

Urbanization and Interethnic Relations in Africa*

Daniel N. Posner
Department of Political Science
University of California, Los Angeles
dposner@polisci.ucla.edu

May 11, 2026

Abstract

This study leverages data from urban and rural Kenya to learn about the impact of urbanization on interethnic relations, complementing prior work that focused on urban populations alone or that characterized interethnic relations via a narrower set of survey-based outcome measures. While I find evidence that both urban and rural participants expect outgroup members to harbor ethnic biases, I find no evidence for differential altruism or cooperation across coethnics and non-coethnics, either among urban or rural participants—including when coding those participants using a measure that captures the continuous nature of urban-rural connections. I do, however, find evidence of cross-group bias along a dimension of interethnic relations rooted in like/dislike, captured via a different measurement strategy. Moreover, I find that this bias differs across urban and rural domains, with stronger preferences for coethnics among rural participants. The results point to a potential weakening of cross-group bias due to urbanization.

Word count: 9,084

*The author thanks Volha Charnysh and Ryan Enos for their comments; Dario Sidhu for his excellent research assistance; and the staff at the Busara Center for Behavioral Economics, especially Lornah Wanhonme, David Clarence, Irene Ngina, Pauline Wanjeri, Suleiman Amanela, Joseph Kimani, Cyrus Muriithi, and Chaning Jang. The project builds on joint work with Ashley Blum and Chad Hazlett, published separately. Human subjects approval was secured from Maseno University (MSU/DRPI/MUERC/00429/17) and UCLA (IRB17-000775). A preanalysis plan was registered at osf.io/mtqr3.

Introduction

The movement of populations from rural to urban areas is one of the central elements of development, with wide-ranging impacts on economic behavior, political participation, and social preferences. This paper focuses on the effect of urbanization on social preferences—in particular, its impact on interethnic relations. It asks: how are attitudes toward non-coethnics (as reflected in altruism toward members of other ethnic groups, willingness to cooperate, expectations of generosity, and like/dislike of ingroup versus outgroup members—summarized here under the label “ethnic bias”) shaped by exposure to urban life? It addresses this question by contrasting evidence for these attitudes across urban and rural dwellers in Kenya, an ethnically heterogeneous country undergoing rapid urbanization.

There are many reasons to think that attitudes toward non-coethnics may be affected by moving to an urban environment, although the directions of these effects are in many cases unclear and contradictory. Urban environments are more diverse than rural areas, implying that urban residents are in much more frequent contact with members of other ethnic groups.¹ Such contact has been shown to reduce social distance (Scacco and Warren 2018; Mousa 2020; Grady et al. 2023) but also to generate competition and intergroup conflict (Young 1976; Bates 1983). Urban residents tend to be wealthier, better educated, and more exposed to the media than their rural counterparts. These factors again have an ambiguous impact on cross-group relations: they may promote the emergence of cosmopolitan identities and make ethnic affiliations less salient (Lerner 1958; Brooks 2014) or they may heighten the importance of ethnic group memberships by increasing competition over jobs and public resources (Melson and Wolpe 1970; Bates 1983).² Differences in the principal modes of economic activity in urban and rural areas (Green 2023), in the proximity and salience of traditional leaders (Koter 2013), and in the strategies that elites use to mobilize political coalitions (Kim and Horowitz 2022) have also been hypothesized to generate differences in the strength of ethnic ties and the tenor of intergroup relations in urban and rural settings.

The variety of channels through which urbanization may impact interethnic relations,

¹Intergroup contact can be both broad (i.e., exposure to non-coethnics in one’s neighborhood) and deep (i.e., inter-ethnic marriage), with likely different impacts on outgroup bias (Chakraborty et al. 2024).

²Urbanization may also have the effect of neither heightening or reducing the salience of ethnic differences but shifting the dimension of ethnic identity that matters, from narrower “tribal” identities to broader regional or linguistic identities (Gluckman 1960; Posner 2005; Green 2023).

along with the inconsistency in the theoretical expectations about the direction of these effects, has given rise to a number of investigations into this relationship. These studies vary in the specific outcomes they consider, the empirical approaches they employ, and the findings they report. For example, [Robinson \(2014\)](#) and [Kramon et al. \(2022\)](#) draw on survey data to study the importance that urban and rural dwellers attach to their ethnic identities, finding weaker ethnic identification among urban residents. [Nathan \(2016\)](#) uses survey and administrative data to test whether ethnic voting diminishes with urbanization, finding that the tendency for urban voters to support coethnic candidates and parties varies with local demography and socioeconomic status. [Kim and Horowitz \(2022\)](#) show that urban voters penalize politicians for making overt ethnically-based appeals and that this tendency increases with length of residence in the city, suggesting that rural voters may be less likely to punish ethnic pandering. [Kramon et al. \(2022\)](#) and [Lyon \(2023\)](#) employ survey data to show that interethnic trust is lower for city dwellers than for individuals living in rural areas. [Barriga et al. \(2023\)](#) and [Peterson \(2016\)](#), whose approach is the closest to the one adopted here, employ economics games to study some combination of interethnic trust, inter-ethnic cooperation, and differential generosity to ingroup and outgroup members among rural and urban Africans. Whereas [Barriga et al. \(2023\)](#) find no evidence for ethnic bias in either urban or rural populations, [Peterson \(2016\)](#) reports evidence of ethnic bias in locations ranging from rural villages to large cities. [Berge et al. \(2020\)](#) also use economics games to investigate ethnic bias in Nairobi, Kenya. Although they do not explicitly study differences across urban and rural areas, they do find evidence that recent urban migrants display more cross-group bias than longer-term urban residents—a finding suggestive of stronger ethnic bias in rural areas, and of the potentially bias-reducing effects of moving to an urban location.

The present study joins these efforts to understand the impact of urbanization on interethnic relations. Like some of the papers just discussed, I use economics games—including a Dictator Game (DG) and a Public Goods Game (PGG)—to study differential generosity and cooperation across coethnic and non-coethnic partners. Economics games have an advantage over survey-based approaches to capturing interethnic relations because they provide incentivized, behavior-based measures of cross-group interaction rather than self-reported affinities, beliefs, and attitudes, often about hypothetical situations. The study also departs from prior research in several noteworthy ways.

First, in measuring the character of interethnic relations, I supplement the findings from the economics games with the results from an Affect Misattribution Procedure (AMP), an experimental task borrowed from social psychology that captures affect or like/dislike (Payne et al. 2005), a distinct and important dimension of a person’s orientation toward ingroup versus outgroup members (Russell 2003). Beyond the additional outcome that the AMP allows us to capture, including it alongside the DG and PGG also makes it possible to investigate whether social desirability bias (in the context of the present study, relating to the desire to try to hide one’s favoritism toward coethnics and/or disdain toward non-coethnics) may be affecting the results in the economics games. This is because the AMP, as a speeded task, makes it harder to conceal one’s biases than in the DG or PGG, which provide opportunities for participants to reflect on their choices before conveying their decisions (Hazlett and Berinsky 2018; Blum, Hazlett and Posner 2021). A finding of ethnic bias in the AMP alongside a finding of no ethnic bias in the DG or PGG would suggest that participants may be taking advantage of the opportunity to censor their responses in the two behavioral games.

Comparing the DG/PGG and AMP results across urban and rural participants also provides insight into whether social desirability bias may vary across urban and rural populations. A finding that it does would hold potentially important implications for our understanding of how urbanization shapes intergroup relations. For example, it might point to the impact of urbanization lying less in the reduction of cross-group prejudices than on the socialization to trying to hide such prejudices in public interactions. The net effects of these two scenarios would be the same in terms of expressed prejudice, but the implications for behavior when one’s actions are not publicly visible—for example, in the voting booth, in private business transactions, or in other non-public settings—would be quite different.

A second departure from prior work involves moving beyond a dichotomous distinction between urban and rural. Recent evidence suggests that 70-90% of urban household heads in Nairobi were born in a rural area (Emina et al. 2011; Onyango, Cruch and Owour 2021). While some of these urban dwellers can be considered fully “urban,” in the sense of having completely disconnected themselves from rural life, others retain strong links to their rural homes (Gugler 2002; Robinson and Riedl 2025). Many migrants to urban areas engage in frequent contacts with their rural kin, plan to return someday to live in their

village, and associate mainly with coethnics from their home region. Such migrants are likely less “urban” in their orientations and behavior than individuals born in an urban area who marry someone from a different ethnic group and maintain few connections to their ancestral village—despite the fact that all such participants would be coded as urban in most prior analyses. The same is true for the “ruralness” of participants recruited from rural areas. A resident of a rural district who frequently visits and communicates with urban-based family members or who previously lived in an urban part of the country will likely be less “rural” in their orientations and behavior than someone who never traveled out of their village and has few connections with urban-based kin—even as both would be treated as rural in most studies.³ To better capture the non-binary nature of “urbanness” and “ruralness,” I draw on information about the connections that urban and rural participants maintain with their rural and urban kin to create continuous measures of these characteristics, which I use to supplement the dichotomous coding of participants based on the places where they were recruited. Among urban participants, I also create a measure of “exposure” that captures the extent to which the participant has contact with members of other ethnic groups. I re-analyze all of my results using these continuous measures of urban and rural life.

Third, I go beyond documenting cross-group altruism (via the DG) to investigate *expectations* of altruism using an additional economics game, the Choose Your Dictator (CYD) game. Analyzing expectations of outgroup members’ behavior is important insofar as the salience of ethnicity in politics and social life may depend as much or more on how one thinks non-coethnics will behave than on one’s own predispositions toward non-coethnics. For example, voters disinclined to support ingroup candidates simply because they share an ethnic tie may nonetheless vote along ethnic lines if they believe that others will only support coethnics or that politicians will favor members of their own groups (Posner 2005). Nathan (2016) makes precisely this argument in his study of urban Ghana, where he shows that ethnic voting is driven by expectations of ethnically biased allocations of goods rather than by the importance people attach to ethnic differences themselves.⁴ Incorporating the results from the CYD game allows me to evaluate whether urban and rural populations differ in their expectations of ethnic bias by non-coethnics. If they do, it will suggest that

³This is not just an African phenomenon. Many US survey respondents identify as “rural” even though they do not live in rural areas (Trujillo 2024).

⁴Drawing on data from urban India, Spater (2022) shows similarly that outgroup prejudice is unrelated to coethnic voting.

an important impact of urbanization may lie in how it alters expectations of how outgroup members will behave rather than simply in changes in one’s own attitudes or behavior toward non-coethnics.

To preview my findings, I find little evidence for differential altruism or willingness to cooperate across ethnic lines among either urban or rural populations—both when treating those populations as distinct monolithic blocks based on the locations from which participants were recruited and when I adjust the urban/rural coding to account for the connections that urban participants maintain with rural kin, and vice versa. These null findings validate the results reported in prior work using similar economics games to measure ethnic bias in urban Kenya (e.g., [Berge et al. \(2020\)](#), [Blum, Hazlett and Posner \(2021\)](#), [Barriga et al. \(2023\)](#)) but broaden them to show that they extend to rural populations as well.

Notwithstanding the absence of evidence for cross-group bias in the DG and PGG, I find consistent evidence for ethnic bias in the AMP, with weaker dislike for outgroup members among urban than rural participants. The implication is that urbanization may be associated with a lessening in outgroup bias, at least for the particular dimension of bias that the AMP captures. To the extent that the different results across the AMP and the economics games stem from the AMP’s ability to uncover biases against non-coethnics that participants are able to conceal in the DG and PGG, the implication is that cross-group biases do in fact exist in the study population, even if they are not well captured using the traditional economics games-based measurement tools.⁵ The finding of stronger out-group bias in the AMP among rural than urban participants suggests that the inclination to hide one’s positive bias toward coethnics and/or negative bias toward non-coethnics may actually be weaker in urban areas, running against the hypothesis of urban socialization to conceal distasteful outgroup biases.

I also find evidence for expectations of in-group favoritism in the DG, even though study participants are not in fact making more generous allocations to coethnics. This finding, which holds equally among urban and rural participants, points to an important aspect of how ethnicity affects social and political life: by shaping expectations of how others will behave as much, and perhaps more, than influencing how others actually do

⁵For a similar finding, see [Blum, Hazlett and Posner \(2021\)](#). [Berge et al. \(2020\)](#), who find no evidence of ethnic bias in the DG and PGG, find evidence of ethnic bias using an Implicit Attitude Test, a speeded task similar to the AMP in its ability to pick up biases that participants may wish to conceal.

act.

Taken together these results have important implications for the differences in interethnic relations in urban and rural Africa and, by extension, hold important lessons regarding the impact of urbanization. It would appear that exposure to urban life does not alter altruism toward coethnics versus non-coethnics, willingness to cooperate along ethnic lines, or expectations about the generosity of ingroup versus outgroup members. But it does seem to be associated with changes in the like or dislike that individuals feel toward ingroup versus outgroup members. The fact that such ingroup bias is weaker in urban areas suggests that urbanization may reduce social distance and foster intergroup cohesion—even if not through the outcomes that prior studies have measured.

Research sites and sample

Estimating the causal effect of urban migration is challenging because the kinds of people who migrate from rural to urban areas are likely different from the kinds of people who do not. This means that any differences we observe across urban and rural populations could stem either from the impact of exposure to urban life or from the distinctive individual characteristics of the people who chose to migrate.⁶ Researchers seeking to estimate the impact of urbanization have dealt with this selection problem in several ways.

Perhaps the strongest response in the political science literature is offered by [Kramon et al. \(2022\)](#), who leverage data from the Kenya Life Panel Survey to follow individuals over time as they migrate from rural to urban areas, using individual-level fixed effects to partially address concerns about selection into migration. However, even this strategy has its drawbacks. First, it requires longitudinal data gathered over many years that, in the case of the [Kramon et al. \(2022\)](#) study and most others like it, is collected within a single rural setting and then employed to follow survey respondents who migrate.⁷ This limits the ability to generalize to migrants from other rural areas, for whom the impact of urban exposure may differ due to ethnic affiliation, socioeconomic status, distance from the city to which they have migrated, or other factors. A second drawback is that the individual

⁶For a particularly compelling demonstration of the consequences of selection bias in studies of urban migration, see [Hicks et al. \(2020\)](#).

⁷Because the leverage in these studies comes from the subset of original survey respondents who migrate, the effective sample size in these analyses tend to be small. In the [Kramon et al. \(2022\)](#) study, just 23 percent of respondents moved to one of the country’s two major urban areas across survey rounds. In [Beegle, De Weerd and Dercon \(2011\)](#), another prominent study using a similar panel approach, only 138 individuals were observed in both rural and urban areas.

fixed-effects approach employed in such studies does not allow for estimating effects among non-movers, whose response to the urban environment may have been different from the responses of those who selected into migration.

An alternative approach, employed in [Berge et al. \(2020\)](#), [Kim and Horowitz \(2022\)](#), [Lyon \(2023\)](#), and [Robinson and Riedl \(2025\)](#), is to study urban populations alone and make inferences about the effects of urban life by comparing recent migrants and long-term urban residents. While this approach sidesteps the problem of selection into migration (as everyone in the study, or their parents, has already chosen to migrate), it shares the drawback of the longitudinal approach in omitting non-movers from the analysis. It also assumes that length of residence in the urban domain is correlated with the intensity of the urban “treatment.”

The strategy adopted here (as in [Bannon, Miguel and Posner \(2004\)](#), [Green \(2013\)](#), [Robinson \(2014\)](#), [Peterson \(2016\)](#), [Barriga et al. \(2023\)](#), and [Lyon \(2023\)](#)) is to compare outcomes across urban and rural populations. While selection bias can be addressed to some degree in this approach by controlling for the individual characteristics that differ across urban and rural participants, this is an imperfect solution—and the results presented below should be interpreted with this caveat in mind. The contribution here, therefore, is not to overcome the challenges to studying the causal effects of urbanization on interethnic relations but to offer insight into those effects by, for one of the first times, using economics games to measure ethnic bias across urban and rural populations in an African setting.⁸

To sharpen the focus on politically relevant inter-ethnic differences, and also to maximize study power, I limit the sample to members of the Kikuyu and Luo ethnic groups, the two most politically relevant and historically antagonistic groups in Kenyan national politics ([Hornsby 2013](#); [Africa Report 2018](#)). All out-group partners presented to Kikuyu participants in the economics games are Luo, and all out-group partners presented to Luo participants are Kikuyu. The study was conducted a decade after the violence surrounding Kenya’s disputed 2007 presidential election and shortly before another election in which the leading candidates were again Kikuyu and Luo. Ethnic tensions between members of these groups were thus high at the time of the laboratory sessions.

⁸The only prior studies that employ economics games to measure ethnic bias across both urban and rural populations are [Peterson \(2016\)](#) and [Barriga et al. \(2023\)](#). [Bannon, Miguel and Posner \(2004\)](#), [Robinson \(2014\)](#), [Green \(2013\)](#), and [Lyon \(2023\)](#) draw on survey data to compare the strength of ethnic identification among urban and rural Africans, but this is a much narrower aspect of inter-ethnic relations than the outcomes treated here.

I recruited participants from Nairobi (specifically, the multi-ethnic compounds of Kibera and Kawangware, $N=318$) and the rural districts of Kirinyaga ($N=171$) and Siaya ($N=164$), which are overwhelmingly populated by Kikuyus and Luos, respectively. The locations of the sites are shown in Figure 1. This sampling strategy makes it possible to compare relations across members of the same two ethnic groups living in both urban and rural settings.

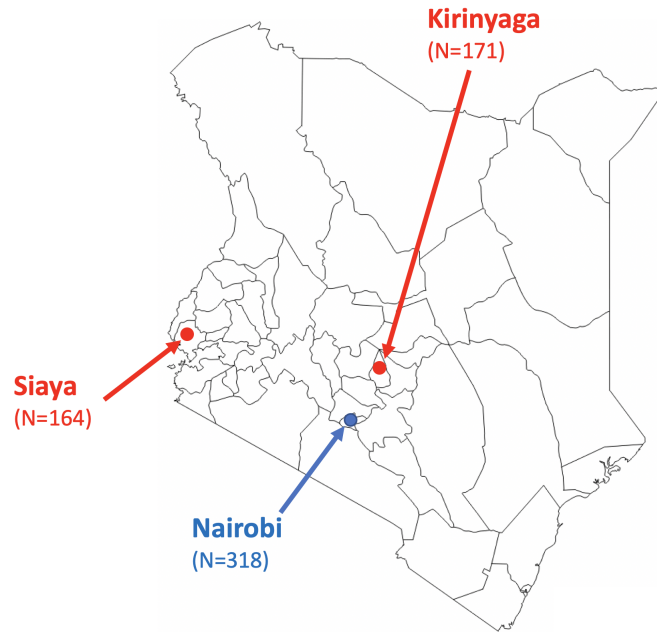


Figure 1: Research sites

The urban sample, in addition to being limited to equal numbers of Kikuyu and Luo participants, was stratified by age, gender, and length of residence in Nairobi. The two rural samples were stratified by age, gender, and frequency of contact with Nairobi. As expected, the urban sample is better educated and somewhat younger than the rural sample (see Appendix Table A1.) For this reason, as pre-registered, I control for age and education in the analyses that follow.

Subjects participated in the study activities in all three research sites using a mobile lab facility run by the Busara Center for Behavioral Economics. In Nairobi, participants were recruited from the Busara Center’s pre-existing subject pool.⁹ In the two rural districts, participants were recruited each day from different non-adjacent villages in order to reduce the likelihood (otherwise high) that people in the community would talk with one another,

⁹To ensure comparability with the two rural sites, where the Busara Center had not previously operated, only participants who had not participated in a prior experimental study were admitted.

which might have caused subjects in later sessions to have been briefed regarding the purpose of the research.

The lab sessions took place in June-July 2017. Given research indicating that proximity to national elections may affect the salience of ethnic identity (Eifert, Miguel and Posner 2010), the lab sessions were scheduled more or less simultaneously in all three sites to equalize the distance from the coming general elections, which were held on August 8.¹⁰ Participants were paid 200 KSh (about \$2) for attending the session, plus an additional 50 KSh if they arrived on time. They also received a reimbursement for their transport expenses, as well as whatever they earned from their play during the experimental games. Payments were made via MPESA mobile money transfer after the conclusion of the lab session, both to reduce opportunities for learning during the course of the experiment and to prevent participants' knowledge of their accumulating earnings from altering their risk preferences as the trials progressed.

Measuring interethnic relations

My main measure of interethnic relations is the difference in the mean contributions across coethnic and non-coethnic partners in the DG and PGG. Economics games like the DG and PGG offer powerful tools for measuring cross-group bias and have been increasingly adopted by researchers studying intergroup relations in Africa and beyond.¹¹ In the DG, participants are given a sum of money (in this case KSh 100) and asked to divide it between themselves and a receiver, keeping whatever portion they choose. The share of the endowment they voluntarily give away serves as a measure of altruism (Kahneman, Knetsch and Thaler 1986). In the PGG, participants are paired with two other players. They are given a sum of money (again, KSh 100) and asked to decide how much to keep for themselves and how much to give to a group fund in which contributions are doubled and then shared equally among all three players. Each player gets to keep their share of the group fund, plus whatever they decide not to contribute. Because contributions in the PGG depend on expectations of what the two other players will contribute, they serve as

¹⁰The Nairobi lab sessions took place from July 7-19; the Kirinyaga sessions from July 7-17; and the Siaya sessions from July 8-15.

¹¹In addition to the already cited studies, prominent examples in the African context include Habyrimana et al. (2009), McCauley (2014), Burbidge and Cheeseman (2017), Jeon, Johnson and Robinson (2017), and Haushofer et al. (2023).

a measure of cooperation.¹² Ethnic bias can be measured using these games by observing whether participants are more altruistic in the DG when paired with a coethnic versus a non-coethnic partner, and more cooperative in the PGG when the two other players are coethnics versus non-coethnics.

Each lab session began with an on-screen literacy test, which involved distinguishing between positive and negative words in the participant’s local language. Participants also completed a set of Raven’s matrices, designed to measure cognitive ability (Raven 2008). Instructions for the games and tasks were provided on the computer screens in both English and Swahili (in Nairobi) and in Kikuyu and Dholuo (in Kirinyaga and Siaya, respectively), in addition to being read aloud in the appropriate language by a trained lab leader. For parts of the session where subsets of participants received different instructions or where they completed games or tasks at their own pace, participants put on headphones and received instructions via recordings in Swahili, Kikuyu, or Dholuo, as appropriate.

In the DG, participants were shown a photograph of the receiver and given his first name, age, education, and hometown, with the hometown selected so as to provide an indication of the receiver’s likely Kikuyu or Luo origins (see Figure 2).¹³ Hometown was used as a cue for ethnicity rather than the labels “Kikuyu” or “Luo” because using overt ethnic labels was potentially inflammatory in the pre-election context in which the study was conducted.¹⁴ Using such labels also risked signaling that the experiment was about ethnic discrimination, which may have generated social desirability bias. The inclusion of information about the receiver’s age and education was meant to further obscure the ethnic focus of the study.

The DG was played twice: once with a receiver from a coethnic hometown and once with a receiver from a non-coethnic hometown, with the order randomized. Participants were informed that the receiver was an actual person to whom the money they allocated would be given, and money was in fact transferred to the person shown in the picture at

¹²The incentives facing players in the PGG, as well as the interpretation of its results, are complex. For useful discussions, see Ledyard (1995); Camerer and Fehr (2004) and Habyarimana et al. (2009).

¹³All of the images used in the DG and other games were of men in order to reduce the number of factors that might affect participants’ response to the stimulus. The photographs were pre-screened to be ethnically ambiguous so that the hometown cue, rather than phenotypical characteristics, would influence the participant’s perception of the receiver’s ethnic group membership. The name, age, and education were fixed to each photograph, but the hometown was randomized across participants. Each participant saw each photograph and associated profile just once. Across all of the sessions, each photograph was presented equally often as belonging to someone from a Kikuyu or Luo hometown.

¹⁴Berge et al. (2020) employ a similar strategy of cueing ethnicity via hometown and report that the vast majority of participants could correctly identify the intended ethnic backgrounds of their partners based on the hometowns ascribed to them.

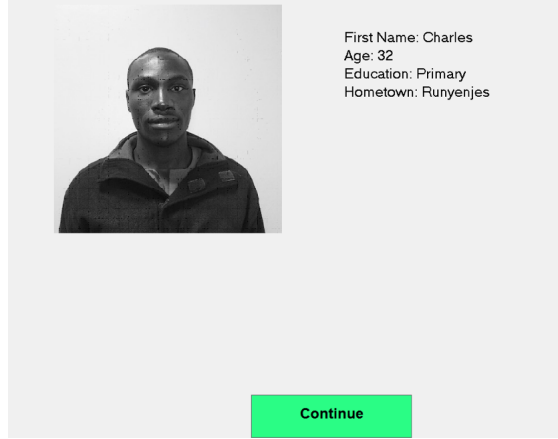


Figure 2: Information Provided About the Other Players in the DG and PGG

the end of the experiment.

As in the DG, participants in the PGG were shown photos and provided the names, ages, education levels, and hometowns of the other players, as shown in Figure 2. Participants played three rounds of the PGG: one with a coethnic group in which both of the other players were from coethnic hometowns, one with a non-coethnic group in which both of the other players were from non-coethnics hometowns, and one with a mixed group. The order in which these three rounds were played was randomized. As in the DG, the other players were real and were paid based on the outcome of the game, as determined by the participant’s contribution and contributions made by the other players in anonymous PGGs they played at the time their photos were taken, prior to the experimental sessions.

The outcome of interest in the DG is the percentage of the endowment transferred to the partner; in the PGG it is the percentage contributed to the group fund. I estimate the following pre-registered equation:

$$Y_i = \beta_0 + \beta_1 Coethnic_i + \beta_2 X_i + \epsilon_i \quad (1)$$

where Y_i is the transfer (in percent of the endowment) by participant i to his/her partner(s). $Coethnic_i$ is an indicator variable for being paired with a coethnic partner (or, in the PGG, being in a coethnic group),¹⁵ X_i is a vector of pre-registered control variables (including gender, ethnicity, age, years of education, and Raven’s score), and ϵ_i is the error term, clustered by individual. The main test of ethnic bias is the hypothesis $\beta_1 \neq 0$.

To test for the hypothesis that ethnic bias varies across urban and rural settings, I

¹⁵The omitted category in the DG is being paired with a non-coethnic partner. In the PGG the omitted category is being in a non-coethnic group (rounds with mixed groups are not analyzed).

estimate the following equation:

$$Y_i = \beta_0 + \beta_1 Coethnic_i + \beta_2 Urban_i + \beta_3 Coethnic_i * Urban_i + \beta_4 X_i + \epsilon_i \quad (2)$$

where $Urban_i$ indicates whether participant i was recruited from an urban location (in this study, from Nairobi). The hypothesis of interest is that $\beta_3 \neq 0$.¹⁶

Main Results

The results for the DG are reported in the first four columns of Table 1. On average, DG participants give about a third of their endowment to their partner. However, there is no statistically significant difference in the amount they give to coethnics versus non-coethnics (column 2)—i.e., no evidence of differential altruism along ethnic lines. Nor is there evidence that urban participants are more generous than rural participants, or that ethnic bias is present in either the urban or rural domains alone (column 3). Adding controls for gender, ethnic group, education, age, and Raven’s score does not change these null findings (column 4).

We see similarly null results for the PGG (columns 5-8). On average, participants contribute about half of their endowment to the group fund. However, we see no differences in average contributions when the two other group members are coethnics versus non-coethnics, and no differences in cooperation across urban or rural participants.

Table 1: Dictator and Public Goods Game Results

	Dictator Game				Public Goods Game			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Intercept	0.328*** (0.010)	0.320*** (0.012)	0.319*** (0.013)	0.373*** (0.037)	0.469*** (0.013)	0.486*** (0.015)	0.489*** (0.018)	0.554*** (0.049)
Coethnic		-0.009 (0.014)	-0.008 (0.019)	-0.008 (0.019)		-0.004 (0.018)	-0.011 (0.025)	-0.011 (0.025)
Urban			0.018 (0.019)	0.027 (0.020)			-0.041 (0.025)	-0.019 (0.026)
Coethnic * Urban			-0.004 (0.027)	-0.004 (0.027)			0.014 (0.036)	0.014 (0.036)
N	1310	1310	1310	1310	1310	1310	1310	1310
Covariates				yes				yes

Outcome is the share of the endowment given to the other player (DG)/contributed to the group fund (PGG). Covariates are gender, ethnic group, education, age, and Raven’s score (demeaned). *p<0.1; **p<0.05; ***p<0.01.

¹⁶In the pre-analysis plan, I also pre-registered the intention to investigate whether priming participants to the possibility of election violence affected either levels of ethnic bias or differences in ethnic bias across urban and rural participants. I found no impact of priming to violence, so I do not report these analyses in the paper.

The absence of ethnic bias in the urban sample in either the DG or the PGG accords with the findings reported in [Berge et al. \(2020\)](#), [Blum, Hazlett and Posner \(2021\)](#), and [Barriga et al. \(2023\)](#), who all find no evidence of cross-group bias in these games in similar, Nairobi-based samples. The fact that these null results extend to the Kirinyaga and Siaya subsamples suggests that the absence of ethnic bias is not simply an urban phenomenon— notwithstanding the suggestive evidence in [Berge et al. \(2020\)](#) that it might be.¹⁷ It would appear not to be the case that generosity toward non-coethnics and willingness cooperate with people from different—and even historically antagonistic—ethnic groups is something learned through exposure to the urban environment. The absence of bias against outgroup members is equally apparent in both the urban and rural samples.

Social desirability bias

A significant concern in studies using economics games to measure cross-group prejudice is the susceptibility of such games to social desirability bias ([Cilliers, Dube and Siddiqi 2015](#); [de Quidt, Haushofer and Roth 2018](#); [Hamilton, Hazlett and Posner 2025](#)). In a context, like Kenya, where social norms discourage discrimination across group lines, participants may have incentives to try to hide their prejudices toward outgroup members by altering their outwardly visible behaviors. These incentives that will only be heightened in a laboratory setting where participants know that their allocation decisions are being observed. A weakness of economics games like the DG and PGG is that they provide opportunities for participants to reflect before making their decisions, allowing them to censor their responses so as not to violate social norms against outgroup discrimination. To the extent that participants take advantage of this opportunity, it may generate null findings, even if cross-group biases are present.

Investigating this possibility is important for confirming the validity of the null results presented thus far. It is also important for providing insight into whether social pressure against outgroup discrimination may vary in urban and rural areas. As noted earlier, a finding that it does would suggest that differences in bias across the two domains might be driven less by reductions in outgroup prejudice due to exposure to urban life than by

¹⁷The divergence between these findings and the significant results reported in [Peterson \(2016\)](#) may stem from the modified version of the DG he employs, which involves asking participants to choose between two applicants for a small business loan. The set-up forces a binary choice between ingroup and outgroup applicants and prevents the dictator from keeping any of the endowment.


socialization to try to conceal such prejudice in publicly visible behavior. That would be an important insight.

My strategy for addressing these questions is to supplement the evidence for ethnic bias collected via the DG and PGG with evidence from a third experimental task, the Affect Misattribution Procedure (AMP). Developed by [Payne et al. \(2005\)](#), the AMP is a speeded, priming-based task that is much less vulnerable to social desirability bias than economics games like the DG and PGG. The AMP works by leveraging the tendency for associations triggered by a particular stimulus to spill over and become associated with a second stimulus—in the case of the AMP task, an ambiguous image with no affective content. Because participants are unaware of the influence that the first stimulus has on their response to the second, researchers can learn about the participants’ affective reaction to the first stimulus by asking them about their judgments about the second. This makes the AMP a valuable tool for measuring outgroup prejudice in a context where participant may wish to hide their biases.¹⁸ Given the absence of cross-group bias found in the DG and PGG, a finding of ethnic bias in the AMP would be consistent with the presence of underlying biases that participants are motivated and able to hide in the two economics games. A finding of bias in the AMP in the urban but not rural samples would be consistent with differential motivation to hiding one’s prejudices in these two domains—potentially due to the socialization against open displays of ethnic bias in the urban setting.

In the version of the AMP employed in the present study, participants are first shown a priming word for 200 ms, which they are told to ignore, and then shown an innocuous image (a Chinese character) for 500 ms, followed by a white noise mask. The participant is asked to rate the image as pleasant or unpleasant by pressing the appropriate key on a keyboard. The priming word is a Kikuyu or Luo hometown in one version of the task (AMP-Town) and a region (“Central” or “Nyanza,” the former strongly associated with Kikuyus; the latter with Luos) in the other (AMP-Region). The idea is that whatever positive or negative affect the participant may have toward coethnics or non-coethnics (as triggered by the priming word) will be reflected in their response to the Chinese character—all without the participant being aware that they are providing information about their disposition toward ingroup and outgroup members. A diagram depicting the flow of an

¹⁸For a further discussion of the AMP and other misattribution-based tasks as tools for measuring ethnic bias, as well as a comparison of results using the DG, PGG, and AMP in a similar setting, see [Blum, Hazlett and Posner \(2021\)](#).

AMP trial is provided in Figure 3.



Source: [Blum, Hazlett and Posner \(2021\)](#).

Figure 3: The Affect Misattribution Procedure (AMP)

Each version of the AMP contained 30 trials, with the priming words randomized. As pre-registered, the results presented here combine the town and region versions and report outcomes averaged across all 60 trials. For each participant, I calculate an “ethnic effect” measuring how responses differ across trials depending on whether the participant received a coethnic or non-coethnic prime. I then take the average of these effects across all participants as my measure of ethnic bias.

The AMP results are presented in Table 2. Unlike in the DG and PGG, we see strong effects of ethnicity. The significant positive estimates for the intercept term indicate that participants are significantly more likely to say that they find the Chinese character pleasant when they are first primed to be thinking about a coethnic (column 1). Since the affect attached to the Chinese character is mediated by the priming to think about a Kikuyu or Luo hometown or region, the interpretation is that participants have more positive affect toward the ingroup than the outgroup.

Viewed alongside the null results in Table 1, these findings point to a one of two different implications. The first is that an underlying bias against non-coethnics exists but is not detected in the DG or PGG, possibly because these measurement tools afford participants the ability to hide their biases.¹⁹ The second is that the AMP is simply measuring something different from the DG and PGG, and that cross-group biases exist with respect to what the AMP measures (like/dislike) but not with respect to what the

¹⁹The underlying bias could equally be positive bias towards coethnics, rather than negative bias against non-coethnics. It is impossible to know from the experimental design which is driving the effects we see.

Table 2: Affect Misattribution Procedure Results

	(1)	(2)	(3)
Intercept	0.121*** (0.014)	0.132*** (0.017)	0.134*** (0.027)
Urban		-0.033 (0.029)	-0.060* (0.031)
<i>N</i>	377	377	377
Covariates			yes

Outcome is the “ethnic effect,” the estimated probability of finding a Chinese character pleasant after exposure to a coethnic prime minus the estimated probability of finding the Chinese character pleasant after exposure to a non-coethnic prime. As pre-registered, results of the AMP-Town and AMP-Region are combined. Sample is limited to participants who correctly answered half or more of the questions in the six question literacy test administered at the start of the lab session and whose reaction times were ≤ 200 ms, as pre-registered. Covariates are gender, ethnic group, education, and age. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

DG and PGG measure (altruism and cooperation, respectively).²⁰

From the standpoint of understanding the impact of urbanization on interethnic relations, the more important result in Table 2 is the suggestive evidence for a weakening of the prejudice toward outgroup members among urban participants, as indicated by the negative sign on the *urban* variable in columns 2 and 3. The effect is small and of borderline statistical significance, and emerges only in the analysis that includes controls for gender, ethnic group, education, and age.²¹ But it suggests that the negative affect toward non-coethnics may diminish somewhat with exposure to urban life.

The AMP findings also speak to the question of whether urbanization may lead to socialization toward trying to hide one’s biases against outgroup members. If that were the case, we would have expected to see stronger evidence of outgroup bias in urban areas in the AMP than in the DG and PGG. The fact that we see the opposite weighs against this hypothesized effect of urban exposure.

Variation within the urban and rural domains

Thus far, my strategy for investigating whether urbanization is associated with changes in interethnic relations has been to compare outcomes among participants recruited from Nairobi and in the rural districts of Kirinyaga and Siaya. The assumption underlying this

²⁰For an elaboration of this argument, see Blum, Hazlett and Posner (2021) and Hamilton, Hazlett and Posner (2025).

²¹The fact that the findings are statistically significant in the specification that includes the covariates is reassuring, as this is the specification that at least partially addresses concerns about selection into migration by controlling for individual characteristics that vary on average across urban and rural participants.

approach is that participants from Nairobi can be considered urban and participants from Kirinyaga and Siaya can be considered rural. However, as noted, this dichotomization belies significant variation in the linkages that urban dwellers maintain with rural family members, and vice versa.²²

To better capture this within-category variation, I collected data on the connections that urban study participants maintain with their rural kin and the connections that rural study participants maintain with urban-based family members. These include questions asked of Nairobi participants about whether they have a place outside of Nairobi they consider home, how often they visit that place, how frequently they send money to family members living there, how well they know what is happening there, whether they have property or investments there, and whether they would like to be laid to rest there. These questions were inspired by survey items developed in [Robinson and Riedl \(2025\)](#), who explore the multiple forms of rural connection that urban dwellers maintain with their rural kin. To get a sense of the connections that rural-based study participants have to urban areas, I asked questions about whether they have family members living in Nairobi or another urban part of the country, how frequently they were visited by (or themselves visited) those urban-based family members, whether family members in Nairobi send them money, and whether they have ever lived in an urban area. I combine the responses to create pre-registered averaged z-score indexes of “urbanness” and “ruralness” for urban and rural study participants, respectively.²³ For participants recruited in Nairobi, higher values of the “urbanness” measure indicate weaker attachment to their rural home; for participants recruited from the two rural sites, higher values of the “ruralness” measure indicate weaker connections to urban life. The left and middle panels of [Figure 4](#) plot the distributions of the two indexes. The broad variation we see in both figures clarifies what is lost when we fail to acknowledge the variation in linkages that urban participants maintain with rural areas, and vice versa.²⁴

²²The dichotomization also belies the fact that cities come in different sizes, with different implications for what the urban experience entails, a point emphasized in [Lyon \(2023\)](#).

²³The full set of questions used in creating these indexes, along with descriptive statistics for each component, is provided in [Appendix B](#).

²⁴Quite apart from the variance in urbanness and ruralness shown in [Figure 4](#), the *levels* of urban-rural connections revealed in the data are striking. As reported in [Appendix B](#), 76 percent of urban participants said they have a specific place outside Nairobi that they consider “home.” Seventy percent report communicating with someone in that place at least once a week, and 66 percent report planning to live there permanently in the future. Seventy-nine percent characterize being in good standing as a member of their home village as important, very important, or essential. Meanwhile, among rural participants, 84 percent report having family members living in Nairobi or other urban areas, 87 percent report having been visited by these urban family members at least once in the past two years, and 50 percent say they




Figure 4: Variation in “Urbanness,” “Ruralness” and “Exposure”

Accounting for these continuing connections is important for understanding the impact of urban or local location on people’s attitudes and behavior—including bias toward non-coethnics. But, for urban migrants, moving from a rural to an urban place of residence is about more than just removing oneself from one’s rural community. It also involves exposing oneself to a diverse environment in which interactions with people from other groups become commonplace. This is a critical feature of urban life that is not well captured in the “urbanness” measure. To test for the potentially different impact of this distinct implication of urban residence, I also create a separate index of “exposure” built from questions about the frequency of personal contact with members of the ethnic outgroup, the ethnic diversity of the neighborhood in which the participant lives, and whether they or their parents are married to a person from a different ethnic group.²⁵ The distribution of the “exposure” index is plotted in the right panel of Figure 4.

In Tables 3, 4, and 5, I revisit the results presented earlier for the DG, PGG, and AMP, replacing the dichotomous urban/rural coding with the continuous measures to test whether ethnic bias varies not just with urban or rural location but, within each of the urban and rural subsamples, with the strength of participants’ urban/rural connections (and, within the urban sample, exposure to non-coethnics). Although using a measure that captures the continuous nature of urbanness, ruralness, and exposure is almost certainly an improvement on prior efforts to investigate the impact of urbanization, the updated estimates continue to indicate no impact of urbanization on either altruism, as measured via the DG (Table 3), or cooperation, as measured via the PGG (Table 4).

We do, however, see more interesting results when we turn to the AMP. Unlike in Table

communicate with them at least once a week. These results, which echo the findings in many other studies (e.g., Gugler (2002); Potts (2010); Robinson and Riedl (2025), among others) further underscore the folly of viewing urban and rural Africans as occupying distinct and disconnected worlds.

²⁵See Appendix B. The “exposure” index was not pre-registered.

Table 3: Dictator Game Results, Revisited

	Urban sample				Rural sample	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.354*** (0.053)	0.353*** (0.053)	0.351*** (0.053)	0.351*** (0.053)	0.384*** (0.053)	0.383*** (0.053)
Coethnic	-0.011 (0.019)	-0.010 (0.019)	-0.011 (0.019)	-0.011 (0.019)	-0.008 (0.019)	-0.006 (0.020)
“Urbanness”	0.015 (0.025)	0.042 (0.033)				
Coethnic * “Urbanness”		-0.055 (0.046)				
“Ruralness”					0.021 (0.014)	0.026 (0.020)
Coethnic * “Ruralness”						-0.009 (0.027)
“Exposure”			-0.015 (0.017)	-0.015 (0.023)		
Coethnic * “Exposure”				-0.001 (0.033)		
<i>N</i>	640	640	640	640	670	670
Covariates	yes	yes	yes	yes	yes	yes

Outcome is the share of the endowment given to the other player. Covariates are gender, ethnic group, education, age, and Raven’s score (demeaned). *p<0.1; **p<0.05; ***p<0.01.

Table 4: Public Goods Game Results, Revisited

	Urban sample				Rural sample	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.565*** (0.069)	0.565*** (0.069)	0.568*** (0.070)	0.568*** (0.070)	0.541*** (0.070)	0.540*** (0.070)
Coethnic	0.004 (0.025)	0.003 (0.025)	0.004 (0.025)	0.004 (0.025)	-0.011 (0.026)	-0.009 (0.026)
“Urbanness”	0.038 (0.032)	0.036 (0.044)				
Coethnic * “Urbanness”		0.003 (0.060)				
“Ruralness”					0.002 (0.019)	0.006 (0.026)
Coethnic * “Ruralness”						-0.009 (0.036)
“Exposure”			0.010 (0.022)	0.006 (0.031)		
Coethnic * “Exposure”				0.009 (0.043)		
<i>N</i>	640	640	640	640	670	670
Covariates	yes	yes	yes	yes	yes	yes

Outcome is the share of the endowment contributed to the group fund. Covariates are gender, ethnic group, education, age, and Raven’s score (demeaned). *p<0.1; **p<0.05; ***p<0.01.

2, where I pooled the results for all participants and tested for differences across urban and rural individuals via a dummy variable, in Table 5, I analyze the urban and rural samples separately. This separation reveals that the ethnic effect shown in Table 2 was driven entirely by rural participants, who are statistically significantly much more likely to demonstrate stronger affect toward coethnics than non-coethnics. There is no ethnic effect among urban participants. This finding reinforces the suggestive result highlighted earlier regarding the reduction in affect-related bias against outgroups (or the reduction in positive bias toward ingroups) that seems to occur with urbanization. The fact that increases in “urbaness” among urban participants are associated with further reductions in such outgroup bias (although not significantly so) and that increases in “ruralness” among rural participants are associated with further increases in outgroup bias (although, again, not significantly so), is consistent with this interpretation.

Table 5: Affect Misattribution Procedure Results, Revisited

	Urban sample		Rural sample
	(1)	(2)	(3)
Intercept	0.073 (0.052)	0.064 (0.051)	0.164*** (0.031)
“Urbaness”	-0.070 (0.059)		
“Ruralness”			0.012 (0.024)
“Exposure”		0.103*** (0.037)	
<i>N</i>	128	128	249
Covariates	yes	yes	yes

See notes to Table 2. Covariates are gender, ethnic group, education, age, and Raven’s score (demeaned). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5 also reveals another interesting result: that greater exposure to non-coethnics among urban residents is associated with greater affect-related ethnic bias. As noted in the paper’s introduction, contact with ethnic others has been hypothesized to have quite different impacts. Some researchers, following Allport (1954), have found contact to reduce social distance and prejudice (Scacco and Warren 2018; Mousa 2020; Grady et al. 2023). Others have pointed to its impact on fostering competition and conflict, and hence increasing social distance (Bates 1983). The results reported in column 2 are consistent with the latter hypothesis.²⁶

²⁶A challenge with interpreting the exposure measure is that individuals select into exposure with

From ethnic bias to *expectations* of ethnic bias

The results presented thus far offer no evidence of differential altruism toward ingroup versus outgroup members and reveal no differences in willingness to cooperate along ethnic lines among either urban or rural participants. It would be wrong, however, to interpret the absence of evidence for ethnic bias in the DG and PGG as suggesting that ethnicity is not salient in social or political life. As noted, the salience of ethnicity can stem equally from *expectations* of ethnically biased behavior, or the suspicion that others harbor hidden ethnic prejudices, even if individuals are unwilling to openly express ethnic bias themselves or if others do not in fact harbor ethnic prejudices. Assessing whether interethnic relations vary across urban and rural populations thus requires moving beyond measures of differential altruism and cooperation, or even affect, as measured by the AMP, to examine measures that capture expectations of ethnic bias in others. To do this, I use a third economics game, the Choose-Your-Dictator (CYD) Game, introduced in [Berge et al. \(2020\)](#), which provides a measure of expectations of differential altruism in coethnic and non-coethnic interactions.

In the CYD, participants are presented with information about two different people and asked to choose one of them to be “their” dictator in a DG or to let the computer decide. They are told that the dictator they choose will be given KSh 100 and asked to divide that sum between themselves and the participant.²⁷ One of the two potential dictators is presented as a coethnic and the other is presented as a non-coethnic. As in the DG and PGG, participants are shown photos of the two potential dictators and provided with basic information about their age, education, and hometown, with the hometown serving as the cue about ethnic group membership (see Figure 5).²⁸

The CYD is played twice. In the first, “anonymous,” version, participants select a dictator after being told that the potential dictators do not have any information about them. This version of the CYD serves as a measure of expected differential altruism across the two dictators. In the present study, it tests whether participants expect Kikuyu dictators to be more or less generous than Luo dictators. In the second, “profiled,” version,

others. Those who dislike ethnic outsiders can choose to minimize their contact with them. However, this process would lead to an expected negative relationship between exposure and out-group bias, which is the opposite of what we find here.

²⁷Because the dictator whose profile is selected is not present in the lab session, the amount the participant receives is determined by the amount the dictator gave to an anonymous receiver in a DG played at the time the dictator’s profile photo was collected, at the beginning of the study.

²⁸As in the DG and PGG, the photos are of men pre-screened to be ethnically ambiguous.

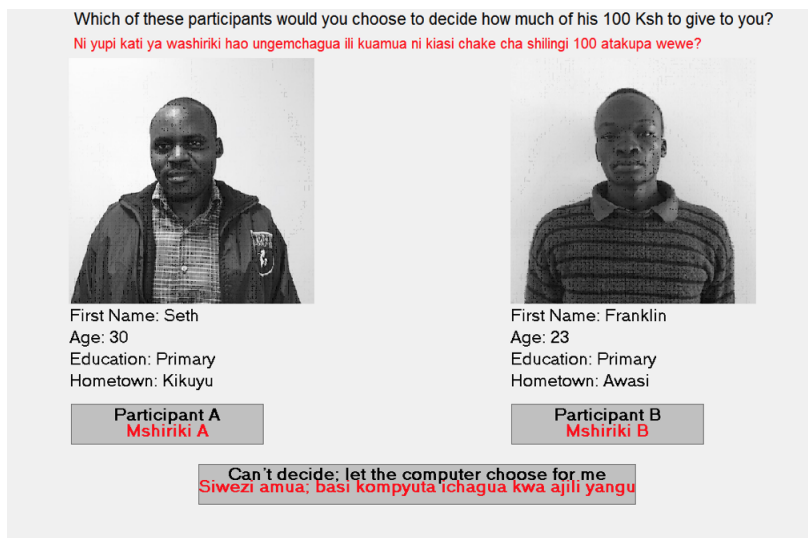


Figure 5: The Choose-Your-Dictator game

participants select a dictator after being told that the dictators will be provided with the participant’s own basic information profile (i.e., their age, years of education, and hometown). In this profiled version of the game, the participant’s choice depends not just on the expected altruism of Kikuyu and Luo dictators *per se* but on the expected differential altruism of each type of dictator toward someone of their own ethnic type. This version of the CYD thus provides insight into expectations of coethnic bias by others.²⁹ If we seek insight into whether expectations of ethnic bias vary across urban and rural participants, this is the measurement tool we should use.

The CYD results are presented in Table 6, where the intercept term reflects the share of the total rounds in which participants select a coethnic to be their dictator rather than a non-coethnic or let the computer decide. Across all specifications study participants are significantly more likely to choose a fellow group member to be their dictator. In the anonymous version, the interpretation is that participants expect that coethnic dictators will be more generous to receivers in general. In the profiled version, the interpretation is that Kikuyus expect fellow Kikuyus to be more generous to them than Luos, and Luos expect fellow Luos to be more generous to them than Kikuyus. The DG results presented earlier reveal that members of neither group are in fact more generous to receivers of their own group. But the expectation, as reflected in the choices in the profiled CYD, is that they will be. Notwithstanding this strong and consistent result, the insignificant coefficients on the *urban* dummy in columns 2, 3, 5 and 6 suggest that there is no difference in the

²⁹Berge et al. (2020) provide a formalization.

strength of these expectations across urban and rural participants. While expectations of ethnic bias are present, they do not appear to weaken or intensify with exposure to urban life.

Table 6: Choose-Your-Dictator Game Results

	Anonymous version			Profiled version		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.440*** (0.019)	0.460*** (0.027)	0.604*** (0.103)	0.478*** (0.020)	0.493*** (0.027)	0.751*** (0.104)
Urban		-0.041 (0.039)	-0.036 (0.041)		-0.030 (0.039)	-0.047 (0.041)
<i>N</i>	655	655	655	655	655	655
Covariates			yes			yes

Outcome is a dummy variable taking the value of 1 if the participant choose the coethnic to be their dictator rather than the non-coethnic or having the computer decide. Covariates are gender, ethnic group, education, age, and Raven’s score (demeaned). *p<0.1; **p<0.05; ***p<0.01.

Conclusion

This study leverages data collected in urban and rural populations in Kenya to learn about the impact of urbanization on interethnic relations, complementing prior work that focused on urban populations alone or that sought to characterize interethnic relations via a narrower set of survey-based outcome measures. While I find evidence that both urban and rural participants expect outgroup members to harbor ethnic biases, I find no evidence for differential altruism or cooperation across coethnics and non-coethnics, either among urban or rural participants—including when coding those participants using a measure that captures the continuous nature of urban-rural connections. I do, however, find evidence of cross-group bias along a dimension of interethnic relations rooted in like/dislike, captured via an AMP. Moreover, I find that this bias differs across urban and rural domains. Specifically, I find that the differential preference for coethnics over non-coethnics is stronger among rural participants, suggesting a potential weakening of cross-group bias due to urbanization.

Given the challenges associated with estimating the impact of urbanization due to selection into migration, the findings reported here can only be suggestive of the effects of urban living on interethnic relations. Even so, the use of multiple measures of cross-group bias and expectations of bias among members of the same politically relevant ethnic

groups living in rural and urban domains provides valuable insight into how moving from a rural to an urban area can alter how members of different communities interact with one another, with important implications for understanding development.

References

- Africa Report. 2018. "The ethnic rivalry that holds Kenya hostage.". 30 January.
- Allport, Gordon. 1954. *The nature of prejudice*. New York: Basic Books.
- Bannon, Alicia, Edward Miguel and Daniel N. Posner. 2004. "Sources of ethnic identification in Africa." Afrobarometer Working Paper No. 44.
- Barriga, Alicia, Neil T.N. Ferguson, Nathan Fiala and Martin Alois Leroch. 2023. "Ethnic cooperation and conflict in Kenya." *Journal of Behavioral and Experimental Economics* 106:1–12.
- Bates, Robert H. 1983. "Modernization, ethnic competition, and the rationality of politics in contemporary Africa". In *State Versus Ethnic Claims: African Policy Dilemmas*, ed. D. Rothchild and V.A. Olorunsola. Westview Press.
- Beegle, Kathleen, Joachim De Weerdt and Stefan Dercon. 2011. "Migration and economic mobility in Tanzania: Evidence from a tracking survey." *Review of Economics and Statistics* 93(3):1010–1033.
- Berge, Lars Ivar Oppedal, Kjetil Bjorvatn, Simon Galle, Edward Miguel, Daniel N. Posner, Bertil Tungodden and Kelly Zhang. 2020. "Ethnically biased? Experimental evidence from Kenya." *Journal of the European Economic Association* 18(1):134–164.
- Blum, Ashley, Chad J. Hazlett and Daniel N. Posner. 2021. "Measuring ethnic bias: Can misattribution-based tools from social psychology reveal group biases that economics games cannot?" *Political Analysis* 29(3):385–404.
- Brooks, David. 2014. "The real Africa." *New York Times*, 8 May.
- Burbidge, Dominic and Nic Cheeseman. 2017. "Trust, ethnicity and integrity in East Africa: Experimental evidence from Kenya and Tanzania." *Journal of Race, Ethnicity and Politics* 2:88–123.
- Camerer, Colin F. and Ernst Fehr. 2004. "Measuring social norms using experimental games: A guide for social scientists". In *Foundations of Human Sociality*, ed. Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr and Herbert Gintis. Oxford University Press.

- Chakraborty, Anujit, Arkadev Ghosh, Matt Lowe and Gareth Nellis. 2024. “Learning about outgroups: The impact of broad versus deep interactions.” Unpublished paper.
- Cilliers, Jacobus, Oeindrila Dube and Bilal Siddiqi. 2015. “The white-man effect: How foreigner presence affects behavior in experiments.” *Journal of Economic Behavior & Organization* 118:397–414.
- de Quidt, Jonathan, Johannes Haushofer and Christopher Roth. 2018. “Measuring and bounding experimenter demand.” *American Economic Review* 108(11):3266–3302.
- Eifert, Benn, Edward Miguel and Daniel N. Posner. 2010. “Political competition and ethnic identification in Africa.” *American Journal of Political Science* 54(2):495–510.
- Emina, Jacques, Donatien Beguy, Eliya M. Zulu, Alex C. Ezeh, Kanyiva Muinda, Patricia Elung’ata, John K. Otsola and Yazoumé Yé. 2011. “Monitoring of Health and Demographic Outcomes in Poor Urban Settlements: Evidence from the Nairobi Urban Health and Demographic Surveillance System.” *Journal of Urban Health* 88:200–218.
- Gluckman, Max. 1960. “Triaxialism in modern British Central Africa.” *Cahiers d’Etudes Africaines* 1(1):55–70.
- Grady, Christopher, Rebecca Wolfe, Danjuma Dawop and Lisa Inks. 2023. “How contact can promote societal change amid conflict: An intergroup contact field experiment in Nigeria.” *Proceedings of the National Academy of Science* 120(43):1–8.
- Green, Elliott. 2013. “Explaining African ethnic diversity.” *International Political Science Review* 34(3):235–253.
- Green, Elliott D. 2023. *Assimilation and industrialization: Explaining ethnic change in the modern world*. Cambridge University Press.
- Gugler, Josef. 2002. “The son of the hawk does not remain abroad: The urban-rural connection in Africa.” *Africa Studies Review* 45(1):21–41.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the dilemmas of collective action*. Russell Sage Foundation.

- Hamilton, Jennifer A., Chad Hazlett and Daniel N. Posner. 2025. “Can a distraction task reveal hidden bias in economics games?” Unpublished paper.
- Haushofer, Johannes, Sara Lowes, Musau Abednego, David Ndeti, Nathan Nunn, Moritz Poll and Nancy Qian. 2023. “Stress, ethnicity, and prosocial behavior.” *Journal of Political Economy, Microeconomics* 1(2):225–269.
- Hazlett, Chad J. and Adam J. Berinsky. 2018. “Stress-testing the affect misattribution procedure: Heterogeneous control of affect misattribution procedure effects under incentives.” *British Journal of Social Psychology* 57(1):61–74.
- Hicks, Joan Hamory, Marieke Kleemans, Nicholas Y. Li and Edward Miguel. 2020. “Reevaluating agricultural productivity gaps with longitudinal microdata.” Unpublished paper.
- Hornsby, Charles. 2013. *Kenya: A history since independence*. I.B. Tauris.
- Jeon, Sangick, Tim Johnson and Amanda Lea Robinson. 2017. “Nationalism and social sanctioning across ethnic lines: Experimental evidence from the Kenya-Tanzania border.” *Journal of Experimental Political Science* 4(1):1–20.
- Kahneman, Daniel, Jack L. Knetsch and Richard H. Thaler. 1986. “Fairness as a constraint on profit seeking: Entitlements in the market.” *American Economic Review* 76(4):728–741.
- Kim, Hue-Sung and Jeremy Horowitz. 2022. “Rejecting ethnic pandering in urban Africa: A survey experiment on voter preferences in Nairobi, Kenya.” *Political Research Quarterly* 75(4):1240–1254.
- Koter, Dominika. 2013. “Urban and rural voting patterns in Senegal: The spatial aspects of incumbency, c. 1978-2012.” *Journal of Modern African Studies* 51(4):653–679.
- Kramon, Eric, Joan Hamory, Sarah Baird and Edward Miguel. 2022. “Deepening or diminishing ethnic divides? The impact of urban migration in Kenya.” *American Journal of Political Science* 66(2):365–384.
- Ledyard, John. 1995. “Public goods: A survey of experimental research”. In *Handbook of Experimental Economics*, ed. Alvin Roth and John Kagel. Princeton University Press.

- Lerner, Daniel. 1958. *The passing of traditional society: Modernizing the Middle East*. Free Press.
- Lyon, Nicholas. 2023. "Unpacking urban: City size, interethnic trust, and prejudice in Africa." Unpublished paper.
- McCauley, John F. 2014. "Pentecostalism as an informal political institution: Experimental evidence from Ghana." *Politics and Religion* 7(4):761–787.
- Melson, Robert and Howard Wolpe. 1970. "Modernization and the politics of communalism: A theoretical perspective." *American Political Science Review* 64(4):1112–1130.
- Mousa, Salma. 2020. "Building social cohesion between Christians and Muslims through soccer in post-ISIS Iraq." *Science* 6505:866–870.
- Nathan, Noah. 2016. "Local ethnic geography, expectations of favoritism, and voting in urban Ghana." *Comparative Political Studies* 49(14):1896–1929.
- Onyango, Elizabeth Opiyo, Jonathan Cruch and Samuel Owour. 2021. "Migration, Rural-Urban Connectivity, and Food Remittances in Kenya." *Environments* 8(92):1–16.
- Payne, B. Keith, Clara Michelle Cheng, Olesya Govorun and Brandon D. Stewart. 2005. "An inkblot for attitudes: Affect misattribution as implicit measurement." *Journal of Personality and Social Psychology* 89(3):277.
- Peterson, Brenton D. 2016. "Social identification and in-group favoritism: Measuring the strength of group ties." Unpublished paper.
- Posner, Daniel N. 2005. *Institutions and ethnic politics in Africa*. New York: Cambridge University Press.
- Potts, Deborah. 2010. *Circular migration in Zimbabwe and contemporary Sub-Saharan Africa*. Oxford: James Currey.
- 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*, ed. John Raven and Jean Raven. Royal Fireworks Press.

- Robinson, Amanda Lea. 2014. "Nationalism versus ethnic identification in Africa: Modernization, colonial legacies, and the origins of territorial nationalism." *World Politics* 66(4):709–746.
- Robinson, Amanda Lea and Rachel Beatty Riedl. 2025. "Beyond the urban-rural divide: Urban duration, rural connections, and sociopolitical (re)orientation in Kenya." *Journal of Modern African Studies* .
- Russell, James A. 2003. "Core affect and the psychological construction of emotion." *Psychological Review* 110:145–172.
- Scacco, Alexandra and Shana S. Warren. 2018. "Can social contact reduce prejudice and discrimination? Evidence from a field experiment in Nigeria." *American Political Science Review* 112(3):654–677.
- Spater, Jeremy. 2022. "Exposure and preferences: Evidence from Indian slums." *American Journal of Political Science* 66(2):302–317.
- Trujillo, Kristin Lunz. 2024. "Feeling out of place: Who are the non-rural rural identifiers, and are they unique politically?" *Political Behavior* 46:2215–2239.
- Young, Crawford. 1976. *The politics of cultural pluralism*. University of Wisconsin Press.

Appendix

A Additional tables and figures

Table A1: Descriptive Statistics

	Urban sample (Nairobi) (%)	Rural sample (Siaya/Kirinyaga) (%)
Gender (Male)	45.9	46.4
Education		
Some/completed primary	36.3	60.9
Some/completed secondary	47.9	33.4
Post-secondary	14.4	5.7
Speaks English fluently	67.6	44.3
Speaks Swahili fluently	31.3	37.8
Age		
<= 25 years old	45	23
<= 35 years old	72	56
Monthly income (KSh)		
Median	5,000	5,750
<=1,000	20.3	5.4
<=5,000	54.1	49.5
<=10,000	74.4	79.5
<=15,000	86.5	90.1
Religion		
Catholic	33.2	30.8
Anglican	16.3	37.3
Pentecostal	22.5	7.3
Media		
Gets news from TV daily	51.6	26.4
Never gets news from TV	13.8	43.3
Gets news from newspapers daily	9.9	2.3
Never gets news from newspapers	43.4	78.8
Gets news from radio daily	65.4	82.9
Never gets news from radio	9.3	2.3
Gets news from internet daily	37.2	8.3
Never gets news from internet	44.2	82.1
Kikuyu-Luo relations		
Intermarriage, as a % of total marriages	19.2	1.8
No personal contacts with other group in past month	10.4	75.1
<= 5 personal contacts with other group in past month	28.7	94
Trust members of other group		
Not at all	12.7	29
A very great deal	19.2	4.9
Find members of other group		
Not at all peace-loving	21.8	39.4
Very peace-loving	35.6	15.5
People in neighborhood are from many different groups	98.9	n/a

Table A2: Dictator Game Results

	Mean offer	Median offer	Rate of positive offers	Average positive offer
Urban ($N = 320$)				
DG, coethnic	32.6	30	0.82	39.8
DG, non-coethnic	33.7	30	0.85	39.5
Difference ^a	-1.14	0	-0.03	0.26
Rural ($N = 335$)				
DG, coethnic	31.1	30	0.78	40.0
DG, non-coethnic	31.9	30	0.76	42.1
Difference ^a	-0.76	0	0.02	-2.10*
Urban-rural difference ^b	-0.39	0	-0.06*	2.36

^a coethnic – non-coethnic

^b $(\text{coethnic} - \text{non-coethnic})_{\text{urban}} - (\text{coethnic} - \text{non-coethnic})_{\text{rural}}$

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A3: Public Goods Game Results

	Mean contribution	Median contribution	Rate of positive contributions	Average positive contribution
Urban ($N = 320$)				
PGG, coethnic group	45.2	40	0.85	53.3
PGG, non-coethnic group	44.8	45	0.85	52.5
Difference ^a	0.35	-5	-0.01	0.80
Rural ($N = 335$)				
PGG, coethnic group	47.9	50	0.83	57.9
PGG, non-coethnic group	48.9	50	0.82	59.6
Difference ^a	-1.07	0	0.01	-1.72
Urban-rural difference ^b	1.42	-5	-0.01	2.53

^a coethnic group – non-coethnic group

^b $(\text{coethnic group} - \text{non-coethnic group})_{\text{urban}} - (\text{coethnic group} - \text{non-coethnic group})_{\text{rural}}$

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A4: AMP Results

	Mean ethnic effect	SE	z-score	p-value
Urban ($N = 128$)	0.099	0.023	4.26	0.000
Rural ($N = 249$)	0.132	0.017	7.84	0.000
Urban-rural difference	-0.033			

Outcome is the “ethnic effect,” the estimated probability of finding a Chinese character pleasant after exposure to a coethnic prime minus the estimated probability of finding the Chinese character pleasant after exposure to a non-coethnic prime. As pre-registered, results of the AMP-Town and AMP-Region are combined. Sample is limited to participants who correctly answered half or more of the questions in the six question literacy test administered at the start of the lab session and whose reaction times were ≤ 200 ms, as pre-registered.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A5: Choose-Your-Dictator Game Results

	% choosing coethnic	% choosing non-coethnic	% letting computer decide
Urban ($N = 320$)			
CYD, non-profiled	0.42	0.38	0.20
CYD, profiled	0.46	0.35	0.19
Rural ($N = 335$)			
CYD, non-profiled	0.46	0.36	0.18
CYD, profiled	0.49	0.34	0.17
Urban-rural difference			
CYD, non-profiled	-0.04	0.02	0.02
CYD, profiled	-0.03	0.01	0.02

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

B Urban/rural connections

To investigate heterogeneity in within the urban and rural samples, I created three averaged z-score indexes:

- an index of *ruralness*, capturing the extent to which rural residents have experience living in urban areas and/or have connections with urban kin
- an index of *urbanness*, capturing the extent to which urban residents have maintained connections to their rural homes
- an index of *exposure*, capturing the extent to which urban residents have contact with members of the ethnic outgroup

Tables B1, B2, and B3 provide the descriptive statistics for the components of each index. Components are reverse-coded where appropriate to ensure that all measures point in the direction of weaker urban connection in the “ruralness” index and weaker rural connection in the “urbanness” index.

Table B1: Components of the “Ruralness” Index

	%
Has family members living in Nairobi/other urban area	84
Were visited by these urban family members during past 2 yrs	
Never	13
Once	23
2-3 times	35
>3 times	29
Visited their urban family members during past 2 yrs	
Never	49
Once	20
2-3 times	19
>3 times	12
Communicate with urban family members	
Every day	5
At least once a week	45
At least once a month	78
Urban family members send them money	56
Have ever lived in Nairobi/other urban area	53
(Among those who did) for ≤ 1 yr	33
(Among those who did) for ≤ 5 yrs	79

Table B2: Components of the “Urbanness” Index

	%
<i>Connection to Nairobi</i>	
Length of residence in Nairobi	
≤ 10 yrs	35
≤ 20 yrs	67
Have a sense of belonging in Nairobi	96
Plan to vote in upcoming election in Nairobi	92
<i>Connection to a rural place outside of Nairobi</i>	
There is a specific place outside Nairobi they consider “home”	76
Length of residence in that place	
≤ 5 yrs	32
≤ 10 yrs	42
≤ 20 yrs	82
Communicate with someone in that place	
Every day	22
> once/wk	70
Visited that place in past 2 yrs	
Never	6
Once	19
2-3 times	33
>3 times	42
Someone from that place visited them in Nairobi in past 2 yrs	
Never	18
Once	31
2-3 times	28
> 3 times	23
Send money to family members/others living in that place	
Never	29
Occasionally	57
Frequently	14
Know what is happening that place	
Almost everything	18
The most important things	41
Very little	35
Has property/investments in that place	38
Plans to live in that place permanently in future	66
After this life on earth, would like to be laid to rest	
In Nairobi	30
In their home village	53
In their partner’s village	17
Being in good standing as a member of one’s home village is	
Essential	8
Very important	25
Important	46
Not important	21

Table B3: Components of the “Exposure” Index (Urban Participants Only)

	%
Married to a spouse from another ethnic group*	18.9
Parents are intermarried**	6.3
Personal contact with member of other ethnic group during past month	
Never	9
1-5 times	17
6-10 times	23
11-20 times	9
> 20 times	42
Ethnic diversity of neighborhood	
One group	1
Many different groups	99

*Among participants who are married.

**Due to a data collection error, we do not have information about the father’s ethnicity. This measure is based on whether the participant’s stated ethnicity is different from their mother.

C Deviations from the preanalysis plan

This appendix reports deviations from the pre-analysis plan (<https://osf.io/mtqr3/overview>).

- I'm not sure what I said about this in the PAP, but in the lab sessions participants played anonymous and profiled version of both the DG and PGG. However, I don't analyze the anonymous versions in the paper. So I'm just not mentioning that. Should I?
- The PAP has a lengthy discussion of the conditional impact of priming to the possibility of election violence, which I don't discuss in the paper.
- I don't analyze results of the PGG with mixed groups.
- I added the "exposure" index to analyze alongside the pre-registered "urbanness" measure.