

Chapter One

Introduction: Diversity and Collective Action

Why are some communities able to generate public goods for their members and others not? Why are some communities blessed with good schools and health care, adequate sanitation, safe drinking water, and low crime, while others are cursed with insecurity, squalor, and pervasive ill health? Wealth is obviously a factor, as prosperous communities can afford to hire police officers, 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 extremely low levels of public goods provision in some very wealthy ones. This is because 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 or not it can buy better public services by simply writing a check.

But if a community's cooperative capacity is central to the generation of public goods, then what accounts for this cooperative capacity? A growing literature suggests that part of the answer may lie in its degree of 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 can impede the provision of public goods. What this literature does less well, however, is to advance and test an explanation for *why* ethnic diversity might have this negative effect. This book aims to fill that gap. It takes the negative relationship between ethnic diversity and public goods provision as its starting point—and, in fact, reconfirms it in yet another case—but then moves beyond this association to probe its sources.

The Problem of Public Goods Provision

Everyone would prefer to live in a community with safe streets, good schools, clean water, accessible health care, and reliable garbage collection. The problem is that many would also prefer to consume these public goods without contributing any of their own time or money to providing them. Public goods present communities with a classic collective action problem. They are, by definition, non-excludable: that is, they can be consumed and enjoyed by everyone in the community irrespective of

whether or not they contributed to their provision.¹ The problem is that this leaves everyone with an incentive to leave the hard work of providing them—the costly contributions of time or money or labor—to others. And if everyone “free rides” in this way, then the public goods will never be provided.

To make the problem concrete, consider the example of Kenneth and Alex. Both own shops on the edge of town where Kenneth sells pots and pans and Alex sells sugar and cooking oil. Their shops are located at the end of a road so heavily potholed that most drivers are reluctant to make the trip to patronize their businesses. The question is: who, if anyone, will fill the potholes? Kenneth knows that if he lies low while Alex works on the potholes by himself, he will reap benefits from Alex’s hard work and Alex will have no way to prevent him from doing so. Kenneth thus has an incentive to shirk. Alex, of course, makes the same calculation. A likely result is that neither will work on the potholes and their businesses will have fewer customers than they would have if they had been able to cooperate.²

The difficulty of organizing community members to clean up garbage, maintain drainage channels, sustain neighborhood watch groups, make contributions to local schools and clinics, and provide other kinds of local public goods stems from the same kind of dilemma that Kenneth and Alex face. As long as community members think in this way, garbage will go uncollected, drainage channels will remain choked with debris, criminals will continue to roam freely, and schools and health care facilities will remain under-funded. Individual incentives to shirk serve to undermine the collective wellbeing.

Note that one of the key details in Kenneth and Alex’s story is that they could not count on the local government to fix the road. As far back as Hobbes, one of the principal justifications of government has been to supply public goods that individuals fail to provide on their own. This does not, of course, mean that governments always do this. Indeed, governments quite often cannot be counted on to deliver the public goods that community members desire. Sometimes this is because they lack the resources or the capacity to provide them. Equally frequently, it is because the

¹ In the classic definition, a second criterion of public goods is that they are non-rival. That is, one person’s consumption of the good does not diminish the ability of others to consume it as well. For the obstacles to public goods provision that are of interest to us here, however, non-excludability is the more important aspect.

² This informal description of the collective action problem actually describes a family of distinct strategic situations enumerated by game theorists. To see this, assume that each of the two players is less happy with both not contributing than with both contributing. Then, if each individual would prefer not to contribute no matter what the other does then the problem is a Prisoner’s Dilemma; if each would rather contribute if and only if the other contributes, then the problem is an Assurance game; if each would rather contribute if and only if the other does *not* contribute, then the problem is a Chicken game.

community has failed to exert sufficient pressure on the government to act on its behalf: indeed, access to good government is itself a public good. So, either way, the community has to solve a collective action problem, either by substituting its own efforts for the absent efforts of the government, or by organizing itself to lobby the government to do its job.

As a general matter, the conflict between individual and collective incentives makes public goods difficult to provide. At the outset, however, we noted that these collective action problems seem to be even more difficult to solve in ethnically diverse settings. 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 co-ethnics—seem to have an advantage in achieving collective ends? Before we proceed to answer this question, we review the evidence linking diversity and the failure of collective action and explore the arguments that have been provided 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. They require frequent maintenance to keep them functioning. Some communities succeed in maintaining their infrastructure while others fail. To find out why, Asim Khwaja (2004) 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 was still being satisfied, and the amount of maintenance work that had been 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 10 percentage 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 critical infrastructure likely has severe consequences for people's wellbeing.

Working in a very different environment, Edward Miguel and Mary Kay Gugerty (2005) studied the determinants of school funding and quality in 84 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 achieve collective action to raise funds for these purposes. Controlling for a wide range of socioeconomic, geographic, and demographic controls (including the possibility that, by generating in-migration, high school quality might increase community-level heterogeneity) Miguel and Gugerty find that communities at average levels of ethnic diversity generate 20 percent lower school contributions per pupil than entirely homogeneous communities. As in the Khwaja study, this is a substantively large 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. Okten and Osili (2004) demonstrate that community-level diversity in Indonesia is negatively related to contributions of labor, materials and money to neighborhood irrigation associations, security arrangements, rice cooperatives, and local health centers.³ Bardan (2000) reports a negative association between caste heterogeneity and the successful maintenance of local irrigation schemes in India; Banerjee et al (2005) and Banerjee and Somanathan (2004) report a similar relationship with respect to the provision of schools, public transport, electricity, health centers, and water projects. Dayton-Johnson (2000) provides evidence for the negative impact of social diversity on the success of community irrigation schemes in Mexico. Baland et al (2001) extend this work to the management of forestry resources in Nepal, where they, again, find a negative association between diversity and community-level collective action.

These findings are not limited to developing countries. Alesina, Baqir and Easterly (1999) present evidence that diversity is associated with lower levels of funding for schools, roads, sewers and trash pick up in U.S. municipalities. Their findings corroborate those of Poterba (1997) and Goldin and Katz (1999), who report negative associations between school funding and ethnic diversity in U.S. cities and school districts, and, recently, by Putnam (2007) who finds social diversity to be negatively associated with trust, social capital, and a variety of measures of public goods provision. Vigdor (2004) finds that racially diverse communities had lower response rates to the 2000 U.S. Census, which, because the allocation of federal funds to communities is related to their response rates, he interprets as a failure of collective action. Videras and Bordoni (nd) find that more ethnically heterogeneous zip codes in New York and New Jersey have lower levels of environmental

³ Their finding is confirmed by Nugent and Swaminathan (2005), who use similar data from Indonesia and reach the same conclusion about the impact of diversity on contributions to village health posts.

protection. Li (2005) extends this work to Europe, where, using data from the European and World Values Surveys, she finds that ethnic heterogeneity is negatively related to tax compliance.

The relationships identified in these studies are all statistical relationships and thus provide evidence of correlation rather than causation. Nevertheless, while in many specific cases heterogeneity will not be associated with poor public goods provision, taken together, these studies provide support for the proposition that on average diversity impedes the provision of public goods. Summarizing the literature on the topic, Banerjee et al go 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” (2005: 639).

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 do venture hypotheses about the mechanisms that might be at work, no study to date has offered a systematic enumeration of the possible causal channels through which ethnic diversity might undermine public goods provision along with a test of the comparative explanatory power of each.⁴ 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 or no 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, this book departs from much work in social science, which typically concerns itself with *whether*, rather than *why*, one thing causes another. The norm in the literature is to focus attention on establishing that x causes y . Answering the question of *why* this relationship obtains is usually secondary. Some kind of explanation is typically ventured, but these tend to be *post hoc* attempts to account for 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 of the reasons for the relative inattention to testing causal mechanisms is because, in many cases, simply establishing that the causal relationship exists is achievement enough. The

⁴ Okten and Osili (2004), Alesina and LaFerrara (2005), and Bowles and Gintis (2005) are partial exceptions in that they lay out a number of possible mechanisms.

literature that we focus on here offers a good example. It is filled with studies that provide suggestive evidence for the link between ethnic heterogeneity and collective action failure, but, as we suggested above, 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 because 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—is hard. This is particularly so in the study of ethnicity, since there exists no generally accepted framework that one can take “off the shelf” to think analytically about how ethnic identities might shape individual behavior, and thus how a community’s ethnic diversity might affect 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 is 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 (i.e., a set of actors), a technology (i.e., a set of strategies available to each of these actors), and preferences (which describe the way that individuals value the outcomes that result once all the actors select their strategies).⁵ The overall outcome of the game can be written as a function of the actions taken by all the players, typically predicted by theorists using a solution concept.

This simple game description provides us with three first order channels to investigate in order to identify how individuals’ ethnic identities might affect the outcome of their interactions. In principle, ethnicity could affect a person’s behavior by changing the population playing the game, the strategies available to the players, and their preferences, or by changing the selection of strategies, conditional on these factors. Assuming that ethnicity is predicated on existence (that is, we do not treat the population 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: preferences mechanisms, technology mechanisms, and strategy selection mechanisms.

⁵ The description we use here is that of a game in normal or strategic form. Games in extensive form use more primitives; those in characteristic function form use fewer.

Each of these families is composed of 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. We describe each in turn.

[Table 1.1]

To illustrate the logic of the *preferences* mechanisms, consider again the example of Kenneth and Alex. Recall that one reason why both men might shirk is that each one hopes that the other will do the work for them, and each one fears that, if he shows up, he will be the “sucker”—the one who provides all the work but then has to share the fruits of his labor. In some cases, this problem corresponds exactly to what game theorists call a “Prisoners’ Dilemma.” In such cases, no matter what the other player does, each player would rather stay at home. As a result, the potholes get left as they are even though both would rather see them filled. This unhappy outcome follows in part from the assumption that neither individual benefits from improvements in the others’ welfare. But what if Kenneth liked not just to see his own business flourish but also to see Alex’s shop thrive? If this were the case, then Kenneth might be willing to volunteer his labor for the road mending project regardless of what Alex does. Of course, the converse is equally true: if Kenneth actually liked to see things go badly for Alex, then he would be less likely to bear the costs of the road repair efforts, and the likelihood that the potholes would be filled would *decrease*. As the example makes clear, having positive (or negative) preferences for the welfare of others can help solve (or exacerbate) the collective action problem (or change one type of collective action problem into another). To the extent that ethnicity is correlated with *other regarding preferences*—as many theorists of ethnicity assume it is⁶—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.

We have just described mechanism I.1 in Table 1.1. We see from the table that there are other ways in which preferences might affect collective action success as well. To illustrate the second mechanism, which we term “*Preferences in Common*,” suppose that the local government in

⁶ For example, see Horowitz (1985). Horowitz takes his cue from a rich tradition in social psychology that emphasizes how a sense of belonging to a particular group, no matter how arbitrary the definition of group, leads individuals to demonstrate an in-group bias, even if it comes at a cost to their own welfare (Tajfel 1974).

Kenneth and Alex's town was, in fact, responsive to community members' requests for assistance. If Kenneth and Alex could coordinate their lobbying efforts on the same problem (such as the repair of the road leading to their shops) then the likelihood of a positive response from the local government (and a mended road) would increase. But if the two men were unable to coordinate their lobbying (for example, if Kenneth focused his efforts on convincing the local government to extend bus services to his home, while Alex, who owns a car but has school-aged children, focused his energy on lobbying the government to improve the quality of the local school), then the likelihood that either Kenneth or Alex would get what he wants would be diminished. Note that here the positive (or negative) outcome derives not from the degree of concern Kenneth and Alex have over each others' welfare but from the degree to which their preferences about the kinds of outcomes that should be accorded the highest priority are correlated. If co-ethnics are more likely to have preferences in common over different public goods—as, again, some theorists of ethnicity assume⁷—then ethnic diversity can imply a diversity of preferences or tastes that, as the example illustrates, can make collective ends more difficult to achieve.⁸

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. Kenneth and Alex may share similar goals but refuse to work with one another because of mutual antipathy—an antipathy that can exist even if each one cares for the other's welfare. Mechanism I.3 emphasizes explicitly these procedural features of collective action. 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. For example, in one recent study of what is known as the “disease-threat” model of inter-group attitudes, scholars found that women displayed lower propensities to work with out-group members in periods in which they were vulnerable, specifically during their first trimester of pregnancy (Navarrete, Fessler and Eng 2007). While political scientists tend to focus more on the consequences of collective

⁷ Bates (1973) points out that a salient feature of ethnic groups is that they are often geographically concentrated. This, he suggests, can induce a commonality of interests over outcomes that have a geographic component, such as the location of public investments. In a model of educational choices, Miguel (1999) proposes that individuals have preferences over the *type* of education they receive (for example, the language of instruction or the language in which parent-teacher meetings are conducted) that is correlated across co-ethnics. This can lead to an unwillingness to invest in education in more heterogeneous communities simply because the public good that is produced is of lower value to some individuals than others.

⁸ Note that strictly speaking, other regarding preferences and preferences over process can formally be treated as special cases of a commonality of preferences over “outcomes” broadly defined. Because other regarding preferences and preferences over process have received special attention in the literature, we treat them separately here.

action than 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).

Consider now the *technology* mechanisms. Here, ethnicity affects the “toolbox” of strategies available to players. Consider again the example of Kenneth and Alex. Suppose that rather than fix the road themselves, Kenneth proposes that they instead work together to write a proposal to the roads department to request that it repair the road for them. Under what conditions will Alex agree to invest his time in such an enterprise? One consideration may be whether or not he thinks that he and Kenneth can work effectively together. A factor that he is likely to consider in this regard is whether or not he and Kenneth possess a common reservoir of cultural material—language, experience, and understandings about modes of interacting—that they can draw upon to facilitate their interactions and, in the language of economics, reduce their transaction costs. To the extent that they do, the feasibility of their collaboration will be higher and Alex’s willingness to be a part of it will increase. By providing individuals with a reservoir of common cultural material, shared ethnicity can serve as a technology that facilitates coordination and collaboration within the community.⁹ And since the ethnic “technology” is not shared across group lines, it will only reduce transaction costs (and facilitate public goods provision) in more ethnically homogeneous settings.¹⁰ This is the “efficacy” mechanism (II.1) from Table 1.1.

Sharing a common language or culture is not the only tool that might reduce transaction costs and facilitate Kenneth and Alex’s joint activity. The second technology mechanism we identify is the “readability” mechanism. The idea here is that co-ethnics may be better able to read cues about the position or intentions of potential partners. In our example, Alex may be more likely to agree to cooperate with Kenneth if he knows that Kenneth is smart or organized or has experience in the relevant area. Similarly, he might be more disposed to help Kenneth if he thinks that Kenneth is in need. Some of these attributes will be clearly observable (for example, whether Kenneth is fit enough for potholing filling), but others may not be (for example, whether Kenneth is smart or willing to work hard). To the extent that sharing an ethnic background with a person makes a potential cooperating partner more “readable” co-ethnicity can be an important tool. In principle

⁹ Deutsch’s (1966) conceptualization of ethnic groups as “communities of communication” whose cohesiveness is a function of how quickly and effectively they can transmit complicated messages captures exactly this notion.

¹⁰ A recent paper by Spolaore and Wacziarg (2006) employs a variant of the technology mechanism in accounting for income differences across countries. The authors use genetic distance between populations as a proxy for the ease with which new technologies can diffuse across borders. In particular, they suggest that the horizontal diffusion of development may be impeded by language and other cultural differences.

then, it is possible that individuals take actions to benefit co-ethnics disproportionately, not because they care more for them, but because they are more confident that they can better target their beneficence.¹¹

The last two technology mechanisms identified in Table 1.1 both relate to the structure of interactions between individuals. Ethnicity may facilitate cooperation by providing a social network that potential cooperating partners can draw upon to increase the success of their collaboration. We distinguish between two ways in which such networks can matter. The first, which we term the *periodicity* mechanism, suggests that social networks function in such a way that co-ethnics interact with each other with greater frequency. If this were the case then, plausibly, co-ethnics might better be able to sustain cooperation though the threat of future non-cooperation in the event that cooperative actions are not reciprocated.

There is a second, closely related way in which networks might matter that makes fuller use of the structure of social networks. If an individual is “close” to another on a social network then the network might be used to collect information about a potential cooperating partner’s unobservable skills or experience (or deficiencies) or to facilitate the sanctioning of a cooperating partner in the event that he reneges on his agreement (for example by spreading information about the person’s untrustworthiness so that others will know to be wary of the person in the future).¹² The existence of the social network may also help increase the severity of the sanction that can be applied to a partner who reneges on a promise by increasing the value of the resource (the network) to which the partner’s access can be denied.¹³ To the extent that network ties are stronger between co-ethnics—that is, that they are more mutually “findable”—these benefits of networks may work disproportionately to the advantage of co-ethnic pairs.

The third family of channels operates through what we call *strategy selection* mechanisms. While the previous mechanisms we examined highlighted ways that people may be playing different

¹¹ A literature in economics on “statistical discrimination” explores how observable markers (such as ethnic identities) might be used by economic actors as a shortcut for assessing unobservable characteristics, such as trustworthiness or competence (Foster and Rosenzweig 1993; Fafchamps 2004). Here, we take an additional step suggesting that the ability to engage in such discrimination may itself depend on group membership.

¹² For a discussion of ethnic networks as information conduits, see Platteau (1994) and Landa (1994). Many of the insights about the importance of networks for the enforcement of contracts in environments with no sovereign authority build on work by Granovetter (1985) who described the “embeddedness” of business relations in a social context.

¹³ This is how the ultra Orthodox Jewish community in New York City is able to transact diamond sales on credit. The benefits of cheating or stealing are offset by the enormous social cost of ostracism from the community (Bernstein 1992; Richman 2006). On the importance of group sanctions, see also Kandori (1992), Greif (1993), Greif (1994), and Fafchamps (2004).

games when they interact with no-ethnics and non-co-ethnics, strategy selection mechanisms posit that people play the same game differently depending on the identity of their partners. To illustrate how these mechanisms operate, imagine that Kenneth and Alex are, again, contemplating a Saturday morning of pothole filling. When facing this collective action problem, it may be the case that contributing to the public good makes sense if and only if the other person also contributes—for example, if the task is necessarily a two-man job. In such cases, the rationality of contributing depends entirely on an individual's beliefs about the likelihood that the other person will contribute. For one pair of players, each might expect the other not to contribute and in response will not contribute, thereby confirming that their expectations are correct. For another pair, each might expect the other to contribute and in response will also contribute, thus confirming jointly their expectations. In such settings, it is possible that individuals expect one type of person (say a co-ethnic) to contribute, but not another type of person (a non-co-ethnic). The success of cooperation then is due entirely to the ways that individuals condition strategies on identity, not on preferences or technology.

Mechanism III.1 supposes that this dynamic is at play. Even if it is the case that for a given interaction each individual would prefer not to contribute, independent of the actions of the other, cooperation may still be supported through the systematic use of reciprocity norms over time. Specifically, it is possible that a group of players supports cooperation by playing strategies in which each cooperates as long as the other does too. But should one fail to cooperate, that individual will be punished by being excluded from future group benefits as the other player stops cooperating. This “tit-for-tat” solution to the collective action problem may exist in some groups but not in others. In some cases, each player might expect that all other players are playing tit-for-tat and do so in response; in other cases, each player might expect that no one will cooperate under any circumstances and so no one does. Like the static example, this strategy selection mechanism has a self-fulfilling aspect to it; individuals play cooperative strategies because they expect that certain types will also play cooperative strategies with them.

The difference between the strategy selection and technology channels is subtle, particularly in the area of sanctioning. It is plausible to assume that sanctioning norms may exist among co-ethnics but not among non-co-ethnics (or exist more strongly in the former than the latter). Indeed, large anthropological literature on norms of sharing and reciprocity in small communities suggests

that this is the case.¹⁴ As a result, people might play different strategies depending on whether or not they are interacting with members of their own ethnic community. If Alex and Kenneth are co-ethnics, they will cooperate; if they are non-co-ethnics, then they will shirk. Is this evidence of a technology or a strategy selection mechanism? To answer the question, note that all of the mechanisms that we group under the umbrella of “technology” share the characteristic that they facilitate collective action among co-ethnics by making it possible for co-ethnics to do something—work together productively, collect information about potential cooperating partners, or sanction free riders—that non-co-ethnics cannot do, or cannot do as easily. The strategy selection mechanism depends on no such difference. It posits merely that individuals will play different strategies when playing with co-ethnics than they will with non-co-ethnics, even though there may be no disparity across the two kinds of pairings in terms of the technologies that are available to them to promote social cooperation. To illustrate, suppose that we observe that Frederick and Jonas (co-ethnics) are more likely to cooperate on a collective task than Richard and Simon (non-co-ethnics). A “technology” story might be that Frederick and Jonas’s greater cooperative capacity stems from the fact that they are better *able* than Richard and Simon to sanction one another in the event that one of them shirks. A “strategy selection” story, on the other hand, would first stipulate that both pairs are equally able to sanction one another but that only Frederick and Jonas would actually do so. In the strategy selection mechanism, the different outcome derives not from the existence of some cooperation-facilitating tool that co-ethnics uniquely possess. It arises instead from the existence of norms or social institutions that make the available tool (for sanctioning, information gathering, or whatever) more likely to be used. The technology story is that everyone might like to sanction defectors but only co-ethnics can. The strategy selection story is that everyone can sanction defectors but only co-ethnics will.

Although we recognize that these three broad families of mechanisms may interact in complex ways—technologies may be endogenous to preferences and vice versa; even with multiple equilibria, strategy selection may be partly determined by preferences and by beliefs about the preferences of others—we believe them to be exhaustive. In particular, they subsume the hypotheses advanced in the literature for the negative relationship between ethnic diversity and

¹⁴ Colonial-era anthropology underscored the autonomy of rural, tribal societies, which exhibited “a high degree of self-sufficiency at a near-subsistence level” and “a distinctive language, culture, and sense of identity.” (Southall 1970, p. 28). And while modernization gave rise to a multiplicity of new social identities (Mitchell 1956), tribal societies retained a strong hold on newly “detrribalized” urban dwellers (Southall 1970), especially when interacting with a diversity of ethnic groups outside of the workplace in growing peri-urban zones (Epstein 1958).

public goods provision. For example, arguments advanced by Alesina et al (1999) and Alesina and LaFerrara (2005) emphasizing the correspondence between ethnic groups and preferences for particular kinds of public goods fall under our “preferences” channel. Hypotheses proposed by Besley et al (2003), Miguel and Gugerty (2005), Ghosh and Ray (1996), Fafchamps and Minten (2002), and Richman (2006) about the role that networks play in facilitating punishment fall under our “technology” mechanism, as does Okten and Osili’s (2004) hypothesis about the lower transaction costs of co-ethnic interactions. The strategy selection account is frequently invoked to explain the greater collective action we observe in homogeneous than non-homogeneous settings—particularly in situations where multiple equilibria exist and players must coordinate their choices to maximize their payoffs (as, for example, in an iterated Prisoners’ Dilemma, a Battle of the Sexes, 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.¹⁵ The norm among U.S. Senators that they will lend their support to a colleague’s bill in exchange for the promise of reciprocated support at a later date (Mayhew 1975) is an example, as 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 Kenneth and Alex are co-ethnics and we observe that they are able to cooperate to mend the road to their shops when non-co-ethnic pairs in the same situation fail, we will have no way of knowing whether the roots of their successful cooperation lie in a sense of altruism they feel toward one another, a calculation that their collaboration is likely to be particularly easy or fruitful, a knowledge that if the other shirks he can be easily sanctioned, or the existence of a norm that makes not reciprocating a co-ethnic’s cooperative overture, or not punishing a failure to reciprocate,

¹⁵ Evolutionary models of reciprocity include, for example, Gintis (2000) and Bowles and Gintis (2000). A rich literature in economics offers a range of theoretical models advancing mechanisms that drive reciprocal behavior (Rabin 1993, Fehr and Schmidt 1999, Bolton and Ockenfels 2000, for example).

unthinkable. Sorting out whether the pattern we observe is best explained by a preference, 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 great weakness lies in their questionable external validity—that is, in the uncertain ability of researchers to extrapolate from their experimental findings to the larger world whose behavior patterns they seek to explain. However, the ability that they afford researchers to isolate and test the explanatory power of competing explanations that are difficult or impossible to disentangle in real world settings makes them an extremely powerful tool.

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 the analyses were anthropological, focusing on aspects of German culture that might have facilitated the rise of the Nazis and the implementation of their “Final Solution.” Others were sociological, emphasizing 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 others delved into the organizational apparatus of the German state. The best of these studies involved detailed analyses of German history, society and culture based on interviews, primary sources, and careful archival research. Yet perhaps the most compelling explanation for the willing participation of “regular” Germans in the execution of six million Jews came from a study that involved no field research, no interviews, and not a single German source.

To help understand how the Holocaust could have happened, Psychologist Stanley Milgram recruited 40 subjects to participate in an experiment (Milgram 1974). 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) would play 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 actually administered), who, as the voltage increased, (pretended to) shriek in pain and beg for the experiment to be stopped. If the subject expressed a desire to cease

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.” If the subject continued to insist that the experiment should be stopped after all four commands by the experimenter to continue, then the experiment was halted. Milgram’s incredible finding was that more than 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 conscience—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 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 (Axelrod 1984). 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 Prisoners’ 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. The main finding was that a strategy of “tit for tat”—starting off by cooperating and then simply responding in kind to whatever one’s partner does—is optimal. This winning strategy is precisely the strategy we described in our discussion of strategy selection mechanisms. Although it might seem that nothing could be 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.

Our decision to employ the methodology of experimental games stems from a conviction, 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 co-ethnics and non-co-ethnics in laboratory games. We might have gone another route. We might have elected instead to conduct an in-depth 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 have decided to use experimental games.

The experimental approach we employ provides considerable leverage that is not available from alternative methods. The games that the subjects played were designed specifically to test the distinct mechanisms outlined above (we describe them in detail in Chapters Four and Five). Subjects played multiple rounds of each game with randomized matching—sometimes with co-ethnics, sometimes with non-co-ethnics—allowing us to rule out confounding factors. Our strategy for assessing the explanatory power of the distinct mechanisms is to examine how the differences in play among co-ethnics and non-co-ethnics varied across games.

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.¹⁶ This literature has its roots in the pioneering work of Kahneman, 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 try to maximize their earnings, they 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 these initial contributions has been to investigate the

¹⁶ For good reviews, see Roth (1995) and Camerer (2003).

implications of relaxing the rather unrealistic assumptions that underlay 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 behaved differently from college students, who in terms of age, education, and other potential determinants of social behavior, comprise a highly unrepresentative sample of the broader population (Smith 2000; Burks et al 2001).¹⁷ A growing number of studies have attempted to test the robustness of standard experimental findings outside the university lab.¹⁸ The most ambitious attempt in this direction is the recent *Foundations of Human Sociality* project, in which researchers played standard experimental games in fifteen small scale societies around the world (Henrich et al 2004). Other examples include Carpenter, Daniere, and Taskahashi (2004), Greig and Bohnet (2006), and Bahry and Wilson (forthcoming).

A second branch of research has attempted to test not 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 vary with gender (Eckel and Grossman 1996, 1998; Andreoni and Vesterlund 2001), race/ethnicity (Eckel and Grossman 2001; Fershtmann and Gneezy 2001), age (Harbaugh et al 2003), and even beauty (Andreoni and Petrie 2005; Eckel and Wilson 2005). Another set of studies have considered not just “main effects”—i.e., 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*—i.e., 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 (2004), Eckel and Wilson (2005), Ferraro and Cummings (2005), Fershtman and Gneezy (2001), Gil-White (2004), and Petrie (2003).

The third (and newest) branch directly takes up the issue of external validity by explicitly applying the findings derived from experimental games to explain patterns of behavior observed outside of the laboratory. For example, Karlan (2005) exploits information gathered in trust and public goods games to predict patterns of savings, repayment, and default, up to one year later,

¹⁷ For a key critique of the reliance on college sophomores that motivated many of these studies, see Sears (1986).

¹⁸ For a good review, see Cardenas and Carpenter (2004).

among participants in a Peruvian micro-credit organization. Carpenter and Seki (2005) demonstrate that Japanese fishing crews composed of fishermen that exhibit greater degrees of conditional cooperation and are more willing to disapprove of shirking tend to be more productive. Others are now using experimental games to assess changes in levels of community efficacy and trust in response to programs of decentralization and democratic empowerment in Sierra Leone and Liberia.

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, examining instead how subjects behave in a context of substantive interest, the slums of urban Kampala. We share with the second branch a concern for how patterns of play vary with the characteristics of the players and the homogeneous or heterogeneous nature of the interaction. 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 pairing 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 whether players are most appropriately classified as “egoists” or “non-egoists.” 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 Ostrom (2000), 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 (e.g., Bolton and Ockenfels 2000; Bowles and Gintis 2004; Fehr and Schmidt 1999). If players of different types respond in different ways to a given treatment, then an aggregation problem may arise. Average behavior across types may mask systematic features of play taking place within them. To avoid this pitfall, we distinguish in our analyses between two different types of players: those whose behavior is consistent with preferences in keeping with the neoclassical model (who we term “egoists”) and those who exhibit higher levels of general altruism (who 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 full experimental control over these types. These all exist prior to our experiments. What we do have control over, however, is the ways in which 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 insist on a methodology that, though standard in other types of research in the social sciences, is surprisingly uncommon in experimental work: we recruit a pool of subjects that are representative of the underlying population whose behavior patterns in the real world we are 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 it allows us, the researchers, to make inferences from the behavior of our sample to the population from which our subjects are drawn. Matching the subject and underlying populations is especially important 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 slums of Kampala, Uganda. Kampala is a good site to test the impact of diversity on public goods provision. First, the city is not only extremely ethnically diverse but is 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, the devolution of responsibility for social service provision over the past decade from the central government to the financially strapped elected local councils (LCs) has meant that the supply of many local public goods—including security, garbage collection, and the maintenance of storm drains—has become a purely local affair that depends almost entirely on the voluntary contributions of local community members (Golooba-Mutebi 2003; Onyach-Olaa 2003)—a subject to which we shall 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.

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 combined high levels of ethnic diversity with low levels of public goods provision. While this puts us

in a strong position to evaluate the connection between diversity and the underprovision of public goods, it leaves us without a benchmark for each of the various mechanisms that we test. Thus, while we can determine whether or not a given mechanism is playing an important role in linking diversity to collective action failure, we cannot be certain that the findings will apply equally well in less heterogeneous settings. To make such a comparison, we would have had to replicate our analyses in a broader cross-section of communities.

Our study area was comprised of the four adjacent parishes (LC2s) of Mulago I, Mulago II, Mulago III, and Kyebando.¹⁹ We refer to them in this book collectively as Mulago-Kyebando. Together, these parishes contain approximately 70,000 people from which we drew a random sample of 300 subjects.²⁰ Interviews with community leaders and residents in Mulago-Kyebando (reported in Chapters 2 and 6) confirm both the difficulties of providing local public goods in this community and the highly ethnically fragmented nature of their populations.

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 fact renders our results more salient for examining interactions such as participation in relatively anonymous urban settings and in regional or national politics, rather than, for example, in village politics.

A second choice involved whether to examine interactions in which ethnicity is highly salient versus those in which it is 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. As a consequence, our results speak most directly to the ways in which identities condition everyday interaction. We cannot rule out the possibility that individuals might behave very differently in settings in which ethnicity is specifically emphasized as a salient category of social interaction, as in political campaigns, riots, or other forms of political violence.

Of course, we made many other choices that matter for the interpretation of results. For example, we used individuals' decisions about how to allocate money as a proxy for all decisions of

¹⁹ A fuller discussion of Kampala's administrative structure is provided in Chapter 2.

²⁰ We describe our sampling procedures in Chapter 3.

value. We also focus in most cases on horizontal interactions among individuals rather than interactions conditioned by hierarchy or other organizational features. We highlight some of these additional issues throughout our presentation of the games and discussion of the results.

Outline of the Book

The book is organized as follows. Chapter Two 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 Three 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 or other outsiders—categorize the members of their own community. Chapter Three 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 Four, in many respects the core 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, we demonstrate that stories about preferences cannot account for the failures of collective action in urban Kampala. Co-ethnics exhibit no tendency to value more highly the welfare of members of their own group, nor do members of particular ethnic groups seem to care about different sorts of outcomes. We also find only weak evidence that co-ethnics are more effective at working together on joint tasks. Differences in the ability of homogenous and heterogeneous communities to act collectively appear to stem instead from two sources. Co-ethnics possess two key technological advantages: (i) co-ethnic partners are more “readable” in that individuals are better able (and believe they are better able) to identify unobservable characteristics of their in-group members and (ii) they are more “findable” in that they can be more efficiently located through the use of social networks. Successful cooperation among co-ethnics also is shaped

by a strategy selection mechanism: we find strong evidence that co-ethnics cooperate at higher levels in anticipation of threats of sanctions from other members of their group.

The powerful evidence in support of both ethnic technologies and ethnic strategies requires that we push our analysis further in Chapter Five. Is it the case that, even conditioning on how easily they are found and sanctioned, individuals still act in a more reciprocating manner when playing with co-ethnics? If so, this would provide evidence for a distinctly co-ethnic norm. Or is it simply *because* they are more easily found by co-ethnics that individuals reciprocate more when paired with someone from their ethnic group? If so, this might constitute evidence that co-ethnic cooperation arises from the uneven application of a universal norm. While 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 leads us to a more precise answer to the motivating question. It is not simply the case that the co-ethnic advantage is one of greater mutual findability. Rather, our subjects appear to observe a distinctly co-ethnic norm of cooperation. While these specifically co-ethnic norms appear to account for the co-ethnic advantage, they nonetheless coexist alongside a powerful universal norm of cooperation that emerges when the threat of punishment is present for all pairings. This suggests that although reciprocity norms are stronger within groups than across them, it is possible to foster norms that span ethnic groups in highly diverse societies.

Chapter Six draws us out of the experimental setting into the community of Mulago-Kyebando in order to investigate whether the patterns we observe in the games are reflected in behavior observed outside the laboratory. We begin with a deeper exploration of our subject population, examining the stories they tell us about how they understand the games, showing how they link play in the lab to situations they confront in their everyday lives, and looking to see how their game behavior correlates 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 26 local councils, many of whom point to their inability to police non-contributions as the single most important impediment to community mobilization. A careful examination of the structure of political life in our study area illuminates why the absence of cross-group reciprocity norms is so damaging: the LC system remains dominated by the Baganda, an ever-shrinking plurality ethnic group that has not succeeded in forging strong ties with Kampala's new entrants. 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 the concluding chapter.

Table 1.1: Mechanisms Linking Ethnicity to Collective Action Success

I Preference Mechanisms		
I.1	<i>Other Regarding Preferences</i>	Co-ethnics may be more likely to take each other's welfare into account
I.2.	<i>Preferences in Common</i>	Co-ethnics may have correlated preferences over outcomes
I.3	<i>Preferences Over Process</i>	Co-ethnics may prefer the process of working together
II Technology Mechanisms		
II.1	<i>Efficacy</i>	Co-ethnics may be able to function together more efficiently
II.2	<i>Readability</i>	Co-ethnics may be better able to gauge each other's characteristics
II.3	<i>Periodicity</i>	Co-ethnics may engage each other with greater frequency
II.4	<i>Findability</i>	Co-ethnics may be more able to track each other down
III Strategy Selection Mechanisms		
III.1	<i>Reciprocity</i>	Co-ethnics may be more likely to punish each other for failure to cooperate