PREPRINT: AUGUST 2021

To appear in *Proceedings of the Royal Society B*

We Predict a Riot: Inequity, Relative Deprivation and Collective Destruction in the Lab

Guillaume Dezechache^{a,b,&}, James M. Allen^{a,&},
Jorina von Zimmermann^a & Daniel C. Richardson^{a,*}

^a*Department of Experimental Psychology, University College London, London, United Kingdom;*

^b*Université Clermont Auvergne, CNRS, LAPSCO, Clermont-Ferrand, France.*

*Corresponding Author:

Daniel C. Richardson dcr@eyethink.org

&Equal contribution

ABSTRACT

Riots are unpredictable and dangerous. Our understanding of the factors that cause riots are based on correlational observations of population data, or post hoc introspection of individuals. To complement these accounts, we developed innovative experimental techniques, investigated the psychological factors of rioting, and explored their consequences with agent-based simulations. We created a game, 'Parklife', that physically co-present participants played using smartphones. In two teams, participants tapped on their screen to grow trees and flowerbeds on separate but adjacent virtual parks. Participants could also tap to vandalise the other team's park. In some conditions, we surreptitiously introduced inequity between the teams so that one (the disadvantaged team) had to tap more for each reward. The experience of inequity caused the disadvantaged team to engage in more destruction, and to report higher relative deprivation and frustration. Agent-based models suggested that acts of destruction were driven by the interaction between individual level of frustration and the team's behaviour. Our results provide insights into the psychological mechanisms underlying collective action.

Keywords: riots, relative deprivation, social identification, collective action, Parklife

INTRODUCTION

Riots – defined as a ‘violent demonstration or clash of more than 100 citizens involving the use of physical force’ [1] – have long been a central topic in the social sciences and public debates due to their societal consequences. In the London riots of 2011, five people died, many more were injured, and property damage worth more than £200 million [2].

Why do people engage in riots? Politicians may caricature rioters as criminal-minded individuals [3], but scientific investigation has shown the reasons to be varied. Perhaps the leading psychological explanation for riots has been Relative Deprivation Theory (RDT) [4]. RDT stated that when people perceive a difference between what they have and what they believe they deserve, they feel relatively deprived. As the perceived disparity grows, so do frustration and resentment, increasing the likelihood for engagement in collective violence. This theory has contributed to explaining puzzling cases of collective violence, such as the bread riots of the 18th century, in which people may not have suffered from starvation, but engaged in collective violence, possibly because what drove them to act was the relative (i.e., as compared to other groups in the society) and not absolute level of deprivation [4].

Though popular, RDT has been subjected to much criticism [5,6]. Sociological work has failed to reveal a clear association between deprivation indices (such as economic deprivation) and riots [6,7]. For example, economically-deprived neighbourhoods were not more riotous than others during the 1960s racial disorders in the US [8]. Economic deprivation is not the same as relative deprivation (the latter is subjective and not necessarily linked to economic hardship), but the idea that relative deprivation cannot explain the emergence of riots has since been a major assumption in the field.

These criticisms, however, are themselves limited by a conceptual flaw and the practical problem of directly measuring relative deprivation and riot participation. Past tests used imperfect proxies for relative deprivation, such as aggregate (rather than individual) and objective (rather than subjective) measures of deprivation [10]. Imperfect proxies for rioting were also used, such as occurrences of disorders in a given neighbourhood, a measure which cannot be linked to specific individuals whose level of relative deprivation is known. People’s reported willingness to participate in a riot was measured, rather than actual participation (e.g., [11,12]). Even though there is evidence that relative deprivation (and particularly, group or fraternal relative deprivation or the extent to which one feels she belongs to a deprived group [13]) can be associated with willingness to join protests [14] or endorse violent actions [15], it has not been shown that it causes people to *engage in collective destruction*.

In recent literature, RDT and its relationship with hostile aggression has sparked new empirical study [13]. For example, Greitemeyer & Sagioglou [16,17] found that participants told that they were of lower socio-economic status (SES) behaved more aggressively than participants told they were of higher SES. At the group level, only one study has – to our knowledge – directly investigated the relationship between the experience of inequity and collective hostile aggression. Abbink, Masclet & Mirza [18] showed that small groups of participants treated with inequity coordinated to deprive the favoured groups of their earnings. In this case however, the outcome of the violent behaviour is instrumental (i.e., to reduce others’ earnings) and may not correspond to core motivations for hostile aggression. Additionally, feelings of relative

deprivation were not measured. This is needed to draw conclusions about any relationship between relative deprivation and destructive collective action. As such, a direct test of the association between the subjective feeling of relative deprivation and actual participation in collective violence is, to our knowledge, missing.

Besides relative deprivation and associated feelings of frustration, social identification (i.e., how much we feel we are part of the same group as others) plays a major role in the emergence of riots. Riots emerge through changes in social identification, beliefs that there is a shared problem and that collective action can prove efficient to change things [3,19,20].

In this work, we sought to experimentally examine the causal role of relative deprivation in the emergence of riot-like phenomena, as well as to measure the importance of social identification in the making of collective action. We developed a lab-based experiment to identify causal factors, and modelled their interactions in agent-based simulations.

The experiment used a group interactive computer game, Parklife, that large numbers of physically co-present participants can play simultaneously using smartphones or tablets. In contrast to many games used in psychology, Parklife can be played in the physical presence of other participants. Parklife resembles popular ‘time management’ apps in which players expend effort to develop resources, and are rewarded by a growing world. Participants are randomly placed into two teams. By varying the comparative effort to create “park features”, we induce feelings of relative deprivation in the disadvantaged team. We therefore have two conditions: the equal game with no difference between the teams, and the unequal game where the disadvantaged team must work harder for equal rewards. By allowing (and measuring) vandalism of the other park across teams (advantaged and disadvantaged) and conditions (equal and unequal), we tested the effect of relative deprivation on vandalism.

More specifically, in the experiment, participants are invited to the lab and to sit around a table. They are randomly assigned to two teams, and two empty parks are displayed on a large shared screen, visible to all (Figure 1). By tapping on a circle on their device, participants contribute towards their team’s total work, which is displayed by a rising bar at the side of the park. Once a second, the average number of team members working is calculated (thereby reducing the impact of different team sizes), and added to the team’s bar (visible to the side of each team’s park). It takes 4 units of work per team for the bar to reach the top, at which point a park feature (e.g., a bench or flower bed) is built. If all members of a team are working, it thus takes 4 seconds to build a feature. Participants also have the option of switching their efforts from ‘do’ to ‘undo’. Tapping ‘undo’ contributes towards a second bar on the screen. When the undo bar has filled, a feature in the park of the other team is vandalised, appearing on screen to be broken, and as before it takes 4 units of vandalism to destroy the park of the other team. A third option is available to players, that is ‘to do nothing’. As the aim of the game is to build as many park features as possible, regardless of the state of the other park (as announced by the instructions), switching to ‘undo’ or to not tap at all seem irrational. As the state of the other team’s park has no direct bearing on a team’s success, participants are taking away effort from improving their own team’s park when tapping ‘undo’ or doing nothing. Aggressive acts are operationalized as the individual and collective decisions to tap ‘undo’ rather than ‘do’ or just staying idle.

In each session, participants played two games of Parklife, each 3 minutes long. In one of the games, the two teams had to do equal amounts of work to be rewarded with a park feature (the equal game). In the other game, the unequal game, one of the teams had to tap twice as much to be rewarded with each feature (i.e., the disadvantaged team must now produce 8 units of work to build a feature), thereby producing an inequity of reward between the teams. This inequity was not announced to participants. We reasoned that this structural difference in the game (which participants experienced as they play) would induce feelings of relative deprivation in the disadvantaged team and cause its members to engage in more acts of vandalism against the opposing team as compared to the other (advantaged) team. Note also that the costs of vandalism remain equal in both teams, irrespective of the type of games (unequal or equal). Although less costly than building, the vandalism option remained more costly than staying idle, a low-cost and non-aggressive option, explicitly available to all players.

For players on the disadvantaged team, their efforts to build a park returns fewer rewards. Switching to vandalism might be a way to affect a change in the game with lower effort. But crucially there are different outcomes too. Players' switching to vandalising make the choice to abandon the alternative (the more effortful action of building the park), and instead decide to engage in lower effort, anti-social action against the other team (even though that indirectly harms their own park too). The effort difference is a feature of our paradigm that may parallel real-world situations. The conditions under which participants make that sort of choice (abandoning one type of effortful action to invest in antisocial actions) are precisely what we want to investigate here.

After each game, we privately polled participants about their emotional state, their feelings towards each other, the opposing team, and the game that they played. We counterbalanced several game features between groups such as the order of the equal and unequal games and whether the disadvantaged team was red or blue.

Our experimental paradigm captures a key aspect of riots: they are a complex *emergent* phenomenon. We then employed agent based models to provide insight into complex behaviour that unfolds over time [21]. These models tested a number of hypotheses related to the central mechanisms that drive emergent phenomena in collective behaviour – frustration and relative deprivation, social identification and social norms – and to see how they interact with inequity and relative deprivation produced by our game.

Like many other psychology experiments, the experience of our study is not exactly like experiences in the everyday world. However, we hypothesize that the psychological mechanisms at work are the same. Unlike other studies on relative deprivation and collective violence, our paradigm measures actual, direct and face-to-face collective destructive efforts (rather than intention to join a protest, or individual aggressive behaviour), and unlike other studies of inter-group competition (e.g., [22]), our study allows groups to act in real time in a more realistic scenario. Such collective destruction would be despite incentives to maintain collective constructive efforts, or the effortless option to just 'do nothing'.

Figure 1. Schematic of a Parklife experiment. After being familiarised with the Hive interface ('familiarisation'), participants were placed into teams at random ('team assignment'). They played a Parklife game for 3 minutes (central screen), answered a series of questions ('survey questions'), and then played a second game with a new series of questions.

RESULTS

We ran 19 experimental sessions with a total of 203 participants. Internet connectivity issues caused some data to be excluded, if, for example, the server could not assign the participant to a team, the participant was not active for the whole duration of the session, or if they had to re-connect to the server. This left us with 171 participants across 19 sessions with full data.

In the figures throughout, the disadvantaged team's actions are shown in red and the advantaged team in blue (though in the experiment, team colours were counterbalanced). For the equal games (when the teams are equitably rewarded for their work), the teams are shown in grey.

Figure 2 shows the actions that the teams took during the games, and the consequences they had for the parks, in terms of park features that were built and destroyed. As we predicted, in the so-called 'unequal' games (when one team has to work more than the other to produce a park feature), the disadvantaged teams built fewer park features and vandalised more of the other team's park features.

Figure 2. Averaged time-course of the probability of actions taken by participants (left) and the number of park features affected as a consequence (right) for advantaged and disadvantaged teams. At the top are work actions and features built, on the bottom are vandalise actions and features vandalised.

To test our hypothesis on the link between inequity, relative deprivation and collective destruction, we analysed individual's actions during the course of the games. Our key dependent variable was the vandalism rate: the number of participants' taps to undo (i.e., to vandalise the other team's park) as a proportion of the total number of taps that they made in the game. We analysed the vandalism rate as a function of game equality (equal vs. unequal), team membership (whether the individual was in the disadvantaged vs. advantaged team for the unequal game), and game order (whether they played the unequal game first or second). Figure 3 shows the observed vandalism rates for the advantaged and disadvantaged participants, in equal and unequal games.

Figure 3. Distributions of the observed vandalism rates for the advantaged (top) and disadvantaged (bottom) participants, when they were playing in the equal (grey) and unequal (coloured) games. Mean rates in dotted lines. On the right are distributions of Bayesian estimations of the differences between equal and unequal games, with 95% of distribution shaded. The percentage of this distribution that is greater than zero is known as the Maximum Probability of Effect (MPE), which directly quantifies the probability that the manipulation condition had an effect on behaviour. Here, it shows that there was only strong evidence of a difference between games for the disadvantaged teams.

We used Bayesian mixed models to quantify the evidence that each of our experimental factors influenced participants' vandalism rate. Mixed models are able to account for the effect of individual participants being nested in a particular group, and the Bayesian approach avoids some of the problems associated with null hypothesis testing [23], or the need for a formal 'stopping rule'. Our models employed weakly informative priors that were scaled following the standard procedures (for full details and model specification see below).

In Figure 3, to the right of the observed data, we show the distributions of the estimated differences between our experimental conditions. The percentage of this distribution that is greater than zero is known as the Maximum Probability of Effect (MPE), which directly quantifies the probability that the manipulation condition had an effect on behaviour. We report below MPEs for each of our experimental factors and contrasts within levels. The Bayesian approach favours quantifying the strength of evidence in this way, rather than simply reporting whether or not an (arbitrary) threshold of significance has been passed. Having said that, researchers generally suggest that an MPE of above 90% or 95% can be thought of as 'strong evidence' [24]. In Figure 3, we can see a grey area that corresponds to 95% of the estimate distribution.

When this interval does not cover zero, it can be seen as strong evidence for a difference between conditions. In addition to these Bayesian analyses, we ran frequentist analysis, which produced a corresponding pattern of results (see Supplementary Material).

As predicted, there was an increase in vandalism rates for the disadvantaged teams when they were in the unequal game compared to the game where they were treated equitably (MPE=99.7%). This was not the case for the advantaged teams, where there was no evidence of a difference between game types (MPE=62.1%). There was evidence of a main effect of a higher amount of vandalism overall in the unequal games compared to the equal ones (MPE=96.3%), but no evidence of a higher amount of vandalism between the teams across all game types (MPE=52.2%). There was no evidence for more vandalism when unequal came first (MPE=70.4%), and no other factors had a significant impact on the proportion of vandalism across conditions.

As we were interested in how the experience of inequity changed participants' experience, they privately reported how they felt about the game and each other after both the equal and unequal game, on a number of custom-made items (see below Figure 4, left). To quantify their response to inequity, we subtracted their answers following the equal game from the answers following the unequal game, and compared these difference scores between advantaged and disadvantaged teams. We ran a Bayesian mixed model for each item, with the factors of team and game order, nested in experimental group. From these runs, we generated distributions for the estimates of the difference scores (Figure 4, centre) and calculated MPEs for each.

Inequity had a very different effect on participants' explicit ratings depending on whether they were advantaged or disadvantaged. For all questions, MPEs were greater than 99.9%, except the two questions that asked if participants would act pro-socially towards the other team by sharing a reward (MPE=80.8%) or building a bridge between the parks (MPE=90.3%). Critically, the disadvantaged teams felt that they deserved more park features compared to what they have got and that their team was unfairly disadvantaged, suggesting they experienced a shared fate with others. They also reported higher frustration, anger and resentment in unequal vs. equal games as compared to the advantaged teams. Interestingly, items linked to social identification (whether people felt their team demonstrated its strength and identification with the team) were higher in advantaged than disadvantaged teams.

Figure 4. Responses to post-game survey items in the equal (circle) and unequal (triangle) games. For example, the first item shows that (compared to the equal game) in the unequal game the disadvantaged teams were less ‘happy’, whereas the advantaged teams were more ‘happy’. To the right are Bayesian estimates of differences between advantaged and disadvantaged teams in the size and direction of these shifts in response inequity. In the first item, for example, 95% of this distribution (shown in shaded area) does not include zero, and so there is strong evidence that there is a reliable difference between teams.

COMPUTATIONAL MODEL

Our experimental work identified in the unequal game a set of causal factors that produced acts of collective destruction. To understand how these factors may interact, we used computational modelling. As individuals can interact with a number of others across different teams in continuous time, we build an Agent Based Model (ABM) [25,26]. The model's goal was to evaluate two key hypotheses: Individuals become frustrated through comparisons with better placed individuals or groups; Also, and social identification and norms of behaviour play a role in escalating cycles of conflict. Our ABM was designed to quantify – in the language of Parklife – whether individuals vandalise at random, or whether it is due to social comparison and frustration, the creation of norms in each team, or some combination of factors.

To summarise the model's conclusions, both relative deprivation (in the form of park differences) and other participant's behaviour is key to the increased proportion of vandalism in the disadvantaged team in the unequal condition (more details may be found in the supplementary material). We find there is a bias towards information from a participant's own team, and that participants engage in coordinated behaviour, distributing vandalism and work across the participants within their own team. Alongside this, we find positive evidence that those in the disadvantaged condition are not behaving rationally, i.e., only vandalising through boredom or to make something happen, and instead are responding to both park differences and the behaviour of the other participants within the game.

Our model is designed to mimic Parklife as closely as possible: each individual within the model is placed in a team in either the equal or unequal condition, plays the game for 180 seconds, and can work or vandalise. The model agents have access to the same information as the participants in Parklife, i.e., the number of features in the parks, and the number in each team working or vandalising at any one time. Agents may decide to work or vandalise based on park differences or the behaviour of others. This information is taken from both teams, and which team agents focus on when deciding whether to work or vandalise is biased (the details of which are described shortly). Finally, the output of the model is the proportion of agents on each team vandalising in each second.

We conclude that participants spend the majority of their time focusing on the state and behaviour of their own team. However, over time park differences increase, and so on the minority occasions that cross park comparisons are made, those in the unequal, disadvantaged condition become frustrated, and vandalise. In the full model, our findings suggest that players focus mainly on their own team, and coordinate their behaviour by performing the opposite function of those on their team (i.e. if many team mates are working, they vandalize, and vice versa). Participants balance team behaviour between working and vandalizing, providing evidence of coordinated behaviour across the teams.

To further test the motivations and mechanisms for vandalism in Parklife, we ran two simpler versions of the model: (the *frustration-only* and *asocial* models) in which individuals do not pay attention to others' behaviour. In the *frustration-only* model we remove the importance of social norms. The *asocial* model is designed to test if participants were simply tapping randomly or performing a cost-benefit analysis in keeping park differences to a minimum: if the latter those in the disadvantaged team in the unequal game would choose to vandalise rather than work with a higher probability, as this decision reduces the effort to change the state of a park/reduce park differences.

Comparing our different models using Bayes factors, we find positive evidence for the full model over the asocial model (Bayes factor = 9.4) and we therefore conclude that individuals are behaving by neither tapping at a base rate, or tapping at an increased constant rate in only the unequal, disadvantaged condition in order to keep park differences to a minimum. We found that there is also strong evidence for the full model over one the frustration-only model (Bayes factor = 12.8), therefore showing the importance of social norms and team behaviour in Parklife.

Our use of the term ‘frustration’ is not a commitment to any specific model of anti-social behaviour. Emotional states are hard to identify, and therefore in this work we are using the catch-all term ‘frustration’ for the emotion driving vandalism within Parklife.

Table 1. Descriptor variables for the full model posteriors.

Parameter	MAP	Median	Minimum	Maximum	MPE
Frustration level	379	383	237	995	1
Imitation strength	-1.20	-1.00	-1.73	-0.01	1
Base tapping rate	0.62	0.56	0.32	0.78	1
Team bias	0.74	0.72	-0.94	0.97	0.99

Figure 5. Ten example runs of the model using MAP point estimates of the posteriors for each team and game condition (blue lines are a single simulation, red lines are the empirical data, both averaged over 19 runs).

DISCUSSION

What are the psychological mechanisms that make people engage in violent collective behaviour? One view has been that when people feel they have less than deserved, they experience relative deprivation. This causes frustration and hostile aggression [4]. This explanation was questioned based on a lack of substantial empirical evidence [10,28]. Here, we used a lab based experimental approach to investigate behavioural and psychological responses to inequity, an approach which allowed us to directly measure actual collective destructive behaviour in response to perceived inequity. We found that the experience of being treated with inequity can lead to acts of collective aggression in a disadvantaged group, associated with reports of being unfairly treated together with one's own team. In our experiment, hostile behaviour took the form of damaging another team's park. This behaviour was also detrimental to the individuals themselves, as they were spending time vandalizing the opposition rather than improving their own park, or simply doing nothing. This suggests that these acts of collective destruction were not a cold, purely rational strategy to succeed at the task. Indeed, violent responses were associated with feelings of frustration, deprivation and of being treated unfairly.

Agent-based modelling confirmed this, and added further insight into how these psychological factors interact with additional factors associated with the emergence of riots in the literature, such as social identification and the importance of group norms [19]. Our best fitting model shows that, while frustration is a key explanatory variable for the increase in vandalism in disadvantaged groups, social identification and other's behaviour play an important role. In this model, participants begin focused on their own parks, with low initial frustration. However, on the relatively rare occasions when the disadvantaged group compares themselves to the advantaged group, their frustration and aggression levels spike. But individuals continue to act in relation to their team mates, ensuring that labour is distributed amongst vandalising and working. This interaction between individual frustration and team behaviour captures our experimental data better than frustration alone. The mechanisms demonstrated here may have wider ramifications for the study of collective behaviour, and in particular the effect that through coordinated action in a group setting, division of labour emerges naturally across each team.

One important distinction in the literature has been between individual (or personal) and collective (or fraternal) relative deprivation [13], with the finding that the latter may be necessary for collective action to occur. In Parklife, participants coordinate (as our models show) but it remains open whether they also explicitly understand that they are taking part in a collective action. The responses to the post-game questionnaire suggest they do as, e.g., members of the disadvantaged team felt their team was unfairly disadvantaged in unequal games, with a reduction of this belief in the advantaged team. This should be more directly addressed using additional post-game self-report items.

Of course, tapping a given button in a virtual game is not equivalent to deciding to riot in the real world. It was not the goal of Parklife or our models to accurately simulate a real riot. Although the stakes and context are different, we contend that the same psychological mechanisms that turn frustration into violence are at work. Experimental approaches such as these can help investigate social behaviours that are difficult to study in the real world with precision or control. Our approach offers a number of advantages over other non-social and non-realistic approaches. Participants in Parklife meet with real participants and interact with them during the game; what's more, they are meant to produce a virtual but perceivable object (a park),

rather than allocating virtual resources they may have great difficulties representing. Our approach thus combines the benefits of tightly controlled experimental methods, with the advantages of realism, physical co-presence and interaction.

It was not the goal of Parklife or our models to accurately simulate a real riot. Psychologists have learnt much about group processes, for example, using methods such as the minimal group paradigm. In typical experiments, participants are assigned to a group by an arbitrary or random criteria, and then asked to allocate abstract resources to in- and out-group members [39]. Real life social identities are not formed like this, of course, and in real life, we do not allocate resources to each other like that. Yet, the minimal group paradigm turns on the same psychological *mechanisms* that operate in the real world, and so provides insight into real world behaviour. Similarly, the goal of Parklife is not to recreate the circumstances of a riot, but to create a game that turns on the same mechanisms of social identification and relative deprivation.

Our results add to the evidence that essentialism – the notion that riotous crowds are simply made up of violent people – is an inadequate explanation. Since people were randomly assigned to experimental conditions, collective violence can be produced by the situation alone.

The game mimicked situations in which contributing to building an item is more costly than destroying it (the disadvantaged team had to work twice more to build, but vandalizing was equally costly between the teams). This asymmetry could contribute to the finding that members of the disadvantaged team prefer to vandalize in unequal games. Evidence from both the fact that participants choose to vandalise rather than do nothing, and our modeling led us to argue that the asymmetry between the two teams, introduced by contributing to building an item being more costly than destroying it (the disadvantaged team had to work twice more to build, but vandalizing was equally costly between the teams), results in relative deprivation, and it is this process which results in the increased vandalism. Of course, ruling out all processes through which this may occur is difficult, and any change in reward structure would result in an asymmetry of costs, introducing its own possibilities. This said, it is a necessary step to experimentally test this in the future, systematically and independently varying the relative costs and benefits of building and destroying in a fully factorial design, to test its impact on the observed individual and collective behaviour.

Social identification, social membership and histories are indisputable in the making of real world riots, and riots involve groups that are already socially structured, or at least, circumstances that favour shared identities and social norms [27,29]. Yet, minimal groups of randomly assigned teams have enough shared sense of identity to perceive their group as being treated unfairly and respond with aggression. In future work, we will explore how teams with different social membership and histories will respond differently to the experience of inequity.

Riots are a paradigmatic example of emerging collective behaviour, and as such, may be seen as a form of collective action. Although we have not explicitly differentiated between individual- or group-level relative deprivation ([13]) within Parklife, we observed the emergence of group level behaviour through the interactions of individuals. Although decisions are made on the individual level, information is received on the group-level, as only information on the parks and the total working or vandalising on each team is available. Within a dynamic group game such as Parklife, it is therefore important to consider behaviour across all levels. Through this, our results may have broader importance for group behaviour in other

circumstances, both in humans and in other species. We find that coordinated behaviour emerges from an interaction of competition and cooperation, and that this is true even with relatively weak social identification (teams are only allocated at random, and no existing social identities are considered). When individuals are placed in a group environment with few restrictions, we still observe collective group behaviour. Finally, this work demonstrates the importance of inequity in the emergence of coordinated behaviour.

Unequal allocations have societal consequences [30]. In addition to the economic and practical difficulties caused by poverty, there might be a pervasive *psychological* response to real and perceived inequity. These correlations have been observed historically at the population level, and we understand little of the psychological responses to inequity at the behavioural level. Our experiment and simulations provide evidence for one piece of this puzzle, showing that the experience of social inequity and the behaviour of others make people more likely to engage in acts of collective aggression.

METHODS

Code

The code to reproduce the analyses can be found at:

https://osf.io/kqc3t/?view_only=e90128144fce449b9cd09ef35c83c7e2

Ethics

We obtained ethical approval from the UCL Research Ethics Committee (Approval ID Number: 3828/003).

Subjects

Since we were to employ Bayesian analysis (that does not require a pre-determined sample size or articulated stopping rule), we collected as much data as we could within a specific time window. Our sample size has therefore not been pre-registered. We tested 203 participants (129 females) in 19 groups. Participants were recruited from the SONA system of the University College London. They were between 18 and 55 years old ($M = 20.97$, $SD = 4.57$). They were compensated 5 GBP for their participation, or given course credit. Our goal was to run participants in groups of 10. But since it was challenging to recruit and ensure the attendance of exactly the same number of participants each session, we ran opportunistically with whoever came to each session, resulting in a range of group sizes from 4 to 23 ($M = 9.63$, $SD = 6.28$). Pilot work suggested that within this range group size did not have a systematic effect on vandalism rates. Internet connectivity issues caused some data to be excluded. This left us with 171 participants across 19 sessions with full data for analysis.

Procedure

Before taking part in the experiment, subjects were asked to fill in an online questionnaire. This questionnaire measured a number of psychological traits. Upon arrival, subjects were seated in a room, around a table, and were asked to fill in a participation consent form. On a voluntary basis, participants could be equipped with a wristband to measure physiological indices (data to be reported elsewhere). Participants were instructed to join the website ‘thehive.sc’ with their smartphone or a tablet we provided. On this website, they provided a subject number we assigned to them, as well as basic demographic information (gender and geographical area). Participants were then presented with a dot on their screen

device that they could drag around with a finger. We directed their attention to a central display which showed the dots of all participants moving in real time.

Participants were then randomly assigned to two teams, which was indicated by the dots on their device and on-screen changing colour. When the subjects were in odd numbers, one team received one more player but participants were all told this would not affect the game outcome as work effort is scaled by the system to accommodate different team numbers. They were then instructed that they would play 2 games of Parklife, each lasting about 3 minutes, and would answer some questions afterwards. They were asked not to speak during the game. The interface (Figure 1) was explained to them, and they tried it out in a one-minute practice game while we pointed out the information onscreen, and how the bars indicated their team's current actions (Figure 1). By the end of the practice game we ensured that participants understood how to play.

Participants were asked not to speak during the game. They could see each other and potentially communicate non-verbally. Because teams were assigned randomly, they did not know which other participants were on their team. Explicit communication or planning within a team was difficult, and we did not observe any attempts.

Participants then played two Parklife games – one equal and one unequal (order counterbalanced between experimental sessions), together with whether the red or the blue team were advantaged or disadvantaged. In the unequal game, the amount of presses on the 'do' button necessary to generate a park feature was 2 times higher in one group than the other. This was not announced to the participants. The amount of presses on the 'undo' button to dismantle a feature in the park of the other team remained equal between groups. Between games, participants were randomly reassigned to red and blue teams, in an attempt to reduce carry over effects from one game to the next. Following each game was a set of custom-made survey questions (see Figure 4). Participants indicated their agreement with statements by moving a dot on their screen across a Likert scale (Figure 1). Finally, participants were asked debriefing questions, told the aims of our study, and thanked for their participation.

Statistical analysis

Our mixed models used fixed effects for the participants team (advantaged / disadvantaged), the game equality (equal / unequal) and game order (equal first / unequal first). There were random effects for the experimental group and the participant, with random intercepts. We used R (v. 3.4.3 [31]) and the package rstanarm (v. 2.18 [32]), employing weakly informative priors that were scaled following the standard rstanarm procedure (full priors are reported in the Supplementary Material). From 4000 samples, we generated estimates of the posterior distributions of the model parameter coefficients, which quantify the strength of the evidence that each experimental condition influenced behaviour in a consistent way. Below we report the estimates of the differences between experimental conditions, using the package psycho (v. 0.3.7 [33]).

We fitted a Markov Chain Monte Carlo, details of which can be found in Supplementary Material. Using the formula notation in the R stats package, the full model was specified as:

$$\text{Vandalism rate} \sim \text{team} * \text{game equality} * \text{game order} + (1 \mid \text{group}) + (1 \mid \text{participant})$$

The model had an explanatory power of around 53.29% (Median Absolute Deviance [MAD] = 0.043, 95% Confidence Interval = [0.44, 0.61], adjusted $R^2 = 0.29$). Supplementary Material gives the full parameter estimates of the model with Median, Median Absolute Deviance (MAD), 95% Confidence-Interval (CI-CI+), Maximum Probability of Effect (MPE) and Overlap for each term.

In addition to these Bayesian analyses, we ran frequentist analysis using more conventional mixed models. These produced a corresponding pattern of results and can be seen in the Supplementary Material.

ACKNOWLEDGMENTS & FUNDING

This study has been funded by a grant from the Nuffield Foundation (The Psychological Roots of Societal Self Harm, 42868) awarded to DCR, and a British Academy Newton International Fellowship (NF 171514) and Alumni Follow-on Funding awarded to GD. GD also acknowledges the support received from the Agence Nationale de la Recherche of the French government through the program “Investissements d’Avenir” (16-IDEX-0001 CAP 20-25).

AUTHOR CONTRIBUTIONS

G.D., J.V.Z and D.C.R. designed the study; G.D. and J.V.Z. collected the data; J.M.A. and D.C.R. analysed the data; J.M.A developed the agent-based modelling; G.D., J.M.A. and D.C.R. wrote the first version of the manuscript; All Authors revised the manuscript and approved the final version.

REFERENCES

1. Braha D. 2012 Global civil unrest: contagion, self-organization, and prediction. *PloS One* **7**, e48596.
2. Bencsik P. 2018 The non-financial costs of violent public disturbances: Emotional responses to the 2011 riots in England. *J. Hous. Econ.* **40**, 73–82.
3. Stott C, Drury J. 2017 Contemporary understanding of riots: Classical crowd psychology, ideology and the social identity approach. *Public Underst. Sci.* **26**, 2–14.
4. Gurr TR. 2015 *Why men rebel*. Routledge.
5. Tilly C. 1971 Reviewed Work: Why Men Rebel by Ted Robert Gurr. *J. Soc. Hist.* **4**, 416–420.
6. McPhail C. 1971 Civil disorder participation: A critical examination of recent research. *Am. Sociol. Rev.* , 1058–1073.
7. Miller AH, Bolce LH, Halligan M. 1977 The J-curve theory and the black urban riots: An empirical test of progressive relative deprivation theory. *Am. Polit. Sci. Rev.* **71**, 964–982.
8. Spilerman S. 1976 Structural characteristics of cities and the severity of racial disorders. *Am. Sociol. Rev.* , 771–793.
9. Snyder D, Tilly C. 1972 Hardship and collective violence in France, 1830 to 1960. *Am. Sociol. Rev.* **37**, 520–532.
10. Brush SG. 1996 Dynamics of theory change in the social sciences: Relative deprivation and collective violence. *J. Confl. Resolut.* **40**, 523–545.
11. Muller EN. 1972 A test of a partial theory of potential for political violence. *Am. Polit. Sci. Rev.* **66**, 928–959.
12. Walker L, Mann L. 1987 Unemployment, relative deprivation, and social protest. *Pers. Soc. Psychol. Bull.* **13**, 275–283.
13. Smith HJ, Pettigrew TF, Pippin GM, Bialosiewicz S. 2012 Relative deprivation: A theoretical and meta-analytic review. *Personal. Soc. Psychol. Rev.* **16**, 203–232.
14. Wright SC, Taylor DM, Moghaddam FM. 1990 Responding to membership in a disadvantaged group: From acceptance to collective protest. *J. Pers. Soc. Psychol.* **58**, 994.
15. Obaidi M, Bergh R, Akrami N, Anjum G. 2019 Group-Based Relative Deprivation Explains Endorsement of Extremism Among Western-Born Muslims. *Psychol. Sci.* **30**, 596–605. (doi:10.1177/0956797619834879)
16. Greitemeyer T, Sagioglou C. 2016 Subjective socioeconomic status causes aggression: A test of the theory of social deprivation. *J. Pers. Soc. Psychol.* **111**, 178.
17. Greitemeyer T, Sagioglou C. 2017 Increasing wealth inequality may increase interpersonal hostility: The relationship between personal relative deprivation and aggression. *J. Soc. Psychol.* **157**, 766–776.
18. Abbink K, Masclet D, Mirza D. 2018 Inequality and inter-group conflicts: experimental evidence. *Soc. Choice Welf.* **50**, 387–423.

19. Stott C, Ball R, Drury J, Neville F, Reicher S, Boardman A, Choudhury S. 2018 The evolving normative dimensions of ‘riot’: Towards an elaborated social identity explanation. *Eur. J. Soc. Psychol.* **48**, 834–849.
20. Reicher SD. 1984 The St. Pauls’ riot: An explanation of the limits of crowd action in terms of a social identity model. *Eur. J. Soc. Psychol.* **14**, 1–21.
21. Smith ER, Conrey FR. 2007 Agent-based modeling: A new approach for theory building in social psychology. *Personal. Soc. Psychol. Rev.* **11**, 87–104.
22. Aaldering H, Böhm R. 2019 Parochial Versus Universal Cooperation: Introducing a Novel Economic Game of Within- and Between-Group Interaction: *Soc. Psychol. Personal. Sci.* (doi:10.1177/1948550619841627)
23. Kruschke JK. 2010 What to believe: Bayesian methods for data analysis. *Trends Cogn. Sci.* **14**, 293–300.
24. Sorensen T, Hohenstein S, Vasishth S. 2016 Bayesian linear mixed models using Stan: A tutorial for psychologists, linguists, and cognitive scientists. *Quant. Methods Psychol.* **12**, 175–200. (doi:10.20982/tqmp.12.3.p175)
25. Davies TP, Fry HM, Wilson AG, Bishop SR. 2013 A mathematical model of the London riots and their policing. *Sci. Rep.* **3**, 1303.
26. Epstein JM. 2002 Modeling civil violence: An agent-based computational approach. *Proc. Natl. Acad. Sci.* **99**, 7243–7250.
27. Drury J *et al.* 2020 A social identity model of riot diffusion: From injustice to empowerment in the 2011 London riots. *Eur. J. Soc. Psychol.* **n/a**. (doi:10.1002/ejsp.2650)
28. McPhail C. 1994 The dark side of purpose: Individual and collective violence in riots. *Sociol. Q.* **35**, 1–32.
29. Scacco A. 2009 *Who Riots? Explaining Individual Participation in Ethnic Violence*. Columbia University.
30. Wilkinson RG, Pickett KE. 2009 Income inequality and social dysfunction. *Annu. Rev. Sociol.* **35**, 493–511.
31. R Core Team. 2019 *A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing; 2012.
32. Goodrich B, Gabry J, Ali I, Brilleman S. 2016 *rstanarm: Bayesian applied regression modeling via Stan*. See <https://mc-stan.org/rstanarm>.
33. Makowski D. 2018 The psycho package: an efficient and publishing-oriented workflow for psychological science. *J. Open Source Softw.* **3**, 470.