

SOCIAL CATEGORIES AND BUSINESS ETHICS

David M. Messick

In this article, I want to draw attention to one strand of the complex web of processes that are involved when people group others, including themselves, into social categories. I will focus on the tendency to treat members of one's own group more favorably than nonmembers, a tendency that has been called *ingroup favoritism*. The structure of the article has three parts. First I will offer an evolutionary argument as to why ingroup favoritism, or something very much like it, is required by theories of the evolution of altruism. I will then review some of the basic social psychological research findings dealing with social categorization generally, and ingroup favoritism specifically. Finally, I will examine two problems in business ethics from the point of view of ingroup favoritism to suggest ways in which social psychological principles and findings may be mobilized to help solve problems of racial or gender discrimination in business contexts.

The Evolution of Altruism

Social psychologists strive to understand the social nature of humans through scientific means. One possible source of illumination stems from considering what constraints the process of natural selection might have placed on our social nature. (Crawford, Smith, & Krebs, 1987). Specifically, how might natural selection have shaped us with regard to our willingness to offer self-sacrificial help to other humans? This is a deep question that Donald Campbell (1975) raised in his presidential address to the American Psychological Association nearly two decades ago. Campbell's answer was that biological evolution failed to equip us for the social demands of a complex urban society, and that a process of cultural evolution, including the evolution of moral norms, institutions, and cultures, compensated for our biological shortcomings. Campbell's conclusions rested heavily on the work done by biologists exploring routes for the evolution of altruism (see Hamilton, 1964; Trivers, 1971; Wilson, 1980). The cornerstone of his argument was the following principle: altruists who dispensed their altruistic gifts to any and all members of a population would be displaced by nonaltruists who, while enjoying the benefits provided by the altruists, declined to bear the cost of providing benefits to others. Campbell acknowledged that Triver's (1971) concept of reciprocal altruism was realistic, but, Campbell claimed, it was not "true" altruism because it was based on future payback.

There was a puzzle associated with Campbell's analysis that attracted me. The fundamental principle states that *saints*, shall we say, who furnish altruism equally

to all, will not prevail in a dichotomous population of saints and *rogues*, who never help others. However, Trivers shows that *pragmatists*, individuals who help only others who would help them, will prevail in a dichotomous population of pragmatists and rogues. So if we rank the three types in terms of fitness, we get pragmatists over rogues, and rogues over saints. The puzzle arises when one realizes that in a dichotomous population of saints and pragmatists, the two types would be indistinguishable since the only feature that distinguishes them is the way they treat rogues. We get a weak intransitivity among the three types that can only be resolved by investigating a trichotomous population containing all three types.

In 1991 I published an analysis of such a trichotomous population (Messick, 1991). I referred to the practice of exclusively helping other helpers as “group-based altruism.” Altruism was defined as making a choice that provided a benefit, *b*, to another at cost, *c*, to one’s self. Rogues never make that choice, saints always do, and pragmatists pay the cost for other pragmatists and for the saints, but not for the rogues.

The task was to calculate the fitness of the three types as a function of the cost/benefit structure and the density of the three types in the population. The analysis led to the following generalizations: when the population was dense with saints, rogues thrived. When the population was dense with pragmatists, both pragmatists and saints thrived. (In fact, if the density of pragmatists was sufficiently great, the saints actually do better than the rogues, reversing the fundamental principle that rogues are more fit than saints.) An exploration of the dynamics of this population revealed that the only evolutionary stable population—that is a population that (a) is reachable and (b) will resist invasion and mutation to a different state—is a population of pragmatists.

This analysis was characterized by two basic principles that characterize all other theories of the evolution of altruism with which I am familiar. I take these principles to be necessary constraints on the type of altruism that could evolve in humans (and in any other species). The principles are *efficiency* and *discrimination*. Efficiency simply means that there must be an “economic” wisdom to the altruism; specifically, the magnitude of the benefit afforded the recipient(s) must exceed the cost paid by the altruist to generate the benefit. If the cost exceeds the benefit (appropriately measured), the rogues will do better than the helpers, be they saints or pragmatists. However, if the benefit-to-cost ratio is rich, pragmatists and saints will thrive. To put the matter somewhat differently, acts that provide substantial aid to others at low cost to the helper will be favored by natural selection.

The second principle, discrimination, is more important for the purposes of this article. The discrimination principle means that pragmatists will provide benefits to some members of the population and not to other members. Self-sacrificial altruism will occur, but it will not be universal. Ambivalence will exist, in that pragmatists will have to know something about potential recipients before deciding whether or not to help them; a discrimination must be made. If errors are made in the classification of others, the advantage of the pragmatists is reduced. (Furthermore,

my analysis indicates that it is more costly to aid a rogue erroneously than to withhold aid from a saint or pragmatist.) But the point I want to stress is that the only type of altruism that survives is a limited altruism, one that offers help to some and denies it to others.

A postscript to this principle is that the most desirable form of altruism, saintly altruism, can only survive or thrive in an environment of non-saintly pragmatists. Turning the other cheek can prevail only when most will not turn the other cheek. Pragmatists, in other words, make saints' existence possible.

Discrimination is implemented in different ways in different models of the evolution of altruism. In kin selection, for instance, discrimination is based on genetic lineage (Maynard-Smith, 1964). The trade-off is between the cost-benefit ratio and the coefficient of relationship, which is a measure of the likelihood that two individuals will share a gene. The closer the relationship, the more costly the act that could be supported by natural selection. Put otherwise, all things equal, the closer the relationship, the more likely one is to find altruistic helping.

Trivers' (1971) important work on reciprocity assumed that individuals would help others who would also be helpers if called upon. It is clear that he had in mind species that lived in close proximity and that could have the chance for multiple repeated interactions. Trivers' concern was with the development of historical relationships in which the profitable exchange of aid would be fostered and nourished, and he worked out sufficient conditions for this exchange. In Trivers' theory, only other helpers get helped. Non-helpers do not. If nonhelpers can successfully mimic helpers, however, they might get help even though they do not give it. Helpers' advantage exists only to the extent that they can prevent successful mimicry. Much of Trivers' analysis is devoted to the dynamics of deceit, detection, and the "moralistic aggression" against cheaters that prevents cheating from eroding the helpers' advantage.

Wilson's (1979, 1980) theory of group selection allows geography to perform the discrimination. In his model, helpers are non-randomly distributed geographically, so that there are helper-dense regions and helper-thin regions. Even if the benefits are randomly and indiscriminately distributed within one of these regions, which Wilson calls "demes," the clustering itself guarantees an unequal distribution. Those residing in the helper-dense demes--by definition helpers--will be more generously rewarded than those abiding in the helper-lean demes--by definition rogues. Thus the discrimination works globally, but not locally--not within a deme. The interesting paradox of Wilson's theory is that, under the right circumstances, rogues may surpass pragmatists in every deme, because there is no discrimination within a deme, but not surpass them globally. Over the whole population, pragmatists will do better than rogues. Messick & van de Geer (1981) present a simple numerical illustration of Wilson's paradox and relate it to several other statistical paradoxes.

The evolution of our own species inevitably involved a fusion of all three mechanisms: kin, reciprocity, and geographical segregation. Kin plays a central role

in virtually every human society and some things are known about kin recognition generally and in humans (Porter, 1987). Early group living also encouraged the exchange of benefits among non-genetically related individuals, providing a fertile ground for the evolution of reciprocity-based norms. Individuals could learn who the helpers and the rogues were and behave accordingly.

I am not an anthropologist and do not want to speculate on the ways in which early human groups were organized. The point I want to make is that kin selection and reciprocity require an act of discrimination. They require that the condition of discrimination, the unequal dispersal of help, be affected by judgment. The differentiation of kin from non-kin, or the separation of rogues from helpers, was accomplished by a perceptual or judgmental act. The capacity for this judgment, I suggest, is the foundation on which a fundamental feature of our social nature is grounded. That feature involves the tendency to group people into social categories and, further, into ingroups (my group) and outgroups (not my group).

The transition from this abstract discussion of the mechanisms for the evolution of altruism to the real world of social categorization is simple. People who perceive themselves to be members of the same group—members of an ingroup—are more likely to help each other than people who are members of different groups. I will describe some of the evidence for this claim in the next section. The very perception of common category membership is sufficient to induce helping behavior. Common group membership, when made salient, is a proximal cue that renders us pragmatists—we are more likely to help other ingroupers than outgroupers. Further, by precisely the same mechanism, other ingroups are more likely to reciprocate the help. The ingroup specifies the boundaries within which preferential treatment is the norm. In the next section I describe some psychological properties of this process.

Before proceeding, I would like to mention four other points. The first is that the results of my analysis depend on the *differential* distribution of payoffs to helpers (those who would reciprocate) and nonhelpers (those who would not), not on the absolute level of reward. In other words, the analysis could as easily have been done in terms of harming nonhelpers in ways that helpers were not harmed. The framing of the outcomes is not important to the theory although it may be to people. What is important and essential to the success of altruism is the *difference* in the payoffs.

The second point is that the theory is mute with regard to the *content* of the group distinction. It says nothing about what constitutes the ingroup, or the outgroup. While there are some natural (biological) categories, like kin structures, there is nothing in the theory that restricts the results to such biological groupings. The process would work just as well, for instance, if saints and pragmatists had dots on their noses and the pragmatists were programmed to help those whose noses had dots. The cues signaling group membership, in other words, may be quite shallow.

However, the cues signaling group membership should not be easily imitated. As Trivers (1971), Frank (1988), and others have noted, if a cue that evokes benefits from others is easily imitated by individuals who are not members of the group, and who would not, therefore, be inclined to help, then there is nothing to prevent non-

members from imitating the cue to ride free on the beneficence of others. Either the cue must be difficult or costly to forge, or forgeries must be detectable in order for the group membership to be a useful signal. Easily mimicked cues will simply cease to be informative.

Third, there may be *multiple* ways in which ingroups can be defined. In one situation, a dot on the nose may be the essential cue; in another, it may be hair color. It could be nationality, religion, employer, school affiliation, or gender. Some cues, like those of gender, race, and age, are difficult to fake. Others, like corporate or school affiliation, are more easily falsified. The analysis only illuminates the conditions under which a limited group-based self-sacrifice would be favored by natural selection. It does not limit the number of ways in which groups could form.

Finally, the model I described is mute with regard to selection mechanisms. There is no genetic element in the model linking group-based altruism to genetic spread: the argument is strictly based on the expected values of the different strategies. Thus, the model does not require any assumption of biological determinism for its validity. The selection mechanisms may be genetic (and I find it hard to imagine that there would not be a biological element involved) but they may also involve cultural mechanisms as Campbell (1975) argued. It is difficult for me to imagine how one level of evolution could proceed without involving the other. Biology influences culture and vice versa. I mention this point to deflect the criticism that the model I presented is merely another bit of biological determinism that may be used to excuse or justify socially unjust arrangements. The argument is not inherently biological, and even if it were, it would not constitute an excuse or justification for social discrimination.

Social Categorization

One of the things that people do automatically and at many levels is create categories. Our sensory organs map a fraction of the electromagnetic spectrum into the experience of color; our brains “interpret” a storm of stimuli as patterns of solid objects that vary in shape and distance from us; and people everywhere classify the social world into useful categories such as children, parents, women, strangers, or the elderly. We use social categories to refer to nations, races, genders, organizations, professions, and religions. Our category memberships are important features of our social identities. The use of categories, with associated linguistic labels, is different from processing the individuating information we encounter when interacting with a person. Categories of which we ourselves are members (ingroups) have special properties, but before discussing these properties, I will mention some of the general consequences of social categorization.

Tajfel & Wilkes (1963) conducted a pivotal early experiment on the cognitive effects of categorization. They had subjects judge the length of a series of lines whose lengths differed by equal amounts between adjacent stimuli. When the judged lengths were plotted against the objective lengths, the function was a simple

linear relationship. Tajfel & Wilkes had a different group of subjects makes precisely the same set of judgments, only this time the shorter lines were given the label "A" and the longer lines were called "B." This labeling divided the continuum into two categories. The judgments of the subjects in this condition were markedly different from those in the control condition. The slope of the function obtained by plotting the judged length against the objective length was shallower within both the A and B groups of lines, but there was a large jump between the largest A and the smallest B. What seems to have happened as a result of the categorization is that the perceived differences within the two categories were reduced, and the difference between the adjacent stimuli in different categories was enhanced. Perceptually, categorization increased intracategory similarity and decreased intercategory similarity. This distortion in the perception of the stimuli was called the *accentuation effect*.

A number of psychologists asked whether a similar phenomenon occurs in judgments of people. In one illustrative study, Wilder (1978) had subjects watch a videotape of four persons who were deliberating about a mock lawsuit. One of the persons on the videotape expressed an opinion and the subjects were asked to predict the opinion of one of the other persons. What was varied was the extent to which the speaker and the target of the judgment were portrayed as members of the same group. Some of the subjects saw the four sitting at the same table and wearing the same group tag, whereas others saw the four people sitting at four different tables and wearing individual name tags. In the "group" condition, subjects ascribed attitudes and beliefs to the target person that were more similar to those of the speaker than the subjects in the "individual" condition. Perception of a common group membership evoked judgments of greater intragroup similarity than ungrouped individuals.

Groupings can also influence the accuracy with which information about people is remembered. Taylor, Fiske, Etcoff, and Ruderman (1978) showed a videotape of an interracial group discussing a decision problem. Observers of the videotape were later asked to recall what had been said and who had said it. One of their findings was that there were more intra-racial errors of recall than interracial errors. In other words, subjects remember the race of the person who made a statement better than the individual identity of the person. Observers are more likely to misattribute a statement to a same-race person than to a different race person. The rate of intracategory errors is greater than intercategory confusions.

Messick and Asuncion (1993) explored yet a different process by which categorization can create erroneous perceptions of groups. Will Rogers is reported to have claimed that the migration of Oklahomans to California during the drought of the 1930s raised the average intelligence of both states. If the immigrants were below the average intelligence level in Oklahoma and, at the same time, above the mean in California, the phenomenon that the joke is based on could, in fact, have been true. The combined intelligence of both states has not changed, however, since the same people are merely moving about. The movement creates the *illusion* that the collective intelligence has improved.

To demonstrate that gender groupings could produce such illusions, we did an experiment in which we showed subjects a list of names, paired with payments that these persons presumably earned in another experiment. Some were women's names and some were men's. In each list, all women either earned more or less than all men, and women were either in a numerical minority or majority. When women were the better paid majority, 9 of the 12 names were female, and all females had higher earnings than any of the males. When women were in the higher paid minority, only the three top earners had female names. To create this latter condition, the 6 intermediate earnings (Will Rogers' "immigrants") were changed from female to male names. After the list was shown to the subjects, it was withdrawn and they were asked to estimate the average payment for the total group. We hypothesized that subjects would rely on subgroup means in forming their estimates, especially when they were asked to judge these subgroup means before estimating the mean of the whole list.

The data supported our hunch. When the majority was higher paid than the minority (so that the intermediate numbers were included with the superior category), the mean estimated payment for the whole group was significantly lower than when the superior category was the minority. In this latter case, both the women's and the men's mean payment was higher than in the former case despite the 12 payments being precisely the same.

These studies illustrate a few of the ways that merely grouping people into categories can have consequences for perceptions and judgments about them and for memory of what they did and said. A novel aspect of categorization arises when the perceiver is a member of one of the groups (ingroup) and not of another (outgroup). Much of the research that has been conducted on this topic has employed highly artificial means of categorizing people. This artificiality is not merely a consequence of the impoverished imagination of social psychologists, it is a crucial feature of the experimental setting. For instance, a standard way of categorizing people is to display a series of pairs of modern paintings, one of each pair by Paul Klee, the other by Wassily Kandinski, and to ask the subjects to express a preference for one painting from each pair. Subjects are then randomly, independent of their actual choices, told that they are either a Klee-preferrer or a Kandinski-preferrer. It is hard to imagine groups that are more artificial than ones involving common preferences for artists whose existence the students were completely unaware of prior to the experiment.

The artificiality is important because there was a controversy within social psychology about necessary conditions for the creation of social hostility and discrimination. One school, led by Muzifer Sherif (1967), argued that a necessary condition for the arousal of intergroup hostility was a real or perceived conflict of interest between groups. In opposition to this "realistic group conflict theory" was a school of thought, led by Henri Tajfel (1970), which claimed that the mere categorization of people into groups with no history or importance was sufficient to create intergroup discrimination. The importance of these artificial categories is

precisely that they are meaningless and content-free. We do not have pre-existing stereotypes for Klee- or Kandinski-preferring people. Thus discriminatory consequences of mere categorization, if found, could not be attributed to past interactions with members of the ingroup or outgroup or to previous exposure to commonly held stereotypes. Tajfel (1970) suggested that differential treatment of ingroups and outgroups may result from a generalized “norm of outgroup discrimination.” This term has unfortunate connotations that we will come to shortly.

The fundamental finding of Tajfel (1970) and his collaborators Billig, Bundy, & Flament (1971) was that when a group of subjects, British school boys in this early work, was categorized with respect to their preference for either Klee or Kandinski, and when they were given the opportunity to make allocations of potentially valuable points to other members of their class whose personal identities were not known, they tended to make the allocations in ways that provided greater payoffs to anonymous ingroup members than to anonymous outgroup members. This result has been called *intergroup bias*. In this study, and in hundreds that have replicated this basic finding, the allocator gets nothing him or herself so the question of selfishness *per se* is not at issue. The allocations are made to anonymous others.

The intergroup bias is one of the most frequently replicated effects in social psychology. People allocate higher payoffs to members of their own (trivial) group than to members of the other group. A close analysis of subjects’ choices in this “minimal group situation” suggests that while subjects are motivated to maintain a difference between the benefits of their group and the other, they do not try to make this difference as large as they might. They tend to strike a balance between providing fair or equal outcomes to the two recipients, on the one hand, and giving a bit more to the ingroup member on the other. Most subjects’ choices reflect a compromise between these motives.

The fact that trivial categorization of people into groups evokes a mild discriminatory response may be a behavioral reflection of the process I referred to earlier as group-based altruism. The mere establishment of ingroups and outgroups seems to create a differential in beneficence across the category boundary. The model provides a context within which one can see an evolutionary basis for intergroup discrimination. From the point of view of the model, the beneficence differential is a necessary condition for self-sacrificial altruism to survive, either in its saintly form or in its more pragmatic form. If one views this differential in terms of the good that is provided to the ingroup members (relative to the outgroup members) one can, like Campbell (1965), view altruism as an “ethnocentric” motive. If one views the difference in terms of the payoff of which the outgroup members are deprived, it may appear as harm.

The minimal group paradigm does not involve having subjects make decisions that influence their own outcomes, so this paradigm does not allow one to examine the issue of self-sacrificial altruism. However, ample data exist to fill in this gap. One particularly interesting study was conducted by Sole, Marton, & Hornstein (1975), who were interested in the circumstances under which people would provide

anonymous aid to a perfect stranger. The experiment involved leaving an addressed and stamped letter in a park in New York City where it could be found by pedestrians. The help in this case was of little cost, entailing picking up the letter, closing it, and depositing it in a mail box. One of the variables that was manipulated in this experiment was the address on the envelope. Half the envelopes were addressed to a Jewish organization and the other half to an Arabic one. Importantly, the park was in a Jewish neighborhood so the subjects were, by assumption, Jewish. One of the findings of the study was that more letters were mailed when the address was to the Jewish rather than the Arabic organization. In all cases, the senders of the letter were strangers to the subject and both organizations were bogus, so the subjects could not have had prior knowledge of the goals of the organization. There are many other studies that support the generalization that people are more likely to help ingroup than outgroup members although many other factors are also usually involved as well.

Categorization *creates* the conditions that I examined in the group-based altruism analysis. Categorized into ingroups (into the same group), whatever their content, people are more likely to exchange benefits than if they were categorized into different groups. Categorization and the beneficence differential associated with it is therefore one process that would increase the fitness of those whose social repertoire included it. I will come back to this shortly.

The categorization of people into ingroups and outgroups does more than create a beneficence differential. It also creates a difference in the perceived homogeneity of the two groups. Outgroups tend to be seen as more homogeneous than ingroups. The tendency to see “them” as being all alike while “we” are diverse and variegated has been called the *outgroup homogeneity effect* (Quatrone & Jones, 1980; Quatrone, 1986; Simon, 1992). The comment that, “If you’ve seen one _____, you’ve seen them all,” is never made by people about their ingroups. Like most empirical phenomena studied by psychologists, this tendency to see outgroups as more homogeneous than ingroups may, under some circumstances, be reversed (Simon, 1992). However, the empirical generalization seems valid across a wide range of contexts.

There are several sensible reasons for this effect. First, we generally have more experience with and exposure to ingroup than outgroup members. Hence we have a better opportunity to learn that they come in all shapes and sizes. Second, the ingroup, by definition, contains “me,” who is a very special member. To the extent that we contrast ourselves against others to maintain some degree of personal distinctiveness (Brewer, 1988), our presence in the ingroup makes it seem like a more varied collection than the outgroup. Finally, people might be motivated to make more accurate judgments about ingroups than outgroups. The consequences of misjudgments or errors may be more severe and immediate, encouraging more careful processing of information about those with whom we interact more frequently, who are, of course, more likely to be ingroup than outgroup members.

There are a number of interesting implications of this tendency. First, if outgroup members tend to be all alike, then it is more efficient to characterize them with group level features, i.e., to stereotype them, than ingroup members. It is not important to know individuating information about outgroup members if they are perceived to deviate only slightly from the group stereotype or prototype. Thus, outgroups can be more easily stereotyped than ingroups.

One of the consequences of stereotyping a group is that information about particular members becomes relatively unimportant. The category label is all that is needed to guide interactions with outgroup persons. This phenomenon even seems to characterize the level of information that is provided about people. I once noted (Messick & Mackie, 1989) that it was common practice in the *Los Angeles times*, when reporting results of international sports competitions, to provide only the nationality of non-American contestants, but to provide the hometowns of American competitors. I can now report that this is also the policy of the *Chicago Tribune*. We learn more personal details about Americans than about non-Americans. Park & Rothbart (1982) have provided experimental evidence that people encode and recall more detailed information about ingroup than outgroup members, suggesting that newspapers merely provide information at the level of informativeness that people prefer.

To return to the beneficence differential, so long as one can only observe the *difference* between allocations to ingroup and outgroup members, the framing of the difference is ambiguous. We cannot say if ingroups members are “overpaid” or if outgroup members are “underpaid.” However, in some cases it is possible to measure people’s behavior to persons whose ingroup status is unknown. Under these circumstances, it is possible to compare behavior toward ingroups and behavior toward outgroups independently to the neutral standard of behavior toward someone with no group identification. The existence of such an external standard or benchmark allows one to address the question of whether intergroup discrimination involves *ingroup favoritism*, treating ingroup members better than the standard, *outgroup derogation*, treating outgroup members worse than the standard, or both. In an important review of the experimental literature, Brewer (1979) concluded that there was more empirical support for the process of ingroup favoritism than outgroup derogation. In short, there was scant evidence that people wanted to deprive outgroup members of benefits, relevant to a neutral standard, but there was strong evidence that they wanted to benefit ingroup members. It is for this reason that I judged Tajfel’s (1970) term, “norm of outgroup discrimination” unfortunate. The locus of the differential may not fall on the outgroup but on the ingroup.

It is not my goal to present a comprehensive review of research on this topic, but I will describe in some detail two recent studies that are relevant to the issue. Perdue, Dovidio, Gurtman, & Tyler (1990) explored the evaluative implications of ingroup words (*we, us, our, ours,*) and outgroup words (*they, them, their, theirs*). I will briefly describe two of the three experiments that they report in order to share their results and illustrate the methodology of cognitive social psychology.

The hypothesis that motivated these experiments is the idea that ingroup language, in particular the pronouns *we*, *ours*, and *us*, evoke more positive cognitive connotations than corresponding outgroup pronouns like *they*, *theirs*, and *them*. In their first experiment, the researchers employed a standard semantic conditioning task, in which subjects read a sequence of pairs of words. The idea behind semantic conditioning is that if words are repeatedly paired with each other they will come to share their semantic meaning through simple association. Thus if ingroup words have more positive connotations than outgroup words, other words with which ingroup words are paired should be rated more positively than words paired with outgroup words.

The words to be paired with the ingroup and outgroup words were three-letter nonsense syllables like *xeh* and *yof*. In a sequence of 108 trials, one of these nonsense words was paired 20 times with either ingroup words (*we*, *us*, and *ours*) or outgroup words (*they*, *them*, and *theirs*). (The pairings were made in both ways, of course, to control for possible differences in the desirability of the two nonsense words.) The subjects were told during the presentation of the pairs to indicate as quickly as possible which word in the pair was a real word. Following the presentations of the 108 pairs, the key nonsense syllables were rated by the subjects on a scale ranging from *pleasant* to *unpleasant*. The crucial findings were that the word paired with the ingroup pronouns was rated significantly more pleasant than the word paired with the outgroup pronouns. Furthermore, while the ingroup words were rated significantly more pleasant than words paired with controls (*he*, *she*, *his*, *hers*), the ratings of the outgroup words did not differ from those of the controls. In other words, pairing with ingroup pronouns increased pleasantness ratings, but pairing with outgroup pronouns did not decrease pleasantness ratings, relative to a control condition.

In the third experiment reported by Perdue et al (1990), the idea that ingroup words were special was further explored using a semantic priming technique. The basic idea here is that information presented below the threshold for conscious recognition can nevertheless influence the processing of subsequent information. If the priming stimulus is cognitively associated with a subsequent stimulus, the processing of the latter stimulus will be facilitated. If the priming stimulus is negatively associated with the latter stimulus, the processing of the latter input will be inhibited. The facilitation or inhibition can be measured by the speed with which subjects respond to the stimulus. Facilitation naturally leads to faster responding and inhibition leads to slower responses.

In this experiment, subjects were presented with words that described either houses or people. The words of both types were either positive (*roomy* or *trustworthy*) or negative (*drafty* or *cruel*). Subjects watched a sequence of words that were flashed on a computer screen with the instructions to indicate as quickly as possible, by pushing one of two keys, whether the word was positive or negative. All the “person words,” the only ones of concern here, were preceded by the cue “PPPPPP” to indicate that the coming stimulus word described a person. However,

prior to the PPPPPPP cue, one of three primes was flashed on the screen for 75 ms, a duration that is too short for the prime to be recognized. The primes were either *we*, *they*, or a neutral prime, *xxx*. The question was whether the primes would influence the speed of deciding whether the test adjectives were positive or negative.

The results of this experiment were, first, that positive adjectives were processed more quickly than negative. The more interesting result for the current purpose is that the ingroup prime (*we*) facilitated (speeded up) the processing of positive adjectives and inhibited (slowed down) the processing of negative adjectives relative to the neutral prime. The outgroup prime (*they*), on the contrary, produced no measurable differences from the neutral prime. As in the first experiment, we had positive properties that differentiated it from neutral words, but *they* did not.

Business Ethics

I have outlined the contours of a psychological phenomenon that is relevant to a number of issues in business ethics. Social categorization and the associated processes of intergroup discrimination create pervasive problems in modern management. The first problem is the issue of racial discrimination in mortgage lending. At stake here is the allegation that some banks discriminate against African-American and Hispanic mortgage loan applicants, making it more difficult for persons from these groups to secure financing for homes, thus depriving them of the privilege of home ownership. The basic data that have been used to support this accusation were collected under the Home Mortgage Disclosure Act that require banks to record and to report the race or national origin of the loan applicants as well as data on the disposition of loan applications.

The data reported by Conner & Smith (1991; 1992) display a pattern in which minority (specifically, black and latino) applications are denied with greater probability than white applications. For instance, in 1991, the HMDA data indicate that 26.4 percent of the black applications for government-backed loans were denied in contrast to 16.3 percent of the white applications for the same type loans. The denial rates for conventional loans were 37.6 percent and 17.3 percent respectively. Black applicants were more than twice as likely as whites to be denied.

These data are not facial evidence of discrimination, however. Race is correlated with many factors that are relevant to the credit-worthiness of an applicant including income, geography, and credit history. To assess the impact of these and other factors on loan decisions, the Boston Federal Reserve Bank conducted an analysis of lending data in the Boston area (Munnell, Brown, McEneaney, & Tootell; 1992) that attempted to account for an assortment of additional factors that could influence lending decisions and that would not be based directly on race or ethnicity. The result of this study was that when other measures of credit-worthiness were accounted for, the gap between the rejection rates for white and minority applicants was reduced but not eliminated. Race and ethnicity made a difference that was independent of economic measures of credit-worthiness. Munnell et al (1992)

concluded that about 17 percent of minority applicants would be denied loans, in contrast to 11 percent white, even if they had the same loan-to-value ratios, credit histories, obligation ratios, and property characteristics as the white applicants in aggregate.

When thinking about discrimination and the unfairness that is associated with it, many people, like Tajfel, think about it as the denial of opportunity, property, or assets to a group that is entitled to the resource. Discrimination is thought to be the withholding of benefits from the target group or the assignment of negative outcomes to that group. In the case of mortgage lending, the view that is often expressed is that qualified minority applicants are being deprived of mortgages and hence of the opportunity for home ownership.

This inference is not supported by the finding of a difference in denial rates. The difference *may* be the result of depriving qualified minorities of loans, or, as the ingroup favorability hypothesis suggests, it may be the result of approving loans to unqualified white applicants. The social psychological research that I discussed in the previous section suggests that the main channel taken by intergroup discrimination is through favoring the ingroup, not by harming the outgroup. Applied to the current issue, the hypothesis emerges that the difference in denial rates may result more from the awarding of loans to unqualified white applicants than from the denial of loans to qualified minority applicants.

However, since there is no absolute standard for assessing whether an applicant is or is not qualified for a loan of a given magnitude, it is not possible to differentiate these hypotheses in a simple and compelling manner. From the point of view of fairness, individuals from different groups appear to be treated differently but the motivation for the differential treatment is obscure. I am not qualified to say whether there is a moral distinction to be made between discrimination that results from favoring ingroups and that which stems from depriving outgroups, but there are important psychological differences between these processes.

One of the most important differences has to do with the level of consciousness at which the discrimination occurs. Recent research by Gaertner & Dovidio (1986) indicates that many whites disapprove of racism at the same time that they themselves have negative attitudes about racial minorities like blacks. This configuration of attitudes, labeled *aversive* racism, indicates that people do not want to think of themselves as racially prejudiced or biased. Moreover, people will avoid actions that would tend to indicate that they have negative views toward minorities. To illustrate, Gaertner & Dovidio (1977) put white subjects in a situation in which they presumably witnessed an emergency in which a second subject in the experiment was injured by a falling bookcase. (Naturally the emergency was staged and the second subject was uninjured!) The experimenters varied the race of the victim, and also the number of other witnesses to the emergency. When the subject was the sole witness, she nearly always responded appropriately--she called the experimenter, sought help, etc.--regardless of the race of the victim. However, when there were other witnesses, the subject helped the white victim more frequently than

the black victim. Gaertner & Dovidio (1977) explain the result in the following way. If the subject is the only witness, a salient reason that she might have for not helping the black is her race. This reason is not personally acceptable, so the subject helps. However, when there are several witnesses, one can explain one's failure to help by assuming that another witness has done so. Thus subjects can avoid ascribing negative racial attitudes to themselves for not helping, and the frequency of helping the black victim drops.

In the case of mortgage lending, denying a loan to a clearly qualified minority applicant would not only be poor business, it might arouse the uncomfortable self-perception that the loan officer was racially biased. The concept of aversive racism suggests that loans will be denied when there is good reason to do so--when there are material problems with the loan application. However, when a white applicant presents an application with an identical or similar problem, the loan officer may well respond differently. First, the officer may discount the seriousness of the problem and overlook it. Alternatively, the officer may coach the applicant on how to correct the problem. Perhaps describing assets differently may create a better impression, or explaining why there was a late payment on a credit card bill may make the infraction seem less ominous. In these cases, the reasons the officer would have for helping may not have anything to do with race at all. Race is usually not salient when two members of the same race interact. It is more likely that a loan officer would offer some extra help because she liked the applicant, felt sympathy for the person, or because she thought the applicant was deserving.

The process that I am suggesting is one in which race *per se* may not enter the mind of the loan officer as a criterion for making a decision. The black applicant's loan was denied for good reason, and the officer made an admirable extra effort to help the white applicant. He or she would be offended at the accusation of discriminating against minorities. Nevertheless the decisions that such a person might make would display precisely the pattern of denials that the Boston Fed. Study found.

This hypothesis, like most others, suggests that sound applications, both minority and white, will be approved, and that poor applications of both types will be denied. In other words, discrimination will be found among the ambiguous loan applications, those which are neither clear winners nor obvious losers.

One final implication of the ingroup favoritism hypothesis of racial discrimination in mortgage lending concerns institutional changes that many banks have mandated to combat differential treatment of white and minority customers. Many banks have responded to the racial discrepancy by establishing review committees whose job it is to evaluate each rejected minority loan application. The presumption is that such a review process will detect those minority applications which were inappropriately rejected, giving the bank the chance to reverse the rejection before the applicant is notified. (A secondary motive for such review committees may be to make the loan officers more sensitive to possible differences in their evaluative criteria, but no bank has publicly acknowledged that this is a

reason for the review process.) The problem with this review process is that it seems to be based on the assumption that the difference in denial rates stems from the rejection of qualified minorities. Our alternative hypothesis implies that these review committees will find little to reverse when they review denied minority applications because the denied applications will deserve to be denied. Where a review process should focus its attention, according to our hypothesis, is on the marginal white applications that are approved.

Differentiating one of these hypotheses from the other is difficult and perhaps impossible because there is no absolute standard against which to compare an application to assess its merit. Furthermore, there is no reason why both ingroup favoritism and outgroup derogation cannot be happening simultaneously. The mere existence of two hypotheses does not require that they be mutually exclusive or exhaustive. There are other possible explanations for the difference in rejection rates that have little to do with racial attitudes. For instance, if loan officers are more effective in getting loans approved for personal acquaintances than for strangers, and if most of their acquaintances are white, then the (unintended) impact of these friendship networks will be to benefit white applicants more than minority applicants.

The hypothesis that I have proposed is a benign interpretation of discrimination in that it does not suppose that the discriminator hates or dislikes the victim. The discriminator may not even be aware of the discriminatory consequences of his or her behavior. I do not want to suggest that all discrimination is so benign; there are clearly cases where discriminatory attitudes and behavior are malicious, obvious, and all too conscious. These cases are probably in the minority, however, and their existence should not mislead us into believing that all discrimination is so virulent.

The ingroup favoritism hypothesis also applies to hiring and promotion practices. The consequences of loan decisions do not impact the decision maker directly in most cases, whereas decisions about hiring and promotion often do alter the social composition of the environment in which decision makers work. (Mortgage bankers might be thought of as being like subjects in the minimal group paradigm in that they are making allocations of someone else's resources to people other than themselves.) It might be expected, therefore, that intergroup discrimination will take a more conspicuous form in the later than the former cases. This is to say, category-based beliefs, e.g., stereotypes, are likely to play a more explicit role here than in mortgage lending, and there is ample evidence that employment discrimination is a major problem in many organizations, including government agencies. Recent accounts in the popular press have identified insurance companies, accounting firms, police and fire departments, the Federal Bureau of Investigation, and, most recently (as of this writing), the Immigration and Naturalization Service (INS) as organizations with serious problems of racial or gender discrimination.

Because discrimination is more conspicuous in these cases, social psychologists have played an active role identifying and combating it. Fiske, Bersoff, Borgida,

Deaux, & Heilman, (1991) described the testimony that was prepared and presented by an expert witness in the case of *Hopkins v. Price Waterhouse* (1989). This case involved a gender discrimination suit brought by Ann Hopkins against Price Waterhouse after Ms. Hopkins was turned down for a partnership. The case eventually went to the Supreme Court, and it represents the first time that evidence based on psychological research on sex stereotyping was heard by the Supreme Court. (The first time that social science evidence was used in a race discrimination case was in 1954 in *Brown v. Board of Education* which dealt with racial discrimination in education.)

Fiske et al (1991) suggest that stereotyping is made more likely by several factors. (The authors address the problem of gender discrimination in their article, but there is no reason to think that the factors described would not characterize discrimination against other groups as well.) One important factor is that the victimized group is in a small minority and has typically not been found in the roles where discrimination is found. At the time when Hopkins was being considered for partnership, for instance, only 7 of the 662 partners were women, and she was the sole woman of the 88 people who had been proposed for partner that year.

A second factor is that the criteria (for hiring or promotion) are ambiguous, vague, subjective, or contradictory. Under such circumstances, beliefs about ingroups and outgroups can influence judgments without being conspicuous. A corollary is that stereotyping is more likely when information about target persons is ambiguous or vague. As we mentioned in regard to the mortgage lending problem, discrimination is not likely to occur when applications are either very poor or very good. It is the mixed applications with both positive and negative features that will invite the use of discriminatory beliefs and information. In the Hopkins case, the reason that was given for Hopkins denial was that she had problems with "interpersonal skills." It is instructive that interpersonal skills are much more ambiguous and hard to define than bookable hours or other quantitative measures of business success. (In this case, interpersonal skills seem to have meant a lack of stereotypical femininity.)

Yet another factor is a perceived lack of fit between the person's category and the occupation. Eagly & Mladinic (1994) summarize evidence that general views of women are often more positive than general views of men. However, prejudice arises against women in male-dominated contexts. Women, stereotypically, do not have the aggressiveness or toughness to be leaders, and women who do display these qualities are often disliked for their lack of femininity. This pattern creates a no-win bind for women managers. They can behave in a nurturant and stereotype-consistent fashion (and confirm the "too nice to succeed" stereotype) or they can behave aggressively (and be condemned for lacking social-code for feminine-skills).

The stereotype of male leadership is supported partly by normal memory processes. Evidence that men are better leaders than women comes partly from the differential ease with which we can think of exemplary male as opposed to female

leaders. Male and female examples are differentially available to memory. Availability has long been recognized as a cue that can bias and distort judgments, and in this case the bias is exaggerated by the fact that women have historically been excluded from positions of leadership. The fact that there have been more male than female leaders of all stripes also means that it is easier to think of more male than female examples of inept or evil leaders. Differential availability of examples of male leaders is evidence only of the fact that historically there have been more male than female leaders.

When a group has not been represented in an occupation or role, the newcomers may create discomfort among those who traditionally occupied the role. Whether they are women in traditionally men's jobs, or blacks in jobs that have been historically white, the initial entrants may be perceived as intruders. Historically, attractive jobs that command resources and power have been held by white males, so the prototypical problem is one of white men resisting either women or ethnic minorities as newcomers. Some of the motivation for the resistance is the perception that the ingroup (male and/or white) quality of the organization would be destroyed by the presence of the newcomers.

In the Hopkins case, Fiske (1993) suggests that the partners of PW saw themselves as an exclusive gentlemen's club in which women were guests who were expected to obey the rules laid down by the men. In a similar vein, a *New York Times* article quotes John J. Washington, a black special agent in the Immigration and Naturalization Service, as stating, "The problem is that the good old boys still treat this agency like their clubhouse. And they still don't want to let us in" (Sontag & Engelberg, 1994, p. 1). The desire to exclude newcomers is motivated by the wish to preserve the culture and fellowship of the organization, not necessarily by a wish to harm others. The goal is to protect the ingroup.

But favoring an ingroup, as we all know, may result in harm to the outgroup. Exclusion on the basis of race or gender is also illegal in many situations, including those pertaining to hiring and promoting. Why are leaders of organizations not more concerned about the consequences of discrimination? One answer to this question is that the motivation behind the discrimination appears positive, not negative. The reasons for maintaining things as they were (governed by white males) have to do with preserving valued aspects of an organization, with promoting and continuing cultural traditions that were successful, with respecting and honoring the past of the organization.

Fiske (1993) offers a novel answer to this question by suggesting that leaders are also overburdened with information and have little attention or cognitive capacity left for such matters. Attention, Fiske notes, is directed toward those who control resources and outcomes, toward powerful people. Outsiders, almost by definition, are not powerful people. The partners at PW did not need women to be successful, but women who wanted to be successful at PW did need the partners. Fiske's analysis of the distribution of attention in organizations corresponds remarkably well to Jackal's (1988) descriptions of the corporations he observed and the

executives he interviewed. Those in powerful ingroups have little motivation to look outside.

This observation raises the question of whether ingroup favoritism characterizes groups that lack power as well as those that have it. Ely (1994) summarizes evidence that only powerful, high-status groups display ingroup favoritism, and Eagly & Mladinic (1994) report weak evidence, at best, of gender-based ingroup favoritism. It may further be that ingroup favoritism is maximal to groups that are less broad and diffuse than gender groups. Brewer (1993) suggests that the tension between individual and group identity results in maximal identification with groups of intermediate size. These studies raise questions about the ubiquity of ingroup favoritism, questions that must be addressed by more research.

I have suggested that the tendency to divide the world into ingroups and outgroups and to display preferential treatment to members of ingroups is a fundamental feature of our social nature. I think that a strong case can be made that this tendency is basic to the development and maintenance of prosocial self-sacrificial actions, whether they are for the benefit of our families, our countries, the organizations in which we work, or the professional groups we voluntarily join. The differential treatment of ingroup and outgroup members may be interpreted as positive, when the focus is on self-sacrificial actions taken on behalf of ingroup members, or negative, when the focus is on the failure to help outgroup persons.

I do not want to leave the impression that such behavior is unalterable simply because it derives from a fundamental part of our social nature. Extreme behavioral flexibility is also part of our psychological nature, and ingroup favoritism, like many other basic processes, can be controlled and shaped. Thus, I would like to conclude by suggesting several measures that can be taken to combat the destructive effects of ingroup favoritism in organizations.

As I said earlier, there is ample evidence that group-based beliefs (e.g., stereotypes) are most influential when other information for evaluation is ambiguous, contradictory, subjective, or vague (Fiske & Taylor, 1991). Preferential ingroup treatment may be reduced by making criteria explicit, objective, and public. One of the first steps that can be taken in this direction is to maintain accurate data that can signal whether discrimination is a problem. The Home Mortgage Disclosure Act, for instance, requires that banks collect and disclose data on race and ethnicity. Without these data, judgments of the equality of treatment of white and minority loan applicants would be subjective and idiosyncratic. The data that are now available may not be completely understood, but they provide a focus for research to identify the cause or causes of the difference in denial rates.

Also, as I mentioned earlier, we would expect that preferential treatment will be more likely to occur with regard to loan applications that are not clear winners or losers. It is with the harder ones, those that have positive and negative qualities, that we should expect to see the influence of group-based beliefs; ambiguous applications can support either a positive or a negative decision. To the extent that it is feasible, we should strive to develop standards and rules for combining the

different criteria that can be used more or less mechanically. The development of quantitative procedures for assessing the quality of an application does not guarantee that discrimination will not occur; it only makes it more difficult for it to occur through the unequal application of qualifying standards. Elster (1992) has discussed how the use of quantitative indices for the evaluations of students in exclusive private universities has altered the way in which discrimination is practiced in these institutions.

Criteria for promotion and hiring can be made as objective and quantitative as possible. Judgments that one “does not fit in” or that one “does not have the right stuff” are vague and subjective. Such appraisals are often equivalent to saying that one is “different,” which may be true and irrelevant at the same time. As Judge Gesell stated in his decision on *Hopkins v. Price Waterhouse*, “A far more subtle process [than the usual discriminatory intent] is involved when one who is in a distinct minority may be viewed differently by the majority because the individual deviates from an artificial standardized profile.” (Quote from Fiske et al, 1991, p. 1052). A considerable amount of research has been done on procedural justice in regard to performance appraisals, and a common finding is that employees believe that the information used for these purposes should be relevant (Lind & Tyler, 1988; Sheppard & Lewicki, 1987). Race, gender, and other categorical information is usually irrelevant. It will not always be possible to eliminate subjective elements from judgments about the suitability of people for employment, but people can learn to be cautious in making such judgments.

One way to reduce discrimination is through affirmative action programs. Banks that want to insure against racial discrimination can hire minority loan officers. (Naturally, proposals to hire more women or minority loan officers may generate resistance among white male loan officers, so both the hiring and the lending problems must be faced simultaneously with this approach.) This strategy may reduce discrimination against minority applicants, but might it also lead to “reverse” discrimination against white applicants by minority loan officers? This is a possibility. Let us assume for the moment that people cannot avoid some degree of preferential treatment of other people who are like themselves, and further assume that race is one dimension on which people judge others to be like or unlike themselves. The question is whether it is fairer to have loan officers of many races and ethnic backgrounds, so that the preferential treatment will be broadly distributed, or to populate the ranks of the loan officers with people of mainly one category? I would argue that the former is a more equalitarian and hence preferable arrangement.

There is one basic premise that flows from the ingroup favorability idea and that is the notion that ingroup members will be viewed and treated more positively than non-ingroup members. Gaertner and his colleagues have noted that this formulation implies that some ways of reducing discrimination may be preferable to others (Gaertner, Dovidio, Anastasio, Bachman, & Rust, 1993; Gaertner, Rust, Dovidio, Bachman, & Anastasio, 1994). Specifically, there is a difference between

eliminating group differentiation completely, so that each person is seen as an individual and not a group member, on the one hand, and extending the superordinate group category so that everyone is treated as an ingroup member. In a situation where people are perceived as ingroup or outgroup members, eliminating the category structures and having everyone treated as individuals, assuming that this is possible, will reduce intergroup bias by reducing the evaluations of the persons who were formerly ingroup members. Extending the ingroup designation, so that previously classified outgroupers now become ingroupers, *increases* their evaluations. Extending ingroup identity, or the *common ingroup identity* approach, results in an increase in the evaluation of others whereas the individuation approach leads to a decrease in the evaluations of the ex-ingroupers. Gaertner, Mann, Murrell, & Dovidio (1989) reported data that are consistent with this hypothesis.

From the point of view of organizations, the common ingroup identity position suggests that it is desirable to create a superordinate category to which everyone belongs than to try to reduce bias by ignoring social categories altogether. This later approach, called the color-blind method in some contexts, fails to exploit the positive aspects of ingroup favoritism, the willingness of people to aid other ingroupers. The common ingroup identity approach promotes the tendency for mutual aid among the ingroup. Indeed many corporations strive to make the corporate identity a vivid component of employees' social identity through portrayals of the firm as a "family," through the creation of athletic teams to compete against other firms' teams, through internal publications that praise efforts of the firms' employees, and through many other techniques to develop loyalty and commitment to the company (see, for instance, Jackal, 1988). However, these efforts are vulnerable to strategic exploitation--to the creation of loyalty and a sense of belonging to an important ingroup, for the sole purpose of evoking sacrifices from the members. In the case of corporations, the economic goal would be to maintain a labor force that is committed, willing to make sacrifices, that displays high morale (and a correspondingly low turnover rate, rate of absenteeism, and drain on healthcare resources). The creation of a strong sense of corporate identity among employees, in short, should lower labor costs.

At this point, I return to the first part of this article and reiterate that the model of group-based altruism which I have proposed as the wellspring of ingroup favoritism requires vigilance against imposters and sham requests for sacrifice. Just as the willingness to aid ingroupers is psychologically fundamental, so is the associated skepticism about requests for help and the need to discriminate between people who truly would reciprocate from those who merely pretend that they would. Corporations that attempt to instill a prefabricated sense of a common ingroup among employees for purely economic reasons may find that their efforts are met with resistance, skepticism, and possibly downright mockery. This response would be predicted, for instance, if some members of the firm were widely perceived as being unwilling to make the types of sacrifices that were being asked of ordinary employees. The corporate ingroup can only have psychological meaning if all

members of the firm, from the CEO on down, display a shared willingness to sacrifice for each other. Otherwise the calls for sacrifice will be perceived as bogus and unjust.

Conclusion

I have suggested that intergroup discrimination may be deeply linked to the human tendency to help others. Since self-sacrificial helping can improve evolutionary fitness only if it is denied to non-helpers, the differentiation of helpers and non-helpers is crucial. Targets of helping (and of non-helping) are defined, in part, through a process of categorization. Among other things, categories, whether linked to biology or not, provide a cognitively efficient way to sort the world. Sorting into ingroups (us) and outgroups (them) is especially useful. We are more likely to help ingroup members (and they us) than outgroup persons. Proximal psychological mechanisms that differentiate ingroups from outgroups may not involve conscious considerations of category differences, as I suggested with regard to mortgage lending. Or they may explicitly be category based as with gender stereotypes in the context of hiring and promotion where historical ingroups (white men) may feel threatened by newcomers. Finally, I mentioned a few measures that can be taken to reduce category-based decision making in organizations. None of the measures will be trivial to implement. Possibly the most effective solution that could be imagined, striving to create a common ingroup identity in order to preserve the social supportiveness of ingroups, is likely to be fragile. It could easily be undermined, for example, by evidence that some members of the group are immune from the sacrifices they are requesting from others.

REFERENCES

- Brewer, M. B. (1979). "In-group bias in the minimal intergroup situation: A cognitive-motivational analysis," *Psychological Bulletin*, 86, 307-324.
- Brewer, M. B. (1988). "A dual process model of impression formation," in T. K. Srull & R. S. Wyer (Eds), *Advances in Social Cognition* Vol 1 (Hillsdale, NJ: Erlbaum).
- Brewer, M. B. (1993). "Social identity, distinctiveness, and in-group homogeneity," *Social Cognition* 11: 150-164.
- Brewer, M. B. (1994). "In-group favoritism: The subtle side of intergroup discrimination," paper presented at the Conference on Behavioral Research and Business Ethics, Northwestern University, August, 1994.
- Campbell, D. T. (1965). "Ethnocentric and other altruistic motives," in D. Levine (Ed), *Nebraska Symposium on Motivation* Vol 13 (Lincoln: University of Nebraska Press).
- Campbell, D. T. (1975). "On the conflicts between biological and social evolution and between psychology and moral tradition," *American Psychologist* 30: 1103-1126.

- Conner, G. B. & Smith, D. S. (1991). "Home mortgage disclosure act: Expanded data on residential lending," *Federal Reserve Bulletin* 77: 859-880.
- Conner, G. B. & Smith, D. S. (1992). "Expanded HMDA data on residential lending. One year later," *Federal Reserve Bulletin* 78: 801-824.
- Crawford, C. Smith, M. & Krebs, D. (1987). *Sociobiology and Psychology* (Hillsdale, NJ: L. Erlbaum).
- Eagly, A. H. & Mladinic, A. (1994). "Are people prejudiced against women? Some answers from research of attitudes, gender stereotypes, and judgments of competence," (1-35) in W. Stroebe & M. Hewstone (Eds.) *European Review of Social Psychology* Vol. 51 (London: John Wiley).
- Elster, J. (1992). *Local Justice* (New York: Russell Sage).
- Ely, R. J. (1994). "The effects of organizational demographics and social identity on relationships among professional women," *Administrative Science Quarterly* 39: 203-238.
- Fiske, S. T. (1993). "Controlling other people: The impact of power on stereotyping," *American Psychologist* 48: 621-628.
- Fiske, S. T., Bersoff, D. N., Borgida, E., Deaux, K., & Heilman, M. E. (1991). "Use of sex stereotyping research in Price Waterhouse v. Hopkin," *American Psychologist* 46: 1049-1060.
- Fiske, S. T. & Taylor, S. E. (1991). *Social Cognition* (2nd Edition) (New York: McGraw-Hill).
- Frank, R. (1988). *Passions Within Reason* (New York: Basic Books).
- Gaertner, S. L. & Dovidio, J. F. (1977). "The subtlety of white racism, arousal, and helping behavior," *Journal of Personality and Social Psychology* 35: 691-707.
- Gaertner, S. L. & Dovidio, J. F. (1986). "The aversive form of racism," in Gaertner, S. L. & Dovidio, J. F. (Eds.) *Prejudice, Discrimination, and Racism* (61-89) (Orlando, FL: Academic Press).
- Gaertner, S. L., Dovidio, J. F., Anastasio, P. A., Bachman, B. A., & Rust, M. C. (1993). "The common ingroup identity model: Recategorization and the reduction of intergroup bias," in W. Stroebe & M. Hewstone (Eds.) *European Review of Social Psychology* Vol. 4 (London: John Wiley).
- Gaertner, S. L., Mann, J. A., Murrell, A., & Dovidio, J. F. (1989). "Reducing intergroup bias: the benefits of recategorization," *Journal of Personality and Social Psychology* 57: 239-249.
- Gaertner, S. L., Rust, M. C., Dovidio, J. F., Bachman, B. A., & Anastasio, P. A. (1994). "The contact hypothesis: The role of common ingroup identity on reducing intergroup bias," *Small Group Research* 25: 224-249.
- Hamilton, W. D. (1964). "The genetical evolution of social behaviour," I. *Journal of Theoretical Biology* 7: 1-16.
- Jackal, R. (1988). *Moral Mazes* (New York: Oxford University Press).
- Lind, E. A. & Tyler, T. R. (1988). *The Social Psychology of Procedural Justice* (New York: Plenum Press).

- Maynard Smith, J. (1964). "Group selection and kin selection," *Nature* 201: 1145-1147.
- Messick, D. M. (1991). "The evolution of group-based altruism," in R. Selten (Ed.) *Game Equilibrium Models I* (304-328) (Berlin: Springer-Verlag).
- Messick, D. M. & Asuncion, A. G. (1993). "The Will Rogers Illusion in judgments about social groups," *Psychological Science* 4: 46-48.
- Messick, D. M. & Mackie, D. M. (1989). "Intergroup relations," *Annual Review of Psychology*, 40: 45-81.
- Messick, D. M. & van de Geer, J. P. (1981). "A reversal paradox," *Psychological Bulletin* 40: 45-81.
- Munnell, A. H., Brown, L. E., McEneaney, J., & Tootel, G. M. B. (1992). "Mortgage lending in Boston: Interpreting HMDA data," Federal Reserve Bank of Boston, Working Paper 92-7.
- Park, B. & Rothbart, M. (1982). "Perception of outgroup homogeneity and levels of social categorization: Memory for the subordinate attributes of in-group and out-group members," *Journal of Personality and Social Psychology* 42: 1051-1068.
- Perdue, C. W., Dovidio, J. F., Gurtman, M. B., & Tyler, R. B. (1990). "Us and Them: Social categorization and the process of intergroup bias," *Journal of Personality and Social Psychology* 59: 475-486.
- Porter, R. H. (1987). "Kin recognition: Functions and mediating mechanisms," in C. Crawford, M. Smith, & D. Krebs (Eds.) *Sociobiology and Psychology* (175-204) (Hillsdale, NJ: L. Erlbaum).
- Quattrone, G. A. (1986). "On the perception of a group's variability," in S. Worchel & W. A. Austin (Eds.) *Psychology of Intergroup Relations* (25-48) (Chicago: Nelson-Hall).
- Quattrone, G. A. & Jones, E. E. (1980). "The perception of variability within in-groups and out-groups: Implications for the law of small numbers," *Journal of Personality and Social Psychology* 38,: 141-150.
- Sheppard, B. H. & Lewicki, R. J. (1987). "Toward general principles of managerial fairness," *Social Justice Research* 2: 1161-176.
- Sherif, M. (1967). *Group Conflict and Co-operation* (London: Routledge and Kegan Paul).
- Simon, B. (1992). "The perception of ingroup and outgroup homogeneity: Reintroducing the social context," in W. Stroebe & M. Hewstone (Eds.) *European Review of Social Psychology* Vol. 3 (1-30) (Chichester: John Wiley).
- Sole, K., Marton, J., & Hornstein, H. A. (1975). "Opinion similarity and helping: Three field experiments investigating the bases of promotive tension," *Journal of Experimental Social Psychology* 11: 1-13.
- Sontag, D. & Engelberg, S. (1994). "Black officers in I. N. S. push racial boundaries," *New York Times* (P. 1), October, 30.

- Tajfel, H. (1970). "Experiments in intergroup discrimination," *Scientific American* 223: 96-102.
- Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). "Social categorization and intergroup behavior," *European Journal of Social Psychology* 1: 149-178.
- Tajfel, H. & Wilkes, A. L. (1963). "Classification and quantitative judgment," *British Journal of Psychology* 54: 101-114.
- Taylor, S. E., Fiske, S.T., Etcoff, N. L., & Ruderman, A. J. (1978). "Categorical and contextual basis of person memory and stereotyping," *Journal of Personality and Social Psychology* 36: 778-793.
- Trivers, R. L. (1971). "The evolution of reciprocal altruism," *Quarterly Review of Biology* 46: 35-57.
- Wilder, D. A. (1978). "Reduction of intergroup discrimination through individuation of the outgroup," *Journal of Personality and Social Psychology* 26: 1361-1374.
- Wilson, D. S. (1979). "Structured demes and trait-group variation," *American Naturalist* 113: 606-610.
- Wilson, D. S. (1980). *The Selection of Populations and Communities* (Menlo Park, CA: Benjamin Cummings).

I want to thank numerous colleagues for their comments, suggestions, and criticisms of the initial version of this paper. First I want to thank Dawn Elm, Edwin Hartman, and Donna Wood who were the discussants for this paper at the Ruffin Conference. I also want to thank the many participants who raised interesting issues during the conference. Marilyn Brewer, Robyn Dawes, Alice Eagly, Robyn Ely, and Samuel Gaertner get special thanks for reading the paper and sending me their thoughts about ways in which I could improve it. Finally Katie Shonk has polished my prose and rendered the paper more readable. The paper has benefitted greatly from all this help, but any remaining errors, infelicities, or silliness I claim for my own.