In S. Levy and M. Killen (Eds.) Intergroup. Relatonships: An integrative developmental and social psychology perspective. Oxford University Press.

6

The Development of Intergroup Social Cognition

Early Emergence, Implicit Nature, and Sensitivity to Group Status

Mahzarin R. Banaji, Andrew Scott Baron, Yarrow Dunham, and Kristina Olson

As we write this chapter in the fall of 2006, the United States is engaged in a war in Iraq that is far deadlier and more complex than even its supporters had imagined. In early December 2006 Israel and Lebanon entered into air and land warfare that escalated while the UN scrambled and editorials hollowly demanded a cease-fire. Dozens more intergroup conflicts rage all over the world, with lives lost every minute. The scope would increase exponentially if the list were to include other forms of harm humans inflict on each other—economic, sociopolitical, and psychological—that have their origins in conflict between groups, be their differences real or imagined.

If we turn the clock back about 6 million years, we stand at a moment in evolution when anatomically modern humans separated from the great African apes (Wolpoff, 1999). Since then the remarkable achievements of the species, from the invention of the first tools to the technological and information revolutions of the present century, do not include what might have been a fairly easy achievement—a sensible solution to group-based conflict. Yet that is simply not the case. Even in the 21st century, among the greatest threats to the survival of the planet is the damage that group-based actions provoke in "us" and in "them." How are we to understand the continuing state of pervasive and enduring intergroup conflict?

Analyses of wars, genocide, holocausts, and pogroms show that they are not restricted to small segments of society or particular regions of the world. Indeed, group-based harm is not the result of a "few bad apples," as some would like to believe, for it occurs in surprisingly virulent forms even in societies that take pride in their "civilized" way of life (e.g., Europe in the 20th century). Stanley Milgram's analyses of

obedience to authority failed to find a single demographic variable that protected groups from collapsing under the direction of an authority figure (Milgram, 1963), and the various sciences of group behavior have confirmed numerous times the strong proclivity of ordinary humans toward intergroup conflict even without much provocation.

For the longest time, concepts of attitude (preferences) and belief (stereotypes) have seemed alien to the very idea of systematic scientific inquiry, in part perhaps because these aspects of mind veer too closely to questions of ethical and moral import and create the unease of "is" versus "ought" questions—the former squarely within the boundaries of science, the latter not. Moreover, such questions are at their core questions of *preferences* and *values* and hence are not merely matters of right and wrong but of good and bad.

An evolving science of entities such as attitudes, beliefs, and values, especially as they are expressed in the early years of childhood, are our concern in this chapter on intergroup conflict. We stay with questions of "is" and focus on children because the "is" of the development of intergroup cognition can shed light on the stubborn problem of intergroup conflict as it originates in the minds of 21st-century humans.

In the last few decades discoveries of the bounded nature of rationality (Simon, 1957; Tversky & Kahneman, 1974) and the unconscious operation of social cognition (Bargh, 1996; Fazio, Jackson, & Dunton, 1995; Greenwald & Banaji, 1995) have allowed an honest appraisal of the way in which human minds create outcomes that deviate from their intended goals (Uleman & Bargh, 1989). From these traditions of research we have discoveries that fly in the face of received opinion about human nature, lay beliefs, and their assumed rationality. If we agree with Patrick Moynihan (1994) that "Everyone is entitled to their own opinion, but not their own facts," no matter how surprising we find the facts about our nature, they deserve our attention.

From the writings of biological anthropologists like Richard Wrangham and David Hamburg (Wrangham & Peterson, 1996; Wrangham, Wilson, & Muller, 2006; Hamburg, 1991; Hamburg & Hamburg, 2004), we learned of the brutal warfare that chimpanzees, who are among our closest primate relatives, are capable of inflicting on others of their own species. To understand intergroup behavior, we begin by acknowledging this aspect of our evolutionary history. It may tell us what we are likely to encounter when we attempt to understand human behavior, especially in the context of groupings based on religion, race, nationality, geography, political ideology, and language, to name just a few. In doing so we will not only entertain the deep origins of our predisposition to conflict but also necessarily set high standards as we work on the "ought" side of the equation (e.g., by realizing that bringing about a permanent state of intergroup peace is likely to be a nontrivial enterprise).

Research on human behavior confirms some of the pessimism evoked by the primate data. Robert Putnam's (2006) study shows that as ethnic diversity in a geographic region increases, social trust decreases. For example, whites in Los Angeles are less trusting of their white neighbors than those same persons would be if they lived in less diverse Montana. Our own work on the unconscious biases of those who consciously intend to be unbiased adds to this evidence (Nosek, Banaji, & Greenwald, 2002). In addition, research shows that, unbeknownst to us, subcortical brain regions involved in the learning of emotion show stronger activation in response to out-group than in-group faces (Cunningham, Johnson, Raye, Gatenby, Gore, & Banaji, 2004; Phelps, O'Connor, Cunningham, Funayama, Gatenby, et al., 2000). Moreover, Olsson, Ebert, Banaji, and Phelps (2005) have demonstrated that fear of in-group members dies away at a faster rate than fear of out-group members. The fact that we spontaneously separate us from them is well documented in adults even though little work has explored the developmental trajectory of such tendencies.

On the other hand, another side of human nature often gets overlooked when we are discussing this grim view. This other side has two components, one that shows a cooperative and peace-loving predisposition that suggests great adaptability in the face of new demands. Humans have equally close primate cousins who evolved to produce nonhierarchical and peaceful means of resolving conflict.

Observations of the changes in male primates raised in aggressive groups and then transferred to pacifist groups is instructive (de Waal & Johanowicz, 1993). Rhesus macaques, who are aggressive and hierarchy enhancing when transferred to the more peaceful and cooperative stump tail macaques, seem to alter their behavior within a few months (de Waal & Johaniwicz, 1993). In spite of millennia of genetic proclivity toward aggression, coupled with a lifetime of

socialization toward aggressiveness, these cousins of ours can be transformed in a couple hundred days to adapt to a new demand to be peaceful. If all of this can happen in a species that has less access to the tools of reason and goal setting, the possibilities for human change are enormous. Over the past decade data derived from implicit measures of cognition show the unexpected malleability of seemingly rigid mental associations between social groups and attributes (Akalis, Nannapaneni, & Banaji, 2007; Blair, 2002). These results offer a different view of capacity for change, one that engenders greater optimism about behavior that follows more closely in line with intention, at least if the disparity between intention and action is raised to a level of awareness that was previously unavailable.

A topic as complex and as pressing as intergroup conflict in humans needs all the expertise it can get. Understanding the simple in order to grasp the complex is accepted scientific practice, and here we take as a given that we can learn much about the nature of intergroup perception, attitudes, beliefs, and actions from the minds of young children (Aboud, 1988; Aboud, Dovidio, Glick, & Rudman, 2005; Bigler, Jones, & Lobliner, 1997; Killen, 2004). Developmental psychology offers unique insight into the origin of end-state processes and how they develop from simpler components, thus serving as constraints to theoretical models of end-state (adult) social cognition. In particular, developmental analyses can test predictions derived from theories of adult social cognition (Killen & McKown, 2005). For example, as we discuss elsewhere in this chapter, the study of development can address the question of whether the implicit evaluative system is essentially a "slow-learning" system searching the environment for regularities or whether it is an automatic and intrinsic feature of categorization and hence does not rely on a slowly developing history of learning (see Baron & Banaji, 2006; Dunham, Baron, & Banaji, 2006; especially, Dunham & Banaji, under review).

Broadly speaking, adult social cognition often places great emphasis on a particular history of learning (Devine, 1989; Rudman, 2004; Rudman & Goodwin, 2004). Whether in the form of beliefs, preferences, norms, heuristics, or priors, the body of knowledge that makes up the core of adult social cognition assumes that such knowledge is acquired over the course of a lifetime, starting in infancy. Insofar as children attend to and manipulate incoming information in ways that are unique, the prior knowledge underlying adult cognition is often the result of developmentally specific processes. Thus acquisition studies that investigate adult populations may be unable to shed light on either the actual ontogenetic history of acquisition or the structure of that prior knowledge.

Moreover, although our experiences undoubtedly affect our attitudes and preferences today, substantial research has demonstrated that preferences for the familiar (Zajonc, 1968) and for the in-group (Tajfel, 1982) can emerge within a matter of seconds. Therefore, implicit preferences for particular social groups may emerge quite early in development even in the absence of much social learning or numerous personal experiences. We suggest that these considerations make attention to development a necessary part of any complete picture of human cognition as it unfolds in a social world, and intergroup cognition is the topic we address here.

Just as human intergroup conflict has its roots in the phylogenetic record, so, too, do these roots extend into our ontogenetic history. Here we focus on the relatively early development of intergroup cognition, including feelings that reside outside conscious awareness or conscious control. In this chapter we report on a few related threads of work that have emerged in our lab over the past few years.

The Backdrop

We build on the work of others who for decades have documented children's emerging cognitions about social groups through observational studies, classroom interventions, and experimental analyses (Aboud, 1988; Ruble & Martin, 1998; Bigler, Jones, & Lobliner, 1997; Killen, Lee-Kim, McGlothlin, & Stangor, 2002; Levy & Dweck, 1999; Piaget, 1929). From several existing traditions of research we have learned that children's intergroup attitudes tend to be pro-in-group at a very young age and that this bias decreases with age (see Aboud, 1988, for a review). This age-related decline in in-group bias, which is measured with a "self-report" assessment (i.e., the measures require conscious introspective access on the part of the respondent, regardless of whether reporting about the self), is difficult for us to interpret, however. This is because the developmental decline in in-group bias occurs around the time that children become aware that reporting negative attitudes toward others who are different violates social norms and coincides

with increased concern about self-presentation more broadly (Eisenberg, Losoya, & Guthrie, 1997), increased desire to be part of a group (Rutland, Bennett, & Sani, 2004), and an increased desire to avoid the appearance of deviating from group norms (Abrams, Rutland, & Cameron, 2003).

Beginning with the well-known doll studies of Clark and Clark (1947), researchers have charted the emergence of racial evaluation, usually by measuring children's preference for in-groups and out-groups as represented by dolls, pictures, or photographs (e.g., see review by Brand, Ruíz, & Padilla, 1974). More recent research with varied paradigms (Aboud, 1988) has converged on the same conclusions, at least with regard to North American and British majority children: Race bias emerges early, as young as 3 or 4 years of age, and appears to peak in middle childhood, when it then undergoes a gradual decline throughout adolescence (Aboud, 1988; Augoustinos & Rosewarne, 2001; Davey, 1983).

Yet this picture is difficult to reconcile with behavioral developments during this same period, a point brought out starkly by research on school demographics and friendship patterns. Over the same period that expressions of dislike of other racial groups appear to decline, the rate of cross-race interactions and interracial friendships, even in racially mixed schools, declines precipitously (Moody, 2001; Smith, 2003) as schools undergo a dramatic process of self-segregation. And of course, several parallel literatures have firmly established the presence of pervasive race-based discrimination in adults in areas as diverse as housing (Turner, Ross, Galster, & Yinger, 2002), employment (Deitch, Barsky, Butz, Chan, & Bradley, 2003), access to quality healthcare (Williams, Neighbors, & Jackson, 2003), and education (Farkas, 2003). What we have is a prima facie paradox between increasingly positive attitudes and ever more negative (or unchanging) behavior.

A potential reason for this attitude-behavior disconnect is that the inquiry into racial attitudes discussed earlier relied on direct questioning, in which children are asked to report on their own attitudes or make deliberative choices between options, from which their attitudes are inferred. For example, children may be asked to describe their race attitudes or to rate or sort race-based photos. Such methodologies, while a valuable source of insight into children's conscious thinking and reasoning, assume that children are able to honestly report their beliefs and attitudes despite the fact that anonymity can increase candor in adults (Evans & Miller, 1969) and even in children (Rutland, Cameron, Milne, & McGeorge, 2005).

Far more indefensible is the assumption inherent in self-reported statements that human beings, children included, can introspectively report on the contents of their mind—specifically values, attitudes, and beliefs about oneself and others. Adult cognition provides numerous examples of the failure of this assumption, most prominently the well-established dissociation between explicit and implicit forms of memory (see review by Schacter, 1987). More recently, social psychologists have proposed a similar division for other constructs, including self-esteem, stereotypes, and, most important for our purposes, attitudes (Devine, 1989; Greenwald & Banaji, 1995). By using a variety of techniques (see Banaji, 2001), we can now measure less conscious ("implicit") forms of attitude, which relate to a wide range of subtle behaviors such as impressions of facial threat (Hugenberg & Bodenhausen, 2003), subtle verbal behavior and overall friendliness toward out-group members (Dovidio, Kawakami, & Gaertner, 2002), and trait ratings of out-group members (Olson & Fazio, 2004).

By measuring these relatively unconscious aspects of the mind, we can see that implicit preferences underlie acts of everyday discrimination (Bazerman & Banaji, 2004); indeed, in the context of discrimination, Poehlman, Uhlmann, Greenwald, and Banaji (2007) have conducted a meta-analysis that concludes that implicit attitudes are often *better* predictors of discriminatory behavior than are their explicit counterparts. Hundreds of studies of adults over the past 25 years have raised the possibility that implicit social cognition, including intergroup cognition, may drive the paradox observed in children's intergroup behavior discussed earlier.

Experiments on the Development of Implicit Intergroup Cognition

Our understanding of children's emerging intergroup attitudes and stereotypes (like those of adults) has until recently been limited by the available methods—those that have allowed access to self-reports of likes and dislikes (preferences) and beliefs about children's own and other social groups. The limitation of explicit measures (e.g., self-report) is not that they are

worse measures per se; they are simply unable to provide information about aspects of mind that they do not reach. For this reason, our own work makes use of self-report measures, but in a self-conscious manner, in order to help us understand both the divergence between conscious and unconscious cognition and aspects of conscious cognition that are direct representations of preferences and decisions about social groups.

For our purposes, we take the distinction between the conscious and the unconscious to be fundamentally about awareness and control. Unconscious processes are inaccessible through introspection and generally elude our awareness; they may also allow awareness but are difficult to control. Conscious processes afford more opportunity for reflection and even deliberation, and both the processes and their output (e.g., preferences, beliefs, attitudes) permit our awareness and control.

For the past 25 years, indirect measures of cognition have become a staple in the measurement of adult social cognition, and in the last few years they have also been used to study children's social cognition. Using an ambiguous situations task (AST), McGlothlin, Killen, and Edmonds (2005) and McGlothlin and Killen (2006) measured children's interpretations of ambiguous situations as a window into their racial attitudes. McGlothlin et al. (2005) found evidence of the early emergence of such bias among European American children and of its mediation by intergroup contact (McGlothlin & Killen, 2006).

The Implicit Association Test

Other measures have relied on a statement of implicit social cognition as the "introspectively unidentified (or inaccurately identified) traces of past experience that mediate favorable or unfavorable feeling, thought, or action toward social objects" (Greenwald & Banaji, 1995, p. 8). Using the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), Sinclair, Dunn, and Lowery (2005) investigated the relationship between parents' racial attitudes and those held by their children, who were 9 to 11 years of age, and found that such a relationship indeed exists. The degree to which children identify with their parents predicts the strength of the correlation between parents' racial attitudes and those of their child, reinforcing the idea that parents do indeed influence their children.

Using a modified version of the IAT, Rutland et al. (2005) found that, among British children, implicit race bias was present and was not subject to self-presentational concerns, unlike their explicit counterparts. Like others in the field (McGlothlin et al., 2005; Sinclair et al., 2005; Rutland et al., 2005; McGlothlin & Killen, 2006), at the time we began our research we were surprised that implicit or automatic measures of evaluation did not have a greater presence in the developmental literature. This discovery led us to develop and incorporate more implicit measures of bias that, by their nature, are relatively uninfluenced by social desirability concerns.

The IAT measures the relative strength of association between paired concepts (e.g., "white American" and "good"; "black American" and "bad"). In a typical IAT, participants are first asked to classify target stimuli into two categories using two response keys on a computer. For example, participants might be presented with male and female names one at a time and instructed to press the "E" key when a male name is presented and to press the "I" key when a female name is presented. Next, participants are presented with attribute stimuli and similarly asked to classify them into two categories. For example, they might be presented with "good" words (e.g., nice, pleasant) and "bad" words (e.g., mean, unpleasant) and asked to press the "E" key in response to good words and the "I" key in response to bad words. For the next set of trials the strength of the association between the target concepts and the attribute concepts is measured.

Specifically, for half the trials, participants are instructed to classify one target concept (male) and one attribute concept (good) using one response key ("E") and to classify the other target concept (female) and the other attribute concept (bad) using another response key ("I"). For the remaining half of the trials, the pairings change, and participants are now instructed to use one response key to classify the target concept "male" and the attribute concept "bad" and to use the other response key to classify the target concept "female" and the attribute concept "good."

The IAT measures the relative strength of associations between these paired concepts by evaluating the speed with which participants are able to accurately classify these stimuli. The underlying assumption is that the stronger the association between two concepts, the faster and more accurate participants should be at classifying those stimuli when they share the same response key. For more details on the methodology of the IAT, see Greenwald, Nosek, and Banaji

(2003). (If you are interested in a demonstration of the task, visit www.implicit.harvard.edu.)

The Child IAT

In our lab we spent several years developing a child-friendly version of the Implicit Association Test (Child IAT; see Baron & Banaji, 2006, for details). One key challenge was developing a test that young children could successfully complete despite their variable computer skills, reading abilities, and well-known attentional limitations (Piaget, 1929, 1970; Kirkham, Cruess, & Diamond, 2003). With certain changes to the adult IAT we were able to preserve the properties of the test that make it a well-suited measure of implicit association while enabling children as young as age 5 to participate.

We made several modifications to the traditional IAT to make it available to young children. First, we used picture stimuli to represent the target concepts (e.g., pictures of boys and girls to represent gender, pictures of African American and European American children to represent race). Next we used an auditory presentation method to present children with the attribute stimuli. Specifically, children heard prerecorded "good" and "bad" words (e.g., nice, mean, happy, sad) over headphones. This served to reduce artificial differences in response rates and accuracy that might arise for children with varied reading abilities. We also replaced a keyboard with large response switches that facilitate responding in younger children, who typically lack fine motor control. The last major change was to have an experimenter sit with each participant and carefully explain the instructions throughout the task. See Baron and Banaji (2006) for a full description of the methods for the Child IAT.

In our exploration of children's implicit attitudes, we sought answers to three interrelated questions: (1) Can such methods detect implicit attitudes at age 5, the earliest age at which the task can be administered? (2) What is the developmental course of implicit attitudes? (3) Is the developmental course of implicit and explicit intergroup attitudes divergent?

Race Baron and Banaji (2006) constructed a child-friendly version of the IAT, and with it we measured the implicit race attitudes of white 6- and 10-year-olds, and adults in the United States. Child participants were recruited from a predominantly middle-class European American suburb of Boston and were tested in their elementary school. Adult participants were recruited from Harvard University and were tested in our laboratory. Twenty-seven 6-year-olds (mean age = 6;1, females = 13), thirty 10-year-olds (mean age = 10;2, females = 15), and 22 adults (mean age = 19, females = 12) participated in this study. All of the participants were European Americans, and all of the children ($N = 57$) showed evidence of having formed implicit attitudes toward the two groups (gender differences are reported only where we observed them).

Remarkably, the 6-year-olds showed implicit attitudes that favored their racial in-group, and the magnitude of this bias did not differ from that of 10-year-olds and adults, which suggests that, for European American children living in a predominantly white community, implicit bias remains relatively stable throughout development. On the measure of explicit race attitude participants were asked to self-report whom they liked more on a series of forced-choice trials during which a picture of European American children and African American children were presented side by side. With this measure we replicated the well-documented pattern of high expression of in-group bias among 6-year-olds, a reduction in such expression by age 10, and a complete absence of any in-group favoring attitude in adults. This study confirms that implicit intergroup attitudes emerge as early as age 6 and shows that, unlike their explicit counterpart, they remain stable throughout development (Baron & Banaji, 2006). See Figure 6.1.

From such evidence we learned that IAT obtained attitudes can be measured in children as young as 5 years of age (the age of our youngest participants) using a very similar method as that used to measure implicit social cognition in adults (see Baron & Banaji, 2006, for details on how the two versions of the IAT diverge). It also charts the emergence of a disjunction between implicit and explicit attitudes over the course of development, driven by marked change in explicit attitudes with age and the striking invariance of implicit attitudes over the same developmental span (mean IAT scores for each cohort were compared, and no differences were observed).

These preliminary results leave several questions unanswered. The most prominent one, and the one to which we now turn, concerns the generality of these findings. Is the pattern charted in such a study (i.e., that of increasing egalitarianism at the explicit level, coupled with sustained in-group favoring bias at the implicit

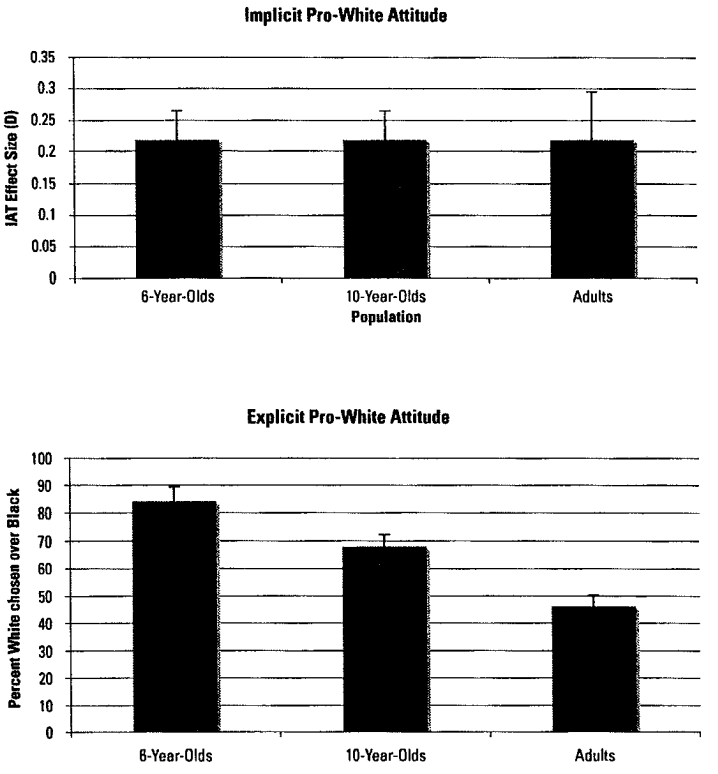

Figure 6.1. Implicit race preference. A positive score indicates a preference for the white social group relative to the black group.

level) common to all racial/ethnic attitudes, or is it confined to the specific population we studied? Is it, for example, a pattern rooted in the particular sociohistorical context of the United States, such that it is limited to the attitudes of the white majority toward a particular stigmatized group? Certainly social learning accounts of attitude formation (Aboud et al., 2005) would assume so. Yet it is also possible that the pattern is indicative of a more general property of the implicit preference system. For example, another possibility, suggested by work in the minimal group paradigm (Tajfel, 1982), as well as more recent work on the pervasiveness of automatic evaluations (Bargh, Chen, & Burrows, 1996), is that merely categorizing others as members of an out-group is enough to incite automatic bias. If so, the stability found in our white American sample might have echoes in many other groups.

Dunham et al. (2006) explored race attitudes in native Japanese children and adults. The rural population tested had little or no direct contact with racial out-groups, which allowed the group to also serve as a first-pass attempt to evaluate the effect of contact on the initial formation of racial attitudes. We measured attitudes toward the in-group category "Japanese" by contrasting it to two out-groups, "white" and "black." The study tested several related hypotheses. First, would the pattern found among European Americans surface in this quite different population half a world away? Second, would these results reveal a central role for direct contact and exposure, which differed so markedly across these two groups? Finally, would the cultural status of the out-group (white vs. black) play a principal role in the development of these implicit race attitudes?

The first author, a native English speaker fluent in Japanese, tested all of the Japanese participants. The children were recruited from a rural school, and the adults, from the surrounding community. Participants completed two child IATs (see Baron & Banaji, 2006, for details), one of which measured their implicit attitudes toward white and Japanese individuals and the other toward black and Japanese individuals.

Overall, our findings reveal substantial similarity between the Japanese and American samples. Japanese 6-year-olds (N = 38, mean age = 6;3, females = 18) and 10-year-olds (N = 34, mean age = 9;10, females = 17), just like American children, showed early and robust implicit in-group bias, which was largely stable over the course of development. Indeed, just as in the American sample, implicit anti-black bias emerged early and did not show developmental change. Anti-white bias also emerged early and retained the same magnitude. However, unlike the anti-black bias, it showed developmental change in undergoing a gradual moderation. Japanese adults (N = 20, females = 8) showed more pro-white bias than did Japanese children (Dunham et al., 2006).

These results confirm what we had learned from the American sample: Early presence of intergroup bias is present in this sample as well. Furthermore, direct contact and exposure do not appear to be necessary conditions for the emergence of implicit bias in view of the fact that these Japanese participants had little to no exposure to racial out-groups. Finally, while social status differences between out-groups indeed influence implicit attitudes, the impact appears to emerge relatively late, certainly not before age 10. On the other hand, explicit biases in the Japanese population show the now-familiar pattern of in-group preference followed by a steady shift toward egalitarianism (i.e., participants self-report equal preference for their racial in-group and a racial out-group), just like the American sample.

With regard to the cross-cultural generalizability of intergroup preferences, we asked two questions. Would the pattern generalize to another majority population in a different cultural context (i.e. the Japanese context), and would the pattern generalize to a *nonmajority* population within the American context? The study in Japan addressed the first question. To examine the second question, we explored the emergence of implicit and explicit race attitudes in a Hispanic American population in Texas (Dunham, Baron, & Banaji, 2007).

Children (N = 234, mean age = 8;7, range = 5;6–12;2, SD = 1;7) were recruited from a bilingual elementary school where Spanish was the first language of most of the students; adults (N = 71, mean age = 20) were recruited from the surrounding neighborhood to match the children on demographic attributes to the extent possible (low SES [social economic status] living in a predominantly Spanish-speaking urban community). Again we examined race attitudes toward two out-groups, the European American majority and an African American minority. This allowed us to examine whether and when the social status of groups implicitly determines partiality for one's own and other groups.

We found that, like European American and Japanese children, Hispanic American children showed early and robust implicit bias that favored the in-group over African Americans. However, a bias toward Hispanics (in preference to European Americans) was entirely absent from the youngest participants we tested (5-year-olds). That is, for the first time we found a case in which early implicit bias *is not present* at this age and occurs only when the out-group was a majority, higher-status out-group, European Americans. Both of these patterns were relatively stable over age. An implicit bias favoring Hispanics (over African Americans) showed a modest increase in strength, but the bias toward Hispanics over European Americans remained absent into adulthood.

These results convey a poignant message about the nature of implicit intergroup relations. While our prior work suggested a relatively automatic tendency to form implicit intergroup bias, that finding must now be confined to majority or advantaged groups: Minority, disadvantaged children, like adolescents (Baron, Shusterman, Bordeaux, & Banaji, 2004) and adults (Nosek et al., 2002) from such groups, do not implicitly favor their own group, thereby revealing the early entry of implicit knowledge of their own group's lesser status. Interestingly, like adults, these children *explicitly* favor their own group. Thus for minority group attitudes toward the majority, the dissociation between implicit and explicit attitudes occurs very early, as early as we can currently test. However, like majority groups, children from disadvantaged minorities show implicit bias against other minority groups, which suggests that, like majority children, they are quite capable of exhibiting an implicit in-group preference.

Gender Unlike the situation with regard to race, ethnicity, religion, and nationality, children have a very different type of experience with gender categories. Decades of research have documented the early emergence of gender bias among toddlers (Katz, 1983; Katz & Kofkin, 1997; Ruble & Martin, 1998) using observational and self-report measures. Research on gender bias in

adults has been conducted with both implicit and explicit measures (Rudman & Goodwin, 2004; Carpenter, 2000).

Contradicting a long-held belief that attitudes toward women are negative, Eagly and Mladinic (1989) showed that negative stereotypes of women notwithstanding, attitudes toward women are positive. In recent years, the effect has been replicated with implicit measures of attitude. Indeed, even on speeded-response latency measures of preference, stronger positive attitudes were revealed toward the categories "female," "women," and "mother" than to "male," "man," and "father" (Carpenter, 2000).

Interestingly, gender preferences reported on explicit measures of preference appear to favor women equally strongly in both males and females. However, in contrast to the data from explicitly reported attitudes, the implicit preference data show a large effect of participant sex: Females show significantly stronger female preference than do males. In other words, females show a strong in-group preference, whereas males do not; instead their data reveal weak preference for female or neutral attitudes depending on the subtype of gender groups specified (Carpenter, 2000).

How are we to think about the development of implicit gender attitudes? Should we expect them to emerge early and follow the same developmental course as their explicit counterparts? Or will they remain relatively stable like implicit race attitudes? Alternatively, given the greater presence of both girls and boys in each others' lives and the gender segregation that occurs in friendship patterns, should we expect that contact will predict the level of in-group bias?

Baron, Dunham, and Banaji (2007) measured European American children's implicit gender attitudes in two studies and found that children 5–12 years of age show strong in-group favoritism on measures of implicit and explicit attitude. However, around the onset of puberty, boys' implicit in-group preference begins to decline until it is either absent or substantially weaker. In contrast, girls at these same ages show relative stability or a slight increase in their levels of implicit in-group preference (Baron et al., 2007).

Importantly, these data reveal similar levels of implicit in-group liking among male and female children. This effect is most pronounced among our youngest samples. The observed pattern of results among the adolescents closely resembles findings from the adult literature on gender attitudes. We are currently exploring the reason for this change in males' implicit attitude but not for females' implicit attitude. One possibility is that the association between males and violence increases around the same years that puberty begins, and this belief (i.e., that males are more violent) causes a reduction in males' implicit in-group preference.

Throughout the studies we have discussed, the data show (a) implicit in-group preference, (b) moderation of implicit out-group attitudes by the status of the group, such that culturally advantaged (European American) groups produce stronger implicit positivity than do less advantaged (African American) groups for Hispanic American children, (c) the presence of social group preferences despite a lack of direct contact (Japanese children show pro-white attitudes [rather than pro-black] in the absence of direct contact with either), and (d) a developmental decline of explicit bias accompanied by the relative stability of implicit bias. Taken together, these results suggest that children's implicit social attitudes emerge by age 5 and are modulated in part by an in-group bias.

Experiments on Preferences for the Lucky and the Advantaged

To develop the abilities that make up social cognition, children need to be able to see demarcations between categories, assign attributes to those categories, and behave along the trajectories those group attributes suggest. By analyzing these processes we can learn about social cognition, the development of moral reasoning (who is good, who is bad, and why), the similarities between adult and child cognition, and the generalization from individuals to social groups and back to individuals.

In the previous section we described an implicit favoring of particular social groups such as those based on race or gender. But one might also ask, is there a general, all-purpose orientation that favors those who are viewed as receiving more (advantaged) or as simply lucky? From such category knowledge, what else flows?

Olson and colleagues (Olson, Banaji, Dweck, & Spelke, 2006; Olson, Dunham, Dweck, Spelke, & Banaji, 2007) conducted a program of research to look at these preferences, and we summarize some of the key results here. The experiments asked for explicitly reported preferences expressed by those who experienced positive and negative events that may or may not have been

within their control. We use the term "preference for the lucky" to refer to a preference for people who experience positive random (lucky) events (e.g., finding $5 on the sidewalk) rather than for those who experience negative random (unlucky) events (e.g., getting splashed by a passing car). Next we describe a series of studies that demonstrates this preference and explain how it may lead to intergroup bias.

Our first demonstration of preference for the lucky was very simple. We presented 32 participants (mean age = 6, SD = 1, females = 18) with short vignettes about targets who had experienced either lucky or unlucky events (Olson et al., 2006). We asked them to indicate how much they liked each one on a smiley-to-frowny-face scale and averaged the ratings for the lucky and unlucky actors. We found that children preferred the lucky to the unlucky and that the magnitude of this preference was quite large (Cohen's d = 1.07).

Within the same study we also included targets who performed good and bad actions that were clearly intentional (e.g., helping a teacher, telling a lie). Such instances allowed us to confirm that, while children assumed that a person who performed an intentional good deed was clearly preferable to one who had merely experienced a lucky event (and likewise for the negative case), the lucky-over-unlucky preference was nevertheless present.

Do children believe that those who experience lucky events are inherently good? That is, that the lucky will engage in intentionally good actions in the future? Conversely, do they believe that those who experience unlucky events are intrinsically bad and will engage in intentional harm (the obverse of Piaget's immanent justice question)?

To test this question, we showed predominantly middle- and upper-middle-class American children 4–12 years of age (N = 58, mean age = 7 years, SD = 2 years, females = 19) pictures of two targets (Olson et al., 2007). One of the targets was described as having experienced a lucky event while another was described as having experienced an unlucky event. We then asked subjects which of these two targets they believed would perform an intentionally good action (e.g., help the teacher) or intentionally bad action (e.g., hit a classmate). Consistent with a generalized preference for the lucky, we found that children believed the lucky target would also be more likely to perform an intentionally good action and that the unlucky actor would be more likely to perform an intentionally bad action. Finding $5 on the street makes one more likely to be a helper of teachers; being hit by a pine cone on the head makes one seem more likely to hit a classmate.

Together these studies show that those who are lucky not only receive children's favorable attitude in the form of greater liking but are also believed to be intrinsically good (i.e., capable of purposive good actions). We found that these preferences exist in children ranging in age from 4 to 12 years with no significant differences across age, surprising as that might seem.

Olson et al. (2006) next tested the generalizability of the effect. Are such actions socially contagious such that they rub off on those associated with the lucky or unlucky target? Studying friends and family, we presented 43 predominantly middle- and upper-middle-class American children (mean age = 6 years, SD = 1 year, females = 21) with members of two social groups of five members each who were designated by an arbitrary feature such as shirt color (Olson et al., 2006). Three members of Group A experienced unique lucky events, and three members of Group B experienced unique unlucky events. The remaining two members of each group were described in a neutral manner (e.g., Jane likes oatmeal).

Children were then shown a sixth new member of each group but were given no information about the new individual other than t-shirt color and were asked, "Which one do you like better?" Children consistently preferred the new member of the lucky group to the new member of the unlucky group. We refer to this phenomenon as *evaluative contagion* to capture the idea that the initial good or bad incident that involved the target in turn affected the evaluation of those who came to be associated with them.

Following the same logic as in the preference studies, we next asked whether children (ages 4–12) not only preferred the associates of lucky people to the associates of unlucky people but also drew inferences about their morality as expressed in predictions about their future actions (Olson et al., 2007). We replicated the prediction study described earlier and this time included a variation: We described a lucky and an unlucky target and then asked which of their siblings was more likely to perform an intentional good or bad event. We found that the predominantly middle- and upper-middle-class participants (N = 78, mean age = 7 years, SD = 2 years, females = 48) were just as willing to say that the sibling of a

lucky actor would perform an intentional good action as they were to say that the lucky actor herself would do so. Similarly, participants expressed the belief that the sibling of an unlucky actor would engage in intentional bad behaviors.

Together these studies show quite clearly the robust nature of preference for the lucky in a variety of settings. Favorable attitudes emerge based not only on the uncontrollable or random events that happen to targets but also on those that happen to someone else who is related. This result acquires deeper meaning when we recognize that good luck and ill fortune are not randomly distributed. Some groups (e.g., poor people and those who are otherwise disenfranchised) tend to experience many more unlucky events (e.g., they may live in regions that are prone to natural disasters). Insofar as children prefer groups that experience more lucky than unlucky events, this preference may serve as a basis for acquiring or maintaining negative attitudes toward those who are already disadvantaged.

Two additional questions presented themselves: At what age does this bias favoring the lucky emerge, and how culturally specific is it? In our first attempt at examining the onset of this preference we included groups of 3- ($N = 16$, 5 female), 3½- ($N = 27$, 15 female), 4- ($N = 27$, 18 female), and 4½-year-old ($N = 16$, 8 female) children to complete a simple preference task (Olson et al., 2007). We presented these preschoolers with pairs of targets, one who had experienced a lucky event and one who had experienced an unlucky event. We then simply asked, "Who's nicer?" Children aged 3½ and older were significantly more likely to think the lucky actor was nicer than the unlucky one, but those aged 3 did not differ from chance, which suggests that, when using this method, luck is not the basis of preference, at least for those who are younger than 3½ years of age.

To test the presence or absence of cultural invariance in a preference for the lucky, we studied a sample of children ($N = 23$, mean age = 5 years, SD = 1 year, females = 10) living in Japan. We chose a Japanese sample because Japanese adults (and those in other East Asian countries) have been shown to exhibit fewer dispositional attributions than Americans (Masuda & Kitayama, 2004; Morris & Peng, 1994). Therefore, since one way to exhibit a preference for the lucky is to make a dispositional attribution for the event (e.g., "He wouldn't have been splashed if he hadn't been dumb enough to stand so close to the edge of the street"), we reasoned that children in a culture less inclined to make quick dispositional attributions might not show this preference.

We found, however, that Japanese children showed this preference for the lucky just as much as American children (Olson et al., 2007). Additionally, another group of Japanese children ($N = 89$, mean age = 5 years, SD = 1 year, females = 50) were presented with this same group preference task and revealed results nearly identical to those of their American peers. A preference for the lucky does not seem to be an exclusively American or even a Western inclination. Instead, it appears in the behavior of children even in countries known to differ in the attributional style of adults (e.g., favoring situational or dispositional explanations, see Masuda & Kitayama, 2004; Morris & Peng, 1994).

In sum, this program of research has demonstrated that (1) children prefer lucky to unlucky individuals, (2) they extend this preference to the family and group members of the lucky and the unlucky, (3) this preference is visible by age 3½ and continues into older ages, and (4) it exists in cultures otherwise known to encourage different ways of thinking about personal responsibility and situational influences. Although this preference for the lucky may seem innocuous, it is exactly such an orientation that may create and maintain social hierarchies and intergroup preferences.

Conclusion

If our species were to be observed by an intelligent other, surely among the features that would stand out about us would be our penchant for drawing group lines. If the residents of East Haven believe that the citizens of West Haven do not line up their garbage cans on the curb correctly, such stereotypes might be comical were it not for the fact that exactly the same process is invoked in conflicts that are far more severe and consequential, when the lines are drawn not between two small towns but between large groups separated on the basis of religion, culture, geography, race/ethnicity, class, or language.

That we are attuned to group differences in excruciating detail such as minor variation in skin color or height is interesting. That we act on them with regard to whom we favor or harm is perplexing. Fine distinctions in a woman's skin color in India will fetch different amounts of dowry, which shows that we place a monetary value on pigmentation. In the United States a

couple of added inches to a man's height may influence his chances to be president or CEO.

Again, we smile when we hear about such irrational human behavior, but it indicates just how easily superficial characteristics come to be noticed and used even if they have no inherent value. If the citizens of West and East Haven ascribe importance to real or imagined differences and characterize them as good and bad, how are we to escape Shia and Sunni, Hutu and Tutsi, southern and northern, Catholic and Protestant, black and white, lower caste and upper caste distinctions?

Surprised as we seem to be about the existence of intergroup conflict when we are not ourselves involved (did the Hutu and Tutsi genocide make sense to anybody else?), we are nevertheless aware that the seeds of such thinking lie in each and every mind and from a very early age. From this understanding we reach a conclusion that the mind sciences have given us recently: that learning to overcome our predilections is not going to be spontaneous or easy and we will be able to do so not by wishing for inspiration from a higher power but by comprehending the role of evolution, cognitive development, social learning, and government and powerful institutions in controlling them.

Robert Putnam's (2006) work has shown that, contrary to what the liberal individual in the United States may wish to believe, the greater the ethnic diversity in a geographic region of the United States, the lower the sense of social trust individuals in that region will experience. Increase in diversity likely leads to more (not less) fear and to more (not less) distrust of one's neighbors (even those who may be from the same ethnic group), and to more (not less) of a hunkering down. Working out the conditions under which contact yields negative social ramifications will be just as important as identifying the conditions under which contact yields positive ramifications (see Pettigrew & Tropp, 2006, for a review of the positive influences of intergroup contact).

Our own work using the IAT as a measure of race bias in adults shows that the greater the ethnic diversity in a region, the greater the anti-black bias in whites. Our research also demonstrates that this depressing result is moderated by close friendships. The greater the number of friends from the out-group, the lower the bias against that group, at least in the case of black-white race relations in the United States. In research we conducted with Phelps's group at New York University (Olsson et al., 2005) we found that classically conditioning white and black Americans to fear members of their own and an out-group led to significant differences in the degree to which fear was extinguished: Fear toward one's own group members died out more easily than learned fear toward members of the out-group. The only individuals who did not distinguish between in- and out-group fear were those (of both groups) who reported having had a romantic relationship with a member of the other group. Much more evidence on these matters is needed, but we are willing to make a bet: Such research is going to continue to show that intergroup love and hate, as well as fear and calm, are deeply rooted orientations of the human psyche. The early expression we see in young children is only further evidence, and it would be ostrichlike to deny this predisposition.

With regard to gender conflict, the data show how hard was the entry of women into America's symphony orchestras. As they entered in larger numbers after blind auditioning became the norm, the orchestras believed themselves to be performing less and less well (Allmendinger & Hackman, 1996). Change in group structures is not always easy, even when the outcome is known (on rational grounds) to be attractive: Expanding the choice of musicians from one half of the population to the entire population *has* to produce a selection of better musicians.

Other stories about human and primate behavior illustrate our capacity to adapt and learn and to do so quickly. As for us, only a few years ago, many if not most Americans could not imagine working alongside a gay coworker. That attitude has transformed significantly in a brief time (Yang, 1997). We cannot say exactly what caused this change, but it is likely that the increasing number of uncloseted gays, the news media (which have shown increased positive and even empowering representations of the group), and laws that enforce gay rights are responsible.

With regard to intergroup conflict, our understanding of both the intercept (baseline) and slope (possibility for change) must be reoriented, based on the evidence. Our sense of the intercept must change if we are to recognize our capacity for harm, even irrational, self-destructive harm. The intercept cannot be set to a hopeful low. On the other hand, the slope (i.e., the ability to change) cannot be assumed to be flat. Change in the matter of intergroup conflict, if imagined appropriately and taking into account the "is" of human nature and responsiveness to social

demands, can point to a different "ought" than one that proceeds without an accurate "is" in place. The job of a science of intergroup behavior is to provide that "is" over and over again in order for it to be persuasive. And here, the need for research to show what the capacities for change are and how to deploy them is in sorry shape and in much need of greater progress.

We conclude by suggesting that the mind and social sciences provide the tools and techniques by which we can understand those processes as they arise early in our development as humans. Insights from studies on the development of implicit social cognition show that we know which groups we belong to early on. Data on infants' remarkable preferences for the familiar demonstrate that, even before a plausible conscious sense of self as a member of a group arises, the ability to tell familiar from unfamiliar creates a basis on which group membership may easily be bootstrapped (Sangrigoli & de Schonen, 2004; Kelly, Quinn, Slater, Lee, Gibson, et al., 2005).

We also know that implicit, in-group-serving biases develop early and importantly seem to largely remain invariant over time, even as explicit attitudes change. Such biases are found in every culture and are highly sensitive to the group's cultural standing. It is remarkable that children who belong to socially disadvantaged groups do not favor their own group as much as those who belong to dominant groups; such a "system justifying" (Jost & Banaji, 1994) result is found in children as young as 5 years of age.

In studying intergroup biases we have also attempted to understand a more general bias observed in those only a little older than 3 years of age: the bias to prefer the lucky to the unlucky. This bias appears to emerge from a very simple conviction (perhaps unconscious) that the good that befalls us is causally driven from within us. To the extent that children observe positive events happening to majority group members and negative events happening to minority group members, they might learn, initially quite implicitly, to view disadvantaged minority groups negatively and the majority advantaged groups positively, even when they are members of the group themselves. Although speculative, this proposal is surely interesting in its simplicity.

Finally, the use of real social groups in our earlier studies often conflated preference for the majority group with preference for the in-group, and thus these data cannot easily divulge the mechanism by which such attitudes originate. Studying the development of attitudes toward novel and minimal groups avoids this misapprehension and others that are inherent to contamination from preexisting knowledge.

By teaching children about novel social groups, we can directly control for their exposure to those groups and provide leverage over the following questions: How much exposure to information about the group is necessary before attitudes form? How many exposures to exemplars are necessary before children will stereotype? Do children weigh the primacy of their exposure to category exemplars more heavily than their recent exposure (or the other way around)? Alternatively, are attitudes and beliefs formed by averaging all exposures to exemplars? Do children's attitudes form differently depending on whether they identify with one of the novel groups? Similarly, by placing children in minimal groups, we can ask whether mere membership is sufficient to generate implicit intergroup bias and whether exposure to information about the groups is even necessary. In a series of ongoing studies in our laboratory we are exploring the answers to these questions.

Acknowledgments We thank Melanie Killen, Sheri Levy, and Virginia Borges for their comments on a previous draft. For discussions of the research reported here, we thank members of the Social Cognitive Development Group, the Laboratory for Developmental Studies, and members of the Understand Seminar in the Department of Psychology at Harvard University. We also gratefully acknowledge the grants from the Third Millennium Foundation, The Mind Science Foundation, and Harvard University.

References

Aboud, F. E. (1988). *Children and prejudice.* Cambridge, MA: Blackwell.

Aboud, F. E., Dovidio, J. F., Glick, P., & Rudman, L. A. (Eds.). (2005). The development of prejudice in childhood and adolescence. In *On the nature of prejudice: Fifty years after Allport* (pp. 310–326). Malden, MA: Blackwell.

Abrams, D., Rutland, A., & Cameron, L. (2003). The development of subjective group dynamics: Children's judgments of normative and deviant in-group and out-group individuals. *Child Development, 74*(6), 1840–1856.

Akalis, S., Nannapaneni, N., & Banaji, M. R. (2007). *Do-it-yourself mental makeovers: Self-*

generated associations shift implicit preferences. Unpublished manuscript, Harvard University.

Allmendinger, J., & Hackman, R. J. (1996). Organizations in changing environments: The case of East German symphony orchestras. *Administrative Science Quarterly, 41*(3), 337–369.

Augoustinos, M., & Rosewarne, D. L. (2001). Stereotype knowledge and prejudice in children. *British Journal of Developmental Psychology, 19*(1), 143–156.

Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger III, J. S. Nairne, I. Neath, & A. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117–150). Washington, DC: American Psychological Association.

Bargh, J. A. (1996). Automaticity in social psychology. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 169–183). New York: Guilford.

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality & Social Psychology, 71*(2), 230–244.

Baron, A. S., & Banaji, M. R. (2006). The development of implicit attitudes: Evidence of race evaluations from ages 6, 10, and adulthood. *Psychological Science, 17*(1), 53–58.

Baron, A. S., Dunham, Y., & Banaji, M. R. (2007). *The developmental origins of the asymmetry in implicit and explicit gender attitudes.* Unpublished manuscript.

Baron, A., Shusterman, A., Bordeaux, A., & Banaji, M. R. (2004, January). *Implicit race attitudes in African American and Hispanic children.* Poster presented at the fifth annual meeting of the Society for Personality and Social Psychology, Austin, TX.

Bazerman, M. H., & Banaji, M. R. (2004). The social psychology of ordinary ethical failures. *Social Justice Research, 17*(2), 111–115.

Bigler, R. S., Jones, L. C., & Lobliner, D. B. (1997). Social categorization and the formation of intergroup attitudes in children. *Child Development, 68,* 530–543.

Bigler, R. S., & Liben, L. S. (2006). A developmental intergroup theory of social stereotypes and prejudice. In R. V. Kail (Ed.), *Advances in child development and behavior* (Vol. 34, pp. 39–89). San Diego: Elsevier.

Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality & Social Psychology Review, 6*(3), 242–261.

Brand, E. S., Ruíz, R. A., & Padilla, A. M. (1974). Ethnic identification and preference: A review. *Psychological Bulletin, 81*(11), 860–890.

Carpenter, S. J. (2000). Implicit gender attitudes. Unpublished doctoral dissertation. Yale University.

Clark, K. B., & Clark, M. P. (1947). Racial identification and preference in Negro children. In T. M. Newcomb & E. L. Hartley (Eds.), *Readings in social psychology* (pp. 169–178). New York: Holt.

Cunningham, W. A., Johnson, M. K., Raye, C. L., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2004). Separable neural components in the processing of black and white faces. *Psychological Science, 15*(12), 806–813.

Davey, A. (1983). *Learning to be prejudiced: Growing up in multi-ethnic Britain.* London: Edward Arnold.

Deitch, E. A., Barsky, A., Butz, R. M., Chan, S., & Bradley, J. C. (2003). Subtle yet significant: The existence and impact of everyday racial discrimination in the workplace. *Human Relations, 56*(11), 1299–1324.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality & Social Psychology, 56,* 5–18.

de Waal, F. B., & Johanowicz, D. L. (1993). Modification of reconciliation behavior through social experience: An experiment with two macaque species. *Child Development, 64*(3), 897–908.

Dovidio, J. F., Kawakami, K., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial interaction. *Journal of Personality & Social Psychology, 82*(1), 62–68.

Dunham, Y., & Banaji, M. R. (under review). *The invariance of intergroup bias across the lifespan.* Manuscript submitted for publication.

Dunham, Y., Baron, A. S., & Banaji, M. R. (2006). From American city to Japanese village: A cross-cultural investigation of implicit race attitudes. *Child Development, 77*(5), 1268–1281.

Dunham, Y., Baron, A. S., & Banaji, M. R. (2007). Children and social groups: A developmental analysis of implicit consistency among Hispanic Americans. *Self and Identity, 6,* pp. 238-255.

Eagly, A. H., & Mladinic, A. (1989). Gender stereotypes and attitudes toward women and men. *Personality & Social Psychology Bulletin, 15*(4), 543–558.

Eisenberg, N., Losoya, S., & Guthrie, I. (1997). Social cognition and prosocial development. In S. Hala (Ed.), *The development of social cognition* (pp. 329–363). East Sussex, UK: Psychology Press.

Evans, W. M., & Miller, J. R. (1969). Differential effects on response bias of computer vs. conventional administration of a social science questionnaire: An exploratory methodological experiment. *Behavioral Science, 14*(3), 216–227.

Farkas, G. (2003). Racial disparities and discrimination in education: What do we know, how do we

know it, and what do we need to know? *Teachers College Record, 105*(6), 1119–1146.

Fazio, R. H., Jackson, J. R., & Dunton, B. C. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality & Social Psychology, 69*(6), 1013–1027.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*(1), 4–27.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality & Social Psychology, 74*(6), 1464–1480.

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality & Social Psychology, 85*(2), 197–216.

Hamburg, D. A. (1991). An evolutionary perspective on human aggression. In P. Bateson (Ed.), *The development and integration of behaviour: Essays in honour of Robert Hinde* (pp. 419–457). New York: Cambridge University Press.

Hamburg, D. A., & Hamburg, B. A. (2004). *Learning to live together: Preventing hatred and violence in child and adolescent development.* New York: Oxford University Press.

Hugenberg, K., & Bodenhausen, G. V. (2003). Facing prejudice: Implicit prejudice and the perception of facial threat. *Psychological Science, 14*(6), 640–643.

Jost, J. T., & Banaji, M. R. (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology, 33*(1), 1–27.

Katz, P. A. (1983). Developmental foundations of gender and racial attitudes. In R. L. Leahy (Ed.), *The child's construction of social inequality* (pp. 41–78). New York: Academic Press.

Katz, P. A., & Kofkin, J. A. (1997). Race, gender, and young children. In S. S. Luthar, J. A. Burack, D. Cicchetti, & J. Weisz (Eds.), *Developmental psychopathology: Perspectives on adjustment, risk, and disorder* (pp. 51–74). New York: Cambridge University Press.

Kelly, D. J., Quinn, P. C., Slater, A. M., Lee, K., Gibson, A., Smith, M., Ge, L., & Pascalis, O. (2005). Three-month-olds, but not newborns, prefer own-race faces. *Developmental Science, 8*(6), F31–F36.

Killen, M. (2004). Race in America: How developmental psychologists can contribute to the conversation on race. *Journal of Applied Developmental Psychology, 25*(1), 127–132.

Killen, M., Lee-Kim, J., McGlothlin, H., & Stangor, C. (2002). *How children and adolescents evaluate gender and racial exclusion.* Monographs of the Society for Research in Child Development (Serial No. 271, Vol. 67, No. 4). Oxford, UK: Blackwell.

Killen, M., & McKown, C. (2005). How integrative approaches to intergroup attitudes advance the field. *Journal of Applied Developmental Psychology, 26*(6), 616–622.

Kirkham, N. Z., Cruess, L., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science, 6*, 449–476.

Levy, S. R., & Dweck, C. S. (1999). The impact of children's static versus dynamic conceptions of people on stereotype formation. *Child Development, 70*(5), 1163–1180.

Masuda, T., & Kitayama, S. (2004). Perceiver-induced constraint and attitude attribution in Japan and the U.S.: A case for the cultural dependence of the correspondence bias. *Journal of Experimental Social Psychology, 40*, 409–416.

McGlothlin, H., & Killen, M. (2006). Intergroup attitudes of European American children attending ethnically homogeneous schools. *Child Development, 77*(5), 1375–1386.

McGlothlin, H., Killen, M., & Edmonds, C. (2005). European American children's intergroup attitudes about peer relationships. *British Journal of Developmental Psychology, 23*(2), 227–249.

Milgram, S. (1963). Behavioral study of obedience. *Journal of Abnormal & Social Psychology, 64*(4), 371–378.

Moody, J. (2001). Race, school integration, and friendship segregation in America. *American Journal of Sociology, 107*(3), 679–716.

Morris, M. W., & Peng, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality & Social Psychology, 67*, 949–971.

Moynihan, D. P. (1994). Personal quotation on WNBC in New York.

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration website. *Group Dynamics, 6*(1), 101–115.

Olson, K. R., Banaji, M. R., Dweck, C. S., & Spelke, E. S. (2006). Children's bias against lucky vs. unlucky people and their social groups. *Psychological Science, 17*(10), 845–846.

Olson, K. R., Dunham, Y., Dweck, C. S., Spelke, E. S., & Banaji, M. R. (under review). Judgments of the lucky across development and culture. Manuscript submitted for publication.

Olson, M. A., & Fazio, R. H. (2004). Trait inferences as a function of automatically activated racial attitudes and motivation to control prejudiced reactions. *Basic & Applied Social Psychology, 26*(1), 1–11.

Olsson, A., Ebert, J. P., Banaji, M. R., & Phelps, E. A. (2005). The role of social groups in the persistence of learned fear. *Science, 309*(5735), 785–787.

Pettigrew, T. F., & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality & Social Psychology, 90*(5), 751–783.

Phelps, E. A., O'Connor, K. J., Cunningham, W. A., Funayama, E. S., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2000). Performance on indirect measures of race evaluation predicts amygdala activation. *Journal of Cognitive Neuroscience, 12*(5), 729–738.

Piaget, J. (1929). *The child's conception of the world.* London: Routledge.

Piaget, J. (1970). Piaget's theory. In P. H. Mussen (Ed.), *Carmichael's manual of child psychology* (pp. 703–732). New York: Wiley.

Poehlman, T. A., Uhlmann, E., Greenwald, A. G., & Banaji, M. R. (2007). *Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity.* Unpublished manuscript.

Putnam, R. D. (2006, September 30). *E pluribus unum: The challenge of diversity and community in the 21st century.* Skytte Prize lecture. Uppsala, Sweden.

Ruble, D. N., & Martin, C. L. (1998). Gender development. In W. Damon, E. William, & N. Eisenberg (Eds.), *Handbook of child psychology* (5th ed.): *Vol. 3. Social, emotional, and personality development* (pp. 933–1016). New York: Wiley.

Rudman, L. A. (2004). Sources of implicit attitudes. *Current Directions in Psychological Science, 13*(2), 79–82.

Rudman, L. A., & Goodwin, S. A. (2004). Gender differences in automatic in-group bias: Why do women like women more than men like men? *Journal of Personality & Social Psychology, 87*(4), 494–509.

Rutland, A., Bennett, M., & Sani, F. (2004). The development and self-regulation of intergroup attitudes in children. In M. Bennett & F. Sani (Eds.), *The development of the social self* (pp. 247–265). New York: Psychology Press.

Rutland, A., Cameron, L., Milne, A., & McGeorge, P. (2005). Social norms and self-presentation: Children's implicit and explicit intergroup attitudes. *Child Development 76*(2), 451–466.

Sangrigoli, S., & de Schonen, S. (2004). Recognition of own-race and other-race faces by three-month-old infants. *Journal of Child Psychology & Psychiatry, 45*(7), 1219–1227.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 13*(3), 501–518.

Simon, H. A. (1957). *Models of man: Social and rational; Mathematical essays on rational human behavior in a social setting.* New York: Wiley.

Sinclair, S., Dunn, E., & Lowery, B. (2005). The relationship between parental racial attitudes and children's implicit prejudice. *Journal of Experimental Social Psychology, 41*(3), 283–289.

Smith, P. K. (2003). Play and peer relations. In A. Slater & G. Bremner (Eds.), *An introduction to developmental psychology* (pp. 311–333). Malden, MA: Blackwell.

Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Review of Psychology, 33,* 1–39.

Turner, M. A., Ross, S. L., Galster, G., & Yinger, J. (2002). *Discrimination in metropolitan housing markets (final report).* Washington DC: Urban Institute, Metropolitan Housing and Communities Policy Center.

Tversky, A., & Kahneman, T. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131.

Uleman, J. S., & Bargh, J. A. (1989). *Unintended thought.* New York: Guilford.

Williams, D. R., Neighbors, H. W., & Jackson, J. S. (2003). Racial/ethnic discrimination and health: Findings from community studies. *American Journal of Public Health, 93*(2), 200–208.

Wolpoff, M. H. (1999). *Paleoanthropology* (2nd ed.). Boston: McGraw-Hill.

Wrangham R. W., & Peterson, D. (1996). *Demonic males: Apes and the origins of human violence.* Boston.: Houghton Mifflin.

Wrangham, R. W., Wilson, M. L., & Muller, M. N. (2006). Comparative rates of violence in chimpanzees and humans. *Primates, 47*(1), 14–26.

Yang, A. S. (1997). The polls—trends: Attitudes toward homosexuality. *Public Opinion Quarterly, 61*(3), 477–507.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality & Social Psychology, 9*(2), 1–27.