

Ethnically Biased?

Experimental Evidence from Kenya*

Lars Ivar Oppedal Berge¹, Kjetil Bjorvatn¹, Simon Galle²,
Edward Miguel³, Daniel Posner⁴,
Bertil Tungodden¹ and Kelly Zhang⁵

August 2017

Abstract

Ethnicity has been shown to shape political, social, and economic behavior in Africa, but the underlying mechanisms remain contested. We utilize lab experiments to isolate one mechanism—an individual’s bias in favor of coethnics and against non-coethnics—that has been central in both theory and in the conventional wisdom about the impact of ethnicity. We employ an unusually rich research design involving a large sample of 1,300 subjects from Nairobi, Kenya; the collection of multiple rounds of experimental data with varying proximity to national elections; within-lab priming conditions; both standard and novel experimental measures of coethnic bias; and an implicit association test (IAT). Our tests find very little evidence of explicit or implicit coethnic bias in the behavioral experiments and IAT. These results run against the common presumption of extensive coethnic bias among ordinary Africans and suggest that mechanisms other than coethnic bias must account for the strong associations we see in the region between ethnicity and political, social and economic outcomes.

* We thank Ana Beatriz Aguilar, Felipe Gonzalez, Ranveig Falch, Ruth Carlitz, and Dounia Saeme for excellent research assistance, and the Busara Center for Behavioral Economics, the Economics and Social Research Foundation (ESRF), and The Choice Lab for administrative support. We are grateful to Obie Porteus, Gwyneth McClendon, and Jessica Leight for sharing data. We are grateful for funding from the Governance Initiative at JPAL-MIT, the Peder Saether Foundation, and the Norwegian Research Council. Galle acknowledges support from U.C. Berkeley IBER, Zhang acknowledges support from the NSF Graduate Research Fellowship Grant No. DGE-114747. We have received extremely useful comments from Michal Bauer, Julie Chytilova, Guy Grossman, Johannes Haushofer, Jonas Hjort, Dean Karlan, Steve Krasner, David Laitin, Rohini Pande, Amanda Robinson, Raul Sanchez de la Sierra, and seminar audiences at the Norwegian School of Economics, U.C. Berkeley, Caltech, Harvard, UCLA, USC, Stanford, MIT J-PAL, Columbia, University of Washington, Princeton, Yale, U.C. Merced, Stockholm School of Economics, the Annual World Bank Conference on Africa, the Barcelona GSE Summer Forum, CEGA, and the Pacific Development Conference.

Author affiliations: ¹ Norwegian School of Economics, ² BI Norwegian Business School, Dept. of Economics, ³ University of California, Berkeley, Dept. of Economics, ⁴ UCLA Dept. of Political Science, ⁵ Stanford University Dept. of Political Science.

Contact information: K. Bjorvatn , Kjetil.Bjorvatn@nhh.no.

1. Introduction

Ethnicity has been shown to affect a range of political, social, and economic outcomes in Africa, from voting patterns (Horowitz 1985; Posner 2005; Ichino and Nathan 2013; Conroy-Krutz 2013) to trade and investment decisions (Fafchamps 2000; Robinson 2016) to workplace productivity (Hjort 2014) to public goods provision (Miguel and Gugerty 2005) to favoritism in the distribution of jobs and development resources (Collier and Garg 1999; Burgess et al. 2015; Kramon and Posner 2016).

While the role of ethnicity in shaping politics and policy in Africa is undisputed, the underlying mechanisms remain contested. In particular, the role played by differences in preferences across ethnic groups—either in terms of direct coethnic favoritism, antipathy towards others, or more subtle variation in tastes for certain outcomes (Alesina, Baqir and Easterly 1999; Alesina and LaFerrara 2005; Vigdor 2002)—is controversial. Explanations based on *ethnic preferences* contrast with theories that emphasize different technologies of cooperation within groups due to superior ability to communicate or to impose social sanctions (Miguel and Gugerty 2005), different norms of behavior in within-group and cross-group interactions (Habyarimana et al. 2007), or other *institutional considerations* (e.g. Fearon 1999, Spolaore and Wacziarg 2009; Robinson 2016).

This paper assesses the degree of coethnic bias in preferences, defined as a preference for discriminating in favor of coethnics and against non-coethnics (Becker 1957; Hjort 2014; Horowitz 1985; Vigdor 2002). The investigation is carried out in Kenya, an African country with well-documented and politically salient ethnic divisions. Presidential vote shares largely fell along ethnic lines during the disputed 2007 national elections; in the aftermath, more than one thousand people died and hundreds of thousands were displaced from politically motivated ethnic violence (Gibson and Long 2009). Moreover, there is a history of ethnic bias in policies, including

large-scale distortions in public roads investment favoring the president's ethnic group (Barkan and Chege 1989, Burgess et al. 2015).

To isolate the presence of coethnic bias from other potential mechanisms, we employ lab experiments in an unusually rich research design that involves multiple rounds of experimental data and a large sample of over 1,300 subjects. We measure coethnic bias using both standard experimental games (e.g., Dictator and Public-good games) and a more novel lab activity (the Choose-Your-Dictator game), which captures expectations about the biases of others. In our theoretical model, we explain how each game captures a different aspect of coethnic bias. Then, as a further test, we employ an implicit association test (IAT) to measure underlying bias free from experimenter demand effects.

In addition to measuring the strength of coethnic bias, we explore whether, as an influential literature suggests it might (Bates 1983; Eifert et al. 2010; Hjort 2014), such bias varies with proximity to political competition. To address this issue, we draw our sample from two different lab sessions: one scheduled immediately before Kenya's 2013 general elections, and the other roughly seven months earlier. We also supplement this variation in real-world timing with within-lab priming designed to increase the situational salience of political competition and various dimensions of ethnicity.

Given the frequency with which outcomes in Africa are linked to ethnic preferences, our results are striking. Most of our tests yield no evidence of meaningful coethnic bias. This lack of coethnic bias holds across multiple experimental measures and well-powered statistical tests. Figure 1 summarizes evidence from the Dictator game and the Public-good game. In both cases, the average level of contribution to coethnic and non-coethnic partners is within one percentage point: 35.6% versus 35.4% in the Dictator game (Panel A), and 46.2% versus 46.4% in the Public-good game (Panel C). Participants' beliefs about the contributions of others in these games are similarly unaffected by their ethnic backgrounds, at 49.3% to 48.4% (Panel B) and 53.9% to 53.1%

(Panel D), respectively. Even though our sample is unusually large, none of these differences are statistically distinguishable from zero. In the IAT, the average bias against members of other ethnic groups is just 0.079 standard deviation units, roughly one sixth of the average bias shown by U.S. whites against blacks (Nosek et al. 2007).

The precisely estimated null effects are robust: they hold across all demographic subgroups (including gender, ethnic group, and education level), and across both the lab sessions held close to the national election and those conducted 7-8 months earlier. Furthermore, the null findings persist across a wide range of randomly assigned priming treatments, which were specifically designed to dampen or amplify the situational salience of politicized ethnicity. The core null finding is also not the result of a selective presentation of results on our part: the econometric approach was pre-specified and, in a novel “selective presentation test,” we compare the distribution of statistical significance levels for results contained in the main tables and in the full set of pre-specified results.¹

It is worth being clear that these results should not be taken to suggest that ethnicity plays an unimportant role in our study setting: ethnicity remains a prominent feature of contemporary life in Nairobi—as in Kenya more generally and in many parts of Africa. This is also evident from survey data collected at the end of our experiment, where the large majority of respondents reported they would support their own coethnic frontrunner in the upcoming elections.² Our findings simply suggest that the salience of ethnicity is likely due to mechanisms other than coethnic bias in preferences.

Indeed, notwithstanding the centrality of coethnic bias *in preferences* in both the scientific and broader public debates, there is considerable empirical evidence in favor of other, *institutional* channels driving ethnic divisions in African cases (Miguel and

¹ The full results for the pre-specified analysis are available in Appendix D.

² For instance, approximately 84% of ethnic Luos reported that they would support their coethnic frontrunner Raila Odinga (against only 8% of ethnic Kikuyus), while 65% of Kikuyus stated that they would support their coethnic frontrunner Uhuru Kenyatta (versus just 1% of Luos).

Gugerty 2005; Burgess et al. 2015; Habyarimana et al. 2007, 2009). Moreover, studies that focus on coethnic bias in preferences present mixed results (Carlson 2015; Michelitch 2015; Dionne 2015; Hjort 2014 and Marx et al. 2016). The current study contributes to this emerging literature along multiple dimensions. First, we examine coethnic bias across a range of experimental measures, using both standard games (Dictator and Public-good game), a novel Choose-Your-Dictator game, and an IAT. Second, in previous studies on coethnic bias in preferences (e.g. Habyarimana et al. 2007) it remains unclear if results would differ if experiments had been conducted in closer proximity to national elections, when politicians are the most likely to mobilize along ethnic cleavages and ethnicity has been shown to acquire greater salience (Eifert et al. 2010, Hjort 2014). To address this question, we amplify the situational salience of ethnicity and political competition both across lab rounds, by varying the proximity to Kenya's national election, and within the lab using priming treatments. Finally, to avoid selective presentation of results, our analysis follows a pre-specified set of hypotheses, and we employ larger samples to ensure sufficient statistical power.

Our results thus serve to crystallize and generalize the finding that coethnic bias in African societies is often less pronounced than is widely believed, and to underscore that scholars interested in explaining ethnicity's effects need to look beyond the idea that people simply care more about coethnics than they do about others.

2. A model of ethnic preferences and behavior

In the empirical analysis, we employ lab games to measure individual-level ethnic biases in behavior. The model developed below builds on Capellen et al. (2007, 2013) and describes what can be learned about coethnic bias in preferences from observing individual-level play in the Dictator and Public-good games. Specifically, it reveals that, under the weak assumption that subjects are not perfectly altruistic, a lack of observed coethnic bias in behavior implies no coethnic bias in individual preferences. We present

the baseline model in the context of a Dictator game; afterward, we discuss how the model relates to the Public-good game.³ We elaborate on how the theoretical model maps to the Choose-Your-Dictator game after we lay out that game's design and estimation procedure in Section 3.5.

Let m_i be the individual's fairness ideal. This could be an even split of the endowment, although the exact value does not matter for measuring the degree of observed coethnic bias in behavior. We introduce a coethnic bias in preferences, given by q_i , and an ethnicity indicator variable n_j , which takes the value one if the counterpart is a non-coethnic and zero if a coethnic. The utility of an individual i is:

$$u_i = (1 - t_i) - \frac{(t_i - m_i)^2}{2b(1 + n_j q_i)} \quad (1)$$

where the endowment of money is given by 1, t_i the transfer to the other player, and b_i is the degree of egoism ($b_i \geq 0$), i.e., $b_i = 0$ denotes perfect altruism.

The first term in equation 1 captures the utility component based on what the decision-maker keeps for himself, and the second term is the decision-maker's disutility due to deviations from the fairness norm. The disutility from deviations from this norm depends on three factors: b_i , q_i and n_j . In the limit as $b_i \rightarrow \infty$, the second term tends to zero as egoism increases, and a fully egoistic decision-maker chooses $t_i^* = 0$. At the other end of the spectrum, for fully altruistic individuals ($b_i = 0$), the second term receives all the weight in the decision process, and $t_i^* = m_i$. The degree of coethnic bias in preferences q_i has an analogous effect when the other party is a non-coethnic, i.e., for $n_j = 1$. In that case, the larger is q_i , the lower is the weight given to the fairness term, and the smaller is the transfer.

³ In Cappelen et al. (2007, 2013), individuals make a trade-off between selfish and other-regarding concerns when sharing a sum of money. We extend their original model by introducing ethnic preferences so that it matches the versions of our lab games where subjects have some information about the likely ethnicity of their partners.

For an interior solution, the first-order condition for sharing implies:⁴

$$t_i^* = m_i - b_i(1 + n_j q_i) \quad (2)$$

Sharing with a coethnic (CE) is then given by $t_i^{CE} = m_i - b_i$, and with a non-coethnic (NCE) by $t_i^{NCE} = m_i - b_i(1 + q_i)$. We measure the observed coethnic bias in behavior in the dictator game as $B^D \equiv t_i^{CE} - t_i^{NCE}$, i.e. how much more the decision-maker contributes when facing a coethnic instead of a non-coethnic. This implies that B^D is only a function of the degree of altruism and the degree of coethnic bias in preferences:

$$B^D = b_i q_i \quad (3)$$

A first implication of the model is that a strong observed coethnic bias in behavior ($b_i q_i$) is not necessarily the result of a high coethnic bias in preferences alone (high q_i), as it may be driven primarily by egoism (high b_i). A second implication of the model is central to the interpretation of empirical results: in a population with imperfect altruism ($b_i > 0$) – a realistic pattern which characterizes nearly all populations that have played the Dictator game, including ours – a finding of no coethnic bias in behavior ($B^D = b_i q_i = 0$) implies no coethnic bias in preferences ($q_i = 0$).

We now extend this analysis to the Public-good Game, where we employ an analogous utility structure for respondent i when making a contribution t_i to the group fund. We modify the utility function by making the fairness-norm group-dependent: $m_i^g \equiv f(E_{g,-i}[t])$, where $E_{g,-i}[t]$ is individual i 's expectation about other group-member contributions to the group fund of group g , and we assume $f' > 0$. The expression for an interior solution becomes:

$$\hat{t}_i^* = m_i^g - b_i(1 + n_j q_i)$$

⁴ Note that for $b_i(1 + n_j q_i) \geq m_i$ we have a corner solution, where the decision-maker keeps all the money for herself, $t_i^* = 0$.

After applying an analogous derivation as before, the observed ethnic bias in the public good game, $B^{PG} \equiv \hat{t}_i^{CE} - \hat{t}_i^{NCE}$, becomes

$$B^{PG} = b_i q_i + m_i^{CE} - m_i^{NCE}$$

The difference in contributions in a coethnic versus non-coethnic group reflects both a coethnic bias in preferences (the first term, as above) and in expectations about others' contributions. Critically, if other group members are expected to contribute less in a non-coethnic setting compared to a coethnic setting, this should amplify the observed coethnic bias in terms of Public-good game contributions.

3. Setting, experimental protocol, and research design

3.1. Setting, sample and timing

The study sample draws from two ethnically diverse low-income neighborhoods in Nairobi, Kenya: Kibera, which was a focal point for the 2007-08 post-election violence, and Viwandani. The experiments were carried out at the Busara Center for Behavioral Economics, which also oversaw the recruitment of subjects. Recruitment was limited to members of five of the six largest ethnic groups in Kenya (the Kikuyu, Luo, Luhya, Kamba, and Kisii, listed in order of size in both the Kenyan population and our study sample). Taken together, these five groups make up 82% of the Nairobi population and approximately sixty percent of the population countrywide.⁵ They also contain the two groups with the deepest historical rivalry, Kikuyu and Luo. The only major Kenyan ethnic group not represented in the sample is the Kalenjin, whose members do not reside in large numbers in either of the two recruitment neighborhoods.

There were two data collection rounds in our study, each with a different set of participants recruited using an equivalent sampling protocol. We call the first round,

⁵ Figures for Nairobi are from the 1999 Kenyan census; national figures are calculated from 2009 Kenyan census data from the Kenya National Bureau of Statistics (KNBS).

which took place from July to August 2012 (7-8 months prior to the March 2013 presidential and parliamentary election), the *Non-election round* and the second round, which took place from January to February 2013 (1-2 months prior to the election), the *Election round*. The purpose of dividing the lab sessions in this way was to test the expectation, documented recently by Eifert et al. (2010), that ethnic identities become more salient, and thus that coethnic bias might be more pronounced, during moments of heightened political competition. The macroeconomic environment, as captured in major commodity prices during 2011 to 2013 (see Appendix Figure A1), was relatively stable during the study period, so we can rule out that any differences we observe across the lab rounds are due to changes in the country's economic conditions. The political climate, by contrast, was markedly different. Whereas just 25% (14%) of participants in the Non-election round reported having recently attended a campaign rally (received cash from a candidate), this more than doubled to 61% (31%) among participants in the Election round. The total sample size is 1,362 distinct individuals, with 608 individuals in the Non-election round and 754 in the Election round (see Appendix Table A4 for a comparison of the two subject pools).

Two lab sessions were held each day, each lasting two to three hours and including up to 20 participants. As compensation, individuals were reimbursed for transport and given a "show-up fee" (in addition to their payouts from the games played). The results of each game were not revealed until after the session, at which point payments were made by mobile money transfer. Further details of the experimental lab protocol are contained in Appendix B.

Figure 2 illustrates the structure of the lab activities. Each lab session began with the administration of a 16-piece Raven Progressive Matrices test to measure individual cognitive ability (Raven 2008)—a potentially important factor in laboratory games. We use the normalized scores from this test both to confirm balance across treatment groups and as a control in our analyses.

Each lab session was then divided into two parts, both chronologically and conceptually. In the first, “standard,” part participants were randomly paired with individuals about whom they were provided no information. These partners were individuals who had played in pilot lab rounds and were randomly chosen to receive payouts based on the play of the games. Although participants had no way of knowing the specific ethnicities of the partners they were paired with, they knew that they were recruited from Nairobi’s multiethnic population. Hence, participants could infer that they were unlikely to be coethnics. Even members of the largest ethnic group in our sample (the Kikuyu, who comprise 30.9% of Nairobi residents according to the most recent available Kenyan census data) would have a lower than one-third chance of being paired with a coethnic. Patterns of play in these “standard” games can thus be interpreted as reflecting how participants behave when interacting with non-coethnics.

In the second, “profiled,” part of the lab, participants were paired with individuals about whom they received three pieces of background information: their education, their age, and their home region. The objective was to provide information with which participants could make strong inferences about their partners’ ethnic affiliations without being completely overt about the focus of our experiment. In Kenya, home region is usually a strong indicator of the participant’s ethnic background. The inclusion of information about the partner’s age and education was meant to obscure the ethnic focus of the study and to minimize experimenter demand effects. The order of the three pieces of information (home region, school completion, and birth year) was varied across subjects so as to avoid bias stemming from order effects.

The profiles of the players with whom participants were paired were randomly assigned. As in the standard games, the profiles were drawn from among the individuals who participated in the pilot sessions, only this time we limited the partner sample to those whose home regions provided unambiguous clues about their

ethnicity.⁶ To ensure a sufficiently large number of coethnic matches, we further limited the partner sample to individuals from the two largest ethnic groups, Kikuyus and Luos, groups whose high political relevance insures that many of the non-coethnic pairings in our sample were between members of groups whose relations have been historically fraught.

As laid out in our pre-analysis plan (see Appendix C), and in accordance with the ethnic political coalitions that emerged in advance of the 2013 national election, we categorize Luos together with Luhyas and Kisiis (two other western Kenyan groups long in coalition with Luos in national politics) as coethnics in the analysis, as distinct from the larger Kikuyu ethnic group. To be conservative, we pre-specified that the Kamba would be excluded from the main analysis due to the shifting political alliances of their leaders during 2012 and 2013 (but the results do not differ if they are included and grouped with the Kikuyu, the group that is culturally and geographically closest to them). The main empirical results are also unchanged if we ignore allied ethnic groups and restrict the sample to just Kikuyu and Luo subjects alone.

Each lab session consisted of three main games: the Dictator Game (DG), the Public-good game (PG), and the Choose-Your-Dictator (CYD) game, discussed in further detail below. Within the profiled parts of the DG and PG, the order of the coethnic and non-coethnic games was randomized. The CYD, by contrast, was always played with at least some partner information, although there are two information variants to this game, described below. At the end of the Election round sessions, subjects were also administered an IAT and then asked survey questions to capture self-reported political preferences and attitudes.

⁶ Piloting prior to the study confirmed that the vast majority of participants could correctly identify the intended ethnic backgrounds of their partners based on the home regions ascribed to them.

3.2. Priming treatments

Participants were exposed during the lab sessions to a randomly assigned priming treatment designed to increase the situational salience of (1) political competition, (2) ethnic-cultural differences, (3) ethnic-political differences, or (4) national identity. There was also a control group in which participants were exposed to a neutral prime.

The political competition prime was designed to serve as an experimental counterpart to the real world variation in proximity to elections afforded by our two lab rounds. As with the comparison of the lab rounds occurring 7-8 months and 1-2 months before the 2013 elections, the expectation was that priming participants to political competition would increase coethnic bias in the behavioral games. The rationale for the ethnic primes, which were meant to mimic the daily exposure that individuals have to appeals to shared ethnic culture (the ethnic-cultural prime) or to blatant tribal politics (the ethnic-political prime), was even more straightforward. By priming subjects to think about ethnicity, the expectation was that ethnic differences between players in the games would become more salient, resulting in higher levels of in-group bias (Bowles and Gintis 2004; Shayo 2009). Finally, the purpose of the national identity prime was to explore whether ethnically motivated behavior would be reduced by invoking a common superordinate identity (Kramer and Brewer 1984; Gaertner and Dovidio 2000; Charnysh, Lucas and Singh 2015).

To reduce the likelihood that priming effects would dissipate, priming was implemented at four distinct points during the lab session (see Figure 2). The primes were administered between each set of games as short, three question on-screen quizzes. For example, one of the political competition priming questions asked: *“How many political candidates are running for the Presidency?”* One of the ethnic-cultural priming questions highlighted the country’s ethnolinguistic diversity: *“This greeting comes from which region: ‘Orie’?”* Questions specific to the priming conditions were intermixed with neutral questions, mainly focused on Kenyan popular culture or everyday

life in Nairobi, for example: “How often do you ride a matatu (mini-bus) every week?” Each quiz included two priming questions and one neutral question. Those in the control group were asked only neutral questions. Appendix B provides the full set of priming questions. Exit interviews and focus group discussions confirmed that the priming questions were effective in eliciting concepts of political competition, ethnic-cultural differences, political-cultural differences, and national identity.

3.3. Dictator Game

The Dictator game aims to capture an individual's altruism toward others. Participants were given an endowment of 50 Ksh (\$0.60) – equivalent to roughly an hour's wage for many Nairobi workers – and were asked to decide how much to give away to a partner with whom they were randomly paired. Subjects played both standard and profiled version of the game, with the latter played twice: once with a coethnic partner and once with a non-coethnic partner (with the order randomized). This allows us to compare the level of altruism toward coethnic versus non-coethnic partners.

The standard Dictator game was played in both the Election and Non-election rounds, but due to a programming error in the lab software, only the coethnic version of the profiled game was played in the Non-election round. This means that we can estimate the extent of coethnic bias in the Dictator game only in the Election round. (There was no such coding error for the Public-good game or the Choose-Your-Dictator game, so for those games we are able to measure coethnic bias in both the Non-election and Election rounds, as described below.)

The outcome of interest is the percentage of the endowment transferred to the partner. We first focus on how partner coethnicity affects transfers, and how coethnicity interacts with election proximity, and then estimate the effect of the priming treatments.

The first specification includes data from both the standard and profiled games:

$$Y_{ij} = \alpha + \beta_1 Coethnic_{ij} + \beta_2 Election_i + \beta_3 Coethnic_{ij} * Election_i + \beta_4 Noncoethnic_{ij} * Election_i + \varepsilon_{ij} \quad (4)$$

Y_{ij} is the transfer (in percent of the endowment) by subject i to their partner in game j , $Coethnic_{ij}$ ($Noncoethnic_{ij}$) is an indicator variable for being paired with a coethnic (non-coethnic) partner in game j , $Election_i$ indicates whether respondent i is observed in the Election round, and ε_{ij} is the error term, clustered by individual. The lack of a non-coethnic profiled Dictator game in the Non-election round explains the missing non-interacted $Noncoethnic_{ij}$ term. The estimate of coethnic bias in behavior (B^D) in the Election round is given by the difference in coefficient estimates on the coethnic partner and non-coethnic terms, and the formal hypothesis test is that $(\beta_1 + \beta_3) = \beta_4$. The Election round effect captured in β_2 is potentially of independent interest. As a robustness check, we also estimate a specification with covariates including ethnicity and gender, years of education, and the normalized Raven's test score.

We estimate effects of the priming treatments in two slightly different ways. In the main analysis that we report in the text, we focus on the standard Dictator game in order to provide the cleanest possible test of priming effects. Since participants could reasonably infer that anyone they were paired with was likely to be a non-coethnic, the results of the standard Dictator game provide insight into the effects of priming on levels of altruism vis-a-vis non-coethnics. We estimate the following equation:

$$Y_{ij} = \alpha + \beta_1 Election_i + \sum_{k=1}^3 \beta_{1+k} T_k + \sum_{k=1}^4 \beta_{4+k} (T_k * Election_i) + \varepsilon_{ij} \quad (5)$$

The terms T_1, T_2, T_3 and T_4 are indicator variables for the four priming conditions, namely, the political-competition prime, the ethnic-cultural prime, the national prime, and the ethnic-political prime, respectively. β_1 estimates the difference in the average level of transfers across the Election and Non-election rounds, β_2, β_3 , and β_4 , the average effects of the first three priming treatments in the Non-election round, and β_5, β_6 , and β_7

the differential effect of these primes in the Election round. β_8 is the effect of the ethnic-political prime, which was only administered in the Election round, and thus is only included in interaction with the Election round indicator.

We also analyze the effects of priming using the profiled Dictator game (see Appendix D). This analysis provides an arguably more direct test of the impact of priming on coethnic bias. However, it is theoretically ambiguous if the impact of priming will be stronger or weaker in the standard or the profiled version of the game (see Benjamin et al. 2010). Hence we put more weight on the priming analyses that employ the non-profiled games. In any case, results are similar using both approaches.

3.4. Public-good Game

Participants in the Public-good game were given an endowment of 60 Ksh (\$0.70) and could choose how much to keep for themselves versus how much to contribute to a group fund in which contributions were multiplied by two before being equally shared among the three players in the game. The contribution level thus captures an individual's willingness to share resources to make others better off, and resembles the classical prisoner's-dilemma. Compared to the Dictator game, the Public-good game's framing in terms of a "group fund" is distinct, and the recognition that other actors are also making decisions could trigger reciprocity that might affect behavior.

Before deciding how much to contribute themselves, subjects stated how much they believed each of the other players would contribute. These were non-incentivized elicitations. Subjects received no information about the two other players in the standard part of the session, where once again the other players were randomly drawn from subjects in the pilot rounds. In the second, profiled part of the lab, they were informed about each of the two other players' years of education, age and home region, just as in the profiled Dictator game. There were two types of profiled Public-good games. The first was a "mixed" group, with one coethnic player and one non-coethnic

player, and the second, the “coethnic” group, in which both players were coethnics with the subject. The order in which these were played was randomized. The difference in individual contributions to the group fund across the coethnic group and the mixed group is our measure of coethnic bias. In the pre-analysis plan, we pre-specified both a focus on contributions and an analysis of “contributions minus beliefs,” a measure of conscious free-riding on the part of the subject proposed in Fischbacher and Gächter (2010). For simplicity here, we focus on contributions alone but results are very similar for the contributions minus beliefs outcomes (see Appendix D).

The main econometric specification follows the Dictator game, where $Mixed_{ij}$ is an indicator for an ethnically mixed group, $Coethnic_{ij}$ is a homogeneous coethnic group, and Y_{ij} is the contribution of individual i in game j , and all games (standard and profiled) are pooled in the analysis:

$$Y_{ij} = \alpha + \beta_1 Coethnic_{ij} + \beta_2 Mixed_{ij} + \beta_3 Election_i + \beta_4 Coethnic_{ij} * Election_i + \beta_5 Mixed_{ij} * Election_i + \varepsilon_{ij} \quad (6)$$

We estimate coethnic bias in the Public-good game (B^{PG}) in the Non-election round in an F-test of the hypothesis that $\beta_1 = \beta_2$. We also assess if there is a differential degree of coethnic bias in the Election round by testing if $\beta_4 = \beta_5$. As in the Dictator game, we also estimate the effects of the priming interventions in the standard Public-good game (see Appendix D).

3.5. Choose-Your-Dictator Game

The Choose-Your-Dictator game was designed to capture expectations of differential altruism in coethnic and non-coethnic interactions by measuring whether subjects

discriminate along ethnic lines in their choice of a “leader” responsible for allocating resources, a conception thus broadly linked with the process of electing public officials.⁷

After having already played the Dictator game and become familiar with its rules, subjects (“choosers”) were presented with two randomly drawn profiles of potential “Dictators” and were then asked to choose one to be “their” Dictator in a Dictator game. One of the two drawn profiles was a coethnic, the other a non-coethnic (as before, these were the profiles of actual individuals who had played in the pilot rounds). As in the profiled Dictator and Public-good games, choosers were provided with basic information about both of the potential Dictators’ years of education, age, and birth region. Prior to choosing “their” Dictator, we also elicited the subject’s beliefs about how much each of the Dictators they were choosing between would allocate to them in the game (this is the source of the Dictator game beliefs data presented in Figure 1, Panel B). The subject’s payoff in the game was the amount the selected Dictator actually transferred when he or she played the game in the pilot rounds.

The Choose-Your-Dictator game was played twice. In the first, “standard” version, choosers select a Dictator knowing that the potential Dictators do not have any information about them, i.e., the payoff will be from the standard Dictator game played by these Dictators. This choice depends only on expected differences in average generosity between the two potential Dictators, but not on their differential altruism toward a person of the chooser’s type. In the second, “profiled” version, the choosers select a Dictator knowing that the Dictator will be provided with the choosers’ own basic information profile (i.e., their years of education, age, and home region). The payoff now is from the chosen Dictator’s play in the profiled Dictator game played in

⁷ Adida et al. (2015) develop a similar game. Their experimental protocol differs, however, as the participants go through a face-to face speed-chatting process prior to voting for a group “leader.” Hence the choice of “leader” could be affected by personal characteristics out of the researchers’ control, and the impact of ethnicity may be more difficult to discern.

the pilot rounds.⁸ In this second version of the game, where the Dictator has information about the ethnicity of the recipient, the chooser’s selection depends on both expected differences in altruism by the two Dictators generally and toward a coethnic. The difference in play across the “standard” and “profiled” versions of the Choose-Your-Dictator game thus allows us to isolate the chooser’s expectations of coethnic bias.

To understand the extent of respondent preferences for a coethnic “Dictator”, we implement an ordered logit specification. The dependent variable is Y_{ip} , which takes on the following values for Dictator profiles p in the choice set of respondent i : $Y_{ip} = 0$ if profile p is not chosen by the subject, $= 1$ if the subject is indifferent between profile p and the other profile in its choice set, and $= 2$ if the subject chooses profile p . The “indifferent” option was included after we piloted the game and observed that large shares of participants claimed they were unable to choose between potential Dictators based on the basic profile information alone, an early hint about the low levels of differential altruism across ethnic lines in this population.

We specify the probabilities for observing each outcome value of Y_{ip} as a function of V_{ip} , which can be interpreted as respondent i ’s latent utility for profile p (Woolridge 2001). Using this framework, we carry out maximum likelihood estimation. The main specification estimates the difference between the valuation placed on a coethnic versus a non-coethnic profile across the “standard” and the “profiled” versions of the Choose-Your-Dictator game:

$$V_{ip} = \alpha_p + \beta_1 Coethnic_{ip} + \beta_2 Coethnic_{ip} * Profiled_{ip} + \beta_3 Coethnic_{ip} * Election_i + \beta_4 Coethnic_{ip} * Profiled_{ip} * Election_i + \varepsilon_{ip} \quad (7)$$

⁸ In practice, it was impossible to perfectly match all characteristics from our set of pilot profiles to each and every subject in the actual lab sample, so the matching between the subject and the profile was done on coethnic home region alone. Given that we do not find that the years of education or age in the profiles are significantly correlated with subjects’ choices, we do not believe that this “mismatch” of characteristics other than home region is consequential for the analysis.

Here α_p is the set of profile fixed effects, $Coethnic_{ip}$ is an indicator variable for profile p being coethnic of the subject, and $Profiled_{ip}$ indicates whether the Dictator was playing the profiled version of the Dictator game, and thus had information about the subject's home region. In the Non-election round, β_1 estimates the degree of coethnic bias in the standard Choose-Your-Dictator game (equivalent to the generosity term b_i in our model), and β_2 is the additional degree of coethnic bias in the profiled Choose-Your-Dictator game, which captures the degree of coethnic bias that the subject expects from the Dictator given that the Dictator knows that the subject is a coethnic (equivalent to their expectation of the q_i coethnic preference term in the model). Thus the estimate of β_2 is of central interest. The β_3 and β_4 terms capture any additional coethnic bias in the Election round. The error-term ε_{ip} has a Type I Extreme-value distribution. As a robustness check, we also condition on subject covariates, in this case interacted with the $Coethnic_{ip}$ term (since explanatory variables that do not vary across choice options cannot be estimated in a logit model).

3.6. Implicit Association Tests (IATs)

To capture coethnic bias which participants may not have felt comfortable revealing publicly, we included an Implicit Association Test (IAT) as part of our main lab protocols in the Election round. This tool is commonly used in social psychology to measure racial or gender bias, free of potential social desirability bias and experimenter demand effects. The IAT is premised on the idea that individuals find it easier to respond to concepts that are “strongly associated” (i.e., about which they agree) than for concepts that are weakly associated (Nosek et al. 2007). In our lab, we focused on the degree to which respondents have a bias in favor of their own ethnic group.⁹ Others have used IAT methodologies to assess coethnic bias in Africa (Lowe et al. 2015),

⁹ We also carried out a parallel “national” IAT to capture the strength of feelings toward Kenya, and plan to discuss those results in future studies (see Appendix E).

gender bias in India (Beaman et al. 2009), and racial bias in the U.S. and elsewhere (Greenwald et al. 2003; Bertrand et al. 2005; Ogunnaiké et al. 2010; Rooth 2010).

The specific IAT that we developed provides a measure of a respondent's automatic associations with certain Kenyan ethnic groups. In the "congruence" round of the IAT, individuals were timed to assess how quickly they were able to associate coethnics with "good" traits and non-coethnics with "bad" traits. In the "dissonance" round, coethnics were associated with bad traits and non-coethnics with good ones. A faster response time for congruence tasks compared to dissonance tasks implies that the subject has a more positive attitude toward coethnics (or a more negative attitude toward non-coethnics). Examples of "screen shots" from the IAT are in Appendix B.

The outcome measure for the IAT is the within-respondent normalized difference in average reaction times (ART) between the dissonance (D) and congruence (C) tasks, called the d-score for subject i : $d_i = \frac{ART_i^D - ART_i^C}{\sigma_i}$, where σ_i is the standard deviation of a respondent's reaction times across all items. In keeping with the literature, we interpret a higher d-score as indicating stronger bias. We estimate the average level of coethnic bias in the IAT both with and without the priming treatments.

While the main rationale for including the IAT is to serve as a check on the possibility that our results might be driven by social desirability bias, it is important to recognize that the IAT measures a type of coethnic bias that is conceptually different from the bias measured by the DG, PG or CYD games. Whereas the IAT picks up potentially unconscious positive or negative associations with in- and out-group members, the conscious choices made in the behavioral games may or may not be affected by implicit biases of the sort measured by the IAT. The implication is that any differences we observe in the findings of the IAT and the behavioral games could stem either from the fact that the behavioral games results were in fact altered by social

desirability bias and/or experimenter demand effects or that the two kinds of tasks are simply capturing different types of coethnic bias.

4. Results

4.1. Descriptive statistics

Descriptive statistics for the lab participants are presented in Table 1 (left column). Females are slightly more than half the sample (at 60%). Respondents are 33 years old on average, and they have completed an average of 9.6 years of education (a schooling level that is typical for young Kenyan adults in urban areas but higher than national averages). Twenty-nine percent of the sample report having a continuous source of wage income, 13% report self-employment, and 30% are unemployed (not shown). More than half of participants belong to the two largest ethnic groups in our sample, with 32% Kikuyu and 21% Luo. We present evidence in Appendix A that the randomization across the various priming interventions created treatment groups with similar observable characteristics along these dimensions.

4.2. Estimating coethnic bias

As noted in the introduction, there is no evidence for coethnic bias in either the Dictator game or the Public-good game. Reproducing the result from Figure 1, the top row of Table 1 presents average transfers in both games in the full sample. For both games, the differences in average transfers are small (35.6% versus 35.4% in the Dictator game and 46.4% versus 46.7% in the Public-good game for Coethnic versus Non-coethnic transfers, respectively) and not statistically significant.

Given the large number of hypotheses we investigate, an important concern is the over-rejection of null hypotheses (i.e., “false positives”) due to the problem of multiple inference (Anderson 2008). To mitigate this risk, we present two types of p-values in our analyses: first the regular or “per-comparison” p-value, which is

appropriate when testing a single hypothesis, and second, the FWER-adjusted p-value, which captures the likelihood that at least one true null hypothesis is falsely rejected within a given set of hypotheses, using the free step-down resampling method described in Westfall and Young (1993).

As explained above, the lack of a Non-coethnic arm in the Non-election round lab implies that the relevant test for differential altruism in the Dictator game is the sum of the coefficient estimates on the Coethnic terms in both rounds minus the coefficient estimate on the Non-coethnic term in the Election round. The average coefficient on bias is just 0.6 percentage points and not statistically significant (Table 2, columns 1-2, p-value = 0.67). Our large sample size implies that these are quite precisely estimated zeros: the 95% confidence interval on the overall Coethnic effect in the Election round ranges from -2.2 to +3.4 percentage points. For those in the priming control group (i.e., who received no priming to ethnic identity, political competition, or national identity), the effect is slightly larger in magnitude but not significantly different from zero (column 3, p-value = 0.48). In line with these findings, beliefs about the generosity of a coethnic Dictator are nearly identical to beliefs about non-coethnic Dictators and not significantly different from zero (Figure 1, Panel B, p-value = 0.51).

In the Public-good game, average coethnic bias in the Non-election round is approximately 1.57 percentage points but this difference is not statistically significant (Table 3, columns 1-2, p-value = 0.16). There is actually a significant reduction in coethnic bias in the Election round of the Public-good game, with an average effect of -3.67 percentage points (p-value = 0.027), which runs counter to the existing evidence that ethnic identity and preferences might become more salient closer to elections (Eifert et al. 2010). Taken together, the overall degree of coethnic bias in the Election round is the sum of these two effects, and thus is small and negative (roughly -2 percentage points) and not statistically significant. In the priming control group, there are no significant coethnic bias effects in either lab round (column 3). Subjects also believe

groups members will be equally generous in coethnic and ethnically mixed groups (Figure 1, Panel D). These results are unchanged when the dependent variable is “contributions minus beliefs” (see Appendix D).

Even if overall coethnic bias is close to zero, a natural question is whether there might be heterogeneity in the degree of this bias across subgroups of our sample. This is illustrated in Table 1. The striking pattern that emerges is how little variation there is along the various dimensions that we pre-specified in the pre-analysis plan. For instance, across both gender groups, the difference in average transfers in the Dictator game to coethnics and non-coethnics is less than one percentage point, and among subgroups defined by age, education, Ravens cognitive score, and major ethnic groups, the differences are at most 1.2 percentage points. None of these differences is statistically significant at traditional confidence levels. A similar pattern holds for Public-good game contributions between coethnic and mixed groups, with coethnic bias estimates of at most 1.9 percentage points for any subgroup. Again, none of these differences are significant.

The evidence from the Dictator game and Public-good game are thus both consistent with little to no coethnic bias in this population as a whole, as well as for the major demographic subgroups. Appendix D contains further descriptive statistics and histograms, as well as tests indicating that the null hypothesis of equality of the distributions of Coethnic and Non-coethnic transfers cannot be rejected in either game.¹⁰

The Choose-Your-Dictator game also sheds light on subjects’ expectations about whether others will be differentially altruistic toward them. The overall preference for a coethnic Dictator (captured in the standard Choose-Your-Dictator game) in the Non-

¹⁰ A careful reader of Tables 2 and 3 will note that average transfers in the Dictator game and contributions in the Public-good game are significantly lower in the Election round. Although this finding is intriguing, interpreting it is complicated by the fact that the comparison across the two lab rounds bundles together variation in proximity to the election with a “time effect” that captures all of the other changes that took place in the broader Nairobi (or Kenyan) environment, the study samples, and the Busara Center lab space. We therefore hesitate to read too much into this result.

election round is minimal, with 27% choosing a coethnic versus 22% choosing a non-coethnic, and half simply opting for “indifferent” (Figure 3, left panel). The patterns in the Election round (right panel) are nearly identical. While seemingly small in magnitude, the effect is significant (Table 4, column 1) though not robust to the inclusion of covariates or focusing on the priming control group (columns 2-3).

The more important test from the point of view of understanding coethnic bias is the difference between these patterns and those that emerge in the profiled Choose-Your-Dictator game, where the Dictator is given information about the participant and is therefore perceived by the participant as being in a position to condition his/her generosity on the participant’s ethnicity. In the profiled game in the Non-election round, the proportion of subjects choosing a coethnic Dictator rises slightly, to 32%, as does the proportion who choose a non-coethnic (25%) (Table 1, top row and Figure 3, left panel). As discussed above, the difference in behavior between the standard and profiled games is captured in the coefficient estimate on the Profiled Game * Coethnic term, and this effect is small and not statistically significant (Table 4, row 2); taking into account the issue of multiple hypothesis testing, the FWER adjusted p-value is close to one (at 0.993). Results for the Election round are similar (row 4). There is thus no evidence that subjects are more likely to choose a coethnic Dictator when the Dictator has information about their home region (and thus a way to make an inference about their ethnic background), relative to when the Dictator does not have this information. This is consistent with the previous findings since there is no evidence for coethnic bias in behavior in the Dictator game (Table 2). Expectations and actions regarding differential altruism across ethnic lines are thus aligned among our subjects.

4.3. The impact of priming on coethnic bias

None of our priming treatments had a statistically significant effect on behavior in either the standard Dictator game (Table 5, column 1) or the standard Public-good game

(column 4), and these effects did not differ significantly in the Election round (columns 2 and 5, respectively).

As noted, the ethnic-political prime was only included in the Election round. This is because, after finding no effect of the ethnic-cultural prime in the Non-election round, we hypothesized that this might be because our priming approach was too subtle. Hence, we decided to introduce a more “blatant” ethnic priming treatment in the Election round. In contrast to the ethnic-cultural prime, which sought to enhance the salience of ethnic identity by emphasizing cultural aspects of ethnic identity, this new treatment directly and overtly primed subjects to the link between ethnic identity and political outcomes. For example, one of the questions asked: *“Which of the following ethnic groups controls the largest share of cabinet positions?”* Yet this prime also had no significant effect on transfers in either game (Table 5, columns 2 and 5). In Appendix A, we show that there is also little evidence of priming effects in the profiled Dictator game, profiled Public-good game or Choose-Your-Dictator game. Few of the relevant differential priming effect estimates are statistically significant in the per-comparison sense, and almost none survive the multiple testing adjustment.

The bottom line is that the priming interventions, which aimed to raise the salience of political competition, ethnicity, and national identity, did little to change the degree of coethnic bias exhibited by participants in the lab activities. We interpret this as consistent with our main findings regarding both the lack of evidence for coethnic bias in our sample and the absence of any differences in levels of bias across the Election and Non-Election lab rounds. The results suggest that coethnic bias does not appear to be “just below the surface” in a way that might be triggered by priming subjects to ethnicity or political competition. They also suggest that priming subjects to their membership in a superordinate national identity does not reduce their (already low) levels of coethnic bias.

4.4. Addressing experimenter demand and social desirability biases

A potential concern with the attempt to measure coethnic bias through experimental games is that game behavior may be affected by the desire not to be seen to be discriminating along ethnic lines. We address this issue in several ways.

First, as noted, the laboratory protocols were designed to reduce the likelihood that participants would be cued to our interest in ethnicity. Even the ethnic priming was subtle, presented in the form of a quiz that also included other neutral questions. Further, lab participants were insured that their answers would be kept confidential and seated in private cubicles with headphones.

We collected exit interview data to ascertain how successful we were in masking our interest in ethnic discrimination. During piloting that took place before the start of the main data collection, we asked subjects—both at the midpoint of the lab sessions and at their conclusion—detailed questions regarding their understanding of the main focus of the experiment. The most common response was that the experiment was mainly about economic or business issues. Roughly equal numbers of subjects thought the activities were about education or about “tribe” (i.e., ethnic issues). The proportion that believed that ethnic issues were a focus of the experiment rose slightly by the end of the lab, perhaps due to the nature of the Choose-Your-Dictator game, which was played last, which asks participants explicitly to choose between two other players after being given information about those players’ region of origin (among other characteristics), but this still remained less than a quarter of all subjects. We are therefore confident that the vast majority of participants were not aware of the study’s core research aims (for further details, see Appendix B).

As noted, subjects’ behavior in the IAT also provides a social desirability bias-free check on our main results (as well as a complementary measure of coethnic bias). While we find a statistically significant and positive degree of bias in this test (with a magnitude of 0.079 standard deviation units; p -value < 0.01 ; Table 6, column 1), this

level of bias is considered “small” in the related IAT research literature in psychology (Cohen 1988). For instance, studies of the bias whites hold against blacks in the United States find estimates that are roughly six times as large, in the range from 0.45 to 0.52 (Nosek et al. 2007). In India, the average bias toward female political leaders is notably higher, at 0.11 to 0.15 (Beaman et al. 2009). In Congo, Lowes et al. (2015) report d-score values in a study of coethnic bias of 0.14, nearly twice as large as our estimates. We therefore interpret our findings as providing only weak support for the presence of social desirability bias in our main results or for the salience of an alternative source of coethnic bias rooted in negative associations with out-group members.

With respect to priming, there is some evidence that the political competition treatment increased coethnic bias in the IAT (Table 6, columns 3-4), which might be interpreted as consistent with the findings in Eifert et al. (2010).¹¹ However, this effect is not significant at conventional levels once multiple testing considerations are accounted for (column 5). None of the other priming interventions affected implicit coethnic bias.

4.5. Pre-analysis plan and a “selective presentation test”

A strength of our analysis is that we pre-specified our analytical approach in a pre-analysis plan (AEA Social Science Registry ID# AEARCTR-0000016; see Appendix C). A registered pre-analysis plan helps address concerns of publication bias and data mining. This latter concern is especially salient in a study such as ours given its complex structure, with multiple games, primes, subsamples, lab rounds, and potential tests.

Although the past few years have seen a rise in the use of pre-analysis plans for field experiments (Casey et al. 2012; Miguel et al. 2014), they are far less common in laboratory studies—in part because, as Coffman and Niederle (2015) argue, they may be superfluous in settings where experiments can be easily replicated. While we

¹¹ Note that since the IATs were administered only in the Election round we are unable to assess whether implicit coethnic bias is different in actual proximity to an election, rather than just via experimentally priming participants to the salience of political competition.

wholeheartedly endorse Coffman and Niederle’s emphasis on replication, we view replication studies and pre-analysis plans as complements rather than substitutes. Moreover, in certain settings where lab experimental researchers only have “one shot” at a particular type of data—due to access, cost, or timing—replication is infeasible. Hence, since we examine ethnicity in the unique context of Kenya’s 2013 national election, we view the pre-analysis plan as a crucial component of our study.

Given the space constraints of a journal article, it was not possible to present every result that was pre-specified. This is a common problem for authors of complex, multifaceted research projects. It raises the possibility that we might have selectively emphasized a non-representative subset of results that were more (or less) statistically significant, or more consistent with our theoretical priors. One means—to our knowledge, novel—of assessing whether we might have done this is to undertake a “selective presentation test” in which we plot the distribution of p-values presented in the paper’s main tables (N=86 tests) against the distribution of p-values for all hypothesis tests specified in the pre-analysis plans (N=763 in total).

We do this in Figure 4. The distributions are broadly similar, but the main study tables do appear to slightly over-report statistically significant estimates (those with p-values less than 0.05) for both the unadjusted p-values (Panel A) and the FWER adjusted p-values (Panel B) and to somewhat under-report those with p-values close to one. Consistent with this visual inspection, the Kolmogorov-Smirnov test on the equality of these two distributions is rejected at 95% confidence for both the unadjusted and the FWER adjusted p-values. Of course, since our central finding is a null result, the slightly disproportionate emphasis on statistically significant effects cannot be driving our main conclusions. Rather we think it likely reflects our decision to relegate a large number of null priming treatment estimates and interaction effects to the appendix.

5. Conclusion and broader implications for Africa

Using a much larger sample than prior studies and a richer set of experimental methods and measures, we find little evidence for coethnic bias in a Kenyan sample. These findings present a challenge to both theory and conventional wisdom about why ethnicity matters and suggest that mechanisms other than coethnic bias should be considered more seriously as causes of the political, social and economic outcomes that ethnicity is commonly thought to produce in Africa.

Notwithstanding the importance of these findings, it is reasonable to wonder, given the urban nature of our study location, whether our results should be interpreted as speaking only to the way ethnicity operates in an urban setting. Indeed, our findings are in apparent tension with the results of Hjort's (2014) important study of coethnic bias in rural Kenya. Taking advantage of the random assignment of workers to teams on a flower farm to study whether within-team productivity is lower when those teams are ethnically diverse, Hjort finds strong evidence of discrimination by team members on behalf of coethnics.

There are several possible ways to account for the differences between these findings and our own that have nothing to do with the contrasting urban and rural locations of the two studies. The most immediate is that, although coethnic bias may play a role, the negative diversity effects in Hjort may be caused by other mechanisms. For example, the fact that Hjort shows that shifting to group-based pay on work teams mitigates much of the negative effect suggests that institutional factors are critical. Hjort's design also makes it hard to rule out the possibility that ethnicity matters by providing a technology that facilitates team production. To the extent that these alternative mechanisms are driving Hjort's results, his findings are not contradictory to ours—indeed they would reinforce our point regarding the salience of other channels in explaining ethnicity's effects.

However, it is difficult to completely rule out the possibility that the different findings are products of different degrees of coethnic bias in urban and rural domains. Given the strictly urban nature of our subject pool, our study is not well suited to fully assessing this possibility. However, we can make some progress by comparing the behavior in our sample of subjects who were long-time Nairobi residents with the behavior of those who had moved to the city relatively recently. While only suggestive, this is a meaningful test to the extent that individuals gradually assimilate to local norms over time, such that those who have moved to Nairobi relatively recently have orientations to non-coethnics that are more similar to those who remain in rural areas. Such an interpretation is consistent with research that finds that migrants gradually adapt their social preferences and behaviors over time to the norms in their new locations (Laitin 1998; Henrich et al. 2006; Jang and Lynham 2015).

This analysis was not specified in our pre-analysis plan, and is thus more speculative than our other results. While we would ideally focus on those individuals who had just moved to the city, there are relatively few such individuals in our sample: just 2 percent had lived in Nairobi for less than one year. However, a sizeable proportion (roughly 20%) had moved to Nairobi within 5 years and even more (40%) in the last 10 years. We examine the behavior of this latter group in comparison with that of longer-term urban residents. Consistent with the hypothesis that there might be an important urban-rural divide in coethnic bias, we find statistically significant differences in behavior of these two subsamples. While individuals who have resided in Nairobi for more than a decade exhibit no evidence of coethnic bias in their behavior, more recent migrants are significantly more generous toward their coethnics in the Election round for the profiled Dictator game, giving an average of 4.2 percentage points (standard error 1.8) more to coethnics. Coethnic bias in the Dictator game is even larger among those who had resided in Nairobi for at most 5 years, with an average increase of 6.2 percentage points (standard error 2.4).

Migrants from rural settings thus appear to gradually assimilate to a new set of norms regarding interethnic cooperation and ethnic preferences over time, perhaps due to the far more extensive interactions they have with non-coethnics in Kenya's cosmopolitan capital. Further evidence comes from our political attitudes survey, which shows rising ethnic identification (relative to other dimensions of individual identity) in the 2013 Election round among those who have lived in Nairobi for less than 10 years (coefficient estimate 8 percentage points, significant at 90% confidence) but not among longer-term Nairobi residents.

Obviously, these findings cannot be taken as causal for many reasons. Most important, we cannot rule out that the differences we find are driven by selective urban migration patterns over time rather than the true effect of the length of urban residency. Furthermore, as noted, these analyses were not pre-specified. Nevertheless, the suggestion that altruism vis-à-vis one's coethnics may be weaker in urban than in rural settings is provocative—in part because it would call for a reevaluation of the well-known variant of modernization theory that posits a positive relationship between urbanization and ethnic divisions (e.g., Bates 1983).

Our suggestive findings may also provide new insight into the future trajectory of ethnic divisions in a region whose urban population has swelled from 27% in 1990 to 40% today and is expected to reach 55% by 2050 (United Nations 2014). Others have pointed to the role that expanding democratization may play in dampening ethnic favoritism (Burgess et al. 2015). Gaining a better understanding of the role played by rapid urbanization will also be an important goal for future scholarship on ethnicity in Africa.

References

- Adida, Claire L., David D. Laitin and Marie-Anne Valfort. 2015. "Religious Homophily in a Secular Country: Evidence From a Voting Game in France." *Economic Inquiry* 53(2): 1-20.
- Alesina, Alberto, Reza Baqir, and William Easterly. 1999. "Public Goods and Ethnic Divisions." *Quarterly Journal of Economics* 114(4): 1243–84.
- Alesina, Alberto, and Eliana La Ferrara. 2005. "Ethnic Diversity and Economic Performance." *Journal of Economic Literature* 43(3): 762–800.
- Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103(484): 1481–95.
- Barkan, Joel D., and Michael Chege. 2008. "Decentralising the State: District Focus and the Politics of Reallocation in Kenya." *Journal of Modern African Studies* 27 (03): 431.
- Bates, Robert. 1983. "Modernization, Ethnic Competition and the Rationality of Politics in Contemporary Africa." In Donald Rothchild and Victor Olorunsola, eds. *State Versus Ethnic Claims: African Policy Dilemmas*. Westview Press: 152-171.
- Beaman, Lori, Raghavendra Chattopadhyay, Esther Duflo, Rohini Pande and Petia Topalova. 2009. "Powerful Women: Does Exposure Reduce Bias?" *Quarterly Journal of Economics* 124(4): 1497–1540.
- Becker, Gary S. 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.
- Benjamin Daniel J., James J. Choi and A. Joshua Strickland. 2010. "Social Identity and Preferences." *American Economic Review* 100(4): 1913-28.
- Bertrand, Marianne, Dolly Chugh and Sendhil Mullainathan. 2005. "Implicit Discrimination." *American Economic Review* 95(2): 94–98.
- Bowles, Samuel and Herb Gintis. 2004. "Persistent Parochialism: Trust and Exclusion in Ethnic Networks." *Journal of Economic Behavior and Organization* 55: 1-23.

- Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria and Gerard Padró i Miquel. 2015. "The Value of Democracy: Evidence from Road Building in Kenya." *American Economic Review* 105(6): 1817–51.
- Cappelen A.W., A. D. Hole, E. Ø. Sørensen and B. Tungodden, 2007. "The Pluralism of Fairness Ideals: An Experimental Approach," *American Economic Review*, 97(3): 818-827.
- Cappelen, A. W., J. Konow, E. Ø. Sørensen, & B. Tungodden (2013). "Just luck: An experimental study of risk-taking and fairness." *American Economic Review*, 103(4), 1398-1413.
- Carlson, Elizabeth. 2015. "Ethnic Voting and Accountability in Africa: A Choice Experiment in Uganda." *World Politics* 67(2): 353–85.
- Casey, Kate, Rachel Glennerster, and Edward Miguel. 2012. "Reshaping Institutions: Evidence on Aid Impacts Using a Preanalysis Plan." *Quarterly Journal of Economics* 127(4): 1755–1812.
- Charnysh, Volha, Christopher Lucas and Prerna Singh. 2015. "The Ties That Bind: National identity Salience and Pro-Social Behavior Toward the Ethnic Other." *Comparative Political Studies* 48: 267-300.
- Coffman, Lucas C. and Muriel Niederle. 2015. "Pre-Analysis Plans Have Limited Upside, Especially Where Replications Are Feasible." *Journal of Economic Perspectives* 29 (3): 81–98.
- Cohen, Jacob. 1988. *Statistical Power Analysis for the Behavioral Sciences. Statistical Power Analysis for the Behavioral Sciences.*
- Collier, Paul and Ashish Garg. 1999. "On Kin Groups and Wages in the Ghanaian Labour Market." *Oxford Bulletin of Economics and Statistics* 61(2): 133-151.

- Conroy-Krutz, Jeffrey. 2012. "Information and Ethnic Politics." *British Journal of Political Science* 43: 345-373.
- Dionne, Kim Yi. 2015. "Social Networks, Ethnic Diversity, and Cooperative Behavior in Rural Malawi." *Journal of Theoretical Politics* 27(4): 522-543.
- Eifert, Benjamin, Edward Miguel and Daniel N. Posner. 2010. "Political Competition and Ethnic Identification in Africa." *American Journal of Political Science* 54 (2): 494–510.
- Fafchamps, Marcel. 2000. "Ethnicity and Credit in African Manufacturing." *Journal of Development Economics* 61: 205-235.
- Fearon, James. 1999. "Why Ethnic Politics and 'Pork' Tend to Go Together." Unpublished paper.
- Fischbacher, Urs and Simon Gächter. 2010. "Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments." *American Economic Review* 100 (1): 541–56.
- Gaertner, Samuel L. and John F. Dovidio. 2000. *Reducing Intergroup Bias: The Common Ingroup Identity Model*. Psychology Press.
- Gibson, Clark C., and James D. Long. 2009. "The Presidential and Parliamentary Elections in Kenya, December 2007." *Electoral Studies*, 28(3): 497–502.
- Greenwald, Anthony G., Brian A. Nosek, Mahzarin R. Banaji. 2003. "Understanding and Using the Implicit Association Test: I. An Improved Scoring Algorithm." *Journal of Personality and Social Psychology*.
- 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(4): 709–25.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the Dilemmas of Collective Action*. NY: Russell Sage.


- Henrich, Joseph, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, et al. 2006. "Costly Punishment across Human Societies." *Science* 312 (5781): 1767–1770.
- Hjort, Jonas. 2014. "Ethnic Divisions and Production in Firms." *Quarterly Journal of Economics* 129(4): 1899–1946.
- Horowitz, Donald L. 1985. *Ethnic Groups in Conflict*. Berkeley: Univ. of California Press.
- Ichino, Nahomi and Noah L. Nathan. 2013. "Crossing the Line: Local Ethnic Geography and Voting in Ghana." *American Political Science Review* 107(2): 344-361.
- Jang, Channing and John Lynham. 2015. "Where Do Social Preferences Come From?" *UIHERO Working Paper* No. 2015-8.
- Kramer, Roderick M. and Marilynn B. Brewer. 1984. "Effects of Group Identity on Resource Use in a Simulated Commons Dilemma." *Journal of Personality and Social Psychology* 46(5): 1044-1057.
- Kramon, Eric and Daniel N. Posner. 2016. "Ethnic Favoritism in Education in Kenya." *Quarterly Journal of Political Science* 11(1): 1-58.
- Laitin, David D. 1998. *Identity in Formation: The Russian-Speaking Populations in the Near Abroad*. Ithaca, NY: Cornell University Press.
- Lowes, Sarah, Nathan Nunn and James A. Robinson, J. Weigel. 2015. "Understanding Ethnic Identity in Africa: Evidence from the Implicit Association Test (IAT)." *American Economic Review: Papers & Proceedings* 105(5): 340-345.
- Marx, Benjamin, Thomas M. Stoker, and Tavneet Suri. 2016. "There Is No Free House: Ethnic Patronage in a Kenyan Slum." *Working paper*.
- Miguel, Edward, Colin Camerer, Kate Casey, et al. 2014. "Promoting Transparency in Social Science Research." *Science* 343 (6166): 30–31.
- Miguel, Edward and Mary Kay Gugerty. 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89(11-12): 2325–68.

- Nosek, Brian A., Frederick L. Smyth, Jeffery J. Hansen, et al. 2007. "Pervasiveness and Correlates of Implicit Attitudes and Stereotypes." *European Review of Social Psychology* 18(1): 36–88.
- Ogunnaike, Oludamini, Yarrow Dunham and Mahzarin R. Banaji. 2010. "The Language of Implicit Preferences." *Journal of Experimental Social Psychology* 46(6): 999–1003.
- Raven, John. 2008. "The Raven Progressive Matrices Tests: Their Theoretical Basis and Measurement Model." In *Uses and Abuses of Intelligence. Studies Advancing Spearman and Raven's Quest for Non-Arbitrary Metrics*, (eds.) John Raven & Jean Raven, 17–68. Unionville, NY: Royal Fireworks Press.
- Robinson, Amanda Lea. 2016. "Internal Borders: Ethnic-Based Market Segmentation in Malawi." *World Development* 87: 371-384.
- Rooth, Dan-Olof. 2010. "Automatic Associations and Discrimination in Hiring: Real World Evidence." *Labour Economics* 17(3): 523–34.
- Shayo, Moses. 2009. "A Model of Social Identity with an Application to Political Economy: Nation, Class and Redistribution." *American Political Science Review* 103(2): 147-174.
- Spolaore, Enrico and Romain Wacziarg. 2009. "The Diffusion of Development." *Quarterly Journal of Economics* 124(2): 469-529.
- United Nations. 2014. *World Urbanization Prospects*. New York: United Nations.
- Vigdor, Jacob L. 2002. "Interpreting Ethnic Fragmentation Effects." *Economics Letters* 75(2): 271-276.
- Westfall, Peter H., and S. Stanley Young. 1993. *Resampling-Based Multiple Testing: Examples and Methods for P-Value Adjustment*. New York: John Wiley & Sons, Inc.
- Woolridge, Jeffrey M. 2001. *Econometric Analysis of Cross Section and Panel Data*. MIT Press.


Tables and Figures:

Figure 1: Coethnic Bias in the Dictator Game and Public-good Game


Panel A: Dictator Game Contributions




Panel B: Dictator Game Beliefs



Panel C: Public-good Game Contributions

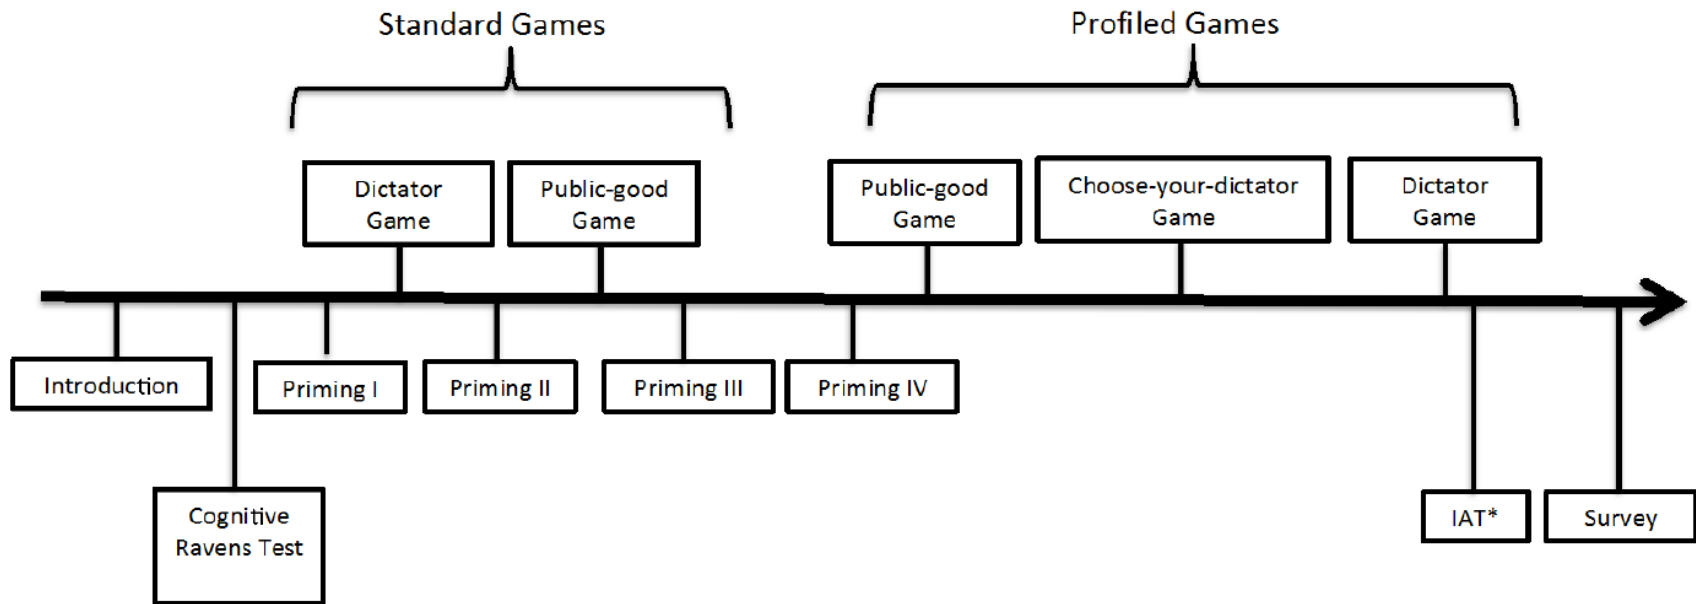


Panel D: Public-good Game Beliefs




Notes: Sample averages and 95% confidence intervals for Dictator game contributions in profiled games for Coethnic vs. Non-Coethnic Transfers (Panel A), beliefs about Dictator game contributions from the profiled Choose-Your-Dictator game (Panel B), for Public-good game contributions in profiled games for Coethnic vs. Mixed Groups (Panel C), and for Public-good game beliefs about others' contributions (Panel D). The Dictator game data in Panels A and B is from the Election round (January-February 2013), the only time the complete profiled game data was collected. The Public-good game data is pooled from both the Non-election round (July/August 2012) and the Election round, since the complete profiled game data was collected in both.

Figure 2: Experimental Laboratory Game Structure



*Indicates that we only have this data for the Election Period.


Figure 3: Coethnic Bias in the Choose-Your-Dictator Game



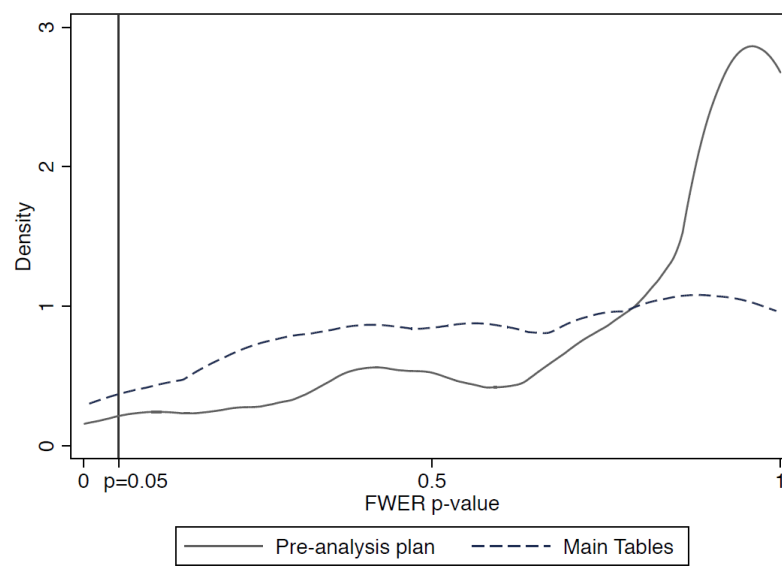
Notes: Sample averages and 95% confidence intervals for Standard and Profiled Choose-Your-Dictator (CYD) games during the Non-election round (left Panel) and Election round (right Panel).

Figure 4: Distribution of P-values in Pre-Analysis Plans versus Main Tables

Panel A: Unadjusted p-values



Panel B: FWER adjusted p-values



Notes: The distribution of p-values from the pre-analysis plan includes p-values for all hypothesis tests discussed in both pre-analysis plans (Non-election round, Election round). The distribution of p-values from the main tables includes all tables (Tables 1 through 6) in this article. The dark vertical line denotes statistical significance at the standard 95% confidence level. Panel A presents p-values unadjusted for multiple testing. Panel B presents FWER adjusted p-values. Kolmogorov-Smirnov tests on the equality of the two distributions yields a p-value of 0.014 in Panel A, and 0.005 in Panel B.

Table 1: Descriptive Statistics on the Sample and Average Game Play

	Dictator Game Transfer (Percent of Endowment)		Public-Good Game Contribution (Percent of Endowment)		Choose-Your-Dictator Game (Dictator Choice)	
	Coethnic	Non-Coethnic	Coethnic Group	Mixed Group	Coethnic	Non-Coethnic
<i>Demographics</i>						
Full Sample	35.6	35.4	46.2	46.4	31.7	25.2
Female (60% of sample)	36.8	35.8	44.9	46.2	27.3	24.3
Male (40% of sample)	33.4	34.6	48.0	46.8	38.0	26.4
Age (mean = 33.0 years)						
Below Median	33.3	32.1	41.2	43.4	32.9	27.2
Median or Above	37.8	38.5	48.0	47.5	30.5	23.2
Education (mean = 9.6 years)						
Below Median	36.5	35.4	46.9	46.2	27.6	26.4
Median or Above	34.4	35.4	45.6	46.6	35.0	24.3
Ravens Score (mean = 0.0)						
Below Median	37.5	37.5	47.3	48.7	28.5	25.2
Median or Above	33.2	32.8	45.2	44.3	34.7	25.2
<i>Ethnic Affiliation</i>						
Kikuyu (32% of sample)	36.3	36.1	45.5	47.4	26.7	25.7
Luo (21% of sample)	33.4	34.5	42.7	43.6	37.2	23.8
Observations	593	593	967	981		981

Notes: Pooled data from the Non-election Round (July-August 2012) and Election Round (January-February 2013). Values are averages, presented in percent terms. For the Dictator game, only the Election Round averages are displayed, as there is incomplete information for the profiled version of the Dictator game in the Non-election Round. For the Choose-Your-Dictator game, the table presents results for the profiled version of the game. Respondents with Kamba ethnicity are excluded from analysis, as pre-specified in the pre-analysis plan.

Table 2: Dictator Game Transfers, in Standard and Profiled Games

	Full Sample		No Prime	FWER p-value
	(1)	(2)	(3)	(4)
Coethnic Recipient	-1.51 (1.22)	-1.44 (1.22)	-1.85 (2.27)	0.518
Election Round	-5.21*** (1.41)	-5.98*** (1.43)	-6.52** (2.60)	0.008
Election Round * Coethnic Recipient	0.70 (1.88)	0.63 (1.87)	-0.83 (3.46)	0.786
Election Round * Non-coethnic Recipient	-1.41 (1.42)	-1.42 (1.42)	-4.54* (2.61)	0.153
Covariates	No	Yes	No	
Observations	2881	2881	748	
Recipient: Coethnic + Election Round * (Coethnic - Non-coethnic) [P-value]	0.60 [0.67]	0.60 [0.67]	1.86 [0.48]	

Notes: The dependent variable is the transfer in the Dictator game (in percent of the endowment). Data are pooled from the Non-election and Election Rounds. Standard errors in parentheses are clustered at the individual level. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Covariates include ethnicity indicators, a gender indicator, education controls, and the Raven's test score. FWER p-values are simulated as described in the pre-analysis plan for column (3). There was no Non-coethnic profile in the Dictator game during the Non-election Round, hence the absence of a direct "Non-Coethnic Recipient" term. The F-test in the bottom row tests the hypothesis that the average level of coethnic bias in the Election round was zero. All specifications exclude ethnic Kamba subjects, as specified in the pre-analysis plan.

Table 3: Public-good Game Contributions, in Standard and Profiled Games

	Full Sample		No Prime	FWER p-value
	(1)	(2)	(3)	(4)
Coethnic Group	2.22*	2.22*	-1.73	0.984
	(1.33)	(1.33)	(2.54)	
Mixed Group	0.65	0.65	0.32	0.988
	(1.18)	(1.18)	(2.29)	
Election Round	-2.97*	-3.95**	-6.03*	0.321
	(1.76)	(1.81)	(3.37)	
Election Round * Coethnic Group	-1.67	-1.68	2.00	0.984
	(1.89)	(1.89)	(3.39)	
Election Round * Mixed Group	2.00	2.00	2.30	0.984
	(1.85)	(1.85)	(3.59)	
Covariates	No	Yes	No	
Observations	2939	2939	763	
Coethnic Group - Mixed Group	1.57	1.57	-2.05	
[P-value]	[0.16]	[0.16]	[0.35]	
Election Round * (Coethnic Group- Mixed Group)	-3.67	-3.67	-0.30	
[P-value]	[0.027]	[0.027]	[0.93]	

Notes: The dependent variable is the contribution in the Public-good game (in percent of the endowment). Data are pooled from the Non-election and Election Rounds. Standard errors in parentheses are clustered at the individual level. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Covariates include ethnicity indicators, a gender indicator, education controls, and the Raven's test score. FWER p-values are simulated as described in the pre-analysis plan for column (3). The first F-test tests the hypothesis that the average level of coethnic bias across both the Non-election round and the Election round is zero; the second tests the hypothesis that the difference in coethnic bias across the Non-election round and the Election round is zero. All specifications exclude ethnic Kamba subjects, as specified in the pre-analysis plan.

Table 4: Choose-Your-Dictator Game Choices, in Standard and Profiled Games

	Full Sample		No Prime	FWER p-value
	(1)	(2)	(3)	(4)
Coethnic	0.25**	0.19	0.071	0.979
	(0.11)	(0.15)	(0.23)	
Profiled Game * Coethnic	0.048	0.048	0.021	0.993
	(0.11)	(0.11)	(0.21)	
Election Round * Coethnic	-0.059	-0.0013	0.15	0.979
	(0.11)	(0.12)	(0.23)	
Election Round * Profiled Game * Coethnic	0.074	0.074	0.095	0.979
	(0.15)	(0.15)	(0.27)	
Covariates	No	Yes	No	
Observations	3924	3924	1020	

Notes: Ordered Logit specification, with dependent variable 0=not chosen, 1=indifferent, 2=chosen. Data are pooled from the Non-election and Election Rounds. Standard errors in parentheses are clustered at the individual level. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include fixed effects for each Dictator-profile (12 profiles in total). The variable “Coethnic” indicates if the dictator profile is a coethnic or not. Covariates include interaction terms of the “Coethnic” indicator with a gender indicator, education controls, and the Raven's test score. FWER p-values are simulated as described in the pre-analysis plan for column (3). All specifications exclude ethnic Kamba subjects, as specified in the pre-analysis plan.

Table 5: Priming Effects in the Standard Dictator Game and Public-good Game

	Dictator Game Transfer (Percent of Endowment)		FWER p-value	Public-Good Game Contribution (Percent of Endowment)		FWER p-value
	(1)	(2)	(3)	(4)	(5)	(6)
Election Round	-6.08*** (1.18)	-7.60*** (2.38)	0.008	-2.61 (1.61)	-6.36* (3.24)	0.321
Political Competition Prime	-0.63 (1.67)	-2.19 (2.37)	0.591	0.26 (2.28)	-1.72 (3.23)	0.826
Ethnic-Cultural Prime	-1.45 (1.67)	-0.52 (2.37)	0.8	1.48 (2.28)	0.32 (3.23)	0.919
National Prime	-2.02 (1.66)	-4.43* (2.37)	0.166	-0.37 (2.28)	-4.70 (3.23)	0.462
Election * Political Competition Prime		3.14 (3.36)	0.675		3.96 (4.58)	0.664
Election * Ethnic-Cultural Prime		-1.88 (3.35)	0.815		2.29 (4.57)	0.84
Election * National Prime		4.82 (3.35)	0.408		8.68* (4.56)	0.321
Election * Ethnic-Political Prime		0.37 (2.38)	0.994		-1.77 (3.24)	0.964
Observations	1211	1362		1211	1362	

Notes: The dependent variables are the transfer in the Dictator game (in percent of the endowment) in columns 1-2 and the contribution in the Public-good game (in percent of the endowment) in columns 4-5. Data are pooled from the Non-election and Election rounds. Standard errors in parentheses are clustered at the individual level. P-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. FWER p-values are simulated as described in the pre-analysis plan for columns 2 and 5. The Ethnic-Political priming, which was only implemented in the Election Round, is not included in columns 1 and 4. All specifications exclude ethnic Kamba subjects, as specified in the pre-analysis plan.

Table 6: Coethnic Bias in the Implicit Association Test (IAT)

	Ethnic IAT: D-score				FWER
	(1)	(2)	(3)	(4)	p-values
Constant	0.079*** (0.027)	0.013 (0.058)			
Political Competition Prime		0.046 (0.081)	0.15** (0.067)	0.15** (0.066)	0.133
Ethnic-Cultural Prime		0.077 (0.085)	0.11 (0.067)	0.082 (0.066)	0.459
Ethnic-Political Prime		0.100 (0.086)	0.10 (0.067)	0.088 (0.066)	0.459
National Prime		0.12 (0.088)	0.11 (0.067)	0.091 (0.066)	0.459
Order Controls	No	No	Yes	Yes	
Covariates	No	No	No	Yes	
Observations	547	547	547	547	

Notes: The dependent variable is the D-score: the difference (in standard deviation units normalized by respondent) between the average response times in the Dissonance and Congruence IAT rounds. Data are from the Election Round only. P-values: * p<0.10, ** p<0.05, *** p<0.01. FWER p-values are simulated as described in the pre-analysis plan for column (3). Order Controls are indicators for the randomized order of (i) the ethnic and national IAT, and (ii) the Dissonance and Congruence rounds within each IAT. Columns 1 and 2 also adjust for the different sampling weights of these randomized IAT order-groups. Covariates include ethnicity indicators, a gender indicator, education controls, and the Raven's test score. All specifications exclude ethnic Kamba subjects, as specified in the pre-analysis plan. Slow-response observations are excluded, as specified in the pre-analysis plan.