

Why Does Ethnic Diversity Undermine Public Goods Provision?

JAMES HABYARIMANA *Georgetown University*

MACARTAN HUMPHREYS *Columbia University*

DANIEL N. POSNER *University of California, Los Angeles*

JEREMY M. WEINSTEIN *Stanford University*

A large and growing literature links high levels of ethnic diversity to low levels of public goods provision. Yet although the empirical connection between ethnic heterogeneity and the underprovision of public goods is widely accepted, there is little consensus on the specific mechanisms through which this relationship operates. We identify three families of mechanisms that link diversity to public goods provision—what we term “preferences,” “technology,” and “strategy selection” mechanisms—and run a series of experimental games that permit us to compare the explanatory power of distinct mechanisms within each of these three families. Results from games conducted with a random sample of 300 subjects from a slum neighborhood of Kampala, Uganda, suggest that successful public goods provision in homogenous ethnic communities can be attributed to a strategy selection mechanism: in similar settings, co-ethnics play cooperative equilibria, whereas non-co-ethnics do not. In addition, we find evidence for a technology mechanism: co-ethnics are more closely linked on social networks and thus plausibly better able to support cooperation through the threat of social sanction. We find no evidence for prominent preference mechanisms that emphasize the commonality of tastes within ethnic groups or a greater degree of altruism toward co-ethnics, and only weak evidence for technology mechanisms that focus on the impact of shared ethnicity on the productivity of teams.

A central question in political science is why some communities are able to generate high levels of public goods—low crime, good schools and health care, adequate sanitation, and clean drinking water—whereas others are not. A wave of recent scholarship, undertaken in both developed and developing countries, has identified ethnic diversity as an important source of variation in these outcomes (e.g., Alesina, Baqir, and Easterly 1999; Khwaja 2002; Miguel and Gugerty 2005). Yet although the negative association between ethnic heterogeneity and public

goods provision is widely accepted—Banerjee et al. (2005) go so far as to describe it as “one of the most powerful hypotheses in political economy” (639)—the specific channel(s) through which this relationship operates remains poorly understood. The literature to date provides a number of plausible hypotheses about the mechanisms that might be at work but little or no research that might permit us to adjudicate among them.

To address the question of why ethnic diversity undermines public goods provision, we identify three analytically distinct families of mechanisms, which we term “preferences,” “technology,” and “strategy selection” mechanisms. Various authors have accounted for the relationship between diversity and the underprovision of public goods by invoking one or more mechanisms from these three families (although without using the labels we introduce here). To identify which of these mechanisms is (are) “doing the work” in linking diversity with the failure of public goods provision, we employ experimental games. We recruit 300 subjects from an area of Kampala, Uganda, characterized by high levels of ethnic diversity and low levels of public goods provision and have these subjects play a series of games, each designed to isolate a different mechanism. Subjects play multiple rounds of each game with randomized matching—sometimes with co-ethnics, sometimes with non-co-ethnics. We then compare patterns of play among co-ethnics and non-co-ethnics across games, interpreting statistically significant differences between each type of pairing within a particular game as evidence for the salience of the mechanism that the game was designed to capture.

The project is novel in several respects. First, our emphasis is not on *whether* ethnic diversity undermines public goods provision, but on *why*. In this respect, the focus of our inquiry runs against the grain in

James Habyarimana is Assistant Professor, Georgetown Public Policy Institute, Georgetown University, 3520 Prospect St., Suite 308A, Washington, DC 20007 (jph35@georgetown.edu)

Macartan Humphreys is Assistant Professor, Department of Political Science, Columbia University, 420 West 118th St., #701, New York, NY 10027 (mh2245@columbia.edu)

Daniel N. Posner is Associate Professor, Department of Political Science, University of California, Los Angeles, Los Angeles, CA 90095 (dposner@polisci.ucla.edu)

Jeremy M. Weinstein is Assistant Professor, Department of Political Science, Stanford University, Encina Hall West, Room 100, Stanford, CA 94305 (jweinst@stanford.edu)

The authors thank Claire Adida, Bernd Beber, Pepine Bulambo, Deo Byabagambi, Elizabeth Carlson, Chris Crabbe, Kenneth Ekode, Nathan Falck, Alex Kunobwa, Sylvester Mubiru, Douglas Musonga, Winfred Nabulo, Ruth Nagawa, Susan Najjuuko, Brenda Nakkazi, Harriet Nambi, Winfred Naziwa, Livingstone Ntensibe, Alex Odwong, Alexandra Scacco, Geoff Sentongo, Simon Ssenyimba, Elizabeth Suubi, Alex Tindyebwa, Sheila Watuwa and Daniel Young for their extraordinary contributions to this project; the Harvard Academy for International and Area Studies, the Russell Sage Foundation, and the Harry Frank Guggenheim Foundation for financial support; Abigail Barr, Robert Bates, Rick Wilson, David Laitin, Kanchan Chandra, Alan Gerber, Don Green, Elisabeth Wood, Sam Bowles, members of the Laboratory for Comparative Ethnic Processes, seminar participants at the University of Michigan, UCLA, UC Berkeley, UC Davis, MIT, Wharton, Northwestern, Oxford, Stanford, Caltech, the EPRC (Kampala), and Yale, and three anonymous reviewers for helpful comments on previous versions of this paper.

contemporary social science research, most of which is concerned principally with establishing causal relationships and only secondarily with examining the various mechanisms that might account for them. Second, in answering the question of why ethnic diversity impedes public goods provision, we do not simply survey the literature for candidate hypotheses to test; instead we identify analytically distinct families of mechanisms to investigate. By examining multiple mechanisms, our approach seeks to identify not only the causal channels that are at work but also those that are not. Finally, in examining these mechanisms we employ experimental games. In doing so, we join a growing group of social scientists—including, increasingly, political scientists (e.g., Bahry and Wilson, forthcoming; Fowler 2006)—who have employed laboratory experiments to study patterns of altruism, cooperation, and trust (for a review, see Camerer 2003). The project reported here advances this research agenda by bringing together in a single empirical strategy many of the most important, recent innovations in behavioral economics. As in the pioneering work of Henrich et al. (2004), we take the experimental laboratory to the field and, as in the work of Bahry and Wilson (forthcoming) and Greig and Bohnet (2006), we randomly draw our sample of subjects from the communities whose ability (or inability) to provide public goods we seek to explain. As in the work of Fershtman and Gneezy (2001), Gil-White (2004), and Whitt and Wilson (2006), we investigate the impact of shared group identities on game play, but we also move beyond this work by using, for the first time in such contexts, subjective indicators of in-group membership.

THREE FAMILIES OF MECHANISMS

There is no well-defined universe of answers to the question of why ethnic diversity might matter for public goods provision. We, therefore, begin our investigation by identifying three analytically distinct families of possible mechanisms. Our identification of these families draws on the game theoretic model of social interaction. In this general model, a social interaction, or game, comprises just three objects: a population (i.e., a well-identified set of actors), a technology (i.e., a set of strategies available to each of these actors), and preferences (i.e., the ways that individuals value the outcomes that result once all the actors select their strategies).¹ The social outcome can then be described as some function of the strategies selected by the actors, typically predicted by theorists using a solution concept. Ethnicity, then, could affect a player's behavior either by changing any of the three primitives or by changing the selection of strategies, conditional on these primitives. Assuming that ethnicity is predicated on existence (that is, we do not treat the population as a function of ethnicity), this simple apparatus leaves us with just three families of mechanisms through which the outcome of

¹ The description we use here is that of a game in normal or strategic form. Games in extensive form use more primitives; those in characteristic function form may use fewer.

the social interaction might be affected by the players' ethnic identities: *preferences* mechanisms, *technology* mechanisms and *strategy selection* mechanisms. These are the three broad families of mechanisms we examine in this paper.

Two major sets of arguments in the literature on ethnic politics emphasize preferences explanations. The first focuses on a *commonality of tastes*. It conjectures that different ethnic groups care about different types of public goods. To the extent that ethnic groups are geographically concentrated, for example, they may have divergent interests over outcomes that have a geographic component, notably the location of public investments (Bates 1973). The different languages and cultures that ethnic groups possess may also provide a source of divergent preferences—for example, for the language of instruction in schools or the religious holidays that should be observed (Miguel 1999). If preferences are correlated with group membership in this way, then ethnic diversity will imply a diversity of tastes, which may cause disagreements about which specific goods should be provided (or where they should be located) that, in turn, lead to their underprovision (Alesina, Baqir, and Easterly 1999; Alesina and LaFerrara 2005).

A second preferences-based explanation focuses not on the nature of the public goods that are to be provided but on the identities of the beneficiaries. This *other-regarding preferences* mechanism can in principle be considered a special case of a commonality of tastes, insofar as the welfare of other players can be viewed as part of the outcome. In practice, however, other-regarding preferences have received special attention in the literature, so we give this mechanism separate treatment here. The argument in this case is that individuals may attach positive utility to the welfare of fellow ethnic group members but no utility (or negative utility) to the welfare of non-group members (Tajfel 1974). That is, people may view themselves as benefiting when fellow group members are made better off but derive no benefit (or feel worse) when members of other ethnic groups experience improvements in their welfare. In the context of such a "taste for discrimination" (Becker 1957), community members may be willing to bear the cost of providing public goods if they believe that most of the beneficiaries will be co-ethnics. Such differential other-regardingness could account for the higher rates of public goods provision that we observe in ethnically homogeneous societies. Arguments of this type are prominent in both the scholarly literature (Poterba 1997; Vigdor 2004) and in journalistic accounts of ethnic behavior as being fundamentally driven by atavistic or primal sympathies (or antipathies) with respect to in-group (and out-group) members.²

² Although these two mechanisms are among the most prominent preferences-based explanations in the literature, they do not exhaust the set. Other explanations might focus on preferences over the *process* of collective action—for example, the pleasure or costs involved with working together with co-ethnics or non-co-ethnics. We thank Elisabeth Wood for pointing this out.

Within the family of technology mechanisms lie explanations that attribute the greater success of public goods provision in ethnically homogeneous communities to the set (toolbox) of strategies for promoting collective action that are available to co-ethnics but not to non-co-ethnics. In our analysis, we focus on two prominent mechanisms of this type, although others exist.³ The first is an *efficacy* mechanism: homogeneous communities may have an advantage in public goods production because they can draw on a reservoir of common cultural material—language, experience, understandings about modes of interaction—that makes it easier for community members to communicate and work together (Deutch 1966; Hardin 1995). This shared cultural material may improve prospects for successful collective action by providing strategy options—for example, the ability to communicate messages—that heterogeneous groups lack (for a similar argument, see Spolaore and Wacziarg 2006). The second technology-based explanation we examine focuses on the mutual *findability* of co-ethnics in social networks. Shared membership in a social network may enable co-ethnics to find, and thus punish, noncooperators (Besley, Coate, and Loury 1993; Miguel and Gugerty 2005). This can create linkages between games (e.g., transforming single-shot interactions into repeat-play games) and, by making strategy options available that do not exist for individuals that are socially isolated from one another, improve the prospects for successful collective action (Bowles and Gintis 2004a; Ghosh and Ray 1996).

The third family of mechanisms posits that individuals will behave differently depending on the ethnicity of the people with whom they are interacting. The logic of such strategy selection mechanisms can be most easily seen in settings with multiple equilibria. Consider, for example, a game in which contributing to a public good is a preferred strategy if and only if one believes that others will also contribute (as in a stag hunt or assurance game). In such circumstances, if all individuals expect that co-ethnics will cooperate and non-co-ethnics will not, then, under the Nash equilibrium solution concept, these expectations will be self-fulfilling. Co-ethnic pairings will indeed cooperate, whereas non-co-ethnic pairings will not. In this case, the ethnic backgrounds of the players can determine the selection of strategies, and thus the outcome, even though it leaves the game itself unchanged.

There are a number of possible strategy selection mechanisms that might be examined.⁴ Here, we focus

³ Two technology mechanisms not examined here but described in Fearon and Laitin (1996) are the greater ease with which co-ethnics may identify each other's type (e.g., their preferences over the outcomes of collective action) and a higher frequency of social interactions among co-ethnics, which can afford greater opportunities for sanctioning co-ethnics (see also Greif 1989, 1994; Ostrom 1990).

⁴ Some treatments of generalized reciprocity suggest that games that from a material perspective are prisoners' dilemmas may be more appropriately modeled as stag hunt/assurance games. As described above, in such settings co-ethnicity could correlate with cooperation simply through beliefs. Another strategy selection mechanism is described in Fearon (1999) in a setting in which ethnic identities can be used to provide focal points that let individuals select among

on just one, the *social sanctioning* mechanism, which we take to be consistent with the most prominent explanation for successful collective action in the political science literature. In repeated prisoners' dilemma games, there exist multiple Nash equilibria, some of which are cooperative and some of which are not. Adherence to a cooperative equilibrium—one in which public goods are produced—relies on expectations that cooperation will be reciprocated and shirking punished. This equilibrium, based on sanctioning, produces a “norm” or “social institution” of cooperation with cooperators and punishment of defectors. In principle, such norms could exist within but not across ethnic groups: if co-ethnics expect that cooperation with co-ethnics will be reciprocated under threat of sanctioning but cooperation with non-co-ethnics not, then public goods provision will be higher in homogeneous communities, where the norm/social institution applies to everyone, than in heterogeneous communities, where it applies only to some potential cooperating partners. The principle that in repeat-play environments different equilibria can be selected conditional on the ethnic identities of partners is embedded, for example, in Fearon and Laitin's (1996) influential account of the existence of norms of interethnic peace. We limit our analysis here to the question of whether individuals condition their strategy selection in this way rather than on the question of how some equilibrium profile of strategies emerges.

The difference between the strategy selection and technology mechanisms is subtle, particularly with respect to the role of sanctioning. The key distinction lies in the fact that technology mechanisms facilitate collective action among co-ethnics by making it possible for them to do something—work together productively, sanction free riders—that non-co-ethnics cannot do, or cannot do as easily. The strategy selection mechanism depends on no such difference. It posits merely that in otherwise identical settings individuals will play different strategies when interacting with co-ethnics than with non-co-ethnics. The strategy selection mechanism does not depend on the existence of some cooperation-facilitating tool that co-ethnics uniquely possess. It focuses instead on the actual choices individuals make, conditional on the available tools for communication, social sanctioning, or whatever. With respect to sanctioning, a technology story might be that everyone would like to punish defectors (a universal norm), but only co-ethnics can (an ethnic technology). In a strategy selection story, everyone can punish defectors (a universal technology), but only co-ethnics will (an ethnic norm).

These three families of mechanisms offer analytically distinct explanations for the negative association between ethnic diversity and public goods provision. The inferential problem, however, is that these different mechanisms yield empirical predictions that are,

multiple equilibria. Axtell, Epstein, and Young (2001) provide a related model in which bargaining strategies emerge as a function of the distribution of markers in a given population.

in many circumstances, observationally equivalent.⁵ Experimental games offer a means of isolating and testing the independent explanatory power of each mechanism.

RESEARCH SITE, SAMPLING, AND EXPERIMENTAL FRAMEWORK

To study the role that these mechanisms play, we take a set of experimental games, typically played in laboratory environments, to the field in Kampala, Uganda.⁶ Kampala is a good site for studying the mechanisms through which ethnic diversity affects public goods provision. The city is not only extremely ethnically diverse but also a place where ethnicity is highly salient in everyday social interactions. Yet although ethnicity matters, the political situation in Kampala today is sufficiently stable and peaceful to permit research on social interactions across ethnic lines. Moreover, the devolution of responsibility for social service provision over the past decade from the central government to the financially strapped elected local councils (LCs) has meant that the supply of many local public goods—including security, garbage collection, and the maintenance of storm drains—has become a purely local affair that depends almost entirely on the voluntary contributions of local community members (Golooba-Mutebi 2003; Onyach-Olaa 2003). Thus, the question of why some communities are able to generate contributions toward public goods and others are not is of real practical consequence in the area we study.

We began the study by conducting a survey of 594 informants at a set of randomly selected sites in Kawempe, the poorest of Kampala's five divisions. In a regression analysis (not shown), we use data from this survey to confirm that the negative relationship between diversity and public goods provision found elsewhere in the literature also exists in urban Kampala.⁷ Apart from allowing us to corroborate this

negative association, the data from the key informant survey also enabled us to select a narrower area within Kawempe from which to recruit our experimental subjects. Because the goal of our project was to identify the mechanisms through which diversity undermined public goods provision (rather than to test whether a relationship exists between diversity and the underprovision of public goods), we deliberately selected a study area that was ethnically diverse and that had generally low levels of public goods provision. The study area was comprised of the four adjacent parishes (LC2s) of Mulago I, Mulago II, Mulago III, and Kyebando. We refer to them in this paper collectively as Mulago-Kyebando.

Our sample of subjects was comprised of 300 individuals randomly selected from Mulago-Kyebando.^{8,9} Whereas the use of random sampling is standard in survey work, it is rare in experimental research of this type. Most studies in behavioral economics employ university students or nonrepresentative samples of volunteers, and the conclusions one can draw from them for outcomes outside the laboratory are limited. There are two major advantages of random sampling. First, by producing a representative sample, it allows subjects to make inferences about other subjects based on their knowledge of the population: they are thus able to form consistent beliefs about the ethnic identities and behaviors of the individuals with whom they are playing. Second it allows us, the researchers, to make inferences from the behavior of our sample to the population from which our subjects are drawn. Matching the subject and underlying populations is especially important given that the outcome under study—how people condition play on the ethnicity of their partners—is not necessarily a universal feature of human behavior, but rather a

the local community had organized efforts in the area of crime prevention and security) on ethnic diversity (measured from 2001 census data) at the level of the 74 LC1s for which we had data. The negative association between ethnic diversity and public goods provision was robust to both the inclusion of controls for migration and wealth (with standard errors clustered at the parish (LC2) in which the LC1 was located) and the substitution of an alternative measure of public goods provision—whether communities have organized collectively to address issues of garbage collection.

⁸ Simple random sampling was used within LC1s. The number of subjects for each LC1, however, was set using targets that diverged modestly from proportionate-to-size in order to oversample the second and third largest ethnic groups. More than 75% of those we contacted agreed to participate in the study—a high rate given that prospective subjects were made aware that full participation would entail attendance at four separate experimental sessions spread out over several weeks. Of those who chose to enter the study, more than 95% attended all four sessions. We attribute this very low attrition rate to the fact that subjects found the games both fun and profitable (largely due to the number of games they played and thus to the number of opportunities they had to earn money). In addition, we contacted subjects the day before each scheduled session to remind them of their appointment.

⁹ An additional 17 players participated in the dictator games, recruited from among the local council chairpersons in Mulago-Kyebando. Our substantive conclusions are unchanged when we exclude these subjects from the analysis. Quantitatively, dropping these cases diminishes our power; nevertheless, we continue to find a significant impact of co-ethnicity on the behavior of egoists in seven of eight specifications.

⁵ As an example, consider the following simple game. Let the population be given by $N = \{1, 2\}$, the strategy sets by $A_i = [\varepsilon, \alpha]$ for some small $\varepsilon > 0$, $\alpha \in [\varepsilon, 1]$ and $i \in N$, and the preferences by $u_i = a_i a_j - \beta(a_i)^2$ for $\beta \in [0, 1]$ for $a_i \in A_i$, $a_j \in A_j$ and $i, j \in N$. Let a player's ethnic identity be given by θ_i . The Nash equilibria are: if $\beta > .5$ (a prisoners' dilemma), then $a_i = \varepsilon$ for all i ; If $\beta = .5$ (a coordination dilemma), then there is an infinite number of equilibria in which both players provide equal contributions, the Pareto optimal equilibrium is (α, α) ; the risk dominant equilibrium is $(\alpha/2, \alpha/2)$. If $\beta < .5$ then the unique Nash equilibrium is (α, α) and hence there is no social dilemma. An example of a case where preferences are structured by ethnicity is: $\beta = 0$ if $\theta_i = \theta_j$, $\beta = .5$ otherwise. Strategy sets (technology) can be structured by ethnicity if, for example, $\alpha = 1$ if $\theta_i = \theta_j$, $\alpha = \varepsilon$ otherwise. Strategy selection could be structured by ethnicity if, for example, players play the Pareto optimal equilibrium whenever $\theta_i = \theta_j$, and play the risk-dominant equilibrium otherwise. For this game, if we find that co-ethnicity is associated with more cooperative outcomes, we cannot know if this is due to a preference, technology, or strategy selection mechanism, or indeed whether some combination of all three is in operation.

⁶ The full protocols for all the experiments described in this paper are available at <http://www.polisci.ucla.edu/faculty/posner/research/hhpw>.

⁷ In our main specification, we regressed a measure of public goods provision (defined as whether, during the last 6 months, residents of

property that may apply in different ways to different populations.

Before playing the experimental games, we collected information from each subject about his or her ethnic group membership, proficiency in first and second languages, and other background characteristics. We also recorded a series of five digital images of each subject, each providing a different level of information that an observer might use to ascertain the subject's ethnic background. In the analyses presented here, we treat all five levels of information equally and distinguish only between situations where players have no information about the identities of their partners and situations where they have some information.

The Public Information Box and General Set-up

Given the empirical strategy adopted in this paper, it was critical to devise a means for subjects to ascertain whether or not the other players in the game were co-ethnics, and to do so without cueing them to our interest in how ethnicity shapes behavior. In three of the games that we describe, subjects interacted face-to-face with one another, so inferences about other players' ethnic backgrounds could readily be made from their appearance, accent, and other visible cues. The other games were played using a computer interface. In these games, players could make inferences about the ethnic backgrounds of the other players from the images of those players that were made available to them in what we call the Public Information Box (PIB). The key attribute of the PIB is that, as its name suggests, the information that it provided about the players in the game was provided publicly. Before each round of each game, all the players in the round were shown the same PIB containing images of all of the players in that round—including themselves—with the images of the players ordered in the same way. Underneath the PIB, each player saw a note indicating which player number he or she was for that round. Figure 1 provides a sample screen shot of a PIB.


Beyond providing the information that could be used by subjects to ascertain the ethnicity of the other play-

ers (an ability that is described in more detail in the next section), the PIB played three roles in the computer-based experiments. First, it made the interaction more realistic by increasing the credibility of the existence of the other players—the importance of which is emphasized by Bohnet and Frey (1999) and Charness and Gneezy (2000). Second, it provided common information: each player was provided not just with information about who the other players were but also with information about what the other players knew about them, and that the other players knew what they knew about the others, and so on. Third, the design of the PIB allowed us to manipulate the anonymity of the players in the game. Compare, for example, Figures 1 and 2. In Figure 1, the middle player's picture is shown to the other players. That player will therefore play the game knowing that the other players can see who he is. But in Figure 2, the middle player's picture is not shown. He still has information about the other players, but he knows that they have no information about him. He will therefore play the game knowing that he is doing so anonymously. Exploiting this manipulation turns out to be extremely valuable for distinguishing behavior motivated by preferences from behavior motivated by strategy.


Although subjects were shown images of (or interacted face-to-face with) the other players in the game, the games were designed to simulate interactions among strangers. Therefore, after viewing the PIB, subjects were asked to report if they knew either of the other players in the round personally. About 5% of all rounds involved subjects that knew one or both of the other players. All results reported in the paper are robust to the exclusion of such rounds.

Each subject played all of the games multiple times but, as they were informed, never played twice with any other player. Furthermore, although players played multiple times (with different partners) they were not given feedback about play until they had completed all games. This limited learning as well as the ability of players to use repeated interaction to establish coordination procedures, norms, or different forms of other-regarding preferences within the context of the game (Crawford and Haller 1990).

FIGURE 1. Public Information Box with Nonanonymous Offerer



Note: Player 2, the offerer, is "seen" by all players. Note that the images used in this figure are for illustration purposes only and are not the images of actual subjects.

FIGURE 2. Public Information Box with Anonymous Offerer

Note: Player 2, the offerer, is anonymous. Note that the images used in this figure are for illustration purposes only and are not the images of our subjects.

Coding Co-ethnicity

The empirical strategy we adopt in this paper depends on our ability to distinguish interactions among co-ethnics from interactions among non-co-ethnics. A simple rule—analogue to the practice employed in most studies of cross-race, -gender, and -ethnic interactions in the experimental literature—would be to code as co-ethnics any pair of players that identified themselves as belonging to the same ethnic category in our pre-experiment questionnaire. Using this rule, we generate what we term a “benchmark” measure of co-ethnicity. However, this benchmark measure runs into the problem that the way a person identifies him- or herself may not correspond with the way he or she is perceived by others. Thus, if two subjects who identified themselves in the pre-experiment questionnaire as members of group *X* were paired in a game, then, under the benchmark measure, this would be coded as a co-ethnic pairing. But if each subject believed that the other was not really a member of group *X*, then they would each behave as if they were playing with a non-co-ethnic, and our inferences about the role of ethnicity in their behavior would be exactly wrong.

To overcome this problem, we employed a relatively straightforward exercise to generate a “subjective” measure of co-ethnicity based on how our subjects *perceived* the ethnic backgrounds of the players with whom they were interacting in the computer-based games. After all the experimental games had been played (and the danger of priming subjects to ethnicity had passed), we showed our subjects a series of images of other subjects (in most cases, players they had been randomly matched with in the computer-based games they had played earlier) and invited them to guess the ethnic identities of the people whose images they were shown. To incentivize the guessers, correct guesses (defined as guessing how the person had identified him or herself in the pre-experiment questionnaire) were rewarded with a small payment. To ensure that everyone had the same prior beliefs about the distribution of ethnic groups in the sample of images, we told the subjects that the ethnic demography of the sample population matched the ethnic demography of Mulago-Kyebando,

and we read aloud a breakdown of shares of the major ethnic groups in Mulago-Kyebando based on 2001 census figures.

We collected data on a total of 15,265 guesses by 274 different subjects. Overall, we found that individuals were able to correctly identify others only about 50% of the time—a result that underscores the inferential problem we would have faced had we defined co-ethnicity based on self-reported identities alone.¹⁰ We used the results of the identification exercise to generate, for every information level, an estimate of the likelihood that an individual of group *A* believes that an individual of group *B* is a co-ethnic.¹¹ The resulting measure (ranging from 0 to 1) provides a measure of “subjective co-ethnicity.” For all games that make use of the PIB, we report results using both the benchmark and the subjective measures of co-ethnicity.

We also parse the results in yet another way to reflect the fact that in Uganda, as elsewhere, ethnic categorization may take place on multiple levels (Mozaffar, Scarritt, and Galaich 2003; Posner 2005) and that we, the researchers, do not know *ex ante* what the salient dimension of cultural cleavage may be for a given interaction. To deal with this issue, we present the results of all games in terms of co-ethnicity defined by shared ethnic group membership and by an additional, broader notion of co-ethnicity based on the region of origin of these ethnic groups (i.e., Center, East, North, West).

¹⁰ This figure is significantly lower than the identification rates of 71% to 89% (depending on information level) found among Los Angeles-based undergraduates participating in a similar exercise (Habyarimana et al. 2004). Such results are difficult to compare across populations, however, due to the differing underlying complexity of the ethnic demography and to the different fine- or coarse-grainedness of the ethnic categories used.

¹¹ Note that by using this rule, an individual *i* is coded as a subjective co-ethnic of individual *j* if individual *j* believes that *i* would code herself in the same group as *j* codes herself. A stricter definition would require that *j* codes *i* in the same group as *j* codes herself under *j*'s own (rather than under *i*'s) classification criteria. Our requirement for a single criterion of “correct” placement to use as a basis for allocating rewards to players precluded us from generating this more precise measure of subjective co-ethnicity.

Examining Play across Types

One of the most important contributions of behavioral economics has been to challenge the neoclassical view that preferences can be viewed entirely as a function of material rewards. Recent research has shown that whereas some individuals behave in ways that are consistent with neoclassical assumptions about the maximization of material gains, others do not: they give away money in anonymous dictator games; they reject low offers in ultimatum games (even at considerable cost to themselves); and, in public goods games, they respond to cooperative overtures even when they would benefit materially from defection and cannot be sanctioned. Ostrom (2000) summarizes this literature by concluding that "a central finding is that the world contains multiple types of individuals, some more willing than others to initiate reciprocity to achieve the benefits of collective action." Recognizing the existence of these different types has important implications for empirical and theoretical work (e.g., Bolton and Ockenfels 2000; Bowles and Gintis 2004b; Fehr and Schmidt 1999). If players of different types respond in different ways to a given treatment, then an aggregation problem may arise. Average behavior across types may mask systematic features of play taking place within them.

To avoid this pitfall, we distinguish in our analyses between two different types of players: those we term "egoists" (whose behavior is consistent with preferences in keeping with the neoclassical model) and "nonegoists" (who exhibit higher levels of general altruism).¹² To classify players into these categories, we use their play in a version of the dictator game (described in more detail below) in which they are provided with two 500 USH coins and asked to allocate the coins among themselves and two partners, with no one player (including the offerer) permitted to receive both coins. A subject who always employs the most selfish strategy available (i.e., he or she keeps one of the two coins for him- or herself in *all* rounds of the game) is coded as an egoist. Otherwise, the subject is coded as a nonegoist. This coding rule yields 124 egoists in our sample (40%) and 182 nonegoists (60%). In the results discussed below, we explore how egoists and nonegoists respond differently to co-ethnic and non-co-ethnic partners.

RESULTS

Preference Mechanisms: Commonality of Tastes and Other-Regarding Preferences

The first preferences mechanism discussed previously focuses on the commonality of tastes within ethnic groups. To probe its plausibility, we use simple survey techniques to test whether preferences over public goods outcomes do in fact correlate with ethnic

¹² In some models, a player's type is defined over behavioral patterns (Bowles and Gintis 2004b); in others they are defined over preferences (Güth and Kliemt 1998; Ostrom 2000). We focus on preferences here because our aim is to generate inferences across different games based on a single definition of types.

group membership. We examine two types of survey questions. First, to what types of public goods do individuals attach the highest priority (security, drainage maintenance, or garbage collection)? Second, how should these goods be provided? For example, should private or public means be used? These issues—the prioritization of projects and the manner of their provision—were identified by community members in interviews and focus group discussions as being among the most salient concerns in Mulago-Kyebando. Although we confirm empirically that there is considerable diversity in attitudes toward these issues, we emphasize that the issues were selected for their general salience and thus may not necessarily represent the most ethnically divisive issues in Mulago-Kyebando.

Our results are reported in Table 1. Using ordinary least-squares regression with dummy variables for each of the major ethnic groups in our sample, we find little evidence that ethnic group differences are associated with differences in either preferences for particular public goods or opinions about how they should be provided. Although the results point to some small differences in opinion, only four (of 54) ethnic group dummies are statistically significant at conventional levels. Moreover, an *F*-test fails to reject the null hypothesis of no systematic variation across ethnic groups for all six questions. Finally, the four group-specific dummies that show up as significant are rendered insignificant across specifications when fixed effects for location are included in the analysis—suggesting that what looks like a (weak) finding with respect to membership in a particular ethnic group is being driven by factors specific to the area in which members of that group happen to predominate (e.g., low-lying terrain generating a concern for well-maintained drainage channels). In short, there is little empirical support in our data for the argument that ethnic groups possess correlated preferences over outcomes in Mulago-Kyebando.

The second preferences mechanism proposes that higher rates of public goods provision in homogeneous communities stems from high degrees of other-regardingness within ethnic groups and the willingness this generates for people to contribute to collective ends when they believe the beneficiaries will be co-ethnics. To test this mechanism, we had subjects play a version of the standard dictator game in which the offerer is anonymous and the receivers' identities are known (Kahneman, Knetsch, and Thaler 1986). In such a game, any observed difference in patterns of play between situations where the offerer and receiver are co-ethnics and where they are from different ethnic groups can be attributed to differential levels of altruism toward in-group and out-group members (Fershtman and Gneezy 2001).

Each round began with subjects (who in this game only play the role of offerer) seated in front of a laptop computer. The screen of the computer showed a PIB containing images of the two other players (the receivers) and a dummy for the subject (as in Figure 2). In front of the computer were three ballot boxes, each located directly below one of the pictures in the PIB. Subjects were given ten 100 USH coins (about 60 cents,

TABLE 1. Variation in Policy Preferences across Ethnic Groups

	How Public Goods are to be Provided					
	First Priority for Public Goods Provision			Preference for fee based garbage collection over free but lower quality provision (1–4)	It's better not to have to pay anything or to volunteer for patrols, even if that means security is low	It's better to have well-maintained drainage channels, even if we have to make contributions of money or labor
	Drainage	Garbage Collection	Security			
Banyankole	-0.10 (0.10)	0.23 (0.13)*	-0.13 (0.12)	0.08 (0.18)	-0.36 (0.19)*	0.15 (0.16)
Bagisu	-0.20 (0.19)	0.00 (0.23)	0.20 (0.22)	-0.07 (0.33)	0.38 (0.34)	-0.01 (0.29)
Bakiga	0.05 (0.13)	0.10 (0.15)	-0.15 (0.15)	0.12 (0.23)	0.17 (0.25)	-0.45 (0.22)**
Banyarwanda	-0.09 (0.14)	0.05 (0.17)	0.05 (0.17)	-0.21 (0.23)	0.25 (0.24)	-0.23 (0.21)
Basoga	0.10 (0.14)	-0.10 (0.17)	0.00 (0.16)	0.34 (0.29)	0.13 (0.32)	0.15 (0.26)
Batoro	-0.00 (0.14)	0.10 (0.17)	-0.10 (0.16)	0.10 (0.23)	0.26 (0.23)	0.08 (0.20)
Banyoro	-0.00 (0.19)	0.20 (0.23)	-0.20 (0.22)	-0.11 (0.28)	0.28 (0.29)	-0.09 (0.25)
Iteso	0.13 (0.17)	0.10 (0.21)	-0.23 (0.20)	0.36 (0.33)	-0.04 (0.34)	0.28 (0.29)
Bafumbira	0.34 (0.13)***	-0.13 (0.16)	-0.22 (0.15)	0.05 (0.23)	0.05 (0.24)	-0.16 (0.21)
Constant	0.20 (0.04)***	0.40 (0.05)***	0.40 (0.05)***	3.21 (0.08)***	1.62 (0.08)***	3.29 (0.07)***
Observations	185	185	185	238	236	235
Adj R-squared	0.01	-0.01	-0.01	-0.02	0.00	0.00
F-Statistic	1.29	0.71	0.75	.49	1.11	1.04
(p-value)	(0.25)	(0.70)	(0.67)	(0.88)	(0.35)	(0.41)

Notes: Ordinary least squares regression. Analysis is limited to ten largest groups. Base category is Baganda. Robust standard errors in parentheses. ***Significant at 1%; **significant at 5%; *significant at 10%.

approximately equal to the per-capita daily income in Uganda) and asked to divide this sum between themselves and the two other players in any way they pleased. They were told to put the amount that they wanted to keep directly into their pocket and to put the amounts that they wanted to allocate to each of the other players into envelopes and to deposit the envelopes into the ballot boxes located directly below the players' pictures. Subjects were told that the envelopes would be delivered to their intended recipients at a later date, which they were. An enumerator manipulated the computer to show the PIB for the given round and handed the subject the money, but then stepped away and waited behind a screen while the subject completed his or her allocation. When the subject was finished making the allocation, he or she signaled the enumerator, who returned from behind the screen and set up play for the next round. In a second version of the game, subjects were given two 500 USH coins which they were asked to allocate among themselves and the two receivers. They were instructed, however, that no player (including themselves) could be awarded both of the coins. The 500 USH denomination game therefore forced subjects to discriminate (although they were

free to discriminate against themselves and to treat their two partners equally)—hence, we refer to this version of the dictator game as the “discrimination game.”¹³

Each subject played multiple rounds (an average of 2.7) of each of the 100 USH and 500 USH denomination games. In all, we have data from 801 rounds (and 1,602 individual choices) in the 100 USH denomination games and 782 rounds (1,564 individual choices) in the 500 USH denomination game. In the 100 USH denomination game, subjects exhibited a substantial degree of altruism. The modal strategy, employed in 25% of the

¹³ We instituted various checks to ensure that our subjects understood the games they were playing. Most importantly, prior to beginning play, subjects were tested on their comprehension of the rules of the game and the set of strategies that were available to them. Subjects that failed this test were given additional instruction until they could explain the game on their own. In addition, we organized a “back-translation” of the games in which an educated Ugandan with no connection to the project met with a group of our subjects and tried to elicit from them sufficient information about the various games they were playing that he could describe the details of the games back to the experimenters. The success of this back-translation exercise gave us confidence that our subjects understood the underlying behaviors that each game sought to assess.

TABLE 2. Average Offers in the 100 USh Dictator Game with Anonymous Offerer

	Benchmark Co-ethnicity		Subjective Co-ethnicity	
	Ethnicity	Region	Ethnicity	Region
In-group member	-11 (11.20)	-10 (9.55)	-9 (15.41)	-4 (15.64)
Egoist	-99 (11.23)***	-100 (11.77)***	-101 (12.22)***	-99 (14.59)***
Egoist & In-group member	1 (16.41)	3 (13.70)	13 (23.63)	3 (24.17)
Observations	1174	1174	1118	1118
R-squared	0.23	0.20	0.22	0.19

Notes: ***Significant at 1%; **significant at 5%; *significant at 10%. Robust standard errors in parentheses. The cells report coefficient estimates from an OLS regression with the offer as the dependent variable. The benchmark ethnicity samples exclude observations in which enumerators indicated that the subject had problems comprehending the exercise logic and observations with no information on the receiver. The subjective ethnicity sample is restricted to observations in which receivers are identified in the identification exercise. Disturbance terms are clustered for each player across all of his or her games and fixed effects for the ethnic group of the offerer are included.

rounds, was to retain 400 USh and to allocate 300 USh to each of the other players. The next most common strategy was to keep 600 USh and to allocate 200 USh to each of the other players (21% of rounds). In the vast majority of allocations, subjects appeared to adhere to the norm that the two receivers should be treated equally. On average, subjects retained 540 shillings and allocated 230 shillings to each of the other players. The modal strategy in the 500 USh denomination game (played in 73% of rounds) was to keep one 500 USh coin and allocate the other to another player. Nonetheless, in 23% of the rounds, subjects allocated both coins to the other players. These offers can be compared with a baseline strategy of random allocation, under which subjects would keep one coin two thirds of the time.¹⁴

We now turn to the question of whether subjects displayed different degrees of other-regardingness toward co-ethnics and non-co-ethnics. Our first finding is that, although we observe high levels of altruism in the 100 USh denomination game, we find no evidence that this altruism was directed more at in-group members than at out-group members. The first row of Table 2 reports determinants of average offers under four different definitions of shared group membership (benchmark versus subjective co-ethnicity, with co-ethnicity defined over ethnic groups and region). In all four, we find that the difference in offers to co-ethnics and non-co-ethnics cannot be distinguished from zero. Here, and for all analyses of the dictator game, we employ specifications that include fixed effects for the ethnic group of the offerer to ensure that the co-ethnic effect is not driven by an ethnic group main effect (which might happen if members of one group were generally more generous and, because of the size of their group, were more likely to be in co-ethnic pairings). We also cluster

¹⁴ Because the discrimination game is played with two partners for each offerer, we stack the data and code a dependent variable that captures whether a given receiver was favored. In stacking the data this way, we double the number of observations in our regressions. Our results, however, are robust to treating each game as a single game.

disturbance terms by the offerer to account for the fact that observations across games for a given player are not independent.¹⁵ We conclude that there is no evidence that subjects in this context exhibit any taste for discrimination along ethnic lines.

Table 2 also explores differences in patterns of play between egoists and nonegoists (row 2). Egoists give approximately 100 USh less to each partner than do nonegoists. Hence, although our definition of egoists is based on play in the discrimination game, the categorization strongly predicts how subjects play in the 100 USh dictator game as well. In row 3 of Table 2, we revisit our main (non-)result regarding differential other-regardingness across ethnic lines, this time restricting the analysis to egoists. The (non-)finding holds: we still find no evidence that subjects are any more altruistic toward in-group members than out-group members.

It is possible that the (non-)finding with respect to ethnic discrimination might be due to the high degree of general inequality aversion among our subjects and the difficulty, given such inequality aversion, of detecting differential altruism in the 100 USh denomination game. Because the 500 USh denomination game makes equitable distributions more costly, it is likely to be more sensitive to even slight preferences for the welfare of in-group members over out-group members. Given the goals of the exercise, we restrict our analysis of the 500 USh denomination game to rounds in which a player was playing (or believed he or she was playing) with one co-ethnic partner and one non-co-ethnic partner, and in which he or she also elected to

¹⁵ We do not employ individual level fixed effects (or a within subjects design) as we want to examine also the impact of player type (egoist, nonegoist) in our baseline specifications. In addition to the controls already described, the (non-)findings reported here are also robust to the inclusion of ethnic group fixed effects for the receiver and individual fixed effects for both the offerer and the receiver, which gives us confidence that the results are not driven by how particular ethnic groups are treated on average or by the play of specific individuals. The results are also robust to the inclusion of a battery of controls for age, education, income, gender, and other characteristics (which, by design, are uncorrelated with the experimental treatment).

TABLE 3. Discrimination Rates in the 500 US\$ Dictator Game with Anonymous Offerer

	Benchmark Co-ethnicity		Subjective Co-ethnicity	
	Ethnicity	Region	Ethnicity	Region
In-group member	-0.03 (0.10)	0.06 (0.08)	0.08 (0.16)	0.08 (0.14)
Observations	222	354	122	160

Notes: ***Significant at 1%; **significant at 5%; *significant at 10%. Robust standard errors in parentheses. Discrimination game analysis is limited to games in which players play with exactly one in-group member and one out-group member and discrimination is observed. Each column reports the marginal effect of co-ethnicity/co-region from a probit estimation. Group fixed effects are included and disturbance terms are clustered across the individual actions of offerers in a given game.

discriminate.¹⁶ In effect, we condition upon egoism and ask: when facing a co-ethnic and a non-co-ethnic, is a player more likely to favor the co-ethnic? In Table 3, we report the marginal effect of co-ethnicity on the likelihood that players discriminated in favor of in-group members. As in our analysis of the 100 US\$ denomination game, we find no evidence of any co-ethnic effect. Even when players elect to discriminate and must choose between a co-ethnic and a non-co-ethnic, they exhibit no propensity to act along ethnic lines.

TECHNOLOGY MECHANISMS: EFFICACY AND FINDABILITY

Whereas the other-regarding preferences mechanism could be examined with a version of a standard experimental game that we took "off the shelf," testing the technology mechanism required that we develop a set of new games. The first game we study, the puzzle game, was used to examine the efficacy mechanism—the idea that, from a technological point of view, co-ethnics simply work better together. The game, played face-to-face rather than via a computer interface, rewards players based on their ability to complete a joint task in which effective communication is a critical determinant of success.

The game was played with a four-piece jigsaw puzzle (see Figure 3). Before play began, subjects were given the opportunity to practice assembling a sample puzzle. They were then randomly divided into pairs and seated on either side of a small table. A low partition between the subjects prevented them from seeing what their partner was doing with his or her hands but permitted communication over the top of the partition. Each player was given two of the four pieces of the puzzle and instructed to work together with his or her partner to complete the full puzzle.

To solve the puzzle, the players each had to work together to decide, based on information they held collectively but not individually, which of their two pieces should go on top and which should go on the bottom. Thus, in the example provided in Figure 3,

¹⁶ In the subjective co-ethnicity analysis, a player was coded as believing that he or she was facing one co-ethnic partner and one non-co-ethnic partner if the difference in his or her estimated beliefs that each of the two partners was a co-ethnic exceeded .5.

the players had to communicate to make sure that, if Player 1 put piece *A* on top and piece *B* on the bottom, Player 2 would place *C* on top and *D* on the bottom. Players were permitted to talk to their partners but they were not allowed to show their partner their pieces. The players were given three minutes to complete the task, at the end of which an enumerator checked to see if the two halves of the puzzle matched. If they did, the players were paid 1,000 US\$ each. Subjects played the puzzle game three times, each time with a different (randomly assigned) partner, and each time with a slightly different puzzle. One of the three puzzles was "easy" (ordering the pieces randomly generated a 50% chance of solving the puzzle correctly), and two were "hard" (the success rate from a random ordering of pieces was 25%).

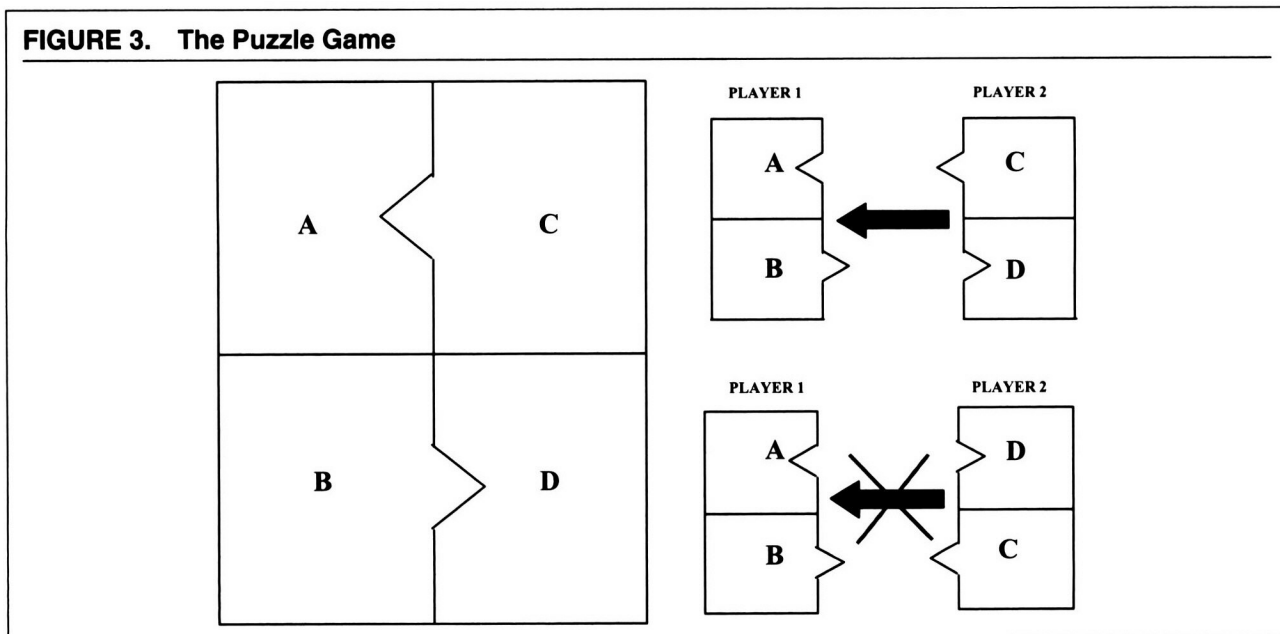
Subjects found the puzzle game to be more difficult than we anticipated. Just 40% of pairs completed the puzzle successfully, although we can reject the null hypothesis that success rates are no better than chance (the associated *p*-value for a one-sample *t*-test is .006). Table 4 summarizes the success rates in terms of the in-group or out-group characteristics of the pairings. Although there is weak evidence that co-ethnic pairings succeeded at a higher rate in completing the task, the difference between co-ethnic and non-co-ethnic pairings (defined by shared ethnicity, regional background, or language group membership) is not statistically significant. The absence of in-group/out-group differences is robust to controlling for round number (i.e., whether it was each player's first, second, or third puzzle). The co-ethnic effect, however, approaches statistical significance (*p*-value of .10) when

TABLE 4. Success Rates in Puzzle Game

	Non-Co-Ethnic Pairing (N)	Co-Ethnic Pairing (N)	Difference (standard errors)
Ethnicity	0.38 (271)	0.48 (73)	0.10 (0.07)
Region	0.40 (227)	0.42 (117)	0.02 (0.06)
Any Common Language	0.45 (20)	0.40 (324)	-0.05 (0.11)

Notes: Table presents results of a two-sample *t*-test. ***Significant at 1%; **significant at 5%; *significant at 10%.

FIGURE 3. The Puzzle Game



we include fixed effects that capture the difficulty of the puzzle.

We were surprised that shared ethnicity did not exert a powerful effect in this highly communication-dependent task. However, our confidence in the weakness of this finding was bolstered by our examination of a second game (results not reported) in which one player, who had been taught how to open a combination lock, was tasked with instructing a second player, from a distance, on how to open the lock after it had been affixed to a chest containing a reward for both players. Success rates in the lockbox game averaged around 60%, but there was no evidence of a difference in the success rates of co-ethnic and non-co-ethnic pairings.

A second variant of the technology mechanism emphasizes how shared group memberships facilitate public goods provision by making individuals more “findable” and thus easier to sanction if they defect from their obligations or commitments (which, in turn, makes them more likely to contribute in the first instance). To test this causal channel, we designed a game that would permit us to measure the extent to which co-ethnics are able to benefit from the existence of a shared social network. To put the “findability” aspect of the network front and center, the task involved finding a stranger.

We began the network game by randomly selecting 148 Mulago-Kyebando residents from outside of our regular subject pool. These were the “targets” that we would send our subjects to try to find. We collected a small amount of background information about each target, including the person’s ethnicity and birthday. We also took the target’s photo with a digital camera. We told the target that, within the next two weeks, someone might come looking for them and that, when they did, the person would ask them for a message. The message was simply the target’s birthday (or, in

case the target did not know his or her birthday, a parent’s name). We asked the targets to provide this message to the person, and gave them 3,000 USh for their participation.

Then, in waves of four or five per day, we sent out 148 “runners,” randomly drawn from our regular subject pool, to find the targets. They were shown the target’s photograph and given a sheet of paper with the target’s name and the parish (LC2) in which the target resided. We instructed the runners to phone us when they found the target and to tell us the target’s message. This provided a check that they had actually found the person they had been sent to find. Runners were given 5,000 USh to defray the cost of hiring a *boda boda* (motorcycle taxi) to take them around the parish, and to pay for the phone call. We told the runners that we would pay them 20,000 USh—a very large sum, equal to about \$12, or more than a half of Ugandan per capita monthly income—if they managed to find the target in three hours and that this sum would decline at a rate of approximately 1,000 USh per hour until it reached zero in 24 hours. This provided runners with a strong incentive to find their assigned targets, and to do so as quickly as possible.

A relatively large number of runners (49 of the 148, or just over 33%) managed to find their targets. More important for our purposes, the success rate among runners whose targets happened to be members of their own ethnic group was significantly higher (43.1%) than that among runners whose targets were from other ethnic groups (27.8%)—see Table 5. This finding is consistent with the hypothesis that ethnic groups constitute social networks that make fellow group members more findable. It also resonates with what several successful runners told us in their postgame debriefs about the strategies they pursued to find their targets. Intriguingly, however, being sent to find a target from the same region had no positive impact on the likelihood

TABLE 5. Success Rates in Network Game

	Non-Co-Ethnic Pairing (N)	Co-Ethnic Pairing (N)	Difference (standard errors)
Ethnicity	0.28 (97)	0.43 (51)	0.15* (0.08)
Region	0.30 (86)	0.38 (61)	0.08 (0.08)

Notes: Table presents results of a two-sample *t*-test. ***Significant at 1%; **significant at 5%; *significant at 10%.

of success. This suggests that the scale at which identity groups are defined is important for determining the relevance of networks. At least in Mulago-Kyebando, the social networks among people who come from the same region would appear not to provide the same advantage in tracking down fellow group members as the social networks built around shared ethnic group membership more narrowly defined.

STRATEGY SELECTION MECHANISMS: SOCIAL SANCTIONING

The strategy selection mechanism we examine attributes higher levels of public goods provision in homogeneous communities to the existence of a norm within ethnic groups that cooperation among co-ethnics should be reciprocated and defection sanctioned. To test for the existence of such a norm—and for its differential strength among co-ethnics and non-co-ethnics—we explored what happened to the offerers' behavior in the dictator game described above when he or she was no longer anonymous (the PIB shifts from the one depicted in Figure 2 to the one depicted in Figure 1). In such a setting, the offerer's behavior can be interpreted as a product of his or her altruism toward the receiver (as it was in the anonymous version of the game) and his or her concern about being seen to violate a social norm requiring cooperation. This latter concern is irrelevant when the offerer is anonymous, but it becomes potentially highly important when the offerer is seen and (in principle) can be sanctioned if his or her actions breach the norm against failing to contribute. To evaluate whether or not such a norm exists, we can therefore compare patterns of play in dictator games where the offerer is and is not seen. To evaluate whether the norm exists more strongly for co-ethnics than non-co-ethnics—the central claim of the strategy selection mechanism that we seek to test here—we can compare the results of anonymous and nonanonymous dictator games played among co-ethnics and non-co-ethnics.

Our decision to interpret the dictator game where the offerer is seen as a situation where the offerer can be punished for the violation of a social norm requires explanation, since the dictator game is a one-shot interaction that, by construction, does not allow for the possibility of punishment within the context of the game. Our logic, following Hoffman, McCabe, and Smith (1996), is that players "are accustomed to

operate in an environment in which there is ongoing social interaction" and as a result "may be concerned about the extent to which their decisions have post-experimental consequences, or that others may judge them by their decisions." Thus, although punishment may not be possible within the game itself, players behave as if it is. This view is supported by the findings of Haley and Fessler (2005), who report strikingly different results in dictator games when players believe they can and cannot be seen, even in a context where the identity of the would-be punisher is undefined.

As with the anonymous offerer version of the dictator game, subjects played both a 100 USH and a 500 USH denomination version of the game. Each subject played approximately two rounds of the former and four rounds of the latter, yielding a total of 672 rounds (and 1,344 individual choices) in the 100 USH denomination game and 1,226 rounds (2,452 choices) in the 500 USH denomination game. Aggregate patterns of play were similar to those found in the dictator game with the anonymous offerer. The modal strategy in the 100 USH denomination version was to retain 600 USH and to allocate 200 USH to each of the other players (employed in 23% of rounds). The next most common strategy was to keep 400 USH and to allocate 300 USH to each other player (22% of rounds). Again, a strong norm of inequality aversion was evident. On average, players retained 548 USH and allocated 226 USH to each of the other players. In the 500 USH denomination game, aggregate results were again similar to those reported for the anonymous offerer version: in 70% of cases, subjects kept one 500 USH coin and allocated the other to another player and in 24% of cases they gave both coins away.

However, as seen in Table 6, there were systematic differences in the way subjects played with co-ethnics and non-co-ethnics, and across different player-types. Nonegoist players offered somewhat *less* to co-ethnics, though this effect is small and statistically weak (row 1). However, when we focus on the behavior of egoists alone, we find a very strong and highly significant co-ethnic effect. Egoists, while offering considerably less on average (row 2), offer significantly more to co-ethnics (row 3)—between 23 and 52 USH more, depending on the particular measure of co-ethnicity we use (calculated by adding the main and interaction effects—rows 1 and 3). Relative to nonegoists, they offer between 31 and 88 USH more to co-ethnics.¹⁷

We find similar results in the 500 USH denomination game (Table 7). Focusing only on the sample of subjects that actually discriminate and only on those rounds where offerers faced a choice between a co-ethnic and a non-co-ethnic, we find strong evidence

¹⁷ In addition to fixed effects for the offerer's group (shown in Table 6), these results survive the inclusion of fixed effects for the ethnic group of the receiver (three of four specifications) and for both the offerer and the receiver as individuals (three of four specifications). As with the results reported in Table 2, they are also robust to the inclusion of a battery of controls for age, education, income, gender, and other characteristics which, by design, are uncorrelated with the experimental treatment.

TABLE 6. Average Offers in the 100 US\$ Dictator Game When Offerers Are Seen

	Benchmark Co-ethnicity		Subjective Co-ethnicity	
	Ethnicity	Region	Ethnicity	Region
In-group member	-19 (11.03)*	-8 (10.94)	-36 (20.26)*	-8 (19.94)
Egoist	-106 (11.73)***	-106 (12.51)***	-114 (12.57)***	-115 (15.11)***
Egoist & In-group member	55 (16.99)***	31 (15.25)**	88 (28.74)***	51 (27.51)*
Observations	951	951	922	922
R-squared	0.23	0.19	0.24	0.20

Notes: ***Significant at 1%; **significant at 5%; *significant at 10%. Robust standard errors in parentheses. Each column reports the coefficient on co-ethnicity/co-region from an OLS regression. The benchmark ethnicity samples exclude observations in which enumerators indicated that the subject had problems comprehending the game logic and observations with no information on the receiver. The subjective ethnicity sample is restricted to observations in which receivers are identified in the identification game. In each case fixed effects for the player's group are included, and disturbance terms are clustered for each player across all of his or her games.

TABLE 7. Discrimination Rates in the 500 US\$ Dictator Game When Offerers Are Seen

	Benchmark Co-ethnicity		Subjective Co-ethnicity	
	Ethnicity	Region	Ethnicity	Region
In-group member	0.12 (0.07)*	0.14 (0.06)**	0.29 (0.14)**	0.25 (0.11)**
Observations	406	608	188	272

Notes: ***Significant at 1%; **significant at 5%; *significant at 10%. Robust standard errors in parentheses. Discrimination game analysis is limited to games in which players play with exactly one in-group member and one out-group member and discrimination is observed. Each column reports the marginal effect of co-ethnicity/co-region from a probit estimation. Group fixed effects are included and disturbance terms are clustered across the individual actions of offerers in a given game.

that subjects discriminate in favor of co-ethnics.¹⁸ Using the benchmark measure, co-ethnicity increases the likelihood that a partner will be favored by 12 percentage points; using our subjective coding, co-ethnicity increases the likelihood of favoritism by 29 percentage points.¹⁹ Results for shared group membership defined by common regional background are similarly large.

The contrast between these findings and those reported in Tables 2 and 3 is striking. When they are making their offers anonymously (egoist), subjects give no more to co-ethnics than to non-co-ethnics. But when they know that they can be seen, they give significantly more. These results provide powerful evidence that cooperation-facilitating norms exist within ethnic groups. To test the findings even more rigorously, we estimate the joint effect of co-ethnicity and offerer anonymity in a single framework (focusing only on egoists). We pool the data from all dictator games and

compare the impact on offerers' behavior of being seen, being paired with a co-ethnic, and the interaction of these two conditions. This interaction term is especially useful, as it provides a direct test of the impact of social norms on within-group cooperation while controlling for the possible effects of within-group altruism. We report our results in Table 8.

For the 500 US\$ denomination game, we find qualitative support for the claim that players exhibit in-group bias if and only if they are observed, although the results do not reach significance at conventional levels (row 3; columns 3 and 4). In the 100 US\$ denomination game, however, we find strong evidence that players discriminate in favor of co-ethnics *if and only if* they can be seen to be doing so (row 3; columns 1 and 2). Taken together, these findings offer strong support for this strategy selection mechanism as an important source of the variation we observe in public goods provision across ethnically homogeneous and heterogeneous settings.

DISCUSSION

Our findings so far provide the building blocks for an explanation for why ethnic diversity impedes the provision of public goods. They suggest that ethnically homogeneous communities have an advantage in

¹⁸ Similarly, in analyses using non-stacked data (not shown), we can reject the null hypothesis that players select between a co-ethnic and a non-co-ethnic randomly at the 99% level for co-ethnicity and the 95% level for co-region using a subjective coding of co-ethnic pairings.

¹⁹ We do not have the power in our sample to test for the robustness of these results to the inclusion of fixed effects for the ethnic group of the receiver (in addition to that of the offerer).

TABLE 8. The Importance of Being Seen

	100 US\$ game (Egoists)		500 US\$ game	
	Benchmark Co-ethnicity	Subjective Co-ethnicity	Benchmark Co-ethnicity	Subjective Co-ethnicity
Offerer is Seen	-7.93 (6.50)	-11.28 (8.13)	-0.08 (0.06)	-0.08 (0.09)
In group member	-17.66 (10.36)*	-10.95 (15.18)	-0.03 (0.10)	0.08 (0.16)
Seen & In group member	45.33 (15.49)***	51.11 (22.36)**	0.15 (0.12)	0.20 (0.21)
Observations	926	892	628	310

Notes: ***Significant at 1%; **significant at 5%; *significant at 10%. Robust standard errors in parentheses. Columns 1 and 2 report the coefficient on co-ethnicity from an OLS regression in the sample of egoistic subjects with group fixed effects and clustering on actions by a given player. Columns 3 and 4 report the marginal effect of co-ethnicity estimated from a probit model with group fixed effects and clustering on actions by a given player in a given game.

providing public goods because ethnic groups possess both norms and networks that facilitate the sanctioning of community members who fail to contribute to collective endeavors.²⁰ Evidence that the threat of sanctioning is especially effective in policing the behavior of egoists only reinforces these results, since egoists are precisely those community members who are most predisposed not to contribute to public goods provision.

In interpreting our findings as evidence for the sanctioning account, we have made two implicit assumptions. First, we have assumed that behavior in games that capture *aspects* of a collective action dilemma also extend directly to collective action problems. It is reasonable to question whether the co-ethnic bias we find in the nonanonymous dictator game is also apparent in a public goods game. Second, we have assumed that offerers' responses to being seen in the nonanonymous dictator game reflect an anticipation of sanctioning even though such sanctioning is not possible within the context of the game. In fact, although the sanctioning mechanism is among the most prominently discussed in political science research, other strategy selection mechanisms that do not depend on sanctioning could be in operation.²¹ It is therefore reasonable to ask

whether players do engage in sanctioning when reciprocity norms are violated and whether the threat of such sanctioning leads to greater cooperation.

To test the plausibility of the first assumption, we had subjects play the well-known prisoners' dilemma (PD) game. In this game, each of two players must decide whether to keep a 1,000 US\$ note or to allocate it to a group project. Money contributed to the group project realizes a return of 50% and then is divided evenly between the two players. Thus if both players allocate their 1,000 US\$ notes to the group project, they will each receive returns of 1,500 US\$. If one player allocates his or her endowment to the group and the other does not, the former will receive 750 US\$ and the latter will receive 1,750 US\$. The players, who can see images of each other in a modified two-person PIB, make their choices simultaneously and are not able to observe each other's actions (payments are aggregated from all rounds played and not paid until the end of the experimental session).

To test the sanctioning assumption more directly, we examine a version of the PD game that adds a third player (the "enforcer," whose image is seen in the upper corner of the PIB) who observes the actions of the first two players and can choose to punish either one or both of them for their behavior (Fehr and Fischbacher 2004). Punishment costs the enforcer 500 US\$ per player to be punished and results in the removal of the punished player's earnings for the round.

Our findings in these games are summarized in Table 9. The baseline rate of cooperation in the PD game with no enforcer is 56% (row 1, column 4). Behavior is quite different, however, across different player types. Egoists are 11 percentage points less likely to cooperate than nonegoists (row 1, column 3). The

²⁰ One interpretation of our findings is that co-ethnics give more when they are observed simply because they know that they can be found; if this were true, the co-ethnic effect would result entirely from a technology correlated with co-ethnicity, rather than a strategy selection mechanism. We looked for evidence of this story in our data by examining whether other types of people that can be found easily (e.g., older people, Muslims) also give more in the dictator game only when they are observed, but found no evidence for such a pattern of behavior. We reserve for future analysis the task of further assessing the relative contributions of the technology (findability) and the strategy selection mechanism (social sanctioning) in supporting cooperation among co-ethnics.

²¹ Our results are, for example, also consistent with the possibility that players are "conditional reciprocators" in the sense that they view collective action problems as stag hunt/assurance games—games in which all-cooperate and all-defect are both Nash equilibria. Under this account, co-ethnics may coordinate on the Pareto optimal equilibrium while non-co-ethnics do not, but, as with the social sanctioning mechanism, achieving such coordination requires knowledge of co-ethnicity and thus the loss of anonymity. One might also imagine a story in which reciprocity norms play no role

at all. It may be the case that individuals choose to cooperate with co-ethnics in nonanonymous versions of the dictator game because the symmetry of the game leads them to believe (irrationally) that their action will lead others to reciprocate, even in the absence of a causal mechanism linking the two players' actions (what some social scientists have referred to as "quasi-magical thinking"). Based on data in the nonanonymous dictator game alone, we cannot rule out these other possibilities.

TABLE 9. Cooperation and Co-Ethnicity in the Prisoners' Dilemma Game

	Enforcer?	(1) Non-Egoist	(2) Egoist	(3) Difference	(4) All Subjects
Likelihood of Cooperation	No	0.61 <i>n</i> = 584	0.50 <i>n</i> = 382	-0.11*** (0.03)	0.56 <i>n</i> = 966
Marginal Effect of Subjective Co-ethnicity	No	0.11 (0.09)	0.29*** (0.11)	0.18 (0.13)	0.18** (0.07)
Likelihood of Cooperation	Yes	0.64 <i>n</i> = 791	0.60 <i>n</i> = 497	-0.04 (0.03)	0.62 <i>n</i> = 1288
Marginal Effect of Subjective Co-ethnicity	Yes	-0.00 (0.08)	-0.02 (0.09)	-0.02 (0.11)	-0.01 (0.06)
Difference Between Row 1 and Row 3		0.03 (0.03)	0.10*** (0.03)	0.07* (0.04)	0.06*** (0.02)

Notes: ***Significant at 1%; **significant at 5%; *significant at 10%. Robust standard errors in parentheses. Column 3, rows 1, 3, and 5: *t*-test of difference of means in columns 1 and 2. Rows 2 and 4: result of a probit (dprobit) model on a continuous identity measure with group specific fixed effects and clustering by individual subject. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Disturbance terms are clustered for each player across all of his or her games.

key question, however, is whether co-ethnics are more likely to cooperate than non-co-ethnics. The answer is that they are. As reported in row 2, column 4, co-ethnics are 18 percentage points more likely to cooperate in the PD game than are non-co-ethnics. As a comparison of columns 1 and 2 of row 2 makes clear, however, this finding is driven largely by the behavior of egoists. As in the nonanonymous dictator game, although egoists are much less likely to cooperate in general, they are much more likely to cooperate when playing with members of their own ethnic group.

What happens once sanctioning becomes possible? Do players anticipate punishment and respond with higher levels of cooperation? Our results suggest that they do. Consistent with the sanctioning mechanism, we find that the introduction of a threat of punishment causes baseline cooperation rates to rise by 6 percentage points (see row 5, column 4). This increase in cooperation is even greater among—and, in fact is largely driven by—the egoist subgroup. Egoists are 10 percentage points more likely to cooperate when there is a threat of sanctioning (see row 5, column 2). We find, however, that the anticipation of punishment among egoists does not appear to depend on the co-ethnicity of the pairing, contrary to what we would predict (see row 4). The likelihood of cooperation by egoists rises both with the introduction of a direct threat of punishment (row 5, column 2) and when they are playing with co-ethnics (row 2, column 2), but we cannot identify a further increase that can be attributed to the interaction of these two conditions (row 4, column 2).

Are players right to anticipate sanctioning, even when punishment is costly? Turning to the behavior of the enforcers, we find that they are. Generally, enforcers punish one third of the time when a player defects and only one time in six when a player cooperates. In keeping with the within-group norms story, we also find that enforcers punish defecting co-ethnics more than they do defecting non-co-ethnics, and that they punish especially when a co-ethnic of the enforcer

defects in a PD game with another co-ethnic (i.e., in a homogeneous trio). Consistent with the theoretical work of Bowles and Gintis (2004b), among others, it is the nonegoists, not the egoists, that engage in punishment: nonegoist enforcers punish defectors in non-homogenous trios with a probability of .35 but punish defectors in homogenous trios with a probability of .58 (the associated *p*-value for this difference in a simple comparison of means test is .01).²²

Evidence from a PD game with and without a third-party enforcer thus provides further support for the existence of in-group reciprocity norms in Mulago-Kyebando: subjects in our sample anticipate punishment; they engage in sanctioning even at a cost to themselves; and they are most likely to punish players that fail to contribute to public goods when the non-contributors are co-ethnics. Although other strategy selection mechanisms cannot be ruled out, the sanctioning account appears to rest on strong empirical foundations.

CONCLUSION

Empirical work from Pakistan to Indonesia and from Kenya to the United States suggests that community-level ethnic diversity impedes the provision of public goods. The aim of this paper has been to determine why.

We have offered an empirical strategy to distinguish among distinct mechanisms that might account for the negative association between ethnic diversity and public goods provision. Employing games from experimental economics as well as more standard survey techniques, and working with a representative sample of subjects from an urban slum, we have generated

²² In a full regression controlling for group fixed effects and clustering by player, the difference in punishment rates between homogeneous and nonhomogenous trios is 22 percentage points with nonegoist punishers. The associated *p*-value is .08. A full table with these results is available from the authors on request.

evidence that successful collective action among homogenous ethnic communities can be attributed to the existence of norms and institutions that police the defection of non-contributors. Our findings suggest that co-ethnics cooperate because they adhere to in-group reciprocity norms—norms that are plausibly supported by expectations that non-contribution will be sanctioned and by an ethnic technology, “findability,” that facilitates sanctioning among co-ethnic pairings. Moreover, by focusing on the play of different “types” of individuals, we have shown that the positive impact of ethnic homogeneity on collective action stems directly from the ability of ethnic ties to induce more cooperative behavior among those individuals who, absent the social connection provided by ethnicity, would be least likely to cooperate.

Just as significant as the mechanisms that appear to be at work are those that are not. We find no evidence that ethnic groups in Mulago-Kyebando have tastes for different kinds of public goods or that individuals exhibit greater degrees of altruism toward co-ethnics. If co-ethnics are more effective at producing public goods, this does not appear to be because they care about the same things or value the welfare improvements of fellow ethnic group members more highly than those of non-co-ethnics. Nor do we find robust evidence that shared ethnicity facilitates the productivity of teams: co-ethnics are not significantly more effective at working together on joint tasks. Thus, although our core findings are consistent with previous work on ethnic diversity and public goods provision (e.g., Besley, Coate, and Loury 1993; Miguel and Gugerty 2005; Yamagishi 2007), our confidence that the success of homogenous communities in urban Kampala derives from in-group reciprocity norms is greatly enhanced by our testing of—and consequent ability to rule out—these alternative mechanisms.

A major implication of our findings is that generating higher levels of public goods provision in diverse communities does not require the segregation of ethnic groups, as many preference-based or technology-based explanations might suggest. Indeed, just the opposite may be needed: policies that promote repeated social interactions and the free flow of information across ethnic lines. Our results suggest that when individuals believe that their behavior is observed by others and that their reputation may influence opportunities for cooperation in the future, social cooperation can happen, even in the face of individual incentives to shirk. The challenge of generating effective cooperation in diverse societies is to make such beliefs equally characteristic of cross-group and within-group interactions.

We have motivated our analysis in this paper with the problem of public goods provision. However, the implications of our study extend to a broad range of outcomes that are similarly dependent on collective action. Scholars who seek to explain why social interaction so often takes on an ethnic character—why groups mobilize along ethnic lines for civil war (Cederman and Girardin 2007), why ethnic cleavages shape the formation of political parties and the distribution of

public resources (Chandra 2004), or why economic activities often take place in ethnic enclaves (Fafchamps 2000)—tend to focus on a class of social dilemmas not unlike the collective action problems experienced by residents in Mulago-Kyebando. The mechanisms we have proposed to account for the relationship between diversity and the failure of public goods delivery, then, are quite general. Whether the same mechanisms we find most salient for understanding the impact of ethnicity diversity on public goods production in Mulago-Kyebando are also salient in other settings, and for other problems, remains an empirical question. The experimental apparatus we have developed here, however, can be easily—and, we believe, profitably—used to answer this question.

REFERENCES

- Alesina, Alberto, Reza Baqir, and William Easterly. 1999. “Public Goods and Ethnic Divisions.” *Quarterly Journal of Economics* 114 (November): 1243–84.
- Alesina, Alberto, and Eliana LaFerrara. 2002. “Who Trusts Others?” *Journal of Public Economics* 85 (August): 207–234.
- Alesina, Alberto, and Eliana LaFerrara. 2005. “Ethnic Diversity and Economic Performance.” *Journal of Economic Literature* 63 (September): 762–800.
- Axtell, Robert, H. M. Epstein, and H. P. Young. 2001. “Emergence of Class Norms in a Multi-Agent Model of Bargaining.” In *Social Dynamics*, ed. S. Durlauf and H. P. Young. Cambridge: MIT Press.
- Bahry, Donna, and Rick Wilson. Forthcoming. “Confusion or Fairness in the Field? Rejections in the Ultimatum Game under the Strategy Method.” *Journal of Economic Behavior and Organization*.
- Banerjee, Abhijit, Lakshmi Iyer and Rohini Somanathan, 2005, History, Social Divisions and Public Goods in Rural India, *Journal of the European Economic Association* 3 (April-May), 639–47.
- Bates, Robert. 1973. *Ethnicity in Contemporary Africa*. Syracuse: Program in East African Studies.
- Becker, Gary. 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.
- Besley, Timothy, Stephen Coate, and Glenn Loury. 1993. “The Economics of Rotating Savings and Credit Associations.” *American Economic Review* 83 (September): 792–810.
- Bohnet, Iris, and Bruno Frey. 1999. “Social Distance and Other-Regarding Behavior in Dictator Games: Comment.” *American Economic Review* 89 (March): 335–39.
- Bolton, Gary E., and Axel Ockenfels. 2000. “ERC: A Theory of Equity, Reciprocity, and Competition.” *American Economic Review* 90 (March): 166–93.
- Bowles, Samuel, and Herbert Gintis. 2004a. “Persistent Parochialism: Trust and Exclusion in Ethnic Networks.” *Journal of Economic Behavior and Organization* 55 (September): 1–23.
- Bowles, Samuel, and Herbert Gintis. 2004b. “The Evolution of Strong Reciprocity: Cooperation in Heterogeneous Populations.” *Theoretical Population Biology* 65 (February): 17–28.
- Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press.
- Cederman, Lars-Erik, and Luc Girardin. 2007. “Beyond Fractionalization: Mapping Ethnicity onto Nationalist Insurgencies.” *American Political Science Review* 101 (February): 173–85.
- Chandra, Kanchan. 2004. *Why Ethnic Parties Succeed: Patronage and Ethnic Headcounts in India*. New York: Cambridge University Press.
- Charness, Gary, and Uri Gneezy. 2000. “What’s in a Name? Anonymity and Social Distance in dictator and Ultimatum games.” Unpublished manuscript, University of California, Santa Barbara.
- Crawford, Vincent P., and Hans Haller. 1990. “Learning How to Cooperate: Optimal Play in Repeated Coordination Games.” *Econometrica* 58 (May): 571–95.
- Deutsch, Karl. 1966. *Nationalism and Social Communication*. Cambridge: The MIT Press.

- Fafchamps, Marcel. 2000. "Ethnicity and Credit in African Manufacturing." *Journal of Development Economics* 61 (February): 205–35.
- Fearon, James D. 1999. "Why Ethnic Politics and 'Pork' Tend to Go Together?" Unpublished manuscript, Stanford University.
- Fearon, James D., and David D. Laitin. 1996. "Explaining Interethnic Cooperation." *American Political Science Review* 90 (December): 715–35.
- Fehr, Ernst, and Urs Fischbacher. 2004. "Third-Party Punishment and Social Norms." *Evolution and Human Behavior* 25 (March): 63–87.
- Fehr, Ernst, and Klaus Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (August): 817–68.
- Fershtman, Chaim, and Uri Gneezy. 2001. "Discrimination in a Segmented Society: An Experimental Approach." *Quarterly Journal of Economics* 116 (February): 351–77.
- Fowler, James H. 2006. "Altruism and Turnout." *Journal of Politics* 68 (August): 674–83.
- Ghosh, Parkishit, and Debraj Ray. 1996. "Cooperation in Community Interaction without Information Flows." *Review of Economic Studies* 63 (July): 491–519.
- Gil-White, Francisco. 2004. "Ultimatum Game with an Ethnicity Manipulation: Results from Kohvdiin Bulgan Sum, Mongolia." In *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr and Herbert Gintis, eds. New York: Oxford University Press.
- Golooba-Mutebi, Frederick. 2003. "Devolution and Outsourcing of Municipal Services in Kampala City, Uganda: An Early Assessment." *Public Administration and Development* 23 (December): 405–18.
- Greif, Avner. 1989. "Reputations and Coalitions in Medieval Trade: Evidence on the Maghribi Traders." *Journal of Economic History* 49 (December): 857–82.
- Greif, Avner. 1994. "Cultural Beliefs and the Organization of Society." *Journal of Political Economy* 102 (October): 912–50.
- Greig, Fiona, and Iris Bohnet. 2006. "Why Women Cooperate with Women and not Men: Evidence from a Slum in Nairobi, Kenya." Unpublished manuscript, Kennedy School of Government.
- Güth, Werner, and Hartmut Kliemt. 1998. "The Indirect Evolutionary Approach: Bridging the Gap between Rationality and Adaptation." *Rationality and Society* 10 (August): 377–99.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy Weinstein. 2004. "Ethnic Identifiability: An Experimental Approach." Unpublished manuscript, UCLA.
- Haley, Kevin, and Daniel Fessler. 2005. "Nobody's Watching? Subtle Cues Affect Generosity in an Anonymous Economic Game." *Evolution and Human Behavior* 26 (May): 245–56.
- Hardin, Russell. 1995. *One for All: The Logic of Group Conflict*. Princeton: Princeton University Press.
- Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis. 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. New York: Oxford University Press.
- Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith. 1996. "Social Distance and Other-Regarding Behavior in Dictator Games." *American Economic Review* 86 (June): 653–60.
- Kahneman, Daniel, Jack Knetsch, and Richard Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review* 76 (September): 728–41.
- Khwaja, Asim. 2002. "Can Good Projects Succeed in Bad Communities? Collective Action in the Himalayas." Unpublished manuscript, Harvard University.
- Miguel, Edward. 1999. "Ethnic Diversity, Mobility, and School Funding: Theory and Evidence from Kenya." Unpublished manuscript, Harvard University.
- Miguel, Edward, and Mary Kay Gugerty. 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89 (December): 2325–68.
- Mozaffar, Shaheen, James R. Scarritt, and Glen Galaich. 2003. "Electoral Institutions, Ethnopolitical Cleavages, and Party Systems in Africa's Emerging Democracies." *American Political Science Review* 97 (August) 379–90.
- Onyach-Olaa, Martin. 2003. "The Challenges of Implementing Decentralization: Recent Experiences in Uganda." *Public Administration and Development* 23 (January): 105–13.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.
- Ostrom, Elinor. 2000. "Collective Action and the Evolution of Social Norms." *Journal of Economic Perspectives* 14 (Summer): 137–58.
- Posner, Daniel N. 2005. *Institutions and Ethnic Politics in Africa*. New York: Cambridge University Press.
- Poterba, James. 1997. "Demographic Structure and the Political Economy of Public Education." *Journal of Fiscal Analysis and Management* 16 (Winter): 48–66.
- Spolaore, Enrico, and Romain Wacziarg. 2006. "The Diffusion of Development." Unpublished manuscript, Stanford University.
- Tajfel, Henri. 1974. "Intergroup Behavior, Social Comparison and Social Change." Katz-Newcomb Lectures, University of Michigan, Ann Arbor.
- Whitt, Sam, and Rick Wilson. 2006. "Fairness and Ethnicity in the Aftermath of Ethnic Conflict: The Dictator Game in Bosnia-Herzegovina." Unpublished Manuscript, Rice University.
- Vigdor, Jacob. 2004. "Community Composition and Collective Action: Analyzing Initial Mail Response to the 2000 Census." *Review of Economics and Statistics* 86: 1: 303–12.
- Yamagishi, Toshio. 2007. "The Social Exchange Heuristic: A Psychological Mechanism That Makes a System of Generalized Exchange Self-Sustaining." In *Cultural and Ecological Foundations of the Mind*. Mark Radford, Susumu Ohnuma, and Toshio Yamagishi, eds. Sapporo: Hokkaido University Press.