Can a Distraction Task Reveal Hidden Ethnic Bias in Economics Games?

Jennifer Hamilton[1], Chad Hazlett[2], and Daniel N. Posner[3]

[1]Assistant Professor, Sam Nunn School of International Affairs, Georgia Institute of Technology, 781 Marietta Street NW, Atlanta, GA, 30332, USA, jhamilton99@gatech.edu

[2]Professor, Departments of Political Science and Statistics and Data Science, University of California, Los Angeles, 4289 Bunche Hall, Los Angeles, CA 90095-1472, USA, chazlett@ucla.edu

[3]James S. Coleman Professor of International Development, Department of Political Science, University of California, Los Angeles, 4289 Bunche Hall, Los Angeles, CA90095-1472, USA, dposner@polisci.ucla.edu

Word Count: 6,447

Abstract

Social scientists widely use behavioral economics games such as the Dictator Game (DG) to measure intergroup biases. Yet, in contexts where ethnicity is thought to be extremely salient, the DG often fails to detect evidence of ethnic bias. One explanation is that the DG permits participants to monitor and alter their behavior where social norms discourage discrimination. We test this explanation by experimentally reducing self-monitoring through a cognitive load manipulation. If self-monitoring masks ethnic bias in the DG, we anticipate more revealed bias under higher cognitive load. We find no evidence that heightened cognitive load increased measured ethnic bias on the DG, suggesting participants are not effortfully suppressing bias that the DG would otherwise reveal. Instead, we argue that norm-consistent behavior may be so routinized that it requires little executive control and/or that intergroup bias in this context arises from dislike, fear, threat perception, or other sources that are not well measured by the DG.

1 Introduction

Discrimination and bias across group lines is a pressing concern in many societies, as both a source and indicator of social conflict. However, if discrimination and bias are to be addressed scientifically, they must first be measured. The most widely used behavioral measure of intergroup bias, the Dictator Game (DG), has a well-documented problem: it routinely fails to detect ethnic bias in settings where such bias is widely believed to be pervasive. We evaluate a novel experimental strategy designed to determine whether this failure stems from participants' ability to monitor and control their own behavior.

The DG involves choosing how much of a small quantity of money should be given to another player, where the features of that player can be changed (we provide a detailed description of the paradigm in Section 2). In the "ethnic DG" used here, each player's ethnic group is signaled directly or indirectly, so that the amount given to co-ethnics vs. non-coethnics serves as a behavioral measure of differential altruism towards ingroup and outgroup members. For this reason, the DG has been widely used where investigators wish to quantify ethnic and other between-group biases.[1]

The empirical question that motivates our work begins with the observation that a large and growing collection of studies suggest that economics games like the DG fail to pick up ethnic biases in settings where cross-group biases are thought to be strongly present. For example, Berge et al. (2020), Blum, Hazlett and Posner (2021), and Barriga et al. (2023) all find no evidence of ethnic bias in either the DG or the Public Goods Game (PGG) when played by large samples of residents of Nairobi, Kenya—a country where ethnicity has historically been highly salient in political and social interactions (Fafchamps, 2000; Miguel and Gugerty, 2005; Hornsby, 2013; Hjort, 2014; Choi, Harris and Shen-Bayh, 2022; Harris, 2024) and where one would have expected bias to have been detected.[2] Habyarimana et al. (2009) also find no evidence of ethnic bias in either the anonymous DG or PGG played in Kampala, Uganda, where ethnic differences are similarly thought to play an important role in everyday affairs (Kasfir, 1976; Kibanja, Kajumba and Johnson, 2012). Both the pilot data for the present project and the results presented below replicate these null findings.[3] Participants who displayed no evidence of ethnic bias in the DG have nonetheless displayed the expected bias when completing implicit tasks, reinforcing the concern that the DG is blind to at least some forms of

[1] A partial list of studies using the DG to measure cross-group bias includes Fershtman and Gneezy (2001); Whitt and Wilson (2007); Habyarimana et al. (2009); Chen and Li (2009); Adida, Laitin and Valfort (2014); McCauley (2014); Mironova and Whitt (2014); Iyengar and Westwood (2015); Peterson (2016); Tanaka and Camerer (2016); Jeon, Johnson and Robinson (2017); Berge et al. (2020); Blum, Hazlett and Posner (2021); Barriga et al. (2023); Haushofer et al. (2023), and Posner (2026).

[2] These null results were found—in both Berge et al. (2020) and Blum, Hazlett and Posner (2021)—even in experimental sessions taking place in the weeks leading up to national elections, when the salience of ethnicity is expected to be heightened (Eifert, Miguel and Posner, 2010).

[3] Haushofer et al. (2023) do, however, find evidence of ethnic bias in the DG in a Nairobi sample quite similar to the ones used in the present study.

otherwise salient cross-group prejudice.[4]

Why is this the case? Setting aside the possibility that there is simply no bias to detect in these settings,[5] two broad explanations stand out. The first, and the focus of this paper, is *controllability*. In many contexts, including Kenya, social norms discourage the open display of prejudice toward others. Individuals also possess a well-documented desire to preserve a conception of themselves as fair, egalitarian, and unbiased (e.g., Steele, 1988; Sherman and Cohen, 2006). Despite the real-money stakes used to incentivize behavior, the DG offers participants the opportunity to monitor and control their allocations to ensure that their decisions in the game are aligned with these social norms and internal expectations. This opportunity arises because the DG asks for a slow, deliberate, and transparent choice: participants see the receiver's profile, understand what an unequal allocation would signal, and have as long as they like to decide how much of their endowment to give to the receiver. As a result, the DG is potentially highly sensitive to social desirability bias and self-censorship (Haley and Fessler, 2005; List, 2007; Levitt and List, 2007; Bardsley, 2008; Zizzo, 2010; Cilliers, Dube and Siddiqi, 2015; de Quidt, Haushofer and Roth, 2018; Blum, Hazlett and Posner, 2021), even when played anonymously (Winking, 2014).

Dual-process theories offer a well-founded basis for this account. These theories distinguish behavior guided by automatic reactions, sometimes referred to as "System 1," from more deliberative and reasoned control, variably referred to as executive control, effortful control, or "System 2" (e.g., Shiffrin and Schneider, 1977; Evans and Stanovich, 2013; Kahneman, 2011). Applied to intergroup bias, the apparent group identity and other characteristics of a person trigger a cascade of semantic and affective associations that can influence the perceiver's assessments and, if left unchecked, realized behavior (Payne, 2001). With time, attention, and effort, however, individuals are often able to override or conceal these impulses in their expressed behavior (Devine, 1989; Payne et al., 2005). On this account, the DG records behavior only after System 2 has had the chance to conceal any underlying bias that might have been expressed absent this control.

Despite this controllability, such underlying biases are worthy of study, both to improve our basic understanding of ethnic bias and because the willingness and ability to control expressions of differential altruism may fail in other settings where the stakes are well beyond the few shillings at issue in the laboratory. In addition, investigators may recognize that, even if differential altruism is not exactly the bias they wish to measure (as we describe next), it may be correlated enough with other biases, such as fear or disliking, that it offers an imperfect but useful proxy measurement. This measurement

[4]These implicit tasks make it difficult for participants to reflect on or control their responses to conceal their biases (see e.g. Greenwald, McGhee and Schwartz, 1998; Payne et al., 2005). Studies reporting measures of bias using both the DG and implicit approaches (e.g., the Implicit Attitude Test (IAT) or the Affect Misattribution Procedure (AMP)) include Berge et al. (2020), Blum, Hazlett and Posner (2021), and Posner (2026). A broader discussion of the incongruence between implicit and explicit measures of cross-group bias is provided in Krosnick, Stark and Scott (2025).

[5]We consider and reject this possibility in Section 5.1.

strategy would be undermined, however, if the correlation between differential altruism and these other types of bias is concealed thanks to the DG's controllability.

A second important explanation for the failure of the DG to capture cross-group biases that are visible outside the lab is a *mismatch* between what the DG could measure, even absent controllability concerns, and the drivers of these prejudices. The DG is constructed to measure differential altruism: a willingness to forgo money for a coethnic that one would not forgo for a non-coethnic. But the cross-group antagonism that informs expectations of bias in Nairobi or Kampala may instead be rooted in dislike, fear, threat perception, or preferences about which groups hold power. Under this explanation, the empirical evidence of null DG effects could be consistent with the intergroup tension and conflict seen outside the lab because that tension and conflict is rooted in a set of biases that the DG is not designed to capture.

These two accounts carry very different implications. If controllability explains the null results noted above, design features that limit participants' ability to edit their behavior, like those we explore here, may be valuable to researchers seeking to estimate ethnic bias absent this control. However, if mismatch concerns explain null results, investigators interested in measuring the biases underlying intergroup conflict should look to other instruments that are better designed to measure the relevant dimensions of bias. Distinguishing between these explanations is thus important. It is also potentially possible: We posit that if participants are suppressing bias, and if that suppression draws on executive resources, then taxing those resources by introducing cognitive load should allow the bias to surface. If bias is indeed unmasked under heightened cognitive load, it supports the controllability account and points towards a remedy. By contrast, a continued absence of evidence for a DG effect under cognitive load would suggest that controllability is likely not the issue, or not the only issue. Such a result would instead point towards the mismatch account (Section 5.5), or, potentially, to concerns either that the procedure failed to substantially degrade effortful control (Section 5.2) or that norm-consistent behavior is so routinized that it requires little to no effortful control at all (Section 5.4).

We implement this test by having participants play the DG during the maintenance phase of a spatial working memory task (the Spatial-Delayed Recall Task, SDRT) at load levels calibrated for difficulty in pilot studies.[6] If self-monitoring is at work and requires executive resources to operate, hidden biases should be revealed when the DG is played during the SDRT, and more so at higher loads, relative to when it is played alone. Several prior studies examine the effect of cognitive load on behavior in the DG (e.g., Cornelissen, Dewitte and Warlop, 2011; Schulz et al., 2014; Hauge et al., 2016; Tinghög et al., 2016). However, these studies focus on the effects of cognitive load on generosity *per se*. We focus instead on its impact on selective generosity toward ingroup versus outgroup members, with

[6]The pre-registration document is available at REDACTED See Appendix F.1 for information about the correspondence between the manuscript and the pre-registration document.

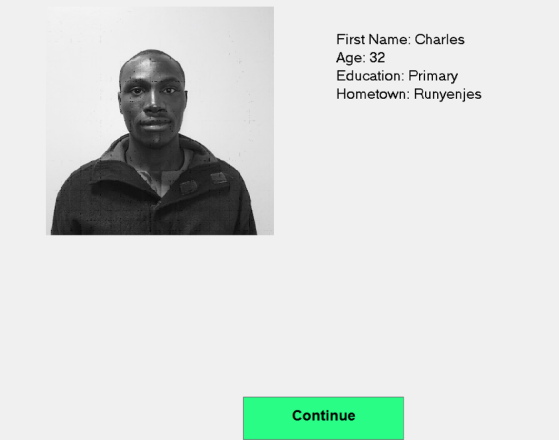
the goal of elucidating the extent to which self-censorship shapes the expression of such differential favoritism.

We find no evidence of differential ethnic bias in the DG when played with versus without the concurrent SDRT, notwithstanding evidence that the SDRT does in fact increase cognitive load. The implication is that participants are not effortfully suppressing bias that the DG would otherwise reveal. We speculate that the reason may lie instead either in the fact that fairness and equity norms are so well-practiced as to not require cognitive effort, or in the blindness of DGs to other dimensions of bias—such as dislike, fear, threat—that are present in the setting we study and that are responsible for the cross group hostilities that inform expectations of what the DG should capture.

2 Our approach

Before describing our strategy for reducing participants' ability to control their responses in the DG, we describe more precisely how the DG is played in our research. Participants ("dictators") are given an endowment of money (approximately 0.5 USD per round, roughly equivalent to the maximal cost of a one-way *matatu* fare within Nairobi) and are asked to decide how much of this money to keep and how much to give away to a partner (the "receiver"). Whatever decision they make is implemented. The receiver is not present in the room with the dictator but is represented by a profile containing their photo, first name, age, education level, and hometown (as in Figure 1).[7] In keeping with deception-free standards, all DG allocations were transferred to the person shown in the photo at the end of the lab sessions.

Figure 1: Profile of receiver used in Dictator Game



[7]Further information on the construction of the profiles can be found in Blum, Hazlett and Posner (2021), whose protocols we adopted. These procedures also follow closely from prior protocols including those of Berge et al. (2020) and Habyarimana et al. (2009).

The key manipulation is provided by the receiver's hometown, which we experimentally vary across hometowns that are recognizable in the context we study as being located in a Kikuyu or Luo part of the country, thus identifying the receiver as almost certainly Kikuyu or Luo.[8] To allow us to test for the impact of having no discernible ethnic cue, one third of the receivers' hometowns are identified as "Nairobi," a place whose multi-ethnic demography provides little information about the receiver's likely ethnic background. The additional information about age and education is included to obscure our interest in ethnic bias, as is standard in experiments of this type.

To reduce the number of factors that might affect participants' responses to the profile, all of the pictures we use are of men. The photos were pre-screened to be ethnically ambiguous, so that the hometown cue would more likely shape the participants' perception of the receiver's ethnic group. Although each participant saw each photograph and associated profile just once, each photo was seen by every participant, attached to different profiles. This guaranteed that, over the course of the study, each face was equally often presented as being from a Luo or Kikuyu hometown or from Nairobi, thus zeroing out the potential effects of phenotypical cues on generosity.[9]

2.1 Depleting System 2: the Spatial Delayed Recall Task (SDRT)

Our strategy for testing whether self-monitoring can explain the consistent null results found in prior work is to have participants play the DG while concurrently engaged in a distraction task that, by inducing higher cognitive load, depletes the executive resources they might otherwise have applied to censoring biased responses they wish to hide. Researchers employ several different strategies for inducing cognitive load (see Deck, Jahedi and Sheremeta, 2021 for a useful review). The approach we adopt, the SDRT (Glahn et al., 2002), has been used for a variety of research and clinical purposes including studies of working memory and ADD/ADHD in children (Lenartowicz et al., 2019; Hazlett et al., 2020) and efforts to document dual process reasoning (De Neys, 2006). We use this approach rather than digit memorization (Gilbert, Tafarodi and Malone, 1993) or an "n-back" task (Kirchner, 1958) because it does not require skilled recall of letters or numbers, which can be challenging for populations with low education. We also prefer the use of visual cues to the common approach of inducing load via time pressure because speeded tasks rely on higher familiarity with computer keyboards and were found to be extremely frustrating for participants, who often disengaged when trials were presented too quickly.

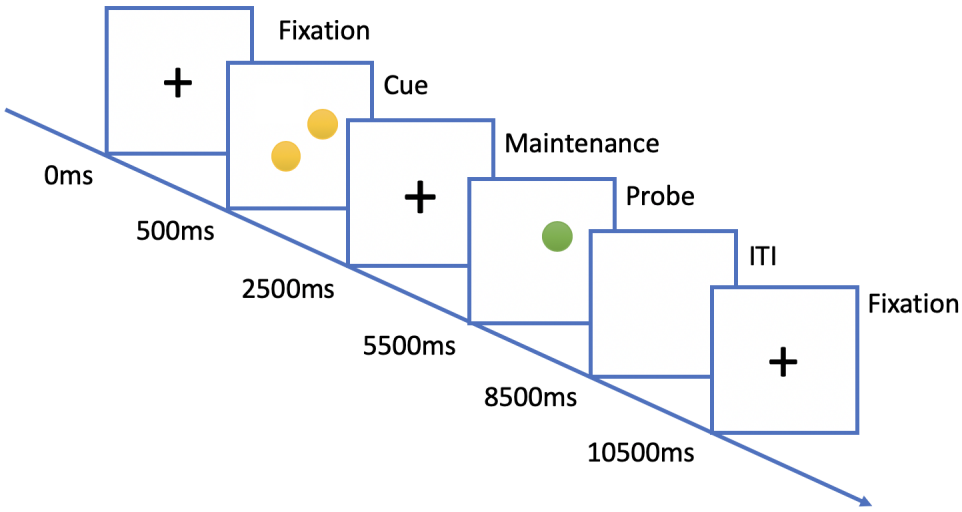
In each trial, participants see either one or two "cue" dots in a variety of positions on the screen, followed by a blank screen (the "maintenance phase"), followed by a "probe" dot at a single location.

[8] Berge et al. (2020), who use the same hometown cues, provide evidence for their effectiveness as indicators of ethnic group affiliation in this context.

[9] On the impact of physical cues on behavior in the DG, see e.g. Rosenblatt (2008).

Participants are instructed to indicate whether the probe dot was in the same location as any cue dot.[10] Successfully completing the task requires that participants maintain in working memory a representation of cue dot location(s). The difficulty of the task (and cognitive load) is greater when there are more cue dots, so including both a one-dot (SDRT-1) and two-dot (SDRT-2) version allows us to test whether different levels of cognitive load have different effects.[11] Figure 2 provides a schematic representation of the SDRT-2.[12]

Figure 2: SDRT schematic



Source: Authors.

2.2 Combining the DG and SDRT

To learn how the distraction task affects allocations in the DG, we ask participants to play the DG while they perform the SDRT-1 or SDRT-2. We then compare the degree of ethnic bias we observe in the standalone DG and when played during the SDRT-1 and SDRT-2. Specifically, in the SDRT-1 and SDRT-2 conditions, participants first see the cue dot(s) of the SDRT, displayed for 2 seconds. Once the cue dot(s) disappear(s) (i.e., in the maintenance phase), a DG profile appears and participants choose how much to give to the receiver, just as they do in the standalone DG. Once they respond, the probe dot appears. The participant is then prompted to respond whether the probe is in a location

[10] As described below, we randomized whether participants were offered a monetary incentive for their accuracy in the dot task.

[11] We chose these difficulty levels after first testing a variety of difficulties on ourselves and lab members, followed by a pilot study that employed SDRT-1 and SDRT-3 combined with the DG as we do here. In the pilot phase, while the response to SDRT-1 was promising, SDRT-3 proved too difficult (responses were correct no more than chance). This prompted us to choose SDRT-1 and SDRT-2 for the full study described here.

[12] The timings, indicating total elapsed time in milliseconds (ms), reflect the ordinary use of this task in isolation. The timing for the maintenance phase is altered when we interrupt the SDRT with the DG, as described below.

previously occupied by a cue dot.[13]

2.3 Estimating the impact of the SDRT on ethnic bias

Participants play the DG at all three levels of cognitive load (no distraction, SDRT-1, and SDRT-2) in randomized order. At each level of distraction, participants play the DG three times: once with a coethnic, once with a non-coethnic, and once with a receiver with a neutral ethnic profile (i.e., from Nairobi), again in a randomized order. With this structure, the degree of bias can be ascertained for each participant at each level of distraction. For example, the ethnic bias under the condition of no induced cognitive load (bias_0) is given by the contribution to the coethnic receiver minus the contribution to the non-coethnic receiver in the standalone DG. Ethnic bias under SDRT-1 (bias_1) and SDRT-2 (bias_2) are defined similarly as the difference between contributions to coethnics and non-coethnics under those levels of distraction. We also pool SDRT-1 and SDRT-2, yielding $\text{bias}_{1,2}$.

Our expectation is that the attentional resources devoted to the SDRT will deplete the resources available for resisting the impulse to discriminate along ethnic lines; as a result, we expect that we will find higher ethnic bias in the DG when performed concurrently with the SDRT-1 and SDRT-2 than in the standalone DG. Our primary outcomes of interest are the contrasts across the allocations under these different treatments, averaged across all participants, all of which we expect to be positive:

$$\text{DistractionEffect}_{1,2\text{vs}0} = \overline{\text{bias}_{1,2} - \text{bias}_0}$$

$$\text{DistractionEffect}_{1\text{vs}0} = \overline{\text{bias}_1 - \text{bias}_0}$$

$$\text{DistractionEffect}_{2\text{vs}0} = \overline{\text{bias}_2 - \text{bias}_0}$$

$$\text{DistractionEffect}_{2\text{vs}1} = \overline{\text{bias}_2 - \text{bias}_1}$$

Nearly every distraction effect we estimate is thus a simple mean over within-participant differences (e.g., for participant i, $\text{DistractionEffect}_{i,1\text{vs}0} = \text{bias}_{i,1} - \text{bias}_{i,0}$). We quantify uncertainty with the standard error of a mean: the standard deviation of the relevant per-participant score divided by $\sqrt{n}$, with t-based confidence intervals.

No clustering or repeated-measures adjustment is needed for these quantities, since reducing each participant to one score removes the only source of within-participant dependence; no regression is used. Two sets of comparisons are instead between participants. First, robustness checks restricted to a single DG round per person (e.g., first-round-only), where a within-person bias score cannot be formed, use two-sample unequal-variance standard errors and are correspondingly less precise. Second,

[13] Note that the duration of the maintenance phase thus depends on the time taken to complete the DG. In designing the task, we faced a tradeoff between allowing participants to endogenously extend the maintenance phase and forcing them to adhere to a fixed, short interval, which would likely have generated a large number of non-responses. We elected to minimize non-responses, at the cost of allowing endogenous variation in the length of the maintenance phase.

the monetary incentive and the permission to discriminate were assigned by lab session rather than by participant (Section 3). We test differences across those arms with randomization tests: holding the participant-level data fixed, we permute the arm labels across the 38 sessions (20,000 draws) and recompute the difference in means. The confidence intervals in Figures 5 and 6 describe variability across participants and do not speak to differences across arms.

Note that exposure to the distraction task need not fully disable effortful control in every participant or in every round. The question is whether distraction by the SDRT is enough to unmask some proportion of bias some of the time, sufficient to measure a difference in the average contributions by level of distraction.

3 Data collection and lab protocols

We recruited 549 Kikuyu and Luo individuals to participate in our study at the Busara Center for Behavioral Economics in Nairobi, Kenya. None had previously participated in an experimental session. We focus on members of the Kikuyu and Luo ethnic groups to sharpen our focus on politically relevant interethnic differences, as they are the two most historically antagonistic and politically competitive groups in Kenya (see e.g., Hornsby 2013). We code participants as Kikuyu when they report speaking Kikuyu and having at least one parent of Kikuyu or Kamba heritage, and we code participants as Luo when they report speaking Dholuo and having at least one parent of Luo or Luhya heritage. In our main analyses, we exclude participants who report having both Kikuyu/Kamba and Luo/Luhya heritage.[14] These restrictions reduce the effective sample from 549 to 516 participants.

Participants began the lab session by familiarizing themselves with the computer keyboard and touchscreen and taking a simple literacy test. They also engaged in training and practice rounds on the DG, SDRT, and combined DG-SDRT. These practice modules were accompanied by comprehension tests to confirm understanding. As preregistered, our main analysis only includes participants who correctly answered at least 5 of the 6 questions on the literacy test, at least 2 of the 3 questions on the DG comprehension test, and at least 3 of the 4 questions on the SDRT practice test. These additional criteria leave us with 77% of the original participant pool (421 of the 549 participants).[15]

After completing all nine of the DG conditions (3 levels of distraction by 3 ethnicity profiles), participants were randomly assigned to play one more round of either the standalone DG or SDRT-2. They then answered a series of multiple choice questions to test their memory of the age, hometown,

[14] This approach diverges from the pre-analysis plan, wherein we specified that main analyses would include participants with mixed Kikuyu-Luo parentage as either Kikuyu or Luo based on language use reported at registration. This change only affects 21 participants, and results are robust to their inclusion.

[15] The literacy screening is important, as completing the task as intended requires the ability to read text stimuli and to follow written instructions. Of the 95 participants excluded due to task competency filters, 61 were dropped due to the literacy filter, with the other two filters excluding an additional 34 participants. We report the results of relaxing these standards below.

education level, and name of the receiver whose profile they had just seen. We included this task to assess whether performing the SDRT might attenuate the processing of receiver profile features, thus potentially reducing correct recall of the ethnic cues and undermining the capacity to discriminate along ethnic lines (see e.g., Lavie (2005)). However, as long as implicit or automatic processing remains largely intact, the reduced recall might not affect a participant's ability to discriminate along ethnic lines. Indeed, the relationship between conscious access to a cue and its influence on behavior is not a simple one. Murphy and Zajonc (1993) found that affective primes shifted evaluations of subsequent targets when presented too briefly to be consciously perceived, but not when presented long enough to be recognized and correctly attributed. Conscious access to a cue is thus not necessary for it to shape behavior, while impaired explicit awareness of that cue could plausibly reduce the activation of social desirability concerns, potentially increasing expressed bias (see e.g., Devine, 1989; Payne, 2001). From that standpoint, the Affect Misattribution Procedure (Payne et al., 2005) is remarkable because it shows the cue's effect on behavior persists even when the cue is clearly visible and participants are instructed to disregard it. Accordingly, diminished recall of profile details under SDRT load may signal that various perceptual and cognitive processes were disrupted, but reduced awareness of the cue need not in itself reduce expressed bias. We nevertheless conducted pre-registered analyses of SDRT effects on cue recall to provide mechanistic evidence of future potential value. We find no substantive or statistically significant differences in recall on most profile characteristics by load condition (see Appendix Figure D.4 and Appendix Table D.2).

The lab session concluded with a post-survey including questions about participants' length of residence in Nairobi, the characteristics of the neighborhood in which they live, their media consumption, their social distance from members of the outgroup, their attitudes toward ethnic intermarriage, and a standard battery of questions measuring motivation to control prejudice (Dunton and Fazio, 1997). We employ this information in the analysis of heterogeneous treatment effects described below.

The post-survey also included a question in which we asked participants what they thought the project was about. Most respondents said they thought the project was about how people make decisions in their everyday lives, testing memory, or learning how much people are willing to share. Although several participants inferred that the project was about how people relate with others in the community, less than one percent inferred that we were interested in discrimination or interethnic relations specifically.

In order to test additional hypotheses about the conditions under which the SDRT should be most distracting and under which demand effects are likely to be strongest, we assigned the 38 lab sessions to one of four treatment arms that varied whether the task instructions for the DG gave participants

explicit "permission to discriminate"[16] and whether participants were provided monetary incentives for their accuracy in the SDRT.[17]

4 Results

4.1 Main findings

Our main, pre-registered results are reported in Figure 3. Contrary to expectations, we find no differences in the sizes of allocations to coethnics and non-coethnics—that is, no difference in ethnic bias—across DGs played with and without the distraction task. These null results are robust to lowering the literacy threshold for inclusion in the sample to just 4 of 6 literacy test questions answered correctly and to removing participants who were determined to be of mixed ethnic background based on the information they provided in the post-survey about their parents' ethnicities (see Appendix B for these and other robustness tests).

These results are based on our pre-registered definition of ethnic bias as the contribution to the coethnic receiver minus the contribution to the non-coethnic receiver. However, ethnic bias may also manifest itself as either positive in-group bias without negative out-group bias, or vice versa. That is, we could consider two alternative measures of bias: amount sent to coethnics minus amount sent to neutral receivers, or amount sent to neutral minus amount sent to non-coethnic. Under these alternatives, we still see no evidence of bias, under any cognitive load condition, with DG allocations always close to 11-12 KSh, about a quarter of the 50 KSh endowment (see Figure 4). As shown in Appendix Figure B.1, the distribution of DG allocations is also very similar across receiver profiles and levels of cognitive load.

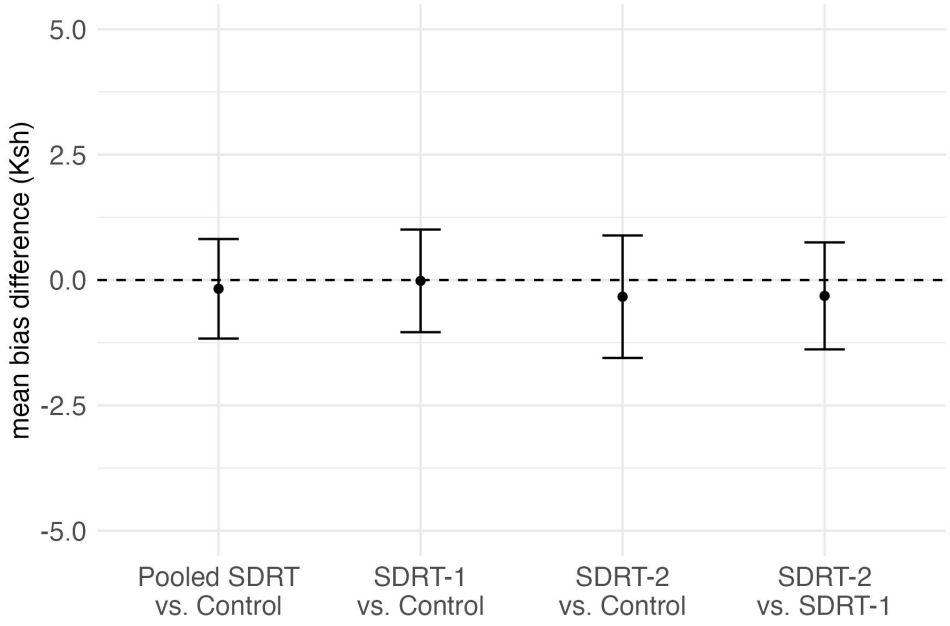
4.1.1 Between-round dependencies

These null findings are based on analyses in which the amounts sent to coethnics and non-coethnics are pooled across all participants in each experimental condition. However, numerous forms of between-round interference can complicate this analysis. One severe form of between-round dependence is anchoring. At the extreme, some participants may make the same allocation in all rounds regardless of the receivers' characteristics. Analyses that pool these non-discriminating participants with more reactive ones will attenuate any average effect, perhaps to the point where it cannot be distinguished

[16] This instruction read: "Participants from around the world have completed this activity. When participants are given information about the other person, this information often leads them to give more or less money. Participants have many reasons to give more or less money to different people, and that is perfectly okay. How you choose to divide the money is entirely up to you, and whatever decision you make is acceptable."

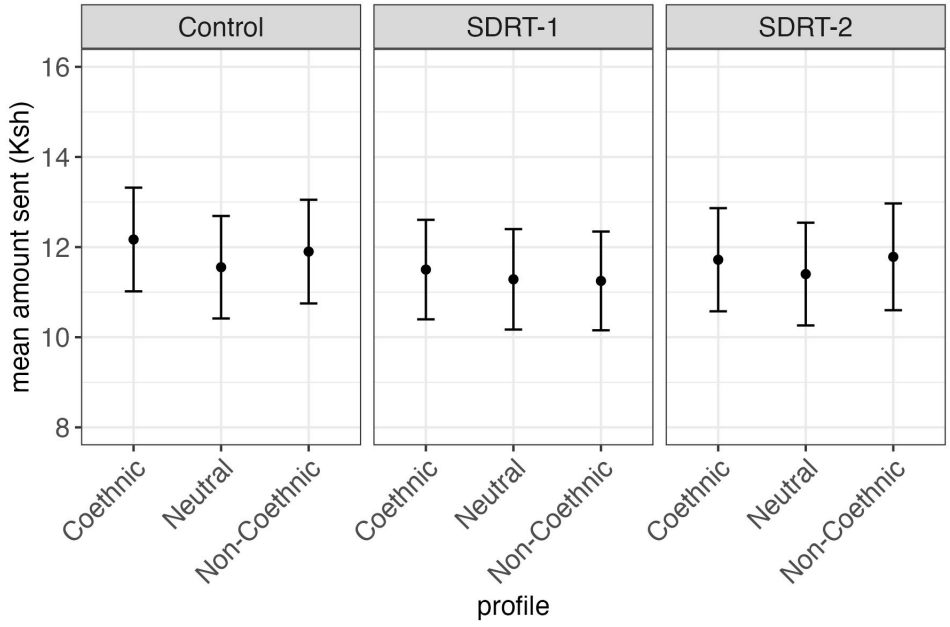
[17] This instruction read: "In addition to any money you keep from the allocation game, you may also earn up to 200 KSh for the memory task depending on how accurate you are." Crossed with the permission to discriminate cue, four session instruction conditions exist. To ensure nearly equal numbers of all four, assignment to each condition type was determined through a consecutive rotation.

Figure 3: Main results: Differences in ethnic bias by condition.



Note: Points are mean within-person bias contrasts, where ethnic bias is a participant's allocation to the coethnic minus the non-coethnic receiver. Each contrast is the bias under the first-named condition minus the second. Bars are 95% confidence intervals. $N = 421$ participants. Numerical values appear in Table A.1.

Figure 4: Main results: Dictator Game allocations by condition and ethnic profile



Note: Points are mean allocations in KSh; bars are 95% confidence intervals. $N = 421$ participants. Numerical values appear in Table A.2.

from zero. We investigate this in two ways. First, we re-estimated our results, dropping the 32% of participants who gave the same amount across all nine DG rounds. As shown in Appendix Figure B.2, our results are unchanged.[18] Second, we examine the data using only the first DG allocation each participant makes, which is counterbalanced to be to either a coethnic, non-coethnic, or neutral (Nairobi) profile. This requires shifting to a between-participants design. While this greatly reduces statistical power, we find no statistically significant difference in bias between the pooled SDRT and control conditions (DistractionEffect$_{12vs0}$=−4.4 KSh, SE=3.4, p=0.20).[19]

A further form of between-round dependence concerns the order in which participants encounter the load conditions. Because each participant plays the DG at every load level, those who happen to play an unloaded round early may use it to deliberate a considered allocation policy. For example, they may decide to give equally regardless of ethnicity. Participants may then carry this simple rule into subsequent rounds even when those rounds are played under load. Distinct from anchoring, which involves repeating the same amount regardless of the receiver, this mechanism involves forming and executing a deliberate, potentially egalitarian rule. Like anchoring, however, it would attenuate the distraction effect we aim to detect. We investigate this by exploiting the randomized order of the load blocks. Among the 295 participants whose first block was loaded, ethnic bias under load is a precisely estimated zero (0.0 KSh, SE 0.37). We also examined whether playing the out-group without distraction first increases the probability of replicating that allocation under load, as this mechanism would imply. However, if anything, we find the opposite: the mean absolute change in the out-group allocation across load conditions was 5.5 KSh for those with no load in the first block, versus 3.8 KSh among those with SDRT-1 or SDRT-2 first. Likewise, out-group allocations under load matched the unloaded allocation 49% of the time among participants with an unloaded first round, versus 65% among those with a loaded first round. This asymmetry is specific to out-group allocations: for co-ethnic receivers, replication did not differ by order (a mean absolute change of 4.8 versus 4.7 KSh). We therefore find no evidence that participants formulate and adhere to a more stable or egalitarian strategy when given the opportunity by an unloaded first round.

The final between-round dependency concern we investigate is the "moral credentialing" effect. Prior work has documented a tendency for participants to become less generous across rounds (e.g., Arechar and Rand, 2022). Generosity in earlier rounds may give participants license to be less altruistic later in the session because they feel they have already established themselves as generous. Such moral credentialing may be particularly pronounced when the behavior in question is not just about altruism

[18] Another form of anchoring involves following a fairness norm of giving half of the endowment to the receiver and keeping half oneself. As shown in Appendix Figure B.1, we see little evidence of this behavior.

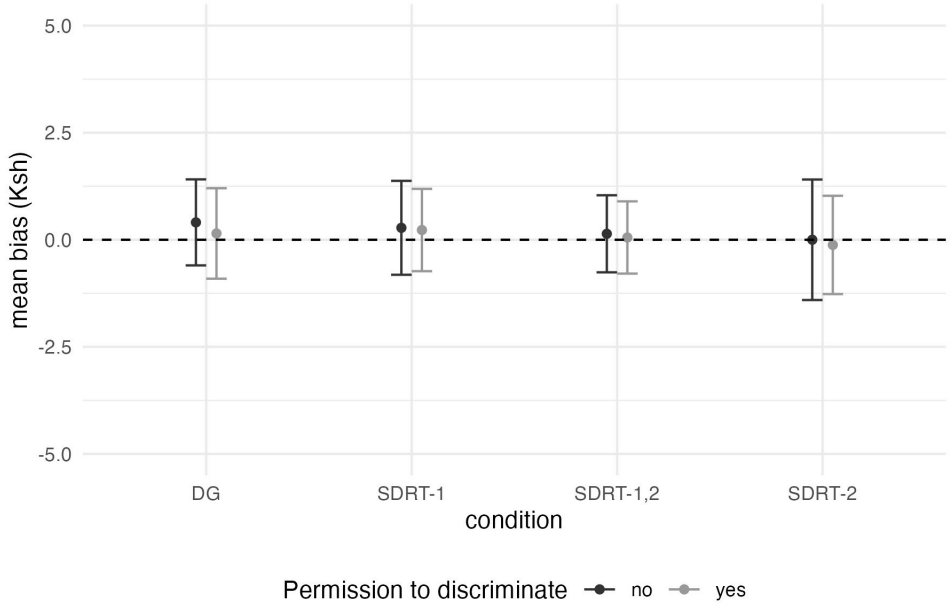
[19] For the other comparisons, we see DistractionEffect$_{1vs0}$=-2.5 KSh, (SE=3.7, p=0.50), and DistractionEffect$_{2vs0}$=-6.6 KSh, (SE=3.8, p=0.09). While the last begins to approach significance, it is in an unexpected direction; this remains insufficient evidence against the null hypothesis of no distraction effect on bias levels in the first DG pairing.

per se but altruism toward someone from an ethnic or racial outgroup—especially in a context where such discrimination is frowned upon (Effron, Cameron and Monin, 2009). If this is the case, we can expect the least bias early in the experiment, with more bias "unmasked" in the final rounds as participants feel they have greater license to express bias. However, when we estimate only the final SDRT condition faced by each participant, we find no evidence of ethnic bias in the DG (see Appendix Figure B.3).

4.2 Heterogeneous effects

Beyond these main effects, we designed the study to test whether the impact of the distraction task varied under different experimental conditions. For example, in keeping with the discussion in de Quidt, Haushofer and Roth (2018), we hypothesized that our findings might differ when we gave participants explicit permission to discriminate. To the extent that social desirability bias stems from the perception that discrimination against outgroup members is unacceptable, such permission might increase ethnic bias in the standalone DG and reduce the bias-revealing effects of the distraction task. As shown in Figure 5, however, we find no differential effects when participants are explicitly told that discrimination is permitted (all $p \geq 0.71$, by the randomization tests described in Section 2.3).

Figure 5: Effect of giving participants permission to discriminate



Note: Points are mean ethnic bias (coethnic minus non-coethnic allocation, KSh) by condition and permission arm. $N = 224$ with and 197 without permission to discriminate. Numerical values appear in Table A.3. Bars are 95% confidence intervals across participants and do not speak to differences across arms (Section 2.3).

We also tested whether our findings differed with the characteristics of the participants themselves.

Specifically, we investigated whether both ethnic bias and the effects of the distraction task on revealing such bias varied across Kikuyu and Luo participants; with length of residence in Nairobi (which might be correlated with a reduction in ethnic bias and/or socialization to the need to hide such bias (Berge et al., 2020; Kim and Horowitz, 2022; Lyon, 2023)); with the diversity of the participant's neighborhood and reported levels of contact with members of other ethnic groups (Enos, 2017); with attitudes toward ethnic intermarriage; and with the participant's motivation to control prejudice, as measured via the Motivation to Control Prejudiced Reactions Scale (MPCRS) (Dunton and Fazio, 1997). As shown in Appendix C, we find no statistically significant differences in bias or distraction effects conditional on these characteristics.[20]

5 Discussion

Contrary to expectations, we find no evidence that the distraction task unmasked ethnic bias. We consider the evidence for and against five possible explanations for this null result: (i) there is little or no cross-group bias to unmask in the setting we are studying; (ii) the SDRT did little to deplete cognitive resources, either because participants exerted little effort on it or because the task was otherwise ineffective at inducing cognitive load; (iii) other features of our experimental protocol suppressed or attenuated the expression of bias; (iv) self-monitoring in this context requires little to no cognitive bandwidth of the type depletable by the SDRT; and (v) behavior in the DG is largely insensitive to the type of bias that exists in this setting.

5.1 No cross-group bias to unmask

While useful to consider, we believe this potential explanation is unlikely. A substantial body of research on ethnic conflict in Kenya documents the historically fraught relationship between Kikuyus and Luos (e.g., Gibson and Long, 2009; Wrong, 2009; Hornsby, 2013). Moreover, manifestations of ethnic bias between these groups are readily observable in real-world settings (see e.g., Choi, Harris and Shen-Bayh, 2022).

Yet, notwithstanding the strong evidence of intergroup antagonism between Kikuyus and Luos in Kenya as a whole, it is possible that relations between members of these groups could be more friendly in our Nairobi study setting, where people have learned to get along and, in many cases, self-selected into living side-by-side with people from other ethnic communities. Although we cannot rule out this explanation entirely, we note that our null effects hold equally among recent migrants to

[20] Appendix C also reports the results of additional pre-registered analyses of heterogeneous treatment effects, none of which reveal meaningful evidence of differences in either ethnic bias or in the effects of varying cognitive load. In the handful of instances where we do detect heterogeneous impacts in some specifications, they are substantively trivial, altering expected allocations by less than 1 KSh, and often in unexpected directions.

Nairobi and longtime residents, and among those living in more and less diverse neighborhoods (see Appendix Figures C.2 and C.3)—both of which weigh against a "Nairobi" effect.[21] The possibility that ethnic bias is weaker in Nairobi is also inconsistent with the results reported in Barriga et al. (2023) and Posner (2026), which find no differences in ethnic bias, as measured via a DG, across subjects recruited in Nairobi and in rural districts.

5.2 The SDRT did not substantially deplete cognitive resources

Next, we consider whether the SDRT may have failed to deplete cognitive resources. This may happen either because participants deployed insufficient effort in order to meet task demands or because the task was insufficiently demanding. We consider evidence for these possibilities below.

First, regarding effort, we draw on data collected at the end of the lab session about how important participants felt it was for them to remember the position of the dots in the SDRT. Although the vast majority of participants indicate that it was "very important" (see Appendix Figure D.1), we do not put too much weight on these findings because participants' responses were almost certainly affected by their desire to say that they had faithfully complied with the task instructions. Turning to behavioral evidence, we consider how often participants responded at all, and how often they were correct when they did. Participants responded to 85% of SDRT-1 probes and 79% of SDRT-2 probes, suggestive of reasonable overall engagement and greater difficulty in the SDRT-2. Among the probes they did answer, 74% (95% CI: [0.71, 0.77]) were correct under SDRT-1, and 59% (95% CI: [0.56, 0.62]) were correct under SDRT-2. These results suggest that respondents are making an effort (responding well above chance), again with SDRT-2 more difficult than SDRT-1.[22]

Several lines of evidence bear on the concern that the task was insufficiently demanding. First, as *prima facie* evidence, participant feedback during our pilot study suggested that participants found the SDRT to be extremely taxing.[23] Second, we asked participants to perform a simple math task during the practice rounds of the standalone DG and the SDRT+DG (i.e., both under cognitive load and not). Performing the math task concurrently with the SDRT reduced accuracy relative to the no-load control, significantly so under SDRT-1 and under the pooled SDRT conditions (see Appendix Table D.1). That a concurrent task degraded performance on an unrelated task is consistent with the

[21] We do find evidence of lower levels of ethnic bias in the SDRT-2 and SDRT-1,2 conditions among participants reporting below median levels of contact with non-coethnics (see Appendix Figure C.4). However this finding runs opposite to the hypothesized relationship, which would have led us to expect greater outgroup bias among participants with less outgroup contact. Also, nearly half of the sample is precisely at the median on this measure, leaving very few participants in the below median category, and undermining our confidence in this finding.

[22] Note that, due to a coding error, all three SDRT-2 probes had "same" as the correct answer. Accordingly the proportion correct under SDRT-2 is influenced by any tendency to simply respond "same" across all three rounds. The SDRT-1 probes included both "same" and "different" correct answers, so do not share this concern.

[23] In the pilot phase, we employed SDRT-1 and SDRT-3, while remaining concerned that many would find SDRT-3 too difficult. The data from the pilot suggests SDRT-3 was indeed too difficult, as accuracy was only at the chance level, and the preliminary unmasking effect we found for the SDRT on DG bias was seen only with the SDRT-1 vs. control.

SDRT imposing meaningful cognitive load.[24]

We next looked to heterogeneity in effort and DG results, evaluating whether participants who performed better on the SDRT, and who may have been investing more effort, reveal more differential outgroup bias. As shown in Appendix Figure D.2, we find no evidence that this was the case.

A similar test leverages the incentive for SDRT performance described earlier to evaluate whether experimentally induced effort is associated with an increased expression of ethnic bias. As shown in the top panel of Figure 6, participants who received a financial incentive for their performance on the SDRT responded to somewhat more of the dot probes than those who did not: 87% versus 82% under SDRT-1 and 81% versus 77% under SDRT-2. Applying the randomization test described in Section 2.3, the SDRT-1 difference is suggestive but not conventionally significant ($p = 0.07$), while the SDRT-2 difference is not ($p = 0.15$). Accuracy among those who did respond was no higher under the incentive: 75% versus 71% under SDRT-1 ($p = 0.28$) and 56% versus 59% under SDRT-2 ($p = 0.42$). The incentive thus appears to have raised engagement with the task modestly at most, and the bottom panel of Figure 6 shows no accompanying increase in expressed bias: ethnic bias does not differ by incentive arm under any load condition, nor does the effect of distraction on it ($p \geq 0.18$ for every such comparison).

5.3 Other features of our experiment mitigated expression of bias

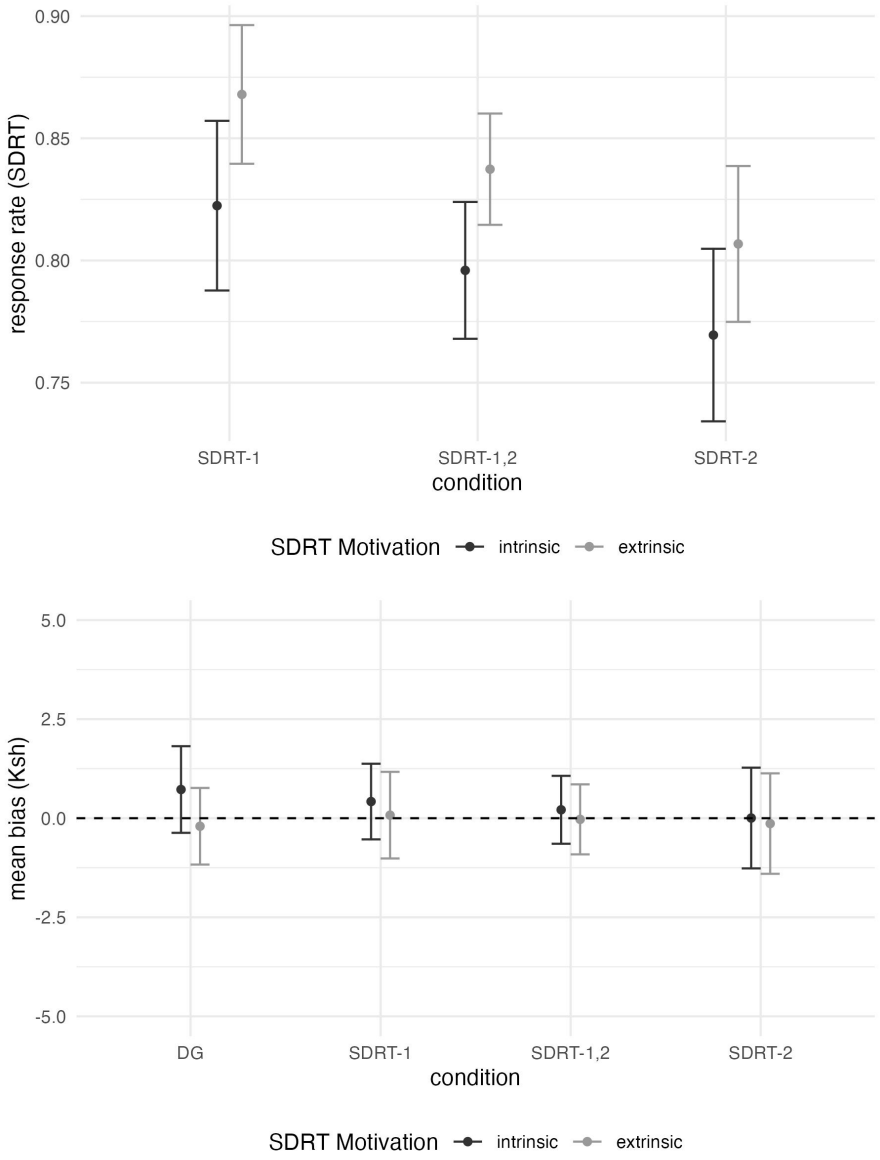
The details of our experimental protocols may have also contributed to the inability of the SDRT to reveal ethnic bias. Following the lead of other studies that use behavioral games to measure outgroup discrimination, we took several steps to hide our interest in ethnic favoritism: we never explicitly mentioned ethnicity in our instructions or study description; we conveyed the receiver's ethnic group membership via a hometown cue rather than by overtly identifying the receiver as Kikuyu or Luo; and we obscured the ethnicity cue within a set of additional pieces of information about the receiver's name, age, and education. These strategies have proved successful in prior studies in hiding researchers' interest in ethnicity and, based on the findings of our post-experimental survey described earlier, appear to have been successful in our study as well.[25] But they may have been ill-advised to the extent that they (perversely) reduced the bias that our distraction task was intended to unmask.

Our decision to embed the ethnicity cue within a set of other profile attributes may have also dulled the impact of the SDRT on expressed ethnic bias by providing alternative criteria on which participants could discriminate. When presented with minimal information about a person or group, respondents often make inferences about other traits based on (often inaccurate) stereotypes, particularly when the

[24] For evidence that increasing cognitive load reduces performance on math problems, see Deck and Jahedi (2015).

[25] For example, Berge et al. (2020), who employ very similar protocols, report that few participants guessed that the researchers were interested in ethnic issues when asked about the study's objectives in post-session debriefings.

Figure 6: Effect of incentivizing SDRT performance on SDRT response rate and resulting bias



Note: Top panel: Proportion of probes with a response by cognitive load condition for extrinsically vs. intrinsically motivated participants; bottom panel: mean ethnic bias (KSh). $N = 207$ extrinsically and 214 intrinsically motivated participants. Bars show 95% confidence intervals across participants for the response rate (top) or mean bias (bottom). As in Figure 5, inferences on the differences in these outcomes by motivation condition rest on the randomization test described in Section 2.3. Numerical values appear in Tables A.4 and A.5.

identities are polarizing. Providing additional characterizing information reduces the need for making such inferences, which can in turn reduce group-based biases (Orr and Huber, 2020; Druckman et al., 2022). Although we find little evidence of bias based on the profile cues we included alongside the receiver's hometown (see Appendix E), we cannot rule out that at least some participants discriminated based on these receiver characteristics and that this behavior weakened the signal we would have

observed had the profiles only conveyed information about ethnic group membership. On the other hand, however, having more elements in the profile might also *increase* outgroup bias by providing alternative, innocuous characteristics that participants may believe will obfuscate their discriminatory choices (Cunningham and de Quidt, 2016; Blum, Hazlett and Posner, 2021). On balance, it is hard to know whether this aspect of our design heightened or dampened our ability to find effects under the distraction task.

5.4 Self-monitoring may not rely on depletable cognitive resources

We turn here and in Section 5.5 to two explanations that we view as most plausible and theoretically consequential. While participants appeared to expend effort on the SDRT and the task may have depleted cognitive resources, it is possible that the processes governing behavior in the DG simply do not draw very much on the executive resources that the SDRT taxes. Decision-making that supports norm-consistent behavior—such as avoiding overt discrimination or treating others fairly—may become highly routinized through repeated social experience. This idea is an old one: James (1890) argued that the function of habit is precisely to remove the need for effortful attention, freeing that capacity for other work. In this view, a well-practiced disposition need not draw on the resources the SDRT aims to deplete. The Social Heuristics hypothesis (Rand et al., 2014) offers a contemporary version of this account, suggesting that individuals develop simple decision rules through their daily experiences, which they continue to apply even when following those rules is not payoff-maximizing (such as in anonymous laboratory interactions where violations cannot be socially sanctioned). Once internalized, maintaining such behavior may require little executive control of the type vulnerable to depletion by the SDRT or other demanding tasks. Experimental work on dual-process reasoning shows that well-practiced heuristic responses often remain intact under conditions of cognitive load (e.g., De Neys, 2006). From this perspective, equal treatment in the DG may reflect an automatized fairness heuristic rather than effortful suppression of bias. Cognitive load would therefore be unlikely to erode such behavior, consistent with our failure to observe increased bias under load.

5.5 Mismatch between the DG and relevant dimensions of ethnic bias

An additional compelling explanation focuses less on the question of self-monitoring and more on the kind of bias to which the DG is sensitive, even in the absence of self-monitoring. As raised in Blum, Hazlett and Posner (2021), the cross-group biases that exist in settings like Nairobi may not be the types of biases that can be expected to be revealed in DG behavior. The DG was designed to capture preferences for altruism versus selfishness toward others (Kanneman, Knetsch and Thaler, 1986). However, ethnic bias may be rooted in disliking, fear, and threat perception, which may be

better detected via other measurement tools. If the conflictual patterns of intergroup behavior that informed our priors about the level of ethnic bias we should have detected in our setting were driven by biases of the latter sort, then this may explain our inability to find cross-group bias in the DG. This conjecture is supported by the fact that, as noted, Berge et al. (2020), Blum, Hazlett and Posner (2021), and Posner (2026), who all work in Nairobi and find no evidence of ethnic bias in the DG, *do* find evidence of ethnic bias in the Implicit Attitude Test or the Affect Misattribution Procedure, measures designed to capture differential affect. Blum, Hazlett and Posner (2021) also find evidence of ethnic bias in a Weapon Misidentification Task, designed to capture fear and threat perception of outgroup members. Taken together, these findings are consistent with the existence of cross-group biases, detectable via other measurement strategies, to which the DG is blind. More broadly, these findings are consistent with literature suggesting that group-based orientations are multi-dimensional, where distinct dimensions require different measurement approaches and have differential validity in predicting specific outcomes.[26] The implication is that researchers seeking to measure the cross-group biases underlying social strife should employ tools other than the DG—not because the DG is susceptible to self-monitoring, but because it is not the right tool for measuring the kinds of biases that are at work in generating intergroup conflict in these settings.

5.6 Conclusion

Accurately measuring intergroup bias is central to understanding and addressing persistent social inequalities. Our study asked whether a distraction task performed concurrently with the DG could "unmask" ethnic biases that are widely believed to exist in this setting, but that the DG has thus far failed to detect. We find no evidence of such unmasking.

Our results point principally toward two plausible interpretations. First, self-monitoring in this context may not rely on the kind of deliberative, resource-intensive processes that are easily disrupted by cognitive load. Norm-consistent behavior—such as equal treatment across ethnic lines in an allocation task—may instead operate as a well-internalized heuristic that persists even under substantial cognitive constraint. Second, the DG itself may be ill-suited to detecting the forms of ethnic bias that are most salient in this setting. If intergroup antagonism is driven primarily by fear, threat perception, or negative affect rather than differential altruism, then the absence of bias in DG behavior does not indicate its absence more broadly, but rather a mismatch between the measurement tool and the underlying construct. Either of these explanations would be consistent with our results, and both may operate simultaneously.

Together, these interpretations underscore the need for conceptual clarity about what different

[26] For example, see literature on Americans' racial orientations including Roos, Hughes and Reichelmann (2019); Davis and Wilson (2024); Kam and Burge (2019); DeSante and Smith (2020).

behavioral measures can and cannot capture, and caution against interpreting null results in the DG as evidence against the presence of meaningful intergroup bias.

References

Adida, Claire L., David D. Laitin and Marie-Anne Valfort. 2014. "Muslims in France: Identifying a discriminatory equilibrium." *Journal of Population Economics* 27(4):1039–1086.

Arechar, Antonio A. and David G. Rand. 2022. "Learning to be selfish? A large-scale longitudinal analysys of Dictator games played on Amazon Mechanical Turk." *Journal of Economic Psychology* 90:1–6.

Bardsley, Nicholas. 2008. "Dictator game giving: Altruism or artefact?" *Experimental Economics* 11(2):122–133.

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.

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 Hazlett and Daniel N Posner. 2021. "Measuring ethnic biases: Can misattribution-based tools from social psychology reveal group biases that economics games cannot?" *Political Analysis* 29(3):385–404.

Chen, Yan and Sherry Xin Li. 2009. "Group identity and social preferences." *American Economic Review* 99(1):431–357.

Choi, Danny, J. Andrew Harris and Fiona Shen-Bayh. 2022. "Ethnic bias in judicial decision making: Evidence from criminal appeals in Kenya." *American Political Science Review* 116(3):1067–1080.

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.

Cornelissen, Gert, Siegfried Dewitte and Luk Warlop. 2011. "Are social value orientations expressed automatically? Decision making in the Dictator Game." *Personality and Social Psychology Bulletin* 37(8):1080–1090.

Cunningham, Tom and Jonathan de Quidt. 2016. "Implicit preferences inferred from choice." *CESifo Working Paper Series* (5704).

Davis, Darren W. and David C. Wilson. 2024. *Racial resentment in the political mind*. University of Chicago Press.

De Neys, Wim. 2006. "Dual processing in reasoning: Two systems but one reasoner." *Psychological Science* 17(5):428–433.

de Quidt, Jonathan, Johannes Haushofer and Christopher Roth. 2018. "Measuring and bounding experimenter demand." *American Economic Review* 108(11):3266–3302.

Deck, Cary and Salar Jahedi. 2015. "The effect of cognitive load on economic decision making: A survey and new experiments." *European Economic Review* 78:97–119.

Deck, Cary, Salar Jahedi and Roman Sheremeta. 2021. "On the consistency of cognitive load." *European Economic Review* 134:1–19.

DeSante, Christopher D. and Candis Watts Smith. 2020. "Fear, institutionalized racism, and empathy: The underlying dimensions of Whites' racial attitudes." *PS: Political Science & Politics* 53(4):639–645.

Devine, Patricia G. 1989. "Stereotypes and prejudice: Their automatic and controlled components." *Journal of personality and social psychology* 56(1):5.

Druckman, James N., Samara Klar, Yanna Krupnikov, Matthew Levendusky and John Barry Ryan. 2022. "(Mis)estimating affective polarization." *The Journal of Politics* 84(2):1106–1117.

Dunton, Bridget C. and Russell H. Fazio. 1997. "An individual difference measure of motivation to control prejudiced reactions." *Personality and Social Psychology Bulletin* 23(3):316–326.

Effron, Daniel A., Jessica S. Cameron and Benoit Monin. 2009. "Endorsing Obama licenses favoring Whites." *Journal of Experimental Social Psychology* 45:590–593.

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.

Enos, Ryan D. 2017. *The space between us: Social geography and politics*. Cambridge University Press.

Evans, Jonathan St. B. T. and Keith E. Stanovich. 2013. "Dual-Process Theories of Higher Cognition: Advancing the Debate." *Perspectives on Psychological Science* 8(3):223–241.

Fafchamps, Marcel. 2000. "Ethnicity and credit in African manufacturing." *Journal of Development Economics* 61(1):205–235.

Fershtman, Chaim and Uri Gneezy. 2001. "Discrimination in a segmented society: An experimental approach." *Quarterly Journal of Economics* 116(1):351–377.

Gibson, Clark C. and James D. Long. 2009. "The presidential and parliamentary elections in Kenya, December 2007." *Electoral Studies* 28(3):497–502.

Gilbert, Daniel T, Romin W Tafarodi and Patrick S Malone. 1993. "You can't not believe everything you read." *Journal of personality and social psychology* 65(2):221.

Glahn, D.C., J. Kim, M.S. Cohen, V.P. Poutanen, S. Therman, S. Bava, T.G.M. Van Erp, M. Manninen, M. Huttunen, J. Lönnqvist et al. 2002. "Maintenance and manipulation in spatial working memory: Dissociations in the prefrontal cortex." *Neuroimage* 17(1):201–213.

Greenwald, Anthony G., Debbie E. McGhee and Jordan L.K. Schwartz. 1998. "Measuring individual differences in implicit cognition: The implicit association test." *Journal of Personality and Social Psychology* 74(6):1464–1480.

Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the dilemmas of collective action*. Russell Sage Foundation.

Haley, Kevin J. and Daniel M.T. Fessler. 2005. "Nobody's watching? Subtle cues affect generosity in an anonymous economic game." *Evolution and Human Behavior* 26(3):245–256.

Harris, Kirk A. 2024. "Ethnic politics and party realignment in African constitutional referendums: Understanding Kenya's 'industry of insults'." *Afrian Affairs* 123(490):55–74.

Hauge, Karen Evelyn, Kjell Arne Brekke, Lars-Olof Johansson, Olof Johansson-Stenman and Henrik Svedster. 2016. "Keeping others in our mind or in our heart? Distribution games under cognitive load." *Experimental Economics* 19:562–576.

Haushofer, Johannes, Sara Lowes, Musau Abednego, David Ndetei, Nathan Nunn, Moritz Poll and Nancy Qian. 2023. "Stress, ethnicity, and prosocial behavior." *Journal of Political Economy, Microeconomics* 1(2):225–269.

Hazlett, Chad, Emilie Campos, Patricia Tan, Holly Truong, Sandra Loo, Charlotte DiStefano and Damla Jeste, Shafali Senturk. 2020. "Principle ERP reduction and analysis: Estimating and using principle ERP waveforms underlying ERPs across tasks, subjects and electrodes." *Under review* .

Hjort, Jonas. 2014. "Ethnic divisions and production in firms." *Quarterly Journal of Economics* 129(4):1899–1946.

Hornsby, Charles. 2013. *Kenya: A history since independence*. I.B. Tauris.

Iyengar, Shanto and Sean J. Westwood. 2015. "Fear and loathing across party lines: New evidence on group polarization." *American Journal of Political Science* 59(3):690–707.

James, William. 1890. *The principles of psychology.* Henry Holt.

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. 2011. *Thinking fast and slow.* Farrar, Straus and Giroux.

Kam, Cindy D. and Camille D. Burge. 2019. "TRENDS: Racial resentment and public opinion across the racial divide." *Political Research Quarterly* 72(4):767–784.

Kanneman, 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.

Kasfir, Nelson. 1976. *The shrinking political arena: Participation and ethnicity in African politics, with a case study of Uganda.* University of California Press.

Kibanja, G.M., M.M. Kajumba and L.R. Johnson. 2012. Ethnocultural conflict in Uganda: Politics based on ethnic divisions inflame tensions across the country. In *Handbook of ethnic conflict*, ed. D. Landis and R. Albert. Springer.

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.

Kirchner, Wayne K. 1958. "Age differences in short-term retention of rapidly changing information." *Journal of experimental psychology* 55(4):352.

Krosnick, Jon A., Tobias H. Stark and Amanda L. Scott. 2025. *Cambridge handbook of implicit bias and racism.* Cambridge University Press.

Lavie, Nilli. 2005. "Distracted and confused?: Selective attention under load." *Trends in cognitive sciences* 9(2):75–82.

Lenartowicz, Agatha, Holly Truong, Giulia C. Salgari, Robert M. Bilder, James McGough, James T. McCracken and Sandra K. Loo. 2019. "Alpha modulation during working memory encoding predicts neurocognitive impairment in ADHD." *Journal of Child Psychology and Psychiatry* .

Levitt, Steven D. and John A. List. 2007. "What do laboratory experiments measuring social preferences reveal about the real world?" *Journal of Economic Perspectives* 21(2):153–174.

List, John A. 2007. "On the interpretation of giving in dictator games." *Journal of Political Economy* 115(3):482–493.

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.

Miguel, Edward and Mary Kay Gugerty. 2005. "Ethnic diversity, social sanctions, and public goods in Kenya." *Journal of Public Economics* 89:2325–2368.

Mironova, Vera and Sam Whitt. 2014. "Ethnicity and altruism after violence: The contact hypothesis in Kosovo." *Journal of Experimental Political Science* 1(2):170–180.

Murphy, Sheila T and Robert B Zajonc. 1993. "Affect, cognition, and awareness: affective priming with optimal and suboptimal stimulus exposures." *Journal of personality and social psychology* 64(5):723.

Orr, Lilla V. and Gregory A. Huber. 2020. "The policy basis of measured partisan animosity in the United States." *American Journal of Political Science* 64(3):569–586.

Payne, B Keith. 2001. "Prejudice and perception: the role of automatic and controlled processes in misperceiving a weapon." *Journal of personality and social psychology* 81(2):181.

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. 2026. "Urbanization and interethnic relations in Africa." Unpublished paper.

Rand, David G., Alexander Peysakhovich, Gordon T. Kraft-Todd, George E. Newman, Owen Wurzbacher, Martin A. Nowak and Joshua D. Greene. 2014. "Social heuristics shape intuitive cooperation." *Nature communications* 5(1):1–12.

Roos, J. Micah, Michael Hughes and Ashley V. Reichelmann. 2019. "A puzzle of racial attitudes: A measurement analysis of racial attitudes and policy indicators." *Socius* 5:2378023119842738.

Rosenblatt, Tanya S. 2008. "The beauty premium: Physical attractiveness and gender in dictator games." *Negotiation Journal* 24(4):465–481–48.

Schulz, Jonathan F., Urs Fischbacher, Christian Thoni and Verena Utikal. 2014. "Affect and fairness: Dictator games under cognitive load." *Journal of Economic Psychology* 41:77–87.

Sherman, David K. and Geoffrey L. Cohen. 2006. The psychology of self-defense: Self-affirmation theory. In *Advances in Experimental Social Psychology*, ed. Mark P. Zanna. Vol. 38 Academic Press pp. 183–242.

Shiffrin, Richard M. and Walter Schneider. 1977. "Controlled and Automatic Human Information Processing: II. Perceptual Learning, Automatic Attending and a General Theory." *Psychological Review* 84(2):127–190.

Steele, Claude M. 1988. The psychology of self-affirmation: Sustaining the integrity of the self. In *Advances in Experimental Social Psychology*, ed. Leonard Berkowitz. Vol. 21 Academic Press pp. 261–302.

Tanaka, Tomomi and Colin F. Camerer. 2016. "Trait perceptions influence economic out-group bias: lab and field evidence from Vietnam." *Experimental Economics* 19(3):513–534.

Tinghög, Gustav, David Andersson, Caroline Bonn, Magnus Johannesson, Michael Kirchler, Lina Koppel and Daniel Västfjäl. 2016. "Intuition and moral decision-making: The effect of time pressure and cognitive load on moral judgment and altruistic behavior." *PLoS ONE* 11(10):1–19.

Whitt, Sam and Rick K. Wilson. 2007. "The Dictator Game, fairness and ethnicity in postwar Bosnia." *American Journal of Political Science* 51(3):655–668.

Winking, Jeffrey. 2014. "Anonymity versus privacy in the dictator game: Revealing donor decisions to recipients does not substantially impact donor behavior." *PloS one* 9(12):e115419.

Wrong, Michela. 2009. *It's our turn to eat.* Fourth Estate London.

Zizzo, Daniel John. 2010. "Experimenter demand effects in economic experiments." *Experimental Economics* 13:75–98.

Appendix

A Numerical values shown in figures

Table A.1: Numerical values underlying Figure 3 (differences in ethnic bias across load conditions). Ethnic bias is a participant's allocation to the coethnic minus the non-coethnic receiver; entries are mean within-person contrasts with 95% confidence intervals. Each contrast "A vs. B" is bias under A minus bias under B.

Contrast	Mean (KSh)	95% CI	N
Pooled SDRT vs. Control	-0.17	[-1.17, 0.82]	421
SDRT-1 vs. Control	-0.02	[-1.04, 1.01]	421
SDRT-2 vs. Control	-0.33	[-1.55, 0.89]	421
SDRT-2 vs. SDRT-1	-0.32	[-1.38, 0.75]	421

Table A.2: Numerical values underlying Figure 4 (mean Dictator Game allocations by load condition and receiver profile), with 95% confidence intervals.

Condition	Profile	Mean amount sent (KSh)	95% CI	N
Control	Coethnic	12.17	[11.02, 13.32]	421
Control	Neutral	11.55	[10.42, 12.69]	421
Control	Non-coethnic	11.90	[10.75, 13.05]	421
SDRT-1	Coethnic	11.50	[10.40, 12.61]	421
SDRT-1	Neutral	11.29	[10.17, 12.40]	421
SDRT-1	Non-coethnic	11.25	[10.15, 12.34]	421
SDRT-2	Coethnic	11.72	[10.58, 12.86]	421
SDRT-2	Neutral	11.40	[10.26, 12.54]	421
SDRT-2	Non-coethnic	11.78	[10.60, 12.97]	421

Table A.3: Numerical values underlying Figure 5 (mean ethnic bias by load condition and permission-to-discriminate arm), with 95% confidence intervals.

Condition	Permission	Mean bias (KSh)	95% CI	N
Control	no	0.41	[-0.60, 1.41]	197
Control	yes	0.15	[-0.91, 1.20]	224
SDRT-1	no	0.28	[-0.82, 1.38]	197
SDRT-1	yes	0.23	[-0.73, 1.19]	224
SDRT-1,2	no	0.14	[-0.76, 1.04]	197
SDRT-1,2	yes	0.05	[-0.79, 0.90]	224
SDRT-2	no	0.00	[-1.41, 1.41]	197
SDRT-2	yes	-0.12	[-1.27, 1.03]	224

Table A.4: Numerical values underlying Figure 6 (top): Proportion responding to dot probe by condition and incentive arm, with 95% confidence intervals.

Condition	Motivation	Response rate	95% CI	N
SDRT-1	intrinsic	0.822	[0.788, 0.857]	214
SDRT-1	extrinsic	0.868	[0.840, 0.896]	207
SDRT-1,2	intrinsic	0.796	[0.768, 0.824]	214
SDRT-1,2	extrinsic	0.837	[0.815, 0.860]	207
SDRT-2	intrinsic	0.769	[0.734, 0.805]	214
SDRT-2	extrinsic	0.807	[0.775, 0.839]	207

Table A.5: Numerical values underlying Figure 6 (bottom): mean ethnic bias by condition and SDRT-incentive arm, with 95% confidence intervals.

Condition	Motivation	Mean bias (KSh)	95% CI	N
Control	extrinsic	-0.20	[-1.17, 0.76]	207
Control	intrinsic	0.72	[-0.37, 1.82]	214
SDRT-1	extrinsic	0.08	[-1.02, 1.17]	207
SDRT-1	intrinsic	0.42	[-0.53, 1.38]	214
SDRT-1,2	extrinsic	-0.03	[-0.91, 0.85]	207
SDRT-1,2	intrinsic	0.21	[-0.64, 1.07]	214
SDRT-2	extrinsic	-0.14	[-1.40, 1.13]	207
SDRT-2	intrinsic	0.00	[-1.27, 1.28]	214

Table A.6: Covariate balance across the randomized ordering of the load blocks (defined by which block came first). Entries are group means, or proportions for binary covariates.

Covariate	Control first (n=126)	SDRT-1 first (n=160)	SDRT-2 first (n=135)	p
Female	0.63	0.68	0.64	0.68
Luo (vs. Kikuyu)	0.48	0.51	0.50	0.86
Years in Nairobi	20.87	21.00	20.51	0.93
Motivation to control prejudice	0.45	-0.02	0.66	0.18
Intermarriage-attitude differential	0.19	0.39	0.36	0.21
Out-group contact differential	0.38	0.63	0.53	0.28
Diverse neighborhood	0.94	0.95	0.95	0.98

The p column reports covariate-by-covariate balance tests (one-way ANOVA for continuous characteristics, χ^2 for binary characteristics). A joint test entering all characteristics together (multinomial logistic regression of ordering on the covariates) gives $p = 0.59$ ($n = 361$).

B Robustness tests

Figure B.1: Distribution of DG allocations, by profile and level of distraction

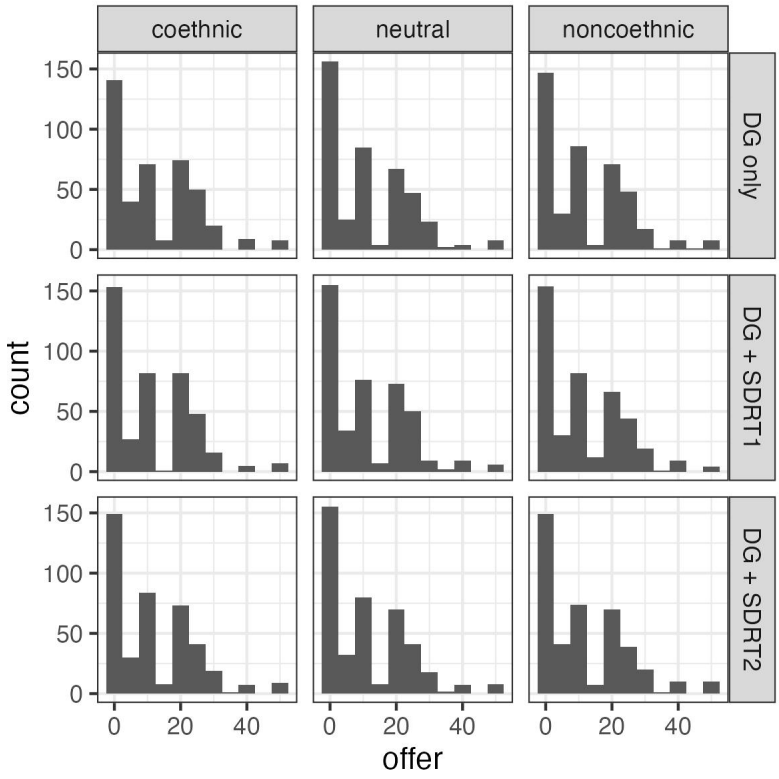


Figure B.2: Differences in ethnic bias, by condition, dropping participants who anchor

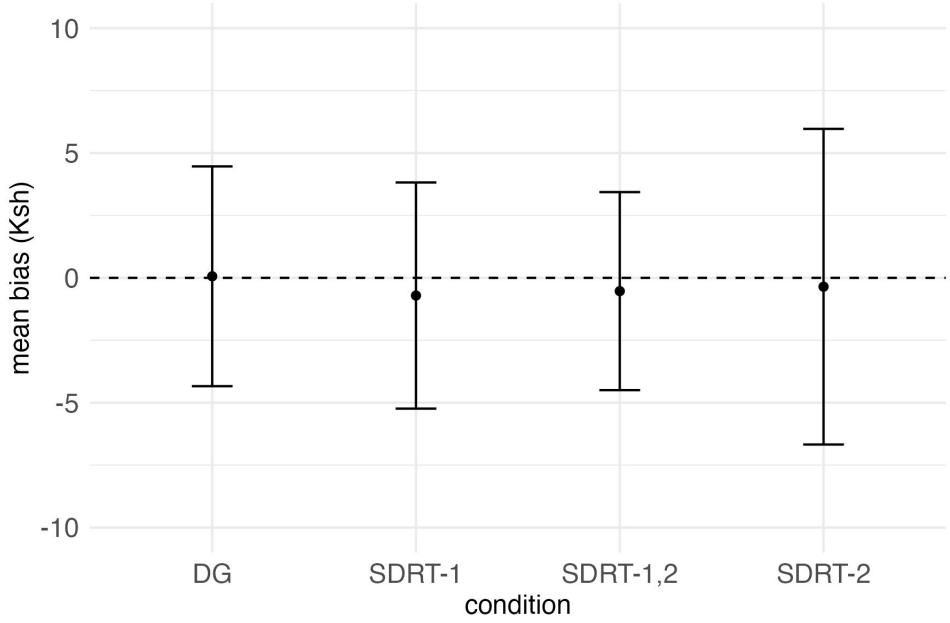


Figure B.3: Differences in ethnic bias, by condition, round three only

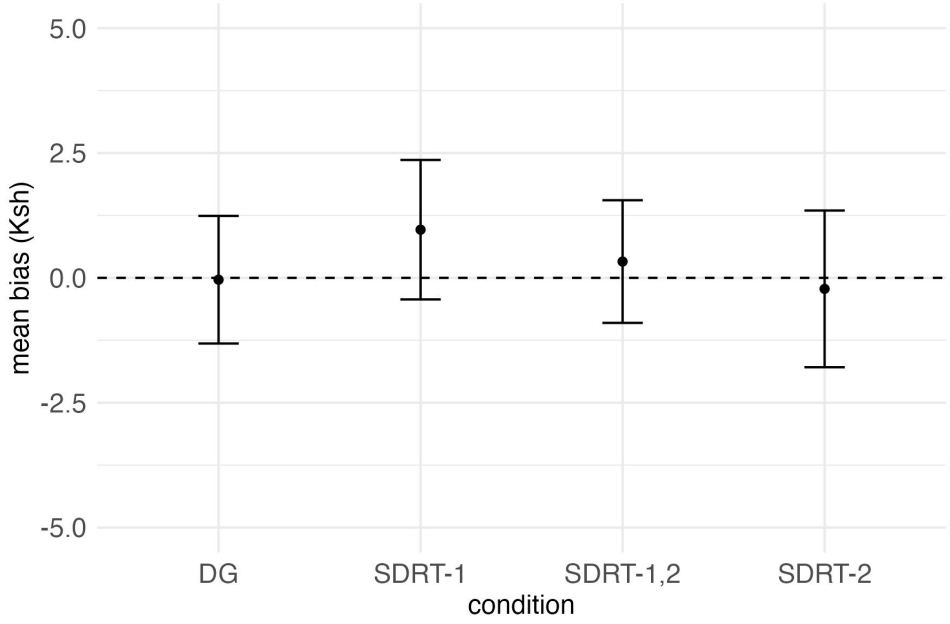


Figure B.4: Differences in ethnic bias, by condition, alternative literacy filter

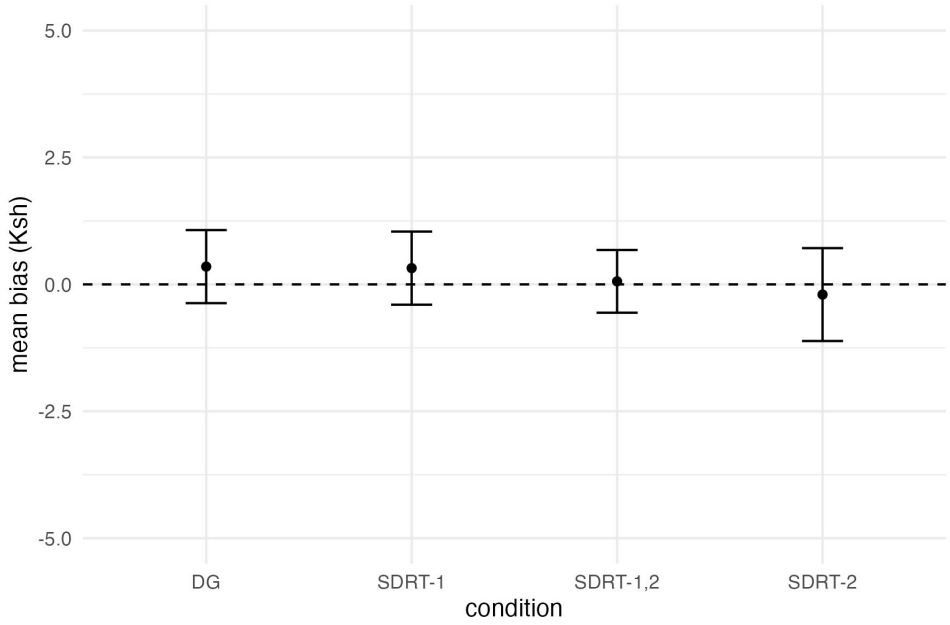


Figure B.5: Differential effects by mixed ethnicity

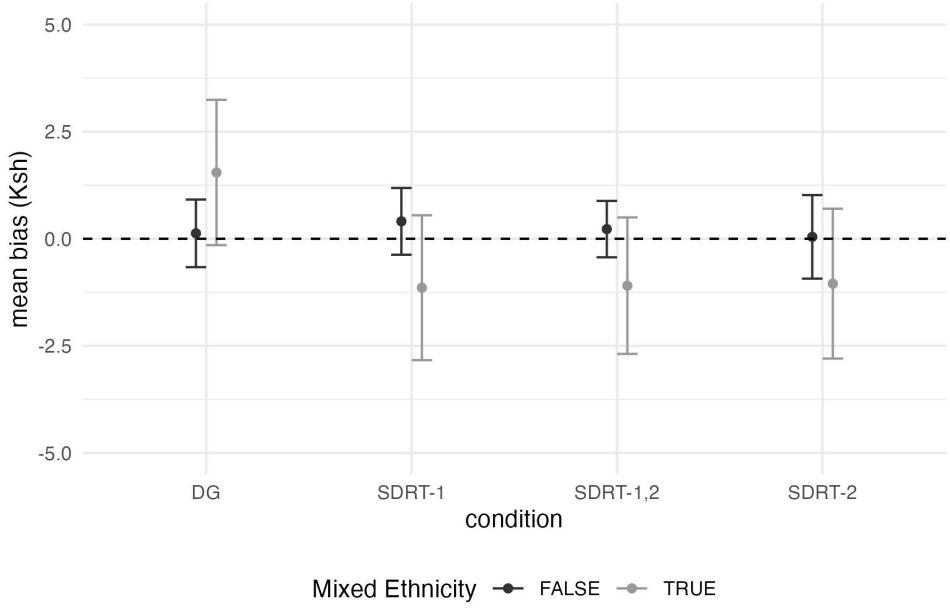


Figure B.6: Differences in ethnic bias, by condition, retaining mixed (K/L) parentage

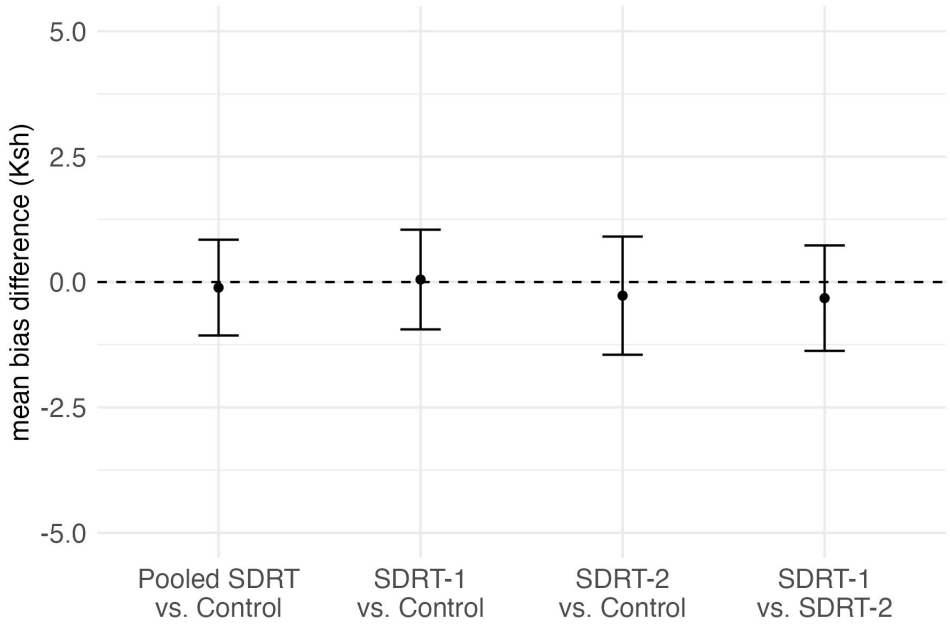
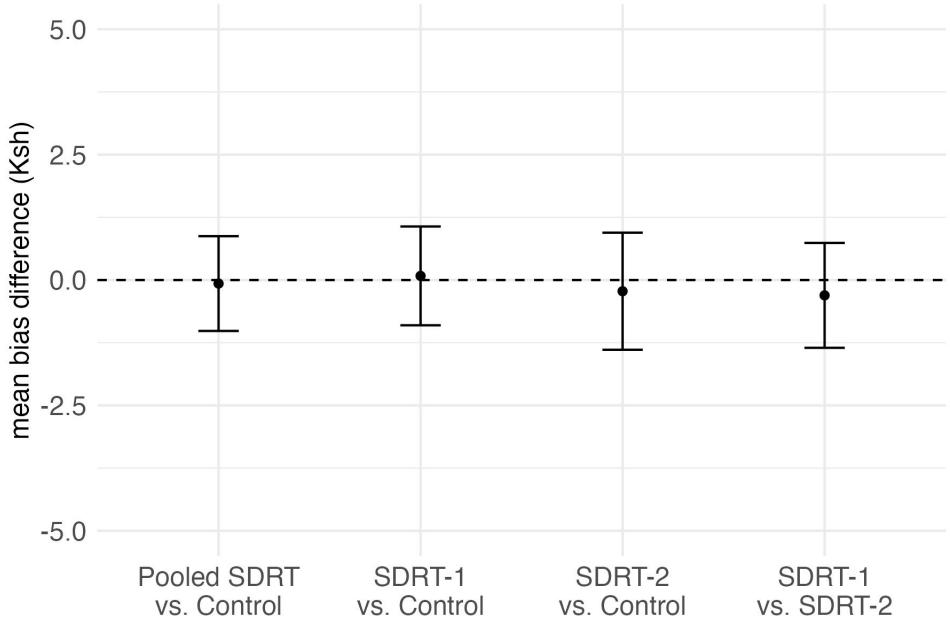


Figure B.7: Differences in ethnic bias, by condition, PAP ethnicity coding (includes mixed K/L parentage)



C Heterogeneous treatment effects

Figure C.1: Differential effects by ethnicity

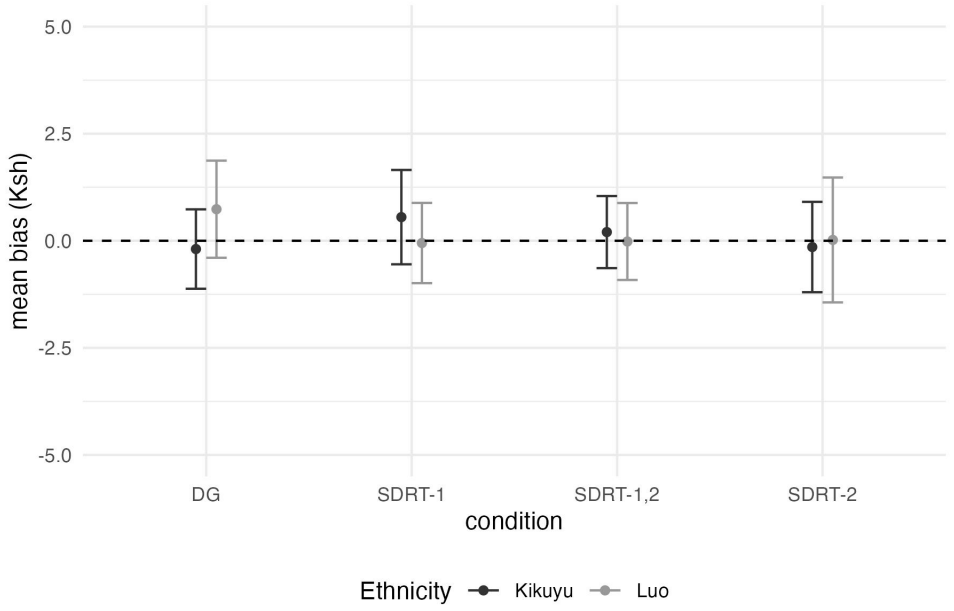


Figure C.2: Differential effects by length of residence in Nairobi

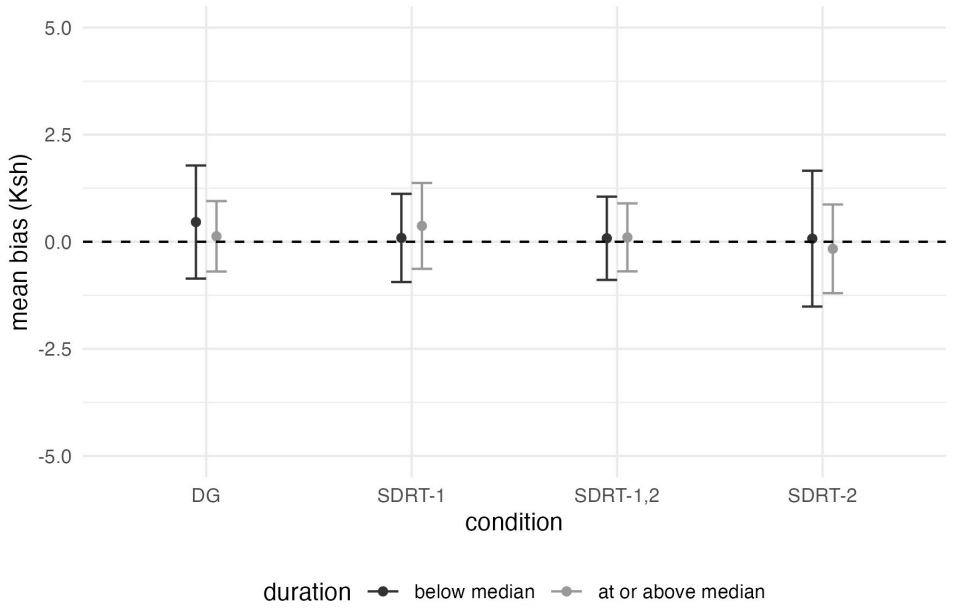


Figure C.3: Differential effects by the diversity of the participant's neighborhood

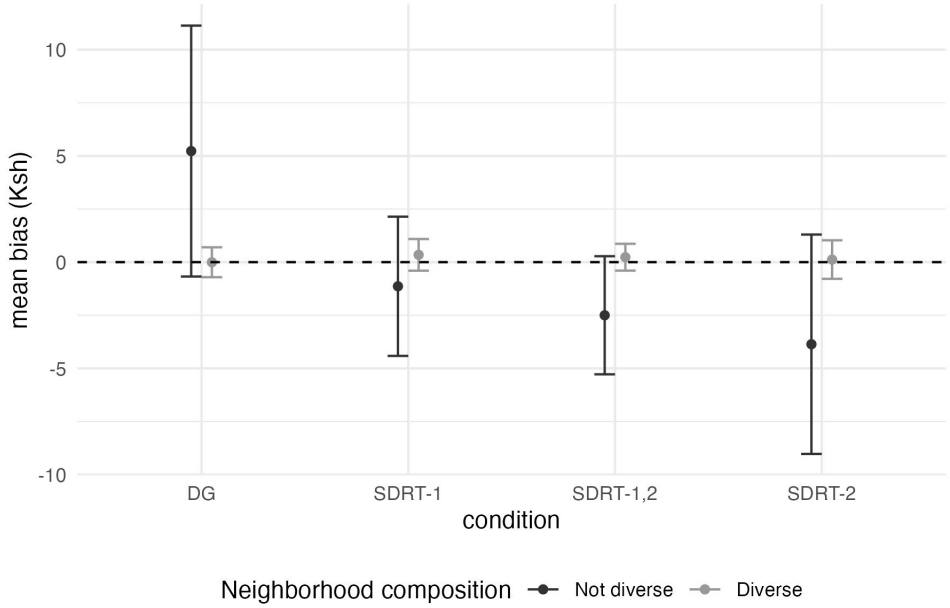


Figure C.4: Differential effects by reported contact with non-coethnics

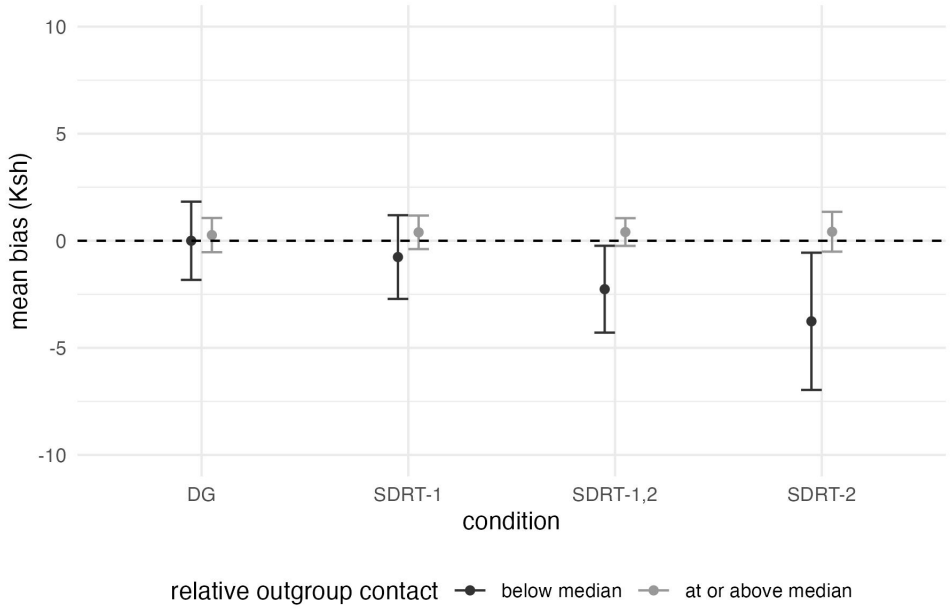


Figure C.5: Differential effects by motivation to control prejudice

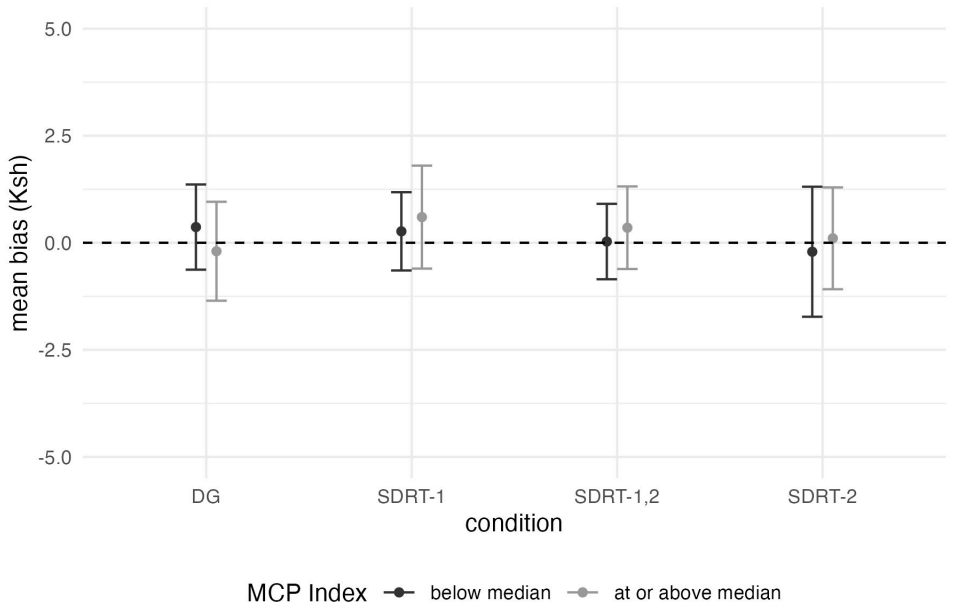


Figure C.6: Differential effects by gender

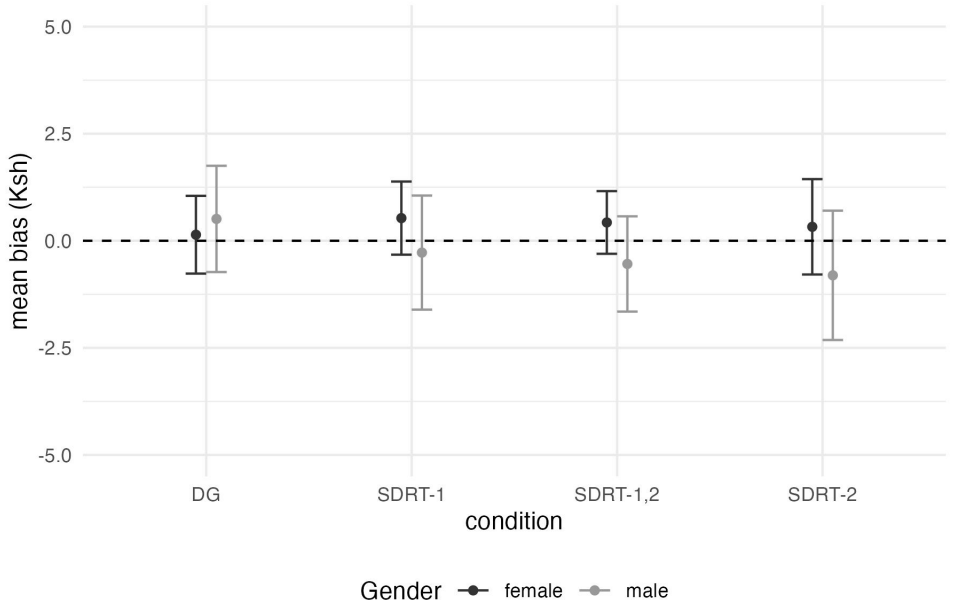
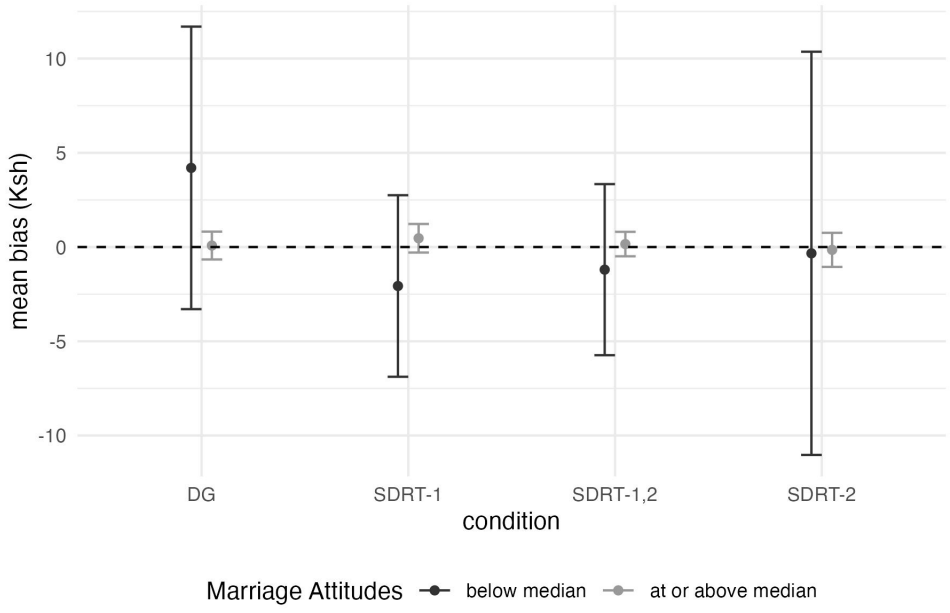


Figure C.7: Differential effects by intermarriage attitudes



D Assessing the impact of the SDRT

Figure D.1: Reported effort in the SDRT

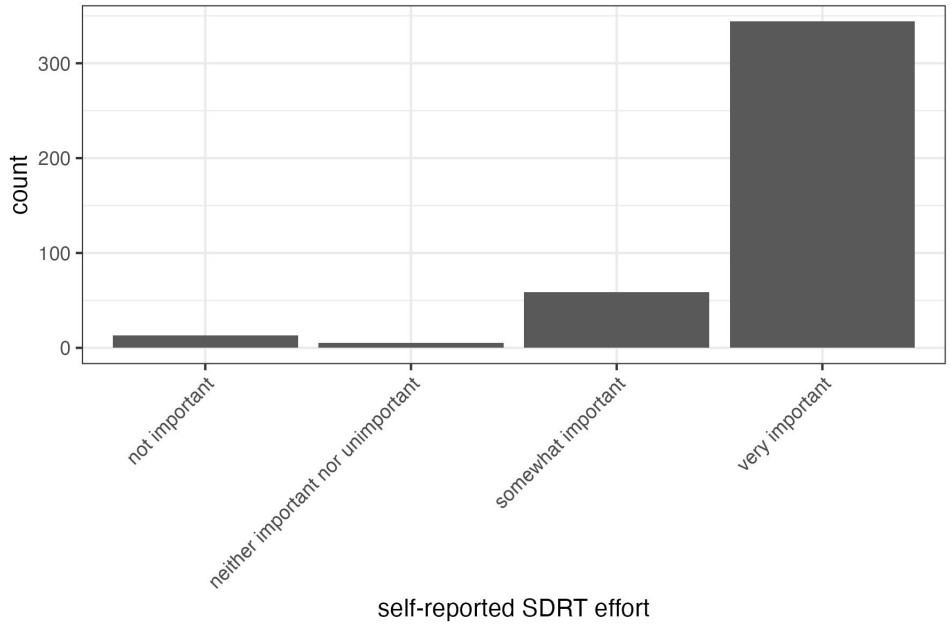


Figure D.2: Differences in ethnic bias by condition and SDRT performance

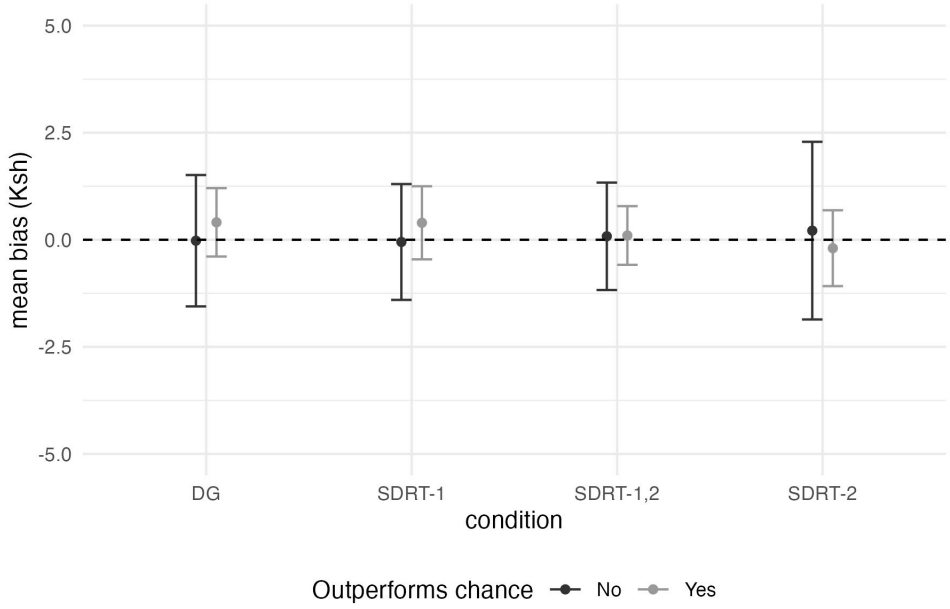


Figure D.3: Differential effects by self-reported SDRT effort

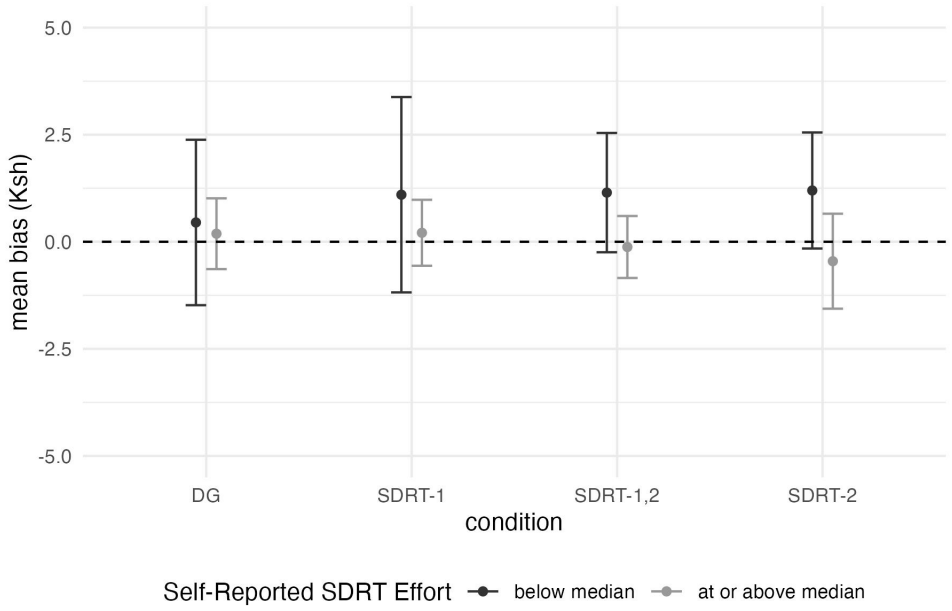


Table D.1: Effect of SDRT on math test

% correct under no load	97.15
% correct under SDRT-1	92.64
% correct under SDRT-2	96.20
Math score$_{1vs0}$	-4.51***
Math score$_{2vs0}$	-0.95
Math score$_{1,2vs0}$	-2.73***

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

Figure D.4: Effect of the SDRT on profile processing

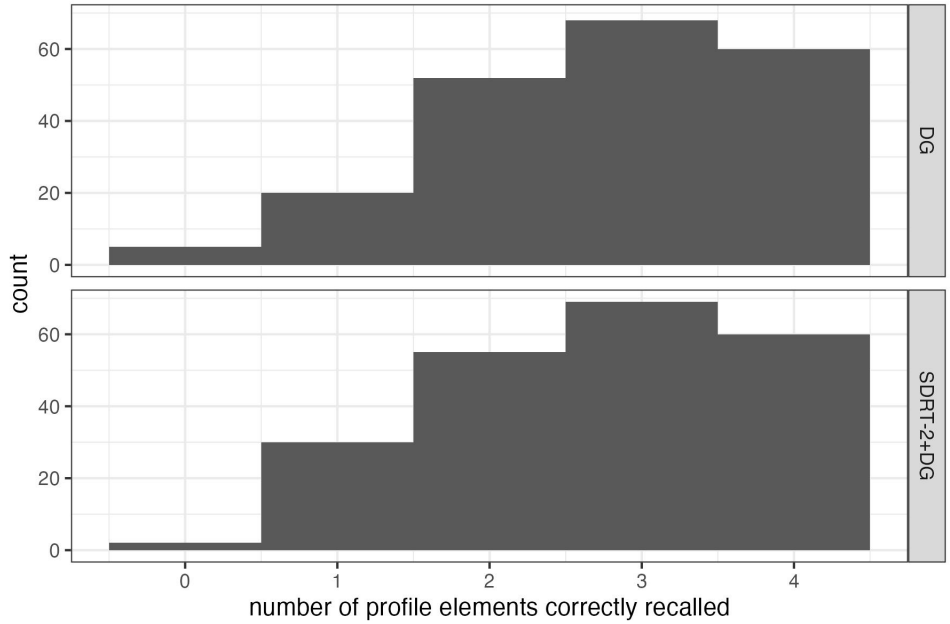


Table D.2: Profile cues correctly recalled, by level of cognitive load

	Name	Age	Education	Hometown
No load	0.659	0.834	0.459	0.820
Load (SDRT-2)	0.639	0.750	0.509	0.819
Difference	0.020	0.084**	-0.050	0.001

Note: With the possible exception of recall for age, there was little difference in the number of profile elements that were correctly recalled by load condition (see Appendix Figure D.4).
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

E Modeling allocations based on receiver profile

We fit a linear regression modeling DG allocations as a function of ethnic status (coethnic as baseline), age, and education with participant fixed effects (omitted from the results) and robust standard errors clustered at the participant level. This model ignores the distraction condition (DG vs. SDRT-1 vs. SDRT-2). Readers should note that coethnicity was the only randomized receiver profile element; name, photo, age, and education were fixed within set profiles. As such, we cannot interpret these coefficients as causal.

Table E.3

	Dependent variable:
	Amount sent to receiver (KSh)
Neutral (NBO)	−0.388
	(0.243)
Non-coethnic	−0.161
	(0.246)
Education	−0.267**
	(0.127)
Age	0.006
	(0.032)
Constant	11.052***
	(1.009)
Observations	3,789
R^2	0.760
Adjusted R^2	0.730
Residual Std. Error	6.154 (df = 3364)
F Statistic	25.180*** (df = 424; 3364)
Note:	*p<0.1; **p<0.05; ***p<0.01

F Pre-registration

We filed a pre-analysis plan for the study at OSF after data collection had occurred and after inspecting a handful of screening-related variables, but before looking at any outcomes. The handful of deviations from the pre-analysis plan are detailed below in Appendix F.1. The pre-analysis plan itself is pasted below as Appendix F.2.

F.1 Deviations from the pre-analysis plan

We deviated from the pre-analysis plans as follows:

- As noted, we pre-specified including participants of mixed Kikuyu-Luo parentage as either Kikuyu or Luo based on the language use reported at registration. However, we decided to exclude such mixed-ethnicity participants. As shown in Appendix B, the results are unchanged when we include these excluded participants.

- We pre-registered the intention to report the impact of the SDRT on allocations to receivers with neutral profiles, in addition to those with profiles identifying them as Kikuyu or Luo. The purpose was to measure effect of cognitive load on pro-sociality. Although we do not directly address this issue in the paper, we shown in Figure 4 that allocations to receivers with neutral profiles do not vary across treatment conditions. Cognitive load would not appear to affect prosociality in general.

F.2 Pre-analysis plan

Uncovering Ethnic Discrimination: Using Distraction to Detect Hidden Bias in Economics Games

Registered at OSF on 1 April 2024

Background

Behavioral economics games such as the Dictator Game (DG) have been widely used to quantify ethnic or other between-group biases. However, in numerous contexts where ethnicity is thought to be extremely salient, the DG has failed to detect any evidence of ethnic bias. One compelling hypothesis is that economics games like the DG are explicit measures that are susceptible to self-monitoring and social desirability concerns. Specifically, participants may alter their behavior to fit social norms that discourage ethnic bias or discrimination, even in cases where implicit ethnic bias is present.

Drawing on dual process models from psychology, we randomly manipulate cognitive load during the DG by having participants concurrently execute a distraction task, the Spatial Delayed Recall Task (SDRT). If self-monitoring explains the apparent lack of ethnic bias under DG, and if self-monitoring can be reduced through such distraction, then we expect evidence of bias to be "unmasked" in the DG when played concurrently with the distraction task. We tested this among 558 Kikuyu and Luo participants in Nairobi, Kenya during October 2022.[27] This is the same location and lab where previous DG studies among these groups have routinely failed to produce evidence of ethnic bias, even where other evidence suggests such bias exists.

Commitments and timing

We composed this pre-analysis plan (PAP) after data collection occurred and screening-related variables have been inspected, but before looking at any outcomes, or even constructing the ethnicity variable that will be required for the analyses.

We report here our planned analysis. We commit to making available the results of all analyses described here, whether in the main paper or elsewhere. The only anticipated exception would be if there is an unforeseen reason why an analysis described here cannot be run. We do not commit to performing only the analyses described here, but we do commit to clearly labeling any analysis not registered here as exploratory or otherwise apart from the planned analyses.

[27] We note that the number 558, included in our original PAP, was incorrect: 3 of these reflect a test run by staff members, and 6 are duplicate records of participants from a session that was aborted and relaunched, leaving 549, the sample size reported in the paper and included in the replication data.

Sample inclusion and exclusion

A central concern for this analysis regards the proportion of the sample we can expect understood the task and implemented it as intended, as we can only expect to learn about the task in this group. Indeed in pilot work, we found such selection to be critical as fewer participants than hoped showed evidence of sufficient comprehension or competencies to have performed the task as intended.

We will use the following as our primary sample for analysis, recognizing that additional analyses that are more or less inclusive may be reported as well. To be included, a participant must get:

- at least 5 out of 6 questions correct on the literacy test
- at least 2 out of 3 questions correct on DG practice math.
- at least 3 of 4 SDRT side-by-side training questions correct.

The above criteria leave us with 82% of the data (450 of the 549). Of the excluded 99, 65 would be dropped due to the literacy filter alone, with the other two only excluding an additional 34 participants.

Coding of ethnicity

- The primary ethnicity coding for an individual will come from their self-reported ethnicity during intake. All study participants identify as either Kikuyu or Luo.
- We also have mother and father ethnicity. If one parent is {Kikuyu, Kamba} and the other is {Luo, Luhya}, we can indicate the participant is of mixed (relevant) ethnicity. This can be used in secondary analysis to determine if the effects look different without these participants included.

Level 1 analyses: Primary effects

This category covers the analyses we commit to (i) reporting and (ii) relying on to make the primary statements regarding our conclusions. The main estimates to be obtained are the **estimated ethnic bias** (point estimate and confidence interval) for the ethnic DG at different levels of SDRT difficulty (0, 1, 2, and 1-2 pooled), and comparisons between such ethnic bias measures as difficulty increases (the "distraction effect"). We note that while increased difficulty presumably increases the distraction effect, if the SDRT is too difficult, participants may "give up". In pilot experiments, we used the SDRT with one dot (SDRT-1) and with three (SDRT-3). The latter was apparently too difficult, as performance on the task was worse than chance. In this experiment we thus attempt the SDRT-1 and SDRT-2, but we remain unsure whether the SDRT-2 difficulty will prove effective or too great.

The approach to estimation is simple: **reduce the data down to the relevant per-participant scores**. Within each participant i, the ethnic bias under SDRT-0 ($bias_0$) is given by the mean contribution to co-ethincs minus the mean contribution to non-coethnics. Ethnic bias under SDRT-1 ($bias_1$), SDRT-2 ($bias_2$), and pooling SDRT-1 and SDRT-2 ($bias_{1,2}$) are defined in the same way. We will show distributions of $bias_0$, $bias_1$, $bias_2$, and $bias_{1,2}$ across participants, and compute their means, $\overline{bias_0}$, $\overline{bias_1}$, $\overline{bias_2}$, and $\overline{bias_{1,2}}$. For these mean biases, we will compute Neyman-style standard errors.

The primary estimates of distraction effects are given by contrasts of these,

$$\text{Distraction}_{1,2|0} = \overline{bias_{1,2}} - \overline{bias_0}$$

$$\text{Distraction}_{1|0} = \overline{bias_1} - \overline{bias_0}$$

$$\text{Distraction}_{2|0} = \overline{bias_2} - \overline{bias_0}$$

$$\text{Distraction}_{2|1} = \overline{bias_2} - \overline{bias_1}$$

Inference requires only two-sample t-tests with unequal variance or the equivalent.[28]

Level 2 analyses: Additional analyses

In this section we make note of additional analyses that are of interest, albeit it secondary. We report here (i) our ex ante intentions to conduct these analyses, (ii) our commitment to report these results publicly, though they may be relegated to an appendix depending on the nature of the publication.

Additional groupings and exclusion. We will repeat the analyses while:

- removing participants of mixed identity.
- including individuals with with literacy scores of 4.

Heterogeneous effects. We will consider the following planned analyses for purposes of studying heterogeneity, whilerecognizing power limitations may limit their informativeness. We will examine heterogeneous effects by:

- permission to discriminate (manipulated by design)
- length of residence (median split),
- gender (male vs. female identification)

[28] Note added in revision: in the analyses reported in the paper, each distraction effect is a mean of within-person differences in bias (e.g., $\text{bias}_{i,1} - \text{bias}_{i,0}$), tested with a one-sample t-test on those per-participant differences (equivalently, a paired t-test). The two-sample unequal-variance (Welch) form is used only for the robustness checks restricted to a single DG round per participant, where the two conditions fall on different participants and a within-person difference cannot be formed.

- Kikuyu vs. Luo;
- diverse nieghborhood (binary)
- incentivized performance (manipulated by design)
- motivation to control prejudice (median split; possibly other splits in expoloratory analyses)
- self-reported effort to perform well on the SDRT (median split)
- mixed ethnicity (between {Kikuyu, Kamba} and {Luo, Luhya}
- upset at a family member marrying out-group minus upset at marrying in-group;
- own attitude towards "intermarriage yourself" (median split)
- contact with people from in-group region minus contact with people from out-group region (median split)
- length of stay in Nairobi (median split, though we will also attempt to estimate the shape of the relationship if any in further exploration)

Robustness tests and exploratory analyses. These will be substantially less powered than the above, but we will look.

- repeat level 1 analyses but restricting the data to the first SDRT task assigned, to avoid spillover concerns. In doing so, the difference in ethnic bias by tasks (SDRT-0 vs. other) can only be estimated at the between subject level, reducing power and efficiency sharply.
- compare level 1 results for those whose first partner is a co-ethnic vs. those whose first partner is non-coethnic. Here, unfortunately, even the ethnic bias measure itself (for any SDRT level) moves to between subject level.

Cognitive and mechanistic effects. To check for evidence that the SDRT task had influence through various processes, we will examine:

- the number of math questions correct under SDRT0 vs. SDRT1 or 2 in the practice rounds to see how cognitive load (as induced by SDRT) affect math ability.
- the effect of SDRT0 vs. SDRT1 and SDRT2 on contribution to neutral profiles (those from Nairobi), to measure effect of cognitive load on pro-sociality.

- comparison of coethnic to "neutral" (Nairobi), and non-coethnic to neutral, on SDRT0, SDRT1,2, and the distraction effect. This aids in distinguishing an "out-group denigration" from an "in-group favoritism" motive.

- whether the SDRT interferes with participants ability to process the ethnic cue: a final DG played at the end of the lab session was randomized to SDRT0, 1, 2, after which we ask participants to recall the characteristics from the recipient profile. Evaluate correct recall of profile characteristics, by SDRT-0, SDRT-1, and SDRT-2.