

Chapter 1

Diversity and Collective Action

JUST OFF the main road that cuts through the slum area behind Kampala's main hospital lies the local council zone (LC1) of East Nsooba.¹ An ethnically mixed neighborhood, it is set at the base of a steep hillside with small, closely spaced houses stretching down the incline and across the swampy valley floor. The houses are simple, with concrete floors and walls and zinc roofs. Few have windows or electricity. None has indoor plumbing. But the biggest hardship for the area's residents comes not from the plainness of their houses or the absence of amenities like electric lights and indoor toilets, but from rainy season flooding, which transforms the neighborhood's unpaved streets into rivers of mud, submerges houses in filth, and leaves putrid standing water that breeds cholera, malaria, and other diseases. Drainage channels designed to carry away the excess rainwater snake throughout the area, but the government has not maintained them for years, and they are too choked with garbage and debris to be of any use. So when the rains come, the floodwaters rise.

To the north of East Nsooba lies Kisalosalalo, an area less prone to flooding but more troubled by crime. Break-ins occur almost every night, and some thefts are accompanied by violence. People feel insecure. The Ugandan government used to sponsor and equip local defense units (LDUs) to patrol neighborhoods like Kisalosalalo. But government support for the LDUs ceased in 2002 when they were incorporated into the formal police system; since then, Kampala's slum areas have been overtaken by crime and violence.

Fredrick Ssalongo has been the chairman of East Nsooba's local council since 1987, when Uganda's system of decentralized administration was initiated. Geoffrey Kashaija has served in the same capacity in Kisalosalalo since 1991.² Both men describe their deep frustration at the unwillingness of the central and local governments to provide basic ser-

vices in their areas. They also describe their disappointment at the inability of members of their communities to work together to help solve some of their areas' problems. A few years back, for example, the Kampala City Council had arranged for a private company to pick up trash in poor neighborhoods, but the trucks stopped coming to East Nsooba because only a handful of the people there were willing to pay the fees. Even without regular garbage collection, the area's flooding problems might be reduced if residents heeded Chairman Ssalongo's calls to clean the channels of debris periodically and to refrain from throwing their trash in the drains. But these efforts, too, have failed. In Kisalosalalo, Chairman Kasaijja's attempts to set up community patrols to deal with the area's crime problem have been similarly unsuccessful. Although everyone in the neighborhood agrees that crime is a critical issue, only a relative handful of residents have been willing to participate in the nightly patrols, and most have been reluctant to contribute funds to remunerate those who have.

The hardships faced by the residents of East Nsooba and Kisalosalalo are like those confronted by many people living in impoverished urban areas and villages around the world. Largely abandoned by their governments, these local communities are left to fend for themselves to provide basic public services like sanitation, flood protection, and security. Often they fail. This is because the services in question are what economists call "public goods" that can be consumed and enjoyed by everyone in the community irrespective of whether they contribute to their provision. As Chairmen Ssalongo and Kashaija have learned, this feature of public goods creates strong incentives for people to attempt to "free-ride" on the contributions of others.³ Residents of East Nsooba calculate that if they lie low while others do the dirty work of extracting the foul-smelling trash from the drainage channels or pay the fees for the private haulage company, they can reap the benefits of less flooding while avoiding unpleasant work and keeping their money. Residents of Kisalosalalo similarly figure that they can enjoy the benefits of reduced crime without contributing to providing it. The problem, of course, is that if everyone makes the same calculation, then the garbage collection, drainage channel maintenance, and public safety will never be provided. The powerful individual incentives to shirk undermine the collective action necessary for providing these goods and thus reduce everybody's well-being.⁴

Given the enormous social costs of such collective action failures, why is it so difficult for communities to find ways to work together to provide critical public services? It turns out that while most cannot find ways to engage in effective collective action, some can. How do we explain this variation? In situations where the government has abdicated responsibility for providing public goods, why are some communities

able to work together to generate public goods themselves while others are not?

Wealth is obviously a factor: prosperous communities can afford to hire security patrols, build clinics, pay teachers, and install water and sewerage systems. But wealth provides, at best, only a partial answer, for we find high levels of public goods provision in some very poor communities and low levels of public goods provision in some very wealthy ones. Public goods provision depends at least as much on the ability of a community to work collectively to solve its problems as it does on whether it can buy better public services by simply writing a check.

If a community's cooperative capacity is central to the generation of public goods, what accounts for this cooperative capacity? A growing literature suggests that part of the answer may lie in the community's ethnic heterogeneity. Studies conducted in communities around the world—in sites as disparate as western Kenya, northern Pakistan, central Mexico, Indonesia, and the United States—have found evidence that ethnic diversity sometimes impedes the provision of public goods by communities. The central goal of this book is to work out why this is so. Why do homogeneous communities—those where interactions are more likely to be among coethnics—seem to have an advantage in achieving collective ends? Before we answer this question, we review the evidence linking diversity and the failure of collective action and explore the arguments advanced by scholars around the world to account for this relationship.

Ethnic Diversity and Public Goods Provision

Evidence from Around the World

Roads, irrigation channels, and other forms of public infrastructure degrade quickly in the harsh weather conditions of northern Pakistan and cannot keep functioning without frequent maintenance. Some communities succeed in maintaining their infrastructure while others fail. To find out why, Asim Khwaja (2008) studied 132 community-maintained public infrastructure projects in 99 communities. He found that project upkeep—as measured by the degree of physical degradation of project infrastructure, the extent to which the project's original purpose is still being satisfied, and the amount of maintenance work carried out over the lifetime of the project—was closely related to the social heterogeneity of the beneficiary community. Controlling for other project- and community-specific factors, such as land inequality, Khwaja found that communities in the first quartile of social heterogeneity scored ten per-

centage points higher on the project maintenance scale than did communities in the third quartile. Given the importance of roads and irrigation in this remote, arid region, the negative impact of ethnic diversity on the maintenance of public infrastructure has severe consequences for people's well-being.

Working in a very different environment, Edward Miguel and Mary Kay Gugerty (2005) studied the determinants of school funding and quality in eighty-four primary schools in western Kenya. As in many parts of Africa, the Kenyan government pays teachers' salaries, but materials such as books, chalk, pencils, paper, desks, and even the physical school buildings themselves are funded through contributions from the local community. School quality is thus closely related to the ability of local communities to act collectively to raise funds for these purposes. Controlling for a wide range of socioeconomic, geographic, and demographic factors (including the possibility that, by generating in-migration, high school quality might increase community-level heterogeneity), Miguel and Gugerty found that communities at average levels of ethnic diversity generated 20 percent lower school contributions per pupil than homogeneous communities. As in the Khwaja study, this is a substantial effect.

The negative impact of ethnic heterogeneity on collective action has also been demonstrated in dozens of other studies conducted in cities and villages around the world. Cagla Okten and Una Okonkwo Osili (2004) demonstrated that community-level diversity in Indonesia was negatively related to contributions of labor, materials, and money to neighborhood irrigation associations, security arrangements, rice cooperatives, and local health centers.⁵ Pranab Bardhan (2000) reported a negative association between caste heterogeneity and the successful maintenance of local irrigation schemes in India; Abhijit Banerjee, Lakshmi Iyer, and Rohini Somanathan (2005) and Banerjee and Somanathan (2004) reported a similar relationship with respect to the provision of schools, public transport, electricity, health centers, and water projects. Jeffrey Dayton-Johnson (2000) provided evidence for the negative impact of social diversity on the success of community irrigation schemes in Mexico. Jean-Marie Baland and his colleagues (2006) extended this work to the management of forestry resources in Nepal, where they, again, found a negative association between diversity and community-level collective action.

These findings are not limited to developing countries. Alberto Alesina, Reza Baqir, and William Easterly (1999) presented evidence that diversity was associated with lower levels of funding for schools, roads, sewers, and trash pickup in U.S. municipalities. Their findings corroborate those of James Poterba (1997) and Claudia Goldin and Lawrence Katz (1999), who reported negative associations between school funding and ethnic diversity in U.S. cities and school districts, and, recently,

those of Robert Putnam (2007), who found social diversity to be negatively associated with trust, social capital, and a variety of measures of public goods provision. Jacob Vigdor (2004) found that racially diverse communities had lower response rates to the 2000 U.S. census; because the allocation of federal funds to communities is related to their response rates, he interpreted this as a failure of collective action. Julio Videras and Christopher Bordon (2006) found that more ethnically heterogeneous zip code areas in New York and New Jersey had lower levels of environmental protection. Xin Sherry Li (2005) extended this work to Europe, where, using data from the European and World Values Surveys, she found that ethnic heterogeneity was negatively related to tax compliance.

The relationships identified in these studies are all statistical relationships and thus provide evidence of correlation rather than of a deterministic causal relation. In specific cases, diversity may not be associated with poor public goods provision; indeed, there is evidence in some contexts that diversity can have powerful positive effects (Page 2007). It might also be the case that poor public goods provision exerts a causal effect on the level of diversity (perhaps by driving down property values and creating opportunities for in-migration), not vice versa. Nevertheless, the consistency of the relationship documented across studies is impressive and provides support for the proposition that, on average, diversity impedes the provision of public goods. Summarizing the literature on the topic, Banerjee, Iyer, and Somanathan (2005, 639) went so far as to refer to the negative relationship between diversity and public goods provision as “one of the most powerful hypotheses in political economy.”

While the literature to date has done a reasonable job of documenting the existence of this empirical regularity, it has been much less successful in accounting for it. Although most analyses venture hypotheses about the mechanisms that might be at work, no study has offered a systematic enumeration of the causal channels through which ethnic diversity might undermine public goods provision, along with a test of the comparative explanatory power of each.⁶ The state of the literature might be summarized by saying that a consensus has emerged about the nature of the association between ethnic heterogeneity and public goods provision, but that the micro-logic of this connection is still poorly understood. The literature provides us with a number of intuitions about the mechanisms that might be at work, but little research that might permit us to adjudicate among them. The goal of this book is to take this next step.

Possible Explanations

In taking on the question of why ethnic diversity impedes collective action (or, on the flip side, why ethnic homogeneity is associated with

more successful cooperative outcomes), this book departs from the typical social science concern with *whether*, rather than *why*, one thing causes another. The norm in the literature has been to focus on establishing that x causes y . Answering the question of why this relationship exists is usually secondary. Some kind of explanation is typically ventured, but such accounts tend to be post hoc attempts to explain a given result (occasionally with suggestive evidence for the plausibility of the preferred explanation) rather than the result of systematic testing of alternative mechanisms.

One reason why relatively little attention is given to causal mechanisms is that simply establishing the existence of a causal relationship is usually achievement enough. The literature we focus on here offers a good example. It is filled with studies that provide evidence for the link between ethnic heterogeneity and collective action failure, but as we noted earlier, only a relative handful employ sufficiently strong identification strategies to give us real confidence in their *causal* claims. Simply establishing that diversity impedes collective action and public goods contribution is thus a real contribution, quite apart from whether the researcher goes on to provide a comprehensive account for why.

A second reason for the relative lack of attention to causal mechanisms is the difficulty of identifying the universe of channels through which a given relationship might operate—a necessary first step in ascertaining which mechanism is actually doing the work. This is particularly so in the study of ethnicity, since there is no generally accepted framework that can be taken “off the shelf” to use in thinking analytically about how ethnic identities shape individual behavior and thus how a community’s ethnic diversity affects its ability to achieve collective ends. To provide such a framework, we turn to game theory, which offers a language and a theoretical apparatus that are particularly well suited to this task.

In the language of game theory, we can describe a social interaction as a game comprising three objects: a population (the set of actors), a technology (the set of strategies available to each of these actors), and preferences (which describe how individuals value the outcomes that result once all the actors select their strategies).⁷ The overall outcome of the game can be written as a function of the strategies selected by all the players, typically predicted by theorists using a solution concept.

This simple game description provides us with a set of first-order channels to investigate in order to identify how people’s ethnic identities affect their behavior and the outcome of their interactions with others. Ethnicity could affect a person’s behavior by changing the population playing the game, the strategies available to the players, or the players’ preferences. In addition, ethnicity could affect the strategies the players decide to use, conditional on these factors. Assuming that eth-

Table 1.1 Mechanisms Linking Ethnicity to Collective Action Success

Preferences Mechanisms	
Other-Regarding Preferences	Coethnics may be more likely to take each other's welfare into account
Preferences in Common	Coethnics may be more likely to care about the same outcomes
Preferences over Process	Coethnics may prefer the process of working together
Technology Mechanisms	
Efficacy	Coethnics may be able to function together more efficiently
Readability	Coethnics may be better able to (or believe they are better able to) gauge each other's characteristics
Periodicity	Coethnics may engage each other with greater frequency
Reachability	Coethnics may be more able to track each other down
Strategy Selection Mechanisms	
Reciprocity	Coethnics may be more likely to punish each other for failing to cooperate

Source: Authors' compilation.

nicity is predicated on existence (that is, we do not treat the set of actors as a function of ethnicity), we are left with three families of mechanisms through which the outcome of a social interaction might be affected by the actors' ethnic identities: preference mechanisms, technology mechanisms, and strategy selection mechanisms. These distinct families of mechanisms can be thought of more generally as different ways in which sharing (or not sharing) an ethnic identity with others may shape the kinds of choices people make in social settings.

Each of these families of mechanisms subsumes multiple, distinct explanations. Because there is no clear way to conceptualize the universe of all possible mechanisms within each family, we examine the major arguments that have been advanced by theorists of ethnic identity, placing them within the appropriate family of mechanisms. Table 1.1 summarizes the collection of leading explanations we have identified and which we will examine in this book. We describe each in turn.

Preferences Mechanisms To illustrate the logic of the *preferences* mechanisms, consider again the situation that Chairman Ssalongo confronts in

organizing community members to maintain the drainage channels in East Nsooba. Recall that one reason why residents might shirk from this task is that each one hopes that others will do all the work, and each one fears that, if he takes part, he will be a “sucker”—the one who provides all the work but then cannot prevent others from sharing in the fruits of his labor. In some cases, this problem corresponds exactly to what game theorists call a Prisoner’s Dilemma. In such cases, no matter what other people do, each individual would rather stay at home (“defect”) than contribute (“cooperate”). As a result, the drainage channel remains filled with trash (and next to useless in the rainy season) even though everyone would rather see it cleared.

This unhappy outcome follows in part from an assumption that is commonly made by game theorists: that people are selfish and do not benefit from improvements in the welfare of others. But what if this assumption is wrong? What if residents in Ssalongo’s zone care not only about the well-being of their own families but also about the quality of life of their neighbors? This would be a situation in which people exhibit *other-regarding preferences*. If this were the case, then residents might be willing to volunteer their labor for drainage channel maintenance regardless of what others do. Of course, the converse is equally true: if people actively dislike their neighbors and gain satisfaction from seeing them in misery, then residents are even less likely to be willing to bear the cost of maintaining the area’s drainage channels. As the example makes clear, having positive (or negative) preferences for the welfare of others can help solve (or exacerbate) the collective action problem. To the extent that people specifically care more about the welfare of individuals in their *own* ethnic group—as some theorists of ethnicity assume⁸—this mechanism would lead us to predict a greater likelihood of collective action success when the members of a community share a common ethnic background.

To illustrate the second preferences mechanism, which we term *preferences in common*, suppose that the local government in Kampala’s urban slums is in fact responsive to community members’ requests for assistance. If residents in Chairman Kasajja’s zone can lobby together for funding to support local defense units in the area, then the likelihood of a positive response from the local government (and greater security) increases. But if community members are unable to coordinate their lobbying—if, for example, some residents want better policing but others prefer that the government allocate its resources to improve local schools or transport infrastructure—then the likelihood that the government will increase its support for local defense units diminishes. Note that here the positive (or negative) outcome for the community derives not from the degree of concern that residents have over each others’ welfare but from the degree to which they share preferences about the

kinds of outcomes that should be accorded the highest priority. If coethnics are more likely to have preferences in common over different public goods—as, again, some theorists of ethnicity assume⁹—then ethnic diversity can imply a diversity of preferences or tastes that, as the example illustrates, makes collective ends more difficult to achieve.

Finally, consider a very different way in which preferences might affect incentives to work together. So far we have focused on the *results* of working together. In fact, however, when choosing with whom to work on a project, individuals may also take account of the *process* of working together. People may share similar goals but refuse to work with one another because of mutual antipathy. The third preferences mechanism, which we call *preferences over process*, emphasizes explicitly these procedural features of collective action.¹⁰ Recent work by social psychologists suggests that racism is associated with feelings of disgust that exert direct effects on the willingness of individuals to work together even for projects of mutual benefit.¹¹ While political scientists tend to focus on the consequences of collective action more than on the processes of participation, recent work on the politics of collective action also finds that a focus on process-oriented motivations can provide additional explanatory power, even in high-stakes settings such as civil war (Wood 2003).

Technology Mechanisms Let us now turn to what we call *technology mechanisms*. The idea here is that ethnicity affects the set of strategies available to players. To make this idea more concrete, imagine that a local council chairman in Kampala decides to try to address the problem of flooding by applying for funds from the government to construct additional concrete drains for his zone (one option that local councils have in fact considered). He might ask two community members to take the lead in drafting the proposal to division headquarters. Under what conditions will these two invest their time in such an enterprise? Presumably, they will be more willing to do so if they think they can work together effectively. In weighing this question, one factor they are likely to consider is whether they speak the same language, share similar experiences, and can draw upon a common understanding about how to work together to facilitate their interactions. To the extent that they do, the feasibility of their collaboration will be greater and their willingness to embark upon it will be greater as well. By providing individuals with precisely this kind of reservoir of common cultural material, shared ethnicity can serve as a technology that facilitates coordination and collaboration within the community.¹² And since any “ethnic technology” is not shared across group lines, it will only facilitate collaborative endeavors—including public goods provision—in more ethnically homogeneous settings.¹³ We term this the *efficacy mechanism*.

Sharing a common language or culture is not the only tool that might

facilitate community members working together. The second technology mechanism we identify is the *readability* mechanism. The idea here is that coethnics may be better able to (or believe they are better able to) read cues about the positions or intentions of potential partners. As a consequence, a given partner (who is a coethnic) may be able to select a strategy conditional on information that is not available to a non-coethnic. This can be highly advantageous. For example, someone might be more willing to work with another person if she can see that he is smart or organized or has experience in the relevant area. Some of these attributes might be clearly observable, but others might not be. To the extent that sharing an ethnic background with a person makes a potential cooperating partner more “readable,” coethnicity can be an important tool. In principle, for example, it is possible that individuals take actions to benefit coethnics disproportionately, not because they care more for them but because they can better target their beneficence.¹⁴

Whereas these first two technology mechanisms emphasize individual-level characteristics that coethnics share, the last two technology mechanisms both relate to the structure of interactions between coethnics more generally—in particular, the ways in which social networks may be structured on ethnic lines. We distinguish between two ways in which such networks can matter. The first, which we term the *periodicity* mechanism, suggests that coethnics interact with one another with greater frequency. This might matter because, as a large empirical and game-theoretic literature shows (Axelrod 1984), the advantages of sustaining cooperation into the future rise relative to the short-run benefits of defection when individuals expect to interact more frequently. Homogeneous groups might be better able to mount joint activities because members are reasonably confident that they will find themselves contemplating collective activities with others on a regular basis in the future. In the specific case of trying to reduce crime, for example, an individual may volunteer her energy for a community patrol knowing that other residents will do the same because, if they do not, she knows that she can withdraw her willingness to be an active participant later on, making other residents worse off.

A second, closely related way in which networks might matter is through what we call *reachability*. Under this mechanism, social networks are used to collect information about a potential cooperating partner’s unobservable skills or experience (or deficiencies) or to facilitate the sanctioning of a person who reneges on his or her agreement (for example, by spreading information about the person’s untrustworthiness so that others will know to be wary of him or her in the future).¹⁵ Note that a key difference between this channel and the previous one is that the reachability mechanism can operate even if repeated encounters are rare.¹⁶ To the extent that network ties are stronger between coethnics,

these benefits may work disproportionately to the advantage of homogeneous communities. Returning to the obstacles faced by local leaders in Kampala's slums, this argument suggests that homogeneous communities may be better able to address shared challenges because they have a tool for disciplining the behavior of residents who do not cooperate. Community residents may contribute to the local patrol or dedicate time to garbage cleanup because they fear the costs to their reputation—which can spread through the homogeneous community, but also more broadly across the ethnic network—if they sit on the sidelines.

Strategy Selection Mechanisms The third family of explanations focuses on what we call *strategy selection* mechanisms. While the previous mechanisms we examined highlighted the ways in which people may be playing different games when they interact with coethnics and non-coethnics, strategy selection mechanisms posit that people play the same game differently depending on the identity of their partners.¹⁷ To illustrate how these operate, imagine that Chairman Ssalongo has organized a few young men to dedicate a Saturday morning to clearing the drainage channels. Will they show up? If it is the case that contributing to the public good makes sense if and only if all the young men participate—for example, if the task cannot be done by one person alone because large pieces of garbage obstruct the flow of water—then the rationality of showing up will depend entirely on an individual's beliefs about what the others are likely to do.

In one group of young men chosen by the chairman, each individual might expect others not to contribute and, in response, will not contribute (thereby jointly ensuring that their expectations are correct). In another group, however, each might expect the others to contribute and, in response, will also contribute (again jointly ensuring that their expectations are correct). If such expectations are related to group membership—for example, if, as a large anthropological literature suggests, people expect someone from their own ethnic group to contribute but not someone from a different ethnic group—then it is straightforward to see how ethnicity might affect the strategies that people select and the likelihood of successful collective action.¹⁸

The difference between the strategy selection and technology channels is subtle, particularly in the area of sanctioning. Assume for a moment that sanctioning norms exist among coethnics but not among non-coethnics (or exist more strongly among the former than among the latter). As a result, people play different strategies depending on whether they are interacting with members of their own ethnic community. Is this evidence of a technology mechanism or a strategy selection mechanism?

To answer the question, we need to know whether the technologies

that facilitate sanctioning (such as reachability and periodicity) are such that reciprocity can only be sustained among coethnics, or whether these technologies could permit reciprocity among both coethnics and non-coethnics but only coethnics employ them to this end. To illustrate, suppose that we observe that two non-coethnics are less likely to cooperate on a collective task than two coethnics. A technology story might be that the non-coethnics fail to cooperate because they interact with insufficient frequency or are insufficiently connected to one another within social networks to make cooperation feasible. In a strategy selection story, on the other hand, though cooperation is feasible, neither individual cooperates because she does not expect the other to cooperate. The technology story is that everyone might like to condition cooperation on cooperation, but only coethnics can. The strategy selection story is that everyone can condition their strategies in this way, but only coethnics do.

Although we recognize that these three broad families of mechanisms may interact in complex ways—for example, the existence of ethnic technologies may reflect a group's preferences, or vice versa, and the choice of strategies may be partly determined by preferences and by beliefs about the preferences of others—we believe the set of families to be exhaustive. In particular, they subsume the major hypotheses advanced in the literature for the negative relationship between ethnic diversity and public goods provision. For example, arguments advanced by Alberto Alesina, Reza Baqir, and William Easterly (1999) and by Alesina and Eliana LaFerrara (2005) emphasizing the correspondence between ethnic groups and preferences for particular kinds of public goods are in line with our preferences mechanisms. The hypotheses proposed by Timothy Besley, Stephen Coate, and Glenn Loury (1993), Avner Greif (1993), Edward Miguel and Mary Kay Gugerty (2005), Parikshit Ghosh and Debraj Ray (1996), Marcel Fafchamps and Bart Minten (2002), and Barak Richman (2006) about the role that networks play in facilitating punishment echo our technology mechanisms, as do the hypotheses advanced by Cagla Okten and Una Okonkwo Osili (2004) and Alessandra Casella and James Rauch (2002) about the lower transaction costs of coethnic interactions and the arguments developed by Michael Bacharach and Diego Gambetta (2001) on the readability of cooperation partners (see also Gambetta and Hamill 2005). And strategy selection mechanisms are frequently invoked to explain the greater collective action we observe in homogeneous settings than in nonhomogeneous ones—particularly in situations where multiple equilibria exist and players must coordinate their choices to maximize their payoffs (as, for example, in an iterated Prisoner's Dilemma game, a Battle of the Sexes game, or an Assurance game). In such situations, the strategy is simply the rule of thumb that has emerged, perhaps through some evolutionary process,

which permits the players to achieve a more desirable outcome.¹⁹ One example is the norm among U.S. senators to lend their support to a colleague's bill in exchange for the promise of reciprocated support at a later date (Mayhew 1975), another is the norm among Italian and Jewish kids in Brooklyn to punish their own rather than risk a spiral of mutual retaliation from members of the other group (Fearon and Laitin 1996), or the norm to "live and let live" that emerged among British and German soldiers positioned within firing distance of one another in the trenches during World War I (Axelrod 1984).

Ethnic diversity might affect public goods provision through any of these three broad channels. From the standpoint of testing the relative explanatory power of each mechanism, the problem is that they all generate the same prediction: lower cooperation in a context of greater social diversity. This makes it impossible to distinguish among them simply by observing the success of homogeneous (or the failure of heterogeneous) communities to provide public goods for their members. Thus, if, in an ethnically homogeneous community, we observe that residents cooperate to pick up the trash, maintain drainage channels, and patrol the streets, we have no way of knowing whether the roots of this successful cooperation lies in a sense of other-regardingness that they feel toward one another; a calculation that their collaboration is likely to be particularly easy or fruitful; or the existence of a norm that makes not reciprocating a coethnic's cooperative overture, or not punishing a failure to reciprocate, unthinkable. Sorting out whether the pattern we observe is best explained by a preferences, technology, or strategy selection mechanism (or by some combination of the three) requires that we find a way to identify situations in which we can rule out the operation of different mechanisms. To solve this difficult inferential problem, we turn to the technique of experiments.

Empirical Strategy

The Experimental Method

Experiments are advantageous because they permit a high degree of control over the factors that might affect the outcome under investigation. Their power comes from enabling researchers to isolate and test the explanatory power of competing explanations that are difficult or impossible to disentangle in real-world settings. Their great weakness lies in their questionable external validity—that is, experiments leave researchers uncertain in extrapolating from their experimental findings to the larger world whose behavioral patterns they seek to explain. Assessing the extent of this trade-off in a given setting is critical to evaluating

the merits of experimental work. While our own strategy involves bringing the laboratory into the field—in an effort to strengthen the external validity of our findings—we also believe that a great deal can be learned in highly controlled settings, as the following examples illustrate.

When evidence of the Holocaust began to emerge after World War II, it gave rise to hundreds of scholarly inquiries into how tens of thousands of “regular” Germans could have been transformed into killers. Some of these analyses were anthropological: they focused on aspects of German culture that might have facilitated the rise of the Nazis and the implementation of their “Final Solution.” Others were sociological investigations that emphasized the nature of German society, the role that Jews played within it, and the rationale that this may have provided for some Germans to turn on their neighbors. Still other inquiries delved into the organizational apparatus of the German state. The best of these studies made detailed analyses of German history, society, and culture based on interviews, primary sources, and careful archival research. Yet one of the most compelling explanations for the willing participation of “regular” Germans in the execution of 6 million Jews came from a study that involved no field research, no interviews, and not a single German source.

In an attempt to understand how the Holocaust could have happened, the psychologist Stanley Milgram (1974) recruited forty subjects to participate in an experiment in the basement of a building in New Haven, Connecticut. The subjects were told that the purpose of the experiment was to study the effects of punishment on learning. They were told that they would play the role of the “teacher” while another subject (in fact a confederate of the experimenter) played the role of the “learner.” After being taken to a different room from the confederate, subjects then were told to administer a series of increasingly large electric shocks to the learner (in fact, no shock was administered), who, as the voltage increased, would (pretend to) shriek in pain and beg for the experiment to be stopped. If the subject expressed a desire to stop the experiment, the experimenter would ask the subject to “please continue.” If the subject continued to resist, the experimenter would push further, first telling the subject that “the experiment requires that you continue, please go on,” then that “it is essential that you continue,” and finally that “you have no choice, you must continue.” Milgram’s extraordinary finding was that about two-thirds of the subjects obeyed the experimenter’s commands and continued administering the electric shocks up to the highest level, despite the fact that they were led to believe that the person receiving the shocks was, by that point, in agonizing pain. By demonstrating the power of an authority figure to compel behavior—even behavior that conflicts with a person’s moral con-

science—the Milgram experiment provided critical insight into how the Holocaust (and also the Rwandan genocide and other instances of mass killing) might have been possible. It also demonstrated the power of experiments to shed light on important social outcomes that are very far removed from the subjects involved or the tasks they are asked to perform.

Robert Axelrod's (1984) research on the evolution of cooperation provides another example of the power of "artificial" games to provide deep insight into real-world phenomena. Axelrod was interested in understanding how cooperation can emerge in a world of egoists—be they nations, people, animals, or bacteria—with no central authority. To gain leverage on this question, he might have conducted an in-depth study of international bargaining. He might have done a careful analysis of collusion among drug companies. Or he might have observed children playing in a park. Instead, Axelrod invited mathematicians, economists, sociologists, and political scientists from around the world to submit strategies to be pitted against one another in a computer-based, iterated Prisoner's Dilemma tournament. Remarkably, he found that all of the most cunning and sophisticated strategies proposed by these strategists were beaten by an extraordinarily simple strategy. His main finding was that a simple strategy of "tit-for-tat"—starting off by cooperating and then simply responding in kind to whatever one's partner does—works extraordinarily well. Although it might seem that nothing could have been further removed from the real-world problems that motivated his research, Axelrod's computer tournament yielded results that have contributed tremendously to our understanding of the evolution of social cooperation.²⁰

Experimental Games

Our decision to employ the methodology of experimental games stems from a belief, shaped by the examples of the research of Milgram, Axelrod, and others, that progress in understanding the mechanisms through which ethnic diversity affects public goods provision can be made by uncovering the essential patterns of behavior that underlie or impede collective action among subjects playing laboratory games. We might have gone another route. We might have elected instead to conduct an in-depth qualitative study of a community that had been particularly successful or unsuccessful in providing public goods. We might have used survey data to identify correlations between observed levels of public goods provision and different kinds of preferences, different types of networks, or different kinds of norms across a broad range of communities. Instead, we decided to use experimental games.

We did this because we believe the experimental approach provides

considerable leverage that is not available from alternative methods. This approach allowed us to study a set of games (described in detail in chapters 4 and 5) designed specifically to test the distinct mechanisms outlined in this chapter. As subjects played multiple rounds of each game with randomized matching—sometimes with coethnics, sometimes with non-coethnics—we were able to rule out many confounding factors and thus to assess the explanatory power of the distinct mechanisms by examining how play among coethnics and non-coethnics varied within and across games.²¹

In turning to experimental games to study the impact of diversity on collective action, we join a growing group of social scientists who have deployed similar techniques to study altruism, cooperation, bargaining, coordination, and trust (for good reviews, see Roth 1995; Camerer 2003). This literature has its roots in the pioneering work of Daniel Kahneman, Vernon Smith, and others who first used experimental methods to challenge neoclassical assumptions about the purely selfish motivations of individuals. Their experiments showed that undergraduate subjects played laboratory games in ways that contradicted the predictions of economic theory. Rather than simply trying to maximize their earnings, these subjects played the games in ways that suggested they were altruistic and cared about the fairness of the allocations to themselves and the other players. These early experiments have been replicated, expanded, and improved upon in thousands of studies, and experimental economics techniques have become accepted and embraced in a growing number of social science disciplines.

The implicit claim in the earliest foundational studies was that human beings are all the same and that we can learn something general about human behavior by studying the decisions of undergraduates in (principally American) college classrooms and experimental laboratories. A major concern of the literature that has followed has been to investigate the implications of relaxing the rather unrealistic assumptions that underlie this claim. To see where our study fits within the larger experimental tradition, it is helpful to review these efforts.

We can identify three broad branches of responses. The first has aimed to “take the laboratory out of the classroom” to investigate whether regular citizens behave differently from college students, who in terms of age, education, and other potential determinants of social behavior make up a highly unrepresentative sample of the broader population (Smith 2000; Carpenter, Burks, and Verhoegen 2005; for a key critique of the reliance on college sophomores that motivated many of these studies, see Sears 1986). A growing number of studies have attempted to test the robustness of standard experimental findings outside the university (for a good review, see Cardenas and Carpenter 2008). The most ambitious attempt in this direction is the recent Founda-

tions of Human Sociality Project, in which researchers set up labs in the field to play standard experimental games in fifteen small-scale societies around the world (Henrich et al. 2004). Other examples include Carpenter, Daniere, and Takahashi (2004), Bohnet and Greig (2008), Marlow and others (2008), and Bahry and Wilson (2006).

A second branch of research has attempted to test whether patterns of behavior vary across subjects that possess different demographic characteristics. The issue of interest in this strand of the literature is not whether average patterns of play vary across societies, but whether they vary across players of different types in the same society. Thus, a number of studies have investigated whether altruism, trust, or reciprocity varies with gender (Eckel and Grossman 1996, 1998; Andreoni and Vesterlund 2001), race-ethnicity (Eckel and Grossman 2001; Fershtman and Gneezy 2001), age (Harbaugh, Krause, and Liday 2003), or even beauty (Andreoni and Petrie 2008; Wilson and Eckel 2006). Another set of studies have considered not just “main effects”—that is, whether women play differently from men, or whether African Americans play differently from Caucasians—but whether behavior is affected by the nature of the *dyad* or *pairing*—that is, whether men behave differently when playing with other men than when playing with women, or whether Caucasians play differently when paired with other Caucasians than when paired with African Americans. Examples of studies in this vein include Burns (2003), Wilson and Eckel (2006), Ferraro and Cummings (2007), Fershtman and Gneezy (2001), Gil-White (2004), and Petrie (2003).

The third (and newest) branch has directly taken up the issue of external validity by explicitly relating the findings derived from experimental games to patterns of behavior observed outside of the laboratory. For example, Dean Karlan (2005) exploited information gathered in trust and public goods games to predict patterns of savings, repayment, and default, up to one year later, among participants in a Peruvian micro-credit organization. Jeffrey Carpenter and Erika Seki (2005) demonstrated that Japanese fishing crews composed of fishermen who exhibited greater degrees of conditional cooperation and were more willing to disapprove of shirking tended to be more productive. Others are now using experimental games to assess how levels of community efficacy and trust respond to outside efforts to inculcate democratic values and institutions in Liberia (Fearon, Humphreys, and Weinstein 2009).

The project described in this book is located at the confluence of these three branches of contemporary experimental economics research. We share with the first branch the characteristic of locating our research outside of a university laboratory as we examine how subjects behave in a context of substantive interest, an impoverished neighborhood of urban Kampala.

We share with the second branch a concern with how patterns of play vary with the characteristics of the players and the homogeneous or heterogeneous nature of the interaction. In making this move, we emphasize two different sources of heterogeneity. The first is based on membership in a given ethnic or regional group—our core interest in the book. By randomizing the partnering of players with one another, we are able to generate variation in the ethnic homogeneity or heterogeneity of the pairings we observe. The second source of heterogeneity is based on differences across individuals within groups. One of the most important contributions of behavioral economics has been to challenge the neoclassical view that all individuals respond in the same way to material incentives. Instead, as summarized by Elinor Ostrom (2000, 138), recent research suggests that “the world contains multiple types of individuals, some more willing than others to initiate reciprocity to achieve the benefits of collective action.” Recognizing the existence of these different types has implications for empirical and theoretical work (see, for example, Bolton and Ockenfels 2000; Bowles and Gintis 2004b; Fehr and Schmidt 1999). If players of different types respond in different ways to a given treatment, then an aggregation problem may arise. Average behavior across types may mask systematic features of play taking place within them. To avoid this pitfall, we distinguish in our analyses between two types of players: those whose behavior is consistent with preferences in keeping with the neoclassical model (whom we term “egoists”) and those who exhibit higher levels of general altruism (whom we term “non-egoists”). Although we examine the different behavior that arises across different types of individuals, we emphasize that, as with other work in this strand of the literature, we do not exercise experimental control over these types. These all exist prior to our experiments. What we do have control over, however, is how these types encounter each other in social settings. In this way, we can ensure that the interactions we examine are free of the selection effects that render much observational analysis so difficult.

With the third branch of recent scholarship using behavioral experiments, we share a commitment to demonstrating the utility of our findings for explaining outcomes outside of the laboratory. To do this, we have insisted on a methodology that, though standard in other types of research in the social sciences, is surprisingly uncommon in experimental work: we used random sampling to recruit a pool of subjects who were representative of the underlying population whose behavior patterns in the real world we were interested in explaining. There are two major advantages of random sampling. First, by producing a representative sample, it allows subjects to make inferences about other subjects based on their knowledge of the population: they are thus able to form consistent beliefs about the ethnic identities and behaviors of the individuals with whom

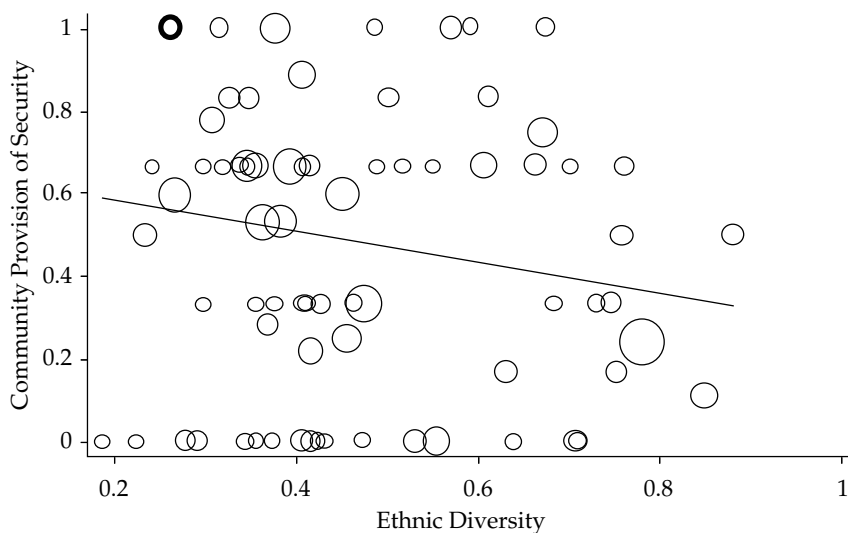
they are playing. Second, random sampling allows us, the researchers, to make inferences from the behavior of our sample to the population from which our subjects are drawn. Matching the subject and underlying populations is especially important when the behaviors under study—in this case, how people condition play on the ethnicity of their partners—are not necessarily a universal feature of human behavior, but rather a property that may apply in different ways to different populations.

The Research Site

We located our research in the poor urban neighborhoods of Kampala, Uganda. Kampala is a good site in which to examine the impact of diversity on public goods provision. First, Kampala is not only an ethnically diverse city but also a place where ethnicity is highly salient in everyday social interactions. Yet while ethnicity matters, the political situation in Kampala is sufficiently stable and peaceful to permit questions about ethnic identifications and attitudes to be asked and for research on social interactions across ethnic lines to be undertaken. Moreover, with the devolution of responsibility for social service provision over the past decade from the central government to the financially strapped elected local councils (LCs), the supply of many local public goods—including security, garbage collection, and the maintenance of storm drains—has become a purely local affair that depends almost entirely on the voluntary contributions of local community members (Golooba-Mutebi 2003; Onyach-Olaa 2003)—a subject to which we return in chapter 2. Thus, the question of why some communities are able to generate contributions toward public goods and others are not is of real practical consequence in the area we study.

In addition, and quite importantly, we were able to confirm in Kampala the negative association between ethnic diversity and public goods provision that has been found elsewhere in the literature. Figure 1.1 plots the relationship between ethnic diversity and public goods provision across seventy-four local communities in Kawempe Division, the poorest of Kampala's five divisions.²² We measured the ethnic diversity of each community using a fragmentation index based on the ethnic demography reported in the 2001 census, and we measured public goods provision in terms of whether, during the previous six months, residents of the local community had organized efforts in the area of crime prevention and security.²³ The data showed a negative, albeit weak, relationship between diversity and public goods provision and made it possible for us to take the next step of accounting for why.²⁴

Apart from allowing us to corroborate the negative association between ethnic heterogeneity and public goods provision, the data from the Kawempe survey also put us in a position to select a narrower (and

Figure 1.1 Ethnic Diversity and Public Goods Provision in Kawempe

Source: Authors' calculations.

Notes: The figure reports the relationship between a community-level measure of ethnic diversity and the likelihood that the community has organized to provide security for its residents in Kawempe, Uganda. Higher levels of ethnic diversity are associated with lower levels of community security provision.

more manageable) area from which to recruit our experimental subjects. Because we were interested in studying *why* rather than *whether* diversity impedes collective action, we deliberately recruited our experimental subjects from neighborhoods of Kampala that combine high levels of ethnic diversity with low levels of public goods provision. Our study area comprises the four adjacent parishes (LC2s) of Mulago I, Mulago II, Mulago III, and Kyebando.²⁵ We refer to them in this book collectively as Mulago-Kyebando. Together, these parishes contain approximately seventy thousand people from whom we drew a random sample of three hundred subjects.²⁶ Interviews with community leaders and residents of Mulago-Kyebando (reported in chapters 2 and 6) confirmed both the difficulties of providing local public goods in this community and the highly ethnically fragmented nature of their populations.

Putting the Strategy in Perspective

As with all methods, it is important to keep in mind the limitations of our empirical strategy. The first limitation flows from our decision re-

garding site selection. We chose to sample our subjects from a set of communities exhibiting high ethnic diversity and low public goods provision. Although this decision put us in a strong position to evaluate the connection between diversity and the underprovision of public goods, it left us without a benchmark for each of the various mechanisms we tested. Thus, while we could determine whether a given mechanism played an important role in linking diversity to collective action failure, we could not be certain that a given mechanism worked more or less strongly in more or less heterogeneous settings. To make such a comparison, we would have had to replicate our analyses in a broader cross-section of communities.

Other limitations follow from the nature of the social interactions captured in our experimental setting. One of the many choices we faced in designing the experiments was whether to aim to replicate settings with largely anonymous interaction or to examine the operation of small-scale communities. We chose the former, sampling subjects from a large area and setting up pairings in which most interactions were anonymous. This anonymity renders our results more salient for examining interactions such as participation in urban settings and in regional or national politics—rather than, for example, in village or elite politics, a context in which subjects are almost certain to know one another.

A second choice was whether to examine interactions in which ethnicity was highly salient versus those in which it was not. We chose to study a setting in which ethnicity was not explicitly salient—in the sense that, for most of our games, we did not suggest to our subjects that ethnicity was a relevant variable to be considered in choosing how to play. Of course, ethnicity was in fact salient to many of our subjects, as evidenced in the results of our games, but ethnic patterns of play were not a consequence of any priming on our part. Thus, our results speak most directly to the ways in which identities condition everyday interaction. We cannot rule out the possibility that individuals behave differently in settings in which ethnicity is specifically emphasized as a salient category of social interaction, as in political campaigns or riots or in contexts where ongoing intergroup violence colors people's expectations about the outcome of cross-group interactions.²⁷

Perhaps the most important limitation of our main empirical strategy is its emphasis on testing mechanisms through the observation of individual choices in two- and three-person experimental games. This approach gives great tractability to our design, but it has the disadvantage of implicitly reducing community- and group-level processes to an aggregation of binary or small-group interactions. We cannot rule out the possibility that homogeneous subgroups of a heterogeneous community act very differently from homogeneous subgroups of a homogeneous community (on such composition effects, see Posner 2004a).

In addition, a rich literature offers theories and empirical accounts of collective action success that emphasize aspects of social behavior that are not easily captured in this experimental setup. For example, building on the seminal work of Mancur Olson (1965), scholars such as Norman Frolich, Joe Oppenheimer, and Oran Young (1971) and Samuel Popkin (1979) have pointed to the centrality of political entrepreneurs—individuals who are willing to bear private costs in order to facilitate the production of public goods. Such entrepreneurs are absent in our study. The approach we adopted also limited our ability to speak to the importance of institutions more broadly (both formal and informal), a prominent topic in the theoretical literature on the origins and maintenance of cooperative behavior (Ostrom 1990; Cook, Hardin, and Levi 2005).

We recognize the fact that the success or failure of collective action in a community may depend on factors that we could not observe directly given our empirical strategy. But we believe nevertheless that a great deal can be learned from observations of binary and small-group interactions. Many of the prominent accounts of how informal institutions shape collective efficacy (for a review, see Sampson, Morenoff, and Gannon-Rowley 2002) depend very much on the notion that such norms exert an effect directly through the choices that individuals make in everyday interactions (for example, whether to help out a neighbor or report a crime). In principle, the effects of informal institutions should be observable in our games, even though such institutions and practices are cultivated at the community level. This question of how to link patterns of individual behavior to aggregate outcomes is clearly an important one—and one to which we return in chapters 6 and 7.

The Outline of the Book

The book is organized as follows. Chapter 2 provides a description of our research site and presents information about the variation in levels of public goods provision across our site and in the region more broadly. We explore how a wide range of public goods—roads, schools, health care, sanitation, and security—are provided in Kampala, document the erosion of the central government's role in providing these goods, and discuss the challenges faced by local communities that seek to organize their members to improve social welfare. Chapter 3 focuses on our main independent variable, ethnic diversity. The chapter provides basic background information about ethnicity in Kampala. In laying the foundation for an analysis of how ethnicity conditions behavior, we emphasize a new understanding of ethnic diversity based on how people themselves—rather than census-takers, analysts, or other outsiders—categorize the members of their own community in different informational set-

tings. The chapter then introduces a method we have developed to measure these subjective perceptions.

Explaining how ethnic diversity undermines the ability of communities to organize the provision of public goods is the objective of chapter 4, which is in many respects the heart of our study. Here we explore the myriad arguments advanced to account for diversity's negative impact, looking for evidence in support of (or against) key preferences, technology, and strategy selection mechanisms. Drawing on the behavior of individuals from Mulago-Kyebando in a series of experimental games, the evidence tells a powerful story about why ethnically homogeneous communities are better able to act collectively than more diverse ones. The answer is simple, yet consistent with what anthropologists have chronicled about African ethnic groups for generations: shared identity makes in-group norms of reciprocity salient. What this means in practice is that individuals cooperate with coethnics at a higher rate than with non-coethnics, not because of biases toward in-group members, but because they expect coethnics to cooperate with them and because they believe that, should they fail to cooperate, they might be punished. This behavior is most prominently displayed by "egoists"—individuals who, absent the reciprocity norms made salient by shared ethnic identities, show a stronger tendency to look out for their own interests at the expense of others. Although our results point to the centrality of the strategy selection mechanism, they also suggest that ethnic technologies are at work: that is, that coethnics engage one another with far greater frequency (periodicity) and are better able to gauge the otherwise unobservable characteristics of fellow group members (readability), and that coethnics may be more effective at accomplishing joint tasks (efficacy) and have access to social networks that facilitate the sanctioning of coethnics who refuse to cooperate (reachability).

Equally important as what we did find is what we did not. Contrary to assumptions in much of the theoretical literature on ethnic diversity and public goods provision, our results suggest that preferences mechanisms cannot account for the failures of collective action in urban Kampala. Individuals in our study showed no tendency to value the welfare of members of their own ethnic group more highly than the welfare of members of other groups. We found almost no variation across groups in their preferences about what public goods should be provided. And we found no evidence that our subjects derived greater enjoyment from working with coethnics than with non-coethnics.

The empirical evidence in support of both ethnic technologies and ethnic strategies requires us to push our analysis further. In chapter 5, we ask: Is it the case that, even conditioning on how easily they are reached and sanctioned, individuals still act in a more reciprocating manner when playing with coethnics? If this is the case, it provides evi-

dence for a distinctly coethnic norm. Or is it simply *because* they interact with coethnics more frequently and are more easily reached by coethnics that individuals reciprocate more when paired with someone from their group? If so, this might constitute evidence that coethnic cooperation arises from the uneven application of a universal norm. Although the distinction between these two mechanisms is subtle, it is essential for resolving the puzzle of why ethnic diversity undermines collective action. Evidence from additional experimental games led us to a more precise answer. More frequent interaction and greater mutual reachability were not the factors that drove coethnic cooperation. Rather, our subjects appeared to observe a distinctly coethnic norm of reciprocity—one that operated alongside a powerful universal norm that emerged when behavior was publicly observable and the threat of punishment by a third player was present. This last finding suggests that although reciprocity norms are stronger within groups than across them, it is possible to sustain cooperative norms that span ethnic groups in highly diverse societies.

Chapter 6 draws us out of the experimental setting and into the community of Mulago-Kyebando to investigate whether the patterns we observe in the games are reflected in behavior that we can measure outside the laboratory. We begin with a deeper exploration of our subject population, examining the stories they told us about how they understood the games, showing how they linked play in the lab to situations they confronted in their everyday lives, and looking to see how their game behavior correlated with their participation in community life. Moving from our subjects to the community, we then probe the sources of collective action failure in Mulago-Kyebando as described by the chairpersons of its twenty-six local councils, many of whom echo chairmen Ssalongo and Kashaija frustration at their inability to police noncontributions to community projects. Moving from the narrow to the more general, we subject our argument to tests of its external validity before turning to the implications of our argument in chapter 7.