

Measuring Discrimination at the Local Level

Neelanjan Sircar (Columbia University & University of Pennsylvania)

Ty Turley (Brigham Young University)

Maarten Voors (Wageningen University & University of Cambridge)

Peter van der Windt (Columbia University & Wageningen University)*

April 20, 2014

Abstract

A robust literature in the social sciences uses experimental games to study discrimination. Often, these games are played between strangers, and players are given little information about the population from which the other players are drawn. We introduce a formal framework and a set of novel dictator games to make nuanced inferences about social discrimination based on differences in knowledge of the social context and the receiver. Specifically, we explore the effects of 1) varying the distribution of the receiver population, and 2) moving from a known distribution of receivers to full information about each individual receiver. We use this framework to study discriminatory behavior in rural Sierra Leone, a context characterized by discrimination based on elite status—a factor widely regarded as an important cause of the civil war. We find that the magnitude of discrimination is driven by a dictator’s knowledge of the receiver, not by her knowledge about the distribution of receivers. Our findings provide important implications for both the measurement and interpretation of discrimination.

*Corresponding author: pv2160@columbia.edu. We thank Paul Richards and Oliver Vanden Eynde, and participants at NEWEPS, BYU, SEEDEC and Oxford University’s CSAE for helpful comments. We thank Martha Ross, Thijs van Bommel, Lizzy van der Wal, Freek van de Wege, Esther Mokuwa and our team of local enumerators for extraordinary field work in Sierra Leone. Thanks also to Cambridge University’s CCI 0512018, ESRC grant #ES-J017620-1 and NWO grant #453-10-001 for financial support. Finally, we thank our Sierra Leonean participants. This study has been pre-registered at: <http://e-gap.org>.

1 Introduction

In developed countries, many exchanges occur through impersonal interactions. In contrast, a large literature on the political economy of development finds that in developing countries, local characteristics such as reciprocity and network structure are important determinants governing behavior (e.g. Udry (1994) and Fafchamps and Lund (2003)). This paper proposes a framework and develops a novel method to understand and measure discrimination among populations in which individuals know each other intimately. Second, we explore possible biases implied by traditional methods to measure discrimination at this *local level*.

Starting with Becker (1957), economists have explored the causes and consequences of discrimination. Because the measurement of discrimination is prone to issues such as social desirability biases, recent empirical work has moved away from survey-based evidence in favor of more behavioral based measures. Two strands of literature stand out.¹ First, audit studies and ‘kindred’ field experiments have evaluated if minorities are treated differently in job application, housing search, vehicle purchases, etc. For example, Bertrand and Mullainathan (2004), randomize the applicant characteristics on resumes to study the impacts for job applications (see also Riach and Rich (2002), List (2004) and Adida et al. (2010)). A second strand has concentrated on laboratory experiments, predominately relying on the attribute-based dictator game (ADG). Discrimination in the ADG is measured as the difference in allocation by a “dictator” (the sender) to different types of “receivers” that are differentiated by an experimental cue on a particular attribute (for example ethnicity, gender, etc). A central concern with these experiments is that they do not allow for the context in which discriminatory behavior is embedded (e.g. Levitt and List (2007)). First, players do not know anything about their counterparts except for that which is revealed by the experimenter. Second, players are often strangers and do not know each other personally. An allocation decision at the local level, however, also depends on other receiver’s attributes and factors such as previous interaction and the social networks in which the individuals are embedded.

This paper derives a formal and experimental framework to study discrimination based upon differences in knowledge of the social context and the receiver. In particular, we focus on how discrimination varies as a function of the distribution of the receiver population, and how this discrimination changes when moving from a setting where only the distribution of receivers is known to a setting with full information about each receiver. Our framework applies to discrimination in a wide array of settings, from situations where little is known about participants (e.g. a rental market) to situations where participants know quite a bit about each other (e.g. daily interaction in a rural village in Sierra Leone). We argue

¹Up to recently regression studies were also popular. These models attempt to interpret the implications of say race or gender in an OLS framework, typically in wage models. These studies often use a Oaxaca-Blinder decomposition to interpret the magnitude of findings. See Neal and Johnson (1996) for one such study.

that the classic ADG is a good measure of discrimination in the former setting but that an adjusted ADG is necessary for empirical inferences about the latter. In this paper we introduce the revealed-receiver attribute-based dictator game (RDG). This game extends the ADG in two ways to capture behavior at the local level: 1) we reveal the identity of the receiver, and 2) play with subjects who know each other well. To empirically illustrate our framework, we conduct a set of RDG and ADG games at the local level. The latter is implemented with three variations—each with a different level of social distance—which allows us to investigate two mechanisms that might explain differences between the ADG and the RDG at the local level: 1) social distance bias, and 2) aggregation bias. The first refers to differences in beliefs about the overall distribution of relevant attributes in the population (e.g. playing with a group of receivers from the village versus from the country). The second refers to the level of uncertainty a dictator has about the relevant attributes of the receiver (playing with a group of receivers versus playing with those same individuals one-by-one).

We implement the experimental games in a set of small and remote communities in rural Sierra Leone where interactions are predominantly local. Because the reach of government is limited, local elites have great authority to organize economic and social activity, including the power to raise taxes, mobilize labor, settle disputes and allocate resources such as land, labor and reproductive opportunities. It is argued that the exclusionary nature based upon social status was an important contributing factor for the 1991-2002 civil war (e.g. Richards (2005)). As a result, our experimental design leverages the importance of social status in rural daily life and takes this as the cleavage for study to understand discriminatory behavior.

Three major results stand out. First, we find large differences between the RDG and the ADG. Compared to the RDG, we find that in certain cases the ADG overestimates levels of discriminating in the local context by up to a factor of almost ten. Second, we show that the difference between both games is mainly driven by the aggregation bias, and not social distance bias. What is important for discrimination in behavior is a dictator’s knowledge of the receiver, not her knowledge about the distribution of receivers. Finally, a major benefit of the RDG is that—in contrast to the ADG—we are able to investigate the importance of characteristics other than only the attribute under study. In fact, the result that discrimination is based upon social status largely disappears when the researcher also controls for receiver and dyad characteristics.

We believe this study investigates a fundamental question concerning how to study human behavior across settings. While current methods allow researchers to isolate underlying behavioral preferences without contamination by social considerations, this abstraction ignores the social dimension of behavior that could easily swamp the importance of personal preferences in highly socially connected societies. In real life, as Levitt and List (2007) rightly argue, behavior also depends on other receiver’s attributes and factors such as previous interaction and the social networks in which the individuals are embedded.

Such factors we expect to be particularly important determinants of discriminatory behavior at the local level; that is, among individuals that know each other intimately and have regular interaction. In this paper we introduce a new tool to the experimental literature that allows for a better extrapolation of lab experimental results into real-world behavior at this local level.

This paper is organized as follows. In the next section we present currently popular methods to measure discrimination, and we discuss how they relate to behavior at the local level. In this section we also we derive the mathematical foundations of our assertions, and set out our theoretical predictions. In section 3 we introduce the Sierra Leonean context and outline our experimental design. Section 4 presents our results. We explore additional benefits of the RDG to learn about behavior at the local level in Section 5. Section 6 concludes.

2 Measuring Discrimination at the Local Level

2.1 Local Discrimination and a Novel Way to Measure It

We define local discrimination as differential behavior based upon a particular attribute of the other *among individuals that know each other*. In this section we relate this concept to the literature on discrimination and popular methods to measure it. We also derive the mathematical foundations of our assertions, and set out our theoretical predictions.

Statistical discrimination takes a prominent position in the literature on discrimination.² A key feature of statistical discrimination is limited information. Firms, for example, have limited information about the skills of job applicants, which gives them an incentive to use easily observable characteristics correlated with productivity (such as skin color or gender) to infer the expected productivity of applicants. It is thus not surprising that also the methods used to measure discrimination are characterized by limited information. For example, in the audit experiments discussed earlier, employers know little more about the job applicant than what is revealed in the resume (Bertrand and Mullainathan (2004)). Limited information also plays a key role in experimental games.³ Because these games

²Classic references are Aigner and Cain (1977), Phelps (1972), and Arrow (1973). The second type of discrimination that has received attention in the literature, although to a much lesser extent, is taste-based discrimination. Taste-based discrimination describes a situation in which individuals have an inherent prejudice towards a certain group of people. This prejudice is part of the utility function and may reflect dislike, anger, or similar emotions (Fershtman and Gneezy (2001)). For example, in Becker (1957) employers have a ‘taste for discrimination’, reflected in their disutility from employing workers that have a certain attribute such as being black or women.

³Fershtman and Gneezy (2001) is one of the few studies that tries to isolate taste-based and statistical motives for possible discrimination. Taste-based discrimination is measured using differential behavior by the dictator in a dictator game. Statistical discrimination is measured using the trust game, where the statistical discrimination reflects the players’ assessments of the differing reactions of members of groups to their actions. Also these games are characterized by limited information. The two ethnic groups under study, Ashkenazic Jews and Eastern Jews, are identifiable only via the name of *unknown* individuals (similarly to

are the focus of this paper, we will discuss this in more detail.

Among experimental games the attribute-based dictator game (ADG) is the workhorse model used to understand behaviors related to discrimination. The game takes a single player, the sender, and asks her to split a fixed sum of money, the endowment, between herself and another person, the receiver. The sender is under no obligation to donate any money to the receiver.⁴ Before the allocation decision the sender is given some experimentally-controlled “cue” regarding the receiver and no other information, and then plays the game at least twice—each time with a different type of receiver. For example, a researcher interested in discrimination based on social status would design an experiment in which the dictator would have to split a fixed sum of money once between herself and a “high status individual”, and a second time between herself and a “low status individual”. The difference in contribution is then a measure of status-based discrimination. Limited information is a central characteristic to these experimental games because the dictator knows little more about the receiver than what the experimenter reveals, which is in most cases only the cue under study.⁵ The ADG has been used to show discriminatory behavior based on various cleavages: gender (Holm, 2000), ethnicity (Whitt and Wilson (2007), Fershtman and Gneezy (2001)), and partisanship (Fowler and Kam, 2007).

This paper argues that neither the concept of statistical discrimination nor our current methods to measure discrimination relate well to behavior at the local level. The principal reason for this assertion is that social interaction at the local level is not characterized by limited information about the other individuals in the population. The intuitive empirical interpretation of the classic ADG, and other measures of discrimination such as the audit studies, is that it measures the extent of discrimination between random strangers with no, or very limited, information about the receivers’ attributes beyond those under study. These studies therefore relate well to the concept of statistical discrimination, and to situations where little is known about the other party beyond a limited amount of information, such as the housing or the job market. In contrast, members of say a small village in rural Sierra Leone, know much more about each other than only the experimenter’s attribute

Ahmed (2010)).

⁴Predictions based on standard Nash equilibria, assuming individuals are purely interested in personal financial gain, are that the dictator will keep all of the money for herself. Empirically, however, one finds that 60% of dictators donate a positive amount of money towards the receiver, with a mean transfer of around 20% of the endowment (Camerer (2003)). This suggests that people have more complex preferences than personal financial gain, what is typically referred to as exhibiting “other-regarding preferences” (Fehr and Schmidt (1999) and Bolton and Ockenfels (2000)). For a critique of this interpretation see List (2007).

⁵The original setup of (attribute-based) dictator games entailed informing the dictator that the receiver was behind a closed door so that no attributes of the receiver affected the decision of the dictator. The goal of many early studies was simply to demonstrate the inadequacy of utility functions that only assume interest in personal financial gain (e.g. Kahneman et al. (1986)). A set of more recent studies do reveal more about the receiver than only the cleavage under study: e.g. (Habyarimana et al., 2007; Binzel and Fehr, 2013; Ligon and Schechter, 2012; Leider et al., 2009; Burnham, 2003; Charness and Gneezy, 2008; Bohnet and Frey, 1999; Croson, 1996). In most cases, however, these games do not mimic the local context with either the receiver not fully revealed, or the subject pool being strangers. See Sircar et al. (2014) for an overview.

under study. Furthermore, behavior between these villagers is guided by previous experiences with each other. Finally, these villager’s behavior also depends on their position in the village’s social network, and how they relate with other third villagers. Levitt and List (2007) argue how by suppressing such factors, the measurement of discrimination by classic experimental games is likely to result in different levels of discrimination than actually present in real life.⁶ We argue that this is the case especially at the local level.

2.2 Differences in the Measurement of Local Discrimination

This paper proposes a novel experimental game to study discrimination at the local level. Specifically, we introduce what we call the revealed-receiver attribute-based dictator game. This game differs from the ADG in two ways: 1) the identity of the receiver is revealed to the dictator, and 2) subjects know each other well. Second, we implement an experimental framework in the field that allows us to investigate in detail the possible biases that may result from using a classic ADG to measure discrimination at the local level. That is, in addition to playing the RDG, all dictators also play a set of attribute-based games. The latter are classic ADGs in that the dictator only knows the cue under study, in our case social status. Each dictator played three variations of the ADG in which we differed the receiver population: the dictator’s contributions towards a random individual from the chiefdom (ADG-C), the village (ADG-V), or one of the other $n - 1$ participants that also play the game (ADG-P). An overview of the experimental framework is given in Table 1, but we leave a detailed discussion of each game to Section 3. By moving game-by-game from the ADG-C, what we argue would be the classic measure of discrimination, to the RDG, our novel measure for local level discrimination, we explore whether different game setups obtain different estimates of discrimination. Furthermore, this experimental framework allows us to distinguish between two mechanisms that might cause differences in behavior between the ADG-C and the RDG: 1) social distance bias, and 2) aggregation bias.

Social Distance Bias. Social distance bias refers to sensitivity of behavior to knowledge about the distribution of receivers. Specifically, in this paper we consider this bias to be the difference between the behavior of the dictator in the ADG with respect to a population frame with which she is less familiar (ADG-C and ADG-V) compared to one from the local context (ADG-P).⁷ To illustrate this bias, imagine the following thought experiment. An experimenter interested in discrimination towards the homeless, might

⁶Levitt and List (2007) propose that divergences may be caused by (i) the stakes of the game (ii) the presence of moral and ethical considerations, (iii) the extent to which one’s actions are scrutinized by others and the nature of such scrutiny, (iv) the subject pool of respondents and (v) the context in which the decision is embedded.

⁷Social distance bias has received considerable attention in the experimental literature: e.g. Etang et al. (2011) and Charness and Gneezy (2008). Note that our description incorporates the definition by Hoffman et al. (1996), p. 654, who define social distance “as the degree of reciprocity that subjects belief exist within a social interaction.”

Table 1: Experimental Framework

		Receiver Population		
		Chieftdom	Village	Participants
Revealed	No	ADG-C	ADG-V	ADG-P
Receiver?	Yes	.	.	RDG

Notes: We implement two sets of games: the anonymous dictator game (ADG) and the revealed-receiver dictator game (RDG). The ADG is conducted with three different receiver populations: a randomly selected individual from the chieftdom (ADG-C), the village (ADG-V) or the other $n - 1$ participants (ADG-P).

decide to conduct an experiment in which the dictator is asked to split an allocation of 5 utils towards herself and a “homeless”, and once again to herself and an individual that is not homeless (a “resident”). A difference in contributions is interpreted as discrimination based upon homeless status. First, assume that the dictator, a resident, lives in a neighborhood (the local level in this example) that has two homeless and two residents. In this setting the dictator prefers to contribute 2 utils to the homeless, because a small contribution makes it unnecessary for the homeless to steal. On the other hand, she does not care about a hypothetical homeless at say the national level; in this case she prefers to keep all 5 utils. Assume that in both cases the dictator prefers to keep 5 utils when the receiver is a resident. This simple example illustrates how the simple cue “homeless” might invite the social distance bias if the experimenter is interested in discrimination at the local level. The experimenter might conclude incorrectly the absence of discrimination, while, in fact, at the local level residents discriminate strongly in favor of the homeless (2 utils). In Table 1 this bias is reflected by moving along the top row: comparing results from classic ADG games with different receiver populations.

Statistical discrimination is central to all three types of the ADG (top row of Table 1): while the distribution of the receiver population is different in each game, all games have in common that the dictator only learns about one attribute of the receiver. The two right-most cells in Table 1 have the same population of receivers, but the RDG is not characterized by statistical discrimination. Because dictators play with revealed counterparts that they know intimately the game is no longer characterized by limited information. The dictator knows more about the receivers than only the cleavage under study. Moreover, other factors that are important for her allocation decision—such as previous experience with the receiver and their relationship inside the social network—are all fully known to her. The RDG therefore gives us the measure for local discrimination: the difference, by cleavage under study, in average over all the individual contributions. If only statistical discrimination is at play, averaging over the individual contributions in the RDG should give the same result as the ADG-P. However, it is possible that there is a difference because of, what we will call, the aggregation bias.

Aggregation Bias Given the same population of individuals about whom the dictator knows all the attributes, the aggregation bias refers to the difference between moving from playing against the group of individuals (ADG-P) to the average of playing against each individual separately (RDG). Classic experimental games are played with a hypothetical player from a group of potential receivers. For example, in the previous two games the donation is to somebody from the group of homeless, and to the group of residents. Classic games assume that each individual in each of these populations has an equal probability to receive the dictator’s donation. At the heart of the aggregation bias is the fact that a researcher cannot directly infer behavior towards an individual from behavior towards a group of individuals. Again, consider our thought experiment. Assume the dictator—one of the two residents—plays the revealed-receiver dictator game. She therefore plays three times: one time she has to split 5 utils between herself and homeless #1, a second time between herself and homeless #2, and a third time between herself (resident #1) and resident #2. In this setup the dictator plays one-on-one with receivers that are known to her (they are all from the same neighborhood). Let’s explore her behavior towards the homeless. Assume that homeless #1 has recently mugged the dictator. The contribution in the RDG might therefore be: 0 utils to homeless #1, and 4 utils to homeless #2. If the dictator contributes 0 utils to the resident, the level of local discrimination is 2. What is the dictator’s contribution if the dictator had played an ADG with these two receivers, where randomly either homeless #1 or #2 receives the allocation? The answer is not necessarily 2 utils. For example, because of the mugging the dictator may dislike contributing to homeless #1 to such an extent, that she is willing to place a lot of weight on the 0 utils allocation, even at the cost of homeless #2. That is, she prefers to punish homeless #1 more than to reward homeless #2. In this case, the contribution to the homeless in the ADG will be less than 2 utils. Similarly, a particularly positive relationship with one individual does not necessarily imply a positive relationship towards the group of individuals with the same characteristic.

The aggregation bias therefore states that the inability of the experimenter to take into account the dictator’s intensity of preferences over the receiver population, may lead to an over or under estimation of discrimination at the local level by classic experimental games.⁸

⁸A well developed literature in psychology has found that behavior towards an individual is different from behavior towards a group (e.g. Hamilton and Sherman (1996)). A particularly prominent research agenda is the “identified victim effect”, which find that individuals behave more generous towards identified victims than towards unidentified or ‘statistical’ victims (Jenni and Lowenstein (1997)). Small (2003) and Kogut and Ritov (2005), for example, find that when comparing dictator behavior when a) the receiver will be a randomly chosen from a group of 20 unknown individuals with b) the receiver is a person that is already chosen (but still unknown to the sender) from a group of 20 unknown individuals, contributions are significantly higher in the latter case. Our empirical measurement of the aggregation bias (the difference between ADG-P and RDG) thus also captures those motivations put forward by psychologists that drive differences in play towards a person versus a group: i.e. individual’s expectation of unity, consistency, and coherence that lead to differences in impression formation (Hamilton and Sherman (1996)). Importantly, and in contrast to our games, the experimental games underlying these concepts are conducted among strangers.

In what follows we introduce our assertions more formally, and show how a traditional attribute-based dictator game leads to biased estimates in the measurement of discrimination at the local level.

2.3 Model and Theoretical Predictions

Social distance relates to differences in the *distribution of “types”* in the receiver population. By contrast, aggregation bias relates to the difference between a dictator knowing the distribution of types and knowing *the type* of person with whom she is interacting. We will formalize these notions by referring to “attributes”, where a person’s type is determined by her relevant attributes. In this section we show that simply giving more context in an ADG (moving along the top row of Table 1) induces ambiguous changes in behavior. On the other hand, we show how moving from a known distribution of receivers (ADG-P) to full information about each individual receiver in that population (RDG) yields systematically different behavior. We first show this for donations and then move to discriminatory behavior.

Consider a sender in the dictator game, and let $u(x, \mathbf{a}), x \in [0, 1], \mathbf{a} \in \mathbb{R}^M$ denote the utility of donating a fraction x of the initial endowment to another player with relevant attributes \mathbf{a} . Beliefs about the distribution of \mathbf{a} over the population will be denoted by the smooth function $F(\mathbf{a})$. In order to simplify the analysis, we assume that u is twice differentiable and concave in x . We further assume that u has a unique interior solution, x^* with respect to x , i.e. $\frac{\partial u}{\partial x} = 0$ for $x \in (0, 1)$ for all \mathbf{a} . For the purposes of illustration, assume, without loss of generality, that all relevant attributes can be summarized by a single variable, $\mathbf{a} = a$ with higher values of a denoting more preferred individuals; thus a can also be viewed as a continuous “type.” The analysis focuses on three quantities of interest: the average contributions in the high social distance ADG (e.g. ADG-C) \bar{x}_H^* and low social distance ADG (e.g. ADG-P) \bar{x}_L^* and $E(x^*)$, the average contribution in the RDG.

Our baseline for calculations follows the expected donation in the RDG. In each round, the player is assumed to maximize her personal utility conditional on the relevant attribute(s) of the receiver, which can be described by the optimal donation function:

$$x^*(a) = \arg \max_x u(x, a) \quad (1)$$

The expected contribution is then given by:

$$E(x^*) = \int_{\mathbb{R}} x^*(a) F'(a) da \quad (2)$$

In contrast, in the classic ADG dictators are asked to share their endowment with an anonymous receiver. Participants form a belief about the characteristics of the receiver. In

what follows the dictator assumes that the receiver is a random person, which mimics our experimental design in Sierra Leone (see Section 3).⁹ In the random individual formulation, the sender believes that a randomly selected individual from the population will receive the donation. The expected utility of donating x is then given by:

$$E(u(x)) = \int_{\mathbb{R}} u(x, a) F'(a) da \quad (3)$$

Since the domain of integration is not a function of x , assuming the existence of an interior solution, the optimal donation in the random individual formulation, \bar{x}^* , satisfies:

$$\frac{\partial}{\partial x} \int_{\mathbb{R}} u(\bar{x}^*, a) F'(a) da = \int_{\mathbb{R}} \frac{\partial u(\bar{x}^*, a)}{\partial x} F'(a) da = 0 \quad (4)$$

In order to interpret the solution for \bar{x}^* , let's consider the quadratic Taylor approximation to the utility function for the player as a function of her donation, expanded around $x^*(a)$, which is the optimal donation to a person with attribute value a :¹⁰

$$u(x, a) = h_0(x^*(a)) + h_1(x^*(a))(x - x^*(a)) - h_2(x^*(a))(x - x^*(a))^2 \quad (5)$$

normalized such that:

$$\int_{\mathbb{R}} h_2(x^*(a)) F'(a) da = 1$$

Note further that the expansion around $x^*(a)$ guarantees $h_1(x^*(a)) = \frac{\partial u(x^*(a), a)}{\partial x} = 0$.

$$\int_{\mathbb{R}} \frac{\partial u(\bar{x}^*, a)}{\partial x} F'(a) da = -2 \int_{\mathbb{R}} (\bar{x}^* - x^*(a)) h_2(x^*(a)) F'(a) da = 0 \quad (6)$$

If h_2 is a constant, then it follows that $\bar{x}^* = E(x^*(a))$. The concavity of u in x implies

⁹Another type of belief is the representative person formulation. In the classic game, senders may behave as if they are donating to an individual with a “representative” set of relevant attributes. Let the player select a representative set of attributes through the function g , which is simply a summary measure of the function F . Behavior in an anonymous game under the representative person formulation can be described by the implicit function: $x^*(g(F(a))) = \arg \max_x u(x, g(F(a)))$. For most natural summary measures, $x^*(g(F(a))) \neq E(x^*)$. For instance, if we let g be the expectation operator, then Jensen's inequality implies x^* is concave in a if $x^*(g(F(a))) \leq E(x^*)$ and it is convex in a if $x^*(g(F(a))) \geq E(x^*)$. Under the assumption that the player uses a representative person formulation, and uses summary operators close to the expectation, differences in anonymous games and full information games can be a useful tool to detect the shape of the optimal donation function vis a vis relevant attributes.

¹⁰Note that we may write the constants of the Taylor expansion as function of $x^*(a)$ due to the fact that x^* is a one-to-one and onto function of a . This is not a necessary assumption of the framework, but it helps with the interpretation.

$h_2 > 0$. It can be shown that $\bar{x}^* > E(x^*(a))$ if:

$$\int_{\mathbb{R}} (h_2(x^*(a)) - 1)x^*(a)F'(a)da > 0 \quad (7)$$

Intuitively, h_2 is a weighting function that may put more or less weight on those with higher values of $x^*(a)$. In order to motivate the result, consider a receiver population of two individuals: one with a low value of the attribute, $a_L = 1$, and one with a high value, $a_H = 2$. Both can be chosen with equal probability. Let $u(x, a) = -(x - x^*(a))^{4-a}$, where $x^*(a_L) = \frac{1}{4}$ and $x^*(a_H) = \frac{5}{8}$. Since $x \in [0, 1]$, this implies that the sender has more intense preferences (or weight) corresponding to a_H . Following a standard expected utility maximization, in the random person formulation, the ADG yields $\bar{x}^* = \frac{1}{2}$.¹¹ We see that the donation in the ADG is closer to $x^*(a_H)$ as compared to $x^*(a_L)$, owing to the differences in the intensity of preferences. The intuitive way to interpret “intensity” of preferences is that if the sender has more intense preferences for a high type, then she prefers to reward a high type more than punish a low type in terms of donations.

A simple set of sufficient conditions is that if $h'_2(x^*(a)) > 0$, then $\bar{x}^* > E(x^*(a))$ and if $h'_2(x^*(a)) < 0$, then $\bar{x}^* < E(x^*(a))$. Taken together, using (5) as an approximation of the utility function of the player, we find that higher (lower) donations in the classic game as compared to the revealed-receiver game implies that the dictator places greater (less) weight upon those individuals for whom the optimal donation is higher. This framework thus shows how, given a receiver population, equivalence in outcomes between the ADG-C and the RDG is a knife-edge result, depending on h_2 being a constant.

We will now show that in contrast simply changing the level of social distance in an ADG induces ambiguous changes in donation behavior. Let $F_L(a)$ and $F_H(a)$ be the beliefs about the distribution of attributes in the low social distance (e.g. ADG-G) and high social distance (e.g. ADG-P) cases, respectively. Adapting (6) and (7), it can be shown that the difference between the donations in the high and low social distance is just:

$$\int_{\mathbb{R}} h_2(x^*(a))x^*(a)(F'_H(a) - F'_L(a))da \quad (8)$$

This equation is likely to have an ambiguous sign. The sign of (8) depends on the relative frequency of each type in each social distance condition. However, the low social distance case tends to a central tendency in the data, whereas the high social distance case (low information) tends to permit uniform distributions (flat priors). Thus it is hard to say that low or high distance conditions will yield uniformly different results.

Applying this logic to discriminatory behavior, below we derive two hypotheses related to the aggregation and social desirability bias. To understand discrimination, consider the attribute distribution of high status types, F_H , and low status types, F_L , with associated

¹¹Note that the local level of donations, as measured by the RDG, is equal to $E(x^*) = \frac{7}{16}$.

expectation operators E_H and E_L and ADG contributions \bar{x}_H^* and \bar{x}_L^* , respectively. Discrimination in the ADG between high status and low status receivers is defined as $\bar{x}_H^* - \bar{x}_L^*$. Similarly, the discrimination of high status and low status receivers in the revealed-receiver game is defined as $E_H(x^*(a)) - E_L(x^*(a))$.

A simple extension of (7) shows that the discrimination in the ADG will be higher than in the RDG if:

$$\int_{\mathbb{R}} (h_2(x^*(a)) - 1)(F'_H(a) - F'_L(a))x^*(a)da > 0 \quad (9)$$

Intuitively, discrimination will be higher (lower) in the ADG if the density of desirable attribute receivers (where $h_2 > 1$) is higher (lower) among high status receivers. A simple sufficient condition for discrimination to be higher in the anonymous game is if $h_2 > 1$ implies $F'_H > F'_L$. Intuitively, as long as the density of high status receiver attributes is not significantly larger than the density of low status receiver attributes in regions where the player puts low weight, the discrimination in the anonymous game will be higher than the revealed-receiver game.

H1 Compared to the ADG-P, measured discrimination will be lower in the RDG. This is because we believe that, for players, high status is strongly correlated to desirable attributes for donation in the game.

H2 The social distance bias in discrimination is likely to be of ambiguous sign.

3 Context and Experimental Design

3.1 Context: Rural Sierra Leone and Social Status

To explore differences between the ADG and the RDG in estimating discrimination at the local level we implement our framework in the seven chiefdoms surrounding Sierra Leone's Gola Rainforest National Park. In rural Sierra Leone interaction takes place at the local level. Villages in our research area, as in most of rural Sierra Leone, are small with an average population size of 222.2 individuals divided over an average of 28.4 houses.¹² It is thus not surprising that villagers know each other.¹³ To illustrate this fact we randomly selected 16 individuals in 96 villages and asked each of them how they are related to the other 15 individuals. We find that on average two randomly paired individuals have a full 25% chance to be family. In fact, only 14% of the individuals respond that none of

¹²Based upon data collected in 2010 in 176 randomly selected villages from the seven chiefdoms surrounding Gola Rainforest National Park. These numbers are in line with successive national censuses (1963, 1984, 2004).

¹³This has also been observed in areas characterized by high levels of population movement (Sircar and Van der Windt (2014)).

the other 15 individuals are family. Not only do villagers know each other, in rural Sierra Leone the arena for social interaction is the village and governance takes place at the local level. One important reason for this is that the weak central government apparatus runs in parallel to a powerful local chieftaincy system. The latter has its origins in the system of indirect rule created by the British. The country’s paramount chiefs serve for life once appointed or elected (by a restricted electorate) and exert considerable control over resource allocation, including land and labor; they operate the local court system that presides outside the capital; and organize the provision of many local public goods (Acemoglu et al. (2013)). For example, some of the most important determinants of rural Sierra Leoneans’ households’ well-being such as basic public goods—road maintenance, communal labor, self-help groups, control of crime and school infrastructure—are not provided by the government, but through collective action at the local level (e.g. (Glennerster et al., 2013; Casey et al., 2013)).

The cleavage under investigation in this paper is social status.¹⁴ We choose this cleavage because of the important role it plays in Sierra Leonean society. Specifically, society in rural Sierra Leone is governed by an intricate system of patron-client relationships between high and low status villagers, where villagers depend on a highly exclusionary set of traditional institutions if they want to access property and gain political rights. The system has historically created a large class of excluded, low status individuals (mostly young men). These individuals cannot access political rights by appealing to the modern state, for it is nearly non-existent in rural areas. But they also cannot do so by appealing to traditional authorities if they lack patronage by those higher up in the hierarchy (Fanthorpe (2001)). Many scholars of Sierra Leone argue that the enforced community labor, harsh fines imposed by chiefs, and the lack of opportunities that accompanied this system created feelings of disenfranchisement and resentment among rural youth, which on its turn was a major contributing factor to the outbreak of civil war (Richards, 2005, 1996; Fanthorpe, 2001; Fanthorpe and Maconachie, 2010; Sawyer, 2008). For example, Humphreys and Weinstein (2008) show that the Revolutionary United Front (RUF) found much popular support under predominantly students and farmers protesting against the recurrent marginalization and humiliation by rural elite. There is evidence that these grievances have persisted in the post-war period (Mokuwa et al. (2011)). Studying the presence and extent of social discrimination is thus of key importance in the Sierra Leonean context.¹⁵

¹⁴We are not the first to investigate experimentally the role of social status for cooperation. For example, Baldassarri and Grossman (2013), based upon a set of lab-in-the-field experiments in Uganda, find that people’s social position is an important factor in explaining variation in levels of generosity.

¹⁵A set of experiments conducted by Cilliers et al. (2014) in Sierra Leone, find that the presence of white foreigners has a positive impact on generosity, especially for players who have a lower social status in their community.

3.2 Experimental Design

During spring 2013, we conducted our experimental games with a total of 736 participants from 46 villages in Sierra Leone.¹⁶ The villages in turn were randomly selected from a larger set of villages located close to the Gola Rainforest National Park in southeastern Sierra Leone. These villages have in common that they are remotely located and have seen little exposure to markets. For individuals in these villages the social arena of interaction is local. We stratified the selection of the 16 players by status in order to obtain the same number of high and low status individuals. Specifically, we first selected the eight highest status individuals in the village—we used the Mende term “Taa Gbakoi”, which typically includes the village chief, town speaker, village imam, women’s and youth leader. From the “Nu Gbamei”—literally “person of nothing” in Mende—we randomly drew eight individuals.¹⁷

In private, each participant was interviewed by a research assistant and completed both the experimental games and a short exit survey.¹⁸ The research activities were implemented by multiple teams, each with one instructor and five research assistants. Much effort was undertaken to make sure the participants understood the games. The research assistants were extensively trained and we ran several pilot tests to ensure that our participants understood the experiments. Also, each participants was informed, among other things, that upon completion of the session they would receive a show-up fee of 1,000 Le (or US\$0.25), plus a payment based on a randomly selected allocation decisions by the player. The total set of activities lasted about half an hour per participant, and on average participants received an amount equal to about half a day’s wage.

Following Table 1, each participant partook in the four versions of the game—ADG-C, ADG-V, ADG-P, and RDG. We randomized the order in which each version was played to prevent contamination of the results by learning effects. For each allocation decision, we ask participants to share part of an endowment of 25 tokens, each worth 100 Le. Related to the ADG the dictator made a total of six allocations decision: one for each status type (high or low status) and three types of prime about the receiver, where the receiver is asked to allocate a share of their endowment to a random and anonymous person from 1) a randomly chosen village in their chiefdom, 2) their village, and 3) the sample participating in the experiments. The exact wording that the enumerators used verbatim per version is listed in Table 2. For the RDG, each dictator made a total of $n - 1$ allocation decisions, each time with a different receiver. Specifically, dictators were told the name of the receiver and, in order to distinguish individuals with the same name, also the receiver’s parents’ names.

¹⁶In all villages we randomly selected 16 players.

¹⁷Local elites are generally referred to as “Nu Muwa” in Mende (literally “big person”). However, it not appropriate to ask for these individuals and we therefore choose to use the term “Taa Gbakoi” which is closer in meaning to a town official. Both concepts are very close though, and there is a considerable overlap between both. Based upon discussions with Paul Richards.

¹⁸The enumerators conducted multiple research activities. In this paper we only report the results of the allocation games and individual surveys.

The exact wording can be found in the bottom row of Table 2. This design allows us to explore differences between the ADG and the RDG. Comparing results from the ADG with different receiver populations allows us to investigate the importance of the social distance bias, while comparing the RDG with the ADG with the same receiver population (ADG-P) allows us to learn about the aggregation bias.

Table 2: Experimental Framework

Game	Message to the dictator before allocation decision
ADG-C	“The person receiving the money is a randomly chosen [Taa Gbakoi (e.g. village chief, an imam, a division head, a societal head, a town speaker, etc)] / [Nu Gbamei (e.g. farmer, youth, etc) not a Taa Gbakoi] from a randomly chosen village in your chieftdom, not from your village.”
ADG-V	“The person receiving the money is a randomly chosen [Taa Gbakoi (e.g., village chief, an imam, a division head, a societal head, a town speaker, etc)] / [Nu Gbamei (e.g., farmer, youth, etc) not a Taa Gbakoi] from your village. He or she may or may not be have come with you today from your village.”
ADG-P	“The person receiving the money is a randomly chosen [Taa Gbakoi (e.g. village chief, an imam, a division head, a societal head, a town speaker, etc)] / [Nu Gbamei (e.g. farmer, youth, etc) not a Taa Gbakoi] from the people from your village that came with you today from your village.”
RDG	“The person receiving the money is [Full Name]. His/her father is [Full Name], and his/her mother is [Full Name].”

Notes: The classic dictator game (ADG) is played in three variations with each participant, and the revealed-receiver dictator game (RDG) is played $n - 1$ times with each participant (i.e. each time with a different receiver).

Finally, our research assistants undertook an exit survey with each individual after their allocation decisions. In this survey we collected additional data on participant characteristics (age, gender, education, etc.). As a result, we thus have this data for both the senders and the receivers. The exit survey also included an extensive set of questions related to dyadic characteristics. Specifically, we asked each player about their relatedness to each of the other players along a set of eight dimensions—family, friends, trust, production, social, etc. In contrast to the ADG, the RDG allows an experimenter to investigate the importance of these individual and network characteristics in explaining discriminatory behavior. We will do so in Section 5.1.

4 Results

In this section we first introduce our participants, and then explore discrimination at the local level. That is, the players’ behavior in the RDG. We then compare these results to behavior by the same participants in the ADG-C, the classic attribute-based dictator game. We find large differences across these two games. Leveraging our experimental framework—

i.e. comparing behavior between the three ADG games, and between the ADG-P and the RDG games—we find that this difference is mainly driven by the aggregation bias.

4.1 Descriptive Statistics and Local Discrimination in Sierra Leone

A total of 736 individuals played the classic and revealed-receiver games: half of them are high status, and the other half low status. In Table 3 we compare both types of participants on a set of key characteristics. Not surprisingly, we find that high status individuals are on average almost 16 years older, more likely to be male, and have larger farms.¹⁹ Unexpectedly, but not statistically significant, we find that high status individuals are poorer (as measured by the number of chickens) and are less likely to be literate.

Table 3: Summary Information by Social Status

	High Status (sd)	Low Status (sd)	Difference (se)
Age	50.26 (15.57)	34.70 (11.32)	15.57*** (1.06)
Gender	0.78 (0.41)	0.52 (0.50)	0.26*** (0.04)
Stranger	0.12 (0.33)	0.18 (0.39)	-0.06** (0.03)
Farm Size	2.46 (1.62)	2.18 (1.80)	0.28** (0.13)
Chickens	4.14 (4.22)	4.53 (7.82)	-0.39 (0.65)
Literate	0.19 (0.39)	0.22 (0.41)	-0.02 (0.04)

Notes: In the first two columns, standard deviations are provided in parentheses. Standard errors related to the difference tests are clustered at the village level and reported in parentheses in the last column. Based on data for 736 participants.

Next, we provide an overview of the contributions in the revealed-receiver game by individuals' social status. The top panel in Figure 1 presents a histogram of the amount contributed by high status dictators to high and low status receivers. High (low) status receivers are indicated by black (gray) bars. The RDG was implemented following a round-robin design where each player makes an allocation decision towards each of the revealed $n - 1$ other players. As a result, we have a total of 5,509 contributions by 368 high status

¹⁹Strangers is a widely used concept in Sierra Leonean society and refers to individuals born in another chiefdom often lacking local land rights, except by marriage affiliation. Farm size is measured as the number of bushels of upland rice planted the previous year.

dictators; of which 2,565 to high status individuals and 2,944 to low status individuals.²⁰ The average contribution is 2.91 (11.6% of endowment) to high status individuals, and 2.62 (10.5%) to low status individuals.²¹ We will return to this difference below. The bottom panel presents the same information but now for low status dictators. In total there are 5,511 contributions by low status dictators; of which 2,936 to high status individuals and 2,575 to low status individuals. The average contribution is 2.11 (8.4% of endowment) to high status individuals, and 2.01 (8.0%) to low status individuals.²² For both types of dictators the modal contribution is zero: of the high status individuals 33% gives zero, and of the low status 40%.

Next, we explore whether there is *within dictator* variation in contributions based on the receivers' social status. To do so the left panel in Figure 2 shows whether high status dictators, on average, contribute less, equal or more to other high status receivers compared to low status receivers. The right panel does the same for low status dictators. The height of each bar indicates the number of dictators that give more/less/equal. Furthermore, to illustrate the size of the difference, the width of the bar is the difference between a dictator's average contribution to high status and low status receivers, averaged over all dictators in that bar. From the left panel we see that high status dictators have a preference to give more to fellow high status individuals. Overall, 51% of the high status dictators give, on average, more to fellow high status players than to low status individuals, while only 26% give more to low status receivers. Furthermore, not only do a larger number of high status individuals give more to other high status individuals, the difference in contribution is also larger: those that contribute more to high status individuals contribute 4.12% (1.03/25) of the endowment more, while those that contribute more to low status individuals only contribute 3.44% (0.86/25) more. Interestingly, we find similar results for low status individuals. Low status dictators have a preference to contribute to high status individuals. While 32% of the low status dictators contribute on average more to fellow low status receivers, 41% of low status dictators contribute more, on average, to high status receivers. Moreover, although the difference is small, those that contribute more to high status receivers contribute more on average than those that contribute more to low level receivers: 3.76% (0.94/25) compared to 3.52% (0.88/25).

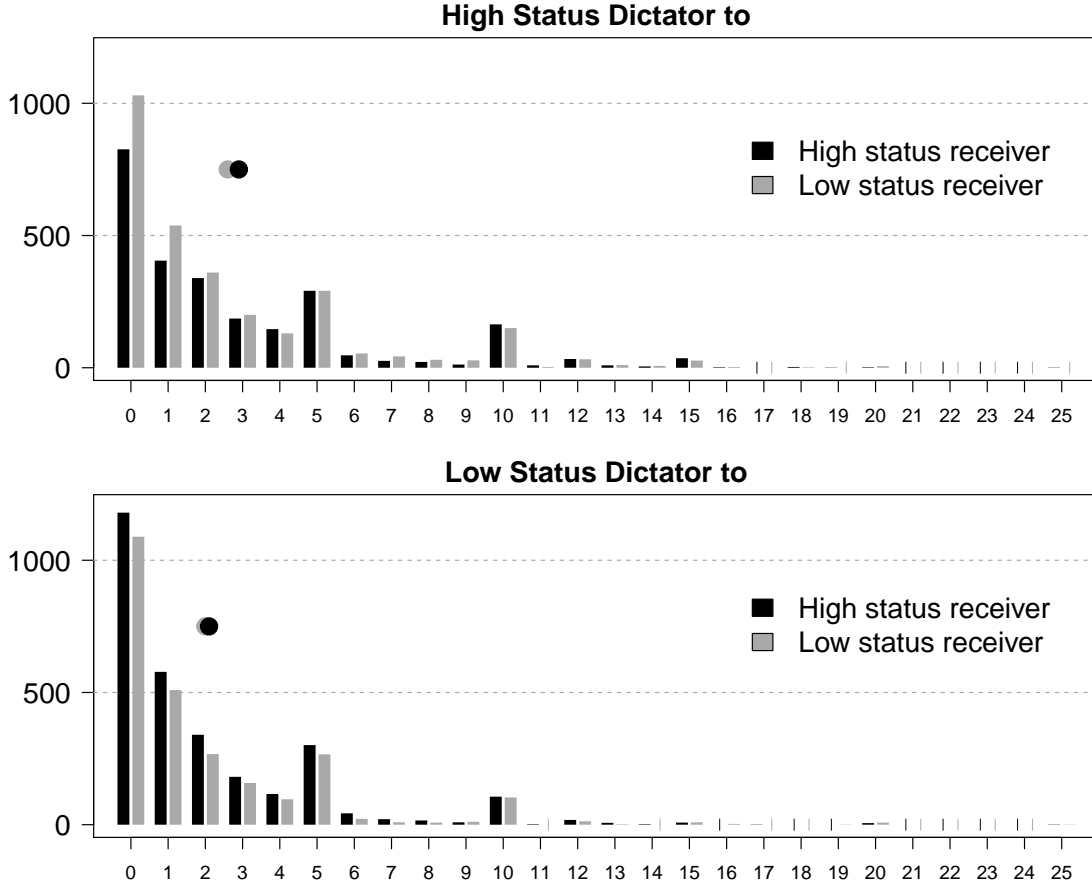
We now move to the measurement of local discrimination in Sierra Leone. Table 4 presents individuals' behavior in the four games, and a comparison between them. The fourth column and the first four rows of Table 4 present average contributions in the RDG separated out by dyad type. For example, the dyad *HL* indicates the contribution by a high status to a low status individual. The standard errors are reported in parentheses and reflect the test whether the average contribution is bounded away from zero. The fourth column and the last two rows of Table 4 show the level of discrimination. We measure

²⁰The number of dyads do not add up to 5,888 (16*368) because individuals do not contribute to themselves. Moreover, in a few cases we miss an observation.

²¹The standard deviations are respectively: 3.56 and 3.39.

²²The standard deviations are respectively: 2.91 and 2.97.

Figure 1: Distribution of Contributions for the RDG, by Dyad




Notes: Contributions are between 0 and 25. Black (gray) bars indicate contribution to high (low) status receivers. Top panel based on 2,565 (2,944) contributions to high (low) status people; bottom panel based on 2,936 (2,575) contributions to high (low) status people. Solid black (gray) point is the average contribution to high (low) status receivers.

discrimination as the difference in contribution between situations where the dictator and receiver are the same status and where they are from different status groups (Fershtman and Gneezy (2001)). In other words—using the notation of this paper—we are interested in the difference between HH and HL , and between LL and LH . Consequently, $HH - HL$ ($LL - LH$) is discrimination by a high (low) status dictator. To account for village main effects, all analyses control for village level fixed effects. Moreover, in order to account for correlation among the $n - 1$ allocations to different receivers by the same dictator and the $n - 1$ donations to the same receiver by different dictators, the standard errors are clustered in two dimensions—by sender and receiver.²³

²³See: Petersen (2008), Thompson (2011), and Cameron et al. (2011).

Figure 2: Distribution of Contributions in RDG Game



Notes: Based on a total of 736 dictators (half high status, half low status). H and low L on the horizontal axis indicate respectively contribution to high and low status receivers. Respectively, the percentages correspond to: 95, 87, 186 dictators (left panel), 121, 99 and 148 dictators (right panel).

Focusing on the fourth column in Table 4 (the RDG game), three results stand out. First, contributions in all four dyads are positive and significantly bounded away from zero. Second, with an average contribution of around 11% of their endowment, high status individuals contribute overall more than those with a low status who contribute on average around 8%. Third, we find evidence of local discrimination based on social status. High status individuals discriminate in favor of members of their own group: contributing 11.3% more to high status than to low status individuals. As we can see from the penultimate row this result is both substantially and statistically significant ($HH - LL = 1.14$, $p\text{-value} < 0.01$). Interestingly, we find a similar result for low status individuals who discriminate also in favor of high status individuals. While the difference is less profound we still find that low status dictators contribute 3.4% more to high status individuals than those of a low status ($LL - LH = -0.27$, $p\text{-value} < 0.05$). These results seem to relate well to more anthropological accounts of rural society in Sierra Leone, which emphasize the importance of social status and patron client relations. We will return to this result in Section 5.1.

4.2 Estimating Discrimination at the Local Level

The principal argument made in this paper is that in contrast to the RDG, the current techniques to measure discrimination (the ADG) are not well-suited for the local level. In this section, we test whether there is indeed a difference between both types of games, and

leverage our experimental design to explore whether this difference can be attributed to the social distance or the aggregation bias (Section 2).

We first investigate behavior as measured by the classic ADG and compare that to the RDG. The first column of Table 4 shows results for the ADG-C: the ADG in which the dictator is told that she is matched with a high/low status individual from her chiefdom. Average contributions in all four dyads are positive and significantly bounded away from zero: contributions are 15.30%, 12.59%, 11.26%, and 13.89% of the endowment for respectively the *HH*, *HL*, *LL* and *LH* dyad. The eighth column tests the difference in contribution between the ADG-C and the RDG directly. We find strong differences in contributions between the two games, with players contributing systematically more in the ADG-C for all four types of dyads. Moving to discrimination, we find strong differences between the RDG and the ADG. The bottom two rows of the eighth column suggest that discrimination is systematically overestimated in the ADG-C. The true level of local discrimination by high status individuals—as measured by the RDG—is 1.14; i.e. high status individuals contribute 1.14% of their endowment more to fellow high status individuals than to those of low social status. An experimenter trying to measure local discrimination with a classic ADG (the ADG-C), however, would conclude incorrectly that local discrimination by high status dictators is a full 2.72% of their endowment. Low status individuals contribute 0.27% more of their endowment to high status individuals in the RDG, which increases to a full 2.63% in the ADG-C. These difference are thus not only statistically but also substantially large, with the ADG-C overestimating levels of discrimination at a factor of 2.4 and 9.7 for respectively high and low status individuals.

The difference we find between the RDG and the classic attribute-based dictator game (ADG-C) can be due to two types of biases: social distance and omitted attribute bias. Our experimental setup in Sierra Leone—illustrated in Table 1—allows us to disentangle both. The social distance bias states that differences between the ADG-C and the RDG are due to the sensitivity of behavior to knowledge about the distribution of receivers (Section 2). Following our design, in Sierra Leone each participant played three anonymous games where we only changed the dictator’s beliefs about the distribution of relevant attributes in the population. That is, we systematically varied the level of social distance from the ADG-C (respondent is from the chiefdom) to ADG-V (respondent is from the village) to ADG-P (respondent is a fellow participant). Differences in game play can then therefore be interpreted as a measure of social distance bias (e.g. Hoffman et al. (1996)). Focusing on Table 4 two results stand out as they relate to the social distance bias. First, in the first four columns we find that by decreasing the level of social distance—that is, moving from the chiefdom prime (ADG-C) to the village prime (ADG-V) to the fellow player prime (ADG-P)—decreases the average dyad contribution in the direction of the RDG. The same holds for the measures of discrimination by the high and low status dictators presented in the bottom two rows. Second, the contributions between the three games, however, do not seem to be discernible different from each other. In fact, moving from the ADG-C to

Table 4: Participant Behavior in the Experimental Games

	Results by Game				Comparison ADGs			Comparison ADG vs RDG		
	ADG-C	ADG-V	ADG-P	RDG	ADG-C ADG-V	ADG-C ADG-P	ADG-V ADG-P	ADG-C RDG	ADG-V RDG	ADG-P RDG
<i>Contribution by dyad</i>										
<i>HH</i>	15.30*** (0.82)	15.17*** (0.45)	14.09*** (0.43)	11.61*** (0.69)	0.13 (0.81)	1.22 (0.75)	1.09 (0.73)	3.63*** (0.72)	3.56*** (0.67)	2.43*** (0.62)
<i>HL</i>	12.59*** (0.71)	12.17*** (0.42)	12.32*** (0.40)	10.54*** (0.67)	0.41 (0.53)	0.3 (0.57)	-0.17 (0.64)	2.03*** (0.62)	1.63** (0.68)	1.78*** (0.58)
<i>LL</i>	11.26*** (0.74)	10.73*** (0.39)	10.38*** (0.38)	8.04*** (0.53)	0.55 (0.63)	0.88 (0.57)	0.33 (0.56)	3.26*** (0.63)	2.71*** (0.56)	2.37*** (0.53)
<i>LH</i>	13.89*** (0.81)	12.55*** (0.44)	12.12*** (0.42)	8.29*** (0.52)	1.34** (0.67)	1.81** (0.75)	0.47 (0.67)	5.56*** (0.69)	4.23*** (0.63)	3.80*** (0.59)
<i>Measure of discrimination</i>										
<i>HH – HL</i>	2.72*** (0.69)	3.00*** (0.74)	1.74*** (0.60)	1.14*** (0.06)	-0.28 (0.96)	0.93 (0.85)	1.27 (0.90)	1.55** (0.67)	1.87*** (0.71)	0.61 (0.57)
<i>LL – LH</i>	-2.63*** (0.66)	-1.82*** (0.61)	-1.75*** (0.60)	-0.27** (0.12)	-0.79 (0.86)	-0.93 (0.84)	-0.14 (0.81)	-2.30*** (0.65)	-1.54** (0.60)	-1.42** (0.60)
N	1,472	1,472	1,472	11,776	1,472	1,472	1,472	11,776	11,776	11,776

Notes: E.g. *HL* indicates the contribution by a high status dictator to a low status receiver. *HH – HL* is discrimination by a high status dictator. Left four columns report average contributions by dyad, and levels of discrimination. Center (right) three columns report differences in play between ADG games (between the ADGs and the RDG). Standard errors reported in parentheses. Contributions are reported as share of endowment. One, two or three asterisks indicate, respectively, 10%, 5% and 1% significance levels.

the ADG-V only the contributions in the *LH* dyad become statistically smaller.²⁴ The change in average contributions when moving from the ADG-V to ADG-P are statistically insignificant for all four dyads. Furthermore, also none of the changes in the measure of discrimination is statistically different across the three anonymous games.²⁵

Next, we investigate the import of the aggregation bias. Given the same population of individuals about whom the dictator knows all the attributes, the aggregation bias refers to the difference between moving from playing against the group of individuals to the average of playing against each individual separately (Section 2). We operationalize this by comparing donations in the RDG with those in the ADG-P. The latter column in Table 4 tests this difference directly. Focusing on the *HH* dyad, for example, the difference between the ADG-P and RDG is large: a full 2.43% of the endowment. That is, average contributions in the ADG-P are on average 21% higher than those in the RDG (p -value<0.01). We find similar results for the other three dyads. Moving to the measurement of discrimination we find that discrimination by high status individuals in the ADG-P is *not* statistically different from that as measured in the RDG. The ADG-P, however, does overestimate levels of discrimination with a factor of 1.5. Discrimination in the RDG by low status individuals is significantly different from the ADG-P. Compared to the RDG, the ADG-P

²⁴On average, contributions in the ADG-P are 1.34% of the endowment lower than in the ADG-C (p -value<0.05).

²⁵Direct tests are provided in the fifth to seventh column in Table 4.

overestimates local discrimination with a factor of 6.5 (-0.27 of the endowment compared to -1.75%, p -value<0.05). A second result is noteworthy. Say that a researcher interested in discrimination at the local level implements an ADG-P. The researcher would incorrectly conclude that the magnitude of discrimination by high and low status individuals is similar: with high (low) status individuals contributing 1.74% (1.75%) more of their endowment to high status players. In fact, discrimination at the local level in Sierra Leone does is not characterized by such symmetry: while high status individuals contribute 1.14% of their endowment more, this is only 0.27% for low status individuals.

To conclude, we find weak evidence that the difference between the ADG-C and the RDG is driven by the social distance bias. In contrast, we find strong evidence in favor of the aggregation bias. That is, the difference between the attribute-dictator game (ADG-C) and the RDG is explained not so much by differences in the receiver population, but, given a receiver population, by the difference between playing with a group of individuals and the average of playing against these participants separately.

5 The Value of the RDG to Understand Discrimination at the Local Level

This paper introduced a novel tool (the RDG) to measure discrimination among populations in which individuals know each other intimately. We also showed how using traditional techniques to understand behavior at this level (the ADG) can provide an experimenter with incorrect estimates. In this final section we illustrate the full use of the RDG when it comes to understanding behavior at the local level. In contrast to the ADG, the receiver is revealed in the RDG, which allows the researcher to explore the role played by receiver characteristics and dyadic relationships. In this section, we will show that these characteristics are important to understand discriminatory behavior at the local level.

5.1 Another Look at Status-Based Discrimination in Sierra Leone

If obtained, the classic ADG is able to explore the role of the dictator’s characteristics and whatever extra receiver information the experimenter provides to the dictator (e.g. the cleavage under study). However, at the local level individuals know much about each other, have had previous experience and are able to position each other in the social network. As a result, receiver and dyadic characteristics are likely to play an important role in shaping discriminatory behavior. A major benefit of the RDG is that we are able to investigate the importance of these factors.

Before moving there, Figure 3 illustrates *within dictator* variation in contributions based on whether the individual plays the ADG-P or RDG. The height of each bar indicate the number of dictators that give more/less/equal across these two games. The width of the


bar indicates the average difference of the contribution by these dictators. Three results stand out. First, we find that for each dyad type around 80% of dictators change their behavior depending on the game. Second, across the four dyads the number of dictators that give less/equal/more is very consistent. Between 40-50% of dictators give on average more in the ADG-P, while only around 30% of the dictators give more in the RDG. Finally, the average difference in contribution is higher for those dictators that give more in the ADG-P ($\text{ADG-P} > \text{RDG}$), compared to those that give less ($\text{ADG-P} < \text{RDG}$). This result is particularly strong for low status dictators. Low status individuals that give less in the ADG-P give only slightly less: on average 1.45% (1.33%) of the endowment when the receiver is high (low) status. In contrast, those that give more in the ADG-P give much more: on average 2.73% (2.53%) when the receiver is high (low) status. We can now add a few additional nuanced points as to why we find differential behavior between the ADG and the RDG (Section 4.2), even when the population of receivers is the same. First, the results are not driven by outliers. Second, the difference in average contribution is driven by the fact that more dictators give more in the ADG-P compared to the RDG. Third, the change in contribution is larger for those dictators that give more in the ADG-P, than those that give less.

In Section 4.2 we found that the level of discrimination in favor of high status individuals is over-estimated in the classic ADG. However, from Figure 3 we know that not all dictators discriminate in favor of high status individuals. Moreover, from Table 3 we know that high and low status individuals differ in many respects other than their social status. We now investigate what other characteristics may explain contribution behavior. To do so we regress sender allocations by high and low status individuals on a range of sender, receiver and dyad characteristics. Table 5 present the results.

Columns ADG-P 1 and RDG 1 report contributions by high and low status dictators for respectively the ADG-P and the RDG. We replicate the result that both high and low status individuals discriminate in favor of high status receivers: the magnitudes are the same as those in Table 4. ADG-P 2 and RDG 2 also take into account a number of dictator characteristics (those presented in Table 3). We find that in the ADG both high and low status men contribute more than women. Also richer low status and illiterate high status individuals contribute more. These results, however, are not corroborated in the RDG. Most importantly, however, we find that the main results related to discriminatory behavior—both the magnitudes and the significance levels—do not change when controlling for dictator characteristics.

A key benefit of the RDG is that it allows the researcher to investigate the importance of factors beyond dictator characteristics. Column RDG 3 in Table 5 adds receiver characteristics. The age of the receiver seems to be a particularly important attribute for contribution, with both high and low status individuals contributing significantly more to older individuals (age is in years). Importantly, we find that not only does the inclusion of receiver characteristic drives down the substantial importance of social status of the

Figure 3: Comparison ADG-P and RDG



Notes: Comparison of individuals' play in ADG-P and RDG. The height of the bar indicates the number of dictators that give more/less/equal. In the ADG-P each dictator plays once against an anonymous high and low status. In the RDG each dictator plays $n - 1$ times. For each dictator, we average contributions by type of receiver. The left two panels add up to 736 dictators, and so do the two right panels. The width of the bar gives an idea of the size of the difference in contribution (contributions are between 0-25).

receiver (the attribute under study), but controlling for receiver characteristics the importance of the receiver's social status is no longer statistically significant for low status dictators.

These result carry over to RDG 4 where we also add dyad characteristics. The age of the sender is still statistically and substantially significant. Moreover, and this should not be surprising, several dyadic characteristics are very important for cooperation. High

Table 5: Exploring the Importance of Receiver and Dyad Characteristics

	ADG-P 1		ADG-P 2		RDG 1		RDG 2		RDG 3		RDG 4	
	H→	L→	H→	L→	H→	L→	H→	L→	H→	L→	H→	L→
HighStatus_R	1.74*	1.75*	1.74*	1.75*	1.14***	0.27**	1.15***	0.27**	0.57**	-0.23	0.61**	-0.17
	(0.92)	(0.91)	(0.91)	(0.91)	(0.06)	(0.12)	(0.08)	(0.13)	(0.27)	(0.28)	(0.28)	(0.27)
Age_S			-0.03	0.02			-0.01	-0.05	-0.01	-0.05	-0.01	-0.06
			(0.03)	(0.05)			(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Gender_S			3.55***	2.15**			1.59	1.52	1.6	1.54	1.42	1.52
			(1.24)	(1.04)			(1.37)	(0.93)	(1.38)	(0.94)	(1.37)	(0.93)
Stranger_S			0.21	-1.06			0.23	-1.87	0.24	-1.88	-0.62	-1.65
			(0.91)	(1.33)			(0.64)	(1.30)	(0.64)	(1.30)	(0.68)	(1.28)
FarmSize_S			-0.04	-0.09			0.56*	-0.10*	0.56*	-0.11*	0.59*	-0.11*
			(0.32)	(0.10)			(0.32)	(0.06)	(0.32)	(0.06)	(0.32)	(0.06)
Chickens_S			-0.06*	0.07**			-0.06	0.06	-0.06	0.06	-0.06	0.06
			(0.03)	(0.03)			(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Literate_S			-0.75**	0.17			-0.25	-0.40*	-0.25	-0.42*	-0.23	-0.42*
			(0.37)	(0.30)			(0.44)	(0.22)	(0.44)	(0.22)	(0.44)	(0.21)
Age_R									0.03***	0.02***	0.03***	0.02***
									(0.01)	(0.01)	(0.01)	(0.01)
Gender_R									0.18	0.34	0.15	0.32
									(0.25)	(0.22)	(0.25)	(0.22)
Stranger_R									0.14	-0.05	0.14	-0.06
									(0.20)	(0.21)	(0.20)	(0.21)
FarmSize_R									-0.01	-0.02	-0.01	-0.02
									(0.02)	(0.02)	(0.03)	(0.02)
Chickens_R									0	0	0	0
									(0.01)	(0.01)	(0.01)	(0.01)
Literate_R									-0.03	-0.16**	-0.05	-0.16**
									(0.06)	(0.08)	(0.07)	(0.08)
Family											0.37	-1.37***
											(0.55)	(0.46)
Farm Share											-0.83	0.1
											(0.87)	(0.53)
Meet											-0.73	-0.01
											(0.76)	(0.61)
Borrow											-0.29	1
											(0.84)	(0.63)
Displaced											1.48*	0.41
											(0.76)	(0.47)
N	735	735	735	735	5,888	5,888	5,888	5,888	5,888	5,888	5,888	5,888

Notes: One, two or three asterisks indicate, respectively, 10%, 5% and 1% significance levels. Regressions control for village fixed effects. Standard errors, clustered by sender and receiver for the RDG regressions, are reported in parentheses.

status dictators contribute more to those individuals with whom they have been displaced. Low status individuals, on the other hand, contribute substantially less to those receivers that are family. The result that stands out, though, is that for low status individuals the social status of the receiver is no longer important for cooperation.

These results illustrate a major benefit of the RDG when it comes to measuring behavior at the local level. At this level we expect individuals to know their partner's attribute such as age. However, age is not, and very likely will never be, an attribute that the researcher would take into account when designing an attribute-based dictator game. That is, an experimenter interested in discrimination by social status is unlikely to tell a random set of dictators "You play with a high/low status individual of age X", and to another set of dictators "You play with a high/low status individual of age Y", etc. However, we find that discrimination in Sierra Leone is not driven by social status per se, but by age. At the local level individuals are very well informed about characteristics such as age. By not taking this attribute into account when measuring discrimination at the local level,

the ADG would have falsely concluded that discrimination in Sierra Leone by low status individuals is driven by social status.

6 Conclusion

Discrimination matters for a wide range of development outcomes including public goods provision (e.g. Alesina et al. (1999)), conflict (e.g. Esteban et al. (2012)) and political decisions (Chandra (2004)).

Recent empirical work that measures discrimination behaviorally have largely relied on audit studies and laboratory experiments. A key characteristic of both approaches is limited information: individuals are often strangers with limited information about each other—often little more than the cleavage under study. The intuitive interpretation of these approaches is therefore that they measure the extent of discrimination between (random) strangers. Such measures thus relate well to situations such as the housing and job market.

This study is motivated by the recognition that many of our interactions take place among individuals that are not strangers. This is particularly the case in developing societies—most farmers sell their crops to middlemen from their village, women obtain the majority of health care from local midwives, villages rely on the local imam to educate their children, and families turn to neighbors for loans, assistance, etc. When interaction takes place at the *local level*, factors such as reciprocity, knowledge of others’ attributes, and the position of individuals within a social network all affect behavior. This study conjectures that classic experimental measures to learn about discrimination may not relate well to a local setting since it fails to account for these factors.

This paper introduces a novel experimental method to measure local level discrimination. Specifically, we move beyond the classic attribute-based dictator game (ADG) and introduce what we call the revealed-receiver attribute-based dictator game (RDG). This game differs from the ADG in two ways: 1) the identity of the receiver is revealed to the dictator, and 2) subjects know each other well.

To explore difference between the ADG and the RDG, we conduct both games with 736 players in a set of rural villages in Sierra Leone. We leverage the importance of an individual’s social status in a village— an important cleavage in Sierra Leone that determines access to resources and reproduction (e.g. Richards (1990)). By having both high and low status players participate in both the ADG and the RDG we are able to test for differences in estimation of discrimination at the local level. Moreover, to explain differences between the ADG and the RDG, we allocate our players to three different versions of the ADG in which we differ the social distance between the dictator and the receiver population: the dictator contributes towards a random individual from the chiefdom (ADG-C), the village (ADG-V), or the other $n - 1$ participants that also play the game (ADG-P). This experimental framework allows us to distinguish whether a difference between the ADG

and the RDG is explained by differences in the receiver population (social distance bias: moving from ADG-C to ADG-P), or, given a receiver population, by the difference between playing with a group of individuals and the average of playing against these participants separately (aggregation bias: ADG-P versus RDG).

We report three main findings. First, our experimental framework illustrates how using traditional techniques to understand behavior at the local level (the ADG) can provide an experimenter with incorrect estimates. Compared to the RDG, we find that the ADG can overestimate levels of discrimination in the local context by up to a factor of almost ten. Second, we show that the difference between both games is mainly driven by aggregation bias, and not social distance bias: what is important for discriminatory behavior is a dictator’s knowledge of the receiver, not her knowledge about the distribution of receivers. Finally, a major benefit of the RDG is that—in contrast to the ADG—we are able to investigate the importance of characteristics other than the attribute under study (in our case social status). In Sierra Leone, we find that receiver and dyad characteristics beyond social status play an important role in explaining discriminatory behavior at the local level.

This study makes two key contributions. First, we introduce a novel tool to measure discriminatory behavior at the local level: among populations that know each other well and have regular interaction. Our results support our interest in incorporating context into lab-in-the-field experiments that aim to measure behavior at the local level. The relative importance of basic individual preferences versus socially determined preferences, and the necessary conditions for individual preferences to swamp or be swamped by social considerations, has received only limited attention in experimental economics. This paper tries to make one step in this direction. Second, currently the literature reports high levels of discrimination in cooperative behaviors—whether based on ethnicity, gender, residential area, etc. In so far as these studies are based upon the classic ADG to measure behavior at the local level, the findings in this paper suggest a less gloomy picture of the world.

References

- Acemoglu, D., Chaves, I. N., Osafo-Kwaako, P., and Robinson, J. A. (2013). Indirect Rule and State Weakness in Africa: Sierra Leone in Comparative Perspective. *Working paper*.
- Adida, C. L., Laitin, D. D., and Valfort, M.-A. (2010). Identifying barriers to Muslim Integration in France. *Proceedings of the National Academy of Sciences*, 107(52):22384–90.
- Ahmed, A. M. (2010). What Is In a Surname? The Role of Ethnicity in Economic Decision Making. *Applied Economics*, 42(21):2715–2723.
- Aigner, D. J. and Cain, G. G. (1977). Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review*, 30(2):175–187.
- Alesina, A., Baqir, R., and Easterly, W. (1999). Public Goods and Ethnic Diversity. *Quarterly Journal of Economics*, 114(4):1243–1284.
- Arrow, K. J. (1973). The Theory of Discrimination. In Ashenfelter, O. and Rees, A., editors, *Discrimination in the Labor Market*, pages 3–33. Princeton University Press, Princeton.
- Baldassarri, D. and Grossman, G. (2013). The Effect of Group Attachment and Social Position on Prosocial Behavior. Evidence from Lab-in-the-Field Experiments. *PloS One*.
- Becker, G. S. (1957). *The Economics of Discrimination*. Chicago University Press, Chicago.
- Bertrand, M. and Mullainathan, S. (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. *American Economic Review*, 94(4):991–1013.
- Binzel, C. and Fehr, D. (2013). Social Distance and Trust: Experimental Evidence from a Slum in Cairo. *Journal of Development Economics*, 103:99–106.
- Bohnet, I. and Frey, B. S. (1999). The Sound of Silence in Prisoner’s Dilemma and Dictator Games. *Journal of Economic Behavior & Organization*, 38(1):43–57.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, 90(1):166–193.
- Burnham, T. C. (2003). Engineering Altruism: A Theoretical and Experimental Investigation of Anonymity and Gift Giving. *Journal of Economic Behavior & Organization*, 50(1):133–144.
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton University Press, Princeton.

- Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust Inference With Multiway Clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.
- Casey, K., Glennerster, R., and Miguel, E. (2013). Reshaping Institutions: Evidence on Aid Impacts using a Preanalysis Plan. *Quarterly Journal of Economics*, 127(4):1755–1812.
- Chandra, K. (2004). *Why Ethnic Parties Succeed*. Cambridge University Press, New York City.
- Charness, G. and Gneezy, U. (2008). What’s in a Name? Anonymity and Social Distance in Dictator and Ultimatum Games. *Journal of Economic Behavior & Organization*, 68(1):29–35.
- Cilliers, J., Dube, O., and Siddiqi, B. (2014). White Man’s Burden? A Field Experiment on Generosity and Foreigner Presence. *Working paper*.
- Croson, R. T. (1996). Information in Ultimatum Games: An Experimental Study. *Journal of Economic Behavior & Organization*, 30(2):197–212.
- Esteban, J., Mayoral, L., and Ray, D. (2012). Ethnicity and Conflict: An Empirical Study. *American Economic Review*, 102(4):1310–1342.
- Etang, A., Fielding, D., and Knowles, S. (2011). Does Trust Extend Beyond the Village? Experimental Trust and Social Distance in Cameroon. *Experimental Economics*, 14:15–35.
- Fafchamps, M. and Lund, S. (2003). Risk-Sharing Networks in Rural Philippines. *Journal of Development Economics*, 71(2):261–287.
- Fanthorpe, R. (2001). Neither Citizen nor Subject? ‘Lumpen’ Agency and the Legacy of Native Administration in Sierra Leone. *African Affairs*, 100(400):363–386.
- Fanthorpe, R. and Maconachie, R. (2010). Beyond the ‘Crisis of Youth’? Mining, Farming, and Civil Society in Post-War Sierra Leone. *African Affairs*, 109(435):251–272.
- Fehr, E. and Schmidt, K. M. (1999). A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics*, 114(3):817–868.
- Fershtman, C. and Gneezy, U. (2001). Discrimination in a Segmented Society: An Experimental Approach. *Quarterly Journal of Economics*, 116(1):351–377.
- Fowler, J. H. and Kam, C. D. (2007). Beyond the Self: Social Identity, Altruism, and Political Participation. *The Journal of Politics*, 69(03):813–827.
- Glennerster, R., Miguel, E., and Rothenberg, A. (2013). Collective Action in Diverse Sierra Leone Communities. *Economic Journal*, 123(568):285–316.

- Habyarimana, J., Humphreys, M., Posner, D. N., and Weinstein, J. M. (2007). Why Does Ethnic Diversity Undermine Public Goods Provision? *American Political Science Review*, 101(4):709–725.
- Hamilton, D. L. and Sherman, S. J. (1996). Perceiving Persons and Groups. *Psychological Review*, 103(2):336–55.
- Hoffman, E., McCabe, K., and Smith, V. L. (1996). Social Distance and Other-Regarding Behavior in Dictator Games. *American Economic Review*, 86(3):653–660.
- Holm, H. J. (2000). Gender-Based Focal Points. *Games and Economic Behavior*, 32(2):292–314.
- Humphreys, M. and Weinstein, J. M. (2008). Who Fights? The Determinants of Participation in Civil War. *American Journal of Political Science*, 52(2):436–455.
- Jenni, K. E. and Lowenstein, G. (1997). Explaining the "Identifiable Victim Effect". *Journal of Risk and Uncertainty*, 14:235–257.
- Kahneman, D., Knetsch, J. L., and Thaler, R. (1986). Fairness as a Constraint on Profit Seeking: Entitlements in the Market. *American Economic Review*, 76(4):728–741.
- Kogut, T. and Ritov, I. (2005). The "Identified Victim" Effect: An Identified Group, or Just a Single Individual? *Journal of Behavioral Decision Making*, 18(3):157–167.
- Leider, S., Mobius, M. M., and Do, Q.-A. (2009). Directed Altruism and Enforced Reciprocity in Social Networks. *Quarterly Journal of Economics*, pages 1815–1851.
- Levitt, S. D. and List, J. A. (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *The Journal of Economic Perspectives*, 21(2):153–174.
- Ligon, E. and Schechter, L. (2012). Motives for Sharing in Social Networks. *Journal of Development Economics*, 99(1):13–26.
- List, J. A. (2004). The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field. *Quarterly Journal of Economics*, 119(1):49–89.
- List, J. A. (2007). On the Interpretation of Giving in Dictator Games. *Journal of Political Economy*, 115(3):482–493.
- Mokuwa, E., Voors, M., Bulte, E., and Richards, P. (2011). Peasant Grievance and Insurgency in Sierra Leone: Judicial Serfdom as a Driver of Conflict. *African Affairs*, 110(440):339–366.
- Neal, D. A. and Johnson, W. R. (1996). The Role of Premarket Factors in Black-White Wage Differences. *Journal of Political Economy*, 104(5):869–895.

- Petersen, M. A. (2008). Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. *Review of Financial Studies*, 22(1):435–480.
- Phelps, E. S. (1972). The Statistical Theory of Racism and Sexism. *American Economic Review*, 62(4):659–661.
- Riach, P. A. and Rich, J. (2002). Field Experiments of Discrimination in the Market Place. *Economic Journal*, 112(483):480–518.
- Richards, P. (1990). Local Strategies for Coping with Hunger: Central Sierra Leone and Northern Nigeria Compared. *African Affairs*, 89(355):265–275.
- Richards, P. (1996). *Fighting for the Rainforest*. Heineman, Oxford.
- Richards, P. (2005). To Fight or to Farm? Agrarian Dimensions of the Mano River Conflicts (Liberia and Sierra Leone). *African Affairs*, 104(417):571–590.
- Sawyer, E. (2008). Remove or Reform? a Case for (Restructuring) Chiefdom Governance in Post-Conflict Sierra Leone. *African Affairs*, 107(428):387–403.
- Sircar, N., Turley, T., Voors, M. J., and Van der Windt, P. (2014). Behavior at the Local Level. *Working paper*.
- Sircar, N. and Van der Windt, P. (2014). Pro-social Behavior in the Context of Rural Migration: Evidence from a Lab-in-the-Field Experiment in the Congo. *Working paper*.
- Small, D. A. (2003). Helping a Victim or Helping the Victim: Altruism and Identifiability. *Journal of Risk and Uncertainty*, 26:5–16.
- Thompson, S. B. (2011). Simple Formulas for Standard Errors that Cluster by Both Firm and Time. *Journal of Financial Economics*, 99(1):1–10.
- Udry, C. (1994). Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria. *Review of Economic Studies*, 61(3):495–526.
- Whitt, S. and Wilson, R. K. (2007). The Dictator Game, Fairness and Ethnicity in Postwar Bosnia. *American Journal of Political Science*, 51(3):655–668.