

Teze Disertace

k získání vědeckého titulu "doktor věd"
ve skupině věd sociálních a humanitních

Anti-social behavior and Discrimination: Selected topics

název disertace

Komise pro obhajoby doktorských disertací v oboru ekonomické teorie a jejich dějiny

Jméno uchazeče: Doc. PhDr. Michal Bauer, Ph.D.

Pracoviště uchazeče Univerzita Karlova v Praze, CERGE a Národohospodářský ústav AV ČR, v.
v. i.

Místo a datum ...20.7.2023, Praha.....

Introduction and summary

The dissertation presents a collection of research papers that focus primarily on anti-social behavior and discrimination, since such behavior can lead to a range of undesirable societal outcomes. They fit in my long-term research agenda which aims to shed light on “behavioural” mechanisms that may contribute to pressing global social issues, including poverty, discrimination and violent inter-group conflicts. To do so, I conduct large-scale field experiments in relevant settings and among disadvantaged groups.

The canonical economic model assumes that individuals are mainly motivated by their own self-interest. This view has been challenged by work in behavioral economics, since many people, as we learnt after decades of research in experimental economics, are not solely own-payoff maximizers. It is widely recognized that the prevalence of prosocial behaviors, reflecting altruism, inequality aversion, reciprocity or social welfare-maximizing preferences are common (e.g., Charness and Rabin 2002; Gintis et al. 2005). These departures from pure self-interest are important, because such individual concerns for fairness and the well-being of other people, at the expense of self-regard, are important in determining a range of desirable societal outcomes – provision of public goods, participation in public life, and cooperation (e.g., Fehr and Gächter 2002; Gintis et al. 2005).

Yet, we also know from casual observations that some people can be surprisingly nasty to other individuals. However, we know little empirically about prevalence and causal determinants of anti-social preferences, i.e., the dark side of human social preferences, defined as a desire to reduce another person’s material payoff for the mere purpose of harming, without creating personal material gain and without fairness justification (e.g., reducing inequality or a reciprocal response to hostility).¹ Experiencing pleasure from being nasty represents another fundamental -- but thus far less studied -- departure from own money-maximizing behavior (Abbink and Sadrieh 2009; Herrmann, Thoni, and Gächter 2008, Ortmann and Zhang 2016)² because it reduces the propensity to cooperate even in situations in which mutual cooperation is an equilibrium for selfish players (Fehr, Hoff, and Kshetramade 2008). While economic agents motivated purely by self-interest are predicted to reduce social welfare only when they can personally gain, the scope for welfare losses and mutually destructive conflicts is magnified when decision-makers derive utility from harming others. Further, nasty behavior typically receives stronger condemnation than behavior guided by

¹ The literature uses different terms when describing the dark side of social preferences, meaning those preferences that place a negative weight on the well-being of other individuals, including: nastiness, spitefulness, anti-social behavior, aggressive competitiveness, and destructiveness. In this dissertation, I refer to this type of preferences as anti-social preferences, nastiness or destructiveness and use these terms interchangeably.

² For recent reviews of the experimental work on anti-social behavior, see Sanjaya (2023); Karakostas, Tran, and Zizzo (2022).

pure self-regard, as indicated by the term “sadism” used by social psychologists to describe behavior motivated by a negative weight placed on the well-being of others (Murphy and Ackermann 2014).

This thesis presents a collection of papers that contribute to understanding of determinants and mechanisms that can contribute to a greater prevalence of anti-social behavior. The main focus is on identifying different factors that can systematically increase individual inclinations to engage in harmful behavior (parental background, economic and health shocks, social context). I also address a methodological question about how to most efficiently measure anti-social behavior integrated into large-scale data collections. In the final part, I describe a project that does not focus on anti-social preferences as the other chapters, but estimates the effects of poverty on self-control in inter-temporal decisions, and thus shares the overall motivation to shed light on another channel that can contribute to perpetuation of inequality and hardship. All papers have a unified methodological approach, since they are based on primary data collections in the field in relevant settings (disadvantaged communities in Eastern Slovakia, children in the Czech Republic, representative online samples in Slovakia and the Czech Republic, rural Uganda, low-income individuals in Kenya), using experimental tools. The advantage of employing the experimental economics toolbox is (i) the ability to more credibly measure anti-social preferences, allowing to rule out the role of variety of confounders that are ubiquitous when researchers aim to identify anti-social preferences using observational data and (ii) the ability to credibly establish the causal effects of studied factors on prevalence of anti-social behavior, by randomly allocating subjects into treatment and control groups.

The first part provides evidence showing that people are more prone to engage in nasty behavior, malevolently causing financial harm to other people at own costs, when they make decisions on behalf of a group rather than when making choices individually on their own. We establish this new behavioral regularity in four large-scale experiments among adolescents (in Slovakia and Uganda), university students and a nationally representative sample of adults in Slovakia -- more than seven thousand subjects in total (N = 7,426). We test several potential mechanisms, and the results suggest that the “destructiveness shift” in groups is driven by lower perception of individual responsibility when individuals make decisions in the group context. The findings have several implications. First, they provide empirical support for the idea that perception of individual responsibility is an important regulator of the dark side of human personality, in line with self-signaling models (Benabou and Tirole 2011; Falk 2020). They illuminate that, in the group context, a decision environment that is ubiquitous in the real world, the perception of individual responsibility can easily be diluted, leading to socially undesirable impacts on behavior. Second, the results help to explain why decisions made by groups tend to be less pro-social than individual decisions. We show that this is partly due to the effect of group context on anti-social behavior, a new behavioral regularity

described in this paper, and partly due to the role of group deliberation, which fosters a “selfishness shift”, a mechanism featured in earlier work (Gary Charness and Sutter 2012; Kugler, Kausel, and Kocher 2012). These findings have potentially important implications for organizations aiming to curb anti-social behavior.

The second part studies scapegoating -- one potentially important trigger of anti-social behavior towards ethnic minorities. Social scientists have long speculated that members of dominant social groups tend to shift punishment for social ills originating within their own group onto innocent members of other, weaker groups. This behavior, termed scapegoating, is sometimes considered to be an important psychological mechanism in the emergence of pogroms, witch-hunts, and large-scale violence. However, anecdotal and historical evidence cannot rule out the role of standard economic motives in harming innocent members of minority groups, and an experimental test of this behavioral phenomenon in controlled environments had been missing so far. This is what we provided in this paper. We have developed a new incentivized task, the Punishing the Scapegoat Game, to uncover how observing harmful actions against members of one’s own group shapes the punishment of innocent individuals. We study the behavior of young adults in Eastern Slovakia, a region experiencing intergroup tensions. We find that when no harm is done, there is no evidence of discrimination against the ethnic minority. In line with the idea that minorities are often treated as scapegoats by dominant groups, we show that when a member of one’s own group is harmed, the punishment “passed” on innocent individuals more than doubles when innocent individuals are from the minority, as compared to when they are from the dominant group. Further, when participants can punish actual Wrongdoers, they punish Wrongdoers from the minority group more harshly than Wrongdoers from the majority group, for the same wrongdoing. These results illuminate how latent biases in preference can transform individualized tensions into a group conflict, dragging minorities into conflicts in a way that is completely unrelated to their behavior.

The third part “Covid-19 and Hostility against Foreigners” is motivated by the observation that harmful behavior against out-group members often rises during periods of economic hardship and health pandemics. We test the widespread concern that the Covid-19 crisis may fuel hostility against people from other nations. Using a controlled money-burning task, we elicited anti-social behavior behavior among a nationally representative sample ($n = 2,186$) in the Czech Republic during the first wave of the pandemic. We provide evidence that exogenously elevating the salience of the Covid-19 crisis increases prevalence of anti-social behavior against foreigners from the EU, USA and Asia. This behavioral response is similar across various demographic sub-groups. Further, we observe zero to small negative effects for both domestic out-groups and in-groups, suggesting that the salience of Covid-19 might negatively affect behavior not only towards foreigners but to other people more generally, though these findings are not

conclusive. While there is a lot of research quantifying the impacts of Covid-19 on health and economic outcomes, this paper provided one of the first pieces of evidence on its impacts in the social domain. The results underscore the importance of not inflaming anti-foreigner sentiments and suggest the need to monitor impacts of the crisis on behavior in the social domain in later stages of the pandemic.

The fourth part explores the role of parental background in the formation of pro-social and anti-social preferences in young children. Our subjects, aged 4–12 years, are classified into social-preferences types based on simple binary-choice dictator games. The main finding is that the children of parents with low education are less altruistic, more selfish, and more likely to be weakly spiteful. It also stands out against the overall development of preferences, as we find children to become more altruistic, less selfish, and less likely to be weakly spiteful with increasing age. The results, supported by a complementary analysis of World Values Survey data, suggest an important role of socialization in the formation of social preferences. Recent work of other researchers has built on these descriptive findings, by providing causal evidence of remedial early-childhood programs in shaping social preferences (Kosse et al. 2020; A. Cappelen et al. 2020).

The fifth part addresses a methodological question concerning measurement of economic preferences: Can a short survey instrument reliably measure a range of fundamental economic preferences across diverse settings? We focus on survey questions that systematically predict behavior in incentivized experimental tasks among German university students (Becker et al. 2016) and were implemented among representative samples across the globe (Falk et al. 2018). This paper presents results of an experimental validation conducted among low-income individuals in Nairobi, Kenya. We find that quantitative survey measures -- hypothetical versions of experimental tasks -- of time preference, attitude to risk and altruism are good predictors of choices in incentivized experiments, suggesting these measures are broadly experimentally valid. At the same time, we find that qualitative questions -- self-assessments -- do not correlate with the experimental measures of preferences in the Kenyan sample. Thus, caution is needed before treating self-assessments as proxies of preferences in new contexts. Further, this study tested the experimental validity of survey preference measures in a new domain, anti-social preferences. We document that survey measures of anti-social preferences only weakly predict incentivized behavior, which strengthens the case for investing resources into gathering incentivized measures in this domain.

The final part studies psychological impacts of poverty on decision-making and self-control. The focus on estimating the impact of hardship on behaviorally revealed preferences relates the paper to the third paper in this collection, but the focus is on a different domain of preferences, namely time preference rather than social preferences. We measured discount rates among farmers in Uganda who made decisions

about when to enjoy entertainment instead of working. To circumvent the role of economic constraints, we experimentally induced thoughts about poverty-related problems, using priming techniques. We find that thinking about poverty increases the preference to consume entertainment early and to delay work. The identified effects on time preference can help to explain persistent nature of poverty.

Part 1: Nastiness in Groups³

This part is motivated by the question whether people become nastier to other people when they act in groups rather than when they act individually on their own. Social scientists have long been concerned about the effects on behavior when individuals act as a part of a group.⁴ In his seminal work on individual behavior in crowds, Le Bon (1895, p. 35), for example, argues that individuals placed in groups lose their sense of personal responsibility and “socially deviant” inclinations may shape behavior: “*Isolated he may be a cultivated individual; in a crowd he is a barbarian—that is, a creature acting by instinct.*” In line with this concern, anecdotal evidence suggests that some people are more prone to engage in vandalism or violence against perfect strangers in collective settings, such as crowds, gangs, or armed groups. Covert obstructionism and sabotage are particularly common in complex organizations and large bureaucracies. Some of this harmful behavior is hard to explain by pure calculated self-regard, and suggests that some people harbor a desire to be nasty and to harm others, even at personal cost to themselves. This part provides evidence of the behavioral importance of the effect of making decisions in a group on greater prevalence of nastiness, and explores its sources.

To measure such behavior, we implement an economic experiment in which participants made decisions whether to financially harm an anonymous experimental counterpart at a small personal cost. We experimentally manipulate whether subjects make choices individually or on behalf of a group, but with no communication among group members. We refer to the difference in behavior as the effect of the group context.

The initial set of lab-in-field experiments focused on adolescents in Slovakia (P.E 1) and in Uganda (P.E. 2), and documents that making decisions in the group context systematically increases the prevalence of nasty behavior in both settings. Then, we probe further among a sample of university students in Slovakia (Experiment 1), in order to shed more light on the underlying mechanism of the observed effects. The results favor the interpretation that the behavioral change caused by the group context is driven mainly by reduced perception of individual responsibility. We also consider other mechanisms and show that motivation to please nasty group members, signaling toughness, action bias, and the role of in-group vs.

³ This paper is co-authored with Jana Cahliková, Dagmara Celik-Katreniak, Julie Chytilová, Lubomír Cingl and Tomáš Želinský. An earlier version of this paper, published as IZA discussion paper and CEPR working papers, was circulated with the title “Anti-social behavior in groups”. Online appendix with supplementary information is provided on my personal websites.

⁴ In Plato’s opinion, democracy involves rule by irrational mobs and for this reason he favored the rule of an enlightened individual (Allport 1968). Alexander Hamilton, James Madison, and John Jay – the first two being members of the US Constitutional Convention – shared a similar concern: “*In all very numerous assemblies, of whatever character composed, passion never fails to wrest the scepter from reason. Had every Athenian citizen been a Socrates, every Athenian assembly would still have been a mob*” (Publius 1948, p. 248).

out-group biases are unlikely to explain the “destructiveness shift” caused by the group context. Finally, in order to gauge the robustness of the main finding, we implement two online experiments among a nationally representative samples of adults in Slovakia (Experiments 2 and 3). We show that the effect holds across a broad spectrum of economic and demographic groups.

Methods

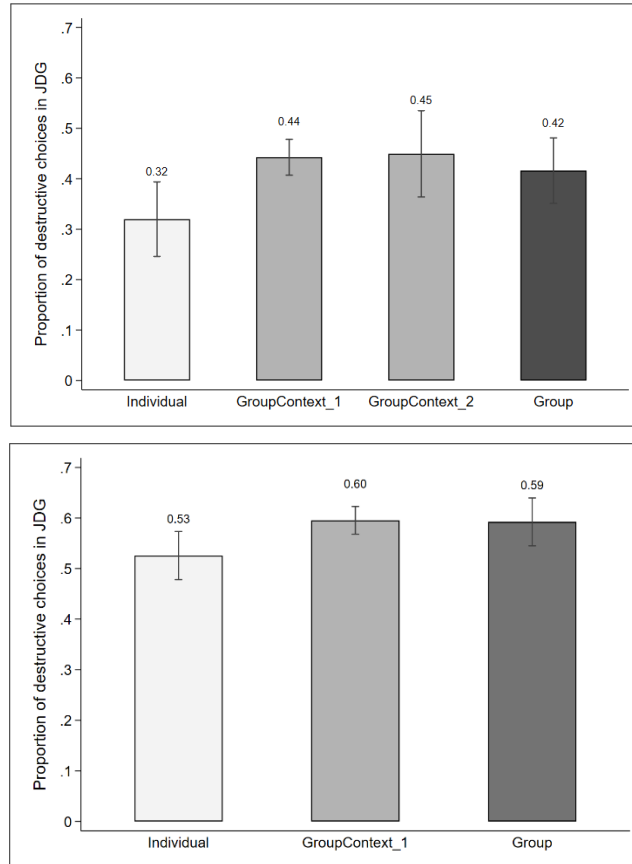
We conducted four large-scale experiments, all studying the influence of making decisions in a group context rather than individually, on the prevalence of nasty behavior. In all these experiments, subjects made decisions in the Joy of Destruction mini-game (Abbink and Sadrieh 2009; Abbink and Herrmann 2011). This task is designed to uncover the dark side of human social preferences. Subjects made a decision whether to pay a small amount of money in order to lower the reward of an anonymous counterpart. Destroying the resources of a counterpart creates greater inequality, is not a response to unkind behavior, and is costly for the decision maker. Therefore, this task helps us to tease out whether people acting in groups destroy more due to greater self-interest or pleasure to be nasty, because the dominant strategy for purely self-regarding (or fairness-minded) individuals is not to engage in destructive behavior, whereas decision-makers that put a negative weight on the utility of the other person can engage in such harmful behavior.

Results

The two initial experiments focus on the behavior of adolescents from disadvantaged regions. This subject pool is a natural starting point for this enquiry, as crime, urban riots and other forms of destructive behavior are more common among youth with lower socioeconomic status, as compared to other demographic groups (Lochner and Moretti 2004; Deming 2011; Bauer, Chytilová, and Pertold-Gebicka 2014). The experiments were conducted in schools in Eastern Slovakia (P.E. 1, N=630) and in rural Uganda (P.E. 2, N=1679). To elicit behavior in a group context, we asked subjects to state preferences about how they wanted their group of three individuals to decide when making a joint decision in the Joy of Destruction game (JDG), prior to discussing a joint decision with other group members. In addition, in Experiment 1, we elicited incentivized decisions, in which individuals made decisions on behalf of their group. We find that making a decision in a group context, rather than as an individual, increases the prevalence of destructive behavior in JDG. Similar effects of the group context hold in both settings, Slovakia and Uganda. In addition to observing the effects in JDG, we find that decision-making in the group context increases defection in PDG. Thus, the group context makes subjects more likely to harm others, both when such action brings them financial

gains as well as when they need to sacrifice their own resources in order to do harm. We refer to this behavioral change in groups as a “destructiveness shift.”

Figure 1. Destructive behavior in groups, Experiment 1 (Slovakia, Adolescents – Upper Panel; Uganda - Lower Panel)



Notes: The proportion of unconditional destructive choices in the Joy of Destruction game. Bars indicate 95% confidence intervals.


The destructiveness shift observed in the pilot experiments is consistent with several plausible mechanisms. First, self-signaling models highlight people’s concern about building a positive self-image of being a moral person as an important regulator of inter-personal behavior (Benabou and Tirole 2011; Benabou, Falk, and Tirole 2018). Decision-making on behalf of groups involves reduced salience of self and creates perception of diffused responsibility, since more individuals are involved in the decision-making and can be attributed the responsibility for the group decision (Darley and Latane 1968; Dana, Weber, and Kuang 2006). Both of these aspects are predicted to reduce self-image concerns, allowing people with a latent intrinsic preference to be nasty to more freely act on such desires. Second, work on

group identity suggests that salient group boundaries may give rise to “in-group/out-group” biases and lead to aggressively competitive behavior, with groups aiming to out-compete other groups in relative terms (Tajfel 1981; Durlauf 1999). Finally, people may care about their social image and what other group members think of them (Benabou and Tirole 2006; Bursztyn and Jensen 2017). When choices on behalf of a group are observable by other group members (or are expected to be discussed with peers), some individuals may choose to act destructively, in order to signal toughness to others or to please anti-social group members.

Therefore, as a next step, we conducted Experiment 1 with additional experimental conditions and measures among 795 university students from Eastern Slovakia, which aims to isolate the role of reduced perception of individual responsibility in groups from other mechanisms. We implemented two experimental conditions in which subjects made decisions on behalf of a group. Together, the patterns provide support for the interpretation that contextual factors that reduce perception of individual responsibility are the main driver of the increased destructiveness when individuals make choices as a part of a group. In Experiment 2, we collaborated with a survey agency and implemented a large online experiment on a nationally representative sample of adults in Slovakia (N=4,243). We again replicate the main pattern – making decisions on behalf of the group systematically increases the prevalence of destructive behavior in JDG and also increases the prevalence of defection in PDG. Further, we show that the effect holds for various sub-groups, in terms of gender, age, education, income level, size of municipality, and political orientation of the respondents. Finally, in Experiment 3 (E.3) implemented among a different sample of adults (N=3,349), we address the concern that the decision situation in JDG may prime conflict and potentially lead to experimenter demand effects. We adapt the task so that it includes an option to act kindly in addition to the option to act destructively. Further, the counterpart is passive and thus cannot be seen as an adversary. We observe the same pattern as in E.1 and E.2. Subjects are more likely to choose the destructive option in *GC_Hidden* than in *Individual*, suggesting that the main effect is not driven by greater sensitivity to cues of conflict in the group setting.

Taken together, the experiments provide evidence of a systematic behavioral regularity: making decisions under a “cover” of group when it is unclear who is responsible for the decision, rather than as an individual, magnifies the prevalence of nasty behavior, and this result holds across a broad range of demographic and economic groups.

Figure 2. Robustness of the Effect of A Group Context on Destructive Behavior in the Joy of Destruction Game --, Experiment 4 (Slovakia, Representative Sample of Adult Population)



Contribution and implications

By providing the novel causal evidence establishing how a group context influences the prevalence of nastiness, this paper contributes to existing experiments on decision-making in groups. Earlier work made important progress in studying situations in which self-interest is pitted against social welfare or moral outcomes, using a range of experimental tasks such as the Prisoners' Dilemma game, the Trust game, and the Dictator game. It has documented that groups or individuals in salient groups often behave less pro-socially than individuals do – they are less willing to sacrifice their own resources to increase social welfare or to achieve fair allocation of payoffs.⁵ A prevailing interpretation of this stylized pattern is that decision-making in groups fosters rational self-regarding behavior, in line with textbook assumptions, meaning that groups are more prone to maximize their own payoff and to disregard the welfare of others (see excellent surveys by Charness and Sutter (2012) and Kugler, Kausel, and Kocher (2012)). Such “selfishness shift”

⁵ Most of this literature is based on comparisons of choices made by individuals and unitary teams/groups (typically composed of three members), in which group members make a joint team decision. Lower pro-social behavior among groups as compared to individuals has been documented in a range of economic experiments - the Dictator game, Ultimatum game, and the Trust game – see, for example, Kugler et al. (2007), Luhan, Kocher, and Sutter (2009) and Charness and Sutter (2012) for more references. Psychologists have extensively studied behavior in the Prisoners' Dilemma game; nearly all of their studies show that groups defect more often than individuals (see Wildschut et al. 2003 for a meta-study).

has been typically attributed to communication among group members, helping to recognize a dominant strategy. The main contribution of this paper is to uncover a new mechanism: the effect of a group context making people more prone to engage in nasty actions. Further, the results provide empirical support for the idea that perception of individual responsibility is an important regulator of the dark side of human personality, in line with self-signaling models (Benabou and Tirole 2011; Falk 2020).

In terms of practical implications, our findings suggest that organizations or public bodies seeking to limit obstructionism and other manifestations of nasty inclinations, may want to create environments that foster perception of individual responsibility. For example, by framing decisions as the choices of specific individuals, rather than presenting them as decisions of the whole units, teams or committees, or by making the ways actions of individual group members translate into group outcomes more transparent.

Part 2: Shifting punishment on minorities: Experimental Evidence of Scapegoating⁶

This part is motivated by the question how can localized social ills spiral into widespread aggressive behavior, often with tragic consequences for whole societies? Although it is widely acknowledged that aggressive behavior against ethnic minorities, including pogroms, often arise unexpectedly and spread quickly even in previously peaceful communities (Fearon and Laitin 2000; Bardhan 2005; Esteban and Ray 2008), little is known empirically about the mechanisms that underlie these *changes* in behavior. In this paper, we study people’s biases in punishment behavior, in response to wrongdoing targeted towards members of one’s own group. A novel aspect is that we focus on situations when punishers cannot punish the wrongdoer but only “bystanders”, as in many real-life situations. Inspired by work in social psychology and political science (Doob et al. 1939; Allport 1954; Girard 1979; Staub 1992; Glick 2005), we take seriously the idea that people may have a tendency to pass revenge on innocent individuals from negatively-stereotyped minority groups. We refer to such behavior as minority scapegoating. To study this behavioral phenomenon, we develop a new experimental task and document that a desire to “pass” the punishment on innocent individuals more than doubles when such persons are from an ethnic minority, as compared to when bystanders are from the dominant group. Such shifting of punishment on minorities violates a fundamental fairness principle (Kant 1965 and classical philosophers) embedded in the legal codes of most

⁶ This paper was co-authored with Jana Cahliková, Julie Chytilová, Gerard Rolad and Tomáš Želinský. It was published in *Economic Journal*, 133 (652):1626–1640. The data collection was supported by a grant from the Czech Science Foundation (17-13869S), and by the Slovak Research and Development Agency (APVV-0125-12 and APVV-19-0329). Bauer and Chytilová acknowledge support from ERC-CZ/AV-B (ERC300851901).

modern societies, i.e. that people should be punished only for wrongs they are responsible for and that they intentionally committed. Demonstrating such bias in punishment preferences of dominant groups is important, because it may drag minorities into conflicts that are completely unrelated to their behavior and transform individualized tensions into group conflicts.

Methods

Identifying people's desire to pass punishment on innocent individuals, and its determinants, is empirically challenging with observational data. First, it is nearly impossible to rule out the role of the standard economic incentives to harm innocent individuals, such as simple self-interested grabbing of resources from them. In addition, in most real-life situations there is an element of uncertainty about who originated the harm. Members of the dominant group may punish innocent individuals from minority groups, because they attribute responsibility for misfortunes to actions of minority groups. A controlled experimental environment allows eliminating these confounding mechanisms, and focusing on whether shifts in punishment onto minorities are embedded in preferences or decision heuristics. Specifically, our aim is to identify an individual desire to engage in harmful behavior towards unrelated parties, in response to experienced hostility against someone one cares about. In light of the literature on scapegoating, the main question is whether indirect revenge becomes stronger when the unrelated party (bystander) is a member of a minority group. Answering this question requires an experimental setup that allows a researcher to (i) measure how people behave when punishment of wrongdoers is inhibited and people can punish only individuals who could not have causally contributed to the original harmful act, (ii) measure punishment responses in one-shot anonymous interactions that are costly for the punisher, and that provide no scope for material benefits of punishment, and (iii) compare behavior towards a weaker minority group and towards the own group.

To address these challenges, we employ a novel experimental paradigm, the *Punishing the Scapegoat Game*. In this game, impartial spectators from the majority ethnic group can impose a monetary punishment on others at own cost, after observing that someone malevolently destroyed the earned income of an individual from their own group. The existing incentivized experiments on punishment of socially undesirable behavior focus exclusively on direct punishment of individuals who make active decisions whether or not to violate a social norm, including the Third Party Punishment Game⁷ on which we build. In contrast, a key element of the Punishing the Scapegoat Game is that we add a fourth person, the passive Scapegoat who does not know the wrongdoer and who is not involved in any way in the original

⁷ The Third Party Punishment game has been originally developed by Fehr and Fischbacher (2004) and later became one of the prominent experimental tasks to study direct punishment of norm-violators, see, for example, Bernhard, Fischbacher, and Fehr (2006); Goette, Huffman, and Meier (2006); Henrich et al. (2010).


wrongdoing. This feature allows us to separate the person who commits a harmful act and a person whom the impartial spectator can punish. We exogenously manipulate information about the ethnicity of the Scapegoat to test whether people are more prone to redirect the punishment onto minority group members, if they cannot punish the Wrongdoer.

We are interested in how behavior of spectators (Punishers) from the majority ethnic group responds to information on whether a Wrongdoer caused harm to a Victim from the majority ethnic group and how serious the harm was. We measure whether the gap in punishment of the Roma Scapegoat as compared to the majority Scapegoat is larger in a situation when the Punishers respond to observing harmful behavior targeting a member of their group than in a situation when no such harm was done. An important feature of using such difference-in-difference approach to identify scapegoating is that we can test whether out-group bias in punishment happens above and beyond out-group bias in circumstances when people do not respond to harm happening to their own group. We study the interaction between real ethnic majority and minority group members. The setting is Eastern Slovakia, a region experiencing high unemployment, where the majority population lives side-by-side with an economically disadvantaged, segregated, and negatively stereotyped Roma ethnic minority, which represents approximately 15% of the local population.

Results

We show that a non-negligible fraction of Punishers (23%) redirect punishment onto Scapegoats (bystanders) when they cannot punish the Wrongdoer. Importantly, the identity of the Scapegoat matters a lot. We find that the interaction between the intensity of harm committed by the Wrongdoer and the Scapegoats' identity has a large and robust effect on the punishment of the Scapegoat. When Wrongdoers harm the Victim, the punishment is twice as severe when the Scapegoat is from the Roma minority than when the Scapegoat is from the majority group. When Wrongdoers do not harm the Victim, ethnic majority Punishers do not discriminate against Scapegoats from the ethnic minority. Therefore, discrimination against the Roma minority Scapegoats *arises* only when Wrongdoers harm the Victim. In other words, the observed out-group bias in punishment happens above and beyond out-group bias in circumstances when people do not respond to harm done to their own group.

Figure 3: Punishment of the Scapegoat, by Scapegoat's ethnicity



Notes: Mean punishment of the Scapegoat, by the ethnicity of the Scapegoat and the harm caused by the Wrongdoer to the Victim. Punishers (and Victims) are from the majority ethnic group. "Scapegoat SAME" indicates that the Scapegoat also comes from the majority ethnic group, while "Scapegoat OTHER" indicates that Scapegoat is ethnic Roma. Differences between the conditions are tested using the Wilcoxon rank-sum test; p-values are presented at the top. The sample is composed of Punishers from the majority group.

In addition to making decisions whether or not to lower the payoff of the Scapegoat, decision-makers also made choices whether to punish the Wrongdoer directly. In this task also, we observe a magnified tendency to engage in revenge: punishers from the majority group punish wrongdoers from the minority more harshly than members of their own group. Together, we establish that the biases in punishment preferences of the dominant group against the minority group we study are not limited to situations when members of the minority are responsible for harmful actions, but, importantly, mark behavior also in situations when it is clear that minority members have nothing to do with the harm done.

Contribution and implications

Our experimental design contributes to existing lab-in-field experiments on punishment of norm violations across group boundaries (Goette, Huffman, and Meier 2006; Bernhard, Fischbacher, and Fehr 2006). Earlier work has made progress in studying group biases in the direct punishment of active norm violators. We contribute by developing an experimental task that sheds light on determinants of the punishment of passive individuals. In this way, we can show that revenge may spill-over and affect behavior not only towards a wrongdoer but also towards unrelated individuals from an ethnic minority. More broadly, the paper

contributes to experiments that study existence of in-group vs. out-group biases based on real-life group attributes, such as ethnicity (Fershtman and Gneezy 2001; Falk and Zehnder 2013; Bauer et al. 2018; Berge et al. 2020; Kranton et al. 2020; Goette, Huffman, and Meier 2006). In-group favoritism and out-group discrimination have also been measured using a minimal-group experimental paradigm, by creating artificial group boundaries in the laboratory, based on, for example, having T-shirts of the same color or sharing preferences for art (Tajfel 1981; Chen & Li 2009). Here we show that measured out-group biases may be strongly influenced by the social decision-environment – they are more pronounced when decisions happen in an environment characterized by wrongdoing against someone from the dominant group, as compared to biases revealed in decisions made in “peaceful” circumstances. Thus, latent discriminatory preferences may be triggered in hostile social environment.

At face value, our findings have several potentially important implications. First, we show that pure observation of injustice and wrongdoing against the individual’s own group activates latent discriminatory preferences, both when treating innocent individuals as well as wrongdoers. This indicates that courts, and other settings in which people make punishment choices, are particularly discrimination-prone environments, in line with evidence of strong biases against minorities in judicial decisions (Alesina and Ferrara 2014; Rehavi and Starr 2014; Shayo and Zussman 2011). Second, the results suggest that ethnic minorities are at greater risk when social problems and unfair behavior become salient features of the societal environment. So far, economists have typically attributed sudden spikes in aggressive behavior towards weaker groups to changes in economic incentives or beliefs about the likelihood of facing a penalty for aggressive behavior (Christopher Blattman and Miguel 2014; Grosfeld, Sakalli, and Zhuravskaya 2018; Miguel 2005), assuming that revealed (anti-) social preferences towards other groups are stable. In our experiment, economic incentives are held constant and thus cannot explain the scapegoating behavior observed. Of course, this does not imply that economic incentives do not play an important role in real-life aggression towards minority groups. However, our evidence suggests that may not be the complete picture. It strengthens the case for taking seriously “behavioral” channels through which deterioration of the social environment may fuel inter-group conflicts.

Part 3: Covid-19 Crisis and Hostility against Foreigners⁸

Motivation

Social scientists have long argued that difficult life conditions imposed upon individuals by external forces that threaten physical wellbeing and safety, such as economic and political upheavals or widespread disease, may create a fertile environment for xenophobia and out-group hostility. Evidence suggests that hostile behaviors and conflicts increase during periods of economic problems (Anderson, Johnson, and Koyama 2017; Miguel, Shanker, and Sergenti 2004; Grosfeld, Sakalli, and Zhuravskaya 2018). In the context of a crisis caused by a contagious disease, a particularly plausible mechanism is that people may form hostile attitudes to members of the groups that are associated with transmission of the disease (Murray and Schaller 2016; O’Shea et al. 2020). In light of this reasoning, the Covid-19 crisis has created an unfortunate but suitable testing ground for exploring whether an important, naturally-occurring shock in the health and economic domains spills over to the social domain and magnifies inter-group animosity. Since Covid-19 originally surfaced in China and spreads across borders via interactions with people from other countries, contemporary commentators have suggested that it may foster prejudice against foreigners, particularly against people from Asia (CNN 2020). For example, Fernand de Varennes, the UN Special Rapporteur, warns that “COVID-19 is not just a health issue; it can also be a virus that exacerbates xenophobia, hate and exclusion.” (United Nations News 2020).

Methods

In this paper, we test causal evidence that greater salience of a health pandemic accompanied by a severe economic shock increases prevalence of harmful behavior towards people living in other countries. Our evidence is based on a large-scale experiment implemented in midst of the Covid-19 crisis. We elicited hostile behavior among a nationally representative sample (n = 2,186) in the Czech Republic, a medium-sized country in Central Europe, while the pandemic was on the rise during its first wave, and the entire population lived under lockdown and border closure. Several features of our experimental design help us to overcome well-known empirical challenges. First, we directly elicit willingness to cause financial harm in a controlled money-allocation task. Subjects make anonymous, one-shot allocation decisions, in which


⁸ This paper is co-authored with Vojtěch Bartoš, Jana Cahlíková and Julie Chytilová and was published in the *European Economic Review* 137 (2021) 103818. We thank PAQ Research and NMS Market Research for implementing the data collection. Bauer and Chytilová thank ERC-CZ/AV-B (ERC300851901) for support and funding of the data collection. Bartoš thanks the German Science Foundation for their support through CRC TRR 190. Online appendix with supplementary information, the datasets and do-file replication files are available in the Harvard Dataverse repository (doi: 10.7910/DVN/XD8OOL).

they can decide to decrease a monetary reward for another person. Since reducing the reward does not result in pecuniary benefits for the decision-maker (or for anyone else), the choice reveals individual willingness to engage in hostile behavior. Second, we exogenously manipulate information about the identity of the recipient of the reward, in order to identify discrimination against foreigners. Third, we randomly assign the participants either to a treatment condition that increased the salience of Covid-related problems and fears, or to the control condition in which Covid-related challenges were not made salient. Random allocation ensures that participants in the treatment and control conditions are comparable in terms of observable and unobservable characteristics, helping to overcome selection issues and concerns about spurious correlation.

Main results

Our main result is that thinking about Covid-19 increases anti-foreigner sentiments, making people more prone to financially harm people from the EU, the USA and Asia. We show that this is a relatively general response, present across various demographic groups. Further, some of our results also indicate that the salience of Covid-19 negatively affects behavior towards other people more generally, since we find zero to small negative effects for domestic out-groups as well as in-groups. Thus, the negative effects may not be limited to behavior towards foreigners, but this evidence should be viewed only as suggestive.

Figure 4. Effect of the COVID-19 condition on allocations in the Help-or-Harm task, by the identity of the recipients.



Notes: Coefficient plots. Bars represent 95 percent confidence intervals. The dependent variable is the amount allocated. Both panels present estimated coefficients of the COVID-19 condition relative to the CONTROL condition.

Contribution and implications

The paper is related to work on “Behavioral immune system” in social psychology which has documented a correlation between greater exposure to (real or perceived) health threats and measures of group biases in explicit and implicit attitudes. For example, in US states with higher rates of infectious diseases, people exhibited greater racial prejudice (O’Shea et al. 2020). A representative survey from the US shows that citizens who felt more vulnerable to contracting Ebola displayed greater prejudice against immigrants in survey questions (Kim, Sherman, and Updegraff 2016). Moving beyond correlations, showing a disease-related picture primes increased prejudice among subjects in the lab (Duncan and Schaller 2009) and among a sample of M-Turk workers (O’Shea et al. 2020). We contribute by providing causal evidence of the impacts of a naturally occurring health pandemic on incentivized behavior among a representative sample.

The evidence presented in this part illuminates that Covid-19 can cause damage in the social domain. Nevertheless, we see results in this part as only an initial step and in the full paper, we discuss several limitations, including the relatively low magnitude of the estimated effects, open questions about potential underlying mechanisms and the fact that the effects were estimated during initial stages of the pandemic. We also outline potential directions for future research in this area.

In terms of potential policy implications, our results underscore the importance of making sure political and other opinion-leaders avoid associating or blaming foreigners and other countries for the Covid-19 crisis. Placing blame as a political strategy can either create or tap into elevated anti-foreigner sentiments, and consequently increase the risk that the health and economic crises will become compounded by unravelling of international collaborations.

Part 4: Parental background and other-regarding preferences in children⁹

Motivation

In a quest to establish whether and when during lifecycle pro-social and anti-social preferences are formed, experimental tools have been increasingly adapted and used to study their development during childhood and adolescence (Sutter, Zoller, and Glätzle-Rützler 2019). The aim of this part is to understand the link between family background and the formation of preferences towards others during childhood. Understanding this relationship is important because altruism helps to establish and maintain cooperative outcomes, while spitefulness and selfishness can undermine cooperation, and thus, these preferences are of relevance for society's welfare and, potentially, for individual success on as well. The existing literature finds that older children take the welfare of others more into account (Fehr et al. 2008; Harbaugh et al. 2003a); this result is confirmed by our study, too, suggesting this is sensitive period during which preferences develop. This paper takes one step forward by studying the role of family background in this process.

Methods and the main findings

Among a sample of 4-12 year-old children ($N = 275$) in the Czech Republic, we collect experimental measures of social preferences, patience, cognitive skills and other observable characteristics. The experimental tasks allow us to classify subjects into different behavioral types, as defined by theory: altruistic, inequality averse, spiteful, and selfish. This data is matched with parental survey data, in order to study the link between parental socio-economic status and social preferences in children.

The main finding is that children of less-educated parents are more selfish, less altruistic, and more likely to exhibit a weak form of spite. In terms of mechanisms that could explain the link, our study provides a few suggestive results. The relationship with parental education is robust to controlling for a large set of child characteristics—age, gender, patience, cognitive skills, and health—as well as sibling composition and class fixed effects. Thus, parental background does not seem to affect other-regarding preferences via peer effects associated with school choice, lower child skills, or worse health. Using the WVS data coming from the same population but a different sample from ours, we find that for less-educated parents it is less important to instil unselfish behaviors in their offspring and that they are less likely to volunteer, suggesting that differences in parental values and socialization could explain heterogeneity in child preferences.

⁹ This paper was co-authored with Julie Chytilová and Barbara Pertold-Gebicka. It was published in *Experimental Economics* 17(1): 24-46. The research was supported by a grant from the CERGE-EI Foundation under a program of the Global Development Network and by the Czech Science Foundation (P402/10/P103 - Chytilová, P402/12/G130 - Bauer).

Implications

Demonstrating that the formation of social preferences is related to disadvantaged socio-economic backgrounds has potentially important economic and social consequences (Fehr and Hoff 2011). The existing evidence shows a strong relationship between low parental background and fundamental life outcomes, such as greater participation in crime or being unemployed (Bowles et al. 2008) and our results suggest that this could be, in part, due to differences in the acquisition of social preferences. The findings strengthen the case for taking seriously the possibility that there may be additional benefits of education interventions targeting disadvantaged children, by closing the gap in pro-social preferences. Interestingly, motivated by findings in this and other papers, recent randomized controlled trials in Germany and the US estimated the causal impact of educational policies and remedial programs on social preferences in children (A. W. Cappelen et al. 2020; Kosse et al. 2020). Specifically, Kosse et al. (2020) show that randomized provision of mentor can close the gap in pro-social behaviour between low SES and high SES children.

Part 5: Using Survey Questions to Measure Preferences: Lessons from an Experimental Validation in Kenya¹⁰

Motivation

Fundamental preferences in the economic domain, such as time discounting and risk preferences, and in the social domain, such as altruism, reciprocity and spitefulness, constitute key elements of individual decision-making. Figuring out ways to accurately measure these preferences among large samples in the field holds considerable promise since doing so may shed light on the sources of vast differences in preferences observed across individuals and societies, and their role in fundamental economic choices and societal trajectories. While measuring preferences using incentivized tasks is generally considered the gold standard, implementing incentivized tasks among large samples outside of the controlled environment of an experimental laboratory is often infeasible, given that they are relatively expensive and time consuming. Consequently, a potentially attractive alternative is to employ survey questions instead of incentivized experiments, but there has long been widespread concern that non-incentivized self-reported survey measures of preferences may not reliably capture real life choices.

To tackle this important methodological trade-off, Falk et al. (2018) have recently developed an innovative short (7-8 minutes) survey module, designed to measure a wide range of economic preferences. It has been implemented among representative samples of subjects in more than seventy countries (Falk et al., 2018), creating the most comprehensive global data set with comparable measures of preferences, namely, the Global Preference Survey (GPS). Measures of preferences in each domain are constructed as a weighted average based on one objective quantitative item -- a hypothetical version of an experimental task -- and one subjective qualitative item that measures self-reported willingness to act in a certain way.

To establish the validity of the survey preference measures, Becker et al. (2016) perform a careful experimental validation of the survey questions, and document that survey measures of preferences do predict choices in incentivized decisions. The validation was conducted among students at the University of Bonn, Germany. Given the wide coverage of the existing GPS data set and the convenience of the survey module in terms of implementation, it has the potential to become a widely adopted instrument for (i) studying differences in preferences across societies and their relationships with economic outcomes, (ii) employing preference measures as control variables when identifying causal effects of other factors correlated with preferences, and (iii) as outcome variables in new randomized controlled trials aiming to uncover the effects of various interventions on individual preferences.

¹⁰ The paper was co-authored with Julie Chytilová and Edward Miguel. It was published in the *European Economic Review*, Volume 127, August 2020, 103493. Michal Bauer and Julie Chytilová thank the Czech Science Foundation for funding of the data collection (17-13869S) and for support of further work on the project (20-11091S).

This part adds to these efforts and aims to be useful in three ways. First, we test the experimental validity of the survey questions outside of a sample of university students from a rich country, by focusing on a sample from the other end of the global distribution of income and education. Our experimental subjects are residents of working class neighborhoods (sometimes referred to as “slums”) in Nairobi, Kenya, a setting with a different set of institutions and economic constraints. The participants are aged between 20-46, with average income of around USD 3 per day, and 54% are unemployed. Establishing the experimental validity of the measures among this subject pool is important for several reasons. Most of humanity lives in low and middle income countries, outside of Western, Educated, Industrialized, Rich and Democratic societies (Henrich et al. 2010), in which the original GPS validation was conducted. Next the GPS module is particularly suitable to be integrated into large-scale follow-up surveys in randomized control trials, which are routinely implemented by development economists (Banerjee and Duflo 2012), often in Africa, and thus knowledge of whether the survey preference measures predict incentivized behavior among low-income individuals in Kenya is a useful input for scholars considering the adoption of these measures.

Second, comparing the results of analogous validations conducted in Germany and Kenya is methodologically interesting, because measures of economic preferences in GPS are derived from both objective quantitative tasks as well as subjective qualitative questions, based on self-assessments.^{11,12} There is a legitimate concern that subjective self-assessments might be understood and interpreted in different ways across countries, which can attenuate their ability to uncover personality traits and complicate cross-country comparisons. For example, the Big Five measures of personality traits, the most widely-used method to measure and classify personality traits in psychology, are based on self-assessments, and recent attempts to validate the Big Five measures have failed to reliably predict the intended personality traits in low- or middle-income countries, in contrast to samples from the wealthy countries for which they were originally developed (Laajaj et al. 2019; Gurven et al. 2013). An advantage of GPS is that, besides self-assessments, it also contains quantitative questions that are arguably less subject to this issue, because they directly define the parameters and nature of the decision and more closely mirror the incentivized experimental task. Thus, we can test whether quantitative questions are relatively more robust predictors of actual incentivized behavior across two diverse settings, as compared to qualitative self-assessments.

¹¹ An example of a qualitative question from GPS would be “Please tell me, in general, how willing or unwilling you are to take risks, using a scale from 0 to 10”, or “How willing are you to give to a charity without expecting anything in return?”

¹² For recent discussions about the importance of replications and other methods that aim to foster credibility of research findings see, for example, Maniadis, Tufano, and List (2014) and Christensen and Miguel (2018).

Third, we place additional emphasis on the types of preferences that are likely to be especially important in settings with low social capital and a history of inter-group conflict, issues that are particularly pressing in low-income countries (C. Blattman and Miguel 2010). While the GPS focuses on measuring preferences relevant for explaining positive aspects of human social behavior, such as generalized altruism and reciprocity, we also assess the experimental validity of survey questions designed to measure the dark side of human social behavior. Specifically, we test the validity of questions designed to uncover anti-social preferences, such as spite, and distinguish between generalized, in-group, and out-group preferences, along both prosocial and anti-social dimensions.

Methods and findings

The sample in our study are 123 subjects from the Kibera neighborhood in Nairobi, Kenya. The participants come from a low-income environment, are between 20 and 46 years of age, more than half are unemployed, half are women and, on average, they have two children. The experiments were implemented in a state-of-the-art experimental economics laboratory in the Busara Center for Behavioral Economics (Haushofer et al. 2014). Subjects were invited to the lab twice, for visits one week apart, where the time gap was introduced in order to minimize spillovers between the survey and experimental measures. During one visit, they made choices in a set of incentivized experiments, while during the other, they answered non-incentivized survey questions. We elicited measures of the following types of preferences: (i) time discounting, (ii) risk preference, (iii) ambiguity aversion, (iv) altruism (generalized, in-group, and out-group), (v) anti-social behavior (generalized, in-group, and out-group), and (vi) positive reciprocity.

We find several patterns. First, we show that *quantitative survey measures* of time preference, attitude to risk and altruism are good predictors of choices in incentivized experiments. This finding reinforces the findings from a similar validation exercise performed among university students in Bonn (Becker et al. 2016), and thus, together, the two studies document the experimental validity of these measures at opposite ends of the global income and education distribution. Second, this study tested the experimental validity of survey preference measures in a new domain, *anti-social preferences*, which is arguably most prone to social desirability biases. We document that survey measures of anti-social preferences only weakly predict incentivized behavior, which strengthens the case for investing resources into gathering incentivized measures in this domain. Third, we find that the *subjective qualitative questions* on preferences do not meaningfully correlate with the experimental measures in the Kenyan sample, in contrast to the German sample. Thus, caution is needed before interpreting these measures as proxies of preferences in all contexts.

Table 1: Correlations between quantitative survey measures and experimental measures

Preference	Quantitative survey item	Kenya: Kibera residents		Germany: Bonn students		
		OLS Coefficient	Correlation	Correlation	Measures	
		(1)	(2)	(3)	(4)	
Time	Staircase measure: 5 interdependent choices	0.33***	0.40***	0.55***	comparable	
	between an early and delayed amount of money	[0.16; 0.50]	[0.24; 0.54]			
Risk	Staircase measure: 5 interdependent choices	0.21**	0.25***	0.34***	comparable	
	between a lottery and varying safe options	[0.03; 0.38]	[0.07; 0.41]			
Ambiguity aversion	Hypothetical choice between a bag with known	0.13	0.13	n.a.		
	and unknown number of balls of different color	[-0.05; 0.31]	[-0.05; 0.30]			
Reciprocity	Hypothetical choice of the amount of a gift given	0.12	0.29***	0.35***	exp. different; survey	
	to a stranger who provided help	[-0.06; 0.30]	[0.12; 0.45]			
Reciprocity (diff)	Hypothetical choice of the amount of a gift given	0.06	0.19**	n.a.	comparable	
	to a stranger who provided help	[-0.12; 0.24]	[0.02; 0.36]			
Altruism	generalized	Hypothetical choice of the amount donated to a	0.41***	0.41***	0.39***	comparable
		charity (out of Ksh3200)	[0.25; 0.58]	[0.26; 0.55]		
	in-group	Hypothetical choice of the amount donated to a	0.33***	0.36***	n.a.	
		charity that helps people in ancestral home area	[0.16; 0.50]	[0.20; 0.51]		
		(out of Ksh3200)				
	out-group	Hypothetical choice of the amount donated to a	0.40***	0.38***	n.a.	
		charity that helps people in other parts of Kenya	[0.23; 0.56]	[0.22; 0.52]		
		than ancestral home area (out of Ksh3200)				
Anti-social behavior	generalized	Hypothetical decision between (3200, 3200) or	0.05	0.05	n.a.	
		(3150, 1600) for self and for another person	[-0.13; 0.23]	[-0.13; 0.22]		
	in-group	Hypothetical decision between (3200, 3200) or	0.07	0.07	n.a.	
		(3150, 1600) for self and for a person from	[-0.12; 0.26]	[-0.11; 0.25]		
		ancestral home area				

out-group	Hypothetical decision between (3200, 3200) or (3150, 1600) for self and for a person from other parts of Kenya than ancestral home area	0.14 [-0.04; 0.32]	0.14 [-0.04; 0.31]	n.a.
-----------	---	-----------------------	-----------------------	------

Notes: Column 1 is an OLS coefficient from a regression of the standardized experimental measure on the standardized survey item. Column 2 displays Spearman correlations between the survey item and the respective experimental measure (one for each preference type, except for reciprocity, where we use two experimental measures). ***, **, and * denote significance at the 1-, 5-, and 10-percent level, respectively. Below each OLS coefficient and Spearman correlation, the table reports 95% confidence interval in the square brackets. Column 3 displays the correlation between experimental and quantitative survey measures from the validation study of Becker et al. (2016) among university students in Germany. Column 4 indicates to what extent measures from our study in Kenya and measures from the German study are comparable.

Implications

What lessons about measuring preferences using survey questions can we draw from the available evidence? First, our results should boost confidence in the ability of *objective quantitative* GPS survey measures of preferences, based on hypothetical tasks, to predict high-stakes incentivized behavior in experiments designed to measure a range of preferences across economically and culturally diverse settings. Second, qualitative survey questions have been shown to do a good job of predicting behavior in incentivized experiments in rich (mostly German) settings (Dohmen et al. 2011; Becker et al. 2016) and a range of real-life behaviors (Bonin et al. 2007; Dohmen et al. 2011). In light of our findings it might be tempting to conjecture that self-assessments are generally unreliable in low-income settings, in contrast to high-income settings. However, since we do not know which factor (participant education levels, exposure to abstract concepts, social desirability biases, culturally-specific ways of interpreting self-assessments, etc.), or which combination of factors, explains the lower experimental validity of self-assessments in the Kenyan context, generalizing from a single study to all low-income environments would still be premature. Rather, our findings highlight the need for more validation studies, ideally a series of comparable validation exercises in a diverse set of contexts across the globe, in order to better understand the characteristics of individuals or societies for which the qualitative self-assessments are informative.

Part 6: Psychological Effects of Poverty on Time Preferences¹³

Motivation

Development economists have long observed that low-income individuals often behave impatiently: many spend surprisingly large shares of their budgets on the consumption of temptation goods, including entertainment and alcohol, do not take advantage of high-return investment opportunities, and repeatedly take out high-interest loans (Banerjee and Duflo 2007; De Mel, McKenzie, and Woodruff 2008; Duflo, Kremer, and Robinson 2011). Using economic experiments, progress has been made in studying the influence of negative income shocks on time discounting, and most of the evidence suggests that having a lower income makes people behave more impatiently (e.g., Tanaka, Camerer, and Nguyen 2010). Documenting such effects of financial pressure on behaviour is important, because they may contribute to a self-reinforcing nature of poverty. It remains an open question whether the effects of low income on inter-temporal decision-making are due only to shifts in economic constraints, such as liquidity constraints, life expectancy and arbitrage opportunities, or whether they reflect changes in behaviourally-revealed time preferences due to psychological constraints.

Recent work on “scarcity” or the “psychology of poverty” has documented that living in an environment of ubiquitous scarcity consumes cognitive resources and adversely affects emotions (Mani et al. 2013; Mullainathan and Shafir 2013; Haushofer and Fehr 2014; Haushofer and Shapiro 2016). However, little evidence exists of the impacts of these factors on economic behaviour, such as productivity, preferences and decision-making (Kremer, Rao, and Schilbach 2019). This paper contributes to this literature by focusing on psychological effects of poverty on time preferences. Motivated by research in behavioural economics, which highlights that delaying gratification and exercising self-control (Bernheim and Rangel 2004; Fudenberg and Levine 2006; Muraven and Baumeister 2000) can be seen as costly mental processes, we test the idea that the cognitive or emotional burdens associated with living in chronic poverty may tax self-control, and thus directly affect time preferences.

Methods

In order to study whether poverty can lead people to behave impatiently through channels other than standard budget-constraints or long-term processes of preference formation, in this part we study the behaviour of extremely poor farmers in rural Uganda who made choices in a controlled longitudinal experiment, in which we elicited time discounting of entertainment and exogenously activated thinking

¹³ The paper is co-authored with Vojtěch Bartoš, Julie Chytilová and Ian Lively. It was published in the *Economic Journal* 131 (638), August 2021, 2357–2382.

about poverty-related problems. The paper offers several empirical innovations. First, we study time discounting of the consumption of a temptation good – watching entertaining videos instead of working. An advantage of implementing the entertainment-discounting task among this population is that it eliminates the role of liquidity and time constraints, and thus mitigates some of the key confounds involved in measuring time *preferences*. Second, to circumvent identification issues and income effects, we directly manipulate concerns about financial difficulties, using priming techniques. Finally, in addition to measuring intertemporal choices, we integrate new tools to monitor attention and information acquisition when participants make decisions. These measures of the *decision-making process* help us to separate the two psychological mechanisms through which poverty may influence inter-temporal decision-making: higher time preference or reduced attention.

Our subjects are 289 adult subsistence farmers in Northern Uganda, who were recruited to perform a tedious manual task for a fixed work period on two dates, one week apart. The subjects were given a budget of “entertainment minutes” which they could use to watch entertaining videos on tablet computers instead of working on an unpleasant task. To elicit discount rates for the consumption of leisure, subjects were asked to allocate minutes of entertainment over the earlier and later work dates, for five different substitution rates.

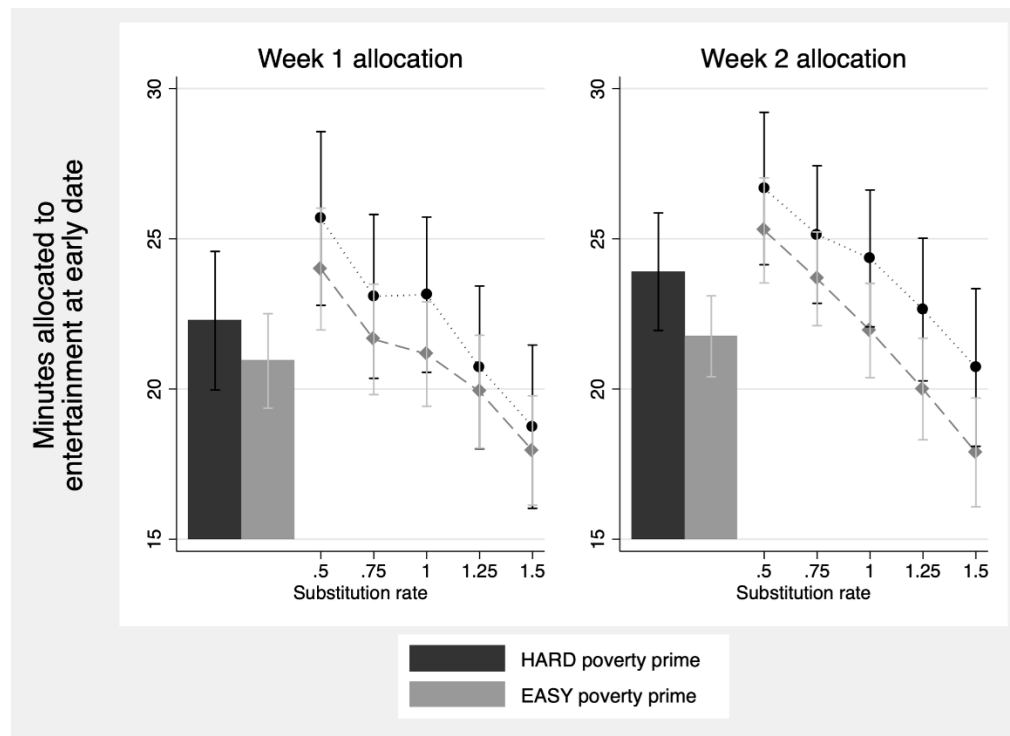
We experimentally induced thoughts about poverty by presenting poverty-related situations to the subjects, as in Mani et al. (2013). Participants were asked how they would go about solving scenarios involving shocks, for example crop damage or a health shock. The scenarios were similar across conditions, except for the severity: half of the subjects deliberated about negative shocks with minor consequences, while the other half considered scenarios with severe consequences before making their inter-temporal choices. Manipulating thoughts about poverty-related concerns resembles priming techniques, a well-established and frequently used method in psychology, and more recently also in economics and finance (Cohn and Maréchal 2016). It refers to mental activation of primed concepts and enables measurement of their pure psychological impact (via cognition and emotions) on behaviour in subsequent tasks. This approach allows us to avoid the confounding influence of liquidity, wealth, access to financial markets, and health, as all of these variables remain unchanged across conditions.

Main results

Our main finding is that thinking about poverty systematically increases preferences for consuming entertainment earlier and delaying work. This effect on discounting is economically meaningful: the poverty-related prime leads individuals to consume 1.7 more minutes of entertainment at an earlier date on a base of 21 minutes. The effect is robust to controlling for a long list of observable characteristics, holds for the whole range of prices of early vs. delayed entertainment, and is slightly stronger when allocation of

leisure/labour has immediate consequences than it is when subjects make allocation decisions one week in advance.

Figure 5: Minutes allocated to entertainment at an early date: by poverty prime, substitution rate, and week



Notes: The thick bars represent choices aggregated over all substitution rates, while the dots indicate choices at the respective substitution rates. Error bars represent 95% level confidence intervals from a regression with standard errors clustered at the individual level.

Further, we show that the effect of the poverty prime on discounting is unlikely to be driven by reduced attention to the task. Using our detailed data on the decision-making process, we find no systematic or significant effects of the poverty prime on decision-making time, patterns of information acquisition, or signs of being distracted while making a decision. We arrive at similar conclusions when analyzing responsiveness to information about the parameters of the choice: the poverty-primed subjects are not less prone to respond to changes in substitution rate, for example.

Based on existing models of time preferences, we describe several mechanisms through which the poverty prime may affect time discounting. We conclude that our empirical findings most closely match the predictions of the costly self-control model of Fudenberg and Levine (2006) and Fudenberg and Levine

(2012), in which individuals are in constant conflict between a short-run self that seeks immediate gratification and a forward-looking long-run self, and in which the parameter capturing costs of self-control is affected by environmental factors, such as anxiety and cognitive load associated with poverty, or the proximity of temptation. An alternative mechanism, denoted in psychology as the myopic-misery hypothesis (Lerner, Li, and Weber 2013; Lerner et al. 2004), is that a negative affect enters individual utility function, and early consumption of entertainment may compensate for this utility loss. While this mechanism can explain the effect of a poverty prime on discounting, it cannot be the sole explanation since it struggles to explain the other patterns we find.

Contribution and implications

The paper contributes to emerging empirical literature testing the psychological effects of poverty on decision-making. Negative income shocks or concerns about income shocks have been shown to reduce cognitive function (Mani et al. 2013; Lichand and Mani 2019). Kaur et al. (2019) show that scarcity of one's own income (but not priming) reduces productivity. Our paper documents impacts of financial anxiety on economic behaviour in a new domain – whether to delay work and enjoy leisure early - and thus may help to explain why sometimes the poor seem to place surprisingly low priority on engaging in income-generating activities, and high weight on consuming temptation goods (Schilbach 2019; Banerjee and Duflo 2007).

Our results speak to a long-standing debate about why the poor behave differently from the rich. “Two-systems” models of individual decision making (Bernheim and Rangel 2004; Fudenberg and Levine 2006) treat decision-making as a result of a strategic interplay between an impulsive agent and a forward-looking agent who can reduce the influence of the impulsive agent only by drawing on a limited budget of cognitive resources. The results indicate that such a two-system model may be a useful way to think about the psychological impacts of poverty. The poor may not necessarily have different hardwired time preferences than the rich, but their impulsive self may more easily affect behaviour due to an increased cognitive load or stress associated with poverty. Also, since such an effect may create a feedback loop between poverty and impatience, our findings provide empirical support for recent efforts to model behavioural poverty traps (see the recent classification of Ghatak 2015), in particular those based on the assumption that poverty directly reduces self-control (Bernheim, Ray, and Yeltekin 2015; Banerjee and Mullainathan 2010).

Further, in terms of methods, we contribute to literature analysing the decision-making process by adapting monitoring tools that have been developed for laboratory experiments, in order to be feasible to implement in a field environment with an important population for which computerized experiments are not suitable. Our effort to gather data on decision-making process is motivated by recent papers, which

have cautioned against automatically interpreting heterogeneity in risky or impatient *behaviour* in experiments as reflecting differences in the underlying *preferences*, since choices may also capture differences in effort and attention, and consequently the quality of the decision-making process more broadly (Andersson et al. 2018; Dohmen et al. 2018). Our approach to addressing this issue is inspired by techniques commonly used in computerized experiments, in which researchers complement choice data with data on the decision-making process. Inspired by mouse-tracking techniques, we use a novel video-recording set-up to obtain detailed measures of decision time, information acquisition and other aspects of the decision-making process in a field setting.

These findings are potentially important for policy. First, if thinking about poverty-related problems directly increases time preferences, then there may be an additional mechanism, besides the standard economic channels, through which even temporary anti-poverty programs may have lasting positive impacts on economic activity and accumulation of assets. In this context, it is noteworthy that a recent series of randomized evaluations of simple unconditional cash transfers finds promising impacts, documenting positive effects on measures of economic activity and human capital investments, but zero or negative effects on alcohol and tobacco consumption (Christopher Blattman and Fiala 2014; Haushofer and Shapiro 2016). Second, perceptions about the sources of inequality have been shown to play an important role in willingness to redistribute from rich to poor (A. W. Cappelen et al. 2007). Negative views on helping the poor are often tied to a presumption that poverty originates in reckless behaviour. Enjoying entertainment while putting off work until later – the choice in our experiment - is frequently featured as an example of such condemnable behaviour. Here we provide unambiguous evidence that the relationship between economic circumstances and (lack of) patience is more complex, by demonstrating that it is, at least in part, driven by poverty damaging the ability to exercise self-control.

References

- Abbink, Klaus, and Benedikt Herrmann. 2011. “The Moral Costs of Nastiness.” *Economic Inquiry* 49 (2): 631–33.
- Abbink, Klaus, and Abdolkarim Sadrieh. 2009. “The Pleasure of Being Nasty.” *Economics Letters* 105 (3): 306–8.
- Alesina, By Alberto, and Eliana La Ferrara. 2014. “A Test of Racial Bias in Capital Sentencing.” *American Economic Review* 104 (11): 3397–3433.
- Allport, Gordon W. 1954. *The Nature of Prejudice*. New York: Addison-Wesley.
- . 1968. “The Historical Background of Modern Social Psychology.” In *The Handbook of Social Psychology*, edited by G. Lindzey and E. Aronson, 1–80. Reading, MA: Addison-Wesley.
- Anderson, Robert Warren, Noel D. Johnson, and Mark Koyama. 2017. “Jewish Persecutions and Weather

- Shocks: 1100–1800.” *Economic Journal* 127 (602): 924–58.
- Andersson, Ola, Hakan J. Holm, Jean Robert Tyran, and Erik Wengstr. 2018. “Risk Aversion Relates to Cognitive Ability: Preferences or Noise.” *Journal of European Economic Association* 14 (June): 1129–54.
- Banerjee, Abhijit V., and Esther Duflo. 2007. “The Economic Lives of the Poor.” *The Journal of Economic Perspectives: A Journal of the American Economic Association* 21 (1): 141.
- . 2012. *Poor Economics*. Public Affairs.
- Banerjee, Abhijit V., and Sendhil Mullainathan. 2010. “The Shape of Temptations: Implications for the Economic Lives of the Poor.” NBER Working Paper n. 15973.
- Bardhan, Pranab. 2005. *Scarcity, Conflict and Cooperation*. Cambridge, MA: The MIT Press.
- Bauer, Michal, Jana Cahliková, Julie Chytilová, and Tomáš Želinský. 2018. “Social Contagion of Ethnic Hostility.” *Proceedings of the National Academy of Sciences* 115 (19): 4881–86.
- Bauer, Michal, Julie Chytilová, and Barbara Pertold-Gebicka. 2014. “Parental Background and Other-Regarding Preferences in Children.” *Experimental Economics* 17 (1): 24–46.
- Becker, Anke, Thomas Dohmen, David Huffman, Armin Falk, and Uwe Sunde. 2016. “The Preference Survey Module: A Validated Instrument for Measuring Time, Risk, and Social Preferences.” IZA Discussion Paper 9674.
- Benabou, Roland, Armin Falk, and Jean Tirole. 2018. “Narratives, Imperatives, and Moral Reasoning.” NBER Working Paper No. 24798.
- Benabou, Roland, and Jean Tirole. 2006. “Incentives and Prosocial Behavior.” *American Economic Review* 96 (5): 1652–78.
- . 2011. “Identity, Moral and Taboos: Beliefs as Assets.” *Quarterly Journal of Economics* 126: 805–55.
- Berge, Lars Ivar Oppedal, Kjetil Bjorvatn, Simon Galle, Edward Miguel, Daniel Posner, Bertil Tungodden, and Kelly Zhang. 2020. “Ethnically Biased? Experimental Evidence from Kenya.” *Journal of the European Economic Association* 18 (1): 134–64.
- Bernhard, Helen, Urs Fischbacher, and Ernst Fehr. 2006. “Parochial Altruism in Humans.” *Nature* 442 (7105): 912–915.
- Bernheim, B. Douglas, Debraj Ray, and Şevin Yeltekin. 2015. “Poverty and Self-Control.” *Econometrica* 83 (5): 1877–1911.
- Bernheim, B Douglas, and Antonio Rangel. 2004. “Addiction and Cue-Triggered Decision Processes.” *American Economic Review* 94 (5): 1558–90.
- Blattman, C., and E. Miguel. 2010. “Civil War.” *Journal of Economic Literature* 48 (1): 3–57.
- Blattman, Christopher, and Nathan Fiala. 2014. “Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda.” *Quarterly Journal of Economics* 129 (2): 697–752.
- Blattman, Christopher, and Edward Miguel. 2014. “Civil War.” *Journal of Economic Literature* 48 (1): 3–57.
- Bon, Gustave Le. 1895. *The Crowd*. New York: Viking (1960).
- Bonin, H., Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde. 2007. “Cross-Sectional Earnings Risk and Occupational Sorting: The Role of Risk Attitudes.” *Labour Economics* 14 (6):

926–37.

- Bursztyn, Leonardo, and Robert Jensen. 2017. “Social Image and Economic Behavior in the Field : Identifying , Understanding and Shaping Social Pressure.” *Annual Review of Economics* 9 (1): 131–53.
- Cappelen, Alexander, John A. List, Anya Samek, and Bertil Tungodden. 2020. “The Effect of Early-Childhood Education on Social Preferences John List Anya Samek Bertil Tungodden.” *Journal of Political Economy* 128 (7): 2739–58.
- Cappelen, Alexander W., John A. List, Anya Samek, and Bertil Tungodden. 2020. “The Effect of Early-Childhood Education on Social Preferences.” *Journal of Political Economy* 128 (7): 2739–58.
- Cappelen, Alexander W, Astri Drange Hole, Erik Sorensen, and Bertil Tungodden. 2007. “The Pluralism of Fairness Ideals : An Experimental Approach.” *American Economic Review* 97 (3): 818–27.
- Charness, G., and M. Rabin. 2002. “Understanding Social Preferences with Simple Tests*.” *Quarterly Journal of Economics* 117 (3): 817–869.
- Charness, Gary, and Matthias Sutter. 2012. “Groups Make Better Self-Interested Decisions.” *Journal of Economic Perspectives* 26 (3): 157–76.
- Chen, Yan, and Sherry Xi Li. 2009. “Group Identity and Social Preferences.” *American Economic Review* 99 (1): 431–457.
- Christensen, Garret, and Edward Miguel. 2018. “Transparency , Reproducibility , and the Credibility of Economics Research.” *Journal of Economic Literature* 56 (3): 920–80.
- CNN. 2020. “A New Virus Stirs up Ancient Hatred,” 2020.
- Cohn, Alain, and Michel André Maréchal. 2016. “Priming in Economics.” *Current Opinion in Psychology* 12: 17–21.
- Dana, Jason, Roberto Weber, and Jason Xi Kuang. 2006. “Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness.” *Economic Theory* 33 (1): 67–80.
- Darley, John M., and Bibb Latane. 1968. “Bystander Intervention in Emergencies: Diffusion of Responsibility.” *Journal of Personality and Social Psychology* 8 (4): 377–83.
- Deming, David J. 2011. “Better Schools, Less Crime?” *The Quarterly Journal of Economics* 126 (4): 2063–2115.
- Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. 2018. “On the Relationship between Cognitive Ability and Risk Preference.” *Journal of Economic Literature* 32 (2): 115–34.
- Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, J. Schupp, and G. Wagner. 2011. “Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences.” *Journal of the European Economic Association* 9 (3): 522–50.
- Doob, Leonard W, Neal E Miller, Orval Hobart Mowrer, Robert Sears, and John Dollard. 1939. *Frustration and Aggression*. New Haven, CT: Yale University Press.
- Duflo, E., M. Kremer, and J. Robinson. 2011. “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya.” *American Economic Review* 101 (6): 2350–90.
- Duncan, Lesley A, and Mark Schaller. 2009. “Prejudicial Attitudes Toward Older Adults May Be Exaggerated When People Feel Vulnerable to Infectious Disease : Evidence and Implications.” *Analyses Of Social Issues and Public Policy* 9 (1): 97–115.
- Durlauf, Steven N. 1999. “The Case ‘against’ Social Capital.” *FOCUS Newsletter for the Institute for Research on Poverty* 20 (3): 1–5.

- Esteban, Joan-Maria, and Debraj Ray. 2008. "On the Saliency of Ethnic Conflict." *American Economic Review* 98 (5): 2185–2202.
- Falk, Armin. 2020. "Facing Yourself – A Note on Self-Image." *Journal of Economic Behavior & Organization*, no. forthcoming.
- Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. "Global Evidence on Economic Preferences." *Quarterly Journal of Economics* 133 (4): 1645–92.
- Falk, Armin, and Christian Zehnder. 2013. "A City-Wide Experiment on Trust Discrimination." *Journal of Public Economics* 100 (April): 15–27.
- Fearon, James D., and David D Laitin. 2000. "Violence and the Social Construction of Ethnic Identity." *International Organization* 54 (4): 845–77.
- Fehr, Ernst, and Urs Fischbacher. 2004. "Third-Party Punishment and Social Norms." *Evolution and Human Behavior* 25 (2): 63–87.
- Fehr, Ernst, Karla Hoff, and Mayursh Kshetramade. 2008. "Spite and Development." *American Economic Review: Papers and Proceedings* 98 (2): 494–499.
- Fershtman, Chaim, and Uri Gneezy. 2001. "Discrimination in a Segmented Society: An Experimental Approach." *Quarterly Journal of Economics* 116 (1): 351–377.
- Fudenberg, Drew, and David K Levine. 2006. "A Dual-Self Model of Impulse Control." *American Economic Review* 96 (5): 1449–76.
- . 2012. "Timing and Self-Control." *Econometrica* 80 (1): 1–42.
- Ghatak, Maitreesh. 2015. "Theories of Poverty Traps and Anti-Poverty Policies." *World Bank Economic Review* 29 (August): S77–105.
- Gintis, Herbert, Samuel Bowles, Robert T. Boyd, and Ernst Fehr. 2005. *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. MIT Press.
- Girard, Rene. 1979. *Things Hidden Since the Foundation Of the World*. Stanford, CA: Stanford University Press.
- Glick, Peter. 2005. "Choice of Scapegoats." In *On the Nature of Prejudice*, edited by John F Dovidio, Peter Glick, and Laurie Budman, 244–61. Blackwell Publishing.
- Goette, Lorenz, David Huffman, and Stephan Meier. 2006. "The Impact of Group Membership on Cooperation and Norm Enforcement." *American Economic Review* 96 (2): 212–16.
- Grosfeld, Irena, Seyhun Orcan Sakalli, and Ekaterina Zhuravskaya. 2018. "Middleman Minorities and Ethnic Violence: Anti-Jewish Pogroms in the Russian Empire." *Review of Economic Studies* forthcoming.
- Gurven, Michael, Christopher von Rueden, Maxim Massenkov, Hillard Kaplan, and Marino Lero Vie. 2013. "How Universal Is the Big Five? Testing the Five-Factor Model of Personality Variation Among Forager–Farmers in the Bolivian Amazo." *Journal of Personality and Social Psychology* 104 (2): 354–70.
- Haushofer, Johannes, Marie Collins, Giovanna De Giusti, Joseph Muiruri Njoroge, Amos Odero, Cynthia Onyago, James Vancel, Chaning Jang, Kuruvilla Maneesh V, and Conor Hughes. 2014. "A Methodology for Laboratory Experiments in Developing Countries : Examples from the Busara Center." Working Paper.
- Haushofer, Johannes, and Ernst Fehr. 2014. "On the Psychology of Poverty." *Science (New York, N.Y.)* 344 (6186): 862–67.

- Haushofer, Johannes, and Jeremy Shapiro. 2016. "The Short-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya." *Quarterly Journal of Economics*, mimeo, 131 (4): 1973–2042.
- Henrich, J., S.J. Heine, A. Norenzayan, and others. 2010. "The Weirdest People in the World." *Behavioral and Brain Sciences* 33 (2–3): 61–83.
- Herrmann, Benedikt, Christian Thoni, and Simon Gächter. 2008. "Antisocial Punishment across Societies." *Science* 319 (5868): 1362–67.
- Kant, Immanuel. 1965. *The Metaphysical Elements of Justice*. New York: Bobbs-Merill Company.
- Kaur, Supreet, Sendhil Mullainathan, Frank Schilbach, and Suanna Oh. 2019. "Does Financial Strain Lower Worker Productivity?" Mimeo.
- Kim, Heejung S, David K Sherman, and John A Updegraff. 2016. "Fear of Ebola : The Influence of Collectivism on Xenophobic Threat Responses." *Psychological Science* 27 (7): 935–44.
- Kosse, Fabian, Thomas Deckers, Pia Pinger, Hannah Schildberg-Hoerisch, and Armin Falk. 2020. "The Formation of Prosociality: Causal Evidence on the Role of the Social Environment." *Journal of Political Economy* 128 (2): 434–67.
- Kranton, Rachel, Matthew Pease, Seth Sanders, and Scott Huettel. 2020. "Deconstructing Bias in Social Preferences Reveals Groupy and Not-Groupy Behavior." *Proceedings of the National Academy of Sciences of the United States of America* 117 (35): 21185–93.
- Kremer, Michael, Gautam Rao, and Frank Schilbach. 2019. "Behavioral Development Economics." In *Handbook of Behavioral Economics*, edited by Douglas Bernheim, Steffano DellaVigna, and David Laibson, 2:345–458. Elsevier B.V.
- Kugler, Tamar, Gary Bornstein, Martin G. Kocher, and Matthias Sutter. 2007. "Trust between Individuals and Groups: Groups Are Less Trusting than Individuals but Just as Trustworthy." *Journal of Economic Psychology* 28 (6): 646–57.
- Kugler, Tamar, Edgar E Kausel, and Martin G. Kocher. 2012. "Are Groups More Rational than Individuals? A Review of Interactive Decision Making in Groups." *Cognitive Science* 3 (4): 471–82.
- Laajaj, Rachid, Karen Macours, Daniel Alejandro Pinzon Hernandez, Omar Arias, Samuel Gosling, Jeff Potter, Marta Rubio-Codina, and Renos Vakis. 2019. "Challenges to Capture the Big Five Personality Traits in Non-WEIRD Population." *Scientific Advances* 5 (7).
- Lerner, Jennifer S, Ye Li, and Elke U Weber. 2013. "The Financial Costs of Sadness." *Psychological Science* 24 (1): 72–79.
- Lerner, Jennifer S, Deborah a Small, George Loewenstein, Jennifer S Lerner, Deborah a Small, and George Loewenstein. 2004. "Heartstrings and Pursestrings: Carryover Effects of Emotions on Economic Decisions." *Psychological Science* 15 (5): 337–41.
- Lichand, Guilherme, and Anandi Mani. 2019. "Cognitive Droughts." Mimeo.
- Lochner, Lance, and Enrico Moretti. 2004. "The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports." *American Economic Review* 94 (1): 155–89.
- Luhan, Wolfgang J., Martin G. Kocher, and Matthias Sutter. 2009. "Group Polarization in the Team Dictator Game Reconsidered." *Experimental Economics* 12 (1): 26–41.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." *Science (New York, N.Y.)* 341 (6149): 976–80.
- Maniadis, Zacharias, Fabio Tufano, and John A List. 2014. "One Swallow Doesn't Make a Summer :

- New Evidence on Anchoring Effects.” *American Economic Review* 104 (1): 277–90.
- Mel, Suresh De, David McKenzie, and Christopher Woodruff. 2008. “Returns to Capital in Microenterprises: Evidence from a Field Experiment.” *The Quarterly Journal of Economics* 123 (4): 1329–1372.
- Miguel, Edward. 2005. “Poverty and Witch Killing.” *Review of Economic Studies* 72 (4): 1153–72.
- Miguel, Edward, Satyanath Shanker, and Ernest Sergenti. 2004. “Economic Shocks and Civil Conflict : An Instrumental Variables Approach.” *Journal of Political Economy* 112 (4): 725–55.
- Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why Having Too Little Means so Much*. Time Books.
- Muraven, Mark, and Roy F Baumeister. 2000. “Self-Regulation and Depletion of Limited Resources : Does Self-Control Resemble a Muscle ?” *Psychological Bulletin* 126 (2): 247–59.
- Murphy, Ryan O., and Kurt A. Ackermann. 2014. “Social Value Orientation : Theoretical and Measurement Issues in the Study of Social Preferences.” *Personality and Social Psychology Review* 18 (1): 13–41.
- Murray, Damian R., and Mark Schaller. 2016. “The Behavioral Immune System.” *Advances in Experimental Social Psychology* 53: 75–129.
- O’Shea, Brian A, Derrick G Watson, Gordon D A Brown, and Corey L Fincher. 2020. “Infectious Disease Prevalence , Not Race Exposure , Predicts Both Implicit and Explicit Racial Prejudice Across the United States.” *Social Psychological and Personality Science* 11 (3): 345–55.
- Publius. 1948. “Federalist No. 55.” In *The Federalist*, edited by M. Beloff, 283–287. Oxford, England: Blackwell (Original work published 1788).
- Rehavi, M. Marit, and Sonja B. Starr. 2014. “Racial Disparity in Federal Criminal Charging and Its Sentencing Consequences.” *Journal of Political Economy* 122 (6): 1320–54.
- Schilbach, Frank. 2019. “Alcohol and - Self-Control : A Field Experiment in India.” *American Economic Review* 109 (4): 1290–1322.
- Shayo, Moses, and Asaf Zussman. 2011. “Judicial Ingroup Bias in the Shadow of Terrorism.” *Quarterly Journal of Economics* 126 (3): 1447–84.
- Staub, Ervin. 1992. *The Roots of Evil: The Origins of Genocide and Other Group Violence*. Cambridge, MA: Cambridge University Press.
- Sutter, Matthias, Claudia Zoller, and Daniela Glätzle-rützler. 2019. “Economic Behavior of Children and Adolescents – A First Survey of Experimental Economics Results.” *European Economic Review* 111: 98–121.
- Tajfel, Henry. 1981. *Human Groups and Social Categories: Studies in Social Psychology*. Cambridge: Cambridge University Press.
- Tanaka, T., C. F Camerer, and Q. Nguyen. 2010. “Risk and Time Preferences: Experimental and Household Survey Data from Vietnam.” *American Economic Review* 100 (1): 557–71.
- United Nations News*. 2020. “COVID-19 Stoking Xenophobia, Hate and Exclusion, Minority Rights Expert Warns,” 2020.
- Wildschut, Tim, Brad Pinter, Jack L Vevea, Chester a Insko, and John Schopler. 2003. “Beyond the Group Mind: A Quantitative Review of the Interindividual-Intergroup Discontinuity Effect.” *Psychological Bulletin* 129 (5): 698–722.