

Chapter 2

Coethnicity and Trust

JAMES HADYARIMANA, MACARTAN HUMPHREYS,
DANIEL N. POSNER, AND JEREMY M. WEINSTEIN

SCHOLARSHIP ON TRUST emphasizes the beliefs individuals hold about actions that others will take.¹ In such accounts, trust is a belief that the other person will take an action in one's own interest, perhaps in response to a trusting action. It is a belief that the other is trustworthy. But where do these beliefs come from? Why are some people trusted in this way and others not?

We examine one of the many answers that have been offered to this question: people are more likely to trust someone from the same ethnic group. This assumption can be found throughout the literature on ethnicity (see Brewer 1981, Cohen 1969; Fearon and Laitin 1996; Landa 1994; Macharia 1988), and finds empirical support in both experimental studies (see Fershtman and Gneezy 2001; Burns 2003; Petrie 2003; Barr 2004; Karlan 2005; chapter 1, this volume) and survey findings.² For example, the Afrobarometer survey project includes a standard question in which respondents are asked how much they trust others, including "people from your ethnic group" and "people from other ethnic groups." While 50.6 percent of respondents indicate that they trust coethnics "somewhat" or "a lot," only 38.3 percent say the same of noncoethnics.³ This chapter probes the sources of this connection between coethnicity and greater perceived trustworthiness. The goal is not to document that people are more likely to perceive coethnics as more worthy of trust—we follow the literature cited in assuming this is the case—but to contribute to an understanding of why this is so.

We examine three possible explanations.⁴ The first is the *other-regarding preferences rationale*, in which trust stems from the belief that the trustee

cares about the trustor. The second is the *incentives rationale*, in which trust derives from the belief that the trustee is motivated to act in the interests of the trustor. The third is the *competence rationale*, in which trust stems from the belief that the trustee is capable of acting in the interests of the trustor. Each rationale provides a different answer to the question of why a person might believe that a coethnic is more trustworthy than a noncoethnic.

A major impediment to figuring out which of these rationales best accounts for the greater expectations of trustworthiness among coethnics is that, though theoretically distinct, the three mechanisms are often observationally equivalent. If we ascertain, through a survey, for example, that a respondent believes that a coethnic is more trustworthy than a noncoethnic, is it because she believes the coethnic cares more about her than the noncoethnic does? Or is it because she simply thinks that the coethnic has stronger incentives, or is better able, to take an action in her interest? It is impossible to know based solely on reported levels of trust.

Our approach to this inferential problem is to use a series of experiments designed to test each of the rationales independently of one another. Specifically, we compare patterns of play among coethnics and noncoethnics across different experimental games, each designed to isolate a single rationale. When coethnics and noncoethnics play differently in a particular game, we interpret it as evidence for the salience of the rationale that the game was designed to capture. In the case of some games, we extract direct statements about beliefs. In others, we infer beliefs from behavior under the assumption that players' beliefs are consistent, on average, with how others behave.

Although this chapter deals with the general question of why people believe coethnics to be more trustworthy, our empirical analysis is grounded in a specific multiethnic setting—that of the urban neighborhoods of Mulago and Kyebando in Kampala, Uganda.⁵ Uganda is a good place to study why ethnicity affects beliefs about trustworthiness. The Afrobarometer findings cited earlier regarding levels of trust for coethnics and noncoethnics suggest a trust gap between in-group and out-group interactions in Africa of about 12 percentage points. In Uganda, the gap is nearly double that size. Whereas 60.9 percent of Ugandans in the round 3 survey reported trusting people from their own ethnic group "somewhat" or "a lot," just 39.4 percent reported equal levels of trust for people from other ethnic communities.⁶

The particular neighborhoods in Uganda we study offer a good laboratory for examining interethnic interactions and the beliefs that shape them. Mulago-Kyebando has been the site of heavy in-migration over the past two decades, driven in part by strong ties to sending areas and the availability of cheap accommodation (and even land) in the area. In a pre-survey of 594 randomly selected individuals in Kawempe, the broader division of the city in which Mulago-Kyebando is located, more than 50 percent of

respondents reported that they had lived in their current neighborhood less than five years. Nearly 80 percent described their community as composed largely of people born outside of Kampala who had moved to the capital. The consequence of this steady stream of migrants has been a dramatic increase in the level of ethnic heterogeneity. Whereas Mulago-Kyeabando were once dominated demographically by the Baganda, the largest ethnic group in Uganda and the group whose historical kingdom is centered in Kampala, Baganda now comprise only a little more than 50 percent of the local population. The ethnic fractionalization value for Uganda as a whole is approximately 0.9, though the figure for Mulago-Kyeabando is somewhat lower given the still dominant place of the Baganda.⁷

Ethnic differences in Mulago-Kyeabando are not only present but also highly salient in everyday social interactions. Many ethnic groups have formed homogeneous associations for accumulating savings, providing access to credit, and meeting other practical needs. In focus groups and interviews with local council leaders, participants frequently used ethnic or regional labels to refer to factions within the community: "the Bafumbira do this," we would commonly be told, and "people from the Northeast do that." Failures of community-level collective action were often explained in terms of the inability of members of different ethnic groups to work together. Although Mulago-Kyeabando's extreme ethnic heterogeneity has never led to intergroup violence, it is nevertheless a central feature of daily life. This is important insofar as it permits us to rule out the possibility that the nonfindings we report stem from the lack of salience of ethnic divisions *per se* within the community.

Our empirical strategy rests on a sequential examination of game play between coethnics and noncoethnics. In the sections that follow, we examine evidence for each of the three rationales described above. First we investigate whether other-regarding preferences are, in fact, structured along ethnic lines. Our strategy is to use a dictator game with an anonymous offerer to examine whether individuals take account of the welfare of in-group members more than out-group members. Such behavior would provide a basis for a belief that coethnics are more likely to have one's interests at heart. Next, we explore whether individuals are incentivized to act in the interests of coethnics. We first draw on survey data to examine whether preferences over outcomes are aligned in ways that would provide a basis to expect that coethnics have incentives to support initiatives consistent with one's own interests. Then we use the results of a dictator game with a *100%*-anonymous offerer to examine whether the prospect of social sanctions affects the actions of offerers or the beliefs of receivers about those offers. The final section addresses the third rationale: that individuals have greater trust in the capacity of coethnics to solve problems with them. We explore this possibility by examining the choices players make about partners in a new game, which we call

the "lockbox game," that is designed to incentivize players to accomplish a joint task. By examining partner selection in this game, we can gain insight into subjects' beliefs about the competence of different potential partners.

Experimental Framework

We took a set of experimental games, typically played in laboratory environments, to the field in Kampala, Uganda.⁸ We began by randomly sampling 300 individuals from Mulago-Kyeabando.⁹ Like the underlying population from which they were recruited, our subjects were extremely diverse. Forty-four percent were Baganda, with declining shares of Banyankole (9 percent), Bafumbira (7 percent), Batoro, Banyarwanda, Bakiga (roughly 5 percent each), and other groups. They were also largely uneducated: 23 percent had not completed primary school, and more than 70 percent had completed at most only some secondary school. Just 17 percent reported holding formal sector jobs; roughly half were either unemployed or working in the informal sector. Fewer than 30 percent owned their dwelling. Thirty-six percent reported not having electricity and 87 percent reported not having piped water in their homes. Ownership of assets such as refrigerators (28 percent), electric irons (53 percent), televisions (51 percent), bicycles (13 percent), and cars or trucks (10 percent) was also very low. Our subjects were also overwhelmingly migrants or children of migrants: just 22 percent reported having been born in Kampala.

Before playing the experimental games, we recorded a series of five digital images of each subject (a headshot and four brief video clips) each providing a different level of information that an observer might use to ascertain the subject's ethnic background. In the analyses presented here, we treat all five levels of information equally and distinguish only between situations in which players have no information about the identities of their partners and situations for which they have some information.

Public Information Box and General Set-Up

All but one of the games we describe were played using a computer interface.¹⁰ Players could make inferences about the ethnic backgrounds of the other players from pictures or video clips made available to them in what we call the public information box (PIB). The key attribute of the PIB is that, as its name suggests, the information provided about the players in the game was provided publicly. Before each round of each game, all the players in the round were shown the same PIB containing images of all of the players in that round—including themselves—with the images of the players ordered in the same way. Underneath the PIB, each player

Figure 2.1 Public Information Box with Nonanonymous Offerer



Source: Author-provided screenshots.


Notes: Player 2, the offerer, is "seen" by all players. Images are for illustration purposes only and are not the images of actual subjects.

saw a note indicating which player number he or she was for that round. Figure 2.1 presents a sample screenshot of a PIB.

Beyond providing the information that players could use to ascertain the ethnicity of the other players, described in more detail in the next section, the PIB played three roles in the computer-based experiments. First, it made the interaction more realistic by increasing the credibility of the existence of the other players (Bohnet and Frey 1999; Eckel and Wilson 2006). Second, it provided common information: each player was provided not just with information about who the other players were but also about what the other players knew about them, and that the other players knew what they knew about the others, and so on. Third, its design allowed us to manipulate the anonymity of the players in the game. Compare, for example, figures 2.1 and 2.2. In figure 2.1, the middle player's picture is shown to the other players. That player will therefore play the game knowing that the other players can see who he is. In figure 2.2, however, the middle player's picture is not shown. He still has information about the other players and knows that they have no information about him. He will therefore play the game knowing that he is doing so anonymously. Exploiting this manipulation turns out to be extremely valuable for distinguishing behavior motivated by other-regarding preferences from behavior motivated by incentives.

Although subjects were shown images of the other players in the game, the games were designed to study interactions among strangers. Therefore, after viewing the PIB, subjects were asked to report whether they knew either of the other players whose images they had just been shown. About 6 percent of all rounds involved subjects who said they knew one or more of the other players. All results reported in this chapter are robust to the exclusion of such rounds. Each subject played all of the games multiple

Figure 2.2 Public Information Box with Anonymous Offerer



Source: Author-provided screenshots.

Notes: Player 2, the offerer, is anonymous. Images are for illustration purposes only and are not the images of our subjects.

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

Coding Coethnicity

The empirical strategy we adopt in this chapter depends on our ability to distinguish interactions among coethnics from interactions among non-coethnics.¹¹ A simple rule—alogous to the practice employed in most studies of cross-race, -gender, and -ethnic interactions in the experimental literature—is to code as coethnics any pair of players that identified themselves as belonging to the same ethnic category in our pre-experiment questionnaire. Using this rule, we generate what we term a benchmark measure of coethnicity. However, this measure runs into the problem that the way a person self-identifies may not correspond with the way he or she is perceived. Thus, if two subjects who identified themselves in the pre-experiment questionnaire as members of group X were paired in a game, then, under the benchmark measure, the pairing would be coded as coethnic. But if each subject believed that the other was not really a member of group X (whether because of imperfect information or different notions of group boundaries), they would each behave as if they were playing with a noncoethnic, and our inferences about the impact of shared ethnicity on their actions would be wrong.

To overcome this problem, we used a relatively straightforward exercise to generate a subjective measure of coethnicity based on how our subjects

perceived the ethnic backgrounds of the players with whom they were interacting. After all the games had been played, and the danger of priming subjects to ethnicity had passed, we showed our subjects a series of images of other subjects (in most cases players they had been randomly matched with earlier) and invited them to guess their ethnic identities. To motivate the guessers, correct guesses (corresponding to how the person had self-identified in the pre-experiment questionnaire) were rewarded with a small payment. To ensure that everyone had the same prior beliefs about the distribution of ethnic groups in the sample of images, we told the subjects that the ethnic demography of the sample population matched that of Mulago-Kyeabando, and we read aloud a breakdown of shares of the major ethnic groups in Mulago-Kyeabando based on 2001 census figures.

To generate estimates of subjective coethnicity, we collected data on a total of 15,265 guesses by 274 subjects. Overall, we found that individuals were able to correctly identify the ethnic backgrounds of others only about 50 percent of the time—a result that underscores the inferential problem we would have faced had we limited ourselves to the benchmark measure. We used the results of the identification exercise to generate, for every information level, an estimate of the likelihood that an individual of group A believes that an individual of group B is a coethnic for each possible information level.¹² The resulting measure (ranging from 0 to 1) provides a measure of subjective coethnicity. For all games that use the PIB, we report results using both the benchmark and the subjective measures of coethnicity.

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

Other-Regarding Preferences Rationale

The first rationale we explore is other-regarding preferences. High levels of other-regarding preferences may produce expectations of trustworthiness by making a potential trustee believe that the trustee has the trustee's interests at heart. If other-regarding preferences are stronger (or believed to be so) between coethnics than between noncoethnics, then this could provide a basis for beliefs about the greater trustworthiness of coethnics.¹³

To test whether such a basis exists, we had subjects play a version of the standard dictator game in which the offerer is anonymous and the

receivers' identities are known (Kahneman, Knetsch, and Thaler 1986). We then attribute observed difference in patterns of play between situations in which the offerer and receiver are coethnics and in which they are not to differential levels of altruism toward ingroup and outgroup members. Because our focus is on detecting ethnic discrimination, the dictator game our subjects played differs from more commonly used versions in two ways. First, we play in trios rather than in pairs. This permits us to evaluate situations where offerers must decide between allocating a sum to a coethnic and a noncoethnic. Second, the sum to be allocated is "lumpy." Instead of being given a continuous endowment to divide among the players, offerers were given two coins and told that no player (including themselves) was permitted to receive more than one coin. This forced the offerer to discriminate between the two receivers, or, if the offerer chose to give away both coins, against him or herself. Hence we refer to this as the discrimination game.¹⁴

Each round began with subjects, who in this game only play the role of offerer, seated in front of a laptop computer. The screen showed a PIB containing images of the two other players (the receivers) and a dummy for the subject, as in figure 2.2. In front of the computer were three ballot boxes, each directly below one of the pictures in the PIB. Subjects were given two 500-USh coins (about US\$0.60, approximately the per capita daily income in Uganda) and asked to divide this sum as described. Subjects were told to put the amount they wanted to keep directly into their pockets and to put the amounts they wanted to allocate to each of the other players into envelopes and deposit them in the ballot boxes below the players' pictures. Subjects were told that the envelopes would be delivered to their intended recipients at a later date, which they were. An enumerator manipulated the computer to show the PIB for the given round and handed the subject the money, but then stepped away and waited behind a screen while the subject completed his or her allocation. When the subject was finished making the allocation, he or she signaled the enumerator, who returned from behind the screen and set up play for the next round.¹⁵

Each subject played multiple rounds (an average of 2.7) of the game. In all, we have data from 782 rounds (1,564 individual choices). The modal strategy (played in 73 percent of rounds) was to keep one 500-USh coin and to allocate the other coin to another player. Nonetheless, in 23 percent of the rounds, subjects allocated both coins to the other players. These offers can be compared with a baseline strategy of random allocation, under which subjects would keep one coin two-thirds of the time.

We now turn to the question of whether subjects displayed different degrees of other-regardingness toward coethnics and noncoethnics. We restrict our analysis to rounds in which a player was or believed he or she was playing with one coethnic partner and one noncoethnic partner, and in which he or she also elected to discriminate.¹⁶ Row 1 of table 2.1 reports

Table 2.1 Discrimination Rates in Dictator Game

	Benchmark Coethnicity	Benchmark Coregion	Subjective Coethnicity	Subjective Coregion
Effect of in-group membership when offerer is anonymous	-0.03 (0.10)	0.06 (0.08)	0.08 (0.15)	0.08 (0.14)
Effect of in-group membership when offerer is seen	0.12* (0.07)	0.14*** (0.06)	0.28** (0.13)	0.24*** (0.11)
Difference	0.15 (0.12)	0.08 (0.09)	0.19 (0.21)	0.16 (0.18)
Observations	628	962	310	432

Source: Authors' calculations.

Notes: Standard errors in parentheses, calculated using weighted OLS regression with weights to account for different assignment probabilities across groups. Disturbance terms are clustered for each player across all of his or her games. Cells in first two rows report average treatment effects on the treated (ATT) using exact matching to average over treatment effects obtained for offerers for each ethnic or regional group. Cells in third row report interaction terms from an OLS regression with weights to take account of different assignment probabilities for different groups.

*significant at 10 percent; **significant at 5 percent; ***significant at 1 percent.

the average effect of coethnicity on the likelihood that players discriminated in favor of coethnics when they played anonymously.¹⁷ We find no evidence of any coethnic effect. Contrary to the expectations of the other-regarding preferences rationale, there is no statistically significant difference in the likelihood that players discriminate in favor of coethnics.

These results do not directly establish what individuals believe about the other-regarding preferences of coethnic and noncoethnic partners. To establish beliefs directly, we would like to have been able to ask our receivers to predict the offers made by coethnic and noncoethnic offerers. However, playing such a back end to this game would have required either asking respondents hypothetical questions or showing receivers pictures of the offerers, who had been told that they were making their allocations anonymously, a violation of the experimental norm of no deception. Nonetheless, the results suggest that other-regarding preferences are not structured along ethnic lines, and any belief to the contrary would be sharply at odds with this observed behavior.

Incentives Rationale

Next we turn to the incentives rationale. An individual may have an incentive to be trustworthy even if he or she does not have the interests of the trustor at heart for a variety of reasons. One possibility is that the

preferences of the two agents over outcomes are aligned—the trustor in this case knows that the trustee will take the right action because that action is also in the interest of the trustee. A second is social norms. In this case, the trustor believes that the trustee will take the right action because he or she knows that the trustee will be violating a norm if he or she does not, a violation that could come at a cost to the violator. In either case, the existence of incentives for the trustee to be trustworthy, if known to the trustor, would lead the trustor to have higher expectations of the trustee's trustworthiness.

Preferences over Outcomes

In a trivial sense, the other-regarding preferences rationale we discussed in the last section may be thought of as an incentive rationale. But other preference rationales may be less trivially linked to incentives. In the classic example provided by Adam Smith, I can rely on a serviceman to deliver goods not because he likes me but because he, like me, benefits from the trade (Smith 1776/1994). What matters is not preferences over the welfare of others but that the welfare of each person is correlated across outcomes. If the preferences of coethnics are relatively more aligned in this way, this could lead to higher levels of trust among coethnics and greater expectations that coethnics will be trustworthy.

To probe the plausibility of this argument for Mulago-Kyeabando, we use simple survey techniques to test whether preferences over public goods outcomes do in fact correlate with ethnic group membership. Insofar as they do, individuals have incentives to trust each other in the joint production of these goods. We examine two types of survey questions. First, to what types of public goods do individuals attach the highest priority (security, drainage maintenance, or garbage collection)? Second, how should these goods be provided, for example, should private or public resources be used? We focus on these issues—how projects are prioritized and provided—because they were identified by community members in interviews and focus group discussions as being among the most salient concerns in Mulago-Kyeabando.

Table 2.2 reports the results from *F*-tests derived from a simple analysis of variance (ANOVA) and from the nonparametric Kruskal-Wallis test, which is more appropriate when the outcome measure is categorical. We cannot reject the null hypothesis of no systematic variation across ethnic or regional groups for any of the six questions. More detailed analysis further suggests that clustering on any given group is minimal and that what clustering exists can be accounted for by locality-fixed effects. In short, we do not find empirical support for the argument that members of ethnic groups have correlated preferences over political outcomes in Mulago-Kyeabando.

Table 2.2 Within-Group Clustering of Policy Preferences

	Ethnicity		Region	
	ANOVA F-test (<i>p</i>)	Kruskal- Wallis χ^2 test (with Ties) (<i>p</i>)	ANOVA F-test (<i>p</i>)	Kruskal- Wallis χ^2 test (with Ties) (<i>p</i>)
I				
II				
First priority for public goods provision	0.25	0.25	0.91	0.91
Drainage	0.70	0.69	0.39	0.39
Garbage collection	0.67	0.66	0.25	0.25
Security				
III				
How public goods are to be provided	0.88	0.66	0.51	0.72
Preference for fee-based garbage collection, over free but lower quality provision				
It's better not to have to pay anything or to volunteer for patrols, even if that means security is low	0.35	0.23	0.73	0.80
It's better to have well-maintained drainage channels, even if we have to make contributions of money or labor	0.41	0.22	0.44	0.48

Source: Authors' calculations.

Notes: Analysis is limited to ten largest groups.

Preferences can be correlated across groups in many ways, of course. We have examined just one and cannot rule out such correlations in general. The key point in assessing these results is that the questions we examine were identified as being locally salient in communities. The main implication is that for these salient issues we cannot find evidence that the failure to coordinate across groups can be attributed to groups wanting different outcomes.

Social Norms

A second source of beliefs about trustworthiness is social norms and the social sanctioning that may take place if those norms are violated. If

individuals know that a norm exists that one should act in a trustworthy way, and that one is subject to sanctioning in the event that they fail to abide by this norm, they are more likely to be trustworthy and increase their expectations of trustworthiness in others who adhere to this norm. In principle, such norms could exist within but not across ethnic groups: if coethnics expect that cooperation with coethnics will be reciprocated under threat of sanctioning but that cooperation with non-coethnics will not, expectations of trustworthiness will be higher among coethnics.

To test for such a norm—and for its differential strength among coethnics and noncoethnics—we explored what happened to the offerers' behavior in the dictator game when the offerers were no longer anonymous: the PIB shifts from the one depicted in figure 2.2 to the one depicted in figure 2.1. In such a setting, the offerer's behavior can be interpreted as a function of his or her relative other-regardingness toward the receiver, as it was in the anonymous version of the game, and his or her concern about being seen to violate a social norm requiring cooperation. This latter concern is irrelevant when the offerer is anonymous, but becomes potentially important when the offerer is seen and in principle can be sanctioned if his or her actions breach the norm against failing to contribute. To evaluate whether such a norm exists, we can therefore compare patterns of play in dictator games where the offerer is and is not seen. To evaluate whether the norm exists more strongly for coethnics than noncoethnics—the central claim of the incentives rationale that we seek to test here—we can compare the results of anonymous and nonanonymous dictator games played among coethnics and noncoethnics.¹⁸

Each subject played approximately four rounds of the nonanonymous version of the dictator game, yielding a total of 1,226 rounds (2,452 choices). Aggregate patterns of play were similar to those found in the dictator game with the anonymous offerer. Subjects kept one 500-US\$ coin and allocated the other to another player 70 percent of the time. They gave both coins away 24 percent of the time.

The effects of coethnicity on offers can be seen in the second row of table 2.1. Again, for our analysis, we focus only on cases in which players actually discriminate (that is, choose to give the coin to one player and not another) when they are or believe they are facing one coethnic and one noncoethnic. In nonanonymous games, we find systematic differences in the way subjects played with coethnics and noncoethnics. Using the benchmark measure, coethnicity increases the likelihood that a partner will be favored by 12 percentage points; using our subjective coding, a zero to one change in subjective coethnicity is associated with a 28 percentage point increase in the likelihood of favoritism. Results for shared group membership defined by common regional background are similarly large, 14 and 24 percentage points respectively.

The contrast between these findings and those reported in the first row of table 2.1 is striking. When they are making their offers anonymously, subjects give no more to coethnics than to noncoethnics. But when they know that they can be seen, they give significantly more. Like the findings presented in chapter 1 of this volume, these results provide evidence that trust-facilitating norms exist within ethnic groups (see also Yamagishi and Mifune 2008). Although we find a significant effect of coethnicity in the nonanonymous games and not in the anonymous games, we cannot distinguish the difference between these marginal effects from zero with confidence (see the second to last row of table 2.1).¹⁹

Beliefs

These results support the claim that individuals take the interests of others into account in response to socially generated incentives. They do not, however, establish that individuals believe that others respond this way. They provide only a basis for drawing a conclusion about beliefs of trustworthiness if we assume consistency of beliefs—that is, that people’s beliefs about trustworthiness follow directly from behavior that is in fact trustworthy. We turn now to a more direct investigation of whether individuals believe that others will act in a trustworthy way. Specifically, we test whether, in nonanonymous dictator games, players expect to receive more from coethnics than from noncoethnics. Because we condition on cases in which play is nonanonymous, this test does not permit us to determine whether the source of these beliefs lies in presumptions about the trustee’s other-regardingness toward the trustee or in incentives that increase the likelihood that the trustee will behave in a trustworthy way. However, if players do expect to receive more from coethnics, then our claim that the threat of punishment is credible when such expectations are not met—the rationale for which we have found empirical support—stands on stronger grounds.

To assess player’s expectations, we examine data gathered during the back end of the nonanonymous dictator game. Subjects played the back end after all dictator games had been completed but, of course, before they learned what they received from each offerer. Capturing an accurate reading of people’s expectations is a challenge. We did not want subjects to provide their guesses directly to our enumerators for fear that that might bias our results. Instead, we gave each receiver the exact endowment of cash that had been provided to offerers in the front end of the game (two 500-US\$ coins) and asked them to indicate how they believed the offerer had allocated the coins among the three players. To do this, we showed them PIBs for each round of each game in which they had been a receiver and, for each one, asked them to place coins in envelopes and put them in the boxes corresponding to the pictures of each player on the

Table 2.3 Receivers’ Expectations in Dictator Game

Marginal Effect	Dictator Game (Back End) (<i>p</i>)	<i>n</i>
Benchmark coethnicity	.12 (.11)	385
Benchmark coregion	.07 (.04)*	684
Subjective coethnicity	.21 (.15)	385
Subjective coregion	.15 (.06)**	675

Source: Authors’ calculations.

Notes: Standard errors in parentheses, calculated using weighted OLS regression with weights to account for different assignment probabilities across groups. Disturbance terms are clustered for each player across all of his or her games. Cells report average treatment effects on treated (ATT) using exact matching to average over treatment effects obtained for offerers for each ethnic or regional group.

computer screen exactly as they thought the offerer had allocated the coins.²⁰ As incentives, players were given 500 US\$ if their guesses about the offerer’s allocation were correct.

In general, our subjects were wishful thinkers. Players guessed that the offerer would keep one coin and give them the other in about half the games, but that the second coin would go to the other receiver only 28 percent of the time. In 17 percent of cases, receivers believed that the offerer would give away both coins, and that the offerer would keep them (against the rules of the game) in about 4 percent.

In table 2.3, we explore whether players conditioned their expectations on the match between the offerer’s ethnic background and their own. We limit the analysis to games in which guessers expected offerers to keep one coin and give the other away, and in which one of the receivers was a member of the offerer’s ethnic group and the other was not. By conditioning on players who always keep one coin, no offerer is more or less generous than another; the question is simply who they chose to benefit. The results suggest that individuals expect coethnic offerers to discriminate in their favor. This effect is especially strong with respect to region. Coregionists are between 7 and 15 percentage points more likely to expect to be benefitted than their counterparts; the substantive magnitudes for coethnic pairings are even larger, but these results do not obtain statistical significance. These findings offer some additional support to the results reported earlier regarding the role of sanctioning and social norms in generating trustworthy behavior: in nonanonymous settings, players favor their coethnics and their coethnics expect them to do so.²¹

Competence Rationale

A third rationale for why some people will be believed to be more trustworthy than others derives from beliefs about their ability to deliver what has been promised. As the myriad stereotypes about the natural abilities of members of particular groups to perform certain tasks attest (that is, Italians are good at making shoes, Nubians make good soldiers, and the like), perceived competence in particular domains is frequently associated with different ethnic groups. However, such stereotyping would not produce higher levels of perceived trustworthiness among coethnics per se. Only if competence is or is perceived to be relational—that is, that individuals believe that coethnics, though perhaps no more competent overall, are better able to work together with them—will beliefs about competence provide a foundation for differential expectations of trustworthiness among coethnics and noncoethnics.

To test the possibility that our subjects had differential beliefs about the abilities of others (and specifically the ability of others to work effectively with them), we created a game in which players had to select from among a set of partners to perform a joint task. Because rewards accrued only to teams that successfully completed the task, players had an incentive to select partners with whom they believed they would be most likely to succeed.

The game we examine here—the lockbox game—involves two people working together to open a combination lock on a box containing cash. Pairs that successfully open the box share the money. Partners were matched by randomly dividing six subjects in a given session into two groups of three: one group was designated to play the role of Player 1, the other to play the role of Player 2. Player 1 is taught how to open the combination lock; he or she must then provide oral instructions to Player 2, who actually manipulates the lock.

After receiving instructions about how the game would be played, one subject was selected at random from the Player 1 pool, shown pictures of the three subjects in the Player 2 pool and asked to select one to be his or her partner in the game. A second subject was then selected at random from the Player 1 pool, shown pictures of the two remaining subjects in the Player 2 pool, and asked to select one to be his or her partner. The final subject in the Player 1 pool had no choice and was simply assigned to play the game with the last remaining subject in the Player 2 pool. In this way it was randomly determined whether a given player faced a choice between one, two, or three potential partners.

Each subject played the game only once, so only half of our subjects played in the position of Player 1. Thirty-one subjects could choose a partner from among three partners, fifty-one from among just two, and sixty-six had no choice. To the extent that they could identify coethnics from looking at their pictures, did subjects select them as partners?

Table 2.4 Partner Selection in Lockbox Game

Case	Share of Players Selecting Ingroup Members		Expected Share Given Random Selection
	Ethnicity (N)	Region (N)	
One coethnic, one noncoethnic	0.55 (11)	0.53 (19)	0.50
One coethnic, two noncoethnics	0.39 (8)	0.33 (12)	0.33
Two coethnics, one noncoethnic	0.6 (5)	0.67 (6)	0.67

Source: Authors' calculations.

The basic results are as follows. First, there is no evidence that competence in this task is relational. Coethnic pairs were no more likely than noncoethnic pairs to succeed in the game: success rates were 0.63 for coethnics ($n = 32$) and 0.63 for noncoethnics ($n = 114$). When group membership is defined in terms of region, success rates are even lower for coethnics, 0.59 ($n = 51$) to 0.61 ($n = 83$) for noncoethnics (difference not significant).²² This is true both when players had a choice in partner selection and when they did not.

Given this pattern, it is perhaps not surprising that players appear not to have believed that coethnics would be better able to work with them. Table 2.4 presents data on partner selection for those cases in which a player could select a partner from among a pool of coethnics and noncoethnics. Although the n is low for all these cases, we see that player selections are as close as possible to random given integer constraints. The same pattern also emerges from a more general analysis (not shown) that uses an alternative-specific multinomial probit model to check for coethnicity effects. This model is appropriate for settings in which an individual chooses one option from a set of alternatives (with the number of options possibly varying) for which each option has distinct characteristics. In this analysis, we looked to see whether any ethnic cues were used. It is possible, and consistent, for example, with findings of Chain Fershtman and Uri Gneezy (2001), that individuals find it optimal to work with some particular ethnic groups though not necessarily their own. Again, we find no evidence to support this view and fail to reject the null that the selection is independent of the ethnic composition of the pool of potential partners.

One possible challenge to our strategy is that players may select partners in part because, independent of success rates, they simply do or do

not prefer working alongside individuals of a given group. If players had preferences to work alongside coethnics, this would bias the results toward finding coethnic selection even if there were no competence advantage of coethnic pairings. That we do not find such selection strengthens our results, to the extent that we expect procedural preferences to be positively correlated with coethnicity.

As with our discussion of correlated preferences over outcomes, there are of course many types of joint tasks we could consider and multiple domains in which differential competence might matter. We cannot infer from these results that there are no competence concerns in other domains. Indeed, in another study, we report more positive evidence for relational effects for a different task for which we do not allow partner selection (Habyarimana et al. 2009). In this sense, our results should be interpreted as absence of evidence for, rather than evidence of the absence of, a competence rationale.

Conclusion

One problem with many approaches to the question of why some people are believed to be more trustworthy than others is that they do not distinguish among competing accounts for the observed outcome. Suppose we observe an individual handing over a monetary contribution to a political campaign in the belief that this money will not be mispent. What is the source of this person's trusting belief? One possibility is that the individual believes that the trustee cares about the welfare of the trustor and so is motivated to act in her interests. Another is that the trustee is motivated to act in the trustor's interests, perhaps because she seeks the same outcomes or perhaps because she expects to be sanctioned if she misuses the funds. A third is that the person believes that the trustee has the skills to spend the campaign funds successfully. These rationales can account for absolute levels of trust but also for differences in trust across groups.

Parsing these explanations is of substantive importance, but is difficult using standard methods. We have used experimental techniques to test these explanations. Our focus was on why people believe that coethnics are more trustworthy than noncoethnics. Our results suggest that, at least in Mulago-Kyebando, the reason is the expectation that, owing to norms of reciprocity that bind more strongly in within-group than in cross-group interactions, coethnics have greater incentives than noncoethnics to respond to trusting overtures in a trustworthy way. The competing rationales find weaker empirical support.

We end with a note of caution and a note of optimism. A skeptical reader might question the generality of our conclusions. The salience of each of the rationales we study, a critic might argue, might depend on the

context of the trusting decision. We concur. Other regarding preferences may plausibly be more important in settings where ethnic differences have been a source of violence. A competence rationale may underlie trusting actions in settings where, unlike in Mulago-Kyebando, no *lingua franca* exists or where success in specific tasks depends more directly on shared group attributes. Experimental protocols that involve explicitly priming subjects to ethnicity (which our protocols do not) may also generate different conclusions about the relative importance of the three rationales.

It is possible, however, that incentives explanations may be of overriding importance for generating trust in a broad class of settings—at this stage, we simply do not have enough evidence to know. To the extent that they are, a hopeful implication is that the distrust that sometimes characterizes cross-ethnic interactions need not be insurmountable. It is difficult to make people care about the welfare of others if they do not do so already. Expectations that collaboration with certain types of people are also not likely to be successful are difficult to change, particularly if they are well founded. But incentives may be more malleable. Formal institutions that penalize people who respond uncooperatively to trusting behavior might generate universal norms about trustworthiness that can displace the ethnic ones that may otherwise predominate.

The authors thank Claire Adida, Bernd Beber, Pepine Bulambo, Deo Byabagambi, Elizabeth Carlson, Chris Crabbe, Kenneth Ekode, Nathan Falck, Alex Kunobwa, Sylvester Mubiru, Douglas Musonga, Winfred Nabulo, Ruth Nagawa, Susan Najjuuko, Brenda Nakkazi, Harriet Nambi, Winfred Naziwa, Livingstone Ntensibe, Alex Odwong, Alexandra Scacco, Geoff Sentongo, Simon Ssenyimba, Elizabeth Suubi, Alex Tindyebwa, Sheila Watuwa, and Daniel Young for their extraordinary research assistance; the Harvard Academy for International and Area Studies, the Russell Sage Foundation, and the Harry Frank Guggenheim Foundation for financial support; and Karen Cook, Jean Ensminger, Henry Farrell, Russell Hardin, Margaret Levi, Susan Stokes, James Walker, and Rick Wilson for their helpful comments on this chapter.

Notes

1. We recognize that in some accounts trust is not a belief: it is possible, for example, to interpret trust as pertaining to the taking of trusting actions. For our purposes here, however, we equate trusting with a belief of trustworthiness.
2. Iris Bohnet and Fiona Greig are an exception (2006). They do not find evidence that coethnicity matters in a trust game played with subjects from a slum in Nairobi, although their sample may exhibit too little ethnic variation to identify an effect.

3. Based on an analysis of pooled data from round 3 of the Afrobarometer, collected in 2005 and 2006. The sample includes 25,397 respondents from eighteen African countries.
4. These explanations parallel those of Margaret Levi and Laura Stoker (2000). Levi and Stoker also identify a fourth rationale, the morality rationale, which hinges on a belief that the trustee adheres to moral values that emphasize promise keeping. We do not have the data to explore this rationale, and we leave it aside in this chapter. We also cannot distinguish the social norms-sanctioning explanation we discuss from an explanation in which trusting beliefs are generated by strong reciprocity within but not between ethnic groups (Gintis 2000).
5. Our study area is comprised of four adjacent parishes (LC2s) in the poorest of Kampala's five divisions, Kawempe: Mulago I, Mulago II, Mulago III, and Kyebando. We refer to them collectively in the text as Mulago-Kyebando.
6. Figures are based on an analysis of round 3 Afrobarometer data for Uganda only ($N = 2,400$).
7. The ethnic fractionalization index measures the likelihood that two people selected at random will be from different ethnic groups. Communities with values of 0 are completely homogeneous; communities with values approaching 1 are extremely heterogeneous (for further details of the demography of our study site, see Habyarimana et al. 2009).
8. The full protocols for all the experiments described in this chapter are available from the authors on request.
9. Simple random sampling was used within local neighborhoods (LC1s). The number of subjects for each LC1, however, was set using targets that diverged modestly from proportionate-to-size to oversample the second and third largest ethnic groups (for detail, see Habyarimana et al. 2009). More than 75 percent of those we contacted agreed to participate in the study. Of those who chose to enter the study, more than 95 percent attended all sessions.
10. The lockbox game is played live. However, the partner selection process we study in that game does involve showing players pictures of other subjects, which is similar to the computer interface we describe here.
11. Given the large number of different ethnic groups in our sample, a treatment of dyadic pairings of groups would have been too complex, so we limit our analysis to the more general difference between coethnic and noncoethnic interactions.
12. Note that by using this rule, an individual i is coded as a subjective coethnic of individual j if individual j believes that i would code herself in the same group as j codes herself. A stricter definition would require that j codes i in the same group as j codes herself under j 's own (rather than under i 's) classification criteria. Our requirement for a single criterion of "correct" identification to use as a basis for allocating rewards to players precluded us from generating this more precise measure of subjective coethnicity.

13. A variant on this mechanism, which some social scientists refer to as quasi-magical thinking, depends on the other-regardingness of the truster toward the trustee. That is, if the truster has the trustee's interests at heart, he or she assumes that the reverse is true as well. By contrast, the mechanism we focus on here depends only on the other-regardingness that the trustee has for the truster, and on the truster's beliefs about the trustee's other-regardingness.
14. We also played a nondiscrimination version of the dictator game with ten coins instead of two. The results, reported elsewhere, are very similar to those described here (see Habyarimana et al. 2007).
15. We instituted various checks to ensure that our subjects understood the games they were playing. Most important, before beginning play, subjects were tested on their comprehension of the rules of the game and the set of strategies available to them. Subjects who failed this test were given additional instruction until they could explain the game on their own. In addition, we organized a back-translation of the games in which an educated Ugandan with no connection to the project met with a group of our subjects and tried to elicit from them enough information about the various games they were playing that he could describe the details of the games back to the experimenters. The success of this back-translation exercise gave us confidence that our subjects understood the underlying behaviors that each game sought to assess.
16. In the subjective coethnicity analysis, a player was coded as believing that he or she was facing one coethnic partner and one noncoethnic partner if the difference in his or her estimated beliefs that each of the two partners was a coethnic exceeded .5. Including the set of games in which there is no discrimination weakens all results marginally but does not affect the conclusions of this chapter.
17. As with all coethnic effects reported in this chapter, average treatment effects are estimated by using exact matching to average over the treatment effects obtained for offerers for each ethnic or regional group. Averaging over differences between treated and untreated subjects within groups ensures that the results are not confounded by main effects that could result from treatment assignment probabilities varying across groups. The results are robust to the inclusion of a battery of controls for age, education, income, gender, and other characteristics (which, by design, are uncorrelated with the experimental treatment).
- Because the game is played with two partners for each offerer, we stack the data and code a dependent variable that captures whether a given receiver was favored. We thus double the number of observations in our regressions. Our results—both those reported here and those in the next section—are robust to treating each game as a single game. In analyses using nonstacked data (not shown), we can reject the null hypothesis that players select between a coethnic and a noncoethnic randomly at the 99 percent level for coethnicity and the 95 percent level for coregion using a subjective coding of coethnic pairings.
18. Our decision to interpret the dictator game where the offerer is seen as a situation where the offerer can be punished for the violation of a social norm

- Macharia, Kinuthia. 1988. "Social Networks: Ethnicity and the Informal Sector in Nairobi." *Institute for Development Studies* working paper 463. Nairobi: University of Nairobi.
- Mozaffar, Shaheen, James R. Scarritt, and Glen Galaich. 2003. "Electoral Institutions, Ethnopolitical Cleavages, and Party Systems in Africa's Emerging Democracies." *American Political Science Review* 97(August): 379-90.
- Petrie, Ragan. 2003. "Trusting Appearances and Reciprocating Looks: Experimental Evidence on Gender and Race Preferences." Unpublished paper. Georgia State University.
- Posner, Daniel N. 2005. *Institutions and Ethnic Politics in Africa*. New York: Cambridge University Press.
- Smith, Adam. 1776/1994. *The Wealth of Nations*. New York: The Modern Library.
- Yamagishi, Toshio, and Nobuhiro Mifune. 2008. "Does Shared Group Membership Promote Altruism? Fear, Greed and Reputation." *Rationality & Society* 20(1): 5-30.