Experimental Evidence from Uganda, South Africa, and the United States on Ethnic Identification and Ethnic Deception*

Adam Harris
University College London

Daniel L. Nielson
Brigham Young University

Lily Medina
WZB Berlin

Clara Bicahlo Maia Correia
WZB Berlin

Michael G. Findley
University of Texas at Austin

Jeremy M. Weinstein
Stanford University

James Habyarimana
Georgetown University

Macartan Humphreys
Columbia University and WZB Berlin

Daniel N. Posner
University of California, Los Angeles

23 August 2018
Prepared for the APSA Annual Conference, Boston, 2018

Abstract

Micro-level theories of ethnic politics commonly assume that individuals can identify the ethnic backgrounds of the people with whom they interact. However, this assumption is difficult to reconcile with anecdotal evidence and with theoretical work that points to the importance of signaling (i.e. communicating or hiding one's ethnic identity) and mimicry (i.e. the ability to deceive another regarding one's ethnic identity) in social interaction. To what extent do placing (i.e., correct identification) and passing (i.e., deceiving others so as to prevent correct identification) occur? This study proposes a theoretical structure of social classification that enables us to examine the specific case of ethnic classification in order to answer these questions and employs three different experiments to assess patterns of placing and passing among subjects recruited from ethnically diverse communities in Uganda, South Africa, and the United States. When confronted with signs of group membership that cannot be easily manipulated in the short term (e.g. first language, surname), we find that subjects are only able to accurately place others into ethnic categories between 28% and 62% of the time Some evidence suggests that co-ethnicity and co-language aid in identification, but these results are fragile and not robust across information conditions. Ethnic identification errors provide the basis for constructing a new measure of group distinctness. We use this measure to examine the potential for passing through the strategic use of signals and find that successful passing is more likely when individuals attempt to pass into groups that are "similar" to their own but "distinct" from the group of the person who is trying to place them. We also find that, contrary to expectations from theories of rational updating, observers do worse when incorporating new information from signals sent by individuals that may, with a known probability, be attempting to pass. On the whole, the findings suggest that ethnicity is both less visible and less sticky than prominent conceptualizations assume.

*Adam Harris is lead author in bringing the various studies together. All other authors listed randomly. This paper builds on an earlier version from Habyarimana, Humphreys, Posner, Weinstein titled "Ethnic Identifiability: An Experimental Approach", which was presented at the annual APSA conference, 2004. The authors thank Claire Adida, Bernd Beber, Pepine Bulambo, Deo Byabagambi, Elizabeth Carlson, Chris Crabbe, Jamie Druckman, Kenneth Ekode, Nathan Falck, Autumn Gardner, Alex Kunobwa, McKenna Lane, Sylvester Mubiru, Douglas Musonga, Winfred Nabulo, Ruth Nagawa, Susan Najjuuko, Brenda Nakkazi, Harriet Nambi, Winfred Naziwa, Ken Noyes, Livingstone Ntensibe, Alex Odwong, Alexandra Scacco, Geoff Sentongo, Simon Ssenyimba, Elizabeth Suubi, Alex Tindyebwa, Sheila Watuwa and Daniel Young for their extraordinary contributions, the Center for International Studies at the University of Southern California and the University of Texas IR Workshops for helpful comments, and the Harvard Academy for International and Area Studies, the Russell Sage Foundation, and the Harry Frank Guggenheim Foundation for financial support.

1 INTRODUCTION

In late 1990, the Habyarimana regime in Rwanda harassed thousands of Tutsis in the run up to what became the Rwandan Genocide of 1994. In his account, journalist Philip Gourevitch reported that many mistakes were made in the identification of Tutsis. One of Gourevitch's Tutsi informants, a physician who figures prominently in his narrative, Odette Nyiramilimo, reported that in 1990 Habyarimana's men came to the hospital to arrest her. In her words, "…I had a colleague who had the same name. She was Hutu and she denied that she was me, but she was much taller than I am and they said, 'There's only one Tutsi doctor named Odette.' So she was imprisoned and tortured, and in 1994 she was again mistaken for a Tutsi, and killed" (Gourevitch 1999, pgs 83-4). Scholars have argued that this is not an isolated incident, finding that, without long-term neighbors to verify ethnic claims, the Interahamwe and other Hutu Power radicals frequently mistook Hutus for Tutsis and killed them indiscriminately (Davenport and Stam 2009).

During the 2008 attacks on immigrants throughout South Africa in which at least 62 people died and hundreds were injured, South Africans used a litmus test of sorts to differentiate between natives and immigrants. Assailants would ask targets to translate the word "elbow" into Zulu. If the person could not provide the Zulu word for elbow, they were assumed to be an immigrant (Hassim et al 2008). The use of this test shows that immigrants may not always be very distinguishable from other-ethnic South Africans. This is surprising because one would think foreign nationals would be among the most identifiable of non-co-ethnics.

Horowitz (2001) provides an example in which a Sri Lankan Tamil is not identified as such by Sinhalese rioters, with life-saving consequences for the Tamil man:

> Sinhalese rioters suspected a man in a car of being a Tamil. Having stopped the car, they inquired about his peculiar accent in Sinhala, which he explained by his lengthy stay in England and his marriage to an English woman. Uncertain, but able to prevent his escape, the rioters went off to kill other Tamils, returning later to question the prospective victim further. Eventually, he was allowed to proceed on his way, even though the mob knew it risked making a mistake, which in fact it had: the man was a Tamil (2001: 130).

Human Rights Watch documents an estimated two thousand backlash incidents targeted at Muslims and people of Arab descent in the aftermath of the September 11 terrorist attacks (Human Rights Watch 2002). More often than not, the victims of these hate crimes turned out not to be Muslims or Arabs at all, but Sikhs, Indians, Pakistanis, Coptic Christians, and, in one case, an Iranian Jew.[1] Other reports have documented the tendency for "brown-skinned men with beards and women with head scarves [to be] seen as 'Muslims'—regardless of their actual faith or nationality" (Kuruvilla 2006).

---

[1] We cannot rule out the possibility that the perpetrators of these anti-Muslim acts simply did not know that the Sikhs, Indians, and Pakistanis, and the others they attacked were not Arab Muslims, in which case the miscodings would not be examples of ethnic misidentification as we treat it in this paper but simply of not being aware that there were different categories into which the would-be victims might be coded.

These difficulties in ethnic identification call into question a bedrock assumption of many studies of ethnic conflict and ethnic identity politics more generally: that ethnic identity is actually identifiable from visible and verbal cues. Indeed, prominent scholarship by Max Weber (1968), Donald Horowitz (1985), and Kanchan Chandra (2004, 2006, 2012), while making clear that ethnicity is not always obvious, also argue that the relative visibility—or at least the identifiability through combined visual and verbal cues—of ethnicity sets it apart from other factors that divide people politically. From theories of in-group sanctioning (Greif 1989; Landa 1994; Fearon and Laitin 1996) to theories that emphasize the ability of ethnic groups to police their boundaries (Barth 1969; Laitin 1995; Fearon 1999; Caselli and Coleman 2002) to theories of ethnic or racial discrimination (Akerlof 1970, 1976; Becker 1971) to experimental treatments of minimal groups (Tajfel, Billig, and Bundy 1971) to theories of ethnic strife and conflict (Caselli and Coleman 2006, Diamond 1987, Horowitz 1985) to theories of public goods provision (e.g. Alesina, Baqir, and Easterly 1999, Alesina and La Ferrara 2005, Habyariman et al 2009, Gisselquist 2017, Gershman and Rivera 2018) and collective action problems (Deaux et al. 2006, Humphreys, Posner, and Weinstein 2002), models of face-to-face ethnic interaction almost always depend on the ability of actors to distinguish accurately between in-group members and outsiders.

The key problem with this broad literature's approach to ethnicity is that group boundaries may be neither crisp nor shared. These boundaries may be subject to contestation, as in disagreements among in-group members about what are the limits of their group. Even if there is consensus on common membership of a collection of individuals in a community, this group may not itself be recognized by out-group members, some of whom may have competing criteria for what constitutes the group. Consequently, statements about whether a person belongs to a given identity group may involve adopting one individual's perspective over another's

We focus on two of these theories: in particular in-group sanctioning and public goods provision. Fearon and Laitin (1996) argue that inter-ethnic conflict can be avoided if each group sanctions its own. However, if a member of group X mistakenly identifies a member of group Y as a co-ethnic (but the group Y individual identifies the other as a member of group X), and X sanctions Y for disapproving behavior, then the group Y individual will be sanctioned by a non-co-ethnic, which could promote conflict. Likewise, with regards to public goods provision, scholars have argued that individuals are more likely to cooperate with those from their own ethnic group. However, as Livy and Harris (2018) argue, if diversity is relatively high, then members of smaller groups are likely to be mistaken as members of the largest group, which leads to even less cooperation than the diversity debit hypothesis (that greater diversity undermines public goods provision) predicts: those of smaller groups, thinking their co-ethnics are non-co-ethnics, will not cooperate, likely leading to even lower public goods provision among smaller groups. On the other hand, there is likely to be more cooperation across ethnic lines between members of the largest group and those from smaller groups believed to be co-ethnics of the largest group. In short, by assuming ethnicity is identifiable, theories of inter-ethnic cooperation overestimate or underestimate the effect of diversity on public goods provision/cooperation depending on the relative sizes of the groups in society and which individuals are difficult to identify.

Scholars from Weber (1968) to Horowitz (1985) to Chandra (2012) have explicitly addressed the identifiability assumption. Chandra perhaps has made the clearest articulation of this assumption (2004, 2006, 2012). She notes that identification errors are possible and yet argues that ethnicity stands apart from other markers of social division due to its "visibility"—it is evident from sight—and "stickiness"—it is hard to change. Ethnic identities, compared to other social identities (i.e. class), are more visible because one can identify another's ethnic identity based on superficial interactions (Chandra 2012), suggesting that the requirements for discernment are minimal. The assumption of identifiability is difficult to reconcile both with anecdotal evidence suggesting large variation in the ability of individuals to identify others within given social schemas (e.g., Kuruvilla 2006) and with research on signaling and mimicry that emphasizes the ability of individuals to use information about their identities strategically (e.g., Gambetta 2005). Even in models inspired by constructivism, such as that found in Caselli and Coleman (2002) in which actors can take costly actions to change their identities, once changed, an actor's identity is assumed to be self-evident to the other players in the game. The assumption that individuals can seamlessly identify the ethnic backgrounds of the people they encounter—or, at the very least, unproblematically distinguish in-group members from out-group members—although common, presents both theoretical and empirical problems. Unfortunately, there has been very little systematic research to date to uncover the informational and strategic conditions under which certain groups or individuals are easily identifiable.

An obstacle to identifying such conditions—and a key reason why research on this topic is not further advanced—is the theoretical complexity of the notion of ethnic identifiability. To say that a person's ethnic background has been "correctly identified" implies the existence of a set of criteria for evaluating the accuracy of a placement decision in a given instance. But there are no such universal criteria,[2] and the appropriateness of *ad hoc* criteria used in any given situation will likely be dependent on the strategic context of the placing. Moreover, even if we have clear criteria for correct placement, determining whether a rate of correct identification under any information condition is 'high' or 'low' depends on base rates, which themselves are likely to be highly sensitive to contextual factors. Beyond these theoretical issues of interpretation, difficulties also arise in observing empirically how individuals classify others, or if indeed such classifications are even made. The result is that we know little about the extent to which "accurate" placing and successful passing occur, or about how the ability to do each of these things depends on informational environments.

[2] The most obvious criterion is the rule that "correct" placing is simply identifying the person with the category that the person him/herself uses to describe him/herself. However, recent research has suggested that self-identification is often challenged, influenced by how others identify an individual, and anything but straightforward (see for example Harris 2018, Dulani et al 2018, others?).

This study seeks to fill this gap by proposing a theoretical structure for making meaningful statements about identifiability. Then, drawing on three separate experiments all with similar designs in urban settings in Uganda (Kampala), South Africa (Johanesburg, Durban, East London), and the United States (Los Angeles), we use experimental methods to explore the determinants of identifiability and successful passing within urban populations in these settings.[3] Specifically, we measure the ability of individuals to place others (whose images they are shown) into the ethnic categories with which those other individuals self-identify. We examine how the characteristics of the person viewing the images, the characteristics of the person whose images are viewed, and the degree of information that the former has about the latter affect the probability of a correct ethnic identification. We then explore the contribution to identification made by physical "signs"—by which we mean aspects of the individual that are not subject to manipulation, at least in the short run—and "signals"—by which we mean information that is more readily under an individual's control and can be used to either communicate or hide their ethnic identity. The latter enables us to analyze the factors that permit individuals to pass as members of groups other than their own.

Our effort to document patterns of placing and passing is more than an empirical exercise. Variation in ethnic identifiability has substantive implications for theoretical and empirical work on ethnic politics. To the extent that ethnic identifiability varies systematically across ethnic groups, their ability to achieve collective ends should vary as well. The fact that passing may be easier for some individuals and groups than for others has implications for the permeability of group boundaries, the ability of groups to police them, the collective benefits that flow from policing, and the effects of diversity on public goods provision. Because the costs and benefits of gathering information about the ethnic identity of individuals may vary across groups, theories of ethnic mobilization, discrimination, sanctioning, and public goods provision must take account of these differences. Moreover, and no less important, given the difficulties in identification, there are potential risks in including this attribute in the basic conceptualizations of ethnicity. In what follows, we attempt to more clearly define the terms of ethnic classification, as an instance of broader social classifications, motivate the experiments, describe the research designs, apply a new measure of ethnic distinctness, and conclude by enumerating the implications of the findings for the study of ethnic politics and for future research.

2 DEFINITIONS

To address these concerns, we propose a notion of social classification that generates measures of identifiability from binary relationships between individuals conditional on arbitrary benchmark classification rules. We begin by defining a social demography. Given a triple $<N, I, f>$ in which N denotes a population, I denotes a set of categories, and f denotes a mapping from N to I (a classification rule), we say that $<N, I, f>$ is

[3] Placing and passing are likely to be quite different in heterogeneous urban and homogeneous rural areas. We focus here on urban areas where diversity is high and thus identifiability is likely more consequential for outcomes.

a "*social demography*" if f is a *function* that places each individual in one and only one category.[4] The restrictions on the placement rule are clearly weak and do not imply any "objective" membership of a category. Membership depends entirely on how f is defined. Examples of f include "i is a member of category j in I if i's mother was a member of that category" or "$f(i)$ is the category into which i places herself when filling out a census form."

In any given context, however, individuals may have personal rules, c_j, for placing player i, given N and I, that may or may not be very different to f. For example, c_j may place an individual i in group B if in j's estimation, i looks like a "typical B" or perhaps if j believes that i is likely to self-place as a B. We exploit this feature in order to construct notions of pairwise identifiability, group identifiability, errors of inclusion and exclusion, and group distinctness. To do so we explicitly privilege some particular rule f and generate measures of correct classification (given f) based on the congruence of such personal rules and f.[5]

Pairwise Identifiability

We define correct identification as follows. Given a social demography $<N, I, f>$, and a rule c—the personal rule that an individual employs to map N to I— we say that the rule *c correctly identifies i* if c and f place i in the same category of I, that is, if $c(i \mid N, I) = f(i \mid N, I)$. In a world of uncertainty, we treat the notion of identifiability as a *probability*. Thus we define the *pairwise identifiability* of i by j under demography $<N, I, f>$ and private placement rule c_j as the probability with which j correctly places i: $g(i,j,c_j \mid N, I) = \Pr[c_j(i \mid N, I) = f(i \mid N, I)]$. For example, if f is the self-placement criterion, c is "player j's best guess about i's ethnicity within a given set of ethnic categories under a given informational context," then we can say that the identifiability of *i for j*, is the probability that j's best guess about i's ethnicity after learning i's name, corresponds to i's self-placement.[6]

Group Identifiability

With this foundation, we can move to more abstract statements about group identifiability. Define the identifiability of group k (under f), for an individual j as the probability with which j codes a random member of group k in group k. Define the identifiability of a given group A, for a given group, B, as the probability

[4] The requirement that each individual falls within only one category is not onerous. It allows for the possibility of membership in categories along multiple dimensions of identity, and it allows for membership in two or more categories along a single dimension of identity. By the same token, should an individual not be classifiable under any category in some set, then she can be classified as unclassifiable in an expanded set of categories. Note further that the notion employed here does not imply any substantive consistency between the categories in I; for example that all elements in I be regional, religious or ethnic groups; indeed to incorporate the possibility that individuals have portfolios of identities, any element in I may be a combination of multiple categories on seemingly disparate dimensions of identity.

[5] Note that the personal placement considered here are also mappings from N to I. implicitly then we exclude from analysis the congruence of rules that map to sets of categories different from I.

[6] Note that whereas others (for example Chandra 2004) argue that identifiability may be a *defining* characteristic of an ethnic category, we treat identifiability as an empirical property of an arbitrary social demography, that may or may not obtain in practice. We use experimental methods specifically to allow us to measure how identifiability varies with informational and strategic context.

that a random member of A (under f) codes a random member of B (under f) as a member of B. As a special case, we refer below to group A's identifiability by N simply as the group's "identifiability."

Signs and Signals

The identifiability of individuals depends then on the information that viewers have at their disposal and the way they choose to act on this information. In many politically relevant contexts, the availability of this information is itself a choice and may be subject to strategic manipulation by the person to be identified. To study this feature of social interaction, we draw on a useful distinction between signs and signals of group membership developed by Gambetta (2005). Although the distinction is imperfect, we use "**sign**" to denote a manifestation of group membership that is beyond an individual's control, at least in the short run, and we use "**signal**" to denote an action taken by an individual in order to communicate membership of an identity group. To apply concepts modified from Gambetta, passing implies a three-way relationship, between the would-be passer (the **portrayer**), the individual that the portrayer is trying to convince (the **subject**), and the type that the portrayer is trying to pass as (the **model**). Signaling theory suggests that the use of signals can lead to successful passing in situations in which the signal is available to portrayer from multiple groups but at different rates. If a signal is available to all individuals then it may not carry any information about group membership; if it is only available to one group, then it cannot be used by another to pass as members of that group (Bacharach and Gambetta 2001).

While signaling can lead to successful passing, its impact on identification will depend on the relative costs to individuals of making different types of errors. An example in the appendix describes the logic in more detail and illustrates how the introduction of the possibility of signaling can result—even in a context of optimal signal extraction—in *either* a rise or fall in the identifiability of a given group and, correspondingly, to *either* a rise or fall in the occurrence of inclusion or exclusion errors. While either type of error can increase based on the distribution of signals and the player's priors, in a context with uniform priors over membership, the introduction of new information should not result in a *net* rise in both types of error.

The formal definition of these notions of identifiability and signs and signals, allows us to make progress in the measurement of identifiability. As emphasized above, and as indicated by the formal expressions we use, a group's identifiability is radically contingent—it is a function of the information available, the characteristics of identifiers and identified, and, crucially, the social demography invoked. This is why, throughout the paper, when we refer to identifiability we do so *for a given social demography* as defined at the beginning of this section. This contingency means that the specifics of our empirical frame are important for interpreting our results. It also means, however, that we can seek to make more general statements about identification processes *as a function of the empirical frame employed*. In particular, we can examine how the identifiability or distinctness of groups change as a function of the set of categories used, the criteria for membership of groups, the population under consideration, and most germane for our work in this paper, of the information made available to—and/or by—individuals.

3 EXPERIMENTS

We now describe three experiments that each fixed a particular social demography and attempted to measure the identifiability of individuals within these demographies. Our design builds on experimental protocols pioneered by social psychologists interested in how individuals assign people to social categories (Allport and Kramer 1946; Secord 1959; Blasovich, Wyer, Swart, and Kibler 1997; Harris 2002) as well as theoretical work from social phsychology investigating how individuals place others into categories (Nosofsky 1986, Terry et al 1999, Fyer and Jackson 2008). The experimental set-up permits us to assess, first, the ability of individuals to *place* others conditional on a controlled set of signs. On the basis of these measures of identifiability, we can describe the propensity of in-group members to make errors of exclusion and inclusion and produce empirical measures of the distinctness of different groups. Our experimental approach also introduces an innovation not included in previous studies: the possibility of strategic signals about identity. From this data, we can examine the potential for *passing* and explore how individuals incorporate the information provided by noisy signals, given the signs already available to them.

General Experimental Protocol Across All Studies

All three experiments share a common structure and build upon one another sequentially, so here we describe the joint setup across all three experiments and delineate relevant differences between them below. As a "benchmark" classification rule fixing a social demography, we used information from a pre-experiment questionnaire asking how the individuals the subjects would later categorize, which we call portrayers, self-identified their ethnic group membership. We then invited subjects to guess the identities of a set of randomly selected portrayers, rewarding them monetarily when their guess corresponded with how the portrayers identified themselves. Importantly, we manipulated experimentally the information about portrayers available to subjects participating in the exercise. For each portrayer, we collected still images and video sequences, each providing the subjects viewing the images or videos with different levels of information that could be used by the subjects to make an inference about the portrayer's ethnic background. The first information was a simple headshot. Subsequent information levels involved short videos in which the portrayers, for example, greeted the camera in their native language, provided other identifying information such as their given or family names, or made a more involved argument for the ethnicity they were claiming, as below.[7]

By exposing subjects to an image or video of a portrayer at a randomly selected level of information and asking them to guess that portrayer's ethnic background, we are able to measure the impact of

[7] Both Isaacs (1975) and Chandra (2004) emphasize the importance of the information contained in a name. In Fershtman and Gneezy (2001), names are employed as the sole marker of ethnic group affiliation in experiments that seek to uncover patterns of inter- and intra-ethnic interaction. Charness and Gneezy (forthcoming) measure the more general impact of names on behavior in dictator and ultimatum games.

informational contexts on ethnic identifiability and to examine how identifiability depends on observer characteristics and ethnic group membership. The structure of the experiments worked as follows (further details, along with the specific proctor instructions, are available in the online appendix). Subjects were told that they would see a random sample of many images/videos of portrayer individuals and would be asked to guess the ethnicity of the portrayers. After (privately) seeing each image or video on a computer screen, subjects were asked to guess how the portrayer would have classified him or herself if asked. For every correct guess, subjects received compensation.

To explore how the introduction of signals impacts identifiability, we developed an additional set of treatments. Portrayers were asked to record further videos in which they would attempt to convince subjects either of their true or of a false ethnic identity. These videos were longer videos (up to 30 seconds in length) and in almost all cases included an explicit statement of the ethnic group to which the portrayer was claiming to belong. Versions were recorded in which portrayers attempted to convince subjects of their true identities (we call these "simulation" videos). Additional videos were then recorded in which portrayers tried to convince subjects of a false identity (we call these "dissimulation" videos). In all three experiments, subjects were instructed that the portrayers may be describing their true ethnic identities or that, alternatively, they may be lying and attempting to pass as a different ethnicity.

Experiment 1: Uganda[8]

Ugandan subjects were drawn from among the country's diverse ethnic groups employing a benchmark demography that uses the ethnic categories and figures contained in the 2000 Ugandan census. The capital, Kampala, which is located in the center of the country and where the experiment took place, is populated by people who (or whose parents) migrated from all regions of the country. Our subjects were drawn from the Kawempe division of Kampala in an area that we call Mulago-Kyebando, which is more diverse than other parts of the country. The Baganda are the dominant ethnic group in Mulago-Kyebando, as they are throughout Kampala. They are followed in the distribution by two sizeable western groups, the Banyankole and the Bafumbira, and one eastern group, the Basoga. Other western groups (including the Bakiga and Batoro) are also present, as are a small number of residents from northern groups.

When we prepared to recruit subjects in Mulago-Kyebando, we were faced with two conflicting goals. One was to draw a fully representative sample. A second (and, unfortunately, incompatible) goal was to maximize statistical power. Given the demographic profile of Mulago-Kyebando, a simple random sample would have resulted in many in-group pairings among the Baganda, many out-group pairings that included Baganda, and very few in-group pairings of any other kind. A simple random sample with randomized matching would have prevented us from observing the number of pairings we needed to work out the impact of shared group membership (independent of patterns of interaction among Baganda).

[8] The Uganda experiment took place in MONTH of 200?

With two worthy goals but only one sample to draw, we selected a middle ground. Using data from a key informant survey, we identified a subset of local council units (LC1s) in Mulago-Kyebando that include somewhat higher proportions of the Banyankole and Bafumbira (the second and third most common ethnic groups). We then set sampling targets for each LC1 such that random sampling techniques within each LC1 would yield a sample population with a slightly smaller share of the largest group in Mulago-Kyebando, the Baganda, and larger shares of the Banyankole and the Bafumbira (and the Banyarwanda as well, a group considered similar ethnically to the Bafumbira). Using this technique, we drew a total sample of 300 subjects from this area.[9] Our sampling strategy yielded, as intended, a slight underrepresentation of the largest group, the Baganda, and an overrepresentation of the smaller groups, notably the Bafumbira and Banyakole.

To ensure that all subjects had the same prior beliefs about the distribution of ethnic groups in the area, we described explicitly the approximate distribution of ethnic groups in Mulago-Kyebando. For every correct guess, participants received 100 USh; their total potential winnings were 5000 USh (approximately three dollars). Two hundred and seventy-four subjects participated in the identification game, producing a total of 15,265 guesses.

Subjects and portrayers were drawn from the same population. For each portrayer, we collected five different images with a digital camera, each providing the subjects viewing the images with different levels of information that could be used to make an inference about the subject's ethnic background. Information level 1 was a simple headshot. Information levels 2 and 3 involved two short videos in which the subject greeted the camera, respectively, in Luganda (Kampala's *lingua franca*) and in the respondent's primary language. Information levels 4 and 5 involved video images, like those in 2 and 3, but in which the subject also provided his or her given and family names. Subjects (128 in total) that took part in the first half of sessions of the identification exercise were asked to record four further videos in which they would attempt in two of the videos to convince subjects either of their true ethnicity (simulation) or in two additional videos of a false identity (dissimulation) for a *model* ethnicity selected by us as the researchers.

In Uganda, the *models* used for the dissimulation videos were selected with probability proportionate to the relative size of a group in the subject population, excluding the group of the portrayer being recorded. This probability is useful in that it results in an approximately equal probability that a given individual is dissimulating conditional upon the group of which he claims to be a member, although it leads to a somewhat lower share of dissimulation videos among those claiming to be Baganda.[10] Importantly, by controlling the

---

[9] Random sampling was conducted using a random compass. Teams entered each LC1 at a central point and chose random directions and a random walking period. At the designated stopping time, they selected the closest house and created a basic roster of household members. Stratifying by gender, a household member was then selected at random and invited to participate in the games. Members were invited whether or not they were present at the time of the selection. In such cases, their names and ages were recorded and these confirmed upon registration. Using this approach, we had a very low non-participation rate: more than 75% of those we selected agreed to participate in the study.

[10] That is, if an individual of group k dissimulates ("Dis") (an event which occurs with probability $.5a_k$) then he claims to be a member of group j ("states j") with probability $a_j / (1 - a_k)$. An application of Bayes' rule then yields: Pr(Dis | States j) = $1 - .5a_j / (.5a_j + \Sigma_{-j} .5 a_k a_j / (1 - a_k)) = 1 - 1 / (1 + \Sigma_{-j}(a_k / (1 - a_k)))$. For example, if all groups are of approximately equal

distribution of models, we are in a position to estimate not simply whether portrayers can pass in general but what the prospects are for passing across each portrayer-subject-model combination.

In all cases, subjects were given a short amount of time to think about how they wanted to present themselves for the simulation and dissimulation videos; beyond being asked to state the group to which they were claiming membership, they were free to use any signals available to them. In addition, players were given an incentive in the dissimulation videos to "try hard"; specifically, they were told that "For every person that you are able to convince that you are "*Y*", you will be paid 500 USh. This means that if you try hard, you could earn up to 4000 USh."

For the second set of sessions, 139 subjects played both the identification game and the "simulation/dissimulation" game. Note that none of the subjects that took part in the identification game, in the first or second sessions, had previously recorded simulation and dissimulation videos. Subjects then viewed simulation/dissimulation videos for the same sets of portrayers. Enumerators advised subjects that (a) about half the time people in the videos were telling the truth and half the time people were trying to pass and (b) that the true distribution of groups was as in the population of Mulago-Kyebando.

Experiment 2: South Africa[11]

Subjects were recruited through random household sampling starting within three townships of South Africa: Daveyton, Umlazi, and Mdantsane. The three townships are located near Johannesburg, Durban, and East London, respectively. We chose these townships because they offer useful variation on South African ethnic identities. In total, there are nine different official black ethnic groups in South Africa: Zulu, Xhosa, Sotho, Tswana, Tsonga, Pedi, Swazi, Venda, and Ndebele. Daveyton is an ethnically diverse township representing each of these official ethnic groups in South Africa. Working in Daveyton allowed us to capture a good proportion of the smaller ethnic groups that are found in South Africa. In contrast to the heterogeneous Daveyton, Umlazi is predominantly Zulu, whereas Mdantsane is primarily Xhosa, thus offering the possibility to compare homogeneous and heterogeneous areas. The ethnic makeup of our sample approximates the ethnic makeup of the black South African population at large. See online appendix for details.

We offered subjects 20 rand (about $2.50 at the time) to come to a lab (at a college in Daveyton, a college in Umlazi, and a primary/secondary school in Mdantsane) to participate. In addition to the base compensation, we also offered further monetary incentives for correctly guessing photos/videos. Subjects could receive up to 64 rand (about $8.20) for their participation, provided they guessed correctly in every case

---

size, then Pr(Dis | States j)=.5. If all groups are small then Pr(Dis | States j) is close to .5. If however there is one larger group then it is in fact more likely that when this group is stated that the truth is being told. To see the logic consider a group X of size 60/100. This group will produce truthful statements about X 60/2=30 times. The remaining groups will provide false statements regarding membership of X at most 40/2=20 times. Thus most statements (at least 60% of statements) about membership in X will be true statements. Given the distribution of types in our sample, the expected Pr(Dis | States "Baganda") is 37% while for all other groups it is 57%-58%. In the empirically sampled set of interactions, Pr(Dis | States "Baganda")=42% while for other groups it is 58%.

[11] This experiment was conducted June through August 2009.

(which none did). On average subjects guessed 28 percent correctly and received 32.3 rand (roughly $4) in compensation.

We exposed subjects to a photo alone or to a photo and a randomly sampled video of a portrayer. After being exposed to the portrayer's photo with no additional information, the subject provided a guess about the portrayer's ethnic identity; the subject then viewed a video of the portrayer and could enter a new ethnicity guess.[12] Subjects repeated the guessing for 22 separate portrayers (44 guesses in all), each of which was presented in random order. In South Africa, 634 subjects participated, making a combined total of 27,896 guesses.

The ethnic portrayers in the photos and videos were recruited the year before based on their ethnic identity, age, and gender. They thus were not drawn from the population of subjects. The group of portrayers closely matches the national demographics of black South Africans, so that there is roughly the same percentage of each ethnic group among our portrayers as there is in the overall population. We informed subjects that the portrayers reflected the underlying ethnic proportions in the South African black population. To accompany the still photo, we filmed ten different videos of each ethnic portrayer. In the videos, the portrayer does one of the following: (1) states first given name, (2) states surname, (3) greets in native language, (4) greets in English,[13] (5) states and argues for true identity, (6) states and argues for true identity with supportive ethnic symbol in background, (7) states and argues for true identity with contradictory ethnic symbol, (8) states and argues for false identity, (9) states and argues for false identity with supportive symbol, and (10) states and argues for false identity with contradictory symbol.[14] As in Experiment 1 in Uganda, we paid portrayers for successfully deceiving subjects in the dissimulation videos in an attempt to motivate them toward greater effort at credible dissimulation.

With the last six videos, we aimed to test the effects of ethnic signals (and symbols), with varying degrees of veracity. For the signal videos, portrayers provided some information about their ethnicity in English. We scripted the signal videos to ensure consistency; portrayers provided their stated ethnicity, their parents' ethnicity, the region from which their family comes, and information about up to three articles of traditional clothing. Importantly, in three of the videos the portrayers stated their ethnicity truthfully (simulation); in the remaining three, the portrayers claimed to belong to a different ethnic group of their choice (dissimulation). As in Uganda, we informed subjects that roughly half the time, portrayers would be stating their true ethnic identity and the other half of the time would be arguing for a false identity.

In South Africa, allowing portrayers to choose the false ethnicity complements the probability-weighted assignment by us as researchers in the Uganda experiment. South African portrayers typically chose the false ethnicity they knew best and therefore were likely to provide stronger signals than if they had been

[12] This procedure differed from the Uganda study, where subjects viewed either a photo or video, but not both.
[13] This also differed from the Uganda study in that the information set was additive (e.g., surname plus greeting in native language).

randomly assigned a false ethnicity for which strong signals may be less known. Results for successful dissimulation in South Africa and Uganda therefore provide a range of possible pass rates enabled by stronger or weaker signals, respectively.

In addition to the basic information provided in the signal videos, we also presented ethnic symbols in the videos. We incorporated the symbols by placing a photo of a cultural symbol linked to a specific ethnic group on the wall behind the portrayer in the video. Testing the effect of these symbols allows us to evaluate whether subjects overlook verbal claims when a physical symbol contradicts the statement. Portrayers could present false information by stating a signal from the wrong ethnicity or by having a symbol from an ethnic group other than their own in the background. Once the subjects finished viewing the photos and videos, we calculated how many correct guesses they made and paid them accordingly.

Experiment 3: United States[15]

As in Experiment 1, each participant in the United States experiment played both roles of subject and portrayer. The participants consisted of undergraduate students drawn from two large universities on the West Coast. The participants were recruited from seven ethnic groups that have large presences on both campuses: African Americans, Arabs, Asians, Caucasians, Indians, Persian/Iranians, and Latino/as. Approximately 54 percent of the participants were recruited through ethnic student associations on each campus. The other 46 percent were recruited from the regular subject population of participants in university laboratory experiments.

We collected a headshot of each portrayer, a video with a very brief greeting, and a second video of the brief greeting and the portrayer's name. A randomly drawn sub-sample of participants recorded three additional videos in which they explicitly stated their ethnic backgrounds, the first two claiming their true ethnicity (simulation) and the third claiming a different ethnicity for which they would attempt to pass (dissimulation) of the portrayer's choice from among the seven.[16] While 120 participants had their images recorded, due to attrition only 96 ultimately participated as subjects in the experiment. Subjects viewed 23 images and videos and were paid 20 cents for each correct guess. In addition to their $5 base participation compensation, the maximum a respondent could have earned was approximately $15.60. On average, respondents earned $11.13.

[14] Conditions 5-6 are similar to the Uganda study, and conditions 7-10 are additions to what was considered in the Uganda study.
[15] The U.S. experiment took place in MONTH of 200?.
[16] We invited thirty-six participants to record the additional images, of whom thirty-two accepted our invitation and had their images recorded.

4 EMPIRICAL RESULTS

4.1 VARIATION IN IDENTIFIABILITY

Tables 1(a)-1(c) report how individuals from different ethnic groups are classified (on average). Each cell entry shows the percentage of viewings in which an individual of a row type is classified as an individual of the column type. A quick look at the off-diagonal elements provides evidence of the extent to which people miscode. If classifications were always correct, this matrix would be an identity matrix—all the diagonals would contain a one and the off-diagonals a zero. The difference between this matrix and the identity matrix is a measure of the degree of misclassification. For example, a Banyankole is correctly classified as a Banyankole 41% of the time (upper left cell). When mistakes are made, Banyankole are most likely to be classified as members of the most populous group, the Baganda (17% of the time). Banyankole are also often thought to be members of other western groups, including the Bakiga (14% of the time), Batoro (8% of the time), and Bafumbira (10% of the time). In a number of cases, individuals from one group are classified in another category at least as often as they are correctly coded. The Bakiga and the Batoro, for example, are guessed to be Banyankole more frequently than they are classified correctly.

1(a): Classifications and Misclassifications in Uganda

	Banyankole	Baganda	Bagisu	Bakiga	Banyarwanda	Basoga	Batoro	Banyoro	Iteso	Bafumbira
Banyankole	0.41	0.17	0.01	0.14	0.04	0.03	0.08	0.01	0.01	0.10
Baganda	0.09	0.70	0.01	0.03	0.03	0.05	0.04	0.01	0.01	0.03
Bagisu	0.13	0.36	0.16	0.04	0.02	0.10	0.06	0.02	0.04	0.06
Bakiga	0.34	0.20	0.01	0.22	0.03	0.03	0.07	0.01	0.01	0.08
Banyarwanda	0.19	0.34	0.01	0.07	0.16	0.02	0.05	0.01	0.01	0.14
Basoga	0.11	0.45	0.02	0.03	0.03	0.23	0.06	0.02	0.02	0.04
Batoro	0.25	0.25	0.01	0.06	0.06	0.02	0.24	0.04	0.01	0.05
Banyoro	0.21	0.35	0.01	0.07	0.05	0.06	0.14	0.06	0.01	0.04
Iteso	0.07	0.33	0.02	0.04	0.02	0.09	0.03	0.02	0.31	0.08
Bafumbira	0.17	0.17	0.01	0.06	0.10	0.04	0.06	0.02	0.02	0.34
Total	0.17	0.48	0.01	0.06	0.05	0.05	0.07	0.02	0.02	0.08

Note: This table reports the likelihood that a row type will be classified (by an arbitrary player) as a column type. For example, the cell in the second row, first column shows that Baganda were classified as Bayankole 9% of the time. The principal diagonal contains the rates of correct classification.

Table 1(b) reports the same results but for South Africa. It demonstrates that the Ndebele and Tswana are the most difficult to identify (portrayers of these groups are only correctly identified 18% of the time); Xhosa portrayers, on the other hand, are the easiest to identify (they are correctly classified 35% of the time). One can also look at the other cells, which report rates of misclassification. For example, this shows that Xhosa portrayers are most often mistaken as Zulu (15% of the time) and that Zulu portrayers are most often mistaken as Xhosa portrayers (13% of the time). These mistakes should be expected given that the two groups speak mutually intelligible languages and have various similar customs. This is also true for the Pedi, Sotho, and Tswana groups. The three groups' languages are mutually intelligible and are all most often mistaken as one of the other groups. These results add validity to our experimental setup: the anticipated mistakes are made and thus the information we are providing appears relevant and accurate for the type of experiment we are conducting. We should worry that the experiment is not well designed if we find that Xhosas and Zulus are infrequently mistaken as each other.

Table 1(b): Classifications and Misclassifications in South Africa

	Ndebele	Pedi	Sotho	Swati	Tsonga	Tswana	Venda	Xhosa	Zulu
Ndebele	0.18	0.04	0.16	0.05	0.13	0.04	0.17	0.09	0.14
Pedi	0.07	0.21	0.16	0.08	0.08	0.09	0.07	0.11	0.13
Sotho	0.07	0.09	0.25	0.08	0.07	0.09	0.05	0.14	0.16
Swati	0.05	0.09	0.07	0.19	0.09	0.10	0.11	0.11	0.19
Tsonga	0.06	0.07	0.10	0.07	0.24	0.07	0.14	0.14	0.1
Tswana	0.08	0.09	0.13	0.07	0.08	0.18	0.10	0.13	0.14
Venda	0.07	0.05	0.08	0.11	0.07	0.05	0.26	0.17	0.14
Xhosa	0.08	0.07	0.08	0.07	0.07	0.06	0.07	0.35	0.15
Zulu	0.08	0.07	0.10	0.07	0.08	0.07	0.09	0.13	0.32
Total	0.08	0.09	0.12	0.08	0.09	0.08	0.09	0.18	0.20

Note: This table reports the likelihood that a row type will be classified (by an arbitrary player) as a column type. For example, the cell in the second row, first column shows that Pedis were classified as Ndebele 7% of the time. The principal diagonal contains the rates of correct classification.

1(c): Classifications and Misclassifications in the United States

	African American	Arab	Asian	Caucasian	Indian	Latino/a	Persian/ Iranian
African American	0.80	0.01	0	0.05	0.02	0.09	0.03
Arab	0.02	0.29	0	0.30	0.05	0.15	0.18
Asian	<0.01	0.01	0.94	0.01	0.02	0.02	<0.01
Caucasian	0	0.02	<0.01	0.83	0.01	0.06	0.08
Indian	0.01	0.18	0.01	0.04	0.52	0.05	0.21
Latino/a	0.03	0.04	0.02	0.18	0.06	0.62	0.05
Persian/Iranian	0	0.18	0.01	0.13	0.12	0.10	0.46
Total	0.09	0.06	0.21	0.31	0.05	0.17	0.10

Note: This table reports the likelihood that a row type will be classified (by an arbitrary player) as a column type. For example, the cell in the second row, first column shows that Arabs were classified as African Americans 2% of the time. The principal diagonal contains the rates of correct classification.

Table 1(c) reports similar results for the United States experiment. In this context, Asians are correctly classified as Asian 94% of the time, the highest of the groups included, and Arabs are classified as Arabs only 29% of the time, the lowest of the groups included. Various interesting patterns emerge: Arabs are incorrectly classified as Caucasian 30% of the time; Indians are incorrectly classified as Arab 18% of the time and as Persian/Iranian 21% of the time; Latinos are incorrectly classified as Caucasian 18% of the time; Persian/Iranians are incorrectly classified as a number of other groups a sizeable proportion of the time.

While Tables 1(a)-1(c) offer a picture of how individuals are misclassified, Tables 2(a)-2(c) provide detail about who does the misclassification. The cells report the frequency of correct identification, averaged across all information levels and for each pair-wise combination of groups. Consider first the final column on the right-hand side of the tables. The numbers in this column reflect how often, on average, individuals in each of the benchmark ethnic categories correctly identify others: about 50% of the time in Uganda, 25% in South Africa, and 60% in the United States. Interestingly, in all three countries there is not much variation across groups in their ability to correctly identify the ethnic backgrounds of others. In Uganda (Table 2(a)), the Banyankole and Batoro are most successful and the Basoga are least successful, but these differences are not substantially large and from a statistical point of view we cannot reject the null hypothesis that there are no differences across groups in their guessing ability. There is significantly more variation across the bottom row of Table 2(a). This row captures the likelihood that an arbitrary member of each of the column types is successfully identified. The Banyoro, for example, are almost never identified correctly as Banyoro. The group that is most commonly correctly identified is the Baganda—68% of the time.

2(a): Correct Guesses in Uganda

	Banyankole	Baganda	Bagisu	Bakiga	Banyarwanda	Basoga	Batoro	Banyoro	Iteso	Bafumbira	Total
Banyankole	0.50	0.70	0.07	0.15	0.16	0.18	0.34	0.10	0.11	0.47	0.49
Baganda	0.39	0.70	0.16	0.20	0.17	0.19	0.23	0.03	0.20	0.27	0.46
Bagisu	0.38	0.71	0.14	0.19	0.09	0.24	0.14	0.05	0.56	0.23	0.47
Bakiga	0.28	0.67	0.04	0.38	0.18	0.21	0.16	0.03	0.31	0.37	0.46
Banyarwanda	0.56	0.58	0.24	0.31	0.24	0.19	0.36	0.00	0.17	0.33	0.45
Basoga	0.34	0.61	0.13	0.13	0.00	0.60	0.23	0.05	0.31	0.46	0.43
Batoro	0.46	0.70	0.11	0.19	0.16	0.37	0.29	0.20	0.35	0.36	0.5
Banyoro	0.30	0.62	0.11	0.33	0.17	0.21	0.18	0.33	0.31	0.32	0.45
Iteso	0.44	0.70	0.30	0.00	0.00	0.15	0.00	0.00	0.33	0.33	0.44
Bafumbira	0.39	0.64	0.13	0.28	0.22	0.24	0.33	0.13	0.27	0.55	0.47
Total	0.41	0.68	0.14	0.21	0.16	0.22	0.25	0.06	0.23	0.35	0.47

Note: This table reports the probability that a row type will correctly classify a column type.

Table 2(b): Correct Guesses in South Africa

	Ndebele	Pedi	Sotho	Swati	Tsonga	Tswana	Venda	Xhosa	Zulu	Total
Ndebele	0.21	0.25	0.30	0.23	0.37	0.20	0.43	0.36	0.33	0.31
Pedi	0.10	0.31	0.25	0.21	0.32	0.17	0.28	0.27	0.25	0.25
Sotho	0.10	0.26	0.25	0.15	0.27	0.14	0.22	0.32	0.34	0.28
Swati	0.07	0.3	0.25	0.15	0.24	0.24	0.28	0.28	0.29	0.26
Tsonga	0.20	0.22	0.26	0.26	0.30	0.19	0.24	0.32	0.30	0.27
Tswana	0.06	0.25	0.20	0.13	0.17	0.12	0.38	0.28	0.29	0.24
Venda	0.50	0.27	0.33	0.00	0.25	0.13	0.40	0.17	0.29	0.24
Xhosa	0.22	0.19	0.22	0.17	0.19	0.18	0.24	0.44	0.29	0.28
Zulu	0.19	0.19	0.26	0.21	0.25	0.18	0.25	0.31	0.37	0.28
Total	0.31	0.25	0.28	0.26	0.27	0.24	0.24	0.28	0.28	0.28

Note: This table reports the probability that a row type will correctly classify a column type.

Table 2(c): Correct Guesses in the United States

	African American	Arab	Asian	Caucasian	Indian	Latino/a	Persian/ Iranian	Total
African American	0.60	0.16	0.80	0.72	0.16	0.57	0.29	0.56
Arab	0.66	0.39	0.84	0.70	0.35	0.45	0.40	0.60
Asian	0.64	0.22	0.85	0.69	0.53	0.54	0.30	0.63
Caucasian	0.65	0.26	0.85	0.76	0.44	0.47	0.39	0.63
Indian	0.69	0.22	0.77	0.61	0.80	0.43	0.34	0.55
Latino/a	0.65	0.20	0.79	0.70	0.20	0.56	0.40	0.59
Persian/Iranian	0.63	0.27	0.87	0.68	0.70	0.61	0.75	0.67
Total	0.64	0.24	0.83	0.72	0.42	0.52	0.39	0.62

Note: This table reports the probability that a row type will correctly classify a column type.

Table 2(b) contains similar data, but for South Africa. Here we find that the rate of successful identification is correlated with group size. The two largest groups (Xhosa and Zulu) are the most successful at identifying others, especially themselves (See Xhosa/Zulu rows). The table also shows that people are best at identifying their co-ethnics (a result we return to in the regression analysis), but even the most successful co-ethnics still do worse than a 50 percent success rate (Xhosas at a 44 percent success rate). The results also show that, on average, Tswanas are one of the most difficult to identify (along with the Swati), even among themselves; Tswanas only successfully identify each other 12 percent of the time and are substantially better at identifying other groups such as Pedi and Zulu. Thus, while co-ethnics perform best, their ability to identify each other is much lower than one would initially expect, and the results come with a few additional surprises.

Table 2(c) reports the same results but for the United States. It demonstrates that, on the whole, ethnic groups are more identifiable than those in the Uganda or South Africa contexts. Subjects had relatively little difficulty correctly identifying Asian and Caucasian students. Subjects, even co-ethnics, had much more difficulty correctly categorizing Arab and Latino students. Yet overall, as these results suggest, and as we will show more systematically in a moment, the ethnic group membership of the subject is by far the most important predictor of identification success.

So far, we have focused only on how groups vary in how easily they can be identified and how well others can identify them, on average. But the variation *within* Tables 2(a)-2(c) is perhaps more interesting. First, compare the numbers in the diagonal of the table (shaded grey) with those that are located off the diagonal. We can see immediately that individuals are more likely to be correctly identified by members of their own group than by those that are not. In all three experiments, reading down each column, the number on the diagonal is consistently one of the highest. Statistically speaking, this pattern is strong: controlling both for the ethnic group of the viewer and the ethnic group of the individual being viewed, co-ethnicity increases the likelihood of correct identification by: 9 percentage points in Uganda, 13 percentage points in South

Africa, and 7 percentage points in the United States.[17] Suggesting that co-ethnicity helps most with successful identification when identification is particularly difficult (given the overall success rate of 28% in South Africa and 47% and 62% in Uganda and the United States, respectively).

There is much to be learned as well from those cases in which other groups seem to do particularly well. Consider the Bafumbira, who correctly identify other Bafumbira 55% of the time. But other groups do quite well also, most notably the Banyankole. As we saw in Table 1(a), the Banyankole are "close" to the Bafumbira in the sense that others often misidentify them as Bafumbira. It appears that not only do co-ethnics have an advantage in identifying members of their own groups, but in identifying individuals in groups with whom they are often confused.

To illustrate two key findings from Tables 2(a)-(c) consider the bar charts in Figures 1 and 2. Figure 1 plots for each country the probability of successful identification of each group and the overall total across groups (the "Total" row from Tables 2(a)-(c)). Figure 2 plots the rates of successful identification of co-ethnics (the principle diagonal from Tables 2(a)-(c)). Groups are ordered from largest to smallest in both figures. Figure 1(a) illustrates that in Uganda, it tends to be the largest and smallest groups that are more easily identified while the moderately-sized groups tend to be more difficult to identify. In South Africa (Figure 1(b)) we see that groups are correctly identified at nearly equal rates. Compared to Figures 1(a) and 1(c), it is clear that groups in South Africa are correctly identified at much lower rates than in Uganda or the United States. And finally, in the United States (Figure 1(c)), smaller groups are more difficult to identify, while Asians and Caucasians are readily identified. In Figure 2, we see different patterns of success when only looking at co-ethnics. In Uganda, with the exception of the Banyankole and Bagisu, there is a monotonic degress in successful identification as group size decreases (Figure 2(a)). Thus, smaller groups have a more difficult time identifying their co-ethnics. In South Africa (Figure 2(b)) and the United States (Figure 2(c)), there seems to be no correlation between groups size and successful identification. In both South Africa and the United States the smallest groups are highly successful in identifying co-ethnics especially compared to the next largest group. Overall, the six plots suggest that there is the most variation in successful identification in Uganda, success is relatively low in South Africa, and success rates are highest in the United States (with a bit of variation as well)

[17] Employing probit models with correct identification as the dependent variable and fixed effects for each ethnic group viewing and each ethnic group being viewed (and clustering over all the guesses of each player), the associated *z* statistic for the coefficient on co-ethnicity is 4.8 in Uganda, 12.2 in South Africa, and 3.4 in the United States.

Figure 1: Rate of Correct Guesses By Target Group

1(a) Uganda

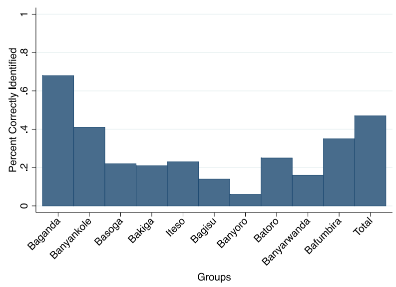

1(b) South Africa

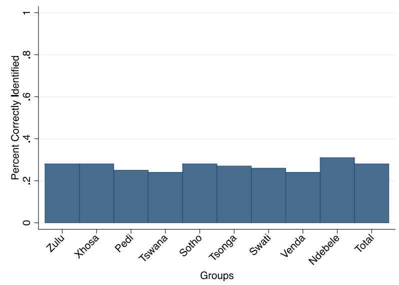

1(c) United States

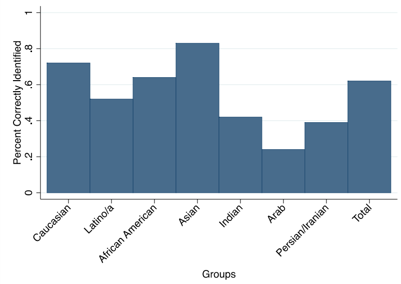

Figure 2: Rates of Correct Identification of Co-Ethnics

2(a) Uganda

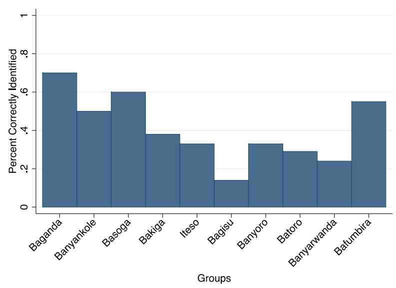

2(b) South Africa

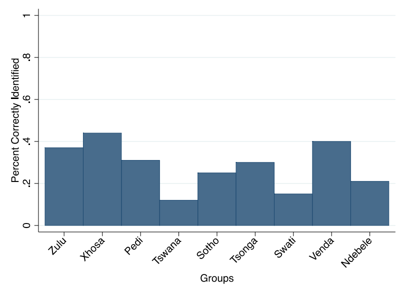

2(c) United States

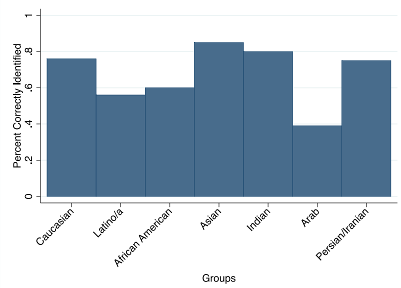

As we noted earlier, any assessment of the degree of correct identification depends on the informational environment. As one might expect, people are much more successful in the identification exercise as they gain access to more information about the people whose group membership they are trying to discern. Beyond the obvious impact of more information, it is natural to ask how different types of information affect the likelihood that individuals code others correctly. Recall that our information levels are not strictly ordered. As shown in the green line in Figures 3(a)-3(c), average levels of identification success move monotonically across the information levels.

In Figure 3(a) for Uganda, overall identification success is higher under information level 3 than under information level 2, indicating that the use of the primary language carries more information than the use of Luganda, the *lingua franca*. Identification success rates are similar at information levels 3 and 4, suggesting that the communication of names in Luganda has a similar impact to the use of a primary language. Success rates are highest at information level 5. This indicates that the information contained in one's name and one's primary language are complementary rather than being direct substitutes. In all, moving from the lowest information level (the headshot) to the highest (the video in which the subject greets the camera and provides his or her full name) leads to a near doubling of success rates, from an average of about 0.29 to 0.59. However, there is important variation across groups in the impact of information.

Figure 3: The Benefits of Information[18]

3(a): Uganda **3(b): South Africa** **3(c): The United States**

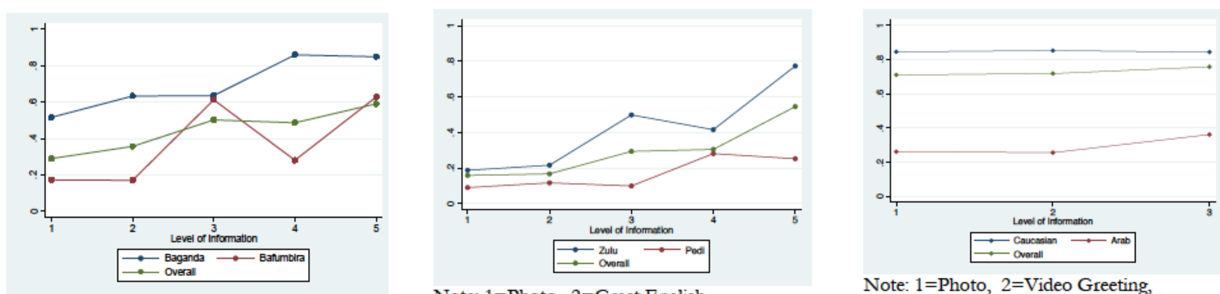

Note: 1=Photo, 2=Greet Luganda, 3=First name, 4=Surname, 5=Greet own language

Note: 1=Photo, 2=Greet English, 3=First name, 4=Surname, 5=Greet own language

Note: 1=Photo, 2=Video Greeting, 3=Video Greeting with Name

In Figure 3(b) for South Africa, overall identification success is higher as information conditions improve. Moving from level 1 to 2, the photograph to a video greeting, the success rate does not improve, but the remaining three levels all suggest higher success rates than the baseline photograph. Levels 3 and 4, though close to each other, are both higher than levels 1 and 2, indicating that giving one's first name and surname provides important evidence to the subjects. Similar to the Uganda result, greeting in one's own language in South Africa leads to the highest overall success rate.

In Figure 3(c) for the United States, there is not much evidence that more information improves the likelihood of successful identification. This could be due to the fact that the three types of information offered in the United States study shifted from a headshot to a brief greeting to the greeting with a name (which compares to the lower levels in Uganda and South Africa where the differences are not as pronounced). If we included additional conditions corresponding to levels 4 and 5 in the United States, it is possible that successful identification rates might improve.

4.2 SIGNS AND SIGNALS

Our data suggest that the ability to place others improves substantially with additional information. Is this still the case when the information available for making a judgment is subject to strategic manipulation by our subjects? In real world situations, where the stakes of misidentification are high, individuals often collect additional information about other peoples' backgrounds before they make decisions about how to assign

[18] In Uganda, the information provided in each condition is cumulative, so condition 3 is greet plus first name, condition 4 is greet, first name, and surname, etc.

people to ethnic categories. Although our experiments are one-sided (only the portrayer speaks in the videos, and no opportunity is provided for the subject to interrogate the portrayer), the simulation/dissimulation (SIMDIS) treatments described above provide an approximation of such real-world interactions and permit us to investigate the ability of individuals to pass as members of groups different from their own.

The main results of the SIMDIS treatments are presented in Tables 3(a)-(c), each divided into sub-tables for SIMDIS treatments in which portrayers were truth-telling versus those in which they were trying to pass as members of other groups. A number of the patterns observed previously obtain here. In Table 3(a)(i) for Uganda, co-ethnics are better able to place each other in the SIMDIS treatment. There are differences as well; for example, there is less variation in the extent to which individuals from different groups are placed successfully. The most striking feature, however, is the number in the bottom right hand corner: when portrayers engage in truth-telling, subjects guess correctly 63% of the time, a marked improvement from the results reported in Table 2(a) for all information conditions, which showed an overall success rate of 47%. This suggests that players are doing better, on average, in the truthful SIMDIS treatments than in the identification game, when portrayers are also telling the truth.

In fact, however, subjects perform significantly worse when portrayers are trying to pass. Table 3(a)(ii) shows an overall success rate of 23%. Given the prior distribution of groups in the population, a strategy of stating "Baganda" for all images would have produced a 42% success rate. In the SIMDIS treatments, however, players were told that truthful signals were provided with a 50% probability. A strategy of believing all signals and using no additional information would provide a 50% success rate. Thus, our subjects did no better than chance when we introduced dissimulation signals to the equation. (Appendix Table A5 reports a regression of successful passing on a host of covariates, providing additional contextual evidence.)

In Table 3(b)(i) for South Africa, as in Uganda, co-ethnics are better able to place each other in the truthful SIMDIS treatment. Again, the numbers in the bottom right hand corners (of 3(b)(i) and 3(b)(ii)) are instructive: when portrayers engage in truth-telling, subjects guess correctly 83% of the time, an even greater improvement relative to Uganda, given the baseline reported in Table 2(b) is 28%. This indicates that players are again doing better, on average, in the truthful SIMDIS treatment than in the identification game, when portrayers are also telling the truth. And, as with the Uganda results, subjects do worse when portrayers are trying to pass. Table 3(b)(ii) shows an overall success rate of 5%, extremely low and indeed much lower than if a subject simply guessed the percentage of the population and got it right half the time (because subjects knew that truthful/false signals were given 50% of the time each). Appendix Table A5 reports a regression of successful passing on a host of covariates, providing additional contextual evidence. (Appendix Table A10 reports a regression of successful passing on a host of covariates, providing additional contextual evidence.)

Table 3(a): Identification Success with Simulation and Dissimulation for Uganda
(Major groups only)

i. Simulation (Truth-telling)

	Banyankole	Baganda	Bagisu	Bakiga	Banyarwanda	Basoga	Batoro	Banyoro	Iteso	Bafumbira	Total
Banyankole	0.63	0.63	0.40	0.64	0.50	1.00	1.00	0.60	0.50	0.80	0.64
	(24)	(60)	(5)	(11)	(10)	(2)	(5)	(5)	(6)	(10)	(138)
Baganda	0.69	0.66	0.50	0.66	0.73	0.80	0.67	0.62	0.48	0.70	0.66
	(70)	(319)	(22)	(40)	(33)	(10)	(54)	(21)	(27)	(53)	(653)
Bagisu	0.64	0.60	.	0.00	1.00	1.00	1.00	0.67	0.50	1.00	0.68
	(11)	(20)		(1)	(2)	(3)	(4)	(3)	(2)	(1)	(47)
Bakiga	0.63	0.65	0.00	0.88	0.00	0.50	0.71	0.00	0.67	0.67	0.61
	(8)	(26)	(5)	(8)	(2)	(2)	(7)	(2)	(3)	(12)	(89)
Banyarwanda	0.88	0.59	0.50	1.00	0.33	0.75	0.75	1.00	1.00	1.00	0.70
	(8)	(32)	(2)	(6)	(6)	(4)	(8)	(1)	(2)	(4)	(73)
Basoga	0.75	0.55	0.00	0.75	0.80	1.00	0.50	.	1.00	0.50	0.58
	(8)	(49)	(3)	(4)	(5)	(1)	(6)		(1)	(6)	(83)
Batoro	0.75	0.61	0.17	1.00	.	1.00	0.75	0.00	0.00	0.83	0.63
	(16)	(31)	(6)	(1)		(2)	(4)	(1)	(1)	(6)	(68)
Banyoro	0.50	0.66	0.00	0.75	0.50	.	0.83	0.00	0.50	0.63	0.61
	(10)	(29)	(2)	(4)	(2)		(6)	(1)	(4)	(8)	(66)
Iteso	0.67	0.71	.	0.50	.	1.00	1.00	0.00	.	1.00	0.71
	(3)	(7)		(2)		(1)	(2)	(1)		(1)	(12)
Bafumbira	0.60	0.57	0.00	1.00	1.00	1.00	0.50	0.00	0.40	0.83	0.60
	(15)	(35)	(3)	(5)	(1)	(1)	(6)	(2)	(5)	(12)	(85)
Total	0.65	0.63	0.33	0.71	0.63	0.85	0.70	0.50	0.51	0.70	0.63
	(191)	(693)	(52)	(90)	(67)	(27)	(116)	(40)	(61)	(131)	(1,468)

ii. Dissimulation (Trying to pass)

	Banyankole	Baganda	Bagisu	Bakiga	Banyarwanda	Basoga	Batoro	Banyoro	Iteso	Bafumbira	Total
Banyankole	0.27	0.42	0.17	0.83	0.33	0.00	0.00	0.00	0.09	0.29	0.32
	(11)	(69)	(6)	(6)	(6)	(4)	(10)	(5)	(11)	(17)	(145)
Baganda	0.18	0.32	0.06	0.19	0.16	0.04	0.02	0.00	0.18	0.19	0.22
	(72)	(332)	(16)	(54)	(43)	(24)	(42)	(21)	(33)	(70)	(707)
Bagisu	0.00	0.15	0.00	1.00	.	.	0.00	0.00	1.00	0.00	0.15
	(2)	(20)	(2)	(1)			(3)	(2)	(1)	(2)	(33)
Bakiga	0.00	0.28	0.25	0.43	0.25	0.67	0.33	0.00	0.00	0.00	0.26
	(5)	(36)	(4)	(7)	(4)	(3)	(6)	(2)	(1)	(5)	(73)
Banyarwanda	0.00	0.19	0.00	0.00	0.00	0.00	0.00	.	0.20	0.10	0.11
	(6)	(32)	(2)	(5)	(6)	(2)	(3)		(5)	(10)	(71)
Basoga	0.33	0.19	0.00	0.00	0.00	.	0.00	0.00	0.33	0.25	0.18
	(9)	(26)	(3)	(2)	(3)		(6)	(3)	(6)	(8)	(66)
Batoro	0.00	0.48	0.00	0.33	0.00	0.00	0.50	.	0.00	0.29	0.29
	(6)	(25)	(2)	(3)	(8)	(1)	(4)		(2)	(7)	(58)
Banyoro	0.20	0.21	0.00	0.13	0.17	0.00	0.00	0.00	0.00	0.20	0.16
	(5)	(28)	(4)	(8)	(6)	(1)	(4)	(1)	(2)	(5)	(64)
Iteso	0.00	0.20	1.00	0.00	0.00	.	0.00	.	0.00	0.00	0.14
	(4)	(10)	(1)	(2)	(1)		(2)		(1)	(1)	(22)
Bafumbira	0.29	0.37	0.00	0.17	0.67	0.00	0.00	0.67	0.50	0.29	0.32
	(17)	(46)	(4)	(6)	(3)	(4)	(3)	(3)	(6)	(7)	(99)
Total	0.20	0.32	0.08	0.22	0.14	0.10	0.07	0.05	0.21	0.18	0.23
	(162)	(696)	(50)	(104)	(91)	(42)	(96)	(43)	(78)	(141)	(1,503)

Note: Cells in this table report the probability (as a percentage) that a typical player of the group denoted in each row will correctly identify a typical member of the column group, given the SIMDIS treatment, either truth telling or trying to pass. The number of experimental observations used to calculate these probabilities are given in parentheses underneath each entry.

Table 3(b): Identification Success with Simulation and Dissimulation Separated for South Africa

i. Simulation (Truth-telling)

	Ndebele	Pedi	Sotho	Swati	Tsonga	Tswana	Venda	Xhosa	Zulu	Total
Ndebele	0.67 (6)	0.81 (21)	0.79 (19)	1.00 (5)	1.00 (5)	0.90 (10)	0.83 (6)	0.90 (41)	0.76 (46)	0.83 (159)
Pedi	1.00 (4)	0.83 (35)	0.75 (28)	0.88 (8)	0.76 (17)	0.62 (21)	0.92 (12)	0.76 (62)	0.74 (80)	0.76 (267)
Sotho	0.83 (6)	0.77 (30)	0.83 (23)	0.71 (7)	0.75 (8)	0.64 (14)	0.80 (10)	0.77 (39)	0.80 (60)	0.78 (197)
Swati	0.25 (4)	0.87 (23)	0.89 (18)	1.00 (6)	0.71 (7)	0.91 (32)	1.00 (2)	0.78 (32)	0.82 (51)	0.83 (175)
Tsonga	0.78 (9)	0.83 (30)	0.90 (20)	1.00 (13)	0.94 (17)	0.86 (29)	0.67 (9)	0.84 (64)	0.87 (68)	0.86 (259)
Tswana	0.50 (2)	0.62 (21)	0.71 (7)	0.50 (4)	0.80 (5)	0.60 (5)	0.80 (5)	0.63 (19)	0.81 (31)	0.70 (99)
Venda	1.00 (2)	1.00 (4)	1.00 (2)	. (.)	1.00 (3)	0.63 (8)	1.00 (1)	1.00 (4)	0.91 (11)	0.89 (35)
Xhosa	0.78 (36)	0.86 (169)	0.86 (130)	0.83 (60)	0.77 (71)	0.82 (137)	0.88 (33)	0.89 (346)	0.89 (410)	0.86 (1,392)
Zulu	0.88 (43)	0.79 (189)	0.83 (132)	0.82 (77)	0.83 (83)	0.78 (181)	0.80 (49)	0.81 (380)	0.87 (438)	0.83 (1,572)
Total	0.80 (112)	0.81 (522)	0.84 (379)	0.84 (180)	0.81 (216)	0.79 (437)	0.83 (127)	0.84 (987)	0.86 (1,195)	0.83 (4,212)

ii. Dissimulation (Trying to pass)

	Ndebele	Pedi	Sotho	Swati	Tsonga	Tswana	Venda	Xhosa	Zulu	Total
Ndebele	0.50 (2)	0 (19)	0 (8)	0 (3)	0.09 (11)	0.06 (18)	0 (4)	0.08 (26)	0.04 (48)	0.05 (139)
Pedi	0 (10)	0.06 (33)	0.09 (23)	0.07 (15)	0.17 (12)	0.04 (28)	0 (6)	0.09 (67)	0.06 (86)	0.07 (280)
Sotho	0 (8)	0.07 (29)	0 (14)	0 (7)	0.09 (22)	0 (15)	0 (8)	0.10 (40)	0.04 (64)	0.05 (207)
Swati	0 (4)	0 (13)	0.04 (24)	0 (10)	0 (11)	0 (20)	0 (7)	0.03 (34)	0.08 (40)	0.03 (163)
Tsonga	0 (5)	0.06 (31)	0.20 (3)	0.22 (9)	0 (14)	0 (35)	0 (9)	0.02 (56)	0.06 (78)	0.06 (267)
Tswana	0 (2)	0 (13)	0 (11)	0 (7)	0 (8)	0 (7)	0.50 (2)	0.04 (23)	0.04 (27)	0.03 (100)
Venda	0 (1)	0 (9)	0 (1)	0 (1)	0 (2)	0 (2)	0 (2)	0 (6)	0.11 (19)	0.05 (43)
Xhosa	0.06 (32)	0.01 (160)	0.03 (128)	0.02 (58)	0.05 (77)	0.04 (135)	0.08 (26)	0.08 (303)	0.05 (414)	0.05 (1,333)
Zulu	0.05 (37)	0.03 (154)	0.07 (151)	0.10 (63)	0.10 (83)	0.01 (155)	0.03 (39)	0.08 (321)	0.06 (448)	0.06 (1,451)
Total	0.05 (101)	0.03 (461)	0.06 (390)	0.06 (173)	0.07 (240)	0.02 (415)	0.04 (103)	0.07 (876)	0.05 (1,224)	0.05 (4,199)

Table 3(c): Identification Success with Simulation and Dissimulation Separated for the United States

i. Simulation (Truth-telling)

Ethnic Group of Respondent (who views subject's image)	Ethnic Group of Subject (whose image is viewed): African American	Arab	Asian	Caucasian	Indian	Latino/a	Persian/ Iranian	Total
African American	50	25	100	100	100	83.33	100	88.37
	(2)	(4)	(13)	(11)	(2)	(6)	(5)	(43)
Arab		66.67	100	100	100	100	100	96.30
		(3)	(8)	(11)	(2)	(3)	(2)	(27)
Asian	80	25	97.37	78.57	100	92.86	77.78	86
	(5)	(4)	(38)	(28)	(3)	(12)	(9)	(100)
Caucasian	81.82	66.67	100	96.67	100	86.39	100	93.55
	(11)	(9)	(38)	(60)	(2)	(22)	(12)	(155)
Indian		40	100	100	100	100	100	83.33
		(5)	(5)	(3)	(2)	(2)	(1)	(18)
Latino/a	66.67	33.33	90.91	72.73	100	81.25	100	80.77
	(3)	(3)	(11)	(11)	(2)	(16)	(6)	(52)
Persian/ Iranian	60		100	75	100	75	100	80.65
	(5)		(6)	(12)	(1)	(4)	(3)	(31)
Total	73.08	46.43	98.32	89.71	100	86.57	94.74	88.50
	(26)	(28)	(119)	(136)	(12)	(67)	(38)	(426)

ii. Dissimulation (Trying to pass)

Ethnic Group of Respondent (who views subject's image)	Ethnic Group of Subject (who is trying to pass as a member of a different group): African American	Arab	Asian	Caucasian	Indian	Latino/a	Persian/ Iranian	Total
African American	40	0	88.89	50		0	42.86	45.45
	(5)	(3)	(9)	(14)		(6)	(7)	(44)
Arab	33.33	25	87.50	28.57	0	50	100	46.43
	(3)	(4)	(8)	(7)	(3)	(2)	(1)	(28)
Asian	40	14.29	87.18	41.67	66.67	18.18	28.57	55.21
	(5)	(7)	(39)	(24)	(3)	(11)	(7)	(96)
Caucasian	50	22.22	84.62	65	0	25	28.57	55.41
	(6)	(9)	(39)	(60)	(5)	(24)	(14)	(157)
Indian		0	100	57.14	0	16.67	100	54.17
		(0)	(6)	(7)	(0)	(6)	(2)	(24)
Latino/a	66.67	0	80	62.50		35.71	66.67	57.97
	(3)	(5)	(20)	(24)		(14)	(3)	(69)
Persian/ Iranian	66.67	33.33	92.31	37.50	50	50	85.71	67.50
	(3)	(3)	(13)	(8)	(2)	(4)	(7)	(40)
Total	48	15.15	86.57	55.56	21.43	25.37	48.78	55.24
	(25)	(33)	(134)	(144)	(14)	(67)	(41)	(458)

Table 3(c)(i), reporting results for the United States experiment, reveals that rates of identification success reach their highest levels for viewings in which the subject sees a truthful simulation video.[19] The

[19] This holds across all ethnic groups except for African Americans. Again, because one key subject in the African American sample was light-skinned, many African American respondents, along with others in the general population,

identification success rate across all viewings reaches 88.5 percent in the truthful simulation sub-sample, exceeding the success rate at even the highest previous levels of information (the video in which the subject provided his or her name), where the success rate was 75.6 percent. Providing subjects with the opportunity to convince respondents of their true ethnic background has a particularly powerful impact on raising identification success rates for those groups that were most difficult to identify at lower levels of information. But providing subjects with the opportunity to *mislead* respondents about their ethnic backgrounds substantially reduces identification success rates. As Table 3(c)(ii) indicates, the overall success rate drops to 55.2 percent with dissimulation. As with identifiability, the ability to pass is something that varies with the ethnic group of the subject. Arabs, Caucasians, Indians, and Latinos are particularly good at passing, reducing the ability of respondents to correctly identify them by between 27 and 39 percentage points. (Appendix Table A14 reports a regression of successful passing on a host of covariates, providing additional contextual evidence.)

5 THEORETICAL AND EMPIRICAL EXTENSIONS

The primary focus of the article is to explore rates of successful identification and to better understand how individuals signal their ethnicity. In addition to these core contributions, we also briefly explore here errors of exclusion and inclusion and group distinctness. We focus our empirical discussion on Uganda only; data and discussion of the other cases can be found in the appendix.

Errors of Exclusion and Inclusion

We can also generate measures that describe the types of errors that individuals make in classification. Consider the question of whether a given individual is an in-group member. We refer to the hypothesis that player *i* is an in-group member for a player *j* as *j*'s ***in-group null*** regarding *i*. Given an in-group null, we say that an **error of exclusion** occurs if *j* incorrectly identifies in-group member *i* as an out-group member. We say that an **error of inclusion** occurs if player *i* is an out-group member who is incorrectly identified as an in-group member. The probability of *not* making an error of exclusion, is simply the identifiability of (in-group member) *i*, for *j*.

Group Distinctness

Finally, we define the ***distinctness*** of groups *A* and *B* for an arbitrary member *j* of group *C*, $D_C(A,B)$, as the (expected) probability that *j* will identify an arbitrary individual from group *A* as a member of group *A,* given that *j* classifies the individual as an *A* or a *B*, plus the (expected) probability that *j* will identify an arbitrary

assumed that the subject was trying to pass as an African American even when the subject provided an explanation for her appearance in the simulation video.

individual from group B as a member of group B, given that j classifies the individual as an A or a B minus 1.[20] The measure ranges from 1 to −1. If members of group A are never confused with each other, then $D_C(A,B)=1$. If guesses are independent of the true identity of the subjects (for example if placement is decided by a coin toss), then $D_C(A,B)=0$. Values between 0 and −1 may arise if players do worse than chance. The score is undefined if players always classify members of A or B as members of groups other than A or B. Note that two groups may not be highly distinct yet individuals from group A may still find particular individuals from B highly identifiable. We illustrate the notion of distinctness and the distinction between group distinctness and individual identifiability with an example.[21]

Consider the case where player i is drawn from group A with probability p, and from group B with probability $1-p$. Assume that A types all possess some characteristic θ, that is held by B types only with probability q. We let X denote the treatment wherein some viewer j observes whether or not i possesses θ. Let j's private placement rule be given by $c_j(i,X) = A$ if and only if j's posterior assessment that i is a member of group A exceeds .5. Using Bayes' rule, j's posterior is given by: $\Pr(i \in A \mid \theta) = p/(p + q(1-p))$ and $\Pr(i \in A \mid 0) = 0$.

What then, is the distinctness of A and B for i under informational treatment X? There are two cases of interest. In the first case $\Pr(i \in A \mid \theta) \leq .5$. In this case, no matter what signal is received, i believes j to be a member of group B. In this case, $D(A,B) = 1 + 0 - 1 = 0$. The signal is too weak relative to the prior to allow the player to distinguish between the groups. In the second case, $\Pr(i \in A \mid \theta) > .5$, and so, upon observing θ, j infers that i is in group A. In this case if i is an A, he will never be classified as a B; however if he is a B then he will be classified as an A with probability q. Therefore $D(A,B) = 1+(1-q) - 1 = 1 - q$. The groups are perfectly distinct only if $q = 0$. Two features of this example are especially noteworthy. The first is that even if distinctness is imperfect, individual members of group B may be perfectly identifiable whenever θ is not observed (but not when it is). The second is that the notion of distinctness is context- rather than simply attribute-dependent, it depends explicitly on an individual's priors; thus A and B, though "objectively" identical may yield different distinctness measures depending on the social demography.

5.1 Errors in Classification and Group Distinctness

We now consider the pattern of errors produced by subjects in our sample. In Table 4, each cell reports the likelihood that an individual of a row type will code someone of a column type as a co-ethnic. If subjects identified the people whose images they were shown perfectly, this matrix would be the identity matrix: each

[20] Although Casselli and Coleman (2002) base a theoretical model of conflict on a similar notion of distinctness, we know of no attempts to measure distinctness in this way empirically.

[21] It is of interest to note that this definition of distinctness is equivalent to the *determinant* of the 2×2 matrix P in which each cell entry, p_{ij}, denotes the probability with which a row type is classified as a column type given that it is classified

cell on the diagonal would be one and all the off-diagonal cells would be zero. In some ways, the matrix is close to the identity matrix. The diagonal is "strong" in the sense that the number on the diagonal is generally the largest number in each row and in each column. This means that a subject from a given group is more likely to say that a co-ethnic is a co-ethnic than is a subject from any other group, and that a subject is most likely to say that a portrayer is a co-ethnic when indeed that portrayer is a co-ethnic.

Table 4 Patterns of Errors: In-Group Coding in Uganda

	Banyankole	Baganda	Bagisu	Bakiga	Banyarwanda	Basoga	Batoro	Bunyoro	Iteso	Bafumbira	Errors of Inclusion
Banyankole	0.50	0.10	0.19	0.44	0.14	0.10	0.23	0.19	0.08	0.14	0.13
Baganda	0.14	0.70	0.27	0.17	0.32	0.46	0.22	0.35	0.25	0.16	0.20
Bagisu	0.00	0.01	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Bakiga	0.20	0.02	0.00	0.38	0.09	0.07	0.00	0.10	0.00	0.05	0.04
Banyarwanda	0.05	0.04	0.00	0.06	0.24	0.07	0.09	0.15	0.00	0.27	0.06
Basoga	0.02	0.05	0.25	0.00	0.00	0.60	0.03	0.00	0.08	0.08	0.05
Batoro	0.07	0.02	0.11	0.17	0.00	0.03	0.29	0.17	0.04	0.04	0.04
Bunyoro	0.00	0.00	0.11	0.00	0.00	0.00	0.05	0.33	0.00	0.00	0.01
Iteso	0.00	0.02	0.10	0.00	0.05	0.00	0.00	0.07	0.33	0.04	0.02
Bafumbira	0.06	0.03	0.08	0.02	0.16	0.00	0.02	0.08	0.04	0.55	0.04

Note: This table reports the likelihood that a row type will classify a column type as a co-ethnic. The final column reports the overall frequency with which people of a row type incorrectly classify out-group members as in-group members.

However, in some ways, the matrix is quite unlike the identity matrix. The diagonal elements are often lower than one, sometimes substantially so. This means that individuals *often* fail to recognize co-ethnics as co-ethnics, they make "errors of exclusion" (these can be computed as 1 minus the shaded value in the diagonal). The Baganda, for example, make this mistake 30% of the time; the Banyarwanda commit errors of exclusion 76% of the time. Across our sample, players incorrectly code in-group members as out-group members about one-third of the time. Mistakes of this form, we find, are most common among smaller ethnic groups. Larger groups have higher prior beliefs about the probability of encountering members of their own group; for smaller groups, encountering a co-ethnic is a rare event. As a result, the possibility that someone is a co-ethnic is discounted, even if he is in fact a co-ethnic.[22]

The matrix also differs from the identity matrix because of its off-diagonal elements. A positive number in an off-diagonal cell means that people are coding out-group members as in-group members. For example, in Uganda Banyankole classify Bakiga as Banyankole about 44% of the time, only slightly less than they do other Banyankole. Other groups that are closely related also make such errors of inclusion. In

among some column type. In this case the determinant is given by $p_{11}p_{22} - p_{12}p_{21} = p_{11}p_{22} - (1-p_{11})(1-p_{22}) = p_{11}+p_{22} - 1$. Thus the distinctness measure is a natural measure of the extent to which this matrix diverges from the identity matrix.

[22] The likelihood of making these errors also depends on the information available. Based only on the information contained in a headshot, the overall likelihood of committing an error of exclusion is 75%, these errors drop to just 15% of viewings at the highest level of information.

contrast with errors of exclusion, the incidence of these miscodings is relatively low (see the last column of Table 4). Overall, they occur just 11% of the time in the Uganda sample. Errors of inclusion tend to be most common among larger groups: the Baganda miscode out-group members as co-ethnics approximately one in every five viewings. Among smaller groups, the incidence is much lower: Bafumbira or Banyarwanda, for example, make this mistake in only 5% of their viewings. More generally, it is clear that the nature of boundary errors is correlated with group size. Whereas errors of exclusion are more common among smaller groups, errors of inclusion are much more common among larger groups. Again, as we would expect, the incidence of these errors declines substantially with increases in information.

5.2 Group Distinctness

In the previous section, we focused on how our subjects miscode one another because the errors subjects make form the basis for an empirical assessment of the distinctness of groups. For each of the pair-wise relationships among groups (i.e. Banyankole-Baganda, Banyankole-Banyankole, and so on), we construct the frequency with which subjects in our sample miscode members of dyads into each other's groups (conditional upon coding into this pair of categories) to create the distinctness measure as described earlier (with our entire sample playing the role of group *C*). The results of this exercise are presented in Table 5.

Table 5 Group Distinctness in Uganda

	Banyankole	Baganda	Bagisu	Bakiga	Banyarwanda	Basoga	Batoro	Banyoro	Iteso	Bafumbira
Banyankole	-									
Baganda	0.60	-								
Bagisu	0.53	**0.30**	-							
Bakiga	**0.14**	0.48	0.77	-						
Banyarwanda	**0.36**	**0.28**	0.80	0.56	-					
Basoga	0.62	**0.27**	0.51	0.75	0.79	-				
Batoro	**0.32**	0.44	0.69	0.56	0.55	0.69	-			
Banyoro	**0.20**	**0.13**	0.79	0.44	0.48	0.44	**0.16**	-		
Iteso	0.79	0.47	0.77	0.83	0.90	0.70	0.86	0.80	-	
Bafumbira	0.47	0.62	0.69	0.58	**0.31**	0.76	0.66	0.59	0.75	-

Note: Entries in the table show the empirical distinctness of row and column types based on the likelihood that a random member of the sample mistakenly classifies a column type as a row type and a row type as a column type. A score of 1 indicates that a row type is never confused for a column type. A score of 0 indicates that placing is no better than chance.

In Table 5, the lowest scores in Uganda are for clusters of groups that share regional connections. The three main eastern groups, the Bagisu, the Basoga and the Iteso are largely distinct from all other groups and from one another; in each case they are most likely to be confused with the dominant group the Baganda, but otherwise have distinctness scores about 0.5 in almost all cases. The situation is different among western groups, however. The Bakiga and the Banyankole are consistently confused for each other (even by their own members). So, too, are the Banyankole and the Banyarwanda, the Banyarwanda and the Bafumbira, and the Batoro and Banyoro. These all have distinctness scores below 0.2. The score of 0.14 for the Banyankole-Bakiga pairing corresponds to a situation where if choosing among the two an individual would code a Banyannkole as a Bakiga and vice versa 57% of time, little better than chance. These measures of distinctness clearly pass a test of face-validity; they quantify, from the aggregation of binary actions, relationships among groups that correspond to historical and geographic patterns.

6 CONCLUSION: IMPLICATIONS FOR THEORIES OF ETHNIC POLITICS

Despite the theoretical complexities associated with the notions of group membership and the "correct" placement of individuals into socially constructed categories, we argue that it is possible to develop a meaningful notion of ethnic identification and group distinctness, based on binary relations between individuals. Using the theoretical apparatus we develop to achieve this end and assigning a privileged position to the self-placement criterion for fixing a social demography, we document that, contrary to the assumption of easy ethnic identifiability implicit or explicit in much micro-level theoretical work on ethnic politics, there is substantial variability both within and between groups in the identifiability of members. Beyond describing the existence of substantial variation, we demonstrate that there is a structure to the successes and errors that people make, and we argue that these systematic errors can be important to models of collective action, at least in contexts where identification based on limited information is germane.

The structure of errors in the identification game allows us to generate a new measure of the distinctness of groups and which we use to help account for variation in rates of passing. In general, we find, individuals are weak at incorporating new information from verbal signals sent by possible dissimulators. The net effect of introducing signals is a decline in identifiability, contrary to expectations from rational updating. Even still, some types of passing are easier than others. Individuals find it much easier to pass when they come from groups that are more difficult to identify on the basis of signs alone and when they attempt to pass into groups similar to their own but distinct from the group of the observer.

Such variation has substantive implications for collective action within and between groups. It suggests one reason why collective action may be easier for some ethnic groups than for others. If identifiability is imperfect, then the ability of groups to police their members will be weakened and the advantage that ethnic groups have for collective action, particularly in anonymous environments, will disappear. To the extent that identifiability varies systematically across ethnic groups, the ability of ethnic groups to achieve collective ends should vary as well. Furthermore, the results suggest that discrimination against some groups may be easier than against others. The fact that *passing* may be easier for some individuals and groups than for others has implications for the permeability of group boundaries and the ability of groups to police them, and for the collective action benefits that come from such policing. All of these features depend on the level of information available to individuals. Information, however, has distinct implications for different types of dyadic interaction. In particular, *the costs and benefits involved in gathering* *information about the ethnic identity of individuals may vary across groups*, with information adding especially in the case of in-group out-group identifications and when subjects are from difficult-to-identify groups.

For scholars, this research has implications for both empirical and theoretical work. For empirical work, this study serves as a demonstration of a methodology for collecting rich information about one important aspect of ethnic structures. Cross-national studies of the implications of ethnic structures on various political and economic outcomes have been hampered by the weakness of data on ethnicity. Much of the data that exists relies on a reified notion of what groups are, but also fails to capture any aspects of ethnic structures beyond the number and relative sizes of groups. Our measures of identifiability and the distinctness of groups provide a richer class of measures of group structure. The methodology employed here also suggests a way of taking seriously the notion that categorization depends on context and moving from there to generating measures of the impact of context on categories.

The arguments developed in this paper also have implications for theoretical work. While the assumption that the ethnic backgrounds of individuals are readily apparent has facilitated theoretical analysis, it has also limited the study of important aspects of ethnic processes that should no longer be overlooked. Variation in identifiability may have real consequences, as the account of ethnic identification failures in Rwanda, South Africa, Sri Lanka, and Burundi provided earlier suggests. Taking these differences seriously—and incorporating them into models of ethnic politics—is a critical next step in producing better theories that link ethnicity to cooperation and conflict.

References

Appiah, Anthony Kwame. 1992.. *In My Father's House*. New York: Oxford University Press.

Akerlof, George. 1970.

Akerlof, George. 1976. "The Economics of Caste and of the Rat Race and Other Woeful Tales." *Quarterly Journal of Economics* 90: 599-617.

Alesina, Alberto, Reza Baqir, and William Easterly. 1999. Public goods and ethnic divisions. *Quarterly Journal of Economics* 114 (November): 1243–84.

Alesina, Alberto, and Eliana LaFerrara. 2005. Ethnic diversity and economic performance. *Journal of Economic Literature* 63 (September): 762–800.

Allport, Gordon and Bernard Kramer. 1946. "Some Roots of Prejudice." *The Journal of Psychology* 22: 9-39.

Anderson, Benedict. 1991. *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. Revised Edition ed. London and New York: Verso.

Appiah. 1992.

Bacharach, Michael and Diego Gambetta. 2001. "Trust in Signs." In Karen Cook (ed.) *Trust and Society*. New York: Russell Sage Foundation.

Bahry, Donna and Rick Wilson. Forthcoming. "Confusion or Fairness in the Field? Rejections in the Ultimatum Game under the Strategy Method." *Journal of Economic Behavior and Organization*.

Barth, Fredrik. 1969. *Ethnic Groups and Boundaries*. Boston: Little Brown.

Benjamin, Daniel and Jesse Shapiro. 2006. "Thin-Slice Forecasts of Gubernatorial Elections" Working Paper.

Blasovich, Jim, Natalie Wyer, Laura Swart, and Jeffrey Kibler. 1997. "Racism and Racial Categorization." *Journal of Personality and Social Psychology* 72: 1364-1372.

Becker, Gary S. 1971. *The Economics of Discrimination.* Chicago: University of Chicago Press.

Caselli, Francseco and Wilbur John Coleman. 2002. "On the Theory of Ethnic Conflict." Typescript. Harvard University.

Chandra, Kanchan. 2001 "Cumulative Findings in the Study of Ethnic Politics." APSA-CP, Vol. 12 (1), 2001: 7-11.

Chandra, Kanchan. 2004. *Why Ethnic Parties Succeed: Patronage and Ethnic Headcounts in India.* New York: Cambridge University Press.

Chandra, Kanchan. 2006. "What is Ethnic Identity and Does it Matter?"*Annual Review of Political Science* 9: 397-424.

Chandra, Kanchan. 2012. *Constructivist Theories of Ethnic Politics.* Oxford University Press.

Charness, Gary and Uri Gneezy. Forthcoming. "What's in a Name? Anonymity and Social Distance in Dictator and Ultimatum games." *Journal of Economic Behavior and Organization*.

Deaux, Kay, Anne Reid, Daniela Martin, and Nida Bikman. 2006. Ideologies of diversity and inequality: Predicting collective action in groups varying in ethnicity and immigrant status. *Political Psychology* 27, no. 1: 23-46.

Diamond, Larry. 1987. Ethnic conflict. *The Journal of Modern African Studies* 25 no. 1 (March): 117-28.

Fearon, James D. 1999. "Why Ethnic Politics and 'Pork' Tend to Go Together." Typescript. Stanford University.

Fearon, James D. and David D. Laitin. 1996. "Explaining Interethnic Cooperation." *American Political Science Review* 90, 4: 715-35.

Fearon, James D., and David D. Laitin. 2000. Violence and the social construction of ethnic identity. *International Organization* 54, 4: 845-877.

Fershtman, Chaim and Uri Gneezy. 2001. "Discrimination in a Segmented Society: An Experimental Approach." *Quarterly Journal of Economic*s 116, 1: 351-377.

Fryer, Roland and Matthew O. Jackson. 2008. A Categorical Model of Cognition and Biased Decision Making." The B.E. Journal of Theoretical Economics 8(1):1-42.

Gambetta, Diego. 2005. "Deceptive Mimicry in Humans." In S. Hurley and N. Chater (eds.) *Perspective on Imitation: From Cognitive Neuroscience to Social Science*. Cambridge: MIT Press: pp. 221-241

Gourevitch, Philip. 1999. *We Wish to Inform You that Tomorrow We Will Be Killed With Our Families: Stories From Rwanda.* NY: Picador Books.

Greif, Avner. 1989. "Reputation and Coalitions in Medieval Trade." *Journal of Economic History* 49 (December): 857-882.

Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein. 2007. Why does ethnic diversity undermine public goods provision? *American Political Science Review* 101, no. 4: 709-25.

Harris, David. 2002. "In the Eye of the Beholder: Observed Race and Observer Characteristics." PSC Research Report No. 02-522. Population Studies Center. University of Michigan.

Hassim, Shireen, Tawana Kupe, and Eric Worby. (Eds.) 2008. *Go Home or Die Here: Violence, Xenophobia, and the Reinvention of Difference in South Africa.* Johannesburg: Wits University Press.

Horowitz, Donald L. 1985. *Ethnic Groups in Conflict*. Berkley, CA: University of California Press.

Horowitz, Donald L. 2001. *The Deadly Ethnic Riot.* Berkeley and Los Angeles: University of California Press.

Human Rights Watch. 2002. "'We Are Not the Enemy': Hate Crimes Against Arabs, Muslims, and Those Perceived to be Arab or Muslim After September 11." 14, 6 (November).

Humphreys, Macartan, Daniel Posner, and Jeremy M. Weinstein. 2002. Ethnic identity, collective action, and conflict: An experimental approach. Paper prepared for Presentation at APSA, Boston, September 2002.

Isaacs, Harold. 1975. "Idols of the Tribe." In Nathan Glazer and Daniel P. Moynihan. *Ethnicity: Theory and Experience*. Cambridge, MA: Harvard University Press: 29-52.

Kosic, Ankica, and Karen Phalet. 2006. Ethnic categorization of immigrants: The role of prejudice, perceived acculturation strategies and group size. *International Journal of Intercultural Relations* 30, no. 6 (November): 769-82.

Jung. 2000.

Kuruvilla, Matthai Chakko. 2006. "9/11 Five Years Later: Typecasting Muslims as a Race." *San Francisco Chronicle*, 3 September.

Laitin, David D. 1995. "Marginality: A Microperspective." *Rationality and Society* 7 (January): 31-57.

Landa. 1994.

Nosofsky, Robert M. 1986. \Attention, Similarity and the Identi_cation-Categorization Relationship." Journal of Experimental Psychology 115(1):39-57.

National Catholic Reporter. 2002.

Posner, Daniel N. 2004. "Measuring Ethnic Fractionalization in Africa." *American Journal of Political Science* 48 (October): 849-863.

Posner. 2005.

Sambanis, Nicholas. 2001. Do ethnic and non-ethnic civil wars have the same causes?: A theoretical and empirical inquiry (part I). *Journal of Conflict Resolution* 45, no.3: 259-82.

Secord, Paul. 1959. "Stereotyping and Favorableness in the Perception of Negro Faces." *Journal of Abnormal Social Psychology* 59: 309-315.

Tajfel, H., Billig, M. G., Bundy, R. P., and Flament, C. 1971. "Social Categorization and Intergroup Behavior." *European Journal of Social Psychology* 1: 149-178.

Terrry, Deborah J., Michael A. Hogg and Katherine M. White. 1999. \The Theory of Planned Behavior: Self-Identity, Social Identity, and Group Norms." British Journal of Social Psychology 38:225-244.

Weber, Max. 1968. *Economy and Society.* University of California Press.