Chapter 5

A Closer Look at Reciprocity

Our inquiry into the nature of interaction among co-ethnics, as compared to non-co-ethnics, has pinpointed a number of places in which the co-ethnic advantage may reside. In the last chapter, we found strong evidence for two technological advantages—co-ethnics are better able to identify and to infer otherwise unobservable characteristics about one another, and co-ethnics can locate each other more easily than non-co-ethnics. Both mechanisms increase the likelihood that co-ethnics will be in a position to find and sanction partners for uncooperative behavior. We also discovered evidence that the strategies people play are conditioned on identity: egoistic players treat co-ethnics better *if and only if* they are seen to be doing so. Taken together, these findings provide the foundation for an answer to our puzzle. Norms of reciprocity built on the threat of sanction may help to account for why more homogeneous communities (where a greater share of interactions are among co-ethnics) are advantaged in the provision of public. Yet a weakness of our evidence is that it derives from behavior in games that capture *aspects* of a collective action dilemma rather than behavior in games in which players are actually facing a collective action problem. Moreover, while we have provided evidence that co-ethnics possess a technology that makes sanctioning easier, we have not provided direct evidence that players actually do sanction, or believe that others will sanction, and respond accordingly.

In this chapter, we push our investigation further. Recognizing that the results we have presented so far are consistent with a variety of strategy selection mechanisms, we look first for evidence that reciprocity norms in Mulago-Kyebando indeed depend on the threat of sanction. To do this, we explore how our subjects play a game that captures the challenge of public goods provision directly. Using a classic public goods problem studied by game theorists, the prisoners' dilemma, we examine how behaviors differ across types of players and how play is structured by co-ethnicity—an important check on the consistency of our results in Chapter 4. Then, within the context of the prisoners' dilemma, we introduce the possibility of sanction by a third player who observes how others play the game. We explore the data for evidence that individuals anticipate punishment, act more cooperatively when sanctioning is possible, and that players do in fact engage in costly punishment. We find strong evidence that they do all of these things.

We then take one additional step in trying to sort out how exactly in-group reciprocity norms function in support of collective action. We ask whether it is the case that, even among players that can be easily found (co-ethnic or non-co-ethnic), individuals act in a more reciprocating manner when playing with co-ethnics. Or is it precisely *because* they are more easily found by co-ethnics that individuals reciprocate more when paired with someone from their ethnic group? Although these two accounts may appear quite similar—both depend on sanctioning behavior enforced through networks—they provide distinct logics for how ethnicity figures in the solution of collective action problems. The first story is purely about strategy—players condition their behavior on the ethnicity of the other player, *ceteris paribus*. The second emphasizes a technology: players condition their behavior on a particular characteristic of others—"findability"—that happens to be correlated with co-ethnicity. According to this logic, we observe a strategy of reciprocation, but not one structured by ethnicity. Disentangling the strategy from the findability story is essential if we are to figure out how co-ethnics achieve successful cooperation in Mulago-Kyebando.

Do Players Cooperate with Co-ethnics because They Fear Sanctioning? Evidence from the Prisoners' Dilemma

The public goods game most studied by game theorists is called "The Prisoners' Dilemma." In the classic account, the story describes the dilemma of two prisoners facing interrogation separately for their role in a crime. Each must decide whether or not to provide information about the other prisoner. The tragedy of the dilemma is that, although both would be better off if they could keep quiet, each has a private incentive to provide information about the other. Dramatic as it may seem, the logic of the prisoners' dilemma is the same as that underlying real-world, everyday decisions faced by community members in Mulago-Kyebando. Many rotating credit and savings organizations—a key source of financial capital for business in the informal sector—have the structure of a public goods game. Individual decisions about whether to report suspicious behavior in the neighborhood or whether to show up for a community initiative to clear drains provide other clear examples.

Consider again the example of the two shopkeepers, Alex and Kenneth, discussed in Chapter 1. They would like to build a road to facilitate the movement of customers and goods to and from their shops, and the road requires an investment from each of them if it is to be functional and reliable. Both Alex and Kenneth have $100 available to be invested in the road. Assume that

every $100 dollars put into the road translates into benefits of just $75, but that these benefits are *public* in the sense that can be enjoyed by both people, whether or not they made an investment.[1] In this case, while the cost of contributing $100 is $100, the total (social) benefits, if both invest, are 2 X $75, or $150. Privately, this is a costly investment, but socially it is profitable. Our challenge is to work out the conditions under which a road will be constructed.

Imagine that both are egoists—they care only about the private benefits that accrue to themselves. If both players invest, the value of the road is $150. Each player receives a net benefit of $50. What if Alex invests in the road and Kenneth does not? In this case, the value of the road is $75. The net gain to Alex is -$25 and the net gain to Kenneth is $75. While Kenneth is a net beneficiary in this scenario, Alex is a net loser—he is what game theorists traditionally call a "sucker" while Alex is a free-rider. In fact, Kenneth is better off than he would have been had he contributed; Alex is worse off than he would have been had he *not* contributed. Outcomes are reversed if only Kenneth invests and Alex keeps his money. If both Alex and Kenneth refuse to invest, no road is built and no one experiences net gains or losses. The basic logic of this situation can be represented in the matrix in Figure 5.1(a). Kenneth's payoffs are in the lower part of each cell.

[Figure 5.1]

The tragedy of games of this form, as has long been recognized by scholars, is that for any decision that Alex makes, Kenneth is better off not investing. But if both players understand this logic, they end up with nothing, even though had they both invested, they would have ended up gaining.[2] To see this more clearly, imagine that Kenneth does not invest. Alex does better by not investing (net gain of 0) than investing (net loss of 25). If Kenneth does invest, Alex still does better by not investing (net gain of 75) rather than investing (net gain of 50). This exercise illustrates that both players have a dominant strategy (not to invest) in this game. Sadly, both shopkeepers decide to keep their money and life goes on without a road.

---

[1] Implicitly we assume that the private return on a dollar invested is –25%. The game we describe is strategically equivalent to any game in which the return is negative, but greater than –50%.

[2] Recall that the prediction of collective action failure is based on the assumption that both Alex and Kenneth are egoistic. If they are not, the game can look quite different. We rule out other-regarding preferences, however, as an explanation for the greater success of co-ethnics in resolving collective action problems in Chapter 4.

There are solutions to this problem.[3] If players replay this game many times with one another, then each can agree to cooperate in expectation that the other will too, but with the knowledge that should he fail to cooperate the other will not cooperate in the future (Axelrod 1984). Thus, although individuals cannot punish one another for failing to cooperate with them in a given instance, they can ensure compliance if they expect to interact with one another in the future. This sort of behavior can emerge more broadly within communities. Alex and Kenneth will prefer to invest in the road if each is afraid that his decision not to invest would be broadcast to friends and relatives—some of whom might respond with a variety sanctions, such as by not inviting them to the next party or reducing the amount of trade credit they are offered. Failing to cooperate can be costly—but only if you are seen to fail to cooperate.

Prisoners' dilemmas can also be resolved if norms of generalized reciprocity are present (Ostrom 1990; Rabin 1993). Players may elect to cooperate with their partner because they like to cooperate with people who cooperate with them. Notice that the threat of sanction does not figure in this story at all. Cooperation is possible because players may have expectations that others will cooperate. With those expectations *and* a preference for cooperating with those who are cooperative, egoistic players can overcome collective action problems without any sanctioning behavior whatsoever. If generalized reciprocity is present within a community, Alex may invest in the road because he has some prior belief that Kenneth will also invest, and Alex likes to repay another person's kindness with kindness.

A third possibility is that players achieve cooperative outcomes because of something that social scientists have termed quasi-magical thinking (Shafir and Tversky 1992). The idea here is that individuals choose to cooperate because they believe that choosing to cooperate will lead others to reciprocate, even if there is no norm that they should or sanction to enforce it. The "magical" part of this thinking is the belief that one's actions exert a causal effect on the actions of others, even if it is not possible that they can, as in a simultaneous prisoners' dilemma. So in the case of Alex and Kenneth, this logic is equivalent to saying: even though both Alex and Kenneth have been asked to contribute at the same time and each has no idea how the other will act, Alex will invest in the road because he believes that doing so will lead Kenneth to invest as well. This may seem a bit far-fetched

[3] Here, we outline a number of plausible and prominent explanations that are consistent with the greater success of co-ethnics in resolving prisoners' dilemmas. These are not, however, an exhaustive set of stories. Depending on how one understands the preferences of the players (egoistic vs. non-egoistic) and the structure of the game (simultaneous vs. sequential, one-shot vs. repeated), other stories can also be consistent with the evidence of co-ethnic success enumerated in the previous chapter.

at first, yet economists and psychologists have identified examples of magical thinking in everyday human behavior.

Chapter 4 already allowed us to rule out a set of preference and technology mechanisms that might account for the greater success of ethnically homogenous communities in resolving collective action problems, including the idea that altruism drives cooperative behavior. So among the strategy selection mechanisms enumerated above, which best describes the logic that underpins successful collective action among co-ethnics in Mulago-Kyebando? Is it the case that norms of reciprocity, sustained by the threat of sanction, exist among co-ethnics but do not extend across group boundaries?

The Prisoners' Dilemma: Results

To examine reciprocity in greater detail, we had subjects play a version of the standard prisoners' dilemma (PD) game (Davis and Holt 1993; Colman 1995; Ledyard 1995; Sally 1995).[4] PD games are models of situations like that faced by Alex and Kenneth, in which non-contributors ("free-riders") cannot be excluded from sharing in the collective benefits gained when others contribute. We also employed a version of the game that involves a third player (the "enforcer") who observes the actions of the first two players in the prisoner's dilemma and can choose to punish either one or both of them for their behavior (for a related game, see Fehr and Fischbacher 2004). We use these games first as a check on the consistency of our findings in Chapter 4. Is the co-ethnic bias observed in non-anonymous dictator games also apparent in public goods games? We then turn to set of strategy selection mechanisms that might account for our findings so far. While the data do not provide us with the power to distinguish among these mechanisms, we can explore games with an enforcer for evidence that players in our sample believe that community members sanction and respond accordingly.

Each round began when subjects were shown a public information box (PIB) containing pictures and/or videos of themselves and a single partner. The amount of information provided about players varied as in previous chapters, although for the PD game the information level was always symmetric (ie. no information for both, headshot for both, and so on). Subjects were then

[4] PD games were played with the same set of subjects that participated in the games described in Chapter 4. The sampling strategy used to select subjects, the randomized nature of pairings and a number of basic protocols regarding

given a 1000 USh note and asked to decide whether to invest the money or to keep it from themselves. In front of each player, we placed two boxes—one delineated with the label "self" and the other "group." To invest, a subject simply placed the 1000 USh note in a sealed envelope and put it into the "group" box.[5] If a player wished not to invest, she placed the 1000 USh note in a sealed envelope and put it in the box marked "self." Any contributions put into the "group" box by individuals within a given pairing were increased in value by 50%. This amount was then divided equally between the two players (note, this division of interest plus capital produces exactly the payoffs we described in our discussion of the PD game with egoistic players above). Each subject played approximately six rounds of this game. In all, we have data from 1,498 rounds (2,996 individual choices).

Half of the rounds involved a "third party enforcer" that had the opportunity to punish players after observing their actions. In these games, the enforcer was given the opportunity to observe whether each player cooperated or defected. Then, with two 500 USh coins provided as an endowment, the enforcer was invited to punish either player (or both) at a cost of 500 USh each. To do this, the enforcer placed the 500 USh coin in a sealed envelope into a box corresponding to the player to be punished. Any player that was punished lost her entire payoff in that game. Of course, if the enforcer chose not to punish at all, she was able to keep the 1000 USh endowment.

The informational treatment (varying how much one knows about his or her partner) is again central to the experimental design. Observing how individuals play with different levels of information and across distinct pairings tells us something about their beliefs about the likelihood that others will cooperate. Subjects may infer from the information provided that a player looks "trustworthy" or perhaps has a trusting voice, increasing the likelihood that they will cooperate. Subjects may also be cued to the ethnic identity or gender of a partner; if subjects believe, for example, that individuals from their ethnic group will behave in a particular way, we will observe this in higher rates of co-ethnic cooperation. Players in the PD game also knew whether or not a third party was observing their play (via an additional information box above on the screen). This enables us explore how the threat of sanctioning conditions play in the games.

---

instruction, checks for comprehension and the manner in which pairings are revealed on the screen (see Appendix C for more details on the exact protocols used for this game) are all identical.

[5] We initially labeled the "group" box as the "bank" box. After finding that subjects mistakenly believed the game was about their patterns of saving (and their trust of the banking system), we changed the name to "group." We made this change after only a day and a half of play (out of ten days). We find no difference in cooperation rates across these two periods.

For the purpose of assessing the consistency of our findings, the outcome of primary interest in the PD game is the frequency of cooperation across pairings and, in particular, the likelihood that two matched players from the same ethnic group will cooperate. As before, in addition to exploring the impact of information and the composition of partners on cooperation rates in the full sample, we examine the extent to which player types (egoist versus non-egoist) affect behavior.

In the PD game, the baseline cooperation rate in games with no information on both players was 52%. This finding is consistent with a number of recent studies that have played this game elsewhere in lab settings (Andreoni 1990; Camerer et al 2004; Sally 1995). In no information games, egoists were less likely to cooperate than non-egoists, but the difference is small and not statistically significant. When the identities of the two players are revealed, the frequency of cooperation increases by four percentage points (on average) to 56%. Egoists and non-egoists do, however, play differently at higher levels of information. Non-egoists cooperate more when they have information about their partners (and their partners have information about them) but egoists do not; under these informational conditions, non-egoists cooperate significantly more than the non-egoists.

But recall from Chapter 4 that egoists *do* respond to information in the two-way dictator game. Specifically, they give more to co-ethnics when they are observed in their actions. Does this pattern emerge when our subjects play the prisoners' dilemma as well? The results of our analysis are reported in Table 5.1.[6]

[Table 5.1]

Three important results add confidence to our earlier findings. First, egoist players generally cooperate at rates lower than non-egoist players. These differences are substantial—11 percentage points in sample of games with information on both players. Second, the coefficients reported in column 3 suggest that players are more likely to cooperate if they are playing with co-ethnics regardless of player type. From a statistical point of view, we can be more confident of this finding for co-region than co-ethnicity; the result is also stronger when we use our measure of subjective beliefs rather than the benchmark co-ethnicity. Most importantly, disaggregating the results for

[6] In principle, the behaviors of our subjects in the PD game were not directly observable by their partners (in contrast to the dictator game where subjects were told exactly how much they received from each offerer). Subjects received total

egoists and non-egoists, we find that the co-ethnic effect is driven *entirely* by self-interested players. Restricting attention only to Column 2, it is evident that the co-ethnic effect is large and (in most cases) statistically significant for egoists. No such effect can be observed among the non-egoists (reported in Column 1).

The final column of Table 5.1 asks a great deal of our data—we ask how much larger the co-ethnicity effect is for egoists as compared to non-egoists (what statisticians call a "cross-partial"). The estimated effect is always positive and above 10 percentage points. But for both measures of co-ethnicity and co-region, it just fails to attain significance. This provides some qualitative (but not robust) evidence that players cooperate more with people from the same region if *and only if* they are egoists.

The PD game is instructive. In contrast to those experiments discussed in Chapter 4, the PD game captures directly the challenge of collective action faced by residents of Mulago-Kyebando. Our findings confirm that co-ethnics resolve collective action problems with greater facility than non-co-ethnics. Consistent with the story we developed earlier, it appears that the success of homogenous pairings comes from their ability to induce egoistic players to act more like non-egoistic players, at least when they interact with co-ethnics. But the results do not yet tell us how ethnic groups discipline the behavior of non-contributors. Is there evidence of sanctioning at work in Mulago-Kyebando, or is co-ethnic cooperation driven an alternative strategy selection mechanism?

Although we cannot directly test the distinct interpretations of reciprocity against one another, we can look for evidence of sanctioning in the PD games played with an enforcer. Table 5.2 reports the impact of the enforcer on baseline rates of cooperation. Levels of cooperation increase from 56% to 64% when an enforcer is present in games with information on both players. This difference is statistically significant at the 5% level. Turning to how the threat of punishment affects play among different types of players, we find that the impact of the enforcer varies across types—egoists increase their rates of cooperation substantially when an enforcer is present, while non-egoists become only slightly more cooperative. The presence of an enforcer increases rates of cooperation among egoists to the same level observed among non-egoists when they play without an enforcer (and only a few points below the rates of the non-egoists with the enforcer). Strikingly, once punishment is permitted within the game, egoists no longer cooperate at significantly lower

---

payoffs from a set of six PD games at the conclusion of the experiment. Our findings suggest, however, that subjects

levels than non-egoists. This result is strongly suggestive: cooperation among our subjects is driven by expectations of sanctioning, especially among egoists.

[Table 5.2]

Turning to the behavior of enforcers, incurring a cost to punish and deter undesirable behavior is also a public good since effective deterrence helps all of those that interact with players facing the threat of punishment. We have already observed that our subjects change their behavior in anticipation of punishment, but do enforcers in fact punish? Moreover, is punishment a response to defection on the part of players?[7] The results presented in Table 5.3 are reassuring.

[Table 5.3]

Enforcers punish defectors and they do so at high rates. The fact that we observe *any* punishment—whether of defectors or not—may be surprising to some. Enforcers bear costs to punish even in the absence of direct benefits for doing so; moreover, enforcers cannot expect to be punished themselves should they fail to punish since their actions go unobserved. Consequently, punishment in this game represents non-rational behavior, at least for egoist players. Yet, players punish more than 25% of the time. In addition, enforcers respond systematically to how subjects play: they punish nearly one third of the time when the front end player defects and only one in six times when the player cooperates.[8] Consistent with the fact that punishment cannot be supported in this game by second order threats of out-of-game sanction, we observe no difference in punishment rates between enforcers that are seen and enforcers that are not seen. Finally, consistent with the logic of rational egoists, the egoists punish at lower rates than the non-egoists. They do sometimes punish cooperators, but at a rate 5 percentage points lower than non-egoists. And although they punish defectors more than they punish cooperators, they do so at rate 11 percentage points lower than non-egoists.

---

appeared to play *as if* their behavior was observable, even when no enforcer was present.

[7] In principle, there is no reason that punishment should only be observed in response to defection. In fact, there is increasing evidence of "spiteful" punishment in games in which punishment is allowed (see Fehr and Gachter 2000; Fehr and Gachter 2002)

[8] These observations of punishment in response to cooperation, although they could in principle be driven by spite or negative other-regarding preferences, could also signal failures of our subjects to understand the game or mistakes in execution of punishment.

The fact that we observe punishment, although difficult to reconcile with egoist incentives, confirms the rationality of how subjects behave in the front end of the PD game: players cooperate in the presence of an enforcer because there is a substantial risk that they will be punished if they do not. The risk of punishment renders cooperation rational even for egoistic players.

Evidence from our analysis of a public goods game thus provides further support for the existence of in-group reciprocity norms sustained by the threat of punishment in Mulago-Kyebando. Subjects in our sample anticipate punishment and respond by cooperating at higher rates. They also engage in sanctioning of uncooperative behavior, even at a cost to themselves. While other strategy selection mechanisms cannot be ruled out, the sanctioning account appears to rest on strong foundations.

Is Collective Action Supported by Ethnic or Universal Norms?

With added confidence that sanctioning sustains the reciprocity norms we observe in the non-anonymous dictator game and the prisoners' dilemma, we now turn to a closer investigation of how these norms work. Is it the case that individuals favor co-ethnics because of a norm of reciprocation within the group? Or do we observe co-ethnic reciprocation simply because in-group members are more easily found? Two distinct understandings of reciprocity are consistent with the findings. We now try to parse them.

To see the logic distinguishing these two stories, let us return to the case of Alex and Kenneth, who now encounter one another for the first time. Alex urgently needs supplies from Kenneth's store but does not have sufficient finances on hand. He asks Kenneth to trust him and provide him with supplies on the promise that he will deliver the money later in the afternoon. Will Kenneth give him the goods? Our evidence so far suggests that Kenneth is more likely to do so if the two are co-ethnics. This co-ethnic effect kicks in once there is sufficient information about their shared identity—for example, we would expect that Alex's request will be more likely to be granted if he comes to the store in person than if he makes the request over the telephone. The trust comes not from Kenneth having more positive feelings for Alex than he would for a non-co-ethnic, but because both know that if Alex fails to repay the loan, Kenneth can find and punish him. Kenneth, too, is sufficiently findable for Alex through social networks that Alex can punish Kenneth if he fails to help out a neighbor in need. The inferential problem is the following: Does the co-ethnic advantage depend on the greater findability of Kenneth and Alex relative to non-co-ethnics? Or,

given a similar level of findability for a non-co-ethnic, would we still observe a co-ethnic advantage with Kenneth fearing sanctioning particularly if he fails to help out a co-ethnic in need? To put it more concretely, would Kenneth feel particularly obligated to assist Alex (his co-ethnic) over an eight-foot tall, orange-haired stranger?

Schematically, the two mechanisms are illustrated in Figure 5.2. If individuals play according to a universal norm in which they cooperate with people that can be easily found, they will, empirically, cooperate more with co-ethnics since co-ethnics are more findable (Channels 1 and 2). If, however, the norm is a specifically ethnic norm (Channel 3), then we will also observe a co-ethnic advantage so long as co-ethnics are findable, but independent of whether or not they are *more* findable than non-co-ethnics.

[Figure 5.2]

To distinguish among these mechanisms, we derive some observable implications from each of the stories. In particular, we search for observable patterns that would be consistent with one story, but inconsistent with the other.

> **Observable Implication 1.** If the co-ethnic advantage works only through the mechanism of findability, then the ease with which someone can be found should predict behavior whether or not players are co-ethnics. If the co-ethnic advantage works though a specifically ethnic norm, then the extent to which players cooperate with or favor non-co-ethnics should not depend on their findability.

> **Observable Implication 2.** If the co-ethnic advantage works only through the strategy selection story—through specifically ethnic norms—then (egoist) players should persist in treating co-ethnics better *even after we remove the advantage co-ethnics have in finding one another.* Enforcers should also condition punishment behavior on the ethnic identities of players. However, if the findability story matters most, then the difference in behavior should vanish once the findability constraint is removed. Moreover, if sanctioning is as easy for co-ethnics as for non-co-ethnics and if reciprocity norms are universal then sanctioning should occur—and be expected to occur—with equal frequency whether or not a player is a co-ethnic.

To test the first observable implication, we generate a measure of the likelihood that an individual can be found using data drawn from the network game. We then examine the relationship between this measure and offers in the non-anonymous dictator game, looking to see whether players that are more easily found give higher offers regardless of whether their partner is a co-ethnic. To explore the second, we look again at the PD game with the third-party enforcer. In this game, we render the costs of punishing co-ethnics and non-co-ethnics equal and remove any advantage that derives from the greater findability of co-ethnics. We examine play to see whether players still condition their actions on co-ethnicity (even when punishment is possible within the context of the game), and whether enforcement is conditioned on ethnicity, as the ethnic norms story predicts.

"Findability" and Cooperation

In the Network Game, we found that group identity matters a great deal for how easily one can be located: individuals are 15 percentage points more likely to find a co-ethnic than a non-co-ethnic. But besides co-ethnicity, other characteristics may render someone more or less findable. Some people may be inherently easier to find no matter who attempts to find them. If we can identify the characteristics of more "findable" people, this provides a key to figuring out whether and how findability shapes cooperative behavior in general.

Individual markers such as height, age, and religious affiliation may make it easier to identify, describe, and locate someone. If so, individuals sharing these highly visible traits may behave differently than individuals that are more difficult to locate. Our exit interviews with players in the Network Game suggest that variation in the visibility of different markers was important as individuals sought to locate their partners in the Network Game. Ethnicity certainly played a role.

One respondent, describing how she selected an area in which to search, said:

> "I knew the name of that person, and I knew her tribe was Munyankole, coming from western Uganda. I went from household to household, asking where maybe her fellow tribesmates were staying. So wherever I find them, or a household speaking her language, I would ask."

Another reported:

> "Because there was a lady there who looked like the person, she was also dark, like she had a Western connection."

One subject emphasized the visibility of particular ethnic markers:

> "From my judgment, she must be a Lugbara, but she tried to make herself lighter… She was dark if you looked at the rest of her body, but her face is light."

In other cases, players used cues based on religion rather than ethnicity. One player reflected on the strategy he pursued as he left our office:

> "I thought about going to the LCs, or going to a Mosque because he was Muslim."

The name, he said, was critical for working out the religious affiliation of his partner. For some players, access to local authorities was more important than information specifically about the partner. One subject described how he first spoke to the LC1 chairperson in his zone:

> "I explained to him what I was doing. He looked at the name, but he did not know the person. He told me to look in a place. Then I talked to another person, who said to look in a particular area. Then I found the person's mother, and then the person right before he left the house."

Stories such as this one suggest that, while some visible markers make people more findable, individuals may also be easier to locate if they are connected to local political institutions. Players in the Network Game often reported that they turned to the LC1 system in an effort to locate their partners. This is perhaps not surprising, as LC1s are responsible for maintaining a list of all community members. In practice, however, these lists are rarely kept. It is possible that the use of LC1s may be an ethnic strategy, given the predominance of Baganda in these communities (and in the leadership of the LC1s). As it turns out, however, players of all ethnic groups were equally likely to visit LC1 chairs when trying to locate others: 61% of Baganda compared to 57% of non-Baganda.

This difference in the proportion using the LC1 system is of the right sign (indicating that Baganda are slightly better able to access LC1s), but it is not statistically significant.

Beyond linkages to the LC system, we might expect that age or the length of residence in a given area may make people easier to find. Gender could matter as well. Of course, not all of these characteristics are observable from simple photographs to the same extent. While it is usually straightforward to establish the gender or age of an individual using visual cues, it is more difficult to identify their religion, migration history, or connections to the LC system.

To identify more systematically the set of characteristics that contribute to an individual's findability, we employ data gathered in a short demographic survey of the subjects selected as targets in the Network Game. The data include information about age, gender, religion, length of residence, and political activity. The relationship between each of these variables and likelihood of being successfully located is given in Table 5.4.

[Table 5.4]

Column 1 presents the marginal effect of each of the variables on the likelihood that a target will be found, taken one at a time. Column 2 provides the estimated effects from a multivariate regression, showing the marginal effect when we control for each of the other determinants. Men are approximately 13 percentage points more likely to be found, although the result is not significant in a bivariate specification; the marginal effect of gender is even larger (19 percentage points) and strongly significant when we control for other factors. Religious affiliation also appears to matter for findability. Muslims are 23 percentage points more likely to be found than non-Muslims and this finding is significant at conventional levels. Older players[9] and those residents that have lived in an area for longer are also easier to locate.[10] Surprisingly, we find no relationship between political activity—as measured by participation in LC1 committees—and how easily one can be located.

To test our hypothesis about findability, we need to know whether people that can be more easily located also behave differently in the dictator game, offering more whether or not the receiver is a co-ethnic. We return to data gathered in the 100 USh denomination dictator game to examine

[9] The choice of the cut-off that we use for age derives from two considerations. First, an examination of the age distribution of targets suggests that "findable" targets (32) are at least three years older than "non-findable" targets (29). Second, the age distribution in the population is skewed. Slightly more than half of Uganda's population is under the age of 15. Moreover, the age distribution of our sample is centered at 26, with only a quarter of the subjects older than 35.
[10] One might imagine the opposite relationship in that recent migrants are often distinguishable by their manner of dress.

whether older people, more established residents, and Muslim offerers are also more generous, and in particular, whether they condition their generosity on being seen (in non-anonymous games) more than non-findable players.[11]

For each observable trait, Table 5.5 presents the average play of individuals in the anonymous game, the non-anonymous game, and the difference (which we interpret as how responsive individuals are to the loss of anonymity). The table also includes average offers for subjects split into two groups on the basis of an aggregate index of findability—an index that takes a value of 1 if the predicted findability of a player exceeds 0.5 and 0 otherwise. The cells report how much is *kept* by players (out of 1000 USh) depending on their characteristics. If the findability story is correct, those individuals that are easier to locate should keep significantly less, especially when they are seen by others (as in the non-anonymous games).

[Table 5.5]

The data provide little support for the argument that findability drives the behavior we observe in co-ethnic pairings. Players that can be more easily located keep 10 USh less (this is driven particularly by older players), but they do *not* condition their play on whether or not they are seen any more than non-findable people. Further, the insensitivity of findable people to information is true both for egoistic and non-egoistic players (results not shown). Information on the offerer reduces how much players keep only for Muslim players; in all other cases, more findable offerers tend not to respond to the loss of anonymity as the findability story would predict.

A closer look at the data from the Network Game thus suggests that ethnic and religious cues provide a useful starting point as players seek to locate their partners. Players are indeed better able to find co-ethnics and to sanction them for bad behavior; knowing this, co-ethnic players that know that they can be seen may anticipate this and offer more (or discriminate in favor of co-ethnics). But the mechanism that explains this behavior is not simply the fact that co-ethnics are easier to find. Other people that can also be found do not adjust their play in the dictator game because they fear out-of-game sanctioning. Understanding how reciprocity works in Mulago-Kyebando requires that we look somewhere else for the underlying mechanism.

[11] We do not examine the discrimination game since our measure of findability is determined by the characteristics of

The Existence of Ethnic Norms of Reciprocity

Perhaps the co-ethnic advantage works instead through specifically ethnic norms. If this is the case, co-ethnics should achieve higher levels of cooperation even after we remove the advantage that co-ethnics have in finding one another (and thus, their greater ability to threaten and implement out-of-game sanctions). Moreover, players should be more likely to punish in games in which a player or players share the same ethnic or regional identity as the enforcer. To explore the evidence for ethnic (as opposed to universal) norms, we return to our analysis of the prisoners' dilemma games that included a third party—a "norm" enforcer.

Evidence from play in the prisoners' dilemma (reported in Table 5.6) is more consistent with the findability than the strategy story. The presence of an enforcer eliminates the co-ethnic edge in cooperation rates among egoists. Yet a comparison of the results in Tables 5.1 and 5.6 reveals a striking finding: the magnitude of the enforcement effect is almost identical to that of co-ethnicity in games without an enforcer. The enforcer effect appears to *substitute* for co-ethnicity. We take this as evidence that *the effects of within-group sanctioning are of the same order of magnitude as the effects of observation of action and punishment*. Universal reciprocity norms appear to be in operation among our subjects and these swamp any specific co-ethnic norms of cooperation.

[Table 5.6]

Patterns of punishment behavior, however, suggest that both universal and ethnic norms are in operation in Mulago-Kyebando. Table 5.7 summarizes the findings. In deciding whether to punish, subjects respond to a strong universal norm of sanctioning when players defect *whether or not they are co-ethnic* (as reported in row 3). This universal norm is powerful. There is also some evidence for a specifically ethnic norm: enforcers are likely to punish whenever they observe defection but they are especially likely to do so when playing with co-ethnics. There is thus evidence for both kinds of norms, although the results in the front end suggest that the universal norm appears considerably more powerful.

[Table 5.7]

---

one player only, and thus cannot help in determining *who* a player should favor in the 500 USh game.

Though smaller in magnitude, the specifically ethnic norm may hold the key to understanding the co-ethnic advantage. Table 5.8 presents average rates of punishment conditional on a given player's defection, with ethnic composition defined using both benchmark co-ethnicity and benchmark co-region.[12] Here, we examine play among both types of enforcers—egoist versus non-egoist—and present average punishment behavior across of the full range of possibilities for the ethnic composition of the trio playing the game. These include: games where all three players come from different ethnic groups (heterogeneous), games in which two of the three players share the same ethnicity (games in with the enforcer shares the same ethnicity one of the other players and where the two players share the same ethnicity), and games where all three players come from the same ethnic group (homogeneous).

In a game with so many different combinations of players, there are a range of possible patterns we might observe. One possibility is that players simply punish in-group members, irrespective of who they were paired with when they defected. Another is that enforcers punish in-group members especially if they defect on other in-group members—this behavior would involve acting to uphold group norms. A third, emphasized by Fearon and Laitin (1996), would find *co-ethnics* punishing *especially* when a player defected on another co-ethnic. A final possibility is that players uphold co-ethnic norms in general, punishing players for defecting on their own co-ethnics, whether or not they are co-ethnics of the enforcer. While our data lack sufficient power to distinguish among these explanations, they do provide some initial evidence about which of these predictions has face validity.

[Table 5.8]

Consistent with our previous results, non-egoists are responsible for much of the punishment we observe; egoists are simply less willing to bear the costs to enforce norms. Moreover, it appears that non-egoists condition their punishment behavior on the identity of the players in the trio.

Further, we find evidence for three propositions. First, non-egoist enforcers punish their own co-ethnics more—by 13 percentage points—than they do non-co-ethnics (row 5). Second, enforcers display some tendency to punish non-co-ethnics when those non-co-ethnics defect on their own co-ethnics, although the result falls just below statistical significance (row 6). Third, the

[12] We do not use imputed measures of co-ethnicity or co-region here because we lack information about the enforcer's beliefs about the likelihood that the two players in the PD game are co-ethnics.

most powerful result obtains when all three members of a trio are co-ethnics: here, non-egoist enforcers punish defectors specifically when they defect on co-ethnics, at a rate higher by 22 percentage points (row 7). While universal norms may predominate, these patterns provide powerful evidence that specifically ethnic norms also condition behavior in Mulago-Kyebando.

Conclusion

We began this chapter with an effort to add confidence to our previous findings. Specifically, we sought to demonstrate that the co-ethnic advantage really exists in public goods games. We also looked for evidence that co-ethnics cooperate more successfully because players fear sanctions should they defect. The results of our additional analyses are strongly supportive of the emerging story. Co-ethnic pairings do succeed in solving collective action problems more readily than do non-co-ethnic pairings. As we found in Chapter 4, the co-ethnic advantage is driven by increased levels of cooperation among egoist players in particular—co-ethnicity induces egoists to cooperate at levels typical of non-egoists. It prevents the tragedy of collective action failure that arises in the presence of self-interested players. Moreover, we show that threats of sanctioning are real. Players correctly anticipate that they will be sanctioned if they fail to cooperate and respond by cooperating at higher rates.

Our ambition, however, was greater than simply corroborating the findings of Chapter 4. We wanted not only to probe the plausibility of our story about reciprocity norms enabling homogenous pairings to cooperate successfully. We sought also to push our analysis further by distinguishing between two rival (though related) accounts of the underlying mechanism that might explain the co-ethnic advantage. Is it the case that reciprocity norms are universal but applied unevenly, because co-ethnics are more easily found than non-co-ethnics? Or is it the case that reciprocity norms act only on individuals when playing with co-ethnics? In this exploratory effort, we met with more mixed results. We failed to find evidence that other subjects that can be easily found also cooperate at higher levels. This suggests that the operation of specifically ethnic norms of cooperation are more important than the findability of co-ethnics, even if these norms are only (or partially) rendered possible by the fact that co-ethnics are indeed more findable. Our examination of the behavior of enforcers confirms this result—enforcers do appear to follow specifically ethnic norms. They punish co-ethnics that defect at higher rates than non-co-ethnics, especially when co-ethnics defect in pairings with other co-ethnics.

However, we find also that this co-ethnic norm exists besides another still more powerful, universal norm. Although enforcers punish co-ethnic defectors at higher rates, they tend to punish a large proportion of subjects that defect (regardless of their ethnic group membership)—far more than would be predicted if we adopted standard, self-interested models of human behavior. This universal norm is important for explaining enforcer behavior and also for understanding levels of cooperation in the prisoners' dilemma game. Although players do cooperate at higher rates with co-ethnics when no enforcer is present, the introduction of an enforcer activates a universal norm of cooperation that outweighs specifically ethnic considerations. Co-ethnic norms are useful for policing the self-interested behavior of egoists; in the absence of a norm enforcer, they are sufficient to raise cooperation levels to those exhibited by non-egoists. But this co-ethnic advantage is erased when universal norms of cooperation are made salient by the presence of an enforcer. The fact that heterogeneous groups do no worse than homogenous groups in resolving collective action problems when an enforcer is present suggests that reciprocity norms can be cultivated across group lines, particularly, in the context of an institutional environment that renders actions observable and punishment possible. How this might matter for addressing the challenges that diversity poses for public goods provision is an issue we turn to in the conclusion.

Figure 5.1: A Prisoners' Dilemma

		Alex: Not Invest	Alex: Invest
Kenneth	Not Invest	*0* *0*	*-25* *75*
	Invest	*75* *-25*	*50* *50*

Figure 5.2: Two Understandings of Reciprocity

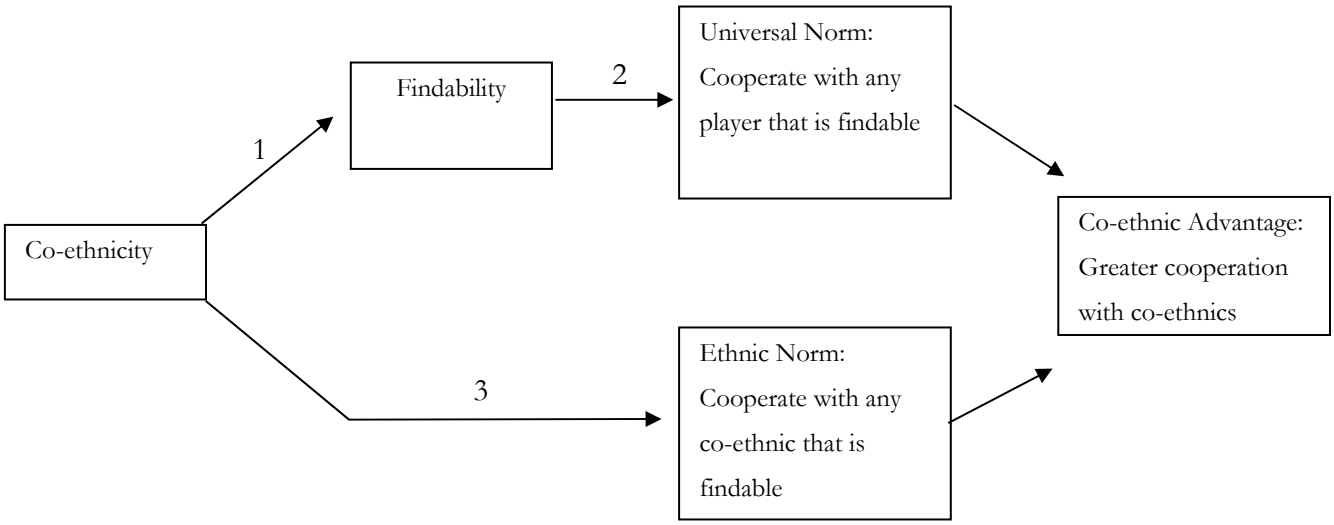

Table 5.1: Co-ethnicity and Cooperation

	(1) Non-Egoist (*se*)	(2) Egoist (*se*)	(3) Total (*se*)	(4) Difference (*se*)
Likelihood of Cooperation	0.61	0.50	0.56	0.11***
	n=584	*n*=382	*n*=966	(0.03)[a]
Marginal Effect of Benchmark	-0.01	0.10	0.04	0.12
Co-ethnicity	(0.05)[b]	(0.07)[b]	(0.04)[b]	(0.08)[c]
Marginal Effect of Subjective	0.10	0.29**	0.18**	0.18
Co-ethnicity	(0.09)[b]	(0.13)[b]	(0.07)[b]	(0.13)[c]
Marginal Effect of Benchmark	0.01	0.14**	0.06	0.11
Co-region	(0.04) [a]	(0.06)[a]	(0.04)[a]	(0.07)[c]
Marginal Effect of Subjective	0.04	0.28***	0.14**	0.13
Co-region	(0.08) [b]	(0.11) [b]	(0.07) [b]	(0.13)[c]

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. Standard errors reported in parentheses.
[a] Differences report result of *t*-test on a binary variable.
[b] Result of a probit (dprobit) model. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Fixed effects for the player's ethnic group are included, and disturbance terms are clustered for each player across all of his or her games.
[c] Result of probit model (dprobit) with the identity measure interacted with the egoist measure. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Fixed effects for the player's ethnic group are included, and disturbance terms are clustered for each player across all of his or her games.

Table 5.2: The Impact of an Enforcer on Cooperation

	(1) Non-Egoistic (*se*)	(2) Egoistic (*se*)	(3) Total (*se*)	(4) Difference (*se*)
Likelihood of Cooperation With	0.61	0.50	0.56	0.11***
No Enforcer	*n*=584	*n*=382	*n*=966	(.03)[a]
Likelihood of Cooperation With	0.64	0.60	0.62	0.04
an Enforcer	*n*=791	*n*=497	*n*=1288	(0.03)[a]
Difference	0.03	0.10***	0.06***	0.07
	(0.03)[a]	(0.03)[a]	(0.02)[a]	(0.05)

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. Standard errors in parentheses.

[a] Differences report result of *t*-test on binary variables.

[b] Result of a probit (dprobit) model. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Disturbance terms are clustered for each player across all of his or her games.

Table 5.3: Punishment Behavior by Information Level and Player Action

	Front End Player Cooperates			Front End Player Defects		
	Non-Egoists	Egoists	Difference	Non-Egoists	Egoists	Difference
Punisher Not Seen	0.12 (75)	0.16 (62)	0.04 (.06)[a]	0.37 (46)	0.26 (34)	-0.10 (.11)[a]
Punisher Seen (*But not seen to punish*)	0.20 (527)	0.14 (360)	-0.05** (.03)[a]	0.37 (337)	0.26 (203)	-0.11*** (.04)[a]
Difference (*se*)	0.08 (0.05)[a]	-0.02 (0.05)[a]	-0.10 (0.07)[b]	0.00 (0.08)[a]	0.00 (0.08)[a]	-0.01 (0.12)[b]

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. Differences report result of *t*-test on binary variables. Standard errors in parentheses.

[a] Differences report result of *t*-test on binary variables.

[b] Result of a probit (dprobit) model. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Disturbance terms are clustered for each player across all of his or her games.

Table 5.4 Correlates of Findability

	Marginal effect on success rates (bivariate relationship)	Marginal effect on success rates (multivariate relationship)
Male	0.13	0.19
	(.08)	(.08)**
Catholic	0.04	0.13
	(0.08)	(0.09)
Muslim	0.23	0.19
	(.10)**	(.11)*
Old (subject is older than 35)	0.32	0.31
	(.08)***	(.09)***
Recent arrival	-0.15	-0.11
	(.08)*	(.08)
LC Committee	0.003	-0.05
	(.10)	(.10)
Observations	148	127
R-squared		0.18

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. OLS regression results, with standard errors in parentheses.

Table 5.5: Average Response to Loss of Anonymity by Findability Traits

	Non-Muslim	Muslim	Young	Old	Recent Arrivals	Established	Findable	Not Findable
Amount Kept in Anonymous Game	535 (539)	573 (73)	552 (470)	496 (140)	540 (259)	538 (353)	533 (294)	545 (315)
Amount Kept in Non-Anonymous Games	551 (423)	533 (66)	555 (385)	523 (104)	547 (207)	550 (280)	544 (215)	551 (271)
Difference (*se*)	16 (14)[a]	-39 (29)[a]	25 (14)[a]	27 (27)[a]	6 (20)[a]	12 (16)[a]	12 (19)[a]	6 (17)[a]

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. Number of observations in parentheses.
[a] Differences and standard errors from a t-test on a binary variable.

Table 5.6: Co-ethnicity and Cooperation with Enforcer

	(1) Non-Egoistic (*se*)	(2) Egoistic (*se*)	(3) Total (*se*)	(4) Difference (*se*)
Marginal Effect of Benchmark	0.02	-0.11*	-0.02	-0.09
Co-ethnicity with an Enforcer	(0.05)[a]	(0.06) [a]	(0.04) [a]	(0.08)[b]
Marginal Effect of Subjective	0.03	-0.07	-0.01	-0.02
Co-ethnicity with an Enforcer	(0.08)[a]	(0.10)[a]	(0.06)[a]	(0.11)[b]
Marginal Effect of Benchmark	0.00	-0.02	-0.00	-0.03
Co-region with an Enforcer	(0.04)[a]	(0.05)[a]	(0.03)[a]	(0.06)[b]
Marginal Effect of Imputed Co-region with an Enforcer	-0.01	-0.07	-0.02	-0.06
	(0.06)[a]	(0.09)[a]	(0.05)[a]	(0.10)[b]

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. Robust standard errors in parentheses.
[a] Result of a probit (dprobit) model. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Fixed effects for the player's ethnic group are included, and disturbance terms are clustered for each player across all of his or her games.
[b] Result of probit model (dprobit) with the identity measure interacted with the egoist measure. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Fixed effects for the player's ethnic group are included, and disturbance terms are clustered for each player across all of his or her games.

Table 5.7: Patterns of Punishment by Player Action and Co-ethnicity

	Non Co-ethnic	Co-ethnic	Difference (*p*)	Non Co-region	Co-region	Difference (*p*)
Player 2 Cooperates	0.17	0.18	0.01	0.17	0.17	0.00
	(718)	(195)	(.03)[a]	(597)	(316)	(0.03)[a]
Player 2 Defects	0.30	0.39	0.09*	0.31	0.34	0.03
	(437)	(119)	(.05)[a]	(368)	(188)	(0.04)[a]
Difference (*se*)	0.13***	0.21***		0.14***	0.17***	
	(0.03)[a]	(0.05)[a]		(0.03)[a]	(0.04)[a]	

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. Number of observations in parentheses.
[a] Differences and standard errors from a t-test on binary variables.

Table 5.8: Effects of Ethnic Composition on Punishment Behavior

	Frequency with which player is punished conditional upon defection (all players seen)							
	Ethnicity				Region			
	Non-egoist	Egoist	Total	Diff (*se*)	Non-egoist	Egoist	Total	Diff (*se*)
Trio heterogeneous	0.35 (144)	0.27 (100)	0.32 (244)	-0.08 (0.06) [a]	0.26 (69)	0.31 (55)	0.28 (124)	0.05 (0.08) [a]
Enforcer and Player 2 share ethnicity	0.40 (35)	0.33 (18)	0.38 (53)	-0.07 (0.14) [a]	0.39 (56)	0.34 (29)	0.38 (85)	-0.05 (0.11) [a]
Observer and Player 3 share ethnicity	0.28 (36)	0.24 (25)	0.26 (61)	-0.04 (0.12) [a]	0.33 (57)	0.23 (39)	0.29 (96)	-0.10 (0.10) [a]
Players 2 and 3 share ethnicity	0.42 (38)	0.23 (22)	0.35 (60)	-0.19 (0.13) [a]	0.50 (60)	0.23 (26)	0.42 (86)	-0.27** (0.11) [a]
Trio homogenous	0.58 (33)	0.17 (12)	0.47 (45)	-0.41** (0.16) [a]	0.47 (45)	0.14 (28)	0.34 (73)	-0.32*** (0.11) [a]
Difference Between Games When Player 2 is and is not a Co-ethnic	0.13* (0.08) [b]	0.04 (0.11) [b]	0.10 (0.07) [b]		0.04 (0.07) [b]	0.00 (0.07) [b]	0.02 (0.05) [b]	
Difference Between Trios with Co-ethnic Front End Players and all other demographies	0.12 (0.08) [b]	-0.05 (0.10) [b]	0.08 (0.06) [b]		0.15** (0.01) [b]	-0.10 (0.08) [b]	0.06 (0.05) [b]	
Difference Between Homogenous Trios and all other demographies	0.22* (0.13) [b]	-0.10 (0.14) [b]	0.14 (0.10) [b]		0.08 (0.11) [b]	-0.12 (0.08) [b]	-0.01 (0.07) [b]	

Notes: *** Significant at 1%; ** Significant at 5%; * Significant at 10%. The averages shown are calculated using the sample of games in which there was positive information on all three players. Number of observations in parentheses.
[a] Differences report result of *t*-test on binary variables. Standard errors in parentheses.
[b] Result of a probit (dprobit) model. Marginal coefficient estimates are reported (at mean values for the explanatory variables). Disturbance terms are clustered for each player across all of his or her games. Standard errors in parentheses.