

IZA DP No. 2272

Why Does Ethnic Diversity Undermine Public Goods Provision? An Experimental Approach

James Habyarimana
Macartan Humphreys
Daniel N. Posner
Jeremy Weinstein

August 2006

Why Does Ethnic Diversity Undermine Public Goods Provision? An Experimental Approach

James Habyarimana

Georgetown University and IZA Bonn

Macartan Humphreys

Columbia University

Daniel N. Posner

University of California, Los Angeles

Jeremy Weinstein

Stanford University

Discussion Paper No. 2272
August 2006

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
Email: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Why Does Ethnic Diversity Undermine Public Goods Provision? An Experimental Approach*

A large and growing literature links high levels of ethnic diversity to low levels of public goods provision. Yet while 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. To gain analytic leverage on the question of why ethnicity matters, we identify three families of mechanisms – what we term preference, technology, and strategy mechanisms. Our empirical strategy is to identify and run a series of experimental games that permit us to examine these mechanisms in isolation and then to compare the importance of ethnicity in each. Results from experimental games conducted with a random sample of 300 subjects in Kampala's slums reveal that successful collective action among homogenous ethnic communities in urban Uganda is attributable to the existence of norms and institutions that facilitate the sanctioning of non-contributors. We find no evidence for a commonality of tastes within ethnic groups, for greater degrees of altruism toward co-ethnics, or for an impact of shared ethnicity on the productivity of teams.

JEL Classification: D71, H41, J15, O10, Z13

Keywords: ethnic diversity, collective action, public goods, field experiments

Corresponding author:

James Habyarimana
Georgetown University
3520 Prospect Street, Suite 308A
Washington, DC 20007
USA
E-mail: jph35@georgetown.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; the Harvard Academy for International and Area Studies, the Russell Sage Foundation, and the Harry Frank Guggenheim Foundation for financial support; Rick Wilson, David Laitin, Kanchan Chandra, Alan Gerber, Don Green, John Okidi, members of the Laboratory for Comparative Ethnic Processes, and seminar participants at the University of Michigan, UCLA, UC Berkeley, UC Davis, MIT, Wharton, Northwestern, Oxford, Stanford, Caltech, and the EPRC (Kampala) for helpful comments on previous versions of this paper.

I. Introduction

In 2003, the Ugandan Police Service withdrew its funding for local-defense units in Kampala's urban slums. Whereas previously the government paid officers to patrol neighborhoods and respond to urban crime, a tightening budget situation meant that the provision of security became a purely local matter – one left to a decentralized system of local councils.¹ In Kifumbira, the chairman of the local council expressed the depth of his concern about problems of criminality: “We have thieves, murders, and lots of people with problems with gambling and opium. We had three murders last month alone.”² Yet community members in Kifumbira have been unable to mobilize their neighbors in response. Amidst rising crime, residents complained bitterly about the lack of a police presence but have done little to address the problems on their own. In nearby Katala, however, the disappearance of government support for policing has not led to a spike in crime, even though it too sits in the middle of one of Kampala's poorest slums. The chairman of Katala zone reported that his community has initiated a community patrol; organized by the elected Parish Security Officer, the patrol involves all members of the local council along with volunteers from the neighborhood. Apart from petty crimes, the chairman explained, “Katala is very, very safe.”³

Although the two communities are similar across a wide range of observables, residents in Kifumbira, unlike their neighbors in Katala, continue to live in fear, choosing to blame the police rather than mobilize community patrols to address rising insecurity. This lack of action cannot be accounted for by ignorance of the problem – indeed, the chairman of the local council showed us scars on his arms, a consequence of his personal effort to apprehend a thief. One characteristic does, however, appear to distinguish Kifumbira from Katala: its level of ethnic diversity. Kifumbira is home to a wide diversity of ethnic groups, while Katala is more homogenous.

The hypothesis that Katala's success and Kifumbira's failure might be linked to Katala's lower level of social heterogeneity resonates with a large and growing literature that associates high levels of ethnic diversity with low levels of public goods provision. Indeed, the failure of Kifumbira's residents to provide basic public goods appears not to be atypical: highly diverse U.S. cities exhibit lower levels of funding for public goods (Alesina, Baqir, and Easterly 1999), contributions to local schools are lower in ethnically heterogeneous East African villages (Miguel and Gugerty 2005), and project maintenance suffers in more socially fractionalized villages in Pakistan (Khwaja 2002).⁴ Yet while the empirical connection between ethnic heterogeneity and the underprovision of public goods is widely accepted, there is little consensus on the

¹ Kampala has a hierarchical administrative structure of local councils (LCs). The city is partitioned into five divisions (LC3s), which are, in turn, divided into parishes (LC2s) and zones (LC1s). LC1s are equivalent in size to villages in rural areas.

² Interview, Kifumbira, July 20, 2005.

³ Interview, Katala, July 7, 2005.

specific mechanisms through which this relationship operates. Some authors emphasize the altruism that co-ethnics feel toward each other and the greater willingness that this provides for individuals in ethnically homogeneous communities to bear the cost of volunteering their time and energy for collective activities. Others point to shared preferences among co-ethnics about the types of public goods that should be provided. Still others stress the role that ethnicity plays as a mechanism for improving the productivity of collective tasks, for example, through shared language and a greater ability to communicate. Finally, some highlight the greater capacity for individuals to sanction in-group members than out-group members and the heightened ability that this provides for enforcing social cooperation in homogeneous settings. All of these mechanisms are plausible. To date, however, no research has established which of these mechanisms is (are) doing the work in linking ethnic heterogeneity with low levels of public goods provision. While much research has been undertaken to establish *that* ethnicity matters – both for public goods provision and for a range of other collective action-dependent outcomes from credit provision to rebel recruitment – considerably less tries to work out *why* ethnicity matters. Filling this gap in knowledge is the purpose of this paper.

To gain analytic leverage on the question why ethnicity matters, we begin by identifying three families of mechanisms— what we term preference, technology, and strategy. Various authors have at times accounted for the relationship between diversity and the underprovision of public goods with reference to each of these analytically distinct channels. Our empirical strategy is to identify and run a series of experimental games that permit us to examine these mechanisms in isolation and then to compare the importance of ethnicity in each. To do this, we recruited 300 subjects from adjacent neighborhoods in Kampala, Uganda that combined high levels of ethnic diversity with low levels of public goods provision. Each subject played a sequence of different games, each designed to isolate and test one of these mechanisms (described in greater detail in the next section of the paper). Subjects played multiple rounds of each game with randomized matching – sometimes with co-ethnics, sometimes with non-co-ethnics. Our strategy is to examine how the differences in play among co-ethnics and non-co-ethnics varies across games.⁵ This allows us to assess the explanatory power of the distinct mechanisms and to examine how they interrelate.

To preview our results, our evidence suggests that the preference-based explanations common in the literature on ethnic processes do not provide much explanatory power. It is not simply the case that co-

⁴ Other studies that also provide evidence for the negative association between diversity and public goods provision include Easterly and Levine (1997); Poterba (1997); Vigdor (2004); Okten and Osilli (2004); and Alesina and LaFerrara (2002 and 2005).

⁵ Other researchers have sought to identify race/ethnicity effects in a different set of experimental games than those described in this paper. Eckel and Grossman (2001) find race effects in ultimatum games; Glaeser et al. (2000) find race effects in a trust game; Fershtman and Gneezy (2001) identify group-specific differences in allocations in the trust game, an effect they argue follows from ethnic stereotypes rather than a taste for discrimination; and Gil-White (2004) explores an ethnicity manipulation in ultimatum games played in Mongolia.

ethnics are more altruistic toward one another or happen to share the same preferences over outcomes. Nor are technological explanations central. Co-ethnics do not appear to cooperate with one another simply because cooperation is easier with people that share a common language or modes of interacting. Instead, we find that strategic concerns dominate: our evidence suggests that co-ethnics cooperate because they adhere to within-group norms and institutions that facilitate the sanctioning of individuals who fail to contribute to collective endeavors. Underlying this strategic explanation, however, is a technological feature: co-ethnics benefit from stronger network linkages among them, linkages which make it possible for individuals to sanction co-ethnics that defect. These results provide some of the first empirical evidence for why ethnic heterogeneity undermines collective action, and for why ethnically diverse communities tend to suffer from such low levels of public goods provision and well-being.

II. Mechanisms

It is a common refrain among scholars of comparative politics that “ethnicity matters.” Measures of ethnic diversity or discrimination are now standard variables in cross-country regressions, hypothesized to account for the onset of conflict (Cederman and Girardin, forthcoming), the instability of democracy (Adsera and Boix 2004), and patterns of economic growth (Easterly and Levine 1997). But what does it mean to say that ethnicity matters? While authors identify myriad ways in which ethnicity appears consequential in the aggregate, no general framework has been used to think analytically about how ethnic identities shape individual behavior. In part, this is because ethnicity is a highly contested term. Indeed, it is difficult to find a definition that captures the central notions of ethnicity employed by major writers on the subject (Chandra 2005). It may also be because general frameworks usually operate at a level of abstraction at which the primitives involve objects such as individuals, rules, or goals, and possibly even “groups,” but not “ethnic groups” in particular. We therefore begin by presenting a general framework of individual behavior which we then use to identify the different channels through which ethnicity might affect behavioral outcomes. We identify three distinct families of mechanisms through which ethnicity might matter. In the section that follows, we show that the various arguments found in the literature about the impact of ethnicity can be grouped into these three general classes of mechanisms.

The general model of individual behavior we draw on is the game theoretic model. In the abstract language of game theory, a social interaction can be described as a game. Such games (we focus here on normal form games) are *triples* comprising (i) a set of players, $N = \{1, 2, \dots, n\}$, (ii) a set of strategies, one for each player, $A = \prod_{i \in N} A_i$, with typical elements $a_i \in A_i$ and $a \in A$ and (iii) a set of preferences over outcomes, $u = \{u_i\}_{i \in N}$, again, one for each player. An abstract game G may be written $G = G(N, A, u)$.

Given any game, the analyst’s behavioral assumptions (typically embedded in a “solution concept”) returns the profile of actions a that the analyst expects the players to take. The set of actions taken in a given game can then be written $a(G(N, \mathcal{A}, u))$. The overall outcome, x , can be then written as a function of a , in effect, the outcome is a summary of the actions taken by all the players.⁶

As an example, consider a game of chess. N consists of two players, \mathcal{A} consists of the set of moves that each can play (i.e., bishops can move diagonally; rooks can move horizontally or vertically, etc.), and u reflects the fact that players prefer combinations of strategies that lead to a win rather than a loss. Together, these three elements form the game of chess. Given this game, alongside the analyst’s behavioral assumptions, we have a description of behavior $a(\cdot)$: for example, white moves his king’s pawn forward two steps, and so on. Given these strategies, x reports whether white or black wins, or whether there is a draw.

Note that this general description of a strategic environment – of a social problem – can be undertaken without any reference to the ethnic identities of players. In practice, however, we might expect that the ethnic identities of players affect their behavior in systematic ways. If so, then such identities can be introduced into the formal description of a game. To see this, let the set of ethnic groups be given by Θ , where $\theta_i \in \Theta$ denotes the ethnicity of player i . Working out the ways in which ethnicity is “responsible” for a given outcome can be thought of as asking: how does variation in the profile of types θ , affect the outcome, x ?

Given our general characterization of the game, there are a small number of “first order channels” to investigate: in principle, ethnicity could enter by changing any of the three primitives, \mathbf{N} , \mathbf{A} or \mathbf{u} , or by changing the selection of strategies, that is, the mapping $\mathbf{a}(\mathbf{G})$. Assuming that ethnicity is predicated on existence (that is we do not treat N as a function of θ), we are left with three mechanisms through which the outcome, x , may be affected by θ :

- The **preferences** mechanism: Ethnicity can alter the preferences players have over outcomes: $u(a | \theta)$.
- The **technology** mechanism: Ethnicity can alter the “toolbox” of strategies available to players: $\mathcal{A}(\theta)$.
- The **strategy** mechanism: Ethnicity can be used to condition the strategies that players select in a given game: $a(\cdot | \theta)$.

As an example, consider the following simple game. Two people can decide whether or not to take some action, such as sharing information or contributing labor to some project. If an individual takes such an action then they both benefit, but their cooperation comes at a private cost. It is socially beneficial to take such an action. How much cooperation should we expect to observe?

⁶ In fact there may be indeterminacy over which actions are chosen and so for some solution concepts $a(G(N, \mathcal{A}, u))$ may be a set of strategy profiles rather than a single set of strategies; with such indeterminacy we can think of x as reflecting the distribution of outcomes given a particular solution concept.

The game can be described formally as:

$$N = \{1, 2\}$$

$$A_i = [0, \alpha], 0 \leq \alpha \leq 1, i \in N,$$

$$u_i = a_i - \beta(a_i)^2, 0 \leq \beta \leq 1, i \in N$$

The expected outcome depends on the solution concept employed. For example, the Nash equilibria of this game are as follows: if $\beta > .5$, then each player weighs her own costs higher than the social gain and the Nash equilibrium involves neither player contributing. In this case, the game is a continuous version of the well known Prisoners' Dilemma. If $\beta < .5$, the reverse is true and the Nash equilibrium involves the maximum level of cooperation, α . In this case, there is no social dilemma. If $\beta = .5$, then there is an infinite number of Nash equilibria in which the two players provide equal contributions. Here, we have a coordination dilemma and different equilibrium refinements may suggest different outcomes; for example, the Pareto optimal Nash equilibrium is (α, α) ; the risk dominant Nash equilibrium is $(\alpha/2, \alpha/2)$.⁷

Now consider the three channels through which the introduction of information about the ethnic identities of players may change the game. The first is that the preferences of players may be structured by their ethnicity. For example, we could have $\beta = 0$ if $\theta_i = \theta_j$, $\beta = .5$ otherwise. In this case, a player discounts the private cost of providing information or making other contributions when she shares the same ethnic identity as her partner.

A second possibility is that ethnic identities can affect the availability of particular strategies to players of different identity types. In this game, ethnicity may determine the extent to which two players are *in fact* capable of communicating or of supporting each other's projects. For example, we could have $\alpha = 1$ if $\theta_i = \theta_j$, $\alpha = 0$ otherwise. Thus, when players do not share the same ethnicity, cooperation is not possible.

A final possibility is that the selection of strategies is structured by the ethnic identities of the players. For example, equilibrium selection could follow a rule: players play the Pareto optimal Nash equilibrium strategy whenever $\theta_i = \theta_j$, and play risk dominant Nash equilibrium strategies otherwise. In this case, players condition their choice of strategy on identities, rendering some equilibria salient only when players share a common ethnicity.

Our efforts to distinguish different families of mechanisms through which ethnic identities condition behavioral outcomes provide a useful heuristic device – a language to help identify distinct processes that are often conflated. In theory and in practice, however, the distinctions are not so clear. A given strategic situation may be represented by many different games and what appears as a technological feature in some

⁷ To see that this is the risk dominant outcome, assume that player i has a uniform prior over the actions to be taken by player j . In that case, her expected payoff from action a_i is $u_i(a_i) = \int_0^\alpha (a_i t / \alpha) dt - (a_i)^2 / 2 = a_i \alpha / 2 - (a_i)^2 / 2$ which is maximized at $\alpha / 2$.

representations may be treated as a preference feature in another. Practical challenges are more severe, however, and they motivate our present study. The existence of multiple mechanisms through which ethnic identities shape strategic behavior creates an inferential problem for empirical work that links ethnic diversity to a range of social outcomes. For example, in the game above, if we find that co-ethnics (players for whom $\theta_i = \theta_j$) typically succeed in working together or sharing information while non-co-ethnics do not, we cannot know if this outcome is the result of ethnic identities shaping preferences, technologies, or strategies, or even if some combination of all three mechanisms is in operation. To do this, we need to identify (or experimentally create) situations in which we can rule out the operation of different mechanisms.

III. Ethnicity and Public Goods Provision

The underprovision of public goods is frequently viewed as a product of collective action failure – either a failure of the community to lobby a government to provide a particular public good or a failure of the community to organize itself to supply the public good on its own. The reason public goods are difficult to provide is because they can be subject to free riding: everyone in the community can enjoy the benefits of having nightly security patrols, mended roads, clean water, and garbage collection even if they do not contribute the time or money to providing these services. Understanding why ethnically diverse communities fail to resolve these collective action problems, while homogenous communities do, involves distinguishing among the three mechanisms described above.

Two arguments in the literature on ethnic politics emphasize preference-based explanations. One focuses on commonality of tastes over outcomes. It is possible, for example, that individuals within a group do not care directly for each other's welfare (or, at least, not more than they do for the welfare of out-group members) but that they nonetheless care about similar things. If so, then ethnic diversity implies a heterogeneity of tastes, which, by creating disagreement about which goods should be provided or where they should be located, leads to the underprovision of public goods. Such arguments are found, for example, in the treatments of Alesina, Baqir, and Easterly (1999) and Alesina and LaFerrara (2005). Bates (1973) argues that a salient feature of ethnic groups is that they are often geographically concentrated – this, he suggests, can induce a commonality of interests over outcomes that have a geographic component, notably the location of public investments. In a model of educational choices, Miguel (1999) proposes that individuals have preferences over the *type* of education they receive (for example, the language of instruction or the language in which parent-teacher meetings are conducted) that is correlated across co-ethnics. This can lead to an unwillingness to invest in education in more heterogeneous communities simply because the public good that is produced is of lower value to some individuals than would otherwise be the case.

A second, preference-based explanation focuses on the identities of the beneficiaries rather than on the outcome itself. It starts from the premise that identity may be correlated with preferences over the welfare of other individuals, rather than over public goods directly. Arguments that link ethnicity to the

presence of such other regarding preferences are in fact quite common. As in the game described in Section II, when other regarding preferences for fellow community members are prominent – that is, when individuals assign positive values in their own utility functions to the welfare of their fellow group members – individuals may be willing to bear the cost of providing the public goods themselves, whether or not other community members who benefit from the goods also contribute to their provision. Such arguments are familiar from 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.⁸ If other regarding preferences are stronger within groups than across group lines (or if individuals assign positive utilities to reductions in the welfare of members of other groups – that is, if they have a “taste for discrimination” (Becker 1957)), then having multiple ethnic groups within a community may lower the willingness of individuals to contribute to collective action and lead to lower levels of public goods provision (Vigdor 2004).

A second class of arguments focuses not on preferences but on the technologies available to co-ethnics. In one study of team production, Hardin (1995) emphasizes the need for individuals to coordinate to achieve jointly beneficial outcomes. Coordination, however, may require that individuals can draw on a reservoir of common cultural material – language, experience, networks, understandings about modes of interaction – that makes it easier for community members to communicate and work together to achieve common ends. Deutch (1966), for example, distinguishes ethnic groups (from other identity groups) as “communities of communication”; their cohesiveness is a function of how quickly and effectively they can transmit complicated messages. By raising community members’ *ex ante* expectations of success in their common enterprise, the existence of the shared cultural material should also increase the willingness of community members to undertake the collective endeavor. Since the “ethnic technology” is not shared across group lines, it will only reduce transaction costs (and facilitate public goods provision) in more ethnically homogeneous settings.⁹

Although there are a variety of ways in which co-ethnicity may provide technological advantages, we focus on only two: the possibility that co-ethnics work more effectively together because they can draw on a common language, and the possibility that co-ethnics have access to dense networks upon which they can draw in resolving social dilemmas. Such networks may provide strategy options to players that are simply not available to individuals that are socially isolated from each other, including the possibility for punishment outside of the context of a single-shot interaction (Ghosh and Ray 1996; Fafchamps and Minten, 2002).

⁸ Social psychological studies, however, suggest that in-group affinities, far from being atavistic, are quite universal and can be constructed even in arbitrary groups (Tajfel 1982).

⁹ A recent paper employs a variant of the technology mechanism in accounting for income differences across countries (Spolaore and Wacziarg 2006). The authors use genetic distance between populations as a proxy for the ease with which new technologies can diffuse across borders. In particular, they suggest that the horizontal diffusion of development may be impeded by language and other cultural differences.

A third set of arguments focuses on how ethnic identities condition the strategies selected by players. As in the example in Section II, many social environments are marked by the existence of “multiple equilibria” – that is, a given solution concept (such as Nash equilibria) produces many possible solutions. In such contexts, ethnic information may be used to select one set of solutions over another.¹⁰ Fearon (1999) and Habyarimana et. al. (2005), for example, demonstrate how ethnic information can be used to provide focal points that let individuals select among multiple equilibria in a version of the divide the dollar game. The model most commonly used to represent public goods production problems is the repeat play prisoners’ dilemma. It is well known that (infinitely) repeated games of this form often admit a multiplicity of equilibria. In such contexts, the selection of equilibria may be made dependent on ethnic information in different ways.

One influential article that sought to explain the norm of interethnic peace (and why it sometimes breaks down) did so by comparing the properties of equilibria that differ in the manner in which strategies are conditioned on the ethnic identity of players (Fearon and Laitin 1996). In this, and related work, co-ethnics may sanction members that fail in some way to contribute to collective goods. The expectation that shirking will be punished *ex post* generates expectations *ex ante* that cooperative overtures will be reciprocated, and this, in turn, makes individuals more willing to contribute to collective activities.¹¹ The implication is that social sanctioning will be generally less feasible (and norms against shirking correspondingly lower) in heterogeneous communities, and that the likelihood that public goods will be provided in such settings will be lower. These equilibria produce what may be termed informal social institutions or norms of reciprocity (Bates 1973; Taylor 1988) that reduce the likelihood of shirking. In practice – a point we return to in the conclusion – the selection of equilibria may itself depend on other mechanisms, for example the possibility that greater information about reputations exists within groups than across group lines (Platteau 1994; Landa 1994) and hence that there are more opportunities for sanctioning in-group members who fail to contribute their fair share to collective activities (Besley et al 1993; Fehr and Gächter 2000; Miguel and Gugerty 2005).

IV. Research Site, Sampling, and Experimental Framework

To study the role that these mechanisms play, we took a set of experimental games, typically played in laboratory environments, to the field in Kampala, Uganda. Kampala is an ideal 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 while ethnicity matters, the political situation in Kampala today is sufficiently stable and peaceful to permit

¹⁰ More generally, it is conceivable that the solution concept adopted by an analyst may be a function of ethnic categorizations.

¹¹ Whether this particular mechanism should be viewed as strategic behavior to avoid a social sanction or as behavior in accord with a norm against shirking is not clear. In any case, they are two sides of the same coin. The sanction is simply the off-the-equilibrium-path event that happens in the event of the violation of the

questions about ethnic identifications and attitudes to be asked and for research on social interactions across ethnic lines to be undertaken.

We began the study by conducting a survey of 594 randomly selected households in Kawempe, the poorest of Kampala's five divisions. In Figure 1, we use data from this survey to plot the relationship between community-level ethnic diversity (using a fragmentation index based on the ethnic demography of Kawempe's 74 LC1s as reported by the subjects) and public goods provision (measured in terms of whether, during the last six months, residents of the LC1 had organized community efforts in the area of crime prevention and security).¹² The data confirm that the negative relationship between diversity and public goods provision found elsewhere in the literature also exists in urban Kampala.¹³

[Figure 1 Here]

Apart from allowing us to corroborate the commonly found negative association between ethnic heterogeneity and public goods provision, the data from the household survey also put us in a position to select a narrower (and more manageable) area from which to recruit our experimental subjects. Given that our goal was to identify the mechanisms through which diversity undermined public goods provision, we deliberately selected a study area that was highly 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.

We drew a random sample of 300 subjects.¹⁴ Working with a representative sample of residents is an important innovation in our research design. Unlike existing studies in behavioral economics that employ

norm, and the norm is the social rule that develops over time as a consequence of the social equilibrium created by the potential for the sanction.

¹² To compute the fragmentation index, we use data gathered from a random sample of households in which we asked individuals to identify the four largest groups and the proportion of the population represented by each. At the level of the LC2, this subjective measure of fragmentation correlates with an ELF produced using census data at 0.66. We do not have census data on ethnic groups at the level of the LC1 and thus use the subjective measure instead.

¹³ Regression results (not shown) confirm the statistical significance of the relationship. The negative association between ethnic diversity and public goods provision is robust to the inclusion of controls for migration and wealth, as well as fixed effects for the parish in which the zone is located. It is also robust to the substitution of an alternative measure of public goods provision – whether communities have organized collectively to address issues of garbage collection. But our goal here (and in this project) is not to empirically identify the relationship between diversity and the underprovision of public goods. We simply want to show that the relationship demonstrated in the literature exists in Kampala, enabling us to focus specifically on the mechanisms that account for the relationship.

¹⁴ 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

university students or non-representative samples of volunteers, our sampling strategy puts us in a position to make inferences from the results of our experiments to the real world that our subjects inhabit.¹⁵ This, we argue, is especially important when the behaviors under study—in this case, how people condition play on the ethnicity of their partners—are not necessarily universal features of human behavior, but rather properties that may apply in different ways to different populations. In the conclusion of the paper, we show how the findings of our experimental games can be linked to differences in public goods provision and levels of organization in the communities from which our subjects were drawn.

In addition to collecting information about ethnic group membership, proficiency in first and second languages, and other background characteristics, we also recorded a series of five digital images including still and video images of each subject with varying levels of information.¹⁶ 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 at least some information.

The Public Information Box and General Set-up

As noted, the empirical strategy in this paper is to have subjects play multiple rounds of multiple games, sometimes with co-ethnics and sometimes with non-co-ethnics, and to compare patterns of play across co-ethnic and non-co-ethnic matchings and across games. Given this approach, it was critical to devise a means for subjects to ascertain whether or not the other players in the game were co-ethnics, without cueing them to our interest in how ethnicity shapes behavior. In two of the games that we will 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 two games were played using a computer interface. In these games, players could make inferences about the ethnic backgrounds of the other players through the information they gleaned from the pictures and/or videos that were made available to them in what we call the Public Information Box (PIB).¹⁷ The key attribute of the PIB is that, as its name suggests,

profitable; in addition, we contacted subjects the day before each scheduled session to remind them of their appointment.

¹⁵ Bahry and Wilson's (forthcoming) study of fairness norms in Russian republics provides a rare but important counterexample to the tradition of employing non-random samples.

¹⁶ The five levels of information are: a simple headshot, two short videos in which the subject greeted the camera in Luganda (Kampala's *lingua franca*) and in the respondent's first language, respectively and two additional images in which the subject also provided his or her given and family names. Since the images provide viewers with increasing levels of information from which they might make subjective inferences about the subject's ethnic background, we can test the effect of a subject's certainty about the ethnic background of his or her partner, a feature we exploit in other work.

¹⁷ Information treatments are now a common part of research in experimental economics. Bohnet and Frey (1999) demonstrate how decreasing social distance (by increasing the amount of information people have about their partners) affects play in dictator games. Charness and Gneezy (2000) introduce family names as an informational treatment. We see our PIB as a natural extension of this work, allowing for the introduction of multiple levels of information and the creation of common knowledge.

the information that it provides about the players in the game is provided publicly: Before each round of each game, all the players in the round were shown the same PIB that contained 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 telling them which player number they were for that round (i.e., “in this round you are player 2”). Figure 2 provides a sample screen shot of the PIB.

[Figure 2 Here]

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. Second, it allowed us to induce common knowledge about the information provided to all of the players in the game: each player was provided not just with information about who the other players were but also about what the other players knew about them, and that the other players knew what the first player 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 2a and 2b. In Figure 2a, the middle player’s picture is shown to the other players. He will therefore play the game knowing that the other players can see who he is. But in Figure 2b, 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. As we will show, 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. On average, about 5% of all rounds involved partners that knew each other. All results reported in the paper are robust to the exclusion of such rounds.

Each subject played all of the games multiple times but 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.¹⁸

Note that the design of the games imposes no structure based on our beliefs about what the ethnic identities of players “really are.” Beyond the selection of our site area (which entailed a deliberate focus on ethnically heterogeneous communities), our sampling of subjects at the LC1 level was undertaken randomly and independent of how subjects identified themselves ethnically. The matchings for the games were also

¹⁸ See Crawford and Heller (1990) for evidence of coordination deriving from repeated interaction.

constructed randomly and without reference to the ethnic identities of players.¹⁹ We identified subjects by ethnicity (based on their self-reported ethnic identifications in our pre-survey) only when we turned to analyze the games. Our interpretation of behavior assumes that subjects could identify whether or not the other players in the game were co-ethnics in this sense – either from the information they gleaned through face-to-face interaction or from the information they were provided in the PIB (an assumption we revisit below).²⁰ However, while we assume that subjects were able to distinguish co-ethnics from non-co-ethnics, we do not assume that we, the researchers, know what the salient lines of ethnic group division are. Thus, for each of the mechanisms we test, we present two different “cuts” at ethnic division: one where we code co-ethnic pairings as those between members of the same “tribe” (as determined by the national census categories) and a second where co-ethnicity is defined in terms of region of origin (i.e., Center, East, North, West).

V. Results

V.1 Preference Mechanisms

Commonality of Tastes

As described above, one of the simplest explanations linking co-ethnicity to cooperation is the argument that members of a given ethnic group have correlated preferences over outcomes. To probe the plausibility of this explanation we use simple survey techniques with our random sample of community members. We examine two types of questions. First, to what types of public goods do these individuals attach the highest priority? Second, how should these goods be provided? For example, should private or public means be used? Asking these questions allows us to examine whether the diversity of preferences over public goods that exists within Mulago-Kyebando is structured along ethnic lines or is a reflection of individual level idiosyncrasies.

When asked to describe their top priority in terms of public goods provision, members of six of the eight major ethnic groups in our sample reported that garbage collection deserved the most attention. In fact, on average, individuals in these six groups ranked their top three priorities identically, with security and the maintenance of drainage channels following garbage collection. Turning to strategies for providing public goods, there was overwhelming support across ethnic groups for a fee-for-service program to collect garbage.

¹⁹ The only exception was the Network Game, where our relatively small sample size required that we abandon random matching in order to insure a sufficient number of co-ethnic pairings.

²⁰ To aid them in making inferences about the ethnic backgrounds of the other players, we induced standardized priors about the distribution of ethnic groups in the subject population. So as not to prime subjects to ethnicity – something we were careful to avoid throughout the experiments (the project itself was advertised, and was known throughout the community, as “The Mulago-Kyebando Community Study”) – we embedded this information in a script that was read to subjects on their first day that provided other information about the subject pool, including the distribution of genders, ages, religions, lengths of residence in Mulago-Kyebando, and occupations.

At the same time, members of most ethnic groups expressed a willingness to contribute time or money to support community efforts to provide security and channel maintenance, if the alternative was worse public goods provision.

The results of our examination of group-specific differences are reported in Table 1. Ordinary least squares regression, with dummy variables for each of the major groups, suggests few ethnic group differences in preferences for public goods or opinions about how they should be provided. Although the results suggest some small differences in opinion, only two group dummies are statistically significant. Moreover, using an F-test, we cannot reject the null hypothesis that all ethnic group dummy variables are equal to zero for each of the six questions. Finally, the two group-specific dummies that show up as significant are rendered insignificant across specifications when fixed effects for location are included in the analysis. In short, there is little empirical support for the argument that ethnic groups possess correlated preferences over outcomes in Mulago-Kyebando.

[Table 1 Here]

Other Regarding Preferences: The Dictator Game with One-Way Information

To test the other regarding preferences variant of the preferences mechanism, we had subjects play a version of the standard Dictator Game in which the offerer is anonymous and the receivers' identities are known (as in Figure 2b) (Kahneman, Knetsch, and Thaler, 1986; for a review of dictator game results, see Camerer 2003).²¹ 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 must be due to differential levels of altruism vis-à-vis 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) being shown a PIB containing pictures and/or videos of themselves and two other players (the receivers). Subjects were then given ten 100 USH coins (about 60 cents, approximately equal to the per capita daily income in Uganda) and asked to divide it between themselves and the two other players any way they pleased. They were told to put the amount that they decided to keep directly into their pockets and to put the amounts that they wanted to allocate to the other players into envelopes which were then deposited into ballot boxes located directly below the pictures in the PIB of the players in question. Subjects were told that the envelopes would be

²¹ Some aspects of the games we employ differ from those commonly reported in the literature. While it naturally would be desirable to be able to compare the results of our games with those reported by other researchers, this goal was secondary to designing games that would permit us to make the clearest tests possible of the competing explanations we were studying. A large literature confirms that the results of experimental games are highly dependent on seemingly small game parameters, including the pro- or anti-social nature of the proctor's instructions and whether players think their behaviors are anonymous (e.g., Hoffman et al. 1994; Haley and Fessler 2005). However, because we are interested in comparing variation in patterns of play *across subjects and across pairings*, whatever biases our game designs may entail will be constant across them, and thus should not undermine our ability to make the inferences we seek to make.

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 his or her allocation, he or she signaled the enumerator, who returned from behind the screen and set up play for the next round.²²

In a second version of the game, they were given two 500 USh coins. This time they were instructed that no player (including the subject him- or herself) 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). As with all of the games described in this paper, the assignment of players into trios was random. Each subject played multiple rounds (an average of 2-5) 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 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. Because the 500 USh denomination games forced subjects to discriminate, altruism was somewhat harder to discern. The modal strategy (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 other players, and in one round a subject, contrary to rules, gave both coins to one of 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.

Although we observe high levels of altruism in the 100 USh denomination game, we find no evidence that this altruism was directed more at co-ethnics than at non-co-ethnics. Columns 1 and 2 of Table 2 report average offers to in-group and out-group receivers in the 100 USh denomination game under two different definitions of group membership (in-group member by ethnicity/tribe and by region of origin). In

²² We instituted various checks to ensure that our subjects understood the games they were playing. Most importantly, after a public session in which instructions were provided, subjects were tested one-on-one to determine whether they were aware of the rules of the game and the set of strategies 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 Uganda with no connections to the project was given the responsibility of soliciting from a focus group of our respondents information about the various games they were playing, how they interpreted them, the set of strategies from which they could choose, and so on. This back-translation provided us with powerful evidence that subjects understood the underlying behaviors each game sought to assess.

each case, we find that the difference in offers to co-ethnics and non-co-ethnics can not be distinguished from zero. We conclude that there is no evidence that subjects in this context exhibit any “taste” for discrimination.

[Table 2 Here]

Given the high degree of inequality aversion that our subjects exhibited, a taste for discrimination might be difficult to identify. Because the 500 US\$ denomination games render equitable distributions more costly, they may force subjects to reveal even slight preferences for the welfare of in-group members over out-group members. In column 4 of Table 2, we report the likelihood that players discriminate in favor of in-group members, restricting our attention to games in which players play with exactly one in-group member and one out-group member and in which the offerer retains one coin and offers the other.²³ The null hypothesis against which we compare our findings is that the co-ethnic will receive one coin 50% of the time (that is, the subject will choose which player to favor at random). Conditioning then only on cases of observed discrimination, we again find no evidence that they do so along ethnic lines.

V.2 Technology Mechanisms

The Lockbox Game

Whereas the other regarding preferences mechanism could be examined with a version of a standard experimental game that we took “off the shelf,” the two games we employ to test the “ethnicity as technology” mechanism are new. The Lockbox Game was used to examine the possibility 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.

In this game, Player 1 was trained in how to open a standard combination Master lock.²⁴ The player was given 20 minutes to practice opening the lock and then tested to ensure that he or she could open it without difficulty. After Player 1 had demonstrated mastery of this task, the lock was affixed to a lockbox containing 6,000 US\$. Player 1 was then matched with a partner (Player 2) and the pair was given ten minutes to work together to open the locked box. The two players were permitted to communicate orally but Player 1 was not permitted to touch the lock.²⁵ The players were instructed that they would share the money in the box if they succeeded in opening it before the ten minutes elapsed and that they would receive an additional bonus of 1,000 US\$ each if they were able to open the box within three minutes.

²³ Note that as a consequence the number of observations in the table is necessarily the same for each in-group and out-group category.

²⁴ These locks are not used in Uganda, and no subject reported ever having seen one.

²⁵ This was accomplished by having Player 1 stand on a balcony out of reach of the lockbox and talk to Player 2 through an open window.

All subjects played the Lockbox Game once, either in the role of Player 1 or Player 2.²⁶ Partners were matched by randomly dividing the twelve subjects in a given session into two groups: one group was designated to play the role of Player 1; the other to play the role of Player 2. After receiving instructions about how the game would be played, one of the six subjects that had been designated to play the role of Player 1 was selected at random. This player was then shown pictures of all six of the subjects in the Player 2 pool and asked to select one to be his or her partner in the game. A second subject was then selected at random from five remaining subjects in the Player 1 pool, shown pictures of the five remaining subjects in the Player 2 pool and asked to select one to be his or her partner. The process was repeated until the final subject in the Player 1 group was simply assigned to play the game with the last remaining subject in the Player 2 pool.

Fifty-nine percent of the pairs in the Lockbox game managed to open the chest within the ten minute time frame and received the 6,000 US\$ reward. Slightly less than one third of these successful pairs were able to complete the task in fewer than three minutes, and thus received the bonus. As the results reported in Table 3 make clear, however, whether or not the two players were members of the same ethnic or regional group cannot account for variation in success rates. Co-ethnic pairings experienced a marginally higher level of success (62.5%) than non-co-ethnic pairings (60%), but this difference is not statistically significant at conventional levels. Differences between players from the same region and from different regions are even smaller, and also statistically insignificant. As language is likely to be important in this context, we also looked at the success rates of pairings that shared the same primary language. We found that sharing a primary language did not matter; the difference that does appear to matter – and to matter critically – was whether or not the pair spoke *any* common language. When they did not, success rates, quite understandably, tumbled to 30%. While this may be considered an ethnic effect, it is one that is relevant for only a very small share (7%) of random pairings in our sample.

[Table 3 Here]

The absence of a difference in success rates between in-group and out-group pairs does not seem to be related to whether or not subjects had an option to choose their partners. If subjects anticipated that communication with their partner would be critical for completing the task successfully, we might have expected that those who were given the opportunity to choose their partners would have selected partners from their own group. We might also have expected these co-ethnic pairings to have achieved higher success

²⁶ Due to uneven numbers of subjects at some of our sessions, seven subjects played the role of Player 1 twice. Our results are robust to the exclusion of these pairings.

rates. Yet we find no evidence that subjects selected their partners along ethnic lines and we observe no increase in success rates when we limit the sample to those that could choose their partners.²⁷

The lack of a technological coordination effect of this form is surprising, especially given the linguistic diversity in Mulgao-Kyebando. We were, however, able to replicate this (non-)finding in the examination of a second, more symmetric game (not shown), in which two players were asked to put together two parts of a jigsaw puzzle independently so that the two parts would ultimately fit together.²⁸ In the Puzzle Game, all subjects were randomly assigned to their partners. The success rate for this game was 40% but, as with the Lockbox Game, there was no evidence that shared ethnicity or regional background had a statistically significant, positive effect on the ability of pairs to complete the task successfully.

The Network Game

A second way in which co-ethnicity may provide a “technology” is through strategy options made available to players due to the existence of shared networks. The mere existence of a network may facilitate the free flow of information about trustworthiness and the sanctioning of individuals who defect from their obligations or commitments.

To study this possibility, we designed a game to measure precisely the network “distance” between two arbitrary individuals in Mulago-Kyebando. We seeded Mulago-Kyebando with 148 randomly selected “targets” – individuals that we would send our subjects to try to find. We collected a small amount of background information about each target, including his or her ethnicity and birthday. We also took their photo with a digital camera. We told them 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. We asked them to tell the person their birthday or in case this was not known, a parent’s name, and we gave them 3,000 USh for participating in the experiment.

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. Each runner was randomly assigned a different target. They were shown the target’s photograph and given a sheet of paper with the target’s name and the parish (LC2) in which he or she

²⁷ The data do suggest, however, that success rates are related to the educational characteristics of the players. Pairs containing players that both had more than a primary school education completed the task successfully 79% of the time, as compared with a 45% success rate for pairs with little education. These results suggest that the game may have been sufficiently difficult that the education effect trumped whatever productivity gains might have been realized from membership in the same ethnic group. Indeed, we find some evidence that, among pairs with high education, shared language group membership significantly raised success rates.

²⁸ In this Puzzle Game, subjects were randomly divided into pairs and seated on either side of a small table. A low divider 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 divider. Each player was given two of four pieces of a jigsaw puzzle and given three minutes to work with his or her partner to assemble the pieces so that they would fit together with his or her partner’s pieces to complete the puzzle. If after three minutes the two

resided. We instructed them to phone us when they found the target and tell us the target’s message. This provided a check that they had actually found the person we had sent them to find. Runners were given 5,000 USh in transport money to defray the cost of hiring a boda boda 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 person in three hours and that this sum would decline at a rate of 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.²⁹

Forty-nine of the 148 runners (or just over 33%) managed to track down their targets.³⁰ However, the success rate among runners whose targets happened to be co-ethnics was significantly higher (43.1%) than the success rate among runners whose targets were from other ethnic groups (27.8%) – see Table 4. Being sent to find a target from the same region had no positive impact on the likelihood of success, suggesting that the scale at which identity groups are defined is important for assessing the relevance of networks. To explore scale effects further, we also report results when an intermediate definition is used (identified in Table 4 as “Aggregate Ethnicity”); co-membership in this broader ethnic community, we find, is associated with greater success but the added effect is substantively small.³¹ These findings suggest that ethnic groups do indeed possess social networks that facilitate the flow of information useful for the tracking down of fellow group members, but that these networks do not extend to communities based on region – or at any rate not those found in Mulago-Kyebando.

[Table 4 Here]

V.3 Strategy Mechanisms

The Dictator Game with Two-Sided Information

To test the institutions/norms of reciprocity mechanism, we again employed a version of the standard Dictator Game. The general set-up is exactly as described above for the Dictator Game with one-sided information. This time, however, we focus on a version of the game in which receivers have (and offerers

halves of the puzzle matched, the players were paid 2,000 Ush. Each subject played the Puzzle Game three times, each time with a different partner, and each time with a slightly different puzzle.

²⁹ Runners who managed to find their targets were extensively debriefed when they came back to the laboratory to be paid. We also debriefed a random sample of 40 of the unsuccessful runners.

³⁰ We surmise that many runners simply pocketed the 5,000 USh transport money and never tried to find their targets.

³¹ “Aggregate Ethnicity” is a measure of co-ethnicity slightly more restrictive than region, but more encompassing than primary tribal identities. To construct it, we combine a set of closely related ethnic groups into their broader ethnic families. The Bafumbira are joined with the Barundi and Banyarwanda; the Banyankole are merged with the Bakiga; and the Batoro are joined with the Bunyoro.

know they have) information about who the offerer is (as in Figure 2a).³² Such information is a necessary condition for the utilization of targeted punishment strategies. We argue that in this version of the game, an offerer's behavior vis-à-vis a receiver can be interpreted as a joint product of the offerer's altruism toward that receiver and the offerer's concerns about being seen to violate social norms and/or opening him or herself up to the possibility of social sanctions. If social institutions are such that co-ethnics are expected to play more co-operative equilibria with each other, then we expect to observe this as differences in patterns of play across Dictator Games with one-sided and two-sided information.

In interpreting this game, it is important to emphasize that the Dictator Game, by construction, does not allow for the possibility of punishment *within the context of the game*. We are interested in whether subjects, entering a laboratory environment, act *as if* punishment is possible or *as if* enforceable norms against shirking apply (see Hoffman et al. 1996 for a discussion of this approach). The Dictator Game with one-sided information discussed in the previous section excludes the possibility of extra-game punishment, since the offerer is anonymous. The Dictator Game with two-way information, however, does not. We employ it to test whether patterns of play in co-ethnic and non-co-ethnic pairings differ when the threat of sanctions and or norm violation obtain.

Each subject played approximately two rounds of the 100 USh denomination game and approximately four rounds of the 500 USh denomination game. This yielded 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. On average, patterns of play were similar to those found in the Dictator Game with one-sided information. The modal strategy in the 100 USh 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.³³ In the 500 USh denomination game, aggregate results were again similar to those reported in the previous section: 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.

Strikingly however, while there were little differences in average play, there were consistent differences in the way subjects played with different kinds of partners. As shown in Table 5, when subjects knew that they could be observed, they offered more to co-ethnic receivers than to non-co-ethnic receivers.³⁴

³² In practice, both pieces of information are provided together since payments are given to receivers as they view the public information box for the corresponding game.

³³ Offers in this version of the game were slightly lower than in the anonymous offerer Dictator Game, but these differences are not statistically significant.

³⁴ Given the distribution of ethnic groups in our sample, the correlation between co-ethnicity and membership in the dominant, Baganda tribe is high. Consequently, we cannot identify an independent effect of co-ethnicity when we include a dummy variable for Baganda. When a Baganda term is introduced, the co-ethnicity measures maintain the same sign but, with the introduction of a highly correlated variable,

In the 100 USh game, this difference – on the order of 24 USh for co-ethnics and 22 USh for people from the same region of the country– is strongly significant, even if it is substantively relatively small (equivalent to about a 10% increase in offers). Similar patterns are observed in the discrimination game: in cases where a player has a choice between favoring a co-ethnic or a non-coethnic and discriminates, she favors the co-ethnic 56% of the time. The same pattern also holds for pairings that match players from the same region. Again, the difference, while small, is in all cases significant at conventional levels.³⁵

[Table 5 Here]

These findings provide strong support for the efficacy of the institutions/norms of reciprocity mechanism in distinguishing patterns of behavior in ethnically homogeneous and ethnically heterogeneous settings. Comparing the results in Table 5 to those in Table 2, we find that there is no evidence of within group bias (along co-ethnic or co-regional lines, in the 100 USh or the 500 USh games) when players make their offers anonymously, and that there is strong evidence of bias when players are observed in their actions. A statistically more difficult test would be one in which, rather than analyzing each game in isolation, we estimate the joint effect of co-ethnicity and information in a single framework. We conduct such a test by examining the interaction between co-ethnicity and the fact that a player is observed. In all cases, we find that the marginal effect of co-ethnicity in the absence of information is close to 0, while the marginal effect of co-ethnicity given information is of the same sign and magnitude as in Table 5. However, while the interactive term is significant at the 95% level for co-regional pairings in the 100 USh game, the associated t -statistic drops below conventional levels in the other three cases. Thus, while the evidence that players exhibit in-group bias *if* they are observed is strong, our confidence that they do so *only if* they are observed is weakened somewhat by our uncertainty regarding their behavior when they are not observed.

VI. Discussion and Interpretation

The results of the games we have analyzed are striking, both for what we find and what we fail to find. We find no evidence for the other regarding preferences or the correlated preferences variants of the preference mechanism. If co-ethnics are more effective at producing public goods, it is not simply because they care about the same things or value more highly the welfare improvements experienced by other group members. Nor do we find evidence for the communication variant of the technology mechanism: co-ethnics are not

significance is lost. We can rule out, however, the possibility that co-ethnic effects obtain only within Baganda pairings: while Baganda give more on average, they are no more likely to give more specifically to co-ethnics.

³⁵ Of the four results in Table 5, three remain significant at conventional levels when we analyze these games using regression analysis and allow for correlations between all decisions made by an individual player. In one case, the discrimination game with co-ethnic pairings, the p value rises to .105, just shy of 10% significance.

more effective working together on joint tasks. However, we do find evidence for one version of the technology mechanism: co-ethnics appear to be more closely linked through network structures in our project area. They are more “findable” than non-co-ethnics. Further, we find compelling evidence that players condition their strategies on the ethnicity of the other players in the game: our findings suggest that ethnically homogeneous communities have an advantage in providing public goods because they make it easier to sanction community members who fail to contribute to collective endeavors.

Our interpretation of the results has implicitly assumed that the subjects in our games can correctly categorize their partners as co-ethnics or non-co-ethnics. If a subject playing the offerer role in the Dictator Game is matched with an in-group member and an out-group member and is faced with a decision about how much money to allocate to each of them, we interpret the subject’s decision as if he or she knew the in-group or out-group status of the other players. But it is quite possible that the subject could not discern, or even misconstrued, the other players’ ethnic group memberships from the information provided about them in the PIB. If this were the case, then the subject would have been playing the game differently from what our interpretation suggests.

To test whether or not this was the case, we invited all subjects back to our laboratory at the very end of the project, after all of the other experimental games had been played, to play what we call the Identification Game. In the Identification Game, subjects were shown images of every player they had been matched with in all of the previous computer-based games and asked to guess the person’s ethnic background.³⁶ We then used the data collected from the Identification Game to reanalyze our results in the Dictator Game with two-sided information, this time coding co-ethnic and non-co-ethnic pairings in terms of the match between the subject’s ethnic group membership (as reported to us in the subject’s pre-survey) and the ethnic backgrounds that the subject *thought* his or her partners possessed. When reinterpreted in this fashion, our results are unchanged. We continue to find that subjects were more generous to co-ethnics than to non-co-ethnics when they know that their actions are being observed, which we interpret as evidence of the operation of within-group norms of reciprocity and/or the greater threat of social sanctions within groups than across group lines.

In our earlier discussion of the rival channels through which ethnicity may operate, we highlighted the difficulty of distinguishing between the strategy and technology mechanisms. Given the results of our investigations, we are confronted by precisely this inferential problem. One might argue that individuals in our sample employ a strategy that is conditioned on ethnicity: a player cooperates with co-ethnics because co-ethnics will be more likely to punish him if he fails to cooperate than would non-co-ethnics in the same situation. An alternative explanation, however, is that the strategies we observe might not be conditioned on ethnicity but simply on “findability.” From this perspective, players cooperate conditional on how easily they

³⁶ The instructions for the Identification Game were the first time in the experiment that subjects were informed that ethnicity was a central variable of interest in the study.

believe they can be found by their partner, whether or not they are co-ethnics. However, it so happens that co-ethnics are easier to find and thus, empirically, cooperation is more commonly conditional on co-ethnicity. In this second case, if ethnicity matters at all the underlying reason is technological not strategic.

The distinction between these two accounts is subtle. However, we have two pieces of information that favor the second account over the first. The first piece of information allows us to work out whether reciprocal behavior continues to obtain even in the absence of social networks. According to the first (strategy) account it should, while the second (“findability”) account suggests it should not. Our test involves allowing punishment to take place *within the context of the game*. In doing so, we remove the advantage that co-ethnics might possess in their ability to use networks to punish deviators in post-game interactions. The game we employ involves a third party that observes actions by two other players engaged in a Prisoners’ Dilemma (Fehr and Fischbacher 2004). In this game, the two players are not able to observe each other’s actions and cannot punish each other. A third player, however, observes the actions of the first two players and can choose to punish either of them for their behavior. We find, consistent with the strategy account, that punishment occurs and that anticipation of punishment incites greater contributions. However, we find that behavior is not structured along ethnic lines: given the possibility of punishment, players do not cooperate more with co-ethnics. The shadow of the future would appear to incite all players to behave more cooperatively, not simply co-ethnic pairings.

The second test allows us to work out whether reciprocal behavior obtains even among non-co-ethnic pairings when the offerer is easily findable. According to the strategy account it should not, while the “findability” account suggests it should. In this case, our test involves using data from the Network game to produce a measure of the findability of players based on basic (but non-ethnic) demographic data: age and gender. We then examine whether the ease with which someone can be found increases the offers that he or she makes in the Dictator Game with two-sided information. We find strong evidence that it does: players that can be easily located offer more even if they are not co-ethnics. In particular, we find that older players are significantly more findable within their communities and are more generous givers (we do not, however, find evidence that they condition their giving on two-sided information more than younger players). This evidence again supports the account of networks as an ethnic technology that facilitates sanctioning over the strategy interpretation.

VII. Conclusion

We began this paper with a focus on the pressing challenge of preventing and responding to crime in Kampala’s urban slums. After the government stopped funding local police, neighbors joined together to organize a successful community policing scheme in Katale, yet failed to act collectively in Kifumbira. It turns out that the relative underprovision of public goods in Kifumbira extends beyond security. While both communities have a network of latrines, only focus group participants in Kifumbira reported that certain

population centers were underserved and needed more latrines. Similarly, both zones are fortunate to have a number of water standpipes (installed mainly by international NGOs), but slightly more than a third are non-functional in Kifumbira while only ten percent are not working in Katale. Facing similar challenges relating to garbage, drainage, and security – public goods that are not provided by the government and which communities must therefore provide themselves – our interviews suggest that local councils have succeeded in mobilizing community participation in Katale while these problems have remained unaddressed in Kifumbira. These differences exist even though both Kifumbira and Katale are located in one of Kampala’s poorest slum areas.

Kifumbira is not unlike other ethnically heterogeneous communities that suffer from inferior public goods provision. From Pakistan to Indonesia and from rural Kenya to the United States, a growing literature suggests that the relationship between diversity and the underprovision of public goods is not simply an artifact of differences in wealth or patterns of residential mobility. It appears that ethnic diversity has an independent impact on the likelihood that communities can organize collectively to improve their welfare. The question is why.

In this paper, we have offered an innovative empirical strategy to distinguish among three plausible families of mechanisms that might explain how ethnic diversity undermines the provision of public goods. 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 evidence that successful collective action among homogenous ethnic communities is attributable to the existence of norms and institutions that facilitate the sanctioning of non-contributors.³⁷ The efficacy of these strategies, the evidence suggests, stems from the fact that co-ethnics enjoy greater proximity to one another through social networks: they are more mutually “findable” – a feature that permits co-ethnics to employ strategies that rely on reciprocation even though, in the context of our games, they are mutual strangers. This finding is consistent with previous work on ethnic diversity and public goods provision (Besley et al 1993; Miguel and Gugerty 2005). Yet our approach also allows us to demonstrate that rival explanations are not at work in this context. We find no evidence for a commonality of tastes within ethnic groups; for greater degrees of altruism toward co-ethnics; or for an impact of shared ethnicity on the productivity of teams. Moreover, we are able to report additional results consistent with our main conclusions: namely that generalized reciprocity is observed when

³⁷ It bears mention that Mulago-Kyebando, while ethnically diverse, is home almost exclusively to Bantu-speaking tribes from western and central Uganda – tribes currently united in a de-facto political coalition that has ruled the country since 1986. Thus the finding of a correlation between high levels of diversity and low levels of public goods provision might surprise some, given the close historical relationships among the different tribes. The fact that social sanctioning mechanisms appear to function only within ethnic groups, even in this area, suggests how high the barriers to inter-group cooperation may actually be in practice. The country’s major ethnic cleavage separates members of Bantu-speaking tribes (resident largely in the South) from descendent of Nilotic tribes (resident in the North). Thus the mechanisms linking ethnic diversity to

barriers to punishment are removed, that higher contributions are made when offerers are findable (even if they are not co-ethnics), and that our results persist even when we account for the fact that individuals are not always certain about the ethnic identities of people they encounter.

Some might reasonably ask whether our subjects' behavior in experimental games captures something about the real world in which they live. Unsurprisingly, given the results we observe in the lab, when asked to indicate how much they trust various groups in their community, nearly 50% of our subjects reported trusting members of their own ethnic group to "a great" or "very great" extent. Only 20% reported trusting members of other ethnic groups to the same degree. Our interviews with LC1 chairmen provide further support for the idea that norms of reciprocity typically operate within, rather than across, ethnic lines. Although local council leaders consistently described the challenge of soliciting voluntary contributions from community members for local public goods, they also highlighted the successful functioning of ethnic associations:

"We have some so-called 'groups' that are organized on a tribal basis. For example, the Bafumbira lend to each other. They are very organized in that way. They have a society that is formalized with executive members, and if you need money and are a Bafumbira, you can go to them. If you need burial money, for example, they apply and can be lent money... If one of them is arrested, they try to raise money to pay his fine." (Lower Nsooba)

Describing a group of people from the Northeast: "They form their own association. When they get a problem, for example, losing a person, they come together; in case of a person having an economic problem, they come together [...] they mobilize funds among themselves to solve the problem." (Kisalosalo)

But in being asked to describe how interethnic cooperation is achieved in the communities in which they live, our interviews also highlighted strategies that have been adopted to minimize the negative impact of ethnic diversity on public goods provision. Explaining how a local credit association avoids the non-contribution problems experienced by the local council, one chairman said:

"Each member contributes said amount each month, and every month a ceremony is held whereby two "brides" are chosen within each unit to receive the total amount. Of that total amount, each has to give Ush 10,000 to the organization, which puts it in a savings account in a bank... The brides are chosen by lottery: the Chairperson of the unit writes "bride" on two pieces of paper, and leaves the rest of them blank... The organization's pool is used to acquire things for the whole group: plastic chairs that they rent out for ceremonies, plates, cups for community gatherings; they are planning on acquiring tents next... There are problems with people not contributing, but only to the extent that the amount paid to the brides is delayed, never to the extent that amounts are not paid... Ethnic distribution of membership is very mixed and hard to tell because it is not information that is collected." (Kiwonvu)

the underprovision of public goods in urban Uganda may be different than those that would operate in a community composed of Southern and Northern ethnic groups (if such a community, in fact, existed).

In the local credit association, individuals are embedded in a repeated game: their identities are known with certainty. Given past contributions and potential gains from cooperation, defection is more costly. Moreover, sustained interaction decreases the technological impediments to sanctioning. The story of the credit association is consistent with our experimental findings, suggesting as it does that generalized reciprocity can emerge as a norm (in multi-ethnic settings) when impediments to sanctioning are removed.

Addressing inefficiencies in public goods provision that result from ethnic diversity, therefore, does not require strategies of greater segregation, as many preference based explanations might suggest. What are needed are policies and strategies that promote successful cooperation across ethnic groups. Sometimes, such cooperation can be achieved by small groups (in credit associations and burial groups) where interactions are structured to facilitate sanctioning. For communities more generally, however, effective cooperation will depend on the creation of social networks among individuals that extend across ethnic groups – networks that allow for effective sanctioning outside of the context of a single-shot interaction. When individuals believe that their behavior is observed by others and that their reputation may influence opportunities for cooperation in the future, norms of reciprocity can be sustained even in the absence of shared ethnic identities.

Figure 1: Ethnicity and Local Public Goods Provision in Kampala





Figure 2: The Public Information Box



2a) Public Information Box with Two-Sided Information. In this case player 2 is “seen” by all players. (note that the faces displayed are those of members of our research team and not of actual subjects)



2b) Public Information Box with One-Sided Information. In this case player 2 is “unseen” by all players. (Note that the faces displayed are those of members of our research team and not of actual subjects)

Table 1: Variation in Policy Preferences Across Ethnic Communities, by Major Ethnic Group

	First Priority for Public Goods Provision			How Public Goods are to be Provided		
	Drainage	Garbage Collection	Security	It's better to have garbage collection, even if we have to pay a fee ¹	it's better not to have to pay anything or to volunteer for patrols, even if that means security is low ²	it's better to have well-maintained drainage channels, even if we have to make contributions of money or labor ³
Ankole	0.067 [0.74]	0.191 [1.90]*	-0.021 [0.21]	0.025 [0.30]	-0.057 [0.82]	0.004 [0.05]
Bakiga	0.101 [0.82]	0.158 [1.16]	-0.054 [0.41]	0.158 [1.43]	0.009 [0.10]	-0.163 [1.53]
Banyarwanda	0.080 [0.67]	0.121 [0.91]	0.112 [0.87]	0.100 [0.93]	0.126 [1.38]	0.046 [0.44]
Basoga	0.040 [0.29]	-0.169 [1.08]	-0.024 [0.16]	0.043 [0.34]	-0.033 [0.31]	-0.011 [0.09]
Batoro	0.003 [0.02]	0.205 [1.60]	-0.035 [0.27]	0.107 [1.03]	0.111 [1.25]	0.053 [0.53]
Bunyoro	0.040 [0.29]	0.195 [1.24]	0.067 [0.44]	0.043 [0.34]	0.149 [1.38]	-0.011 [0.09]
Bafumbira	0.317 [2.95]***	-0.042 [0.35]	-0.138 [1.17]	-0.125 [1.29]	-0.024 [0.29]	-0.129 [1.38]
Constant	0.233 [5.89]***	0.442 [10.08]***	0.388 [8.99]***	0.775 [21.73]***	0.124 [4.09]***	0.829 [24.13]***
Observations	249	249	249	249	249	249
F-Statistic	1.32 [0.24]	1.45 [0.19]	0.41 [0.90]	0.87 [0.53]	0.95 [0.47]	0.72 [0.66]

Ordinary least squares regression. Analysis is limited to ethnic groups with more than 10 individuals in our sample. Baganda is the excluded category. Absolute value of t statistics in brackets. * significant at 10%; ** significant at 5%; *** significant at 1%.

¹Alternative option was: “It is better to have free garbage collection, even if the quality of garbage collection is low.”

²Alternative option was: “It is better to have good security in our zone, even if we must pay for guards or volunteer for patrols.”

³Alternative option was: “It is better not to have to clear debris out of drainage channels, even if this means that we sometimes have floods.”

**Table 2: Average Offers and Discrimination Rates
in Dictator Game with One-Way Information**

	3 Person Dictator Game (100 Shilling Units)			Discrimination Game (500 Shilling Units)	
	Out-group offer (N) (1)	In-group offer (N) (2)	Difference (p-value) (3)	Probability of In-group discrimination (N) (4)	p-value, (null= 50%) (5)
Ethnicity	229 (941)	236 (283)	7 (0.39)	49% (112)	0.61
Region	232 (846)	228 (378)	-4 (0.57)	51% (157)	0.44

Notes: Column 3 presents results of two-sample t-test. Column 5 presents results of a one-tailed binomial probability test.
 *** Significant at 1%; ** Significant at 5%; * Significant at 10%. 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.

Table 3: Success Rates in Lockbox Game

	Co-Ethnic Pairing (N)	Non-Co- Ethnic Pairing (N)	Differenc e (p-value)
Ethnicity	62.50 (32)	60.00 (115)	2.5 (0.80)
Region	60.00 (45)	60.78 (102)	-0.01 (0.93)
Any Common Language	61.87 (139)	30.00 (10)	31.87** (0.05)

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

Table 4: Success Rates in Network Game

	Co-Ethnic Pairing (N)	Non-Co- Ethnic Pairing (N)	Differenc e (p-value)
Ethnicity	43.14 (51)	27.84 (97)	15.30* (0.06)
Region	37.73 (61)	29.89 (87)	7.82 (0.32)
Aggregate Ethnicity	44.23 (52)	27.08 (96)	17.15** (0.03)

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

Table 5: Offers and Discrimination Rates in Dictator Game with Two-Way Information

	3 Person Dictator Game (100 Shilling Units)			Discrimination Game (500 Shilling Units)	
	Out-group offer (N) (1)	In-group offer (N) (2)	Difference (p-value) (3)	Probability of In-group discrimination (N) (4)	p-value, (null= 50%) (5)
Ethnicity	220 (748)	244 (230)	24*** (0.008)	56% (204)	0.05**
Region	219 (662)	240 (316)	22*** (0.009)	56% (271)	0.03**

Notes: Column 3 presents results of two-sample t-test. Column 5 presents results of a one-tailed binomial probability test.
 *** Significant at 1%; ** Significant at 5%; * Significant at 10%. 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.

Bibliography

- Adsera, Alicia and Carles Boix, 2004, "Constitutional Engineering and the Stability of Democracies," unpublished manuscript, University of Chicago.
- Alesina, Alberto, Reza Baqir and William Easterly, 1999, "Public Goods and Ethnic Divisions," *Quarterly Journal of Economics* 114 (November), pp. 1243-1284.
- Alesina, Alberto and Eliana LaFerrara, 2002, "Who Trusts Others?" *Journal of Public Economics* 85 (2), pp. 207-234.
- Alesina, Alberto and Eliana LaFerrara, 2005, "Ethnic Diversity and Economic Performance," *Journal of Economic Literature* 63 (September), pp. 762-800.
- 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*.
- 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), pp. 792-810.
- Bohnet, Iris and Bruno Frey, 1999, "Social distance and other-regarding behavior in dictator games: Comment," *American Economic Review*, 89, pp. 335-339.
- Camerer, Colin, 2003, *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton: Princeton University Press).
- Cederman, Lars-Erik and Luc Girardin, forthcoming, "Beyond Fractionalization," *American Political Science Review*.
- Chandra, Kanchan, 2005, "What is Ethnic Identity and Does it Matter?" *Annual Review of Political Science*, 9, pp. 397-424.
- 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, V. P. and H. Heller, 1990, "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica* 58, pp. 571-595.
- Deutsch, Karl, 1966, *Nationalism and Social Communication* (Cambridge: The MIT Press).
- Easterly, William and Ross Levine, 1997, "Africa's Growth Tragedy: Policies and Ethnic Divisions," *Quarterly Journal of Economics* 112 (November), pp. 1203-1250.

- Eckel, Catherine and Philip Grossman, 2001, "Chivalry and solidarity in ultimatum games," *Economic Inquiry*, 39, pp. 171-188.
- Fafchamps, Marcel and Bart Minten, 2002, "Returns to Social Network Capital Among Traders," *Oxford Economic Papers* 54 (April), 173-206.
- 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 (4), p. 715-35.
- Fehr, Ernst and Simon Gächter, 2000, "Cooperation and Punishment in Public Goods Experiments," *American Economic Review* 90 (September), pp. 980-994.
- Fehr, Ernst and Urs Fischbacher, 2004, "Third-party punishment and social norms," *Evolution and Human Behavior*, 25, pp. 63-87.
- Fershtman, Chaim and Uri Gneezy, 2001, "Discrimination in a Segmented Society: An Experimental Approach," *Quarterly Journal of Economics* 116 (February), pp. 351-377.
- Ghosh, P. and D. Ray, 1996, "Cooperation in Community Interaction without Information Flows," *Review of Economic Studies* LXIII, pp. 491-519.
- Gil-White, Francisco, 2004, "Ultimatum Game with an Ethnicity Manipulation: Results from Kohvdiin Bulgan Sum, Mongolia" in Joseph Henrich et al. (eds.), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies* (New York: Oxford University Press), pp. 260-304.
- Glaeser, Edward, David Laibson, Jose Scheinkman, and Christine Soutter, 2000, "What is social capital? The determinants of trust and trustworthiness," *Quarterly Journal of Economics*, 115, pp. 811-846.
- Habyarimana, James, Macartan Humphreys, Daniel Posner, and Jeremy Weinstein, 2005, "Social Focal Points," unpublished manuscript, UCLA.
- Haley, K.J. and Fessler D.M.T., 2005, "Nobody's watching? Subtle cues affect generosity in an anonymous economic game," *Evolution and Human Behavior*, 26 (3), pp. 245-256.
- Hardin, Russell, 1995, *One for All: The Logic of Group Conflict* (Princeton: Princeton University Press).
- Hoffman, Elizabeth, Kevin McCabe, Keith Shachat, and Vernon L. Smith, 1994, "Preferences, property rights, and anonymity in bargaining games," *Games and Economic Behavior*, 7, pp. 346-380.
- Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith, 1996, "Social Distance and Other-Regarding Behavior in Dictator Games," *American Economic Review* 86 (June), pp. 653-660.

- Kahneman, Daniel, Jack Knetsch, and Richard Thaler, 1986, "Fairness as a constraint on profit seeking: Entitlements in the market," *American Economic Review*, 76, pp. 728-741.
- Khwaja, Asim 2002, "Can Good Projects Succeed in Bad Communities? Collective Action in the Himalayas," unpublished manuscript, Harvard University.
- Landa, Janet Tai, 1994, *Trust, Ethnicity, and Identity: Beyond the New Institutional Economics of Ethnic Trading Networks, Contract Law, and Gift-Exchange* (Ann Arbor: University of Michigan Press).
- 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, pp. 2325-2368.
- Okten, Cagla and Una Okonkwo Osili, 2004, "Contributions in Heterogeneous Communities: Evidence from Indonesia," *Journal of Population Economics* 17 (December), pp. 603-626.
- Platteau, Jean-Philippe, 1994, "Behind the Market Stage: Where Real Societies Exist (Parts I and II)", *Journal of Development Studies* 30 (April and July), pp. 533-577 and 753-817.
- Poterba, James, 1997, "Demographic Structure and the Political Economy of Public Education," *Journal of Policy Analysis and Management* 16 (1), pp. 48-66.
- Spolaore, Enrico and Romain Wacziarg, 2006, "The Diffusion of Development," unpublished manuscript, Stanford University.
- Tajfel, Henry. 1982, "Social Psychology of Intergroup Relations," *Annual Review of Psychology* 33, pp. 1-39.
- Taylor, Michael, 1988, "Rationality and revolutionary collective action," in Michael Taylor (ed.), *Rationality and revolution* (New York: Cambridge University Press).
- Vigdor, Jacob, 2004, "Community Composition and Collective Action: Analyzing Initial Mail Response to the 2000 Census," *Review of Economics and Statistics* 86 (1), pp. 303-312.