NBER WORKING PAPER SERIES

DESTRUCTIVE BEHAVIOR, JUDGMENT, AND ECONOMIC DECISION-MAKING UNDER THERMAL STRESS

Ingvild Almås
Maximilian Auffhammer
Tessa Bold
Ian Bolliger
Aluma Dembo
Solomon M. Hsiang
Shuhei Kitamura
Edward Miguel
Robert Pickmans

Working Paper 25785
http://www.nber.org/papers/w25785

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02134
April 2019

We thank seminar participants at the Santa Fe Institute, Stockholm University, University of California, Berkeley, and the East African Social Science Training Annual Summit for useful comments. This project was supported by the Peder Sather Center for Advanced Study and the Swedish Research Council. The authors declare no competing financial interests. A randomized controlled trials registry entry is available at https://www.socialscienceregistry.org/trials/136. This study has complied with all relevant ethical regulations. The University of California, Berkeley Committee for the Protection of Human Subjects and the Kenya Medical Research Institute Scientific and Ethics Review Unit approved the study protocols. Informed consent was obtained from all participants. All authors contributed equally to this work, and are listed in alphabetical order. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w25785.ack

Destructive Behavior, Judgment, and Economic Decision-making under Thermal Stress
Ingvild Almås, Maximilian Auffhammer, Tessa Bold, Ian Bolliger, Aluma Dembo, Solomon M. Hsiang, Shuhei Kitamura, Edward Miguel, and Robert Pickmans
NBER Working Paper No. 25785
April 2019
JEL No. C91,D74,O1,Q50

ABSTRACT

Accumulating evidence indicates that environmental temperature substantially affects economic outcomes and violence, but the reasons for this linkage are not well understood. We systematically evaluate the effect of thermal stress on multiple dimensions of economic decisionmaking, judgment, and destructive behavior with 2,000 participants in Kenya and the US who were randomly assigned to different temperatures in a laboratory. We find that most dimensions of decision-making are unaffected by temperature. However, heat causes individuals to voluntarily destroy other participants' assets, with more pronounced effects during a period of heightened political conflict in Kenya.

Ingvild Almås
Institute for International Economic Studies
Stockholm University
SE-106 91 Stockholm
Sweden
Ingvild.Almas@nhh.no

Maximilian Auffhammer
Agricultural and Resource
Economics Department
University of California, Berkeley
207 Giannini Hall
Berkeley, CA 94720-3310
and NBER
auffhammer@berkeley.edu

Tessa Bold
Stockholm, Sweden
IIES
tessa.bold@iies.su.se

Ian Bolliger
Energy and Resources Group
University California Berkeley
Berkeley, CA 94720
bolliger@berkeley.edu

Aluma Dembo
Department of Economics
Manor Road Building
Oxford Nuffield College
New Road
Oxford, OX1 1NF
United Kingdom
aluma.dembo@economics.ox.ac.uk

Solomon M. Hsiang
Goldman School of Public Policy
University of California, Berkeley
2607 Hearst Avenue
Berkeley, CA 94720-7320
and NBER
shsiang@berkeley.edu

Shuhei Kitamura
Osaka School of International Public Policy
Osaka University
1-31 Machikaneyama, Toyonaka
Osaka, 560-0043
Japan
kitamura@osipp.osaka-u.ac.jp

Edward Miguel
Department of Economics
University of California, Berkeley
530 Evans Hall #3880
Berkeley, CA 94720
and NBER
emiguel@econ.berkeley.edu

Robert Pickmans
Agricultural and Resource
Economics Department
University of California, Berkeley
207 Giannini Hall
Berkeley, CA 94720-3310
rpickmans@berkeley.edu

Randomized controlled trials registry entray available at https://www.socialscienceregistry.org/trials/1361

Introduction

Scholars have long attempted to understand the role that climate plays in human societies (de Secondat and De Montesquieu, 1748; Huntington, 1922; Lee, 1962), with a recent acceleration of findings enabled by improved data, computing, and statistical techniques (Carleton and Hsiang, 2016). High temperatures are found to influence important societal outcomes, reducing economic productivity (Burke et al., 2015), international trade (Jones and Olken, 2010), hours worked (Graff Zivin and Neidell, 2014), physical activity (Obradovich and Fowler, 2017), law enforcement effort (Obradovich et al., 2018), and fertility (Barreca et al., 2015), while increasing domestic migration (Bohra-Mishra et al., 2014), energy use (Wenz et al., 2017), political turnover (Obradovich, 2017), suicides (Carleton, 2017; Baysan et al., 2018), HIV transmission (Burke et al., 2015), interpersonal violence (Ranson, 2014), and intergroup conflict (Hsiang et al., 2013).

A key remaining challenge is identifying the mechanisms through which high temperatures affect this array of human outcomes (Dell et al., 2014). Two hypotheses are typically proposed, at least informally. First, it is possible that factors *external* to human beings respond to high temperature, which then directly or indirectly affect outcomes. For example, climate-induced loss of agricultural income (Schlenker and Roberts, 2009) has been proposed as a mechanism linking extreme climatic events with decisions to engage in political conflict (Miguel et al., 2004), transactional sex (Burke et al., 2015), or suicide (Carleton, 2017). Alternatively, extreme temperatures could directly influence the *internal mental processes* governing decision-making. In this vein, the medical and psychological literature has posited that heat increases aggression because of its negative effect on emotional affect (Bell and Baron, 1976), because lowered affect is misattributed to anger (Zillmann, 2006), or because of a physiological effect of heat that triggers aggression. In the economics literature, it has been suggested that violent crime and group conflicts become more likely at high temperatures because individual decision-making is altered—through diminished consideration for future consequences, lowered risk-aversion, or heightened risk of strategic miscalculation (Anderson, 1989; Baysan et al., 2018). Determining the relative influence of external factors and internal mental processes is crucial to developing policies or technologies to address the damages caused by high temperatures.

Prior analyses established that cognitive performance, such as deciphering Morse code (Mackworth, 1946) and solving math problems (Graff Zivin et al., 2018), declines at high temperature, and suggest that "aggressive" behaviors—such as sports violence (Larrick et al., 2011), horn-honking (Kenrick and MacFarlane, 1986), swearing on social media (Baylis et al., 2018) and aggression in the lab (Anderson, 1989)—may increase. These prior analyses were narrow in scope, did not always rely on experimental variation, and could not isolate mechanisms.

We conduct the first systematic laboratory evaluation of high temperature effects on the internal mental processes that govern economic decision-making, judgment and destructive behavior. We randomly assign subjects to Hot or Control environments in two parallel and simultaneous experiments conducted in Berkeley, California and in Nairobi, Kenya. We deploy a battery of assessments designed to measure: components of individual cognition, including *precision* (Gill and Prowse, 2012), *fluid intelligence* (Penrose and Raven, 1936), and *reflection* (Frederick, 2005); private preferences and decision-making quality, such as *risk taking* and *rationality* (Eckel and Grossman, 2008), and *patience* and *consistency of time preferences* (Andreoni et al., 2015); social preferences and destructive behavior, including *fairness* (Cherry et al., 2005; Almås et al., 2010), *trust* (Johnson and Mislin, 2011), *public goods provision* (Fischbacher et al., 2001), *joy of destruction* (Abbink and Sadrieh, 2009), *charitableness* (Andreoni, 2006); and individual affect, namely, *happiness* and *alertness* (Russell and Barrett, 1999).

The experiments were run from September 2017 to February 2018. Notably, a controversial national election in Kenya occurred just prior to our experiment and was followed by months of heightened tensions between ethnic Luo (20.8% of Nairobi subjects) and Kikuyu (20.6%) communities, political turmoil, and violent protests (The Economist, 2017) that caused some experimental sessions to be rescheduled (Supplementary Fig. D.8.2). This unplanned situation created a social context that seems to have influenced our findings in unanticipated but important ways, discussed below.

Research Design

The experiment was conducted in Berkeley, California, with $N = 903$ participants (Fig. 1B) and in Nairobi, Kenya, with $N = 1015$ participants (Fig. 1D). The experimental treatment consists of temperature manipulation: in both California and Nairobi, the target temperature in the control room is set to 22°C, while the target temperature in the treatment room is set to 30°C (Fig. 1A,C display actual temperatures). The experiment consists of experimental modules developed and validated in previous studies (Fig. 1E). Participants are incentivized and play over tokens, where the conversion rate is 3.5 tokens to 1 Kenyan shilling in Nairobi and 2.1 tokens to 1 cent in California.[1] Participants are drawn from the University of California, Berkeley for the California sample and the University of Nairobi for the Nairobi sample.[2] Although one should bear in mind that university students are not necessarily representative of the population as a whole (Sears, 1986), the use of student subjects should not pose a problem for the study's external validity per se (Druckman and Kam, 2011).

For each session, we aim to recruit 12 participants who are randomly assigned to either the control or treatment group (a session-condition).[3] At both sites, participants are randomly assigned to the treatment or the control room. To ensure sufficient exposure to treatment for all modules, participants at each location waited 20 minutes in the temperature-controlled laboratory rooms before beginning the experimental modules. During this time, they sat quietly at their seats and received instructions about the experiment. A water bottle, a notepad for writing, and a pen was placed at each participant's desk. Photos of the experimental set up are included in Supplementary Online Materials (SOM) Section D. The entire lab lasted for roughly one hour.

[1] For context, 1 US dollar is equal to about 100 Kenyan shillings. Note that the exchange rate was previously 1.75 tokens to 1 cent in California, but is changed early in the experimental timeline on September 28th, 2017, both to keep at the original target payment of $25 an hour and to sustain recruitment. After all experiments are completed, average earnings in California were around $35, while average earnings in Nairobi were around 2000 Kenyan shillings.

[2] Due to the University of Nairobi temporarily closing due to political events, students were also recruited from Strathmore University for a period of time. Please see Supplementary Online Materials Section D.8 for more discussion.

[3] In practice, while we over-recruit at 50% to better ensure having 12 participants attend, we were not always able to secure 12 participants. We ran sessions in which only 11 participants attended, but we canceled sessions in which fewer showed up.

The experimental modules are by and large standard and well-known in the literature and allow us to measure the effect of temperature on a broad range of preferences, economic decisions and pro-social behavior. In the analyses below, we focus on a set of pre-specified outcomes that we term as "primary outcomes" of interest. We also note more "exploratory outcomes" of interest.

The *Risk* module elicits risk preferences using choices over lotteries (Eckel & Grossman, 2008) and presents two different menus of choices. The *Time preferences* module consists of the standard protocol for eliciting so-called "beta-delta" preferences, namely a convex time budget (CTB) design following Andreoni et al. (2015). By detecting violations of the Generalized Axiom of Revealed Preference, both modules can identify the effect of temperature on the quality of decision-making. In addition, the *Risk* module can also reveal violations of first-order stochastic dominance. The primary outcomes of interest from the Time preferences module include aggregate estimates of beta (time inconsistency) and delta (discounting). Exploratory outcomes of interest include an indicator for the Generalized Axiom of Revealed Preference violation, as well as individual rankings of choices from each menu. The primary outcomes of interest from the Risk module include the participant-chosen variance in outcome from the virtual coin toss featured in the first menu (where choosing a higher variance outcome is defined as being more risk-loving) as well as an indicator for a transitivity violation using both menus. An exploratory outcome of interest is an indicator for first-order stochastic dominance violation.

The *Fairness* module uses a real effort dictator game where participants first earn tokens from a Precision task (see Gill and Prowse (2012) and a short description below). Such earned income generates clear entitlements for the dictator game (Kahneman et al., 1986; Forsythe et al., 1994), where it is to be distributed among a pair of agents with identical earnings.[4] All participants act as dictators and know that their decision is implemented with fifty percent probability (Almås et al., 2010). The primary outcome of interest is the share of the total initial amount of tokens that is allocated to the other participant.

The *Public contribution* module consists of a public goods game with three players, which

[4]Hence the fair outcome given by entitlements is always an equal split.

encourages cooperation by doubling all contributions to a common pool (Fischbacher et al., 2001). In the *Trust* module, one player in a pair is given an initial amount of tokens that can be sent to the other player. As in the public goods game, pro-social behavior is encouraged through a multiplier: any transfer sent is multiplied by three. The other player then decides to return any share of the amount received back to the first mover (including zero) (see Johnson and Mislin (2011) for a similar design). The primary outcome of interest from the Public contribution module is cooperation, defined as the amount of tokens placed into the group fund, while an exploratory outcome of interest is having correct beliefs (which is defined as an indicator variable for whether the individual guesses correctly about another's contribution). The primary outcome of interest from the Trust module is trust, defined as the share of the total initial amount that is sent to the other participant. An exploratory outcome of interest from the Trust module is trustworthiness, defined as the share of tokens sent back to the first participant. We also use the value coming from a 0 to 10 scale of trust in others' intentions as an exploratory outcome.

The *Joy-of-destruction* (JOD) game (Abbink and Sadrieh, 2009) has been studied somewhat less. In this module, participants are informed that everyone has earned different amounts of (up to six) $1 Amazon gift cards in California or Airtime vouchers worth 50 Ksh each in Nairobi from a fluid intelligence task (described below). They are then anonymously matched in pairs and are told how much their partner has earned (X). Each participant can destroy any number of their assigned partner's vouchers between 0 and X. The computer can also destroy some of the remaining vouchers: it destroys nothing with probability 0.5, and any number of the remaining cards with equal probability. The other participant does not know whether her earnings are destroyed because of the computer or the participant's decision. Thus, a participant's purposeful destruction is partly concealed by the computer's random destruction. The game is conducted with gift cards (vouchers), so that actual destruction (and not reallocation) can take place. The primary outcome of interest is a measure of destruction, defined as the proportion of gift cards or vouchers destroyed.

In addition to the incentivized modules mentioned, we measure affect by including two questions: "On a scale from 1-7, with 1 being tired and 7 being alert, how do you feel right now?"

and “On a scale from 1-7, with 1 being sad and 7 being happy, how do you feel right now?”. We also use a slider task to measure what we call precision (Gill and Prowse, 2012), and measure fluid intelligence using Raven’s matrices. The primary outcome of interest from the Precision task is the number of sliders correctly placed within three minutes, and the primary outcome in the Fluid intelligence task is the proportion of puzzles solved correctly. As part of the non-incentivized survey, we use a cognitive reflection test (CRT) consisting of five survey questions (Frederick, 2005). The primary outcome of interest from the CRT is the share of survey questions answered correctly. The secondary outcome of interest is the time spent answering the questions. Finally, we include a demographic survey with questions on trust, body weight, and parental education and income.

Towards the end of the experimental session, after the participants are informed about their earnings, they can choose to donate some of their payments to a charity organization. The primary outcome from the Charitable donation module is the absolute amount they choose to donate.

The experimental design, protocol and core analyses were declared in advance and pre-registered on the American Economic Association RCT registry (available at https://www.socialscienceregistry.org/trials/1361). All results from the pre-registered analysis are presented in the main text or the SOM. Additional *post hoc* analyses are explicitly identified below. All statistical tests in this analysis are two-sided. For details about the design of each module as well as the exact timeline for the experiment and political setting at each site and the pre-analysis plan, please see the SOM.

Subjects were never deceived, and in order to incentivize truthfulness and effort, their performance and choices during the experiment determine their monetary earnings. In the social choice modules, individual rewards depend on interactions with other participants in the same room, anonymized and mediated through a computer terminal. The ordering of modules was fixed across treatments and sites (Fig. 1A,C), although subjects in California generally completed tasks somewhat faster.

Nairobi and California were intentionally selected as experimental sites because the underlying population of university student subjects are distinct and might indicate whether or not the response to thermal stress differs across populations (Supplementary Table F.3.1). Subjects in California are on average slightly younger (20.0 ± 2.7 yrs vs. 22.0 ± 2.4 yrs in Nairobi), more likely to be female (60% vs. 35%), and more likely to have at least one parent with university education (78% vs. 25%) and higher income levels (89% vs. 8% earn $> \$24,000$ US, PPP adjusted). Crucially, *within* each site subjects are randomly assigned to Hot or Control conditions, such that within-site treatment effect estimates are not confounded by differences in individual-level characteristics. Subject populations in Hot conditions are statistically indistinguishable from the Control population in California and differ in Nairobi only by being slightly heavier on average (63.0 ± 9.3 kg vs 61.3 ± 7.8 kg; see Supplementary Table F.3.1).

We use standard regression methods to analyze the outcomes.[5] Specifically, we employ the following specification:

$$Y_i^k = \beta_0^k + \beta_1^k T_i + \beta_2^k Male_i + \beta_3^k Nairobi_i + \varepsilon_i^k \tag{1}$$

where i refers to individual, s refers to site, k refers to the module outcome (e.g., trust), Y_i refers to an outcome of interest, T_i is the treatment indicator for Heat, $Male$ is an indicator for the individual's self-reported gender being male, $Nairobi$ is a fixed effect for the experimental location being Nairobi, and ε_i is an idiosyncratic error term.[6] Treatment effects in Fig. 2 are presented in terms of standard deviation units, normalized relative to participants in the Control condition. We cluster standard errors at the session level. Analysis was also carried out separately on the California and Nairobi subsamples.

To generate more meaningful inference, given the large number of hypotheses we consider, we

[5] An exception is the set of time preference outcomes, for which we employ a non-linear least squares procedure following Andreoni et al. (2015).

[6] Our main specification in our pre-analysis plan includes X_i, which represents a vector of covariates that are unbalanced between treatment and control groups. In our study, only self-reported weight is unbalanced within Nairobi. Our results are robust to the inclusion of this covariate in our regressions (see Online Appendix Section B for these results); X_i has not been included in the regressions generating results below, due to a non-trivial amount of missing data for the weight variable. The regression model for the Charitable donation outcome is slightly different, namely, it includes an additional indicator variable that takes on a value of 1 if the individual donates to an organization based locally (for the California sample) or based in their home region (for the Nairobi sample), as well as a control variable for total lab earnings. Several other regressions involving exploratory outcomes also feature additional covariates, depending on the outcome, please refer to Online Appendix Tables for details.

carry out multiple hypothesis testing adjustments, a practice that is already carried out regularly in other fields such as genomics (Okbay et al., 2016; Hirschhorn and Daly, 2005) and increasingly in economics. In particular, we calculate false discovery rate adjusted "q-values" that limit the expected proportion of rejections within a set of hypotheses that are Type I errors, following Anderson (2008). The use of q-values, which are interpreted analogously to p-values, was also specified in the pre-analysis plan. The use of both a pre-analysis plan that describes the primary hypothesis tests as well as multiple hypothesis testing adjustments should increase confidence in the results.

Results

We find that nearly all components of economic decision-making and judgment are largely unchanged by thermal stress (Fig. 2), at the level of thermal exposure in our experiment. Further, we cannot reject that subjects in Nairobi and California respond similarly to temperature across nearly all outcomes. Specifically, we find no effect of high temperature on our measures of individual cognition, private preferences, decision-making quality, or pro-sociality. Across each of these outcomes, a null result is found in both the Nairobi (Fig. 2 in dark orange) and California (light green) samples, and pooled (black). In the precision task and fluid intelligence measures, individual estimates exhibit p-values < 0.05, but these effects are not statistically significant once we account for a pre-specified multiple hypothesis testing adjustment (q-values).

In terms of magnitudes, the effect sizes themselves are also negligible. These are precisely estimated zeros. In other words, failure to detect significant impacts is not a consequence of small sample size; on the contrary, our study was powered at 80% to detect even modest impacts (of at least 0.2 of a standard deviation). Overall, we conclude that many fundamental elements of economic decision-making and judgment are broadly unaffected by thermal stress despite the large observed declines in self-reported alertness (all samples) and happiness (in Nairobi) (Fig. 22, bottom panel).

There is a statistically significant effect of thermal stress on destructive behavior in Nairobi, and this effect survives the multiple testing adjustment. Specifically, individuals choose to destroy more of the winnings of other players in the JOD module (Abbink and Sadrieh, 2009) (q-value = 0.001). In Nairobi, subjects destroyed an additional 8 percentage points (0.30 s.d.) of other players' earnings in the Hot condition relative to Control – equivalent to an increase of almost 50%. In contrast, there is no increase in destructive behavior under thermal stress in the California sample. In fact, destruction falls slightly among California subjects in the Hot treatment, by roughly half the magnitude of the Nairobi increase. We note that the average share of vouchers destroyed in California (6%, Fig. 3D) is far lower than in Nairobi (18%, Fig 3A).[7]

The strong aggressive response to heat in the Nairobi sample, and the absence of a similar response in the California sample, prompted us to investigate relevant mechanisms, including through analyses that were not pre-specified (see SOM) but that are guided by existing theoretical models of aggression. In particular, psychologists model aggression by means of a general affective aggressive framework, consisting of four stages/mechanisms (Anderson et al., 2000): (i) initial inputs such as personal characteristics and situational variables, e.g., high temperature or aggressive cues; (ii) the impact of these initial inputs on a person's internal state, namely, emotional and physiological arousal, as well as cognition; (iii) appraisal of the situation either automatically or in a more controlled fashion through the lens of the internal state and choice of an appropriate response; and possibly (iv) manifestation of aggressive behavior.

The rich experimental and survey data we collect allows us to examine the extent to which some of these mechanisms operated in our lab setting. First, we find a significant negative impact of high temperature on both happiness and alertness in Nairobi, and on the latter in California. In this regard, it is also interesting to reiterate that we find no direct effect of heat on cognitive

[7] As noted in Bauer et al. (2018), previous work using the joy of destruction game has documented destruction rates among university students ranging from 8 to 40% in the Netherlands (Abbink and Sadrieh, 2009) 10 to 26% in Ukraine (Abbink and Herrmann, 2011). Meanwhile, destruction rates ranged from 23 to 40% among pastoralists in Namibia, with higher levels of destruction among the resource-scarce subsample. Bauer et al. (2018) find destruction rates to be higher among adolescents from rural Uganda (53 to 59%) compared to adolescents from disadvantaged regions in Slovakia (32 to 42%). The differences in destruction rates across countries in this experiment is thus consistent with the available evidence.

performance. This suggests that, at least in our sample, heat may trigger aggression through its effect on affect and arousal rather than cognition.

Despite the absence of a direct effect on cognition, we find an interesting interaction between measures of cognitive ability and the aggressive response to the Heat condition: individuals in Nairobi with below median cognitive scores had a significantly larger ($p = 0.020$) destructive response to heat (12%) than those with above median cognitive scores (4%, Fig. 3C). This suggests that cognitive ability may be an important mediator in the heat-aggression relationship, perhaps because it affects the extent to which "aggression-related thoughts, feelings and behaviour scripts" (Anderson et al., 2000) are accessed or how the situation is appraised. Parental university education (a correlate of socio-economic status), on the other hand, was unrelated to an individuals destructive response to thermal stress in both sites (Fig. 3B,E).[8]

Excitation transfer theory posits that arousal caused by an initial stimulus (here, heat) persists and "can be transferred and added to the arousal produced by a second stimulus" (Eysenick, 2004), i.e., a provocation. Individuals may then (mis-)attribute their arousal exclusively to the second stimulus and respond excessively. More generally, excitation transfer theory suggests that negative arousal from several stimuli may combine and amplify responses, even if the underlying stimuli are different in nature (i.e., heat and provocation/frustration). The Kenyan sample provides a fortuitous testing ground for this conjecture. Kenyan national elections have been controversial in recent decades due to election fraud accusations and resulting ethnic violence (Mebane, 2017). Our experiment occurred during a period of heightened instability immediately following the 8 August 2017 presidential election between Raila Odinga (a member of the Luo ethnic group) and the incumbent Uhuru Kenyatta (a member of the Kikuyu ethnic group; Fig. 5). Kenyatta won the general election, followed by an annulment of the results by the Supreme Court due to election discrepancies (September 1). Kenyatta won a second round of voting that was boycotted by the

[8] A somewhat higher proportion of California lab participants correctly guessed that the purpose of the experiment was to assess the effect of thermal stress, compared to Kenya participants (see SOM Section F.3). However, there are no statistically significant differences in the effect of thermal stress on destructive behavior between those who did and did not figure out the purpose of the experiment (among the subset of our participants ($n = 913$) who received a debriefing questionnaire, see Online Appendix Table E.10.2).

opposition (October 26) and generated accusations of voter intimidation and fraud. Violent clashes resulted, including in Nairobi, with many victims reported to be ethnic Luos supporting Odinga (The Economist, 2017; Kiruga, 2018). This political context was salient during our study, as many experimental sessions in October were postponed because violence caused university closures (see SOM Sections B.1 and C.1 for details). No similar disruption occurred in Berkeley, California during the study.

Specifically, we hypothesize that the stronger aggressive response to heat seen in Nairobi is driven by the ongoing political marginalization and associated frustration of non-Kikuyu ethnic communities. Such a pattern would be consistent with the long-existing frustration-aggression literature, through which frustration leads to more aggressive behavior through negative affect (Berkowitz, 1989), and would also be related to a recent finding that destruction increases with scarcity (Prediger et al., 2014). Consistent with this, we find that Nairobi subjects that identify with an ethnicity other than Kikuyu (the ethnic group of the president re-elected just prior to the experiment) have a much larger increase in destructive behavior under thermal stress (Fig. 4B). Ethnic Kikuyu subjects exhibited no destruction response (Control = 14%, Hot = 12%). The largest response was visible in ethnic Luo subjects, where heat caused destructive behavior to more than double (12% vs. 27%), while other ethnic groups exhibited a response of intermediate magnitude. The difference between Luo and Kikuyu responses is statistically significant (p-value $= 0.003$). Among members of other ethnic groups, those predominantly politically aligned with the opposition show a response to high temperature that is somewhat larger than that found among members of largely pro-government ethnic groups, mirroring the pattern among ethnic Luo and Kikuyu participants (see Fig. 4C, and SOM Section F.7).

Importantly, our results do not suggest that there is anything fundamentally different in the judgment or economic decision-making processes of Nairobi non-Kikuyu participants. The results for all other modules suggest that these subjects exhibit patterns of decision-making that are otherwise statistically indistinguishable from other subjects (see Online Appendix Section E.7); moreover, average rates of destruction in the Control condition are nearly identical across ethnic groups in Nairobi. We hypothesize that if Kikuyu or California subjects were politically marginalized similarly to the recent experience of ethnic Luos in Kenya, they would exhibit a JOD-temperature response similar to Luo participants.

Discussion

The finding that a politically disenfranchised population exhibits an increased preference for harming others under thermal stress might partially explain previous findings that higher temperatures elevate rates of interpersonal violence and social conflict around the world (Hsiang et al., 2013). However, the possible role of an internal mental channel does not preclude a simultaneous influence of external factors, such as altered economic circumstances (Miguel et al., 2004), in generating human conflict via thermal stress.

Overall, our results indicate that, with the exception of destructive behavior, many basic elements of economic choices and decision-making quality are largely unchanged by thermal stress across populations on two continents facing different baseline climates. This suggests that some population-level effects of high temperatures, such as reduced economic productivity (Burke et al., 2015) and labor supply (Graff Zivin and Neidell, 2014), may be predominantly caused by changes to external factors, such as crop physiology (Schlenker and Roberts, 2009) or other non-cognitive components of the economy (Carleton and Hsiang, 2016).

It is possible that different results may be achieved if subjects are exposed over longer a duration to temperatures higher than 30°C. The conditions of this study were constrained by ethical considerations. Further, the collection of biomarker data (Haushofer and Fehr, 2014) could shed light on physiological changes. It is also possible that other non-temperature stressors could similarly lead to increases in destructive and anti-social behavior among socially marginalized groups, and this remains a promising direction for future research.

Our ability to mitigate adverse human impacts of climate may depend on the underlying channels of impact. Traditional policy responses and technologies can effectively manage several external factors, but appear less suited to manage internal mental processes. Our findings on destruction suggest that interventions that dampen the economic consequences of thermal stress—such as expanded rainfall insurance—might not break the link between temperatures and human conflict if it is indeed mediated by internal mental processes. Especially in high-risk settings, other innovations that instead directly disincentivize destructive behavior or insulate individuals from high heat exposure, for instance, through air conditioning (Barreca et al., 2016), may be necessary.

References

Abbink, K. and B. Herrmann (2011). The moral costs of nastiness. *Economic Inquiry 49 (2)*, 631–633.

Abbink, K. and A. Sadrieh (2009). The pleasure of being nasty. *Economics letters 105 (3)*(3), 306–308.

Almås, I., A. W. Cappelen, E. Ø. Sørensen, and B. Tungodden (2010). Fairness and the development of inequality acceptance. *Science 328 (5982)*, 1176–1178.

Anderson, C. A. (1989). Temperature and aggression: Ubiquitous effects of heat on occurrence of human violence. *Psychological bulletin 106 (1)*(1), 74.

Anderson, C. A., K. B. Anderson, N. Dorr, K. M. DeNeve, and M. Flanagan (2000). Temperature and aggression. *Advances in Experimental Social Psychology 32*, 63–133.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American Statistical Association 103 (484)*, 1481–1495.

Andreoni, J. (2006). Philanthropy. In S.-C. Kolm and J. Mercier Ythier (Eds.), *Handbook of the economics of giving, altruism, and reciprocity*, Volume 2, Chapter 18, pp. 1201–1269. North-Holland.

Andreoni, J., M. A. Kuhn, and C. Sprenger (2015). Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior & Organization 116*, 451–464.

Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro (2016). Adapting to climate change: the remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy 124*(1), 105–159.

Barreca, A., O. Deschenes, and M. Guldi (2015). Maybe next month? temperature shocks, climate change, and dynamic adjustments in birth rates. Technical Report 21681, National Bureau of Economic Research.

Bauer, M., J. Cahlíková, D. Celik Katreniak, J. Chytilová, L. Cingl, and T. Želinský (2018). Anti-social behavior in groups. Technical Report 11944, IZA Institute of Labor Economics.

Baylis, P., N. Obradovich, Y. Kryvasheyeu, H. Chen, L. Coviello, E. Moro, M. Cebrian, and J. H. Fowler (2018). Weather impacts expressed sentiment. *PloS one 13*(4), e0195750.

Baysan, C., M. Burke, F. Gonzalez, S. Hsiang, and E. Miguel (2018). Economic and non-economic factors in violence: Evidence from organized crime, suicides and climate in mexico. Technical Report 24897, National Bureau of Economic Research.

Bell, P. A. and R. A. Baron (1976). Aggression and heat: The mediating role of negative affect. *Journal of Applied Social Psychology 6*, 18–30.

Berkowitz, L. (1989). Frustration–aggression hypothesis: Examination and reformulation. *Psychological Bulletin 106 (1)*(1), 59–73.

Bohra-Mishra, P., M. Oppenheimer, and S. M. Hsiang (2014). Nonlinear permanent migration response to climatic variations but minimal response to disasters. *Proceedings of the National Academy of Sciences 111*(27), 9780–9785.

Burke, M., E. Gong, and K. Jones (2015). Income shocks and hiv in africa. *The Economic Journal 125*(585), 1157–1189.

Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature 527*, 235–239.

Carleton, T. A. (2017). Crop-damaging temperatures increase suicide rates in india. *Proceedings of the National Academy of Sciences*, 201701354.

Carleton, T. A. and S. M. Hsiang (2016). Social and economic impacts of climate. *Science 353*(6304), aad9837.

Cherry, T. L., S. Kroll, and J. F. Shogren (2005). The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab. *Journal of Economic Behavior & Organization 57*(3), 357–365.

de Secondat, C. and B. De Montesquieu (1748). *The Spirit of Laws*. Hafner Publishers.

Dell, M., B. F. Jones, and B. A. Olken (2014). What do we learn from the weather? the new climate-economy literature. *Journal of Economic Literature 52*(3), 740–98.

Druckman, J. N. and C. D. Kam (2011). Students as experimental participants. In J. N. Druckman, D. P. Green, J. H. Kuklinski, and A. Lupia (Eds.), *Cambridge handbook of experimental political science*, Volume 1, Chapter 4, pp. 41–57. Cambridge UP.

Eckel, C. C. and P. J. Grossman (2008). Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior & Organization 68*(1), 1–17.

Eysenick, M. W. (2004). *Psychology: an international perspective* (1 ed.). Psychology Press. 684.

Fischbacher, U., S. Gächter, and E. Fehr (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics letters 71*(3), 397–404.

Forsythe, R., J. L. Horowitz, N. E. Savin, and M. Sefton (1994). Fairness in simple bargaining experiments. *Games and Economic behavior 6*(3), 347–369.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives 19 (4)*(4), 25–42.

Gill, D. and V. Prowse (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review 102*(1), 469–503.

Graff Zivin, J., S. M. Hsiang, and M. Neidell (2018). Temperature and human capital in the short and long run. *Journal of the Association of Environmental and Resource Economists 5*(1), 77–105.

Graff Zivin, J. and M. Neidell (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics 32*(1), 1–26.

Haushofer, J. and E. Fehr (2014). On the psychology of poverty. *Science 344*, 862–867.

Hirschhorn, J. N. and M. J. Daly (2005). Genome-wide association studies for common diseases and complex traits. *Nature Reviews Genetics 6*, 95.

Hsiang, Solomon, M., M. Burke, and E. Miguel (2013). Quantifying the influence of climate on human conflict. *Science 341 (6151)*(1235367).

Huntington, E. (1922). *Civilization and climate*. Yale University Press.

Johnson, N. D. and A. A. Mislin (2011). Trust games: A meta-analysis. *Journal of Economic Psychology 32*(5), 865–889.

Jones, B. F. and B. A. Olken (2010). Climate shocks and exports. *American Economic Review 100*(2), 454–59.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1986). Fairness and the assumptions of economics. *Journal of business*, S285–S300.

Kenrick, D. T. and S. W. MacFarlane (1986). Ambient temperature and horn honking: A field study of the heat/aggression relationship. *Environment and behavior 18*(2), 179–191.

Kiruga, M. (2018, January). The ethnic rivalry that holds kenya hostage. `http://www.theafricareport.com/East-Horn-Africa/the-ethnic-rivalry-that-holds-kenya-hostage.html`.

Larrick, R. P., T. A. Timmerman, A. M. Carton, and J. Abrevaya (2011). Temper, temperature, and temptation: Heat-related retaliation in baseball. *Psychological Science 22*(4), 423–428.

Lee, H. D. P. (1962). *Aristotle: Meteorologica*. Harvard University Press.

Mackworth, N. H. (1946). Effects of heat on wireless telegraphy operators hearing and recording morse messages. *British Journal of Industrial Medicine 3*(3), 143–158.

Mebane, Jr., W. R. (2017). Anomalies and frauds(?) in the kenya 2017 presidential election. Working paper.

Miguel, E., S. Satyanath, and E. Sergenti (2004). Economic shocks and civil conflict: An instrumental variables approach. *Journal of Political Economy 112*(4), 725–753.

Obradovich, N. (2017). Climate change may speed democratic turnover. *Climatic Change 140*(2), 135–147.

Obradovich, N. and J. H. Fowler (2017). Climate change may alter human physical activity patterns. *Nature Human Behaviour 1*(5), 0097.

Obradovich, N., D. Tingley, and I. Rahwan (2018). Effects of environmental stressors on daily governance. *Proceedings of the National Academy of Sciences*, 201803765.

Okbay, A., J. P. Beauchamp, M. A. Fontana, J. J. Lee, T. H. Pers, C. A. Rietveld, P. Turley, G.-B. Chen, V. Emilsson, S. F. W. Meddens, S. Oskarsson, J. K. Pickrell, K. Thom, P. Timshel, R. de Vlaming, A. Abdellaoui, T. S. Ahluwalia, J. Bacelis, C. Baumbach, G. Bjornsdottir, J. H. Brandsma, M. Pina Concas, J. Derringer, N. A. Furlotte, T. E. Galesloot, G. Girotto, R. Gupta, L. M. Hall, S. E. Harris, E. Hofer, M. Horikoshi, J. E. Huffman, K. Kaasik, I. P. Kalafati, R. Karlsson, A. Kong, J. Lahti, S. J. v. d. Lee, C. deLeeuw, P. A. Lind, K.-O. Lindgren, T. Liu, M. Mangino, J. Marten, E. Mihailov, M. B. Miller, P. J. van der Most, C. Oldmeadow, A. Payton, N. Pervjakova, W. J. Peyrot, Y. Qian, O. Raitakari, R. Rueedi, E. Salvi, B. Schmidt, K. E. Schraut, J. Shi, A. V. Smith, R. A. Poot, B. St Pourcain, A. Teumer, G. Thorleifsson, N. Verweij, D. Vuckovic, J. Wellmann, H.-J. Westra, J. Yang, W. Zhao, Z. Zhu, B. Z. Alizadeh, N. Amin, A. Bakshi, S. E. Baumeister, G. Biino, K. Bønnelykke, P. A. Boyle, H. Campbell, F. P. Cappuccio, G. Davies, J.-E. De Neve, P. Deloukas, I. Demuth, J. Ding, P. Eibich, L. Eisele, N. Eklund, D. M. Evans, J. D. Faul, M. F. Feitosa, A. J. Forstner, I. Gandin, B. Gunnarsson, B. V. Halldórsson, T. B. Harris, A. C. Heath, L. J. Hocking, E. G. Holliday, G. Homuth, M. A. Horan, J.-J. Hottenga, P. L. de Jager, P. K. Joshi, A. Jugessur, M. A. Kaakinen, M. Kähönen, S. Kanoni, L. Keltigangas-Järvinen, L. A. L. M. Kiemeney, I. Kolcic, S. Koskinen, A. T. Kraja, M. Kroh, Z. Kutalik, A. Latvala, L. J. Launer, M. P. Lebreton, D. F. Levinson, P. Lichtenstein, P. Lichtner, D. C. M. Liewald, L. Cohort Study, A. Loukola, P. A. Madden, R. Mägi, T. Mäki-Opas, R. E. Marioni, P. Marques-Vidal, G. A. Meddens, G. McMahon, C. Meisinger, T. Meitinger, Y. Milaneschi, L. Milani, G. W. Montgomery, R. Myhre, C. P. Nelson, D. R. Nyholt, W. E. R. Ollier, A. Palotie, L. Paternoster, N. L. Pedersen, K. E. Petrovic, D. J. Porteous, K. Räikkönen, S. M. Ring, A. Robino, O. Rostapshova, I. Rudan, A. Rustichini, V. Salomaa, A. R. Sanders, A.-P. Sarin, H. Schmidt, R. J. Scott, B. H. Smith, J. A. Smith, J. A. Staessen, E. Steinhagen-Thiessen, K. Strauch, A. Terracciano, M. D. Tobin, S. Ulivi, S. Vaccargiu, L. Quaye, F. J. A. van Rooij, C. Venturini, A. A. E. Vinkhuyzen, U. Völker, H. Völzke, J. M. Vonk, D. Vozzi, J. Waage, E. B. Ware, G. Willemsen, J. R. Attia, D. A. Bennett, K. Berger, L. Bertram, H. Bisgaard, D. I. Boomsma, I. B. Borecki, U. Bültmann, C. F. Chabris, F. Cucca, D. Cusi, I. J. Deary, G. V. Dedoussis, C. M. van Duijn, J. G. Eriksson, B. Franke, L. Franke, P. Gasparini, P. V. Gejman, C. Gieger, H.-J. Grabe, J. Gratten, P. J. F. Groenen, V. Gudnason, P. van der Harst, C. Hayward, D. A. Hinds, W. Hoffmann, E. Hyppönen, W. G. Iacono, B. Jacobsson, M.-R. Järvelin, K.-H. Jöckel, J. Kaprio, S. L. R. Kardia, T. Lehtimki, S. F. Lehrer, P. K. E. Magnusson, N. G. Martin, M. McGue, A. Metspalu, N. Pendleton, B. W. J. H. Penninx, M. Perola, N. Pirastu, M. Pirastu, O. Polasek, D. Posthuma, C. Power, M. A. Province, N. J. Samani, D. Schlessinger, S. T. I. A. Schmidt, Reinhold, T. D. Spector, K. Stefansson, U. Thorsteinsdottir, A. R. Thurik, N. J. Timpson, H. Tiemeier, J. Y. Tung, A. G. Uitterlinden, V. Vitart, P. Vollenweider, D. R. Weir, J. F. Wilson, A. F. Wright, D. C. Conley, R. F. Krueger, G. Davey Smith, A. Hofman, D. I. Laibson, S. E. Medland, M. N. Meyer, J. Yang, M. Johannesson, P. M. Visscher, T. Esko, P. D. Koellinger, D. Cesarini, and D. J. Benjamin (2016). Genome-wide association study identifies 74 loci associated with educational attainment. *Nature 533*, 539–542.

Penrose, L. S. and J. Raven (1936). A new series of perceptual tests: Preliminary communication. *Psychology and Psychotherapy: Theory, Research and Practice 16*(2), 97–104.

Prediger, S., B. Vollan, and B. Herrmann (2014). Resource scarcity and antisocial behavior. *Journal of Public Economics 119*, 1–9.

Ranson, M. (2014). Crime, weather, and climate change. *Journal of Environmental Economics and Management 67*(3), 274–302.

Russell, J. A. and L. F. Barrett (1999). Core affect, prototypical emotional episodes, and other things called emotion: dissecting the elephant. *Journal of personality and social psychology 76*, 805–819.

Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences 106*(37), 15594–15598.

Sears, D. O. (1986). College sophomores in the laboratory: Influences of a narrow data base on social psychology's view of human nature. *Journal of personality and social psychology 51*(3), 515–530.

The Economist (2017, October). Kenya's fresh election is preposterously flawed. `https://www.economist.com/middle-east-and-africa/2017/10/28/kenyas-fresh-election-is-preposterously-flawed`.

Wenz, L., A. Levermann, and M. Auffhammer (2017). North–south polarization of european electricity consumption under future warming. *Proceedings of the National Academy of Sciences*, 201704339.

Zillmann, D. (2006). Transfer of excitation in emotional behavior. In J. T. Cacioppo and R. E. Petty (Eds.), *Social psychophysiology: A sourcebook*, pp. 215–240. New York: Guilford Press.

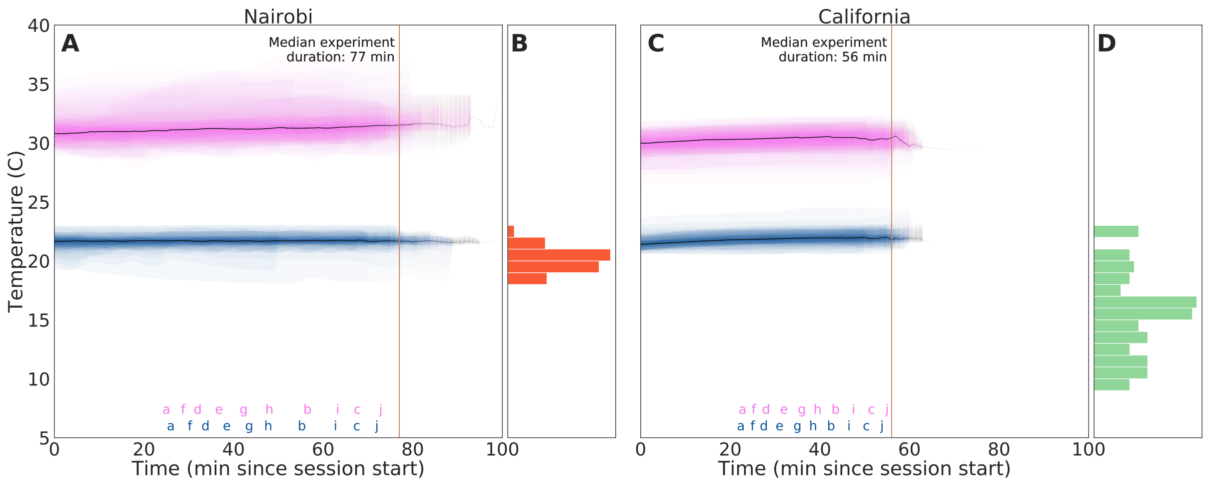

E

Group	Label	Measurement	Description
Individual Cognitive Performance	a	Precision task (Gill and Prowse, 2012)	Productivity, measured through moving sliders
	b	Fluid intelligence (Penrose and Raven, 1936)	Reasoning ability, measured through Raven's matrices
	c	Cognitive reflection (Frederick, 2005)	Ability to override gut response, measured through five questions
Standard Economic Preferences	d	Risk-taking (Eckel and Grossman, 2008)	Participant-chosen variance in outcome from a coin toss, where choosing a higher variance outcome is risk-loving
	d	Rational choice violation I (Eckel and Grossman, 2008)	When coin toss A is preferred to coin toss B, and coin toss B is preferred to coin toss C, but coin toss A is not preferred to coin toss C
	e	Patience (Andreoni et al., 2015)	How one weighs two future periods (time discounting)
	e	Time inconsistency (Andreoni et al., 2015)	How one weighs the present versus the future (present bias)
Social Preferences	f	Fairness (Almås et al., 2010)	Share of tokens earned from precision task given freely to partner
	g	Trust (Johnson and Mislin, 2011)	Share of endowed tokens entrusted to partner, as a measure of trust
	h	Public contribution (Fischbacher et al., 2001)	Share of endowed tokens placed into the group fund, as a measure of willingness to contribute to a public good
	i	Joy of Destruction (Abbink and Sadrieh, 2009)	Share of partner's vouchers or gift cards voluntarily and anonymously destroyed, as a measure of ill-will
	j	Charitable donation (Andreoni, 2006)	Amount of tokens donated to charity, as a measure of charitableness
Emotions and Affect	c	Happiness (Russell and Barrett, 1999)	Self-expressed degree of happiness at the end of the experiment
	c	Alertness (Russell and Barrett, 1999)	Self-expressed degree of alertness at the end of the experiment

Figure 1: Experimental temperatures, timeline, and description. Panels A and C display the median indoor temperature trend during experiments (black) and the distribution of temperature measurements across experiments (colored shading). The opacity of the median line is correlated with the number of experiments that had a measurement recorded at a given time, and the opacity of the shading is correlated with the number of experiments that recorded a temperature measurement falling within a given range. The edge of the shading indicates the minimum and maximum temperatures across experiments at each point in time. The mean temperature experienced by Hot and Control groups was 31.3°C (σ = 0.3°C, n = 77) and 21.7°C (0.1°C, 79) in Nairobi, respectively, and 30.2°C (0.2°C, 76) and 21.8°C (0.2°C, 76) in California. The midpoint of the median start and end times of participation in each experiment module are indicated at the bottom of Panels A and C for treatment (pink) and control (blue) groups, with labels corresponding to those in Panel E. The vertical green line shows the median experiment length in each location. Panels B and D contain session-weighted histograms of daily mean outdoor temperatures on days in which we ran sessions. On average, the mean daily outdoor temperature was 20.0°C (σ = 0.9°C) in Nairobi and 15.6°C (3.7°C) in California. Panel E contains a table of the 14 measurements (outcomes) and brief descriptions of each. These measurements consist of the 12 primary outcomes from the pre-analysis plan plus happiness and alertness.

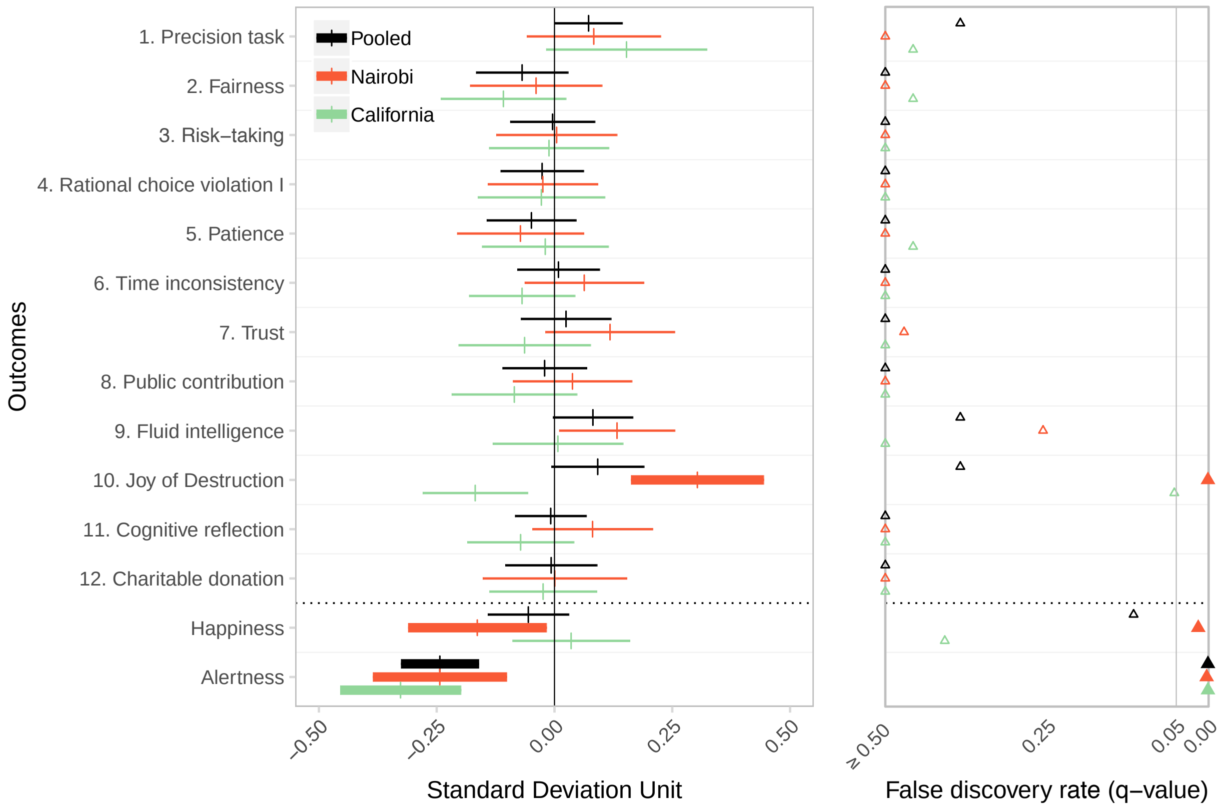

Figure 2: Treatment effects from the main specification in the pooled sample (black; $N = 1,878$ participants), as well as from the Nairobi (dark orange; 1,004) and California (light green; 874) specifications (see SOM Section F for details) in Panel A. These tests were pre-specified in the pre-analysis plan. Treatment effects are presented in terms of standard deviation units for the 12 pre-specified primary outcomes of interest and for the survey questions on participant Happiness and Alertness. For comparability with regression results, individuals who did not respond "Female" or "Male" to the gender survey question were dropped (2% of the sample). Thin standard error bars correspond to the 95% confidence interval around the treatment effect off of the regression for the relevant specification, while thick standard error bars refer to those instances where the multiple testing adjusted False Discovery Rate (FDR) q-value significance level is less than 0.05. Multiple hypothesis testing adjustments are performed on the set of per comparison p-values associated with the treatment across pre-specified primary outcomes, for each specification. Multiple hypothesis testing adjustments are also performed on the set of p-values associated with the treatment for Happiness and Alertness (not pre-specified), for each specification. Panel B records the q-value associated with treatment across each outcome, by specification, where open triangles record $q \geq 0.05$ and closed triangles record $q < 0.05$. The treatment effect for Time Inconsistency is flipped so that a positive treatment effect means more time inconsistency. To graphically display standardized treatment effects for Patience and Time Inconsistency, individual level estimates were developed via non-linear least squares, following Andreoni et al. (2015). Individuals who responded solely with corner cases were dropped (8% of the sample) as well as those who exhibited GARP violations (4% of the sample) for the former calculations. Both individual level estimates for Patience and Time Inconsistency were bottom and top coded to be between 0 and 1 (inclusive) before standardization.

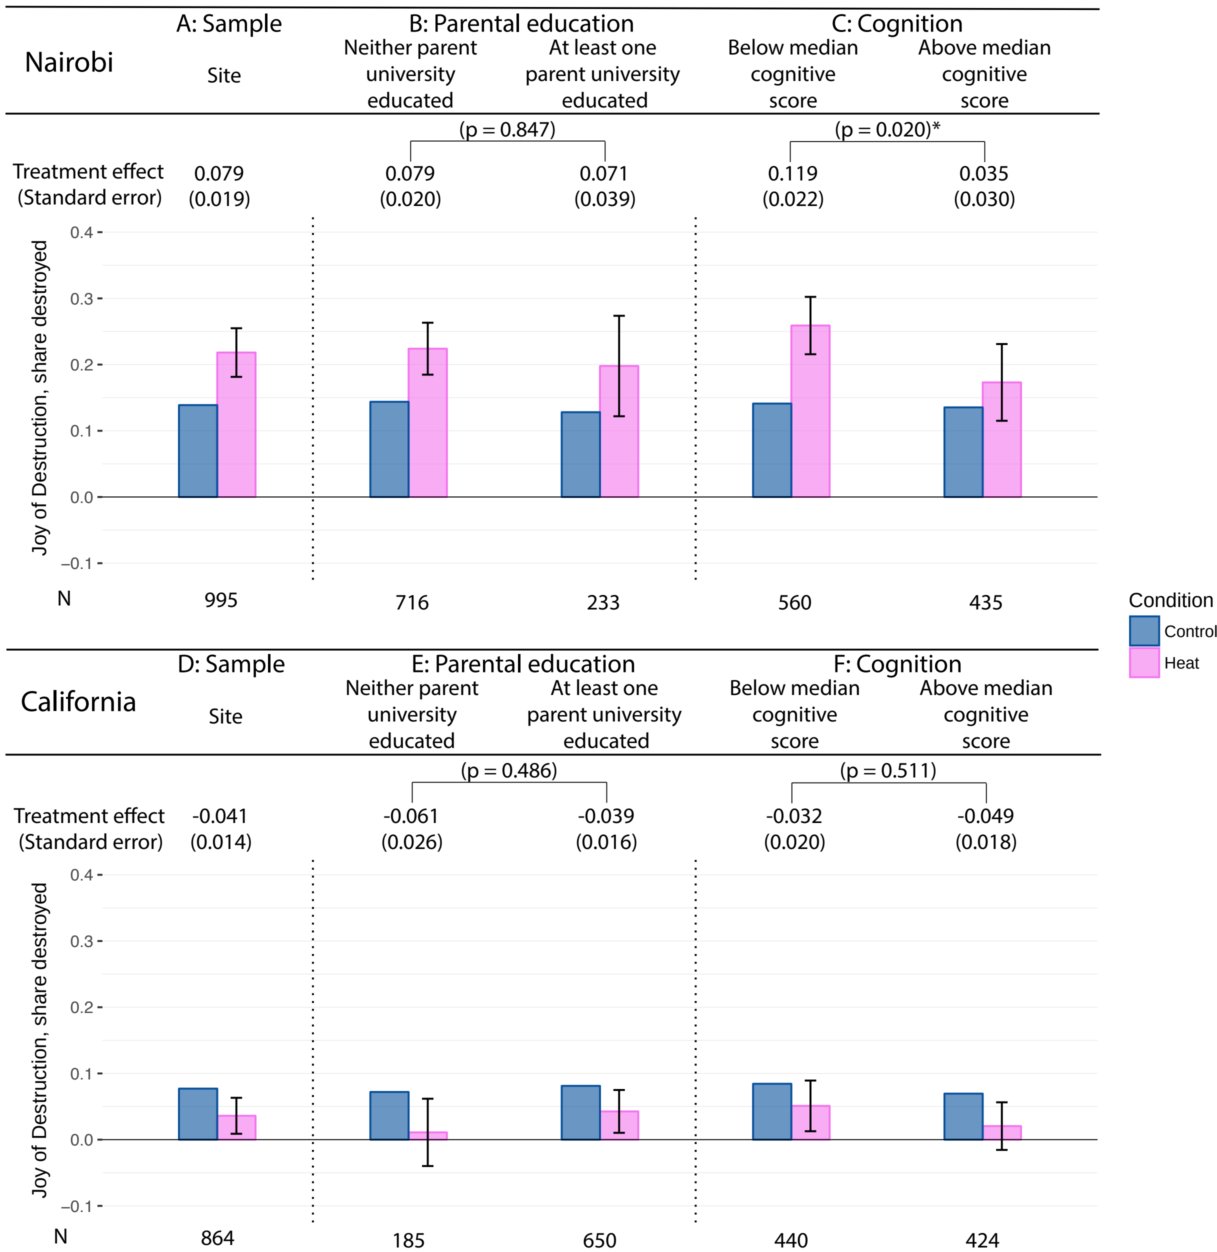

Figure 3: Treatment effects in the Joy of Destruction module in each site. Treatment effects and standard errors are presented above the bars. Blue bars represent the average share of vouchers destroyed among individuals in the control group, while pink bars represent the average share of vouchers destroyed among individuals experiencing the heat treatment, in both the Nairobi sample (Panels A, B, and C) and the California sample (Panels D, E, and F). These sub-group tests were not pre-specified in the pre-analysis plan. Panels A and D provide site-specific average shares of destruction. Panels B and E provide site-specific average shares of destruction by parental university education status (an indicator variable set to one if the participant has at least one university educated parent). Panels C and F provide site-specific average shares of destruction by median "cognition score" status (where "cognition score" is developed by normalizing the sum of individuals' normalized results from the fluid intelligence task and cognitive reflection test). The standard error bars on the pink bars correspond to the 95% confidence interval around the treatment effect off of the main specification regression for the specific subsample. The overarching lines capture the p-value from the interaction effect (from the within-panel regression) between the specific group and treatment, where the leftmost group is the reference group. At the bottom of each panel is the number of individuals in that subsample. For comparability with regression results, individuals who did not respond "Female" or "Male" to the gender survey question were dropped (2% of the sample). $^{*}p < .05$, $^{**}p < .01$.

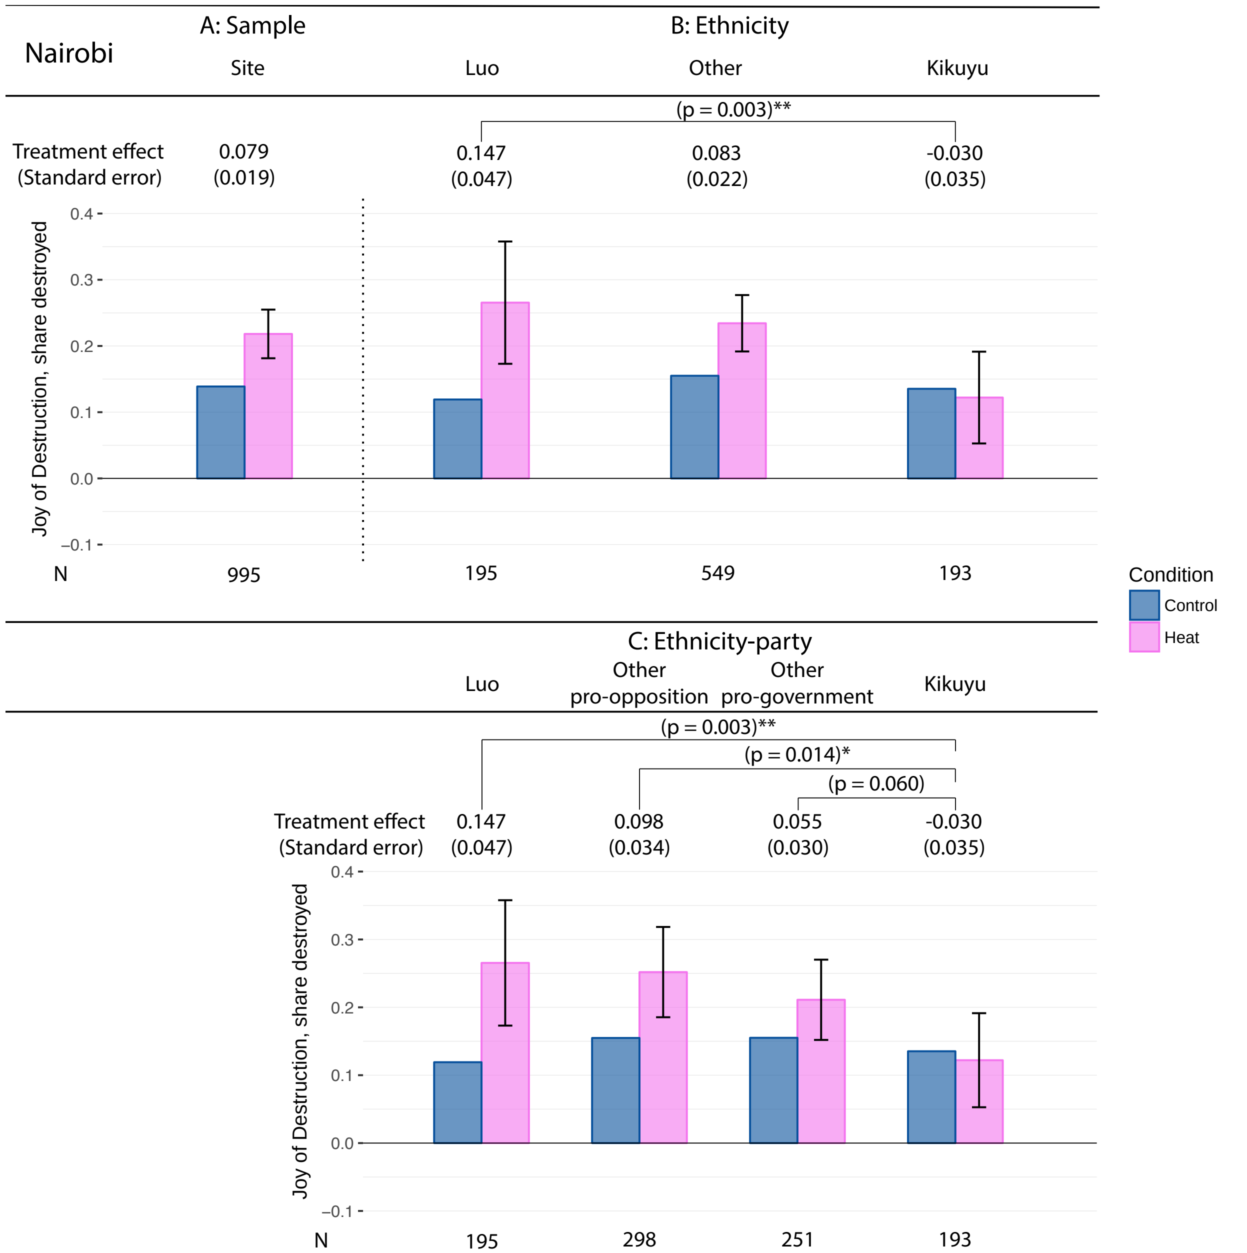

Figure 4: Treatment effects across self-reported ethnicities and ethnicity-party groups in the Joy of Destruction module in Nairobi. Treatment effects and standard errors are presented above the bars. Blue bars represent the average share of vouchers destroyed among individuals in the control group, while pink bars represent the average share of vouchers destroyed among individuals experiencing the heat treatment. These sub-group tests were not pre-specified in the pre-analysis plan. Panel A provides average shares of destruction in Nairobi. Panel B provides average shares of destruction by self-reported ethnicity in Nairobi. Panel C provides a similar breakdown as Panel B but splits other ethnicities into "other pro-opposition" or "other pro-government". "Other pro-opposition" refers to those individuals who self-identified with an ethnicity aligned with ethnic Luos, the ethnicity in opposition in the 2017 election, while "other pro-government" refers to those individuals who self-identified with an ethnicity aligned with ethnic Kikuyus, the ethnicity in power in the 2017 election. The standard error bars on the pink bars correspond to the 95% confidence interval around the treatment effect off of the main specification regression for the specific subsample. The overarching lines capture the p-value from the interaction effect (from the within-panel regression) between the specific group and treatment, where the leftmost group is the reference group. At the bottom of each panel is the number of individuals in that subsample. For comparability with regression results, individuals who did not respond "Female" or "Male" to the gender survey question were dropped (2% of the sample). $^{*}p < .05, ^{**}p < .01$.

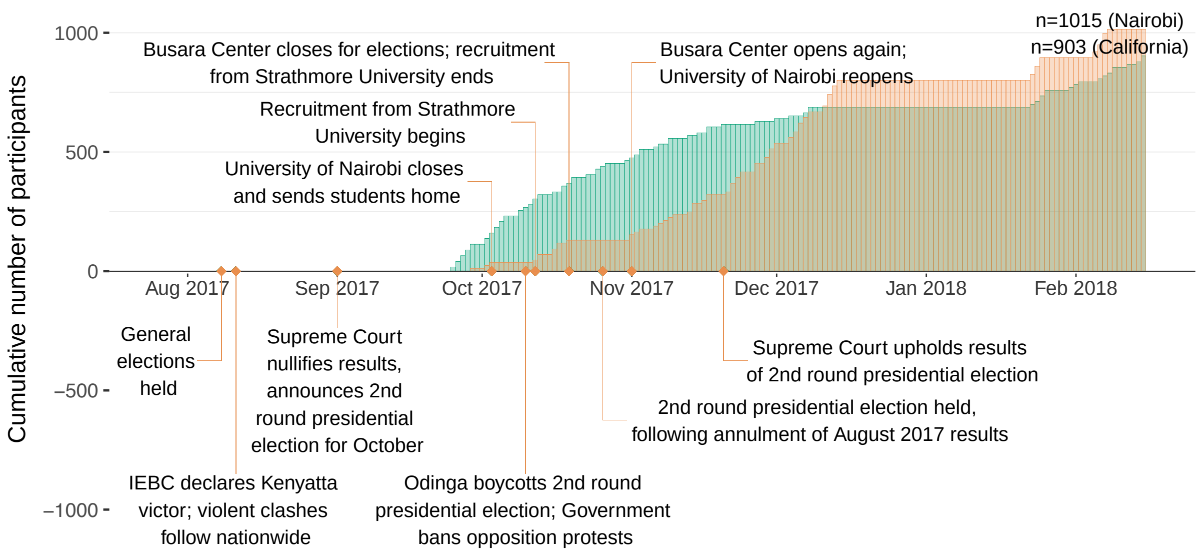

Figure 5: Experimental timelines for Nairobi (dark orange) and California (light green). The colored bars display the cumulative number of experiment participants. Events directly relating to the Nairobi site are included above the timeline, while events relating to the Nairobi political context are included below the timeline (see SOM Sections D.8 and B.1).