position slightly ahead of creationists (Wilson and Sober 1994, Sober and Wilson 1998).

The consensus position was based in part on the argument that a number of animal behaviors previously “explained” on the grounds that they benefited (were adaptive for) a group could now be shown to be favored by the forces of individual-level selection alone and were therefore not “altruistic” at all (Williams 1966). If behavioral predispositions benefited a group of nonkin, it was because they benefited these individuals *and* the actor (in other words, the behavior was mutualistic, not altruistic; to be favored by organism-level natural selection, the relative fitness of the actor still had to be improved). The apparent coup de grâce in the argument against group selection was the more systematic demonstration that traits that provided benefit to a group, provided they were initially widespread, could be sustained (protected from invasion by other strategies) by individual-level frequency-dependent selection alone (Maynard Smith and Price 1973).

For several reasons, group selection is now making a comeback in biology. First, there remains the set of critical evolutionary transitions, such as to sexual reproduction or to sociality, that are difficult to account for without allowing for the operation of selection at higher levels.

Second, models of group selection have moved away from the “islands” formulation associated with Sewall Wright to the structured deme approach pioneered by D.S. Wilson (a deme is a local community of potentially interbreeding individuals). Third, virtually all players today accept the necessity of a gene-centric approach. Finally, it is increasingly recognized that evolutionary game-theory models and the concept of an evolutionarily stable strategy (ESS) are tools for understanding and explaining how evolutionary equilibria are sustained, not necessarily how they originate. Third rail or not, a number of scholars both in and outside biology are saying that the rejection of the empirical likelihood of higher-level selection needs to be reexamined.

Sewall Wright recognized the fitness disadvantage experienced by altruists within each group. But he suggested that if groups were small enough and isolated enough, altruism might by chance—that is, through the mechanism of genetic drift—evolve to fixation within some groups, and that such groups could then outcompete others by persisting longer and colonizing new territories by contributing more dispersers. The conditions necessary for these “island” models to operate continue to be viewed by evolutionary biologists as unlikely to obtain. A major concern is that if there is enough dispersion to permit the colonization of new territories, there is also likely enough to permit dispersing nonaltruists from other groups, or islands, to invade existing altruist groups (Maynard Smith 1964; Wade 1978).

Because of these concerns, greater interest has focused in recent years on haystack or structured deme models, which depend on organisms separating or being separated into groups for part of a life cycle, or perhaps for a period of several generations, and then reentering a global population before again reassorting into groups. Genetic recombination, outcrossing, and infrequent mutation are sufficient to produce variability in individuals.⁴ If the groups are small enough, average group predispositions toward altruistic behavior will vary, even where initial assortment into groups is done randomly (Cavalli-Sforza, Menozzi, and Piazza 1994, 13; Field 2003). Within each group, altruists will lose out in the competition for resources, and in particular the competition to pass on their genes to the next generation. Consequently, their share of each group will fall (or at best stay even in the unlikely event altruists completely dominate a group) throughout the period of time that the group has a distinct existence. But because the behavior of altruists differentially benefits groups in which their frequencies may be relatively higher, the proportion of altruists in the global population may rise in cases where the forces of group selection are stronger than the forces of individual selection.

Examples of circumstances favorable to group-level selection include the life cycle relationships between parasites and other disease vectors and their hosts. Suppose groups of a polymorphic virus with more and less virulent versions invade a number of host organisms, let us say rabbits. Among the rabbits, chance variation will result in varying mixtures of the two forms within each of them. Within each host, the more virulent versions enjoy a fitness advantage and increases in frequency. But organisms infected with high frequencies of the more virulent strain die quickly, before they can infect many other hosts. In contrast, hosts infected with less virulent mixtures live longer, more successfully spreading the disease. The longer-living hosts therefore exercise a greater weight in determining the frequencies of the two versions in subsequent time periods, and each of their “votes” carries a higher proportion of the less virulent strain. Paradoxically, the less virulent strain can be decreasing in frequency within every host, yet increasing over time within the global population.

In fact, such a scenario has been documented among rabbits in Australia in the early 1950s who were confronted by the myxoma virus, introduced by the government to control an exploding population. Initially mortality was very high, but gradually it declined, suggesting that the rabbits were acquiring resistance. Subsequent testing revealed two outcomes, one of which was

surprising. First, the rabbits had indeed, on average, become more resistant to myxoma, which would be expected from individual-level selection operating through differential mortality within the rabbit population. The surprising and unexpected conclusion applied to the viral population. The virus itself became on average less virulent, as measured by extracting blood from rabbits in the wild and comparing samples so obtained with original viral samples stored in a laboratory. This outcome would not be expected if natural selection operated only at the level of the individual virus, as it indeed was within each host. The initial interpretation of this episode as reflecting the operation of group selection was controversial, but most evolutionary biologists, including George Williams, now accept it (Williams and Nesse 1991, 8; Williams 1992). The mechanism is essential to the burgeoning field of Darwinian medicine.

Wright's analysis of group selection relied on the evolution to fixation of altruistic traits within some groups, combined with interdemic competition. The structured deme approach does not require evolution to fixation in any group but emphasizes the necessity of periodic recombining of or migration between groups—called trait groups by D.S. Wilson—in order for group selection to occur. Altruists from the faster-growing more altruistic groups must periodically disperse throughout the global population (see Wade 1978 for discussion).

The possibility that altruists may be declining in frequency within every group and yet rising in frequency in the global population—an apparent contradiction—can be more easily understood with reference to the Simpson paradox (Simpson 1951). The paradox exists when a population divided into groups exhibits a population average that differs from the average of the group averages. A corollary is that under these conditions, changes in the average of the groups may differ from changes in the population average. A compelling example of this was an investigation at the University of California at Berkeley of alleged discrimination against women in admissions to graduate study in the 1970s. Aggregate data showed that admission rates for women were lower than for men. But when administrators looked at the data department by department they found no evidence of discrimination: in each department women were being admitted at approximately the same rate as men.

The explanation for the paradox was that the distribution of applicants by department was not the same for the two sexes. In particular, women were applying disproportionately to departments in which it was more difficult for an applicant of either sex to gain admission. This covariance meant that the averages of the departmental admission rates, roughly equivalent for men and women, differed from the global average admission rates, in which women did less well than men. What was true at the level of each individual group (department)—roughly equal admissions rates—was not true for the entire population of applicants (Dawes 1988, 297).⁵ Similarly, in considering the fate of altruists, what is true for each individual group (altruists are losing out) may not be true for the global population, because groups in which altruists are differentially concentrated may grow more rapidly. The argument at first glance seems akin to the story of the retailer who lost money on every sale but made it up on volume. Upon careful examination, however, the Simpson paradox is not based on such an inherently contradictory claim.

THE ROLE OF GROUP SELECTION IN FOSTERING ALTRUISM TOWARD KIN AND NONKIN

Whereas the Hamilton inclusive-fitness logic as applied to altruism toward kin is broadly accepted, its application to the case of weaker altruistic predispositions toward nonkin, including restraint on first strike, remains controversial. Like Hamilton, recent proponents of group selection, such as D.S. Wilson, interpret the kin selection mechanism as an instance of multilevel

selection, where the group is the family unit. They then proceed as if the phylogenies of altruism toward kin and altruism toward nonkin are similar in all fundamental respects. This, I argue, is a mistake because the biologically altruistic character of sacrifice for kin is largely independent of whether others in a group express similar behaviors. This is not necessarily true of the sociality algorithms, particularly the primary sociality algorithm, because of its role enabling reciprocity, upon which mutually beneficial exchange relations are based.

Certain behaviors expressed toward nonkin are clearly biologically altruistic (and, correspondingly, not rational for narrowly self-interested agents) upon first appearance. Lone cooperators in a sea of defectors harm themselves while benefiting others. If the result of a first encounter is death, it makes no difference what an organism may have been prepared to do should a continuing series of interactions have been initiated. All that matters is that she cooperated, to her detriment and to the advantage of another, on what turned out to be a one-shot game.

The fitness disadvantage of parental sacrifice for children is largely independent of whether such propensities are widely shared in a group. That is not so true of algorithms favoring cooperation among nonkin. The higher the frequency of nonkin cooperators in a population, the smaller the fitness disadvantage of cooperating (assume players are randomly paired to play one-shot prisoner's dilemma games). It is also, however, true that the fitness advantage of the remaining defectors rises with the frequency of cooperators—because a defector is now almost sure to be matched with a cooperator. That is why, if selection is only at the organism level and players play one-shot prisoner's dilemma games, the replicator dynamic will inexorably drive cooperators to extinction.

There is one condition under which cooperative strategies could be sustained by individual self-interest alone, and thus one condition under which populations not subject to selection above the level of the individual could be proof against direct invasion by defectors.⁶ That condition is that players, instead of playing one-shot games with their counterpart, play a sequence of prisoner's dilemma games in which neither player knows in advance when the interaction will end. Under these conditions, it is possible to sustain contingently cooperative behavior as a self-policing rational choice equilibrium. If both parties were playing tit for tat, for example, with the right discount rate neither would have an incentive to deviate from her strategy, since each would be playing a best response to that of the other.

If one can assume indefinite interaction, then, one can tell a story in which group selection is necessary for contingently cooperative strategies to rise from low to higher frequencies, but in which the need for it disappears when such strategies have attained sufficiently high frequency. Indeed, a good deal of evolutionary game theory has been devoted to telling the second half of this story, neglecting, however, the fact that explanations of equilibrium maintenance are not necessarily the same as explanations of origin.

There is, however, an important weakness in this argument as it applies to the real world. In bilateral pairings, *either party typically has the option of terminating the interaction at any time*. Cohabitants can cease living together, long-term supplier-customer relationships can end, veteran employees can be fired, and military alliances can be and have been abrogated overnight. The assumption that the duration of the game is exogenous, while a convenient modeling convention, is as a practical matter unrealistic in the absence of third-party “umpires.” If both players are strictly self-interested and know they can terminate at any time without third-party retribution, they will defect from the get-go, and there will be no continuing interaction.

Experimental results show that the human inclination to play cooperate in a one-shot prisoner's dilemma game with anonymity is in fact substantially greater than the Nash prediction of zero. This means that two players, selected at random, may, to their mutual benefit, be able to initiate and sustain a string of cooperative interactions. But such cooperation is fragile because the pro-

ensity to play cooperate is nowhere near 100 percent. Foraging algorithms reinforce a persisting temptation to play what remains, after all, a strictly dominant strategy.

Since freely chosen interactions cannot realistically be assumed to persist for indefinite durations beyond the control of the parties (even if this were so, this would only make self-interested cooperation possible, not guarantee it), the restraint of defection in a group requires more than the threat that a counterpart will match defection with defection. It requires a set of evolved secondary sociability algorithms that govern behaviors toward third parties. At least two of these have been well documented. First, humans devote remarkable amounts of energy to seeking out and identifying rule violators (defectors)—our obsession with this is so strong that it can distort our ability to think logically (Cosmides and Tooby 1992). Second, a high fraction of humans are prepared to engage in costly punishment of those who defect (Fehr and Gächter 2000). When present at sufficiently high frequency, these two algorithms serve to weaken the attractiveness of defecting in a prisoner's dilemma game. They may weaken it so much that even those so inclined find that the behavior is no longer a dominant strategy. In the limit, because of the threat of widespread third-party punishment, which reduces the net payoff to the defector in a defect/cooperate profile, the prisoner's dilemma may, in fact, cease to be a prisoner's dilemma. That is why, when well established, reciprocal relations among nonkin (the precursor for exchange) appear to be mutually beneficial, and thus to have lost the altruistic character that the same contingently cooperative strategy, expressed within a sea of defectors, surely would have possessed.

Unfortunately, to the degree that the primary prisoner's dilemma may be approximately resolved in this fashion, a set of secondary dilemmas that are just as serious are created. Foraging algorithms, optimized for wealth maximization for the individual organism, counsel against being influenced by either of these secondary algorithms just as much as they counsel against an initial cooperative move in a bilateral pairing. The secondary sociality algorithms also involve action that benefits other group members at cost to self. Consequently they could have been reinforced by natural selection only if such selection occurred above the level of the individual organism.

Note that positing an inclination to punish those who fail to punish doesn't resolve the problem. It is always a strictly dominant strategy to let others punish rule violators, or to let others punish those who fail to punish, or to let others punish those who fail to punish those who fail to punish, and so on. As fast as one prisoner's dilemma is at least hypothetically resolved, another is created. Different societies have used different mechanisms to control defection, but none has eliminated the basic problem posed by the prisoner's dilemma, although some displace it to different levels.

So far I have not talked much about culture and socialization. The evidence is overwhelming that humans can be trained not only to master technologies (means) but also, within limits, to adopt certain ends. We are smart, but we also are apparently biologically prepared to accept indoctrination in group-reinforcing ideologies, particularly when we are young. With remarkable success parents are able to impart to their children not only their own language vocabularies and recipes for making cake but also their own religious or ethical beliefs and practices. This can be done for nonkin as well. It is an essential feature of military training that infantrymen are taught (learn) to throw themselves on a grenade to save other members of their squad. Organized political groups in the Middle East have been quite successful in recruiting and motivating suicide bombers.

The last two examples are particularly stark illustrations of the fact that humans can be taught (or learn through instruction or imitation) to express behaviors that from a level 3 game-theoretic standpoint are strictly dominated strategies. For many in the sociological/anthropological tradition, an appeal to culture or norms is all that is needed to explain human sociality. Cultural explanations, presume, however, that we have this capability to learn or be taught ends as well as

means. But we must then ask how a capability to “learn” strictly dominated strategies ever could have evolved if selection occurred at no higher level than that of the individual.

The one-shot prisoner’s dilemma is a source of great frustration for rational choice theorists. The cooperative profile, which can be attained only if each player plays a strictly dominated strategy, is clearly the most jointly attractive of the three efficient outcomes. The Nash equilibrium, on the other hand, is the only inefficient outcome. Economists often argue that the pursuit of individual self-interest also leads to social benefits. What economists are trying to say here is that behavior that looks as if it is selfish is really mutualistic from a biological standpoint.

But whereas voluntary exchange in an atmosphere where force and fraud are at low levels can be mutually beneficial, such exchange can persist only as the consequence of widespread secondary sociality algorithms whose operation continues to involve the play of strictly dominated strategies. So taking this more encompassing perspective, it is simply wrong to say that we can build a theory of social order based on level 3 rational choice models.

In a market economy, the sociability algorithms help establish a human environment in which individuals are able to pursue the counsel of their foraging algorithms in a context never anticipated by the forces of natural selection that honed them. The genius of markets is that in the limit they appear to turn games with people into games against nature, for which the foraging algorithms were originally selected. That limit is the purely competitive market where prices confront individuals as parametric.

But markets don’t abolish the problem of the prisoner’s dilemma. We may intone that crime doesn’t pay. With sufficient third-party enforcement the statement may, at one level, become true. But at some other level it always pays. The appeal to the foraging algorithms to dissuade us from criminal behavior is ultimately destined to fail. That is why this claim is always paired with the nonconsequentialist argument that criminal acts are wrong. The fact that many humans are receptive to such arguments is, again, reflective of the fact that we can be trained—indoctrinated—to express in our behavior group-beneficial but strictly dominated strategies (to some degree these inclinations are hardwired at birth). Without being able to ensure indefinitely repeated interaction, one simply cannot have a self-policing rational choice equilibrium involving cooperation, even with frequency-dependent selection, and that is why the foraging algorithms alone are insufficient as a foundation for understanding political behavior.

Still, it is important to note that the effect of frequency-dependent selection can be a powerful influence on the relative fitness disadvantages of sociability algorithms in a way that is not true for sacrifice toward kin. The biologically altruistic character of sacrifice for kin is largely frequency-independent. Traits necessary for the development of sociality among nonkin have the property that as they rise in frequency within a population, the within-group selection becomes over time less negative. High frequencies of contingently cooperative behavior are preconditions for reciprocal relations, which enable mutually beneficial exchange.

Altruism toward kin is different in its phylogeny and in its game-theoretic character. Sacrifices for our children do not represent a solution to a prisoner’s dilemma.⁷ It is not really possible—even stretching—to interpret them as an instance of exchange over time (this is the idea behind the misnamed theory of reciprocal altruism; see Trivers 1971). The interaction of a parent with a helpless child lacks the strategic character of an encounter between two armed men, because the material payoff to the parent is unaffected by any action the baby takes. And finally, the genetic payoff to sacrificing for babies is largely independent of how widespread such inclinations may be among other parents in a group.

Whereas the kin selection mechanism has been relatively uncontroversial among evolutionary biologists,⁸ the same cannot be said for group selection explanations of altruistic behavior toward

nonkin. The initial Maynard Smith and Price work (1973) investigating restraints on intraspecific harm, and subsequent development of the concept of an evolutionarily stable strategy, explored the implications of frequency-dependent selection in equilibrium maintenance. Their great insight was to understand and explain how strategies that might benefit a group might be sustained in the absence of any selection at the group level. Initially, a number of theorists—perhaps even Maynard Smith and Price themselves—believed that this conclusion provided the explanation for the persistence *and* the emergence of altruistic behavior as the consequence of individual selection forces alone.⁹ But this is an error that arises from confusing the explanation of forces that may sustain an equilibrium with those responsible for the evolutionary trajectory that led to it.

In human interactions behavior is influenced both by foraging algorithms and by the primary and secondary sociability algorithms augmented by cultural reinforcers. Whereas the foraging algorithms were always reinforced by selection at the level of the organism, the sociability algorithms required the operation of selection above that level. The development of predispositions to sacrifice for kin, particularly offspring, could not have been favored by natural selection operating no higher than the level of the organism. The same is true of the algorithms underlying cooperation among nonkin: the propensity to play cooperate in what might be a one-shot prisoner's dilemma, the inclination to search out rule violators, and the predisposition to engage in costly punishment of them.

Readers will observe the close analogues between rational choice game-theoretic analysis and evolutionary game theory, and the related roles played by indefinitely repeated interaction in the former and frequency-dependent selection in the latter. The assumption of indefinitely repeated interaction is necessary to provide a logically consistent account of how selfish actors might sustain a self-policing nondefection outcome. The assumption of frequency-dependent selection is necessary to explain how prosocial behavioral predispositions might be sustained in the absence of selection above the level of the organism. But in both cases there is a historical question left unanswered. In game theory, how could selfish actors ever have transitioned from one-shot to indefinitely repeated interaction? The obviously unsatisfactory answer is that they never had to (and never have to) because interaction is always indefinitely repeated.

In evolutionary game theory the challenge is to explain how prosocial behavioral predispositions ever could have transitioned from low to higher frequency without the operation of selection above the level of the organism. The equally unsatisfactory answer is that these traits are and have always been at sufficiently high frequency that we need not concern ourselves with problems of origin.

ECONOMISTS AND THE INTELLECTUAL HISTORY OF THE GROUP SELECTION DEBATE

The debate about the empirical likelihood of higher-level selection in human evolution and its likely legacy on our behavioral predispositions has profound implications for the conduct of social science. These are controversies that will benefit from interdisciplinary dialogue, not only within the social sciences but also between the social and natural sciences. For historical reasons, economists may be particularly well situated to contribute to this discourse. First, in comparison with sociologists or anthropologists, economists have been more favorably disposed toward considering biological/genetic explanations of human behavior and cognition. Second, to date only a very few have actually considered the mechanics and implications of the operation of selection pressures at different levels. Although this has resulted in superficial defenses of level 3 modeling as having an evolutionary justification, it may now be an advantage because, in contrast with

biology, economics lacks a legacy of intellectual warfare in this area. It may be easier, consequently, to discuss the issue on its merits.

Although the number of economists aware of the debate has been small, those who have examined the issue have generally concluded that selection at higher levels must have been empirically important at some point in our history. A full-text JSTOR search for “group selection” among articles published in economics journals between 1890 and 1996 turns up just five hits that use the term in its evolutionary sense.¹⁰ One is for an article by the biologist Edward O. Wilson in the *American Economic Review*. The other four hits included one for Paul Samuelson, one for Gary Becker, and two for the late Jack Hirshleifer. Each of these authors has written positively about the likelihood that group-level selection has operated at some time in our evolutionary past (Becker 1976, 284, 294; Hirshleifer 1977, 25; 1982, 30–33; Samuelson 1993).

To this list, which already includes two Nobelists, one should add a third. Although Friedrich Hayek’s argument in *The Fatal Conceit* does not depend on the assumption that biological group selection played a role in human evolutionary history, he was clearly receptive to the likelihood that it did (1988, 25). His analysis is notable in turning the standard Hobbesian dilemma on its head.

Since 1996, other economists, including Herbert Gintis (2000) and Theodore Bergstrom (2002), have treated the likelihood of higher-level selection sympathetically in explorations of human sociality. I do not mean to suggest that all economists, particularly those who have more recently become aware of multilevel selection issues, have necessarily embraced the likelihood of higher-level selection. Indeed, a number, taking their lead from what has appeared to be the consensus view in biology, have not. But as a group economists have tended to be open-minded about the issue when presented with the logic and evidence in support of the hypothesis.

A case in point is Robert Frank. His influential 1988 book, *Passions Within Reason*, followed conventional biological wisdom and dismissed group selection as empirically unlikely. He then attempted to proceed within a modeling environment in which higher-level selection was out of bounds, resulting in a number of internal contradictions in his analysis (see Field 2001, ch. 4). Six years later, however, Frank acknowledged the limitations of his initial approach (Frank 1994, 620).

These data and citations are sufficient to establish two points: first, at least until recently, the number of economists who have considered or even been aware of the possibility of multilevel selection has been small. Second, in contrast to biology, the profession does not harbor a large cohort of individuals passionately prejudiced against the empirical likelihood of higher-level selection. Many remain unaware that it is a possibility, although this group has been shrinking.

As this has been happening, outside of economics Robert Boyd and Peter Richerson (see Richerson and Boyd 2004) have articulated an influential position that subtly undermines much of the progress that has been made to date in understanding the logic of group selection and its likely role in explaining some of the otherwise anomalous experimental results in behavioral economics. Richerson and Boyd argue that cultural group selection, reinforced by conformity norms and transmission through other than parent-to-child routes, has been an important influence on human development. At the same time, they maintain that there has never been any empirically important genetic/biological group selection, and thus there are no significant behavioral legacies of its operation. On this issue Richerson and Boyd essentially reendorse the position adopted by Williams in 1966, one that Williams has subsequently backed away from.

Richerson and Boyd are right to argue, contrary to some economists’ prejudices, that cultural variation is an independent influence on behavioral outcomes, one that cannot be entirely swept back to environmental differences. Thus normative structures and behavior may differ in regions that share a similar resource base or access to similar technologies (see also Field 1991). There are, nevertheless, a number of difficulties with their rejection of an empirical

legacy for biological group selection. These include the continuing need to account for prosocial inclinations prior to the development of cultural capabilities in humans, which Richerson and Boyd date from half a million years ago. Their position and its limitations are discussed in greater detail in Field 2005b and 2006.

CONCLUSION

As noted earlier, level 1 rational choice analysis, which assumes that people act according to their desires, can have no claim to scientific standing. Level 3 analysis, on the other hand (preferences are assumed stable, transitive, and strictly individualistic), is at variance with a broad range of experimental and observational evidence. There has sometimes been a reluctance to accept this evidence on the grounds that we must be universally selfish because Darwin told us so. One purpose of this chapter has been to undercut that argument. If we are to use rational choice as a general approach to building models of human behavior, and if these models are to have any chance of real predictive success, the action will have to be at level 2.

It is often said that economics is superior as a discipline to the other social sciences because it has a theory. If so, what are its essential features? If the distinguishing feature of economic theory is the assumption of universal selfishness, in the sense that in all realms we efficiently advance our material welfare, it is not a very good theory. There is abundant evidence that whenever people sacrifice for their children, exercise restraint on first strike in what might be a one-shot prisoner's dilemma, reject a positive offer in an ultimatum game, make voluntary contributions to a public good, engage in altruistic punishment, or enter the voting booth, we leave money on the table, literally or figuratively. Based on such evidence, traditional sociologists and anthropologists, and many economists, object to the assumption of universal selfishness reflected in level 3 models, and rightly so.

Level 2 models have their problems. To be sure, we face cognitive and computational problems in forming accurate beliefs about the world, and sometimes we do make choices that we would not have made if we had had stable and transitive preferences and maximized appropriately.¹¹ But the assumptions that humans are goal-oriented, that these goals don't change every minute, that our preferences are transitive, and that we try to do the best we can for ourselves, given these preferences and the constraints we face, is indeed a reasonable starting point. When spelled out in this way, it is hard to see anyone objecting to it as a first approximation. It is also difficult, however, to see people manning disciplinary ramparts in defense of the proposition that preferences are stable and transitive.

Selfishness, though, is another story entirely. The superiority economists secretly (and sometimes not so secretly) feel with respect to other social scientists is, I suspect, based less on the belief that we understand that preferences are stable and transitive while others don't, and more on the assumption that we understand that people are universally selfish while others don't. The hardheaded (apparent) realism of this assumption also helps account for some of the distaste, ambivalence, and envy those in other disciplines sometimes experience with respect to economics.

Still, there is unease within the profession: a significant number of economists may practice level 3 but will publicly defend level 2. Nevertheless, if we say that the core of economic theory rests on stability and transitivity, but what we mean is stability, transitivity, and selfishness, then we are ultimately forced to abandon what is *truly* distinctive about the rational choice approach: a commitment, to the degree scientifically possible, to explain collective phenomena with respect to the aggregated characteristics of constituent units. Only by firmly relinquishing the assumption

of universal selfishness can behavioral scientists offer a coherent alternative to the traditional sociological and anthropological approaches to issues of human sociality.

As has been noted, the combination of methodological individualism with the assumptions of universal selfishness is a nonstarter in terms of explaining the origins of human sociality. If we insist on universal selfishness, we are forced to introduce culture as a supraindividual constraining force. If we relax that stricture, however, we are in a position to understand institutions, culture, or norms with reference to the aggregated properties of those individuals whose behavior they organize. These properties include behavioral and cognitive predispositions along with learned associations, conditional probabilities updated using Bayesian statistical methodologies, and the memory capabilities that facilitate this learning.

Language, for example, is broadly recognized as an important facilitative social institution. To understand its significance, however, we are not obliged to appeal to emergent properties. Its origin and operation can be understood entirely with reference to the aggregated properties of those individuals whose communication it organizes. Thanks to the work of Noam Chomsky and his followers, it is widely accepted today that we are born with a set of deep structural rules of grammar hardwired in our brain. This hardwiring explains a set of rules of universal grammar that all 5,000-plus known human languages obey. Vocabularies, on the other hand, vary widely: they are a learned cultural phenomenon. They represent shared expectations about how others will associate strings of phonemes with objects or concepts.

The analogy with respect to social organization is not exact. Shared vocabularies represent from a game-theoretic perspective the solution to a coordination problem. A particular language vocabulary can be viewed as a historically determined selection from a set of multiple equilibria. But it is not possible to fit all aspects of culture into this framework, especially those that are near universal. In addition to the oft-noted incest taboo, all known societies proscribe within-group murder and excessive within-group lying or cheating. Whether it has formal institutions of government or not, any group, if it is to survive or even form in the first place, must contain individuals who restrain their own potential use of force and fraud. Once sustained interaction is established, members must also be prepared to engage in costly punishment of others who don't. The challenge of surmounting the prisoner's dilemma in initiating interaction with nonkin, and of sustaining it once established, is universal, and because of its recurring character, the solution to it likely involves some genetic substrate.

Social solutions to coordination games are Nash, whereas social solutions to the one-shot prisoner's dilemma, the problem of voluntary contributions to public goods, or the problem of costly punishment of norm violators are not.¹² From an evolutionary perspective, some of the cognitive and behavioral machinery we bring to the tasks of dealing with other humans could not have evolved from low to higher frequency without the benefit of higher-level selection. Our historical and contemporary ability to "solve" these dilemmas is based on action inclinations sometimes at variance with those counseled by our foraging algorithms (those celebrated in traditional economic theory).

All social scientists, economists included, are susceptible to self-deception, which can involve projecting a level 3 rational choice interpretation on behavior that cannot possibly support it. There is no clearer instance of this than the goal-oriented rationalizations that obviously intelligent, logical individuals give for voting. The odds of influencing a national election through one's voting behavior are infinitesimal, and the act of voting is nearly impossible to defend using consequentialist reasoning (Field, 2005a).

Yet for many the urge to vote *and* the sense that they are rational level 3 maximizers are so powerful that they will simply cover their ears or otherwise indicate they do not wish to hear more

“economist” arguments about the material benefit-cost ratio of the activity. There would be less discomfort were there greater understanding and acceptance of the proposition that in a number of important spheres of action we have hardwired dispositions to behave in ways contrary to our individual material well-being.

If we get “utility” from voting or from cooperating in prisoner’s dilemma games where it is not in our material interest to do so, it is probably because we are biologically programmed to take pleasure in participating in some forms of collective action, an inclination that is apparently sufficiently strong to trump the logic of free riding and get more than 120 million voters to the polls in a national election in the United States. Socialization and education may explain some of the cross-sectional variation in such behavior, but the high participation rates in countries experiencing democratic elections for the first time suggest that this can be at best part of the story. An unavoidable conclusion is that activities such as voting are expressive, not goal-oriented, if goals are understood in the traditional terms of influencing outcomes. It may still make sense philosophically to adopt the principle of choosing as if our vote would decide the election, so long as we don’t give the principle an instrumental rationale.

Similarly, when educated economists and mathematicians play cooperate in a one-shot prisoner’s dilemma or in a fixed- and known-duration prisoner’s dilemma, as did Armen Alchian and John Williams in the world’s very first such experiment in 1950 (they played cooperate with increasing frequency as the game neared its end) (Poundstone 1992, 106–16), something must be trumping wealth maximization algorithms, even though this trumping may mean that the participants end up with more wealth. An appreciation of the richness and subtlety of evolutionary mechanisms and the operation of natural selection, in particular the possibility and likelihood of selection at multiple levels, is a necessary step in allowing us to advance behavioral science informed by a version of methodological individualism while at the same time avoiding the cul-de-sacs into which level 3 rational choice models have repeatedly led.

NOTES

This chapter was presented at the conference The Evolution of Preferences, April 5, 2003, University of California, Davis, at the Society for the Advancement of Behavioral Economics conference at Lake Tahoe, California, on July 29, 2003, at seminars at Brown and Emory universities on October 22 and 29, 2003, and at the Max Planck Institute, Jena, Germany, December 7, 2004. I am grateful to participants in these sessions and to Avinash Dixit, Herbert Gintis, and Bill Sundstrom for comments and discussion on earlier drafts.

1. I ignore here the elaborations necessary for a full treatment of decision making under uncertainty. Indeed, the focus here is on decisions that do *not* involve uncertain or risky choices or problems in forming correct beliefs about the world.

2. For example, if I believe my opponent is playing tit for tat, my best response as a strictly self-interested player is to cooperate at all stages save the last. But the resulting strategy profile is not Nash, since tit for tat is not the best response to the strategy just described.

3. The vast bulk of genetic endowment—over 98 percent—is identical for all humans. It ensures that we develop as *Homo sapiens*, not as a worm or an elephant. A relatively small fraction of our 3.1 billion base pairs may differ among us: they account for the genetic portion of human variation.

4. Genetic recombination occurs during the reproductive process in diploid species. Pairs of chromosomes trade segments before the final random assignment by meiotic division of one or the other of each recombined pair to produce germ cells (sperms or eggs). The latter process is sometimes referred to as outcrossing: the random selection of a single set of chromosomes from the mother (one from each recombined pair) and a single set of chromosomes from the father (one from each recombined pair) to produce a zygote.

5. Another example from the development economics literature: the capital-labor ratio in each sector may be dropping, but for the economy as a whole it may be rising if production is shifting to the more capital-intensive sectors. Thanks to Avinash Dixit for this example.

6. I ignore here the complication that, for example, a population of tit-for-tat players might be first invaded by the behaviorally indistinguishable all cooperate, for example, and would then be vulnerable to invasion by defectors.

7. Within political behavior I mean to include those economic interactions in which strategic considerations are relevant (in other words, the purview of game theory, as opposed to the analysis of purely competitive markets).

8. One might argue that human parents sacrifice for their children in anticipation of care during old age. Perhaps, but it is unlikely that other mammals who sacrifice for their young are making similar calculations. And humans have no rational (level 3) reasons for expecting anything other than default on these implied intergenerational obligations.

9. Its interpretation as an instance of group-level selection has been neither widely advertised (perhaps due to the group selection controversies) nor seriously disputed (it has, after all, been proposed by the originator of the mechanism).

10. Dawkins, for example, wrote, "There is a common misconception that cooperation within a group at a given level of organization must come about through selection between groups. . . . ESS theory provides a more parsimonious alternative" (1980, 360; cited in Sober and Wilson 1998, 79). According to Maynard Smith (1993), it is Dawkins who suffers from the misconception: ESS doesn't provide such an alternative because it addresses the stability of an outcome, not its origin.

11. JSTOR is an online archive for academic journals. A full-text search covers not just titles but also the text of the articles, including citations. Thus some of the hits may refer to articles that are not themselves included in the archive. Currently twenty-five of the most widely read economic journals are included.

12. If monotonic preferences are included within the maintained hypotheses underlying level 2 theorizing, then evidence that people leave money on the table also poses something of a problem for level 2.

REFERENCES

- Becker, Gary. 1976. *The Economic Approach to Human Behavior*. Chicago: University of Chicago Press.
- Bergstrom, Theodore. 2002. "Evolution of Social Behavior: Individual and Group Selection." *Journal of Economic Perspectives* 16: 67–88.
- Binmore, Ken. 1994. *Game Theory and the Social Contract I: Playing Fair*. Cambridge, MA: MIT Press.
- Cavalli-Sforza, L.L., Paolo Menozzi, and Alberto Piazza. 1994. *The History and Geography of Human Genes*. Princeton, NJ: Princeton University Press.
- Cosmides, Leda, and John Tooby. 1992. "Cognitive Adaptations for Social Exchange." In Jerome H. Barkow, Leda Cosmides, and John Tooby, eds., *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, 163–228. New York: Oxford University Press.
- Dawes, Robyn M. 1988. *Rational Choice in an Uncertain World*. San Diego: Harcourt Brace Jovanovich.
- Dawkins, Richard. 1980. "Good Strategy or Evolutionary Stable Strategy?" In G. W. Barlow and J. Silverberg, eds., *Sociobiology: Beyond Nature/Nurture*. Boulder, CO: Westview Press.
- Edgeworth, Francis Y. 1881. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. London: C.K. Paul.
- Fehr, Ernst, and Simon Gächter. 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review* 90: 980–94.
- Field, Alexander J. 1984. "Microeconomics, Norms, and Rationality." *Economic Development and Cultural Change* 32 (July): 683–711.
- . 1991. "Do Legal Systems Matter?" *Explorations in Economic History* 28: 1–35.
- . 2001. *Altruistically Inclined? Evolutionary Theory, the Behavioral Sciences, and the Origins of Reciprocity*. Ann Arbor: University of Michigan Press.
- . 2003. [Comment on Bergstrom, "Evolution of Social Behavior."] *Journal of Economic Perspectives* 17, 2: 209–10.
- . 2005a. "Dictatorship, Democracy, and Quasi-Magical Thinking." Paper presented at the Western Economic Association Annual Meetings, San Francisco, CA, July 5.
- . 2005b. Review article, "Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small Scale Societies." *Quarterly Review of Biology* 80 (December): 453–59.
- . 2006. "Why Multilevel Selection Matters." *Journal of Bioeconomics* 8: forthcoming.

- Frank, Robert. 1988. *Passions Within Reason: The Strategic Use of Emotions*. New York: W.W. Norton.
- . 1994. "Group Selection and "Genuine" Altruism." *Behavioral and Brain Sciences* 17: 620–21.
- Fudenberg, Drew, and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54: 533–54.
- Gintis, Herbert. 2000. *Game Theory Evolving*. Princeton, NJ: Princeton University Press.
- Keller, Laurent, ed. 1999. *Levels of Selection in Evolution*. Princeton, NJ: Princeton University Press.
- Hamilton, W.D. 1964. "The Genetical Evolution of Social Behaviour I, II." *Journal of Theoretical Biology* 7: 1–16, 17–52.
- Hayek, Friedrich. 1988. *The Fatal Conceit*. Chicago: University of Chicago Press.
- Hirshleifer, Jack. 1977. "Economics from a Biological Perspective." *Journal of Law and Economics* 20: 1–52.
- . 1982. "Evolutionary Models in Economics and Law." *Research in Law and Economics* 4: 1–60.
- Kagel, John H., and Alvin E. Roth, eds. 1995. *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Lorenz, Konrad. 1966. *On Aggression*. New York: Harcourt Brace and World.
- Maynard Smith, John. 1964. "Group Selection and Kin Selection." *Nature* 201: 1145–7.
- . 1993. *The Theory of Evolution*. Cambridge: Cambridge University Press.
- Maynard Smith, John, and George Price. 1973. "The Logic of Animal Conflict." *Nature* 246: 15–18.
- Maynard Smith, John, and Eörs Szathmáry. 1995. *The Major Transitions in Evolution*. Oxford: Freeman.
- Parsons, Talcott. 1937. *The Structure of Social Action*. New York: McGraw-Hill.
- Pinker, Steven. 2002. *How the Mind Works: The Modern Denial of Human Nature*. New York: Viking.
- Poundstone, William. 1992. *Prisoner's Dilemma*. New York: Doubleday.
- Richerson, Peter J., and Robert Boyd. 2004. *Not by Genes Alone: How Culture Transformed Human Evolution*. Princeton, NJ: Princeton University Press.
- Samuelson, Paul. 1993. "Altruism as a Problem Involving Group Versus Individual Selection in Economics and Biology." *American Economic Review* 83: 143–48.
- Simpson, E.H. 1951. "The Interpretation of Interaction in Contingency Tables." *Journal of the Royal Statistical Society B* 13: 238–41.
- Sober, Elliott, and David Sloan Wilson. 1998. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Tooby, John, and Leda Cosmides. 1992. "The Psychological Foundations of Culture." In Jerome H. Barkow, Leda Cosmides, and John Tooby, eds., *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, 19–136. New York: Oxford University Press.
- Trivers, Robert. 1971. "The Evolution of Reciprocal Altruism." *Quarterly Review of Biology* 46: 35–57.
- Wade, Michael J. 1978. "A Critical Review of Models of Group Selection." *Quarterly Review of Biology* 53: 101–14.
- Williams, George C. 1966. *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton, NJ: Princeton University Press.
- . 1992. *Natural Selection: Domains, Levels, and Challenges*. New York: Oxford University Press.
- Williams, George C., and R.M. Nesse. 1991. "The Dawn of Darwinian Medicine." *Quarterly Review of Biology* 66: 1–22.
- Wilson, David Sloan and Elliott Sober. 1994. "Reintroducing Group Selection to the Human Behavioral Sciences." *Behavioral and Brain Sciences* 17: 585–654.
- Wright, Sewall. 1945. "Tempo and Mode in Evolution: A Critical Review." *Ecology* 26: 415–19.