Anti-social behaviour and economic decision-making: Panel experimental evidence in the wake of COVID-19

Paul Lohmann[1,2], Elisabeth Gsottbauer[3,4], Jing You[5,*], and Andreas Kontoleon[6,2]

[1]*El-Erian Institute of Behavioural Economics and Policy, Judge Business School, University of Cambridge, United Kingdom*

[2]*Centre for Environment, Energy and Natural Resource Governance, Department of Land Economy, University of Cambridge, Cambridge, United Kingdom*

[3]*Institute of Public Finance, University of Innsbruck*

[4]*London School of Economics and Political Science (LSE), Grantham Research Institute on Climate Change and the Environment, United Kingdom*

[5]*School of Agricultural Economics and Rural Development, Renmin University of China, China*

[6]*Department of Land Economy, University of Cambridge, UK*

[*]*Corresponding author: School of Agricultural Economics and Rural Development, Renmin University of China, jing.you@ruc.edu.cn*

Abstract

We systematically examine the acute impact of exposure to a public health crisis on anti-social behaviour and economic decision-making using unique experimental panel data from China, collected just before the outbreak of COVID-19 and immediately after the first wave was overcome. Exploiting plausibly exogenous geographical variation in virus exposure coupled with a dataset of longitudinal experiments, we show that participants who were more intensely exposed to the virus outbreak became more antisocial than those with lower exposure, while other aspects of economic and social preferences remain largely stable. The finding is robust to multiple hypothesis testing and a similar, yet less pronounced pattern emerges when using alternative measures of virus exposure, reflecting societal concern and sentiment, constructed using social media data. The anti-social response is particularly pronounced for individuals who experienced an increase in depression or negative affect, which highlights the important role of psychological health as a potential mechanism through which the virus outbreak affected behaviour.

Keywords: Anti-social Behaviour, Coronavirus, Risk Preferences, Time Preferences, Natural Experiment, Panel Data, Social Media Data

JEL Codes: C93, D64, D81, D91, I18

Acknowledgements: We thank the seminar participants at C-EENRG for very helpful comments. We are indebted to our research assistants for supporting the execution of the experiments. Financial support from the Faculty of Economics and Statistics, University of Innsbruck, the Department of Land Economy, University of Cambridge and Renmin University is gratefully acknowledged.

1 Introduction

The novel coronavirus SARS-CoV-2, causing the infectious disease now known as COVID-19, was first reported in China in the city of Wuhan in December 2019.[1] In less than three months, the new virus spiralled into a national health crisis and global pandemic. Governments have faced unprecedented challenges to mitigate the spread of the virus and in response have imposed extensive policies that limit social contact and have mandated forms of preventative behaviour.

What are the immediate impacts on social behaviour and economic decision-making of such an unprecedented public health crisis? For example, how does direct exposure to COVID-19 affect people's proclivity for acting pro-socially, their attitudes towards taking risks or their patience levels. This question is of particular importance as economic preferences have been shown to be an important predictor of people's willingness to adopt emergency measures at the early critical stages of such a health crisis, including social distancing, hand hygiene and wearing of face masks meant to contain the further spread of the virus (e.g. Nikolov et al., 2020; Campos-Mercade et al., 2021; Müller & Rau, 2021). To manage the pandemic effectively, it is vital to understand factors that drive people's willingness to comply with confinement measures, especially those that could be significantly affected by the exposure to COVID-19 itself. To shed light on this question, we use a unique dataset of longitudinal experiments to examine the effect of exposure to COVID-19 on anti-social behaviour and economic decision-making. The experiments were conducted in October and December 2019 right before the outbreak and shortly after the first wave of the pandemic in March 2020, relying on a large sample of university students from Beijing. Students were all based in Beijing during the baseline survey and were spread across 183 cities in China during endline survey in March 2020. Unlike other studies on the impacts of COVID-19, our panel data enables identification not only to a higher degree of internal validity, as we were able to track the change in behaviour and preferences of the same individuals before and after the outbreak, but also to larger external validity with wider geographic, epidemic and socioeconomic representation. More broadly, our paper speaks to a sizeable body of empirical literature assessing if negative shocks (e.g., violent conflicts, natural disasters, economics crisis) can bring about systematic changes in economic decision-making and affect the temporal stability of economic preferences (for an overview of largely mixed findings see,

[1]The authors are not making an assertion as to the global origin of the specific virus, but simply that within China, Wuhan municipal Center for Disease Control and Prevention (CDC) released the first epidemiological alert of 27 cases on 31 December 2019.

Chuang and Schechter (2015)).[2] A related literature focuses on the stability of preferences in relation to acute stress events and scarcity, again producing mixed results.[3]

We also add to a number of papers exploring the impact of COVID-19 on economic and social preferences, including research focusing on risk and time preferences (Angrisani et al., 2020; Bu et al., 2020; Li, Huang, et al., 2020; Drichoutis & Nayga, 2021; Guenther et al., 2021; Harrison et al., 2022), social preferences (Buso et al., 2020; Grimalda et al., 2021; Branas-Garza et al., 2022) and papers assessing various preference measures (e.g. Alsharawy et al., 2021; Bokern et al., 2021; Shachat et al., 2021). Findings from all of the aforementioned studies on the effect of the outbreak of the pandemic on economic and social preferences are mixed. One reason for these mixed findings might be that these papers cover a broad spectrum of research approaches, amongst others, differing in the use of incentivized and non-incentivized preference measures, sampling among a student or a more general population sample and data collection taking place before and after the outbreak or only after the outbreak of COVID-19. In addition, all of the aforementioned studies were conducted in many different countries and cultural contexts.

The closest of these papers to ours are Shachat et al. (2021), Li, Huang, et al. (2020) and Bu et al. (2020) making use of risk and social experimental preference measures elicited among Chinese samples to assess the impact of the first wave of COVID-19 on the stability of these preferences. While we acknowledge that our experimental approach holds many similarities to these papers, there are also some notable differences with respect to research design and identification strategy including the use of a within- instead of between-subject design, additional survey data to study potential mechanisms and a more nuanced analysis with respect to exposure to the virus outbreak through ample geographical variation in virus prevalence. Note that the main sample difference to Shachat et al. (2021) and Bu et al. (2020) is that both studies heavily draw on students located in Hubei province where the majority of Covid-19 cases were reported, while our study relies on geographical variation in students' location across all of mainland China (only 10 participants, i.e. 1.92% of the sample were located within Hubei province).[4]

[2]See, for instance, Voors et al. (2012), Cassar et al. (2013), Becchetti et al. (2014), Cassar et al. (2014), Fleming et al. (2014), Grosjean (2014), Page et al. (2014), Callen (2015), Cohn et al. (2015), Said et al. (2015), Andrabi and Das (2017), Cassar et al. (2017), Falco and Vieider (2018), Hanaoka et al. (2018), Brown et al. (2019), and Filipski et al. (2019).

[3]See, for instance, Haushofer et al. (2013), Delaney et al. (2014), Prediger et al. (2014), Cahlíková and Cingl (2017), Koppel et al. (2017), Aksoy and Palma (2019), Cahlikova et al. (2019), Cettolin et al. (2019), Fehr et al. (2019), and Kettlewell (2019).

[4]We acknowledge that the majority of Covid cases were concentrated in Wuhan during the first Wave of the pandemic in China. Our findings thus complement those of Shachat et al. (2021) and Bu et al. (2020) by exploiting the wider geographic dispersion in our data. Moreover, comparing only the extremely high prevalence rate in Wuhan against zero otherwise could potentially exacerbate the variance of the "intervention" and, thus, the estimated "treatment" effect of Covid-19. The wider geographic coverage of our sample is able to reduce this upward bias. This point, that empirical regularities in the far tails of the distribution tend to disappear, has been re-illustrated recently by (Hamermesh & Leigh, 2022).

We summarize similarities and differences to these studies in Table A1 in the Appendix.

Our paper also links to research on the connection between economic and social preferences and health behaviours, including the willingness to take protective action and the demand for vaccines (Chapman & Coups, 1999; Sutter et al., 2013; Böhm et al., 2016; Galizzi & Miraldo, 2017). For example, recent research with respect to the spread of COVID-19 finds that pro-social preferences and patience positively correlate with personal protective behaviour related to COVID-19, while risk tolerance negatively impacts the willingness to engage in such behaviour (Campos-Mercade et al., 2021; Müller & Rau, 2021).[5] However, the direction of these effects would become unclear if exposure to an acute public health crisis itself could trigger behaviour and preferences to change in different directions for significant segments of the population. On the one hand, evidence from economics and psychology suggests that people exposed to a major crisis event may more likely display selfish and reckless behaviour (Fisman et al., 2015; Fritsche & Jugert, 2017). Anecdotal evidence from panic buying and increased racial discrimination, xenophobia, and riots in response to the coronavirus outbreak is indicative of such behaviour. On the other hand, research also indicates that disaster and crisis lead individuals to engage in widespread altruism and acts of solidarity (Solnit, 2010; Bauer et al., 2016). With respect to the COVID-19 outbreak, the public's willingness to engage in cooperative behaviour including social distancing as well as the formation of neighbourhood networks to assist vulnerable groups speaks to this strand of literature.

The research described in this paper contributes to the aforementioned literatures in several important ways. First, it adds to the body of empirical work testing the theoretical assumption of stable preferences over time. Our study explores the acute effect of a public health crisis on temporal stability of fundamental preferences predictive of economic and social behaviour, including risk aversion, patience, trust, cooperation, altruism, norm enforcement and anti-social behaviour. More importantly, it also explores potential mechanisms behind any changes in preferences and behaviour.

Second, unlike a number of cross-sectional studies, it applies a 'dose-response' difference-in-difference framework to panel data, tracking the same individuals before and after the virus outbreak. Our identification strategy exploits within-individual variation in economic decision-making and exogenous variation in exposure to the virus across 183 cities to assess causal impacts of the crisis. Importantly, our

[5]Notably, many governments have designed their response to COVID-19 around the premise that people are able and willing to engage in pro-social behaviour. The UK's slogan "Stay Home, Protect the NHS, Save Lives" (and other similar informational campaigns in other countries) directly conveys an underlying appeal to pro-social behaviour. In contrast, the Chinese government has largely drawn on wartime rhetoric to enforce strict containment and lockdown policies.

design resembles a type of natural experiment, whereby the individuals' locations (and thus exposure to the virus) are pre-determined by factors unrelated to the virus outbreak. Compared with between-individual identification, it yields more precise estimates by explicitly controlling for individually heterogeneous confounders (fixed effects) to preferential changes.

Third, it combines data from multiple sources and disciplines. To ascertain incentive compatible economic decision-making, we employ well-established experimental protocols as opposed to responses to purely hypothetical behavioural questions. Together with a range of survey variables from economics and psychology, we can capture both behavioural and trait-like characteristics. In order to comprehensively reflect what COVID-19 means for individuals, we go beyond simple epidemiological measures of virus prevalence and exploit information from big-data extracted from Chinese social media to construct two additional measures of virus exposure capturing social concern and sentiment: (1) an innovative index reflecting public concern/anxiety based on internet search volume sourced from China's largest search engine (Baidu Inc.) and (2) a novel index of expressed negative sentiment based on linguistic text analysis of 523,222 tweets posted on the main microblogging platform (Sina Weibo).

Finally, to better purge non-random components of the exposure to the virus in identification as well as to minimise omitted variable problems in regressions, we collected data from various sources including population mobility based on mobile phone check-ins at Baidu Inc. and official air quality information from 1,436 air monitoring stations across China. We also hand-collected and coded city-level lockdown policies on various aspects of life, work, and education from government sources.

We show a substantial and statistically significant increase in anti-social behaviour for those individuals more intensely exposed to the virus outbreak. In contrast, our measures of pro-social behaviour and economic preferences are largely unaffected by the Covid-19 shock. Moreover, our analysis of potential mechanisms suggests that increases in depression and negative affect are likely driving the observed relationship for anti-social behaviour. The indication that mental well-being is likely responsible for the increase in anti-social behaviour can inform better targeted policies and relief programs, including increased attention to mental health issues at the onset of a public health emergency and in turn greater investment into mental health services (Dong & Bouey, 2020; Liu et al., 2020).

Our paper is structured as follows. Section 2 describes the design of study, detailing our outcome, control and mediation variables. Section 3 presents our identification strategy and how we address potential endogeneity concerns. Section 4 describes our empirical strategy, presents sample statistics and outlines how we address attrition. Section 5 presents our results, the sensitivity analysis undertaken, as well as the potential mechanisms explored. Section 6 concludes with a discussion of our main findings and how these relate to the relevant literature.

2 Study Design

The experiment was conducted on a sample drawn from the general student population of universities in Beijing (with the majority of students enrolled at Renmin University) in October 2019.[6] We informed participants about the longitudinal nature of the study and asked them to consent to participate in multiple experimental survey waves.[7] Students were offered a 10 Yuan (1.50 USD) flat-fee payment for participation in the panel study and the opportunity to obtain bonus payments based on their decisions in the experiments. Note that we follow standards in experimental economics to conduct an incentive-based experiment in which participants obtain a monetary reward based on the results of their decisions.[8]

Experimental protocols and surveys were administered online using a survey tool integrated into WeChat, a popular mobile messaging application in China. The average payment per participant was approximately 32 Yuan (5 USD), including a 10 Yuan show-up fee for each wave). Average completion time per wave varied between 15-20 minutes.

Data was collected in three waves. Wave 1 (N=793) was conducted in October 2019 and designed as a baseline survey including questions on participants' socio-demographics which were not repeated in later waves. Wave 2 (N=650) was conducted in December 2019 and Wave 3 (N=539) in March 2020 which comprised elements of both proceeding surveys as well as questions specific to the COVID-19 crisis. Importantly, Waves 1 and 2 were conducted before the outbreak of COVID-19 in China, while

[6]We acknowledge that it is often criticised that student samples do not accurately represent the overall population. However, there is increasing evidence that student samples are appropriate for studying human social behaviour (Exadaktylos et al., 2013; Falk et al., 2013). In addition, we argue that in the case of COVID-19 response, this is a particularly important demographic. Students are subject to a lower risk of suffering severe medical consequences of infection, however, mitigating the spread of the virus relies heavily on low-risk demographics to follow public health guidelines and engage in social distancing. Hence, we believe that studying students' behaviour is highly relevant in the context of COVID-19.

[7]Note that Waves 1 and 2 of the experiment were initially designed and collected as part of a pre-registered experiment on economic decision-making and air pollution. Details at: https://doi.org/10.1257/rct.4856-1.0

[8]While incentivization is the norm for preference elicitation in experimental economics in order to reduce hypothetical bias, we would like to highlight recent results by Hackethal et al. (2022) which suggest that hypothetical bias is rather limited when eliciting risk preferences in online experiments – a setting very similar to ours. Nonetheless, previous literature provides no clear indication if the results of Hackethal et al. (2022) can be extended to social preference elicitation (Gillis & Hettler, 2007; Engel, 2011; Camerer & Mobbs, 2017).

Wave 3 at a point when the epidemic in China had significantly slowed and new cases were close to zero. Figure 1 displays a detailed timeline of events and highlights the spread of the epidemic, indicating daily new confirmed cases of COVID-19 in China.

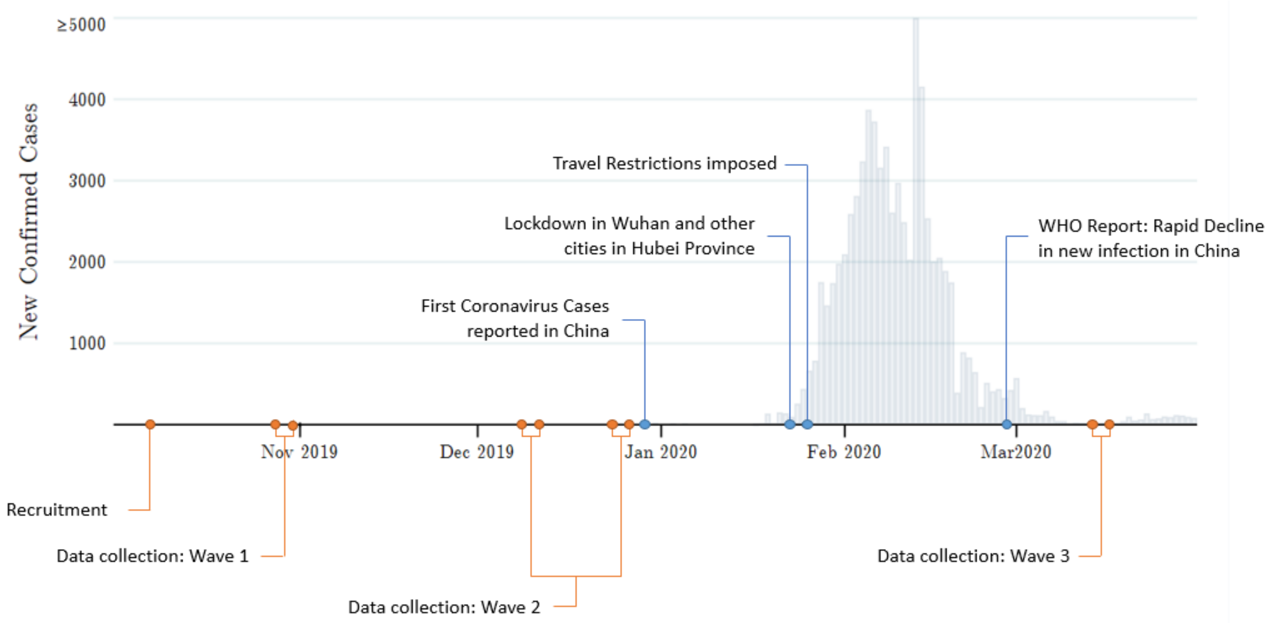

Figure 1: ***Timeline of Data Collection and Virus Outbreak***

Note: Confirmed cases of COVID-19 were obtained from official sources (State Council, provincial governments, and the Chinese CDC)

The experimental modules which were employed across the three waves of the survey consist of well-established experimental games. To measure anti-social behaviour as well as economic preferences on risk and time/discounting (our main outcome variables) we use the following incentivized decision tasks: a joy of destruction game; a take game with and without deterrence; a third-party punishment game; a lottery choice task; an investment game and a convex time budget task. In addition, we use a hypothetical trust game and one-shot public good game to capture other aspects of social preferences (results presented in the Appendix).

The survey waves also included standardized survey modules to obtained relevant socio-economic control variables, but also measures of participant's cognitive functioning and well-being. The latter two sets of measures would serve as potential mechanisms explaining the effects of Covid exposure on our outcome variables. In particular, to measure cognitive functioning, we use a set of Raven's matrices and a five-item self-completion questionnaire to assess participants' momentary level of ego-depletion. For psychological and physiological well-being, we measure self-reported subjective well-being, depressive symptoms, positive and negative affect, sleep quality and general health status of all participants.

Table 1 provides an overview of each survey module. Table A2 in the Appendix provides more detail on the experimental modules and how the outcome and mediating variables were defined. Note that Wave 3 consisted of all experimental modules, while Waves 1 and 2 were made up of sub-sets of these.

To incentivise truthfulness and effort, the majority of tasks were incentivised so that payoff depended on the participant's choices. The incentivized tasks were presented in separate questionnaire parts to participants.[9] In Part I of the questionnaire, participants were presented with the Joy of Destruction Game, the Take Game and the Third-Party Punishment game and subjects were paid on the basis of one randomly selected task from Part I.[10]. The average endowment for each task was approximately 20 Yuan (3 USD). In Part II, participants made a total of 25 decisions across a Lottery Choice Task and a Convex Time Budget Task with significantly increased stakes (payoff between 56 – 140 Yuan, equivalent to 8.60 – 21.50 USD). Participants were informed that 30 students would be selected at random to receive payment for one of their decisions (selected at random) from Part II. In Part III, participants were incentivized to complete the Raven Matrices test, in which we paid participants for each correct answer (out of 9). Note that the trust game and the one-shot public good game were not incentivized in any of the waves.[11]. At the end of each wave, the decision tasks that were used for payment were randomly selected and respondents received their respective payments to their WeChat Wallet on the following day. Note that the time preferences payments were delivered according to the time schedule indicated in the selected decision task (details provided below). All instructions were provided in Chinese and all choices were framed in terms of Chinese Yuan (CNY). The English translation of the instructions is included in Appendix C. In the following subsection we describe the survey modules and key variables used in our analysis in more detail.

[9]Note that participants at no point were provided feedback about the (payoff) outcomes of the different experimental tasks they completed. By doing so, we can exclude learning effects or strategic behaviour of participants, which could potentially bias our results.

[10]We acknowledge that if a task was chosen in which participants were assigned to multiple roles, we applied a 'Pay One' payment procedure, i.e., we randomly draw one role of that task for payment. Such a procedure has been shown to eliminate hedging opportunities and also wealth effects (e.g. Bardsley et al., 2009)

[11]Previous literature indicates that even hypothetical incentives in economic experiments can give accurate results (e.g. Gillis & Hettler, 2007)

Table 1: Panel Survey Modules

Wave	N	Anti-social Behaviour	Risk & Time Preferences	Cognition	Well-being	Pro-social Behaviour
1	793		Lottery Choice Task$, Investment Game		CES-D, General health	Trust Game, Public Good Game
2	650	Joy of Destruction$, Take Game$, Punishment Game$	Convex Time Budget$, Lottery Choice Task$	*Raven*$, Depletion	CES-D, General health, Subjective well-being, PANAS	
3	539	Joy of Destruction$, Take Game$, Punishment Game$	Convex Time Budget$, Lottery choice task$, Investment Game	*Raven*$, Depletion	CES-D, General health, Subjective well-being, PANAS	Trust Game, Public Good Game

Note: Waves 1 & 2 were collected before the COVID-19 outbreak while Wave 3 was collected after. Tasks marked with $ were incentivised. CES-D = Centre for Epidemiologic Studies Depression Scale; PANAS = Positive and Negative Affect Schedule.

2.1 Outcome variables

2.1.1 Anti-social behaviour

The anti-social behaviour module consists of two separate incentivized games to elicit different dimensions of people' willingness to engage in anti-social behaviour. The binary Joy of Destruction (JOD) game provides a measure of nasty behaviour (Abbink & Herrmann, 2011). In this two-player game, participants were anonymously matched in pairs (each with an initial endowment of 20 Yuan) and then faced the binary decision whether to destroy their assigned partner's endowment by half at a cost of 2 Yuan or maintain the status quo. Participants were further informed that, with a one third probability, the other player's endowment will be reduced to 10 Yuan, regardless of their decision. The design of JOD game removes all conventional motivations for anti-social behaviour and further allows destructive behaviour to be partially hidden behind a component of random destruction. The primary outcome variable from this task is a binary indicator identifying individuals that chose to destroy their counterpart's endowment. The Take Game provides a measure of covert anti-social behaviour in the form of stealing or theft (Schildberg-Hörisch & Strassmair, 2012). In this two-player game, participants were anonymously matched and provided unequal endowment (of 10 Yuan or 18 Yuan). The participants then had to decide whether to take from the other player's initial endowment in two different scenarios. In the first scenario, the player could take any amount (between 0 and 18 Yuan) without facing any consequences. In the second scenario, the player could take any amount but faced a 60% probability of being detected, effectively reducing their payoff to 6 Yuan due to a penalty. Note that the game was constructed as such that no losses or negative payments were possible. From this task we obtain two primary outcome variables for our analysis: (1) a continuous measure of taking without and (2) with a risk of being detected.[12]

2.1.2 Risk and time preferences

Risk preferences are obtained using a standard incentivised Lottery Choice Task Eckel and Grossman (2002). In this task, participants had to decide between six lotteries each with a 50% chance of paying a lower or higher amount. Lotteries were increasing in variance, total pay-off and riskiness. Find instructions in the Appendix C4, Part 4.2. Based on the chosen lottery, we obtained the participant's constant relative risk aversion (CRRA) parameter interval. For our analysis, we calculate the CRRA interval midpoint for each participant in a given survey wave, with a higher value indicating greater risk aversion.[13] In an additional task, we obtain a simple measure of risk aversion with the help of a non-incentivized Investment Game based on Gneezy and Potters (1997). Participants could invest part of their hypothetical 20,000 Yuan endowment into a lottery with a winning probability of 50%. The higher the investment, the higher the risk participants are willing to take.

Time preferences are elicited using Convex Time Budgets (CTB) following Andreoni et al. (2015). Participants made 24 consecutive decisions between sooner or later payments, across four different timeframes, with six budget lines for each timeframe. Participants thus faced decisions over payment "today and 5 weeks from today", "today and 9 weeks from today", "5 weeks from today and 10 weeks from today" and "5 weeks from today and 14 weeks from today". The 24 budget lines and instructions are displayed in Appendix C4, Part 4.3. Prior to our main analysis, we estimated the individual-level parameters beta and delta parameters via non-linear least squares following Andreoni et al. (2015). For our main analysis, we utilize the individual-level delta parameter as a measure of patience and construct a binary measure of present bias equal to one if a participant's individual-level beta parameter is smaller than 1.

[12]For more details find the English translation of the instructions in Appendix C.

[13]To calculate the mid-points, we first replaced the infinity value for the lower bound with -1 and for the upper bound with 10.

2.2 Control and Mediation Variables

We collect extensive socio-demographic control variables relevant to the Chinese context including, participants urban or rural origin (or "Hukou" status) and whether participants have siblings ("only child" due to family planning). In addition, we use a set of survey questions to assess cognitive functioning and well-being, which may serve as potential mechanisms for changes in economic decision-making.

To measure cognitive functioning, we used a subset of Raven's Standard Progressive Matrices (Bilker et al., 2012) and pay subjects for each correct answer. Cognitive performance was indexed by the sum of correctly solved matrices (range 0-9). In addition, we assessed participants' momentary and self-reported state of ego-depletion, which reflects an individual's self-control capacity at a given moment, according to ego-depletion theory (Baumeister et al., 1998).[14] This measure was obtained from a modified 5-item Depletion Scale adapted from Twenge et al. (2004) where a higher score indicates higher levels of depletion.[15]

The well-being module consists of a selection of survey questions to capture different dimensions of well-being. We use the Center for Epidemiologic Studies Depression Scale (CESD) 10-item scale as a validated self-reported instrument to measure the prevalence of depressive symptoms (Andresen et al., 1994). Respondents are asked to report the frequency at which they experienced a given mood or symptom during the past week on a four-point scale, ranging from zero ("none of the time") to three ("most of the time"). A depression score is obtained by totalling responses to each of the 10 items (range 0 – 30). Moreover, we construct a binary measure indicating the presence of depressive symptoms for subjects with a depression score of 10 or higher (Andresen et al., 1994). To measure short-term mood on the day of the survey, we use the international short form of the Positive and Negative Affect Schedule (PANAS-ISF) consisting of a 10-item self-reported questionnaire (Thompson, 2007). Using the respective negative and positive affect items (five each), we construct scores for positive and negative affect, where higher scores indicate greater presence of positive or negative mood on the day of the survey (range 5-25). To measure subjective well-being, we focus on three dimensions including life satisfaction, happiness, and meaningfulness of life, where higher scores indicate higher levels of subjective well-being in the

[14]While the theory of ego-depletion is very prevalent in the psychology literature, more recently, it has also attracted attention in economics and there is a growing number of studies which have assessed the impact of self-control depletion on economic preferences (Achtziger et al., 2016; Gerhardt et al., 2017; Achtziger et al., 2018).

[15]The following five items were used: "I feel drained", "I feel calm and rational", "I feel lazy", "I feel sharp and focused" and "I feel like my willpower is gone". Responses were given on a 5-point Likert scale ranging from 1 "not true" to 5 "very true".

respective categories. We also check the general health status of our participants, using responses to a question on their general health condition ranging from 1 to 5, indicating very poor to very good health status on the day of the survey. See Appendix C4-C5, for instructions.

Finally, we also assess participant's pro-social behaviour with some simple hypothetical tasks: A standard one-shot Public Good Game was used to obtain a measure of cooperation (see Appendix 2, Part C2, for instructions. In this game, participants could invest part of their hypothetical endowment (10,000 Yuan) into the production of a public good with a return of 1.6. We also used a Trust Game to obtain a measure of trust in an investment setting (Berg et al., 1995). Participants chose how much of their hypothetical endowment (100 Yuan) to invest into a partner who doubles the investment and decides how much to return. Finally, we used an incentivized Third-Party Punishment Game (Fehr & Fischbacher, 2004) to measure both prosocial behaviour and third-party sanctioning behaviour for violations of a distribution norm (Fehr & Fischbacher, 2004). As in a classic dictator game, players first decided whether to transfer between 0 and 10 Yuan of their 20 Yuan endowment to an anonymously matched recipient. They then took the role of a third party observing another player's transfer decision, with the option to enact costly punishment for each possible transfer amount sent by the observed dictator. In our setting, the third-party observer had an endowment of 10 Yuan and could use any of this amount to punish the dictator by reducing their endowment by a factor of three (e.g., 2 Yuan would reduce the dictator's endowment by 6 Yuan). As players faced multiple decisions, they were informed that one of their choices would be randomly selected for payment. From this game we construct three primary outcomes for our analysis: (1) an incentivised measure of (observed) giving from the dictator game, as a measure of pro-social behaviour or altruism, (2) the amount spent to punish if the observed dictator transfers zero and (3) a binary indicator identifying subjects that were willing to pay any amount to punish a dictator that gave zero.[16]

[16] For our analysis, we selected to explore sanctioning behaviour for the most unequal distribution (i.e., dictator giving nothing to the recipient). We also observe punishment decisions for each of the alternative transfer amounts (2, 4, 6, 8 and 10 Yuan) and now provide a descriptive overview Appendix Figure A1. For robustness, we additionally explore sanctioning behaviour to enforce a 50/50 distribution norm based on the amount participants were willing to punish if the dictator transferred half of their endowment (i.e., 10 Yuan). Our results are robust to this analysis.

Table 2: Summary Statistics

	Mean	SD	Min	Max	N
Anti-social Behaviour					
Joy of Destruction (Destroy = 1)	0.16	0.37	0.00	1.00	1044
Taking (¥)	10.23	6.30	0.00	18.00	1044
Taking with Deterrence (¥)	9.48	6.73	0.00	18.00	1044
Risk & Time Preferences					
Risk Aversion (CRRA midpoint - EG)	2.99	2.78	-0.50	6.73	1566
Risk Taking (GP)	7156.81	4086.02	0.00	20000.00	1044
Present Bias (Yes = 1)	0.67	0.47	0.00	1.00	1026
Patience (δ parameter)	0.98	0.12	0.00	1.00	1026
Pro-social Behaviour & Norm-enforcement					
Cooperation (¥ invested in PGG)	4138.33	3370.25	0.00	10000.00	1044
Trust (¥ invested)	43.55	25.29	0.00	100.00	1044
Dictator Giving (¥)	3.91	3.44	0.00	10.00	1044
Punishment (¥)	0.59	0.49	0.00	1.00	1044
Punish (Punish = 1)	2.08	2.29	0.00	10.00	1044
Cognition and Health					
Cognitive Ability (correct puzzles)	6.70	1.39	1.00	9.00	1044
Depletion (score)	1.68	3.54	-7.00	11.00	1044
Depression (score)	10.60	5.67	0.00	29.00	1566
Depressive Symptomns (Yes = 1)	0.53	0.50	0.00	1.00	1566
Negative Affect (score)	9.81	4.21	5.00	25.00	1044
Positive Affect (score)	12.90	3.41	5.00	21.00	1044
General Health (scale)	3.81	0.80	1.00	5.00	1566
Socio-demographic Characteristics					
Age	19.85	1.53	17.00	29.00	522
Female (%)	0.82	0.39	0.00	1.00	522
Rural Hukou (%)	0.21	0.41	0.00	1.00	522
Only Child (%)	0.64	0.48	0.00	1.00	522
Economics/Finance Major (%)	0.44	0.50	0.00	1.00	522
Year of Study	2.56	1.16	1.00	6.00	522

Note: Table displays the summary statistics for the full sample by pooling data from all three waves (where applicable).

2.3 Summary Statistics

Table 2 presents summary statistics for all outcome variables employed in the analysis, as well as the socio-demographic characteristics of the full sample pooling responses from all waves (if N=1566 the variable was collected in all three waves, if N=1044 the variable was collected in only two of three waves).[17] Most notably, with respect to our main outcome variables related to anti-social behaviour we observe that 16% of participants decided to destroy their counterpart's endowment, and 9.48 (10.23) Yuan were taken, on average, in the Take Game with (and without) deterrence. Both risk measures suggest that the sample was slightly risk averse, and 67% of participants were classified as present biased.

The 'Cognition & Health' section includes potential mechanism influencing decision-making including (i) cognitive ability and depletion, (ii) emotional affect and depressive symptoms and (iii) general health. We also obtained respondents 'Socio-demographic Characteristics' from the baseline demographic survey (N=522) which was conducted in October 2019.

3 Identification

We exploit geographical variation in virus prevalence to estimate the causal effect of virus exposure on economic decision-making (similar to Bu et al. (2020)). Although initial recruitment (prior to the COIVD-19 crisis) took place at Beijing universities in October 2019, students were geographically (and exogenously) dispersed across the country by the time of the 3rd Wave of data collection (14$^{\text{th}}$ – 17$^{\text{th}}$ March). We also collected students' travel history from the end of their academic term to our 3$^{\text{rd}}$ survey wave date. Section 3.2 below provide an in-depth discussion on possible threats to our identification strategy and how these were addressed.

At the time of Wave 3 data collection, 73.4% of participants had travelled outside Beijing and returned to their respective hometowns or family homes to celebrate the Spring Festival (25$^{\text{th}}$ January), the most important national holiday in China. For generations, it has been a very strongly and widely adhered tradition to celebrate Chinese New Year with one's family. Note that the locations where students originate

[17] We acknowledge that a greater proportion of participants in our sample are female (80%). This is due to the fact that around 70% of our sample come from liberal arts colleges/universities where the share of female students constitutes on average more than 60% with the highest share being at 98%. However, related research on preferences and the Covid-19 pandemic in China does not suggest a significant heterogeneity in results by gender Bu et al. (2020). The gender bias of our sample should therefore not preclude a valid interpretation of the presented findings.

from are geographically diverse due to the college admission rule that has been implemented since 1978. The Ministry of Education and provincial governments jointly set up regional admission quotas according to not only provincial socioeconomic and demographic characteristics but also universities' classifications, for the purpose of equal access to higher education across regions and ethnic groups. After the national college entrance exam in every June, the colleges/universities will announce their subject- and province-specific admission quotas according to the general guidelines of the Ministry and the provincial governments. Students submit their applications to colleges/universities according to these quotas and their predicted exam performance during June and August. The academic year starts in early September. Thus, the dispersion of our participants' hometowns has been pre-determined by factors unrelated to those accounting for the distribution of Covid-19 prevalence.[18] Shortly after the Spring Festival, nationwide travel restrictions were imposed, and the university spring term was postponed indefinitely. Effectively, the Ministry of Education restrained all students at the cities and towns they were located in late January 2020. This means, participants in our sample were located in 183 cities across China when the endline survey took place, with varying degrees of virus prevalence when we fielded our 3rd survey Wave.

Importantly (and what uniquely benefits our identification) is that participants' geographic dispersion throughout the virus outbreak was totally unrelated to the COVID-19 crisis nor to different levels of COVID-19 exposure. Hence, our data has characteristics of a natural experiment in that treatment assignment (or in our case students' exposure to different degrees of COVID-19) has been largely determined by exogenous distribution of their geographic locations as a result of the pre-determined universities' admission across regions coupled with the government imposed domestic travel ban. Figure 2 provides a graphical illustration of the locations of our participants and corresponding city-level virus prevalence at the time of the third survey.

[18] Lu et al. (2018) have utilised this exogenously determined admission rule to study the impact of students' experiences in competitive college admissions on their risk preferences.

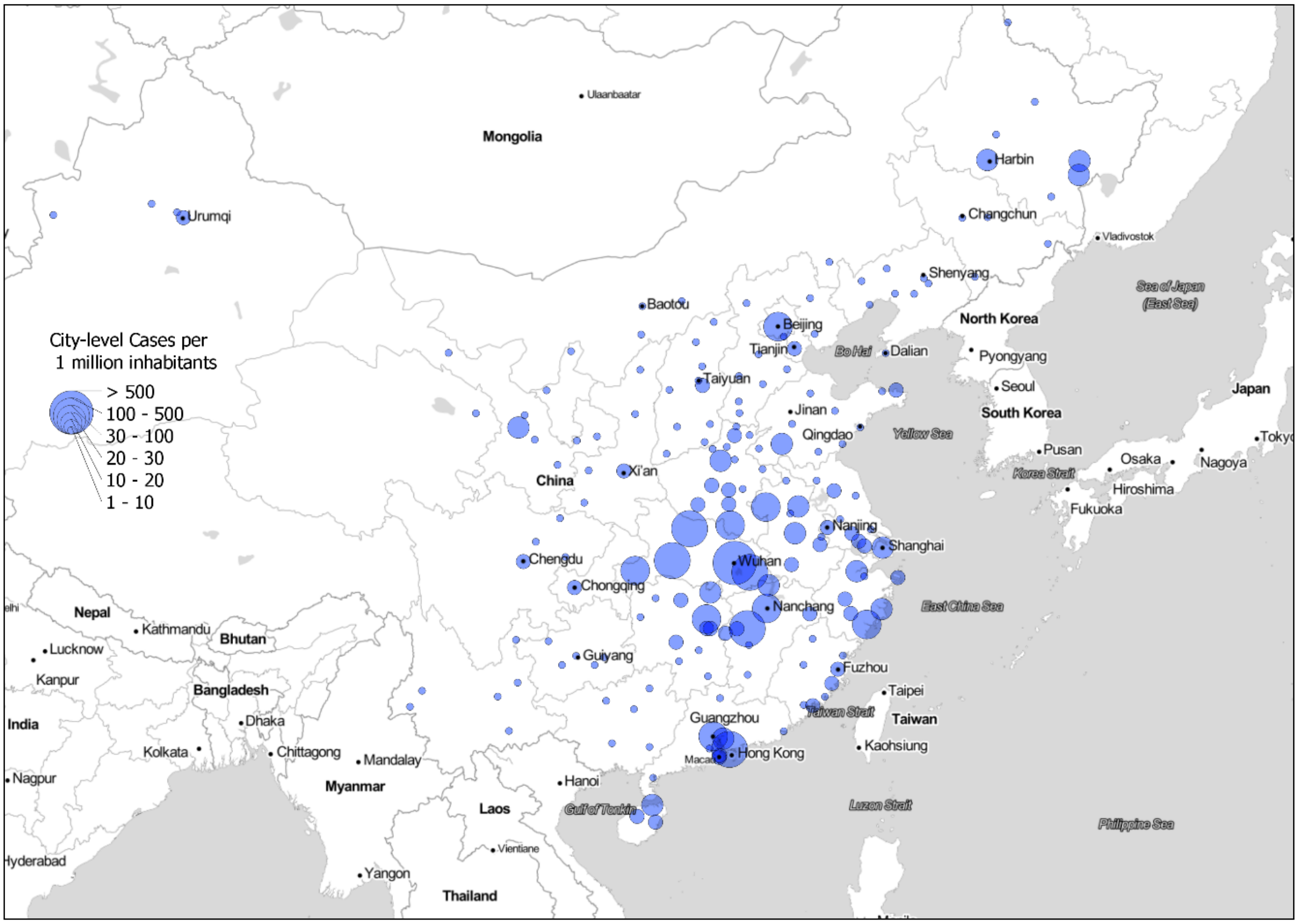

Figure 2: *Survey Participants' Locations during Survey Wave 3 and Virus Prevalence*

3.1 Measures of Virus Exposure

For robustness, we use three key measures of virus exposure.[19] See 3 for a summary. First, we use a standard epidemiological measure of disease prevalence: the logged number of confirmed cases per million inhabitants at the city-level, which we obtained from a variety of official sources including central and provincial governments and the Chinese Centre for Disease Control and Prevention (CDC). We then match cumulative case statistics on the date of the survey in March 2020 with participants'

[19]Admittedly, there are different measures of exposure that one can consider. We primarily use the number of confirmed infections or cases to measure exposure, which is heavily used in the epidemiological literature to model epidemic spread (e.g. Zhao & Chen, 2020) and well-accepted in the economics literature to investigate the effects of epidemics and pandemics on economic outcomes (Flückiger et al., 2019; Aksoy et al., 2020; Gonzalez-Torres & Esposito, 2020). An alternative measure is mortality and number of deaths. While mortality has been used as a measure of severity in papers assessing the long-term effects of a pandemic such as the Spanish flu (Karlsson et al., 2014; Adda, 2016; Aassve et al., 2020), it is not practical for assessing short-term effects due to its little variability at the onset of a pandemic. Epidemiological measures may however not fully capture the extent of exposure to the virus nor the general social 'concern' or 'sentiment' about the epidemic at the time. For robustness purposes we, thus, also use two alternative measures of exposure that capture these dimensions based on internet and social media data.

location. Cumulative case prevalence per million inhabitants in the 183 cities where our participants reside during the March survey fall between 0 and 5658. In our analysis, we use the log-transformed COVID-19 counts.[20]

Second, we construct a novel measure of city-level concern about the virus outbreak based on internet search data from Baidu, the most popular online search engine in China. The Baidu database provides daily population weighted search volume indices for commonly searched (coronavirus related) keywords at the city level. A high value of the Baidu 'concern index' for a certain keyword indicates that many people searched for information on the relevant keyword and cared about the relevant topic. The index has been widely applied in public health research for disease monitoring and prediction (Yuan et al., 2013; Li et al., 2017; He et al., 2018), the measurement of health-related public concern and awareness (e.g. Dong et al., 2019) and more recently also to the COVID-19 outbreak in China (Xiong et al., 2020). We extracted search volume indices for 20 keywords related to general interest searches about COVID-19 (e.g. novel coronavirus) and more specific to symptoms (e.g. dry cough) and personal protective measures (e.g. N95 masks) indexed at the city level (see Appendix Table A3 for a list of all keywords used). To capture overall city-level concern during the virus outbreak, we calculated the sum of all search term indices during the peak of the COVID-19 outbreak, between Wuhan Lockdown (23rd of January) and the date of the survey. We again use the log-transformed Baidu concern index for our analysis. Figure 3 displays the search volume indices between January and April for three popular keywords, as well as our 20-Keyword index.

[20]We use a log-transformation of COVID-19 confirmed cases to deal with skewed data due to the over proportionally large amount of cases reported in Wuhan City and Hubei province where some of our participants were located.

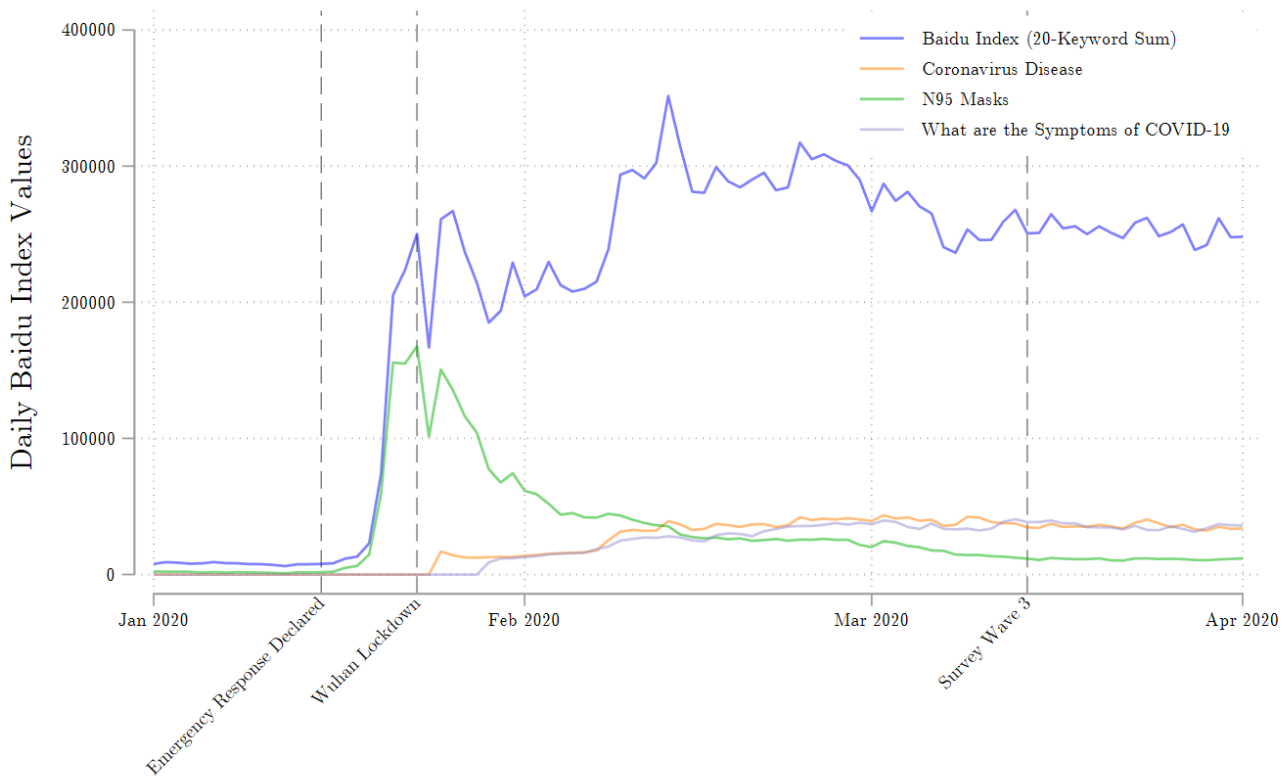

Figure 3: ***Time Series Baidu Index***

Note: The dashed lines indicate the dates on which the CDC announced the highest emergency response level (January 15$^{\text{th}}$), the lockdown of Wuhan (23$^{\text{rd}}$ January) and the date of the third Survey was released (14$^{\text{th}}$ March)

Third, we construct a novel city-level index of expressed (negative) sentiment related to COVID-19 from social media, as online sharing of emotional content specific to COVID-19 has the potential to bring about long-run societal change via emotional contagion (Steinert, 2020). For this, we extracted microblog posts (or tweets) from Sina Weibo, the Chinese equivalent to Twitter and one of the most popular social media platforms in China.[21] First, we extracted 523,222 geotagged microblog posts with the keyword novel coronavirus ('xin guan') which were posted online during the week prior to the third survey wave (from 0:00 am on 7th March to 5:00pm on 14th March). Posts were recorded in 179 of the 183 cities in

[21] A popular view is that an authoritarian regime censors social media. We believe that Chinese social media data provides a particularly interesting and valid source of expressed opinion in China. First, social media is not necessarily censored in an authoritarian regime, as the government can also use it as propaganda or surveillance tools (Qin et al., 2017). Second, the COVID-19 outbreak is a public health crisis and is less sensitive than a political event for the purpose of censorship of public opinions. Third, Sina microblog has been widely used as a reliable tool to analyse and track sentiment dynamics, psychological well-being, public knowledge and opinions, as well as a range of other attitudes towards public issues (e.g. air pollution in Zheng et al. (2019) and COVID-19 (Han et al., 2020; Li, Chen, et al., 2020; Li, 2020).

our sample. Second, we utilized the Linguistic Inquiry Word Count (LIWC) method, an automated text analysis method widely applied in psychology, which measures psychological and linguistic dimensions of written expression (Pennebaker & King, 1999). We employ the simplified Chinese version of the LIWC adapted by Gao et al. (2013). The LIWC text-processing programme uses a set of dictionaries to calculate the percentage of words that express positive and negative emotions for each microblog post. We construct our measure of expressed negative sentiment in a given city by calculating the average share of negative emotions expressed across all posts discussing COVID-19 in the week before the third survey was disseminated. The average score is recoded so that higher values represent greater negative mood (see Table 3 for details).

Both the Baidu search index and the negative sentiment index correlate positively to the infection rate, with the correlation coefficients being r=0.6 and r=0.2, respectively. The distributions across cities of the two indices do not deviate from that of the infection rate. We believe that the two indices provide valid measures of social sentiment, reflecting the intensity of exposure to the virus.

Table 3: Exposure Variables

	N	Mean	SD	Min	Max
City-level Cases	183	47	419	0	5658
City-level Cases (logged)	183	2	1	0	9
Baidu Search Index	183	80443	78138	5669	647294
Baidu Search Index (logged)	183	11	1	9	13
Negative Sentiment Index	179	2	0	1	3

Note: COVID-19 cases are population adjusted at the city-level (per 1 million inhabitants). Baidu Search Index is the city-level sum of search volumes for 20 Keywords related to COVID-19 between 23rd January and 17th March (see Table A3 for individual keywords). Negative Sentiment Index is the city-level average share of negative emotions expressed across all Sina-Microblog posts discussing COVID-19, shared between 7th and 14th March Data Sources: (1) Authors' compilation of official data from the State Council, provincial governments, and the Chinese CDC. (2) & (3) Authors' compilation of Baidu search data and Sina Weibo data.

3.2 Threats to Identification

Our identifying assumption relies on virus exposure being randomly assigned across participants. There are two possible sources of endogeneity that could undermine our identification – students' geographic dispersion and the spread of the virus. We discuss below how both concerns do not apply in our case.

3.2.1 Participants' Location Sorting

As discussed above, the pre-determined university admissions rule benefits our identification strategy by exogenously dispersing our participants across locations, which rules out potential bias from residential sorting. Moreover, the timing of student mobility in January 2020 was pre-determined exogenously by their term dates relative to the Chinese New Year.[22] However, one may be concerned that individuals' adaptive behaviour prior to the event undermines the estimated impact.[23] Specifically, students' travel decisions with respect to timing and destination may be related to how the unfolding disease situation was unfolding. To explore if this is the case, we first look at the descriptive statistics with respect to student movements. We observe that only 33 students (6% of the sample) whose registered Hukou (hometown) was not Beijing actually stayed in Beijing after the academic term had ended in December 2019 and none of these were from Hubei Province, the region most affected by the virus outbreak. During the holidays a small percentage of students normally remains in Beijing for various reasons (e.g., visiting family, internships, selection of civil servants, additional academic commitments). At the time of the 3rd survey wave, only 5 of these students (<1% of the total sample) had remained on the campuses that we surveyed. Hence, the raw data itself suggests that we do not observe any discernible patterns of adaptive behaviour that could undermine our identification strategy. To further investigate this concern, we explore using regression analysis whether students' travel decisions are independent of virus prevalence. First, we regress students' departure dates from Beijing on their initial (baseline) preferences, socio-demographic characteristics, their host university and their destination city. We do not find any significant estimates. Second, we regress a dummy variable equal to one if a student travelled to an alternative destination (i.e. not their hometown) or stayed in Beijing on future virus prevalence in their respective hometowns and a set of province dummies in which their hometown is located. Results show that whether students returned to their hometowns or not is unrelated to future virus prevalence. Moreover, only 28 students ever moved between neighbouring cities after January 23$^{\text{rd}}$ when Wuhan was locked down. In all cases, this was reported as visits to relatives, which is also part of the traditional celebration of the Chinese New Year. Overall, these findings clearly indicate that students did not behave adaptively in terms of mobility in response to the possible outbreak.

[22] There are typically two terms in Chinese education system – autumn and spring terms. The former consists of 17-20 weeks starting from early September till the Chinese New Year. The term dates are pre-determined and released before each academic year in September. The Term dates were not altered on account of the pandemic.

[23] For example, people would migrate out of cities in response to rising risk of adverse events (Brown et al., 2019).

Finally, students' mobility after the initial lockdown of January 2020 was strictly forbidden, and this was retained when the spring term started in late February. The Ministry of Education required all levels of schools to deliver online courses, and all students to stay home. College students were not allowed to return to their colleges.[24] The timing of leaving Beijing and the subsequent restriction on student mobility make their cumulative exposure to the virus situation in their current cities less likely to be individually selected. That said, individual fixed effects in our panel model mitigate any remaining concerns regarding endogenous adaptation.

3.2.2 Dispersion of the Virus

Whilst the initial outbreak in Wuhan can be deemed an unanticipated event, the further dispersion of the virus across China is unlikely to have followed an entirely random pattern. There are two possible confounding factors. First, according to the Chinese Emergency Law, there are four levels of emergency (from 1 (high) to 4 (low)) and from the second to the fourth level, the provincial governments are responsible for designing and implementing policies to control and prevent the infectious diseases. The State Council announced the highest level on 23rd January and adjusted it to Level 2 on 23rd February. Given our sample dates in March, it is likely that provincial policies and implementation affect individuals' most recent exposure (in the time domain) to the disease. Second, conditional on provincial environment, socioeconomic development (e.g., health facilities, population density) and geographic location of the city are plausible factors affecting individuals' exposure to local outbreaks (in the spatial domain). There is evidence which suggests that the virus spread was largely determined by population flows from Wuhan to other cities in China in the days before strict travel restrictions from and to Wuhan were enacted (Kraemer et al., 2020). If confounding factors exist, we need to control for these in our main specification.

To determine which city-level factors might confound our results, we regress our three measures of city-level virus exposure on a set of city-level variables, which have been found to affect the dispersion of COVID-19, most importantly the rate of migration between Wuhan and other cities in China. To calculate population mobility we extracted data on inter-city population flows from the Baidu Migration Database (https://qianxi.baidu.com/) tracking individuals' check-in locations in all Baidu applications (e.g., Baidu

[24]For example, some of our sample universities also required students to report their locations and health information on a daily basis. Mobility out of their current city has been forbidden.

map, search, takeaway, and social media "tieba") through their mobile devices. We use the average population inflow from Wuhan as a share of total immigration to each of our sample cities between 20th and 23rd January. Higher values indicate that a larger proportion of the inflowing population originated from Wuhan, which reflects a greater connectedness between Wuhan and the respective city.

Following recent research (e.g. Becchetti et al., 2020; Pluchino et al., 2020), we further control for city-level population density, the number of hospitals and doctors per million inhabitants, the amount of city-level health expenditure as a share of total fiscal expenditure, GDP per capita and annual average Air Quality Index (AQI) based on daily records of 1,436 air monitoring stations since 2015.[25] The stringency and duration of lockdown may also significantly affect the development of city-level virus outbreaks. We manually collected data of official re-opening dates for shops, restaurants, indoor and outdoor activities, respectively, for each city from 183 municipal governments' official websites and news. We constructed a lockdown duration index as the standardised sum of days all city-level lockdown measures were in place. Finally, given that the provincial governments are responsible for designing and implementing local policy for COVID-19, the dispersion of the virus is likely to be determined by numerous province-level factors, including social and geographic proximity to Wuhan, long-run policies effecting socio-spatial vulnerability of communities, virus-preparedness, and the ability to respond (e.g. province-level measures to mitigate the virus outbreak). Hence, we include province-level fixed effects into the model.

The results of our regressions are shown in Appendix Table A4. We find that for all three city-level exposure variables (*Ln*Cases, Baidu Concern Index and Sentiment Index) a large part of the variation is explained by province fixed effects and the share of immigration from Wuhan during the days prior to the lockdown of Wuhan and the imposition of travel restrictions. With respect to additional city-level factors, population density and GDP per capita are positively, and health expenditure negatively correlated with the Baidu Search Index, long-run air quality and the number of hospitals per capita are both positively associated with the Negative Sentiment Index. Based on this analysis, we include immigration rate, population density, GDP per capita, the number of hospitals, health expenditure and annual average AQI as city-level controls as well as province fixed effects into our main empirical specification, which we discuss next.

[25]The city-level factors are the 2018 data compiled from provincial statistical yearbooks. The AQI is an index of air quality consisting of six key pollutants from 1,436 air monitoring stations across 338 cities, having been set up by the Ministry of Ecology and Environment since 2012. We calculate the annual average AQI for each city based on daily AQI readings between 2015-2019.

4 Empirical Strategy and Attrition

To estimate the effect of virus exposure on social behaviour and economic preferences, we use a generalised Difference-in-Differences Model (DID). It differs from a classic DID model in the sense that the treatment variables in our case are continuous, rather than binary (Wing et al., 2018). Importantly, the panel structure of our data allows us to control for individual unobserved fixed effects and isolate the effects of the exogenous treatment, by comparing the differences before and after the virus outbreak across participants who experienced different levels of exposure to COVID-19. We estimate the following main specification:

$$\mathbf{Y^k}_{\mathbf{ijt}} = \delta Exposure_j + \beta_3 X_{jt} + \eta_i + \lambda_t d_p + \varepsilon_{ijt} \tag{1}$$

where $\mathbf{Y}^k_{ijt}$ is primary outcome k from the experimental modules discussed above for individual i living in city j at time $t.X_{jt}$ is a vector of city-level controls ; η_i represents unobservable time-invariant individual fixed effects; $\lambda_t d_p$ represent a province-specific time trend, given that provincial governments' design and implementation of policies provide sources of variation in city-level exposure to the virus; $Exposure_j$ is a continuous variable of being exposed to COVID-19 (City-level cases, Baidu Concern Index, and Sentiment Index) at the time of survey Wave 3 and ε_{ijt} is the random error term. In this specification, the parameter of interest is the difference-in-differences estimator δ, reflecting the impact on $Y^k{}_{ijt}$ from variations in the intensity of treatment in the post-outbreak period ($Post$). We accommodate for potential serial correlation by estimating clustered standard errors at the individual level.

A key assumption underlying the DID identification strategy is the common trends in outcomes between treatment and control groups in absence of a treatment. Whilst this assumption is not directly testable, we are able to test parallel trends before the virus outbreak for outcome variables which were collected in both the October and December 2019 surveys, including an incentivised measure of risk preferences and two measures of well-being (i.e. depression and general health of participants). This is shown graphically in Appendix Figure A3. We further estimate the difference-in-differences model above using the October and December data on the same three outcome variables as if the outbreak had taken place before the December survey (see Table A5 in the Appendix). Both the visual and formal assessments lead to the conclusion that trends in risk preferences, depression and general health did not differ between treatment and control groups in the months prior to the virus outbreak.

For the remaining outcomes that only appear in either the October or December survey, we repeat the same regressions described above using survey data from each corresponding month. The results indicate that pre-outbreak preferences are uncorrelated with future virus-exposure (see Table A6). We conclude from this exercise that the common trends assumption likely holds in the context of our data.

We are also able to ascertain that individuals did not differ in their socio-demographic characteristics with respect to the degree of virus exposure. Based on the epidemiological measure of exposure (i.e. the number of cumulative confirmed cases at the city-level on the day of the survey adjusted by population), we report summary statistics of basic characteristics of survey participants between those that were severely exposed (the top tercile of virus prevalence), moderately exposed (the middle tercile) and those that were only mildly exposed (the bottom tercile). See Table A7 in the Appendix. We find broad balance across basic demographics including age, gender, year of study, being the only child (significant at the 5% level, chi2-test) and Hukou registration indicating rural or urban origin of participants.

A further concern relates to the potential of differential attrition, which may bias our estimates. Table B1 in Appendix B shows that attrition rates across the three waves in our data as 16% between Waves 1 and 2, 19% between Waves 2 and 3. These rates are comparable to previous research conducted via WeChat surveys (e.g. Chen & Yang, 2019). In Appendix B, we also explore in more detail the patterns of attrition in our data and conduct standard attrition tests. We attempt to address differential attrition in our analysis by applying inverse probability weights (IPW) following Wooldridge (2002). First, we predict the probability (p_i) of being observed in all three survey waves by regressing a dummy variable equal to one if an individual did not attrite, on (1) a constant term, (2) the primary treatment variable (*Ln*Cases) and (3) a rich set of co-variates measured at baseline for all initially recruited participants. Each individual then receives a weight equal to $1/p_i$ in all regressions in the proceeding analysis.

Finally, we address the threat of multiple hypothesis testing and the possibility of false positives by estimating sharpened q-values using the false discovery rate (FDR) procedure (Benjamini et al., 2006; Anderson, 2008). We calculate FDR adjusted q-values for three sets of p-values across all k-outcomes (including three indices) for each of our three treatment variables. We report both conventional p-values and FDR adjusted q-values in all regression output tables.

5 Results

5.1 Short-term effects

We present our main treatment effects based on official COVID-19 infection data. Figure 4 shows the average treatment effect on the treated estimated following equation (1) and corresponding confidence intervals for our incentivised primary outcomes for anti-social behaviour and economic preferences.[26] Prior to estimation, all outcomes were standardized (z-scored) on the mean to allow for a comparison of treatment effects in units of standard deviations across different outcomes.

[26]All regressions are estimated using all participants who took part in all 3 survey waves (N=522). Note that 15 individuals were excluded from the original sample (N=539), who had not completed all three surveys. Additionally, one individual from Macau and one individual from Hong Kong were removed from the sampled due to unavailability of data for control variables. Individual beta parameters could not be estimated for 9 individuals. This resulted in models for present bias to be estimated with 513 individuals and all other models estimated with 522 individuals. The full results of each regression can be found in Appendix Tables A8 – A11. A visualisation of treatment effects for all outcome variables is provided in Figures A4 and A5.

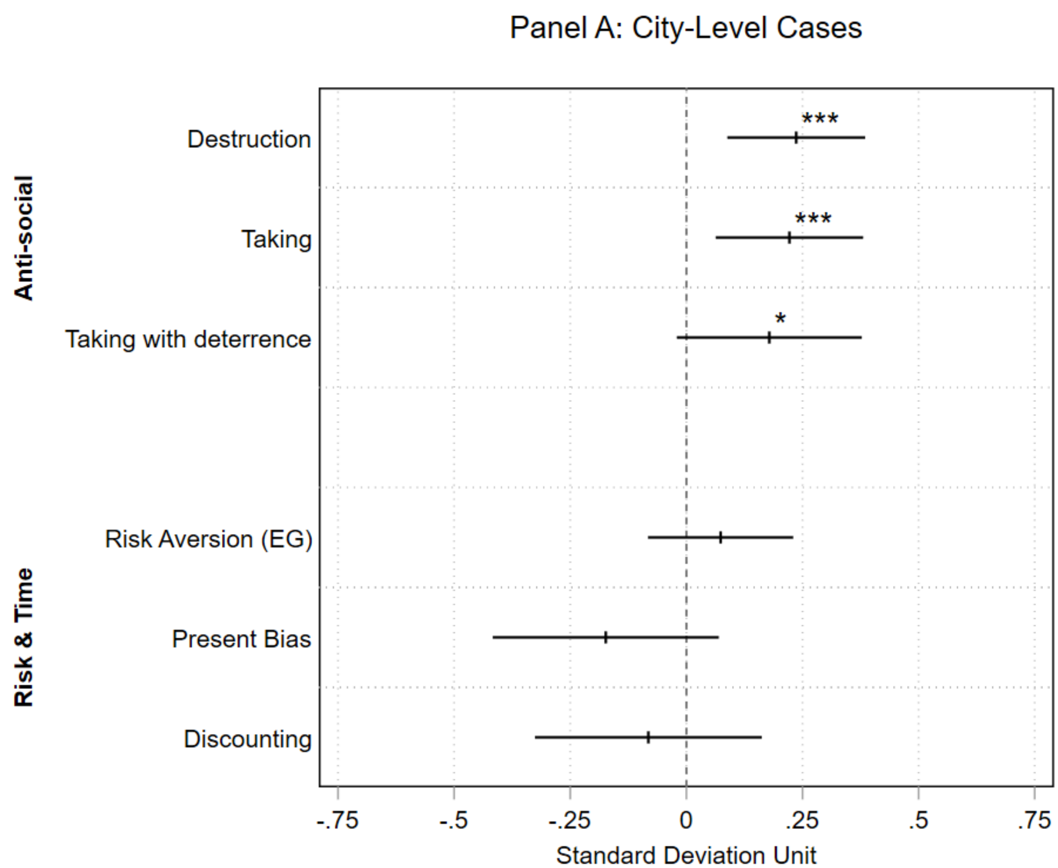

Figure 4: ***Treatment Effects - Official COVID-19 Infection Data***

Note: X-axis plots the estimated coefficient for a 1-unit increase in the Log of Cases per million Inhabitants. Significance stars *** p<0.01, ** p<0.05, * p<0.1 based on p-values estimated using cluster-robust standard errors, clustered at the individual level. All models are estimated with IPW to account for differential attrition. Number of observations N=1044.

Figure 4 shows the estimated treatment effect of a one-unit increase in the log of city-level cases per million inhabitants. We find that people more exposed to virus appear to become more antisocial. Estimates for pro-social, norm enforcement, risk and time preferences are close to zero and show no statistically detectable difference. However, people more exposed to the virus outbreak appear to become more antisocial.

Table 4 provides corresponding difference-in-difference estimates for anti-social behaviour, the only dimension of decision-making which appears to be significantly affected by virus exposure. We observe that the coefficient of interest (post $\times$ LnCases) shows a significant relationship between the intensity of the outbreak and all outcomes. Specifically, we find that individuals destroy more of their paired player's endowment in the Joy of Destruction Game (column 1) and take more in the Take Game without deterrence (column 2). The 0.24 standard deviation increase in destructive behaviour corresponds to an increase of approximately 9 percentage points, which is statistically significant at the 1% significance level. Similarly, we find that more exposed individuals take on average around 7% more of the other player's endowment (0.22 s.d.). Both results are statistically significant at the 1% significance level and remain significant at the 5% level after adjusting for multiple hypothesis testing. In column (3), we find a slightly smaller effect of virus exposure on taking when there is a risk of being detected (0.18 s.d.). which is statistically significant at the 10% level using conventional p-values, but does not survive multiple hypothesis testing corrections.

Table 4: Difference-in-difference analysis: Anti-social behaviour

	(1) Joy of Destruction	(2) Take Game	(3) Take Game (Det.)	(4) Anti-sociality Index
Panel A				
Post × *Ln*Cases	0.237***	0.222***	0.179*	0.212***
	(0.076)	(0.081)	(0.101)	(0.058)
	[0.014]	[0.028]	[0.311]	[0.005]
R^2-Within	0.018	0.038	0.036	0.053
Number of Individuals	522	522	522	522
Observations	1,044	1,044	1,044	1,044

Note: Difference in differences analysis using fixed effects OLS regressions accounting for attrition using Inverse Probability Weighting (IPW). Standard errors clustered at the individual level in parenthesis. Multiple testing adjusted False Discovery Rate (FDR) q-values in square brackets. *Post* × *Ln*Cases is the interaction of logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey with a post-outbreak indicator. All regressions include individual fixed effects, time-varying city-level controls for average immigration rate from Wuhan (20-23 Jan), number of hospitals per million inhabitants, health expenditure as a share of total expenditure, population density, GDP per capita and province-specific time trends. The dependent variable in column (4) is an index for anti-sociality based on the average of the z-scores of all three anti-social outcome variables.
* p < 0.1, ** p < 0.05, *** p < 0.01.

Finally, we also construct a simple index for anti-social behaviour compromising the choices made in all three games of the anti-social behaviour module.[27] Column (4) reports the effects of a unit increase in exposure on the anti-social behaviour index, which confirms our earlier results and show that participants exposed more heavily to the virus significantly increase anti-social behaviours (0.21 s.d.).

[27]The Anti-sociality index is an average of z-scores of Destruction, Taking and Taking with deterrence. The index construction follows Kling et al. (2007).

5.2 Results based on Social Media Data

The following section presents the results of our alternative measures of COVID-19 exposure based on social media data. Panel B in Figure 5 shows the estimated treatment effect of a one-unit increase in the log of the Baidu Index, which captures city-level concern around COVID-19. Panel C shows the estimated treatment effect of a one standard-deviation unit increase in the Negative Sentiment Index. The latter index, constructed using text-analysis of Sina Micro-blog posts discussing COVID-19, provides a measure of city-level (negative) sentiment. Regression results presenting detailed estimations can be found in the Appendix Tables A8-A9.

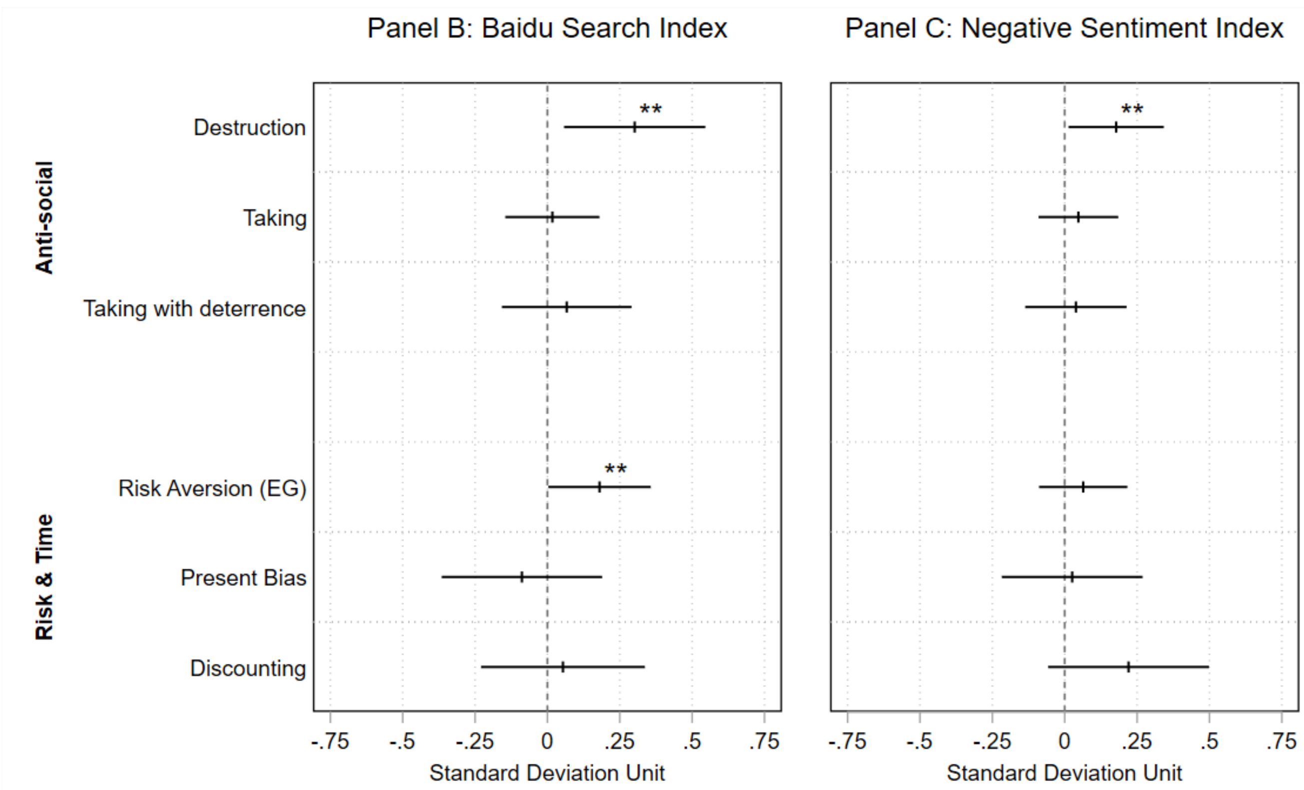

Figure 5: ***Treatment effects - Social Media Data***

Note: X-axis plots the estimated coefficient for a 1-unit increase in the log of the Baidu Index measuring concern about COVID-19 (Panel B) and a 1-unit increase in the Negative Sentiment Index (Sina Weibo) (Panel C). Significance stars *** p<0.01, ** p<0.05, * p<0.1 based on p-values estimated using cluster-robust standard errors, clustered at the individual level. All models are estimated with IPW to account for differential attrition. Number of observations Panel B: N=1044, Panel C: N=1036.

With respect to anti-social behaviour, we observe again a similar yet less pronounced pattern as in the previous section. In Panel B, an increase in city-level concern is associated with a general increase in destructive behaviour (0.28 s.d.), significant at the 5% level. In Panel C, higher levels of negative sentiment led to a general increase in destructive behaviour (0.21 s.d.), significant at the 5% level. However, after adjusting for multiple hypothesis testing, the observed increases in destructive behaviour are no-longer statistically different from zero.

The exposure variables based on social media data reveal additional treatment effects, not found when using infection data. First, we find that higher levels of city-level concern (Panel B) are associated with a significant increase (0.21 s.d.) in risk-aversion (at the 5% level, measured using Eckel & Grossmann's lottery choice task (Risk Aversion – EG). However, FDR adjustments render this finding insignificant. Second, we observe that city-level concern (Panel B) is also associated with a decrease in altruism, which is highly statistically significant at 1% level and remains statistically significant at the 10% level after FDR adjustments.

5.3 Sensitivity Analysis

We perform additional sensitivity analysis to ensure robustness of the results presented in this section. First, we exclude participants located in Hubei Province (N=10) from the analysis to mitigate the influence of potential outliers. Wuhan, the epicentre of the outbreak and surrounding cities in Hubei province were most severely affected by the virus and reported disproportionally high numbers of cases compared to the rest of China. We find that our results are largely unaffected (see Appendix Figure A6).

Second, another indicator of the COVID-19 epidemic which is usually reported is mortality and thus we control for reported mortality at the city-level. In China, COVID-19 related mortality was largely concentrated in Hubei Province, with 50% of the cities in our dataset reporting zero deaths by the date of the survey. Hence, we believe that mortality does not serve as a good indicator for virus exposure per se. Nonetheless, we add mortality (i.e. the cumulated confirmed cases reported on the date of the third survey at the city level) as a control to our baseline specification. Again, we find that our main results are robust to this specification (see Appendix Figure A6).

Finally, as our study sets itself apart from a number of rapidly emerging COVID-19 papers, relying on post-outbreak data between-subject designs, we test whether our results would be affected by how the impact of COVID-19 is identified. We focus on the following comparison: We use the data collected in

March 2020 only and re-estimate equation (1) thereby reflecting a between-subject design relying entirely on post-outbreak data and exploiting variations in individual exposure to the virus. This identification has been widely utilised by existing studies reviewed in the Introduction. Although differences are subtle, the results show that we would underestimate the effect of virus exposure on anti-social behaviour, if ignoring individuals' heterogeneous initial preferences and other time-invariant unobservables (see Appendix Figure A7).

5.4 Potential Pathways

The analysis so far has found consistent evidence that COVID-19 exposure leads to an increase in anti-social behaviour. In this section, we explore a range of potential pathways through which exposure to COVID-19 may be associated with anti-social behaviour. We are especially interested in cognitive and psychological well-being, which have found to be important determinants of anti-social behaviour. We utilise variables measured both before and after the outbreak, capturing components of cognitive ability, psychological and physiological well-being. In addition, we construct a measure of virus-specific subjective risk perception using a set of variables elicited in the post-outbreak survey.[28] Empirically, we use a triple-difference approach to assess potential pathways by estimating separate regressions for each variable using a fully interacted variant of our main equation (1):

$$\begin{aligned} \mathbf{antisocial}_{it} &= \delta(LnCases_j) + \beta_2(post_t \times Z_{it}) \\ &+ \theta(post_t \times LnCases_j \times Z_{it}) + Z_{it} + X_{ijt} + \eta_i + \varepsilon_{ijt} \end{aligned} \tag{2}$$

where $\mathbf{antisocial}_{it}$ is an index of anti-social behaviour[29]; Z_{it} captures the change in cognitive ability, psychological, physiological well-being or virus risk perception measured only in Wave 3. X_{jit} now represents a vector of all time-varying control variables and time-fixed effects contained in equation (1) as

[28] Specifically, we asked respondents to indicate, on a 10-point Likert Scale, the perceived level of risk/threat posed by the virus to (1) themselves, (2) their family and (3) society as a whole. This provides a measure of emotional risk perception. We further asked respondents how they perceived the level of infections at their current location and whether any of their friends or family had been infected with the virus, which captures cognitive risk perception. As all five variables are highly correlated, we conduct a factor analysis to predict an underlying "Virus-Risk Factor" for each individual.

[29] The index is constructed by calculating the average of z-scores of the three tasks of the anti-social behaviour module following Kling et al. (2007).

Table 5: Heterogeneous Impact of Virus Exposure by Potential Mechanism

	Anti-sociality Index						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post × *Ln*Cases	0.215***	0.213***	-0.008	0.192***	0.098	-0.153	0.273
	(0.058)	(0.059)	(0.352)	(0.066)	(0.072)	(0.152)	(0.382)
Post × Virus Risk × *Ln*Cases		0.137**					
		(0.069)					
Post × Cognitive Ability × *Ln*Cases			-0.034				
			(0.053)				
Post × Depletion × *Ln*Cases				0.030*			
				(0.016)			
Post × Depressive Symptoms × *Ln*Cases					0.460***		
					(0.125)		
Post × Negative Affect × *Ln*Cases						0.042***	
						(0.015)	
Post × General Health × *Ln*Cases							-0.013
							(0.094)
R^2	0.115	0.204	0.208	0.235	0.208	0.194	0.201
Number of Individuals	522	522	522	522	522	522	522
Observations	1044	1044	1044	1044	1044	1044	1044

Note: Table reports OLS estimates of equation (2) where the dependent variable is an index of anti-social behaviour. *Ln*Cases is the logged number of confirmed cases at the city-level per million inhabitants reported on the date of the third survey. All triple interaction terms provide estimates for potential heterogeneous effects. Virus-risk Factor is a continuous score of virus-risk perception obtained from a factor analysis of post-outbreak survey responses (score ranges between approximately -2 and 2). All remaining variables were measured both pre and post outbreak: Raven score captures the number of correctly completed puzzles (Score: 0-9, recoded so that a higher score represents less completed puzzles). Depletion is a continuous score for state self-control capacity (Score between -7 and 11: higher score indicating more depletion). Depression is a binary variable that takes the value 1 if an individual has depressive symptoms. Negative affect is a continuous score measuring negative affect (Score: 0-30, higher score indicating higher negative affect); General health is a continuous measure of general health (Likert scale: 1-5, higher score indicating better health). For details on how variables were measured and constructed, see Table A2.

* p < 0.1, ** p < 0.05, *** p < 0.01.

well as their interaction with Z_{it}. In this regression, a positive significant estimate for the triple-difference coefficient θ would suggest that there is a statistically significant difference between individuals who experienced an increase in Z_{it} and those who did not. Table 4 provides an overview of all estimation results from equation (2).

First, we explore whether effects of COVID-19 exposure vary by subjective virus-risk perception. This is particularly relevant, as the perceived threat of the virus may differ largely between individuals, even if they are exposed to the same number of cases at the city-level. Hence, one might hypothesise that only those individuals with higher subjective risk perception change their behaviour in response to increased objective virus exposure. The triple-difference estimate in Table 5, column (2) suggests that this may be the case. There is a statistically significant positive difference between individuals with higher subjective risk perception, which points to the importance of how the virus is perceived. However, this difference is only significant at the 5% level and the estimate of δ remains positive and highly statistically significant, which suggests that differences are not fully explained by subjective risk perception.

A large literature in behavioural economics suggests that cognitive capacity and self-control can affect economic and social decision-making (e.g. Mani et al., 2013; Carvalho et al., 2016; Friehe & Schildberg-Hörisch, 2017). We hypothesise that changes in cognitive capacity may interact with higher virus-prevalence. For example, it has been shown that individuals with low self-regulatory resources (i.e. in a state of ego-depletion), feel less guilt and subsequently show less pro-social behaviour (Xu et al., 2012). In columns (3) and (4) we explore two measures of cognitive capacity. We find no statistically significant difference for individuals who perform worse in a set of Raven's Progressive Matrices (a measure of cognitive ability). We find suggestive evidence of higher anti-social behaviour, significant at the 10% level, for individuals who report higher levels of momentary ego-depletion (evaluated at the time the survey was taken) and hence may have lower self-control capacity.

Next, we explore whether mental health may be driving the observed relationship. Emerging research in psychology shows that COVID-19 is likely to have serious consequences on mental health, resulting in increased levels of depression and other mental disorders (Huang & Zhao, 2020; Pfefferbaum & North, 2020; Raker et al., 2020; Thombs et al., 2020). In turn, research in behavioural economics and cognitive science find that depression and negative emotions and mood are able to impair decision-making in more general terms (Gotlib & Joormann, 2010; Haushofer et al., 2013; De Quidt et al., 2018). Based on such previous evidence, we hypothesise that once an individual's mental health is compromised, he or she may be less likely to care for others and act in a more anti-social manner. In columns (5) and (6) we provide evidence that this might be the case. We find a statistically significant increase in anti-social behaviour for individuals who were subject to greater virus exposure and experienced an increase in depressive symptoms (as measured by the 10-item CESD depression scale) and negative affect (or mood) (measured by the PANAS scale).[30]

Finally, besides psychological well-being, we also check for the effects of physiological well-being using self-reported health status as an indicator (column 7). We find no statistically significant difference, which leads us to conclude that changes in anti-social behaviour are likely to be driven by a deterioration in mental health. In the following section, we discuss interesting directions for the design of public

[30]We also assess whether changes in other dimensions of psychological well-being including happiness and positive mood interact with virus-prevalence. We find no statistically significant difference in anti-social behaviour for individuals with increased positive affect (on the date of the survey), self-assessed happiness, meaningfulness of life and life satisfaction (see table A12 in the Appendix).

health interventions to mitigate compromising social behaviour and mental health.

5.5 Limitations

The following clarifies important limitations of our sample with respect to attrition and generalizability. In addition, we also address the fact that our study's time frame is limited to only the pre- and immediate post-outbreak and first lockdown period in China. Therefore, we are only able to provide estimates on the short-term impact of Covid-19 and our results are unable to speak to literature addressing long-term impacts and the effect of multiple lockdowns.

One of the major problems with longitudinal studies is attrition by introducing possible bias when participants who drop out of a study are systematically different from those who remain in it. We have taken various steps to deal with this potential concern in our dataset (all additional analysis can be found in Appendix B). Amongst others, we formally test for non-random attrition and find that attrition is unrelated to our treatment variable – city-level Covid-19 cases – yet is related to certain participant characteristics measured at baseline. We address attrition by implementing a separate BGWL test and by applying inverse probability weights to all our regressions. We acknowledge, however, that inverse probability weighting is limited in that it can only address attrition based on observable characteristics, and some attrition might still be non-random.

Another caveat is that our study was conducted with a convenience sample of university students from Beijing which is not representative of a more general population sample and therefore our results and suggestions for certain policy interventions should be interpreted considering this specific group. Nonetheless, although university students differ from the general population in certain characteristics (e.g., our sample is significantly younger and better educated and females are overrepresented), other research indicates that student populations exhibit very similar behavioural patterns with respect to social preferences, where university student samples usually provide lower bound estimates of pro-sociality (i.e., they are less altruistic) (Falk et al., 2013; Snowberg & Yariv, 2021).

Finally, our study focuses on the immediate impact of the Covid-19 outbreak, leveraging data from just before the outbreak of COVID-19 and immediately after the first wave was overcome. We are therefore not able to investigate preferences, perceptions and attitudes over longer periods of time or capture the effects of experiencing multiple lockdowns as other studies are able to (e.g. Aragon et al., 2022; Harrison et al., 2022). Noteworthy in the Chinese context is that China has maintained a 'Zero-Covid' strategy, which has resulted in multiple large-scale lockdowns since the initial lockdown which we study. Another notable difference to the initial lockdown is that those later, larger and longer lockdowns have also sparked social unrest among residents and university students living under strict lockdown conditions. Some literature focusing on Covid-19 and social unrest highlights an association between increased emotional stress, anxiety and aggression and the incidence of social unrest (.) This also speaks to our results, as we already observe an increase of negative affect and anti-social behaviour after the first lockdown period in China. Finally, we acknowledge, that our results must be interpreted as short term effects, as we are confined to data from before and immediately after the first lock-down. Our findings thus complement research which utilises longitudinal data and multiple surveys over longer time periods after the first lockdown.

6 Discussion and Conclusion

In this paper, we test whether exposure to the COVID-19 pandemic alters social behaviour and economic preferences of individuals. We exploit a unique experimental panel dataset that enables us to track changes in social behaviour and economic decision-making of the same individuals before and after the COVID-19 outbreak. In order to capture multidimensional responses to the virus outbreak, we construct city-level measures of societal concern and sentiment specific to COVID-19 in addition to standard epidemiological measures of virus exposure (cases per million inhabitants). The novelty of our approach pertains to our within-subject design which controls for unobserved individual characteristics, rich variation in individual exposure to multiple measures of the virus outbreak and the ability to provide insights into the channels transmitting the influence on individual preferences.

Our main finding is that greater exposure to COVID-19 causes an increase in anti-social behaviour. This finding contributes to a growing body of literature exploring how preferences respond to traumatic exogenous shocks and stressful situations such as war, conflict and public health crises. We are able to extend this earlier work by considering the acute effect on decision-making during an unfolding crisis and testing potential pathways through which such an event may influence behaviour, in particular mental health. Bauer et al. (2016) note that negative shocks are likely to have a positive legacy on pro-social behaviour in the long-term in terms of cooperation, altruistic giving and civic participation. This is in line with findings from Grimalda et al. (2021) showing that exposure to COVID-19 is associated with increased altruism measured months after the outbreak. In contrast, our findings show that, in the short-term, anti-social behaviour increases. In addition, we show that anxiety reflected by online search behaviour at the onset of the crisis and negative sentiments further undermine altruism. Our findings of increased antisociality and largely stable prosociality contribute to the literature on social preferences and exogenous shocks, which has largely produced mixed evidence. For instance, Branas-Garza et al. (2022) and Buso et al. (2020) find that prosociality decreased during periods of the first Covid-lockdowns. Others, such as Bokern et al. (2021) using data of multiple waves up to one year after the start of the first lockdown, note some short-term fluctuations yet show by large stability of social preferences measured with the help of a solidarity game. Shachat et al. (2021) provides mixed findings with respect to social preferences, showing greater levels of cooperation and lower levels of trust in their sample.[31]

We also contribute to the literature that examines the stability of risk and time preferences over the course of the Covid-19 outbreak. We find no significant changes in either risk or time preferences caused by exposure to the virus outbreak. Our findings are in line with a number of other studies providing evidence on the intertemporal stability of risk and time preferences (Angrisani et al., 2020; Drichoutis & Nayga, 2021; Guenther et al., 2021; Harrison et al., 2022). Two studies focusing on samples from Wuhan, China provide mixed evidence. Shachat et al. (2021) find increased risk tolerance, while Bu et al. (2020) find decreased risk tolerance in their post-outbreak survey, the latter effect potentially explained by rising pessimistic beliefs rather than changes in general risk preferences. Exploiting within-student changes in preferences, and variation in exposure to the outbreak, Bu et al. (2020) also show that risk taking is irresponsive to the level of virus exposure, which aligns with the findings from our within-subject analysis.

In addition to methodological innovations, our research is further able to elucidate the potential mechanisms driving the relationship between COVID-19 exposure and changes in anti-social behaviour. We find that the effect of virus-exposure on anti-social behaviour is most pronounced for those individuals

[31] We acknowledge that our experimental approach holds many similarities to Shachat et al. (2021), but there are also some notable differences with respect to research design and identification strategy, including the use of a within- instead of between-subject design, additional survey data to study potential mechanisms and a more nuanced analysis with respect to exposure to the virus outbreak through ample geographical variation in virus prevalence. Note that the main sample difference to Shachat et al. (2021) and Bu et al. (2020) is that both studies heavily draw on students located in Hubei province where the majority of Covid-19 cases were reported, while our study relies on geographical variation in student's location with only few students having been located in Hubei.

who experienced an increase in depression or in their negative mood, whereas changes in cognitive ability and ego-depletion do not seem to interact with virus exposure. Our results are related to Belot et al. (2020) providing survey evidence from the early phase of the pandemic in China, documenting that younger people are significantly more likely to report negative effects on mental health. That said, our results from a student sample suggest that the effect on anti-social behaviour is likely to be smaller in a general population sample. Nonetheless, we are not able to rule out that alternative mechanisms exist, which are not explored in this paper. For example, economic stressors are often named as a cause of antisocial behaviour (e.g. Schneider et al., 2016). Due to a lack of specific data on individual economic conditions, we are unable to ascertain whether economic uncertainty or financial insecurity interact with increased virus exposure. Nonetheless, we believe economic stressors to be closely related to the emotional well-being pathway, for which we find robust evidence.

This finding has important and practical implications for policies designed to tackle major public health crisis events. While most government resources usually focus on mitigating the virus outbreak per se, such as in the form of expanding medical treatment for infected people, our results suggest that interventions to provide psychological support are critical in response to such pandemics. In the context of COVID-19 or similar events, investments should therefore also focus on expanding the supply of consultation with mental health professionals in the form of online and smartphone-based psychological support avenues that can reach a wider audience of potentially affected people. Our evidence suggests that such psychological interventions that aim to promote mental well-being should be initiated from the starting point of a major health crises and not follow much later (Duan & Zhu, 2020).[32] In addition to counselling, research from behavioural economics and psychology point out promising light-touch interventions to reduce acute stress and depression and foster pro-social behaviour including the application of mindfulness mediation and mindfulness-based cognitive therapy (Leiberg et al., 2011; Kang et al., 2014; Sun et al., 2015; Iwamoto et al., 2020).[33]

Promoting and galvanizing socially responsible behaviour has been at the core of many governments' COVID-19 response and research shows that pro-sociality predicts health behaviours and compliance with public health guidelines (Campos-Mercade et al., 2021).[34] Our findings suggest that addressing

[32]The advantage of online consulting is that it can be efficiently scaled at low cost and at the same time there is evidence of the effectiveness of digitally provided psychotherapy when compared to face-to-face therapy, in particular when treating acute symptoms of stress and depression (Barak et al., 2008; Andersson et al., 2014; Carlbring et al., 2018).

[33]Again, this intervention has been shown to be effective when delivered online (Spijkerman et al., 2016) and thus lends itself for large-scale application during COVID-19 or similar events.

[34]We also assess whether social behaviour correlates with self-reported compliance with protective behaviour and knowledge

poor mental health, early on during the crisis, may play an important role in avoiding increases in anti-social behaviour and ensuring wide-scale adherence to public health guidelines.

related to the virus. Exact wording of the questions can be found in Appendix C. We find little to no correlation between our indices of pro-sociality, anti-sociality and norm-enforcement and protective behaviour as well as virus knowledge. This, however, is likely caused by a lack of variation in compliance with protective behaviour, with overall compliance being overwhelmingly high amongst our sample population. Regression results are presented in Table A13 in the Appendix.

References

Aassve, A., Alfani, G., Gandolfi, F., & Moglie, M. L. (2020). Epidemics and Trust: The Case of the Spanish Flu. *NBER Working Paper*.

Abbink, K., & Herrmann, B. (2011). The moral costs of nastiness. *Economic Inquiry*, *49*(2), 631–633. https://doi.org/10.1111/j.1465-7295.2010.00309.x

Achtziger, A., Alós-Ferrer, C., & Wagner, A. K. (2016). The impact of self-control depletion on social preferences in the ultimatum game. *Journal of Economic Psychology*, *53*, 1–16. https://doi.org/10.1016/j.joep.2015.12.005

Achtziger, A., Alós-Ferrer, C., & Wagner, A. K. (2018). Social preferences and self-control. *Journal of Behavioral and Experimental Economics*, *74*(April), 161–166. https://doi.org/10.1016/j.socec.2018.04.009

Adda, J. (2016). Economic Activity And The Spread Of Viral Diseases: Evidence From High Frequency Data. *The Quarterly Journal of Economics*, *53*, 891–941. https://doi.org/doi:10.1093/qje/qjw005

Aksoy, B., & Palma, M. A. (2019). The effects of scarcity on cheating and in-group favoritism. *Journal of Economic Behavior and Organization*, *165*, 100–117. https://doi.org/10.1016/j.jebo.2019.06.024

Aksoy, C. G., Eichengreen, B., & Saka, O. (2020). Revenge of the Experts: Will COVID-19 Renew or Diminish Public Trust in Science? *SSRN Electronic Journal*, 1–17. https://doi.org/10.2139/ssrn.3617825

Alsharawy, A., Ball, S., Smith, A., & Spoon, R. (2021). Fear of COVID-19 changes economic preferences: evidence from a repeated cross-sectional MTurk survey. *Journal of the Economic Science Association*, *7*(2), 103–119. https://doi.org/10.1007/s40881-021-00111-x

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association*, *103*(484), 1481–1495. https://doi.org/10.1198/016214508000000841

Andersson, G., Cuijpers, P., Carlbring, P., Riper, H., & Hedman, E. (2014). Guided Internet-based vs. face-to-face cognitive behavior therapy for psychiatric and somatic disorders: a systematic review and meta-analysis. *World Psychiatry*, *13*, 288–295. https://doi.org/10.1080/16506073.2017.1401115

Andrabi, T., & Das, J. (2017). In aid we trust: Hearts and minds and the Pakistan earthquake of 2005. *Review of Economics and Statistics*, *99*(3), 371–386. https://doi.org/https://doi.org/10.1162/REST_a_00638

Andreoni, J., Kuhn, M. A., & Sprenger, C. (2015). Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior and Organization*, *116*, 451–464. https://doi.org/10.1016/j.jebo.2015.05.018

Andresen, E. M., Malmgren, J. A., Carter, W. B., & Patrick, D. L. (1994). Screening for depression in well older adults: Evaluation of a short form of the CES-D. *American Journal of Preventive Medicine*, *10*(2), 77–84. https://doi.org/10.1016/s0749-3797(18)30622-6

Angrisani, M., Cipriani, M., Guarino, A., Kendall, R., & Ortiz de Zarate, J. (2020). Risk Preferences at the Time of COVID-19: An Experiment with Professional Traders and Students. *SSRN Electronic Journal*, (927). https://doi.org/10.2139/ssrn.3609586

Aragon, F., Bernal, N., Bosch, M., Molina, O., et al. (2022). Covid-19 and economic preferences: Evidence from a panel of cab drivers.

Barak, A., Hen, L., Boniel-Nissim, M., & Shapira, N. (2008). A comprehensive review and a meta-analysis of the effectiveness of internet-based psychotherapeutic interventions. *Journal of Technology in Human Services*, *26*(2-4), 109–160. https://doi.org/10.1080/15228830802094429

Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C., & Sugden, R. (2009). Experimental economics. In *Experimental economics*. Princeton University Press.

Bauer, M., Blattman, C., Henrich, J., Miguel, E., & Mitts, T. (2016). Can War Foster Cooperation? *Journal of Economic Perspectives*, *30*(3), 249–274.

Baumeister, R. F., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology*, *74*(5), 1252–1265. https://doi.org/10.1037/0022-3514.74.5.1252

Becchetti, L., Conzo, G., Conzo, P., & Salustri, F. (2020). Understanding the Heterogeneity of Adverse COVID-19 Outcomes: the Role of Poor Quality of Air and Lockdown Decisions. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3572548

Becchetti, L., Conzo, P., & Romeo, A. (2014). Violence, trust, and trustworthiness: Evidence from a Nairobi slum. *Oxford Economic Papers*, *66*(1), 283–305. https://doi.org/10.1093/oep/gpt008

Becketti, S., Gould, W., Lillard, L., & Welch, F. (1988). The panel study of income dynamics after fourteen years: An evaluation. *Journal of Labor Economics*, *6*(4), 472–492.

Belot, M., Choi, S., Jamison, J. C., Papageorge, N. W., Tripodi, E., & Van den Broek-Altenburg, E. (2020). Six-country survey on covid-19.

Benjamini, Y., Krieger, A. M., & Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, *93*(3), 491–507.

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. https://doi.org/10.1006/game.1995.1027

Bilker, W. B., Hansen, J. A., Brensinger, C. M., Richard, J., Gur, R. E., & Gur, R. C. (2012). Development of Abbreviated Nine-item Forms of the Raven's Standard Progressive Matrices Test. *Assessment*, *19*(3), 354–369. https://doi.org/doi:10.1177/1073191112446655

Böhm, R., Betsch, C., & Korn, L. (2016). Selfish-rational non-vaccination: Experimental evidence from an interactive vaccination game. *Journal of Economic Behavior and Organization*, *131*, 183–195. https://doi.org/10.1016/j.jebo.2015.11.008

Bokern, P., Linde, J., Riedl, A., & Werner, P. (2021). *The effect of the COVID - 19 crisis on economic and social preferences*.

Branas-Garza, P., Jorrat, D. A., Alfonso, A., Espin, A., García, T., & Kovarik, J. (2022). Exposure to the Covid-19 pandemic and generosity. *Royal Society Open Science*. https://doi.org/10.31234/osf.io/6ktuz

Brown, R., Montalva, V., Thomas, D., & Velásquez, A. (2019). Impact of violent crime on risk aversion: Evidence from the mexican drug war. *Review of Economics and Statistics*, *101*(5), 892–904. https://doi.org/https://doi.org/10.1162/rest_a_00788

Bu, D., Hanspal, T., Liao, Y., & Liu, Y. (2020). Risk taking during a global crisis: Evidence from wuhan. *Covid Economics*, *5*, 106–146.

Buso, I. M., De Caprariis, S., Di Cagno, D., Ferrari, L., Larocca, V., Marazzi, F., Panaccione, L., & Spadoni, L. (2020). The effects of COVID-19 lockdown on fairness and cooperation: Evidence from a lablike experiment. *Economics Letters*, *196*, 109577. https://doi.org/10.1016/j.econlet.2020.109577

Cahlikova, J., Cingl, L., & Levely, I. V. (2019). How Stress Affects Performance and Competitiveness across Gender. *Management Science*. https://doi.org/https://doi.org/10.1287/mnsc.2019.3400

Cahlíková, J., & Cingl, L. (2017). Risk preferences under acute stress. *Experimental Economics*, *20*(1), 209–236. https://doi.org/10.1007/s10683-016-9482-3

Callen, M. (2015). Catastrophes and time preference: Evidence from the Indian Ocean Earthquake. *Journal of Economic Behavior and Organization*, *118*, 199–214. https://doi.org/10.1016/j.jebo.2015.02.019

Camerer, C., & Mobbs, D. (2017). Differences in behavior and brain activity during hypothetical and real choices. *Trends in cognitive sciences*, *21*(1), 46–56.

Campos-Mercade, P., Meier, A. N., Schneider, F. H., & Wengström, E. (2021). Prosociality predicts health behaviors during the covid-19 pandemic. *Journal of Public Economics*, *195*, 104367.

Carlbring, P., Andersson, G., Cuijpers, P., Riper, H., & Hedman-Lagerlöf, E. (2018). Internet-based vs. face-to-face cognitive behavior therapy for psychiatric and somatic disorders: an updated systematic review and meta-analysis. *Cognitive Behaviour Therapy*, *47*(1), 1–18. https://doi.org/10.1080/16506073.2017.1401115

Carvalho, B. L. S., Meier, S., & Wang, S. W. (2016). Poverty and Economic Decision-Making : Evidence from Changes in Financial Resources at Payday. *The American Economic Review*, *106*(2), 260–284.

Cassar, A., Grosjean, P., & Whitt, S. (2013). Legacies of violence: Trust and market development. *Journal of Economic Growth*, *18*(3), 285–318. https://doi.org/10.1007/s10887-013-9091-3

Cassar, A., Grosjean, P., & Whitt, S. (2014). Social Preferences of Ex-Combatants: Survey and Experimental Evidence from Postwar Tajikistan. *The Economics of Conflict: Theory and empricial evidence*, 231–262. https://doi.org/10.7551/mitpress/9729.003.0012

Cassar, A., Healy, A., & von Kessler, C. (2017). Trust, Risk, and Time Preferences After a Natural Disaster: Experimental Evidence from Thailand. *World Development*, *94*, 90–105. https://doi.org/10.1016/j.worlddev.2016.12.042

Cettolin, E., Dalton, P. S., Kop, W. J., & Zhang, W. (2019). Cortisol meets GARP: the effect of stress on economic rationality. *Experimental Economics*, (0123456789). https://doi.org/10.1007/s10683-019-09624-z

Chapman, G. B., & Coups, E. J. (1999). Time preferences and preventive health behavior: Acceptance of the influenza vaccine. *Medical Decision Making*, *19*(3), 307–314. https://doi.org/10.1177/0272989X9901900309

Chen, Y., & Yang, D. Y. (2019). The impact of media censorship: 1984 or brave new world? *American Economic Review*, *109*(6), 2294–2332.

Chuang, Y., & Schechter, L. (2015). Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results. *Journal of Development Economics*, *117*, 151–170. https://doi.org/10.1016/j.jdeveco.2015.07.008

Cohn, A., Engelmann, J., Fehr, E., & Maréchal, M. A. (2015). Evidence for countercyclical risk aversion: An experiment with financial professionals. *American Economic Review*, *105*(2), 860–885. https://doi.org/10.1257/aer.20131314

De Quidt, J., Haushofer, J., & Roth, C. (2018). Measuring and bounding experimenter demand. *American Economic Review*, *108*(11), 3266–3302. https://doi.org/10.1257/aer.20171330

Delaney, L., Fink, G., & Harmon, C. (2014). Effects of stress on economic decision-making: Evidence from laboratory experiments. *IZA Discussion Paper*, (8060). http://ssrn.com/abstract=2420705

Dong, D., Xu, X., Xu, W., & Xie, J. (2019). The relationship between the actual level of air pollution and residents' concern about air pollution: Evidence from Shanghai, China. *International Journal of Environmental Research and Public Health*, *16*(23). https://doi.org/10.3390/ijerph16234784

Dong, L., & Bouey, J. (2020). Public mental health crisis during covid-19 pandemic, china. *Emerging infectious diseases*, *26*(7), 1616.

Drichoutis, A. C., & Nayga, R. M. (2021). *On the stability of risk and time preferences amid the COVID-19 pandemic*. Springer US. https://doi.org/10.1007/s10683-021-09727-6

Duan, L., & Zhu, G. (2020). Psychological interventions for people affected by the COVID-19 epidemic. *The Lancet Psychiatry*, *7*(4), 300–302. https://doi.org/10.1016/S2215-0366(20)30073-0

Eckel, C. C., & Grossman, P. J. (2002). Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and Human Behavior*, *23*(4), 281–295. https://doi.org/10.1016/S1090-5138(02)00097-1

Engel, C. (2011). Dictator games: A meta study. *Experimental economics*, *14*(4), 583–610.

Exadaktylos, F., Espín, A. M., & Brañas-Garza, P. (2013). Experimental subjects are not different. *Scientific Reports*, *3*, 1–6. https://doi.org/10.1038/srep01213

Falco, S., & Vieider, F. M. (2018). Environmental forces shape risk preferences. (4), 1–27.

Falk, A., Meier, S., & Zehnder, C. (2013). Do lab experiments misrepresent social preferences? The case of self-selected student samples. *Journal of the European Economic Association*, *11*(4), 839–852. https://doi.org/10.1111/jeea.12019

Fehr, D., Fink, G., & Jack, K. (2019). *Poverty, Seasonal Scarcity and Exchange Asymmetries*.

Fehr, E., & Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior*, *25*(2), 63–87. https://doi.org/10.1016/S1090-5138(04)00005-4

Filipski, M., Jin, L., Zhang, X., & Chen, K. (2019). Living like there's no tomorrow: The psychological effects of an earthquake on savings and spending behavior. *European Economic Review*, *116*, 107–128. https://doi.org/10.1016/j.euroecorev.2019.04.004

Fisman, R., Jakiela, P., & Kariv, S. (2015). How did distributional preferences change during the Great Recession? *Journal of Public Economics*, *128*, 84–95. https://doi.org/10.1016/j.jpubeco.2015.06.001

Fitzgerald, J., Gottschalk, P., & Moffitt, R. (1998). *An analysis of sample attrition in panel data: The michigan panel study of income dynamics* (tech. rep.). National Bureau of Economic Research.

Fleming, D. A., Chong, A., & Bejarano, H. D. (2014). Trust and Reciprocity in the Aftermath of Natural Disasters. *Journal of Development Studies*, *50*(11), 1482–1493. https://doi.org/10.1080/00220388.2014.936395

Flückiger, M., Ludwig, M., & Önder, A. S. (2019). Ebola and State Legitimacy. *Economic Journal*, *129*(621), 2064–2089. https://doi.org/10.1111/ecoj.12638

Friehe, T., & Schildberg-Hörisch, H. (2017). Self-control and crime revisited: Disentangling the effect of self-control on risk taking and antisocial behavior. *International Review of Law and Economics*, *49*(264), 23–32. https://doi.org/10.1016/j.irle.2016.11.001

Fritsche, I., & Jugert, P. (2017). The consequences of economic threat for motivated social cognition and action. *Current Opinion in Psychology*, *18*, 31–36. https://doi.org/10.1016/j.copsyc.2017.07.027

Galizzi, M. M., & Miraldo, M. (2017). Are You What You Eat? Healthy Behaviour and Risk Preferences. *B.E. Journal of Economic Analysis and Policy*, *17*(1). https://doi.org/10.1515/bejeap-2016-0081

Gao, R., Hao, B., Li, H., Gao, Y., & Zhu, T. (2013). Developing simplified chinese psychological linguistic analysis dictionary for microblog. *Brain and Health Informatics*, 359–368.

Gerhardt, H., Schildberg-Hörisch, H., & Willrodt, J. (2017). Does self-control depletion affect risk attitudes? *European Economic Review*, *100*, 463–487. https://doi.org/10.1016/j.euroecorev.2017.09.004

Gillis, M. T., & Hettler, P. L. (2007). Hypothetical and real incentives in the ultimatum game and andreoni's public goods game: An experimental study. *Eastern Economic Journal*, *33*(4), 491–510.

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, *112*(2), 631–645. https://doi.org/10.1162/003355397555217

Gonzalez-Torres, A., & Esposito, E. (2020). Epidemics and Conflict: Evidence from the Ebola Outbreak in Western Africa. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3544606

Gotlib, I. H., & Joormann, J. (2010). Cognition and Depression: Current Status and Future Directions. *Annual Review of Clinical Psychology*, *April 27*(6), 285–312. https://doi.org/10.1146/annurev.clinpsy.121208.131305

Grimalda, G., Buchan, N. R., Ozturk, O. D., Pinate, A. C., Urso, G., & Brewer, M. B. (2021). Exposure to COVID-19 is associated with increased altruism, particularly at the local level. *Scientific Reports*, *11*(1), 1–14. https://doi.org/10.1038/s41598-021-97234-2

Grosjean, P. (2014). Conflict and social and political preferences: Evidence from world war II and civil conflict in 35 european countries. *Comparative Economic Studies*, *56*(3), 424–451. https://doi.org/10.1057/ces.2014.2

Guenther, B., Galizzi, M. M., & Sanders, J. G. (2021). Heterogeneity in Risk-Taking During the COVID-19 Pandemic: Evidence From the UK Lockdown. *Frontiers in Psychology*, *12*(April), 1–15. https://doi.org/10.3389/fpsyg.2021.643653

Hackethal, A., Kirchler, M., Laudenbach, C., Razen, M., & Weber, A. (2022). On the role of monetary incentives in risk preference elicitation experiments. *Journal of Risk and Uncertainty*, 1–25.

Hamermesh, D. S., & Leigh, A. K. (2022). "beauty too rich for use": Billionaires' assets and attractiveness. *Labour Economics*, *76*, 102153.

Han, Y., Lam, J. C. K., Li, V. O. K., Guo, P., & Zhang, Q. (2020). The effects of outdoor air pollution concentrations and lockdowns on Covid-19 infections in Wuhan and other provincial capitals in China. (March), 1–7. https://doi.org/10.20944/preprints202003.0364.v1

Hanaoka, C., Shigeoka, H., & Watanabe, Y. (2018). Do Risk Preferences Change ? Evidence From Panel Data Before and After the Great East Japan Earthquake. *American Economic Journal: Applied Economics*, *10*(2), 298–330. https://doi.org/https://doi.org/10.1257/app.20170048

Harrison, G. W., Hofmeyr, A., Kincaid, H., Monroe, B., Ross, D., Schneider, M., & Swarthout, J. T. (2022). Subjective beliefs and economic preferences during the COVID-19 pandemic. *Experimental Economics*, (0123456789). https://doi.org/10.1007/s10683-021-09738-3

Haushofer, J., Cornelisse, S., Seinstra, M., Fehr, E., Joëls, M., & Kalenscher, T. (2013). No effects of psychosocial stress on intertemporal choice. *PLoS ONE*, *8*(11). https://doi.org/10.1371/journal.pone.0078597

He, G., Chen, Y., Chen, B., Wang, H., Shen, L., Liu, L., Suolang, D., Zhang, B., Ju, G., Zhang, L., Du, S., Jiang, X., Pan, Y., & Min, Z. (2018). Using the Baidu Search Index to Predict the Incidence of HIV/AIDS in China. *Scientific Reports*, *8*(1), 1–10. https://doi.org/10.1038/s41598-018-27413-1

Huang, Y., & Zhao, N. (2020). Generalized anxiety disorder, depressive symptoms and sleep quality during COVID-19 outbreak in China: a web-based cross-sectional survey. *Psychiatry Research*, *288*(April), 112954. https://doi.org/10.1016/j.psychres.2020.112954

Iwamoto, S. K., Alexander, M., Torres, M., Irwin, M. R., Christakis, N. A., & Nishi, A. (2020). Mindfulness Meditation Activates Altruism. *Scientific Reports*, *10*(1), 1–7. https://doi.org/10.1038/s41598-020-62652-1

Kang, Y., Gray, J. R., & Dovidio, J. F. (2014). The nondiscriminating heart: Lovingkindness meditation training decreases implicit intergroup bias. *Journal of Experimental Psychology: General*, *143*(3), 1306–1313. https://doi.org/10.1037/a0034150

Karlsson, M., Nilsson, T., & Pichler, S. (2014). The impact of the 1918 Spanish flu epidemic on economic performance in Sweden. An investigation into the consequences of an extraordinary mortality shock. *Journal of Health Economics*, *36*(1), 1–19. https://doi.org/10.1016/j.jhealeco.2014.03.005

Kettlewell, N. (2019). Risk preference dynamics around life events. *Journal of Economic Behavior and Organization*, *162*, 66–84. https://doi.org/10.1016/j.jebo.2019.04.018

Kling, J., Liebman, J., & Katz, L. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, *75*(1), 83–119.

Koppel, L., Andersson, D., Morrison, I., Posadzy, K., Västfjäll, D., & Tinghög, G. (2017). The effect of acute pain on risky and intertemporal choice. *Experimental Economics*, *20*(4), 878–893. https://doi.org/10.1007/s10683-017-9515-6

Kraemer, M. U., Yang, C. H., Gutierrez, B., Wu, C. H., Klein, B., Pigott, D. M., du Plessis, L., Faria, N. R., Li, R., Hanage, W. P., Brownstein, J. S., Layan, M., Vespignani, A., Tian, H., Dye, C., Pybus, O. G., & Scarpino, S. V. (2020). The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science (New York, N.Y.)*, *497*(May), 493–497. https://doi.org/10.1126/science.abb4218

Leiberg, S., Klimecki, O., & Singer, T. (2011). Short-term compassion training increases prosocial behavior in a newly developed prosocial game. *PLoS ONE*, *6*(3). https://doi.org/10.1371/journal.pone.0017798

Li, C., Chen, L. J., Chen, X., Zhang, M., Pang, C. P., & Chen, H. (2020). Retrospective analysis of the possibility of predicting the COVID-19 outbreak from Internet searches and social media data, China, 2020. *Eurosurveillance*, *25*(10), 1–5. https://doi.org/10.2807/1560-7917.ES.2020.25.10.2000199

Li, K. K. (2020). *How Does the COVID-19 Outbreak Affect People ' s Expectation about the Macroeconomy ?*

Li, K. K., Huang, B., Tam, T., & Hong, Y.-y. (2020). Does the covid-19 pandemic affect people's social and economic preferences? evidence from china. *Evidence from China (July 6, 2020)*.

Li, Z., Liu, T., Zhu, G., Lin, H., Zhang, Y., He, J., Deng, A., Peng, Z., Xiao, J., Rutherford, S., Xie, R., Zeng, W., Li, X., & Ma, W. (2017). Dengue Baidu Search Index data can improve the prediction of local dengue epidemic: A case study in Guangzhou, China (B. Althouse, Ed.). *PLOS Neglected Tropical Diseases*, *11*(3), e0005354. https://doi.org/10.1371/journal.pntd.0005354

Liu, S., Yang, L., Zhang, C., Xiang, Y.-T., Liu, Z., Hu, S., & Zhang, B. (2020). Online mental health services in china during the covid-19 outbreak. *The Lancet Psychiatry*, *7*(4), e17–e18.

Lu, Y., Shi, X., & Zhong, S. (2018). Competitive experience and gender difference in risk preference, trust preference and academic performance: Evidence from Gaokao in China. *Journal of Comparative Economics*, *46*(4), 1388–1410. https://doi.org/10.1016/j.jce.2018.05.002

Mani, A., Mullainathan, S., Shafir, E., & Zhao, J. (2013). Poverty impedes cognitive function. *Science*, *341*(6149), 976–980. https://doi.org/10.1126/science.1238041

Müller, S., & Rau, H. A. (2021). Economic preferences and compliance in the social stress test of the covid-19 crisis. *Journal of Public Economics*, *194*, 104322.

Nikolov, P., Pape, A., Tonguc, O., & Williams, C. (2020). Predictors of social distancing and mask-wearing behavior: Panel survey in seven us states. *arXiv preprint arXiv:2009.13103*.

Page, L., Savage, D. A., & Torgler, B. (2014). Variation in risk seeking behaviour following large losses: A natural experiment. *European Economic Review*, *71*, 121–131. https://doi.org/10.1016/j.euroecorev.2014.04.009

Pennebaker, J. W., & King, L. A. (1999). Linguistic styles: Language use as an individual difference. *Journal of Personality and Social Psychology*, *77*(6), 1296–1312. https://doi.org/10.1037/0022-3514.77.6.1296

Pfefferbaum, B., & North, C. S. (2020). Mental Health and the Covid-19 Pandemic. *New England Journal of Medicine*, 1–2. https://doi.org/DOI:10.1056/NEJMp2008017

Pluchino, A., Inturri, G., Rapisarda, A., Biondo, A. E., Moli, R. L., Zappala', C., Giuffrida, N., Russo, G., & Latora, V. (2020). *A Novel Methodology for Epidemic Risk Assessment: the case of COVID-19 outbreak in Italy*. http://arxiv.org/abs/2004.02739

Prediger, S., Vollan, B., & Herrmann, B. (2014). Resource scarcity and antisocial behavior. *Journal of Public Economics*, *119*, 1–9. https://doi.org/10.1016/j.jpubeco.2014.07.007

Qin, B., Strömberg, D., & Wu, Y. (2017). Why does China allow freer social media? Protests versus surveillance and propaganda. *Journal of Economic Perspectives*, *31*(1), 117–140. https://doi.org/10.1257/jep.31.1.117

Raker, E. J., Zacher, M., & Lowe, S. R. (2020). Lessons from Hurricane Katrina for predicting the indirect health consequences of the COVID-19 pandemic. *Proceedings of the National Academy of Sciences*, *117*(23), 202006706. https://doi.org/10.1073/pnas.2006706117

Said, F., Afzal, U., & Turner, G. (2015). Risk taking and risk learning after a rare event: Evidence from a field experiment in Pakistan. *Journal of Economic Behavior and Organization*, *118*, 167–183. https://doi.org/10.1016/j.jebo.2015.03.001

Schildberg-Hörisch, H., & Strassmair, C. (2012). An experimental test of the deterrence hypothesis. *Journal of Law, Economics, and Organization*, *28*(3), 447–459. https://doi.org/10.1093/jleo/ewq015

Schneider, D., Harknett, K., & McLanahan, S. (2016). Intimate Partner Violence in the Great Recession. *Demography*, *53*(2), 471–505. https://doi.org/10.1007/s13524-016-0462-1

Shachat, J., Walker, M. J., & Wei, L. (2021). How the onset of the covid-19 pandemic impacted pro-social behaviour and individual preferences: Experimental evidence from china. *Journal of Economic Behavior & Organization*, *190*, 480–494.

Snowberg, E., & Yariv, L. (2021). Testing the waters: Behavior across participant pools. *American Economic Review*, *111*(2), 687–719.

Solnit, R. (2010). *A Paradise Built in Hell: The Extraordinary Communities That Arise in Disaster*. Penguin.

Spijkerman, M. P., Pots, W. T., & Bohlmeijer, E. T. (2016). Effectiveness of online mindfulness-based interventions in improving mental health: A review and meta-analysis of randomised controlled trials. *Clinical Psychology Review*, *45*, 102–114. https://doi.org/10.1016/j.cpr.2016.03.009

Steinert, S. (2020). Corona and value change. The role of social media and emotional contagion. *Ethics and Information Technology*, (0123456789). https://doi.org/10.1007/s10676-020-09545-z

Sun, S., Yao, Z., Wei, J., & Yu, R. (2015). Calm and smart? A selective review of meditation effects on decision making. *Frontiers in Psychology*, *6*(July). https://doi.org/10.3389/fpsyg.2015.01059

Sutter, M., Kocher, M. G., Daniela, G. R., & Trautmann, S. T. (2013). Impatience and uncertainty: Experimental decisions predict adolescents' field behavior. *American Economic Review*, *103*(1), 510–531. https://doi.org/10.1257/aer.103.1.510

Thombs, B. D., Bonardi, O., Rice, D. B., Boruff, J. T., Azar, M., He, C., Markham, S., Sun, Y., Wu, Y., Krishnan, A., Thombs-Vite, I., & Benedetti, A. (2020). Curating evidence on mental health during COVID-19: A living systematic review. *Journal of Psychosomatic Research*, (January). https://doi.org/10.1016/j.jpsychores.2020.110113

Thompson, E. R. (2007). Development and validation of an internationally reliable short-form of the Positive and Negative Affect Schedule (PANAS). *Journal of Cross-Cultural Psychology*, *38*(2), 227–242. https://doi.org/10.1177/0022022106297301

Twenge, J. M., Muraven, M., & Tice, D. M. (2004). Measuring state self-control: Reliability, validity, and correlations with physical and psychological stress. *Unpublished Manuscript, San Diego State University*.

Voors, M. J., Nillesen, E. E., Verwimp, P., Bulte, E. H., Lensink, R., & Van Soest, D. P. (2012). Violent conflict and behavior: A field experiment in Burundi. *American Economic Review*, *102*(2), 941–964. https://doi.org/10.1257/aer.102.2.941

Wing, C., Simon, K., & Bello-Gomez, R. A. (2018). Designing Difference in Difference Studies: Best Practices for Public Health Policy Research. *Annual Review of Public Health*, *39*(1), 453–469. https://doi.org/10.1146/annurev-publhealth-040617-013507

Wooldridge, J. M. (2002). Inverse probability weighted m-estimators for sample selection, attrition, and stratification. *Portuguese Economic Journal*, *1*(2), 117–139.

Wooldridge, J. M. (2007). Inverse probability weighted estimation for general missing data problems. *Journal of econometrics*, *141*(2), 1281–1301.

Xiong, H., Liu, J., Huang, J., Huang, S., An, H., Kang, Q., Li, Y., Dou, D., & Wang, H. (2020). Understanding the Collective Responses of Populations to the COVID-19 Pandemic in Mainland China. *medRxiv*, 2020.04.20.20068676. https://doi.org/10.1101/2020.04.20.20068676

Xu, H., Bègue, L., & Bushman, B. J. (2012). Too fatigued to care: Ego depletion, guilt, and prosocial behavior. *Journal of Experimental Social Psychology*, *48*(5), 1183–1186. https://doi.org/10.1016/j.jesp.2012.03.007

Yuan, Q., Nsoesie, E. O., Lv, B., Peng, G., Chunara, R., & Brownstein, J. S. (2013). Monitoring Influenza Epidemics in China with Search Query from Baidu. *PLoS ONE*, *8*(5). https://doi.org/10.1371/journal.pone.0064323

Zhao, S., & Chen, H. (2020). Modeling the epidemic dynamics and control of COVID-19 outbreak in China. *Quantitative Biology*, *8*(1), 11–19. https://doi.org/10.1007/s40484-020-0199-0

Zheng, S., Wang, J., Sun, C., Zhang, X., & Kahn, M. E. (2019). Air pollution lowers Chinese urbanites ’ expressed happiness on social media. *Nature Human Behaviour*, *3*(March). https://doi.org/10.1038/s41562-018-0521-2

Appendix A

Tables

Table A1: Summary of studies on the impact of Covid-19 on economic and social preferences in China

Paper	Population	Sample size	Same individual over time	Time span	Games	Inc	Identification	Change Sig
Shachat et al. (2021)	Students from Wuhan University	N=602 across pre- and post Covid-19 samples	No (main sample); Yes (sub sample)	Baseline: 2019/05; Endline: various samples 01/02/03 2020	Dictator Game; Ultimatum Game; Trust Game; Prisoner's Dilemma Game; Stag Hunt Game; Risk attitudes (Holt and Laury, 2002); Ambiguity attitudes	Yes ($)	Main analysis: repeated cross-sectional data (pre-post analysis); Robustness: panel data on sub-sample, N=92	Yes: Prisoner Dilemma [cooperation] (+); Stag Hunt [risky action] (-); Risk aversion in gains (-); Risk tolerance in losses (-); Ambiguity aversion (+)
Bu et al. (2020)	Students from Wuhan University	N=257	Yes (retention 88%)	Baseline: 2019/10; Endline: 2020/02 – early 2020/03	Hypothetical allocation to a risky investment; Stated risk aversion (risk attitudes)	No	Main analysis: heterogenous exposure (Wuhan, Hubai Province or rest of China); Robustness: panel DiD framework	Yes: Risk investment (-); Risk aversion (+)
Li et al. (2021)	Chinese general population	N= 1872 across pre- and post Covid-19 samples	No (pre: 696; post: 1176)	Baseline: 2019/9–2019/12; Endline: early 2020/3	Trust game; Risk attitudes (Holt and Laury, 2002); Time preferences	Yes ($)	Main analysis: repeated cross-sectional data (pre-post analysis)	Yes: Trust [-]; Trustworthy [+]; Risk aversion [+]; Impatience [+]
Our paper	Students from different Beijing universities	N=793	Yes (Retention 68%)	Baseline: 2019/10 2019/12; Endline: early 2020/03	Joy of Destruction; Take Game; Dictator Game with Third-Party Punishment; Trust Game (hypothetical); Public Good Game (hypothetical); Lottery Choice Task; Investment Game (hypothetical); Time preferences	Yes ($), majority of the games	Main analysis: heterogenous exposure + panel DiD framework	Yes: Joy of destruction (+); Take Game (+)

Table A2: Survey Modules

Group	Measure	Description	Variable construct
Anti-social Behaviour	Joy of Destruction (Abbink & Herrmann, 2011)[$ 23]	Binary decision to anonymously destroy a matched player's endowment as a measure of nastiness.	Dummy which takes the value of 1 if the participant decides to destroy another player's endowment at a cost to him/her-self.
	Take Game (Schildberg-Hörisch & Strassmair, 2012)[$ 23]	Share of endowment taken from a matched player as a measure of theft.	Percentage taken from other player's endowment
	Take Game with Deterrence (Schildberg-Hörisch & Strassmair, 2012) [$ 23]	Share of endowment taken from a matched player with a 40% chance of detection resulting in loss of endowment, as a measure of theft with risk.	Percentage taken from other player's endowment
Pro-social Behaviour	Dictator Game (Fehr & Fischbacher, 2004) [$ 23]	Amount of endowment transferred to a matched player (decision observed by third party).	Percentage invested into a public good.
	Trust game (Berg et al., 1995) [13]	Share of hypothetical endowment entrusted to a hypothetical player, as a measure of trust.	Percentage sent to the other player
	Public-Goods Game (low return) [13]	Share of hypothetical endowment contributed towards a public good, as a measure of cooperation in a low and high return scenario.	Percentage given to the other player
Norm-enforcement	Third-party punishment game (Fehr & Fischbacher, 2004) [$ 23]	Amount of costly punishment imposed on a matched player based on the amount transferred by the matched player in a dictator game.	Binary variable: Takes the value of 1 if a participant is willing to punish when the dictator transfers zero credits to the other player. Extent variable: Amount punished at a cost ratio of 1 Yuan for every 3 Yuan deducted.
Risk & Time Preferences	CRRA coefficient (Eckel & Grossman, 2002) [$ 123]	Choice between six lotteries (50/50 odds) increasing in variance, absolute pay-off and riskiness.	Coefficient of relative risk aversion midpoints (CRRA)
	Risk aversion (Gneezy & Potters, 1997) [13]	Share of hypothetical endowment not invested in a lottery (50/50 odds).	Percentage invested into a lottery
	Present Bias (Andreoni et al., 2015) [$ 23]	Individual β parameter derived from 24 budget lines across 4 timeframes	Dummy which takes the value of 1 if present biasedness parameter beta is greater than 1.
	Time Discounting (Andreoni et al., 2015) [$ 23]	Individual δ parameter derived from 24 budget lines across 4 timeframes	Discount rate (parameter delta)
Cognitive Ability & Well-being	Raven's Standard Progressive Matrices (Bilker et al., 2012) [$ 23]	Cognitive ability measured by the number of correctly completed puzzles (out of 9).	Score between 0 and 9.
	Depression (Andresen et al., 1994) [123]	Depression score calculated using the Centre for Epidemiological Studies Depression Scale Short-form (CESD-10).	Continuous variable: Depression score between 0 and 30 (sum of ten items). Binary variable: Takes the value of 1 if depression score is greater than 10.
	Positive Affect (Thompson, 2007) [23]	Assessment of mood on the day of the survey using the international Short-form of the Positive and Negative Affect Schedule (PANAS-ISF)	Positive affect score between 5 and 25 (sum of five items).
	Negative Affect (Thompson, 2007) [23]		Negative affect score between 5 and 25 (sum of five items).
	Life Satisfaction[23]	Self-assessed general life satisfaction	Likert scale between 1 and 5
	Happiness[23]	Self-assessed general happiness (enjoying life)	Likert scale between 1 and 5
	Eudaemonic Well-being[23]	Self-assessed meaningfulness of life	Likert scale between 1 and 5
	Depletion [23]	Five-item depletion scale adapted from Twenge et al. (2004).	Score between – 7 and + 11
	General health[123]	Self-assessed general health status	Likert scale between 1 and 5

Note: [$] Incentivised tasks; [13] Included in Survey Wave 1 and 3; [23] Included in Survey Wave 2 and 3; [123] Included in Survey Wave 1, 2 and 3

Table A3: Baidu Search Terms

No.	English Translation
1	Coronavirus disease (pneumonia caused by the novel coronavirus)
2	Novel coronavirus
3	Real-time Situation of COVID-19
4	The Latest News about pneumonia caused by COVID-19
5	The latest news about COVID-19
6	Coronavirus disease outbreak situation
7	Confirmed cases
8	New cases
9	New cases of pneumonia caused by the novel coronavirus
10	N95 masks
11	How often change n95 mask
12	Antibacterial gel
13	What are the symptoms of pneumonia caused by the novel coronavirus
14	Symptoms of the novel coronavirus
15	Symptoms of coronavirus disease (pneumonia caused by the novel coronavirus)
16	Dry cough
17	What is the temperature of COVID-19
18	Is dry cough a symptom of COVID-19
19	Fever clinic
20	Early symptoms of COVID-19

Table A4: Threats to Identification - Formal Assessment

	*Ln*Cases		Baidu Index		Sentiment Index	
	(1)	(2)	(3)	(4)	(5)	(6)
Immigration Rate (Wuhan)	0.912***	0.844***	0.450***	0.343***	0.173*	0.196**
	(0.116)	(0.121)	(0.091)	(0.089)	(0.098)	(0.097)
Population Density		0.088		0.142*		-0.036
		(0.108)		(0.080)		(0.087)
Number of Hospitals		0.017		-0.024		0.174**
		(0.065)		(0.044)		(0.076)
Lockdown Duration Index		-0.088		-0.045		-0.074
		(0.066)		(0.060)		(0.081)
GDP per Capita		0.128		0.179**		-0.068
		(0.120)		(0.075)		(0.122)
Health Expenditure Share		-0.055		-0.074*		-0.023
		(0.052)		(0.038)		(0.057)
Annual Average AQI		-0.067		-0.007		0.224**
		(0.068)		(0.072)		(0.096)
Constant	3.378***	3.200***	13.316***	13.012***	0.304***	0.312**
	(0.017)	(0.112)	(0.013)	(0.084)	(0.014)	(0.140)
R^2	0.711	0.734	0.443	0.515	0.420	0.473
Province Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	183	183	183	183	179	179

Note: Table shows results from simple OLS regressions to assess city-level determinants of virus exposure. Dependent variables are *Ln*Cases, Baidu Index and Negative Sentiment Index. *Ln*Cases is the logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey. Baidu Index is an index of city-level COVID-19 concern based on Baidu search volume indices for 20 virus-related keywords. Sentiment Index is the city-level average share of negative expressed emotions via social media. All explanatory variables are z-scored and all regressions include province fixed effects. Robust standard errors in parentheses.

* p < 0.1, ** p < 0.05, *** p < 0.01.

Table A5: Difference-in-difference Analysis: Parallel Survey Trends October – December 2019

	(1) Risk Aversion	(2) Depression	(3) General Health
December 2019	0.011	0.624***	-0.140
	(0.105)	(0.111)	(0.146)
December 2019 × *Ln*Cases	-0.010	-0.039	0.036
	(0.035)	(0.039)	(0.051)
Number of Individuals	522	522	522
Waves	2	2	2
Observations	1044	1044	1044

Note: Difference in Difference Analysis using fixed effects OLS regressions to test for pre-outbreak parallel trends between October and December 2019. Dependent variables are standardized (see details of measures in Table A2). *Ln*Cases is the logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey. December 2019 is a dummy referring to the second survey wave (prior to the virus outbreak). Standard errors in parentheses

* p < 0.1, ** p < 0.05, *** p < 0.01.

Table A6: Pre-Outbreak Exposure Analysis

	(1) Co-op	(2) Trust	(3) Altruism	(4) Punish (Binary)	(5) Punish (Extent)	(6) Destruction	(7) Taking	(8) Taking (Det.)	(9) Risk Taking	(10) Present Bias	(11) Discounting
*Ln*Cases	-0.037	-0.022	-0.049	-0.033	-0.058	-0.049	0.026	0.025	0.059	0.067**	0.062*
	(0.050)	(0.037)	(0.039)	(0.055)	(0.045)	(0.030)	(0.029)	(0.027)	(0.036)	(0.034)	(0.035)
R^2	0.002	0.001	0.003	0.001	0.005	0.003	0.001	0.001	0.005	0.006	0.008
Observations	522	522	522	522	522	522	522	522	522	513	513

Note: OLS analysis of pre-outbreak exposure. Standard errors clustered at the city level in parenthesis. *Ln*Cases is the logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey.

Table A7: Basic Characteristics of Participants by Exposure Tercile

	Full Sample	Mildly Exposed	Moderately Exposed	Highly Exposed	P-val
Gender	0.82	0.79	0.86	0.81	0.22
Age	19.85	19.73	19.78	20.03	0.23
Year of Study	2.56	2.47	2.58	2.61	0.82
Hukou	0.21	0.24	0.20	0.18	0.35
Only Child	0.64	0.58	0.63	0.73	0.01
Chronic Illness	0.09	0.08	0.09	0.10	0.81
General Health	3.76	3.75	3.75	3.78	0.77
Observations	522.00	179.00	169.00	174.00	.

Note: Mild, moderate and high exposure categories are based upon terciles of the number of cumulative confirmed cases at the city-level per million population officially reported on the date of the third survey. Mildly exposed (0-7 Cases per million population), Moderately exposed (8-30 Cases per million population), Highly exposed (>30 Cases per million population). P-val refers to the p-value obtained from tests of equality of means across all three categories of exposure using Anova and proportions using chi2-test.

Table A8: Difference-in-difference Analysis: Anti-social Behaviour

	(1) Joy of Destruction	(2) Take Game	(3) Take Game (Det.)	(4) Anti-sociality Index
Panel A				
Post × *Ln*Cases	0.237***	0.222***	0.179*	0.212***
	(0.076)	(0.081)	(0.101)	(0.058)
	[0.014]	[0.028]	[0.311]	[0.005]
Panel B				
Post × Baidu Index	0.301**	0.017	0.067	0.128*
	(0.124)	(0.083)	(0.114)	(0.072)
	[0.123]	[1.000]	[1.000]	[0.299]
Panel C				
Post × Sentiment Index	0.178**	0.047	0.039	0.088
	(0.084)	(0.070)	(0.089)	(0.057)
	[0.887]	[1.000]	[1.000]	[0.887]
Number of Individuals	522	522	522	522
Observations	1,044	1,044	1,044	1,044

Note: Difference in differences analysis using fixed effects OLS regressions accounting for attrition using Inverse Probability Weighting (IPW). Standard errors clustered at the individual level in parenthesis. Multiple testing adjusted False Discovery Rate (FDR) q-values in square brackets. *Post* × *Ln*Cases is the interaction of logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey with a post-outbreak indicator. *Post* × Baidu is the interaction of Baidu Search Index with a post-outbreak indicator. *Post* × Sentiment is the interaction of the Negative Sentiment Index with a post-outbreak indicator. All regressions include individual fixed effects, time-varying city-level controls for average immigration rate from Wuhan (20-23 Jan), number of hospitals per million inhabitants, health expenditure as a share of total expenditure, population density, GDP per capita and province-specific time trends. The dependent variable in column (4) is an index for anti-sociality based on the average of the z-scores of all three anti-social outcome variables.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Difference-in-difference Analysis: Risk & Time Preferences

	(1) Risk Aversion	(2) Risk Taking	(3) Present Bias	(4) Discounting
Panel A				
Post × *Ln*Cases	0.074	-0.144	-0.174	-0.082
	(0.080)	(0.125)	(0.124)	(0.124)
	[0.577]	[0.549]	[0.423]	[0.680]
Panel B				
Post × Baidu Index	0.180**	0.066	-0.088	0.054
	(0.090)	(0.117)	(0.141)	(0.144)
	[0.249]	[1.000]	[1.000]	[1.000]
Panel C				
Post × Sentiment Index	0.064	-0.065	0.026	0.221
	(0.078)	(0.090)	(0.124)	(0.141)
	[1.000]	[1.000]	[1.000]	[0.887]
Number of Individuals	522	522	522	522
Observations	1,044	1,044	1,044	1,044

Note: Difference-in-difference analysis using fixed effects OLS regressions accounting for attrition using Inverse Probability Weighting (IPW). Standard errors clustered at the individual level in parenthesis. Multiple testing adjusted False Discovery Rate (FDR) q-values in square brackets. *Post* × *Ln*Cases is the interaction of logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey with a post-outbreak indicator. *Post* × Baidu is the interaction of Baidu Search Index with a post-outbreak indicator. *Post* × Sentiment is the interaction of the Negative Sentiment Index with a post-outbreak indicator. All regressions include individual fixed effects, time-varying city-level controls for average immigration rate from Wuhan (20-23 Jan), number of hospitals per million inhabitants, health expenditure as a share of total expenditure, population density, GDP per capita and province-specific time trends.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Difference-in-difference Analysis: Pro-social Behaviour

	(1) Cooperation	(2) Trust	(3) Altruism	(4) Pro-sociality Index
Panel A				
Post × *Ln*Cases	-0.034	-0.155	0.003	-0.062
	(0.112)	(0.096)	(0.109)	(0.059)
	[0.959]	[0.353]	[1.000]	[0.577]
Panel B				
Post × Baidu Index	0.139	-0.050	-0.283***	-0.065
	(0.125)	(0.117)	(0.099)	(0.068)
	[1.000]	[1.000]	[0.075]	[1.000]
Panel C				
Post × Sentiment Index	0.140	-0.004	-0.002	0.045
	(0.088)	(0.083)	(0.086)	(0.051)
	[0.887]	[1.000]	[1.000]	[1.000]
Number of Individuals	522	522	522	522
Observations	1,044	1,044	1,044	1,044

Note: Difference in differences analysis using fixed effects OLS regressions accounting for attrition using Inverse Probability Weighting (IPW). Standard errors clustered at the individual level in parenthesis. Multiple testing adjusted False Discovery Rate (FDR) q-values in square brackets. *Post* × *Ln*Cases is the interaction of logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey with a post-outbreak indicator. *Post* × Baidu is the interaction of Baidu Search Index with a post-outbreak indicator. *Post* × Sentiment is the interaction of the Negative Sentiment Index with a post-outbreak indicator. All regressions include individual fixed effects, time-varying city-level controls for average immigration rate from Wuhan (20-23 Jan), number of hospitals per million inhabitants, health expenditure as a share of total expenditure, population density, GDP per capita and province-specific time trends. The dependent variable in column (4) is an index for pro-sociality based on the average of the z-scores of all three pro-social outcome variables.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Difference-in-difference Analysis: Norm Enforcement Behaviour

	(1) Punishment (Binary)	(2) Punishment (Extent)	(3) Norm-enforcement Index
Panel A			
Post × *Ln*Cases	0.070	0.038	0.076
	(0.105)	(0.099)	(0.084)
	[0.680]	[0.947]	[0.577]
Panel B			
Post × Baidu Index	-0.018	-0.088	-0.041
	(0.104)	(0.127)	(0.104)
	[1.000]	[1.000]	[1.000]
Panel C			
Post × Sentiment Index	-0.034	-0.026	0.001
	(0.096)	(0.092)	(0.076)
	[1.000]	[1.000]	[1.000]
Number of Individuals	522	522	522
Observations	1,044	1,044	1,044

Note: Difference in differences analysis using fixed effects OLS regressions accounting for attrition using Inverse Probability Weighting (IPW). Standard errors clustered at the individual level in parenthesis. Multiple testing adjusted False Discovery Rate (FDR) q-values in square brackets. *Post* × *Ln*Cases is the interaction of logged number of cumulative confirmed cases at the city-level per million inhabitants officially reported on the date of the third survey with a post-outbreak indicator. *Post* × Baidu is the interaction of Baidu Search Index with a post-outbreak indicator. *Post* × Sentiment is the interaction of the Negative Sentiment Index with a post-outbreak indicator. All regressions include individual fixed effects, time-varying city-level controls for average immigration rate from Wuhan (20-23 Jan), number of hospitals per million inhabitants, health expenditure as a share of total expenditure, population density, GDP per capita and province-specific time trends. The dependent variable in column (3) is an index for norm enforcement based on the average of the z-scores of three punishment decisions (punishment extent if the dictator gives 0, 2 or 4 Yuan).

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: Mechanism: Improvement in Mental Health

	Anti-sociality Index				
	(1)	(2)	(3)	(4)	(5)
Post=1 × *Ln*Cases	0.215***	0.618	0.024	0.386	-0.157
	(0.058)	(0.403)	(0.430)	(0.457)	(0.354)
Post=1 × Positive Affect × *Ln*Cases		-0.028			
		(0.029)			
Post=1 × Life Satisfaction × *Ln*Cases			0.054		
			(0.113)		
Post=1 × Happiness × *Ln*Cases				-0.049	
				(0.122)	
Post=1 × Meaningfulness × *Ln*Cases					0.077
					(0.089)
R^2	0.115	0.207	0.185	0.194	0.202
Number of Individuals	522	522	522	522	522
Observations	1044	1044	1044	1044	1044

Note: This table reports OLS estimates of equation (2) where the dependent variable is an index of anti-social behaviour. LnCases is the logged number of confirmed cases at the city-level per million inhabitants reported on the date of the third survey. All triple interaction terms provide estimates for potential mechanisms

* p < 0.1, ** p < 0.05, *** p < 0.01.

Table A13: Wave 3 Survey Data Analysis

	(1) Wash Hands	(2) Social Distancing	(3) Stay Home	(4) Use Face Mask	(5) Avoid Touch	(6) Virus Knowledge Index
Pro-sociality	0.158**	0.015	0.052	0.043	0.077	0.030
	(0.069)	(0.046)	(0.043)	(0.036)	(0.096)	(0.193)
Anti-sociality	-0.017	0.070	0.034	0.044	-0.007	-0.099
	(0.067)	(0.043)	(0.045)	(0.032)	(0.103)	(0.210)
Norm-enforcement	0.103	-0.073	-0.093	-0.050	0.035	0.050
	(0.065)	(0.059)	(0.062)	(0.056)	(0.084)	(0.157)
Observations	522	522	522	522	522	522

Note: This table is based on 18 OLS regressions (all coefficient estimates presented in this table come from individual regressions). Dependent variables are based on individual survey responses to questions on frequency of protective behaviours collected in the third survey and an index of virus-related knowledge. Each OLS regressions includes additional controls for age, gender, a dummy for being an only child, hukou registration, general health status, depression score, risk aversion, a categorical variable for political membership, an index for perceived virus-risk and city fixed effects.

* p < 0.1, ** p < 0.05, *** p < 0.01.

Figures

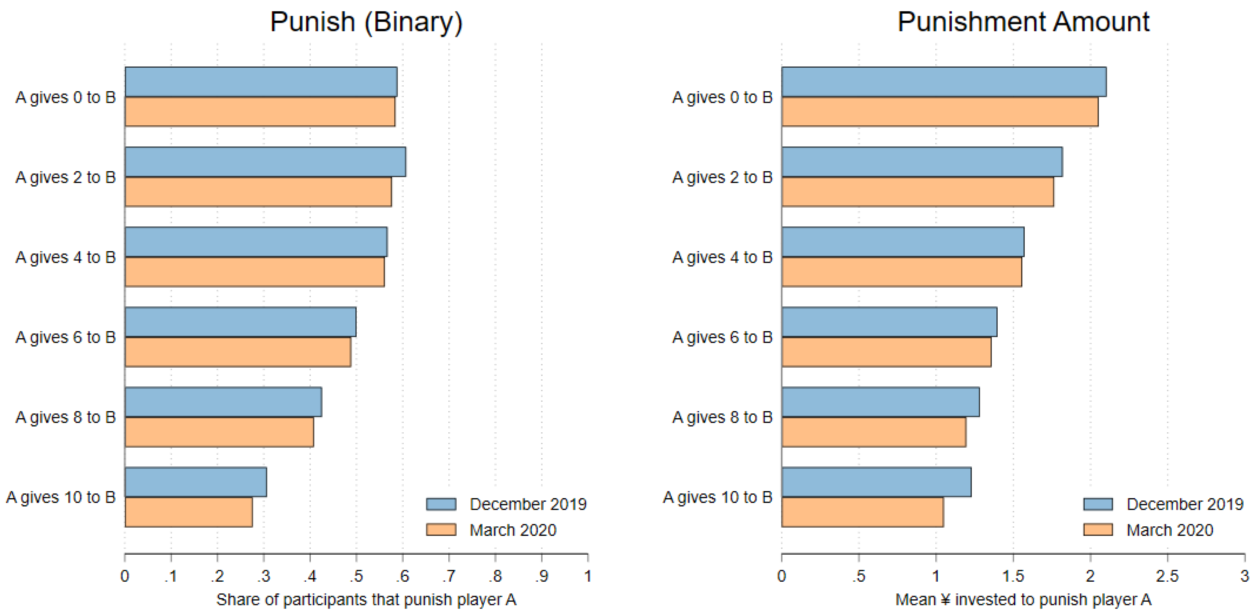

Figure A1: ***Norm-enforcement preferences in pre (December 2019) and post-outbreak waves (March 2020).***

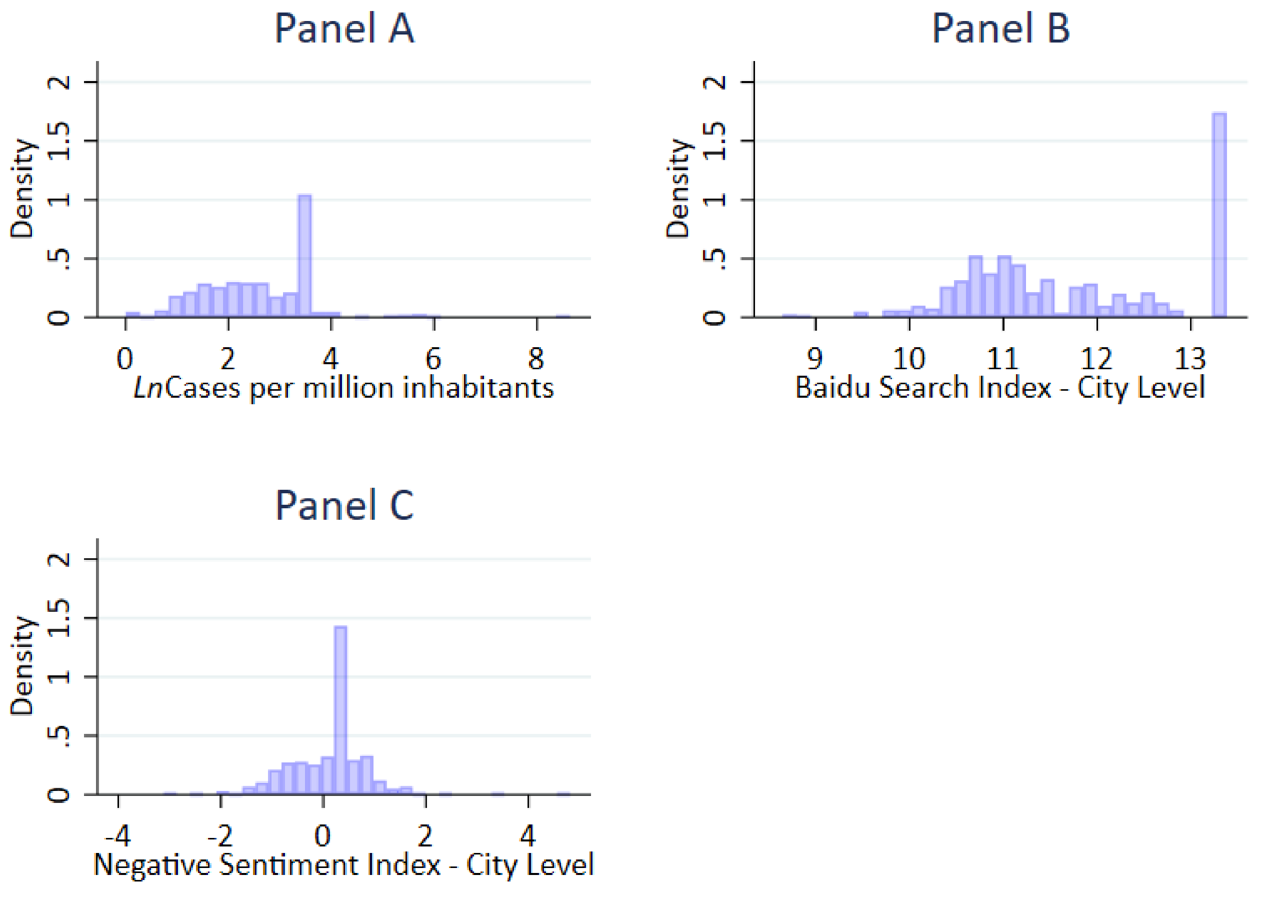

Figure A2: ***Frequency Distribution of Sample Exposure to COVID-19***

Note: Exposure measured as the logged number of cumulative confirmed city-level cases per million inhabitants on the date of the third survey (Panel A) and Baidu Search Index (Panel B) and Negative Sentiment Index (Panel C).

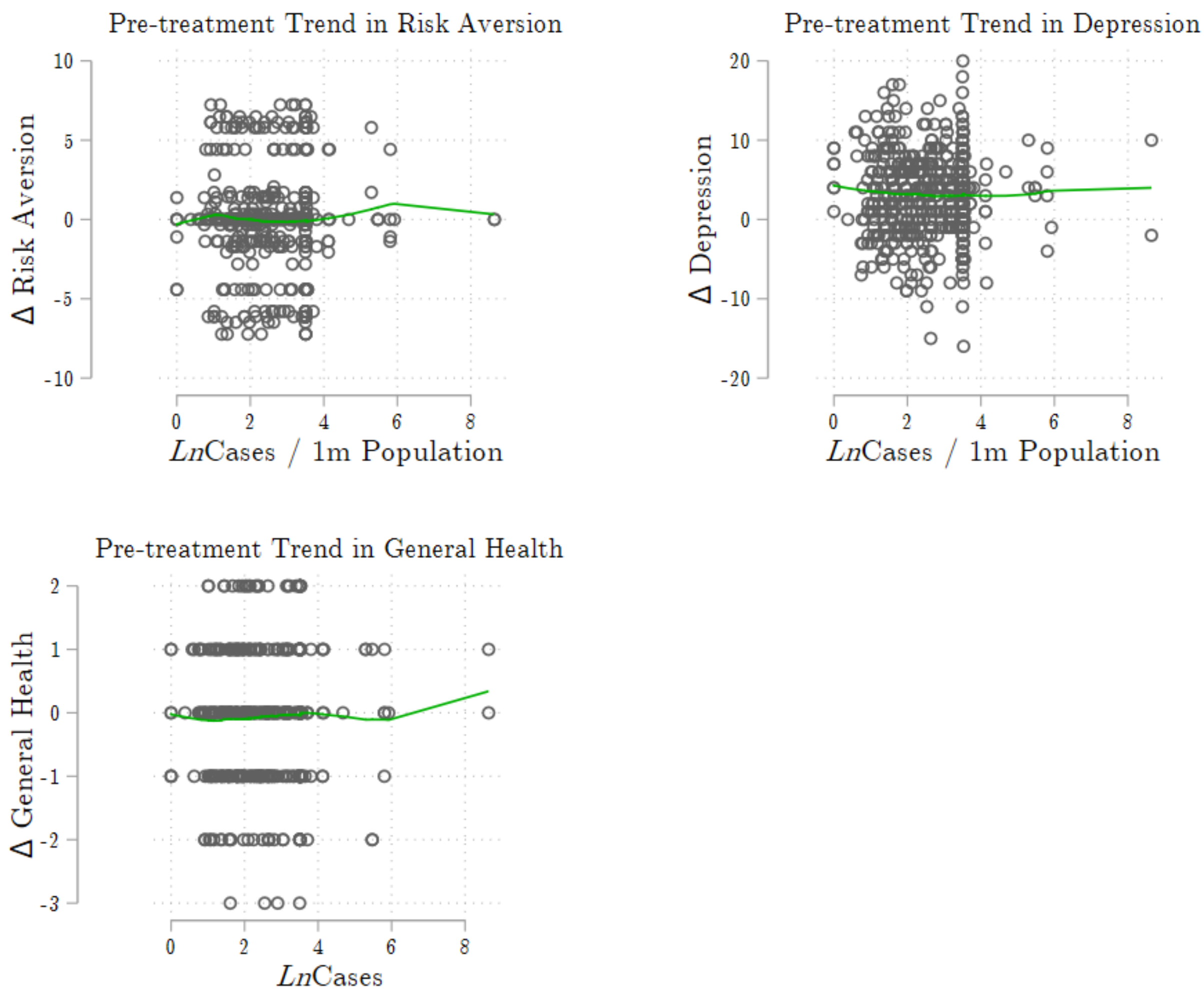

Figure A3: ***Pre-trend Visual Assessment on Survey Outcomes***

Note: Plots show change in Risk aversion measured via the CRRA interval midpoints from a lottery choice task, Depression measured using the Centre for Epidemiological Studies Depression Scale Short-form (CESD-10) and General health assessed via self-reported health condition between wave 1 and wave 2 of the panel survey. Both surveys took place prior to the outbreak of COVID-19.

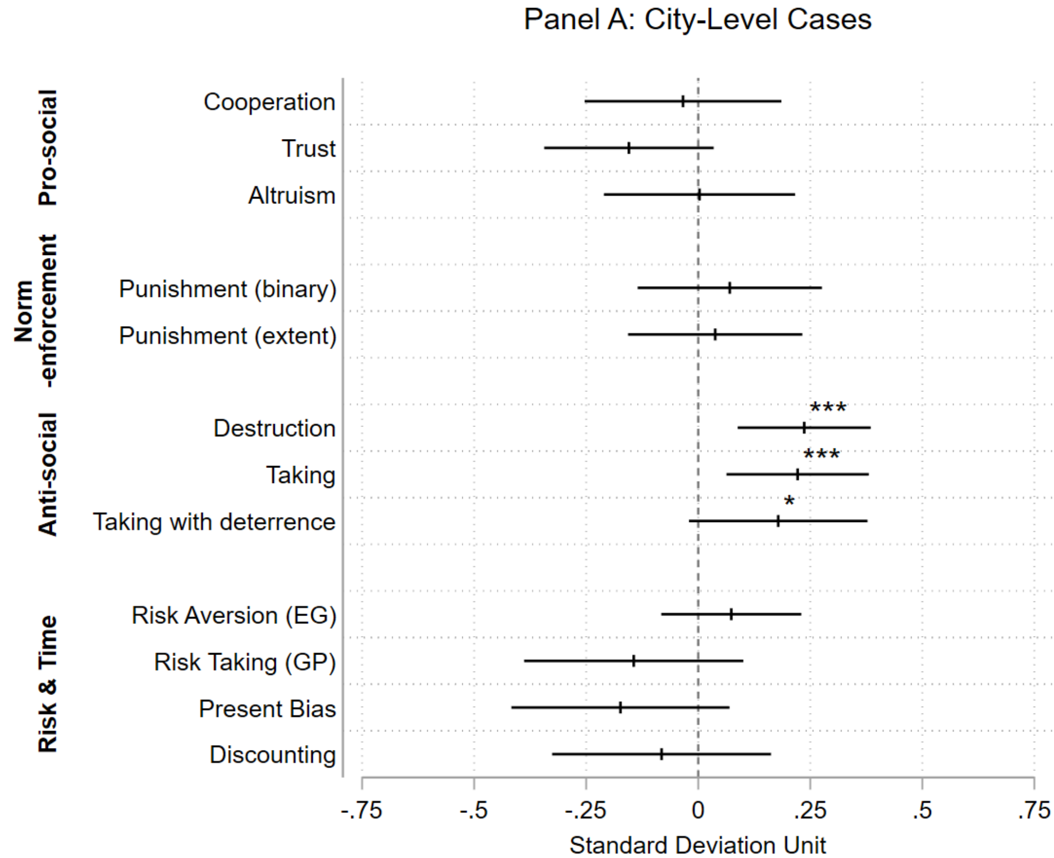

Figure A4: ***Treatment Effects - Official COVID-19 Infection Data***

Note: X-axis plots the estimated coefficient for a 1-unit increase in the Log of Cases per million Inhabitants. Significance stars *** p<0.01, ** p<0.05, * p<0.1 based on p-values estimated using cluster-robust standard errors, clustered at the individual level. All models are estimated with IPW to account for differential attrition. Number of observations *N*=1044.

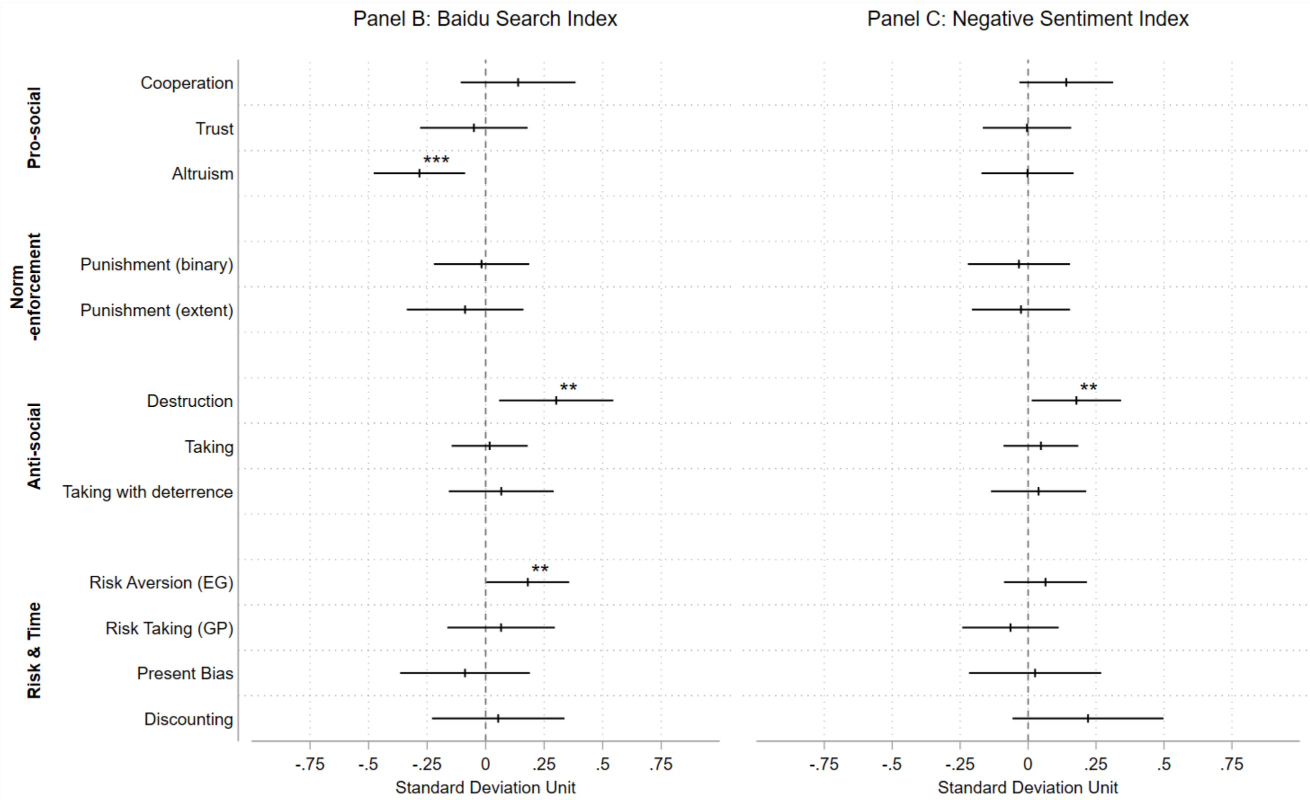

Figure A5: ***Treatment effects - Social Media Data***

Note: X-axis plots the estimated coefficient for a 1-unit increase in the log of the Baidu Index measuring concern about COVID-19 (Panel B) and a 1-unit increase in the Negative Sentiment Index (Sina Weibo) (Panel C). Significance stars *** p<0.01, ** p<0.05, * p<0.1 based on p-values estimated using cluster-robust standard errors, clustered at the individual level. All models are estimated with IPW to account for differential attrition. Number of observations Panel B: N=1044, Panel C: N=1036.

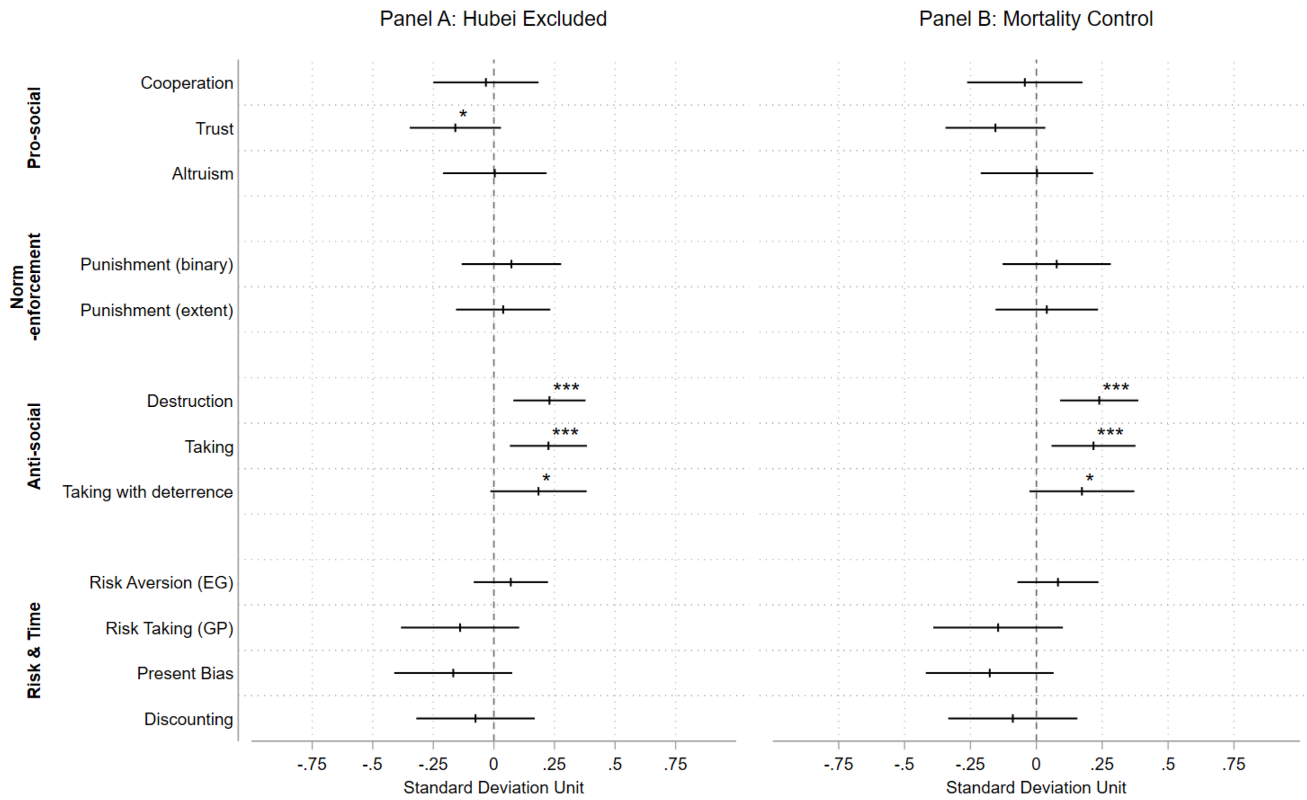

Figure A6: ***Sensitivity Analysis***

Note: Panel A: Treatment effects when Hubei Province participants (10 individuals) are excluded from the analysis (Number of Individuals N = 512). Panel B: Treatment effects when city-level Mortality is entered as a control (Number of Individuals N = 522). X axis plots the estimated coefficient for a 1-unit increase in the Log of Cases per million population. Significance stars (*** p<0.01, ** p<0.05, * p<0.1) based on p-values estimated using cluster-robust standard errors, clustered at the individual level. Individual regression result tables are available upon request

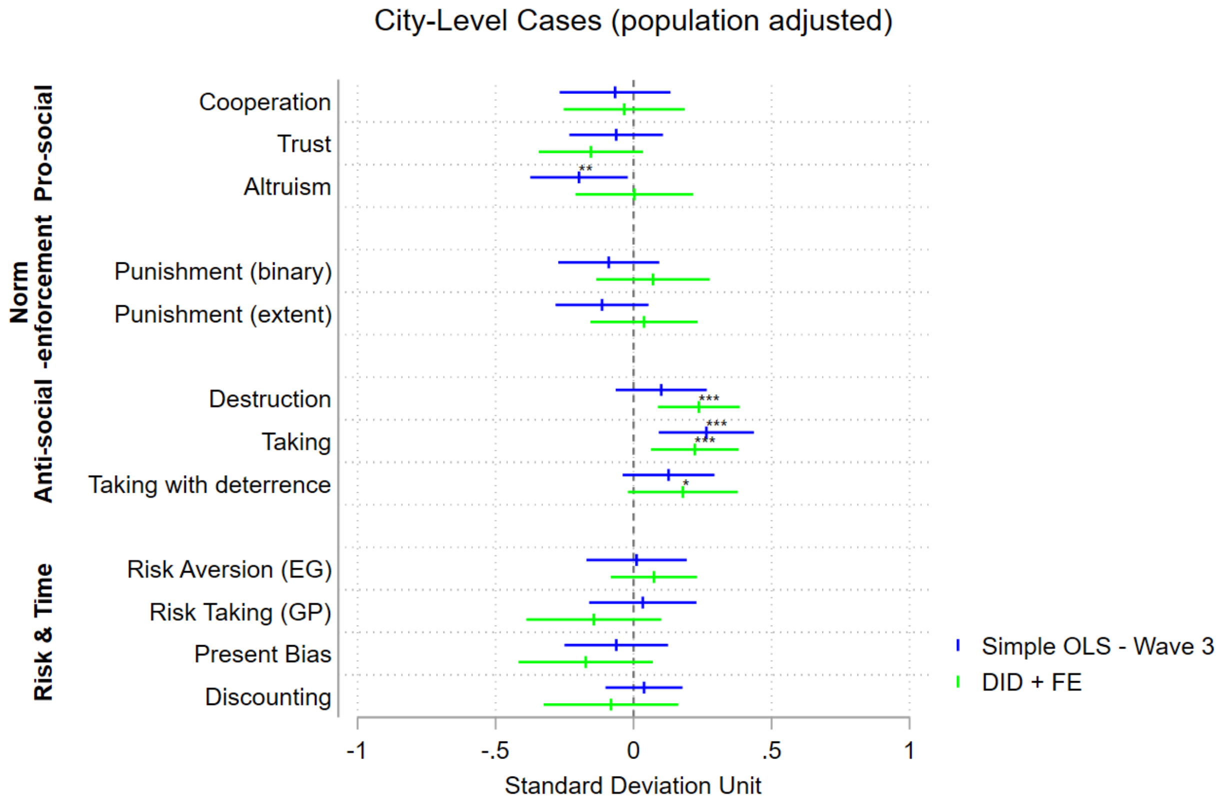

Figure A7: ***Comparing Estimates from simple OLS (Post-outbreak data only) and the preferred DID model in equation*** (1)

Appendix B

Data Collection & Attrition

The experiment was initially designed as a two-wave experiment, with data collection taking place in October and December 2019. The short survey administered in October (Wave 1) was used to build an initial subject pool, with the objective to collect socio-demographic information and key preference measures relevant to our original research question. This information allowed us to implement a stratified randomisation procedure prior to Wave 2. In March 2020 we re-contacted all students from the original subject pool with a follow-up survey (Wave 3), designed around the new objective to assess the stability of preferences after COVID-19.

In all three waves, the entire data collection was conducted via the Chinese messaging app WeChat. Research Assistants were trained to contact students via WeChat, send survey links on pre-specified dates and administer payment directly to participants' WeChat Wallets. Due to the lack of reliable and trustworthy online crowdsourcing platforms in China (such as Amazon MTurk or Prolific), using WeChat is a common procedure to maintain student subject pools for research purposes. However, one can expect high levels of attrition with this form of data collection (see e.g. Chen & Yang, 2019). To minimise attrition, the original study design included an additional prize-draw for ten 100Y bonus payments, for which participants were eligible only if they completed both initial survey waves (1 and 2). For Wave 3, no such incentive was possible.

From the initial sample (N=793) recruited in October 2019, we exclude 3 individuals for which no city location is available, 4 individuals who live in Hong Kong, Macau and Taiwan, and 15 individuals who completed Waves 1 and 3, but skipped Wave 2 (the main experimental survey in the pre-outbreak period). The remaining sample of 771 individuals serves as our starting point for the following attrition analysis. Table B1 shows the number of participants in each wave, the number of attrited individuals as well as the share of attrition for each survey wave. As attrition poses a potential threat to producing unbiased estimation results, the analysis below will carefully consider the potential impact of attrition.

Table B1: Attrition Share across Survey Waves

Wave	Participants	Attrition	Attrition Share
1	771		
2	646	125	16.21%
3	522	124	19.20%

Table B1 shows that attrition rates are high across the three waves in our data (16% between Wave 1 and 2, 19% between Wave 2 and 3), however, comparable to previous research conducted via WeChat surveys.

First, we explore the patterns of attrition in our data by comparing attritors vs. the non-attritors using a rich set of sociodemographic characteristics, heath indicators and economic preferences collected at baseline (Wave 1). We further explore city-level variables (representative of the respondents' hometown) as well as city-level confirmed cumulative cases (log-transformed) reported in the respondents' hometown on 14th March 2020, our primary treatment variable. Tables B2 and B3 present the results from this exercise for attrition between Waves 1 and 2 and 2 and 3, respectively. Columns (1) and (2) show the means of both attrition and non-attrition samples columns (3) and (4) report the difference in means and a p-value derived from a t-test for the equality of means. In Table B2 we first focus on attrition that occurred between Waves 1 and 2.

Table B2: Difference between Attrited and Non-Attrited: Wave 1 to 2

Variable	(1) Non-Attritors	(2) Attritors	(3) Difference	(4) p-value
*Ln*Cases	2.607	2.633	0.026	(0.816)
Pay-off (Wave 1)	26.296	21.376	-4.920***	(0.003)
Cooperation	4,487.288	4,337.112	-150.176	(0.638)
Trust	48.260	47.728	-0.532	(0.831)
Risk Aversion	3.009	2.988	-0.021	(0.940)
Risk Taking	7,805.771	7,990.112	184.341	(0.670)
Age	19.920	20.488	0.568***	(0.001)
Female	0.774	0.704	-0.070*	(0.092)
Rural Hukou	0.204	0.232	0.028	(0.487)
Only Child	0.655	0.664	0.009	(0.843)
General Health	3.771	3.632	-0.139*	(0.097)
Depression Score	9.065	10.096	1.031**	(0.047)
Economics Major	0.455	0.600	0.145***	(0.003)
Chronic Illness	0.091	0.104	0.013	(0.656)
Perseverance	2.532	2.552	0.020	(0.751)
Prosocial Trait	0.221	0.208	-0.013	(0.741)
Competitiveness	12.673	11.376	-1.297***	(0.001)
Immigration Rate (Wuhan)	0.497	0.568	0.071	(0.167)
Hospitals (City)	30.615	30.634	0.019	(0.991)
GDP per Capita (City)	90.283	84.338	-5.945	(0.181)
Health Expenditure Share (City)	0.077	0.081	0.004	(0.167)
Observations	646	125	771	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Statistically significant differences in means of attritors and non-attritors show that there are systematic attrition patterns between the first two survey waves. With respect to sociodemographic characteristics, we observe that individuals who are older, identify as male, major in economics, revealed a lower willingness to compete and earned less in the baseline survey are all more likely to leave the sample at Wave 2. With respect to health variables, individuals with a higher depression scores and a lower general health score are significantly more likely to attrite.

In Table B3 we focus on attrition that occurred between Waves 2 and 3. In addition to the individual sociodemographic characteristics, health variables and city-level variables observed at baseline, we further include variables measured in Wave 2. These variables include our main outcome measures of anti-social and norm-enforcement behaviour as well as additional subjective well-being and health indicators. Here we see that older participants and men are significantly more likely to be in the attrition sample. Attritors are less risk-averse and show slightly more anti-social behaviour in the Take-Game. Individuals with higher depression scores and lower cognitive ability are more likely to attrite.

Table B3: Difference between Attrited and Non-attrited: Wave 2 to 3

Variable	(1) Non-Attritors	(2) Attritors	(3) Difference	(4) p-value
*Ln*Cases	2.596	2.654	0.059	(0.590)
Pay-off (Wave 1)	26.205	26.677	0.472	(0.795)
Pay-off (Wave 2)	32.518	29.390	-3.127	(0.157)
Pay-off Difference (W1-W2)	6.313	2.713	-3.600	(0.215)
Cooperation	4,472.960	4,547.605	74.645	(0.817)
Trust	48.056	49.121	1.065	(0.675)
Risk Aversion	3.008	2.687	-0.321	(0.257)
Risk Taking	7,605.580	8,648.508	1,042.928**	(0.015)
Altruism	3.939	3.774	-0.165	(0.632)
Punishment (Binary)	0.588	0.573	-0.016	(0.753)
Punishment (Extent)	2.103	2.081	-0.023	(0.923)
Destruction	0.157	0.145	-0.012	(0.742)
Taking	9.983	11.089	1.106*	(0.079)
Taking (Deterrence)	9.362	10.500	1.138*	(0.095)
Age	19.849	20.218	0.369**	(0.017)
Female	0.818	0.589	-0.229***	(0.000)
Rural Hukou	0.207	0.194	-0.013	(0.741)
Only Child	0.644	0.702	0.058	(0.223)
General Health	3.699	3.653	-0.046	(0.575)
Depression Score	12.153	13.097	0.944*	(0.088)
Economics Major	0.443	0.508	0.066	(0.188)
Chronic Illness	0.088	0.105	0.017	(0.562)
Perseverance	2.522	2.573	0.051	(0.438)
Prosocial Trait	0.230	0.185	-0.044	(0.285)
Competitiveness	12.795	12.161	-0.634	(0.125)
Raven Score	6.548	6.137	-0.411***	(0.006)
Sleep Quality	7.739	7.597	-0.143	(0.364)
Life Satisfaction	3.404	3.331	-0.074	(0.438)
Immigration Rate (Wuhan)	0.485	0.545	0.059	(0.224)
Hospitals (City)	30.530	30.975	0.446	(0.801)
GDP per Capita (City)	90.826	87.998	-2.828	(0.538)
Health Expenditure Share (City)	0.076	0.077	0.001	(0.768)
Observations	522	124	646	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Following Fitzgerald et al. (1998) we formally test for non-random attrition in the data. Specifically, we explore whether the observable individual- and city-level characteristics are associated with a greater probability of leaving the sample. To do so, we regress an attrition indicator, equal to 1 for attrited individuals and zero otherwise, on the full set of variables measured in the initial survey waves shown in Tables B2 and B3. If attrition is random, the estimated parameters will not be statistically different from zero. Results are shown in Table B4. The dependent variable in column (1) captures attrition at either Wave 2 or 3. The results indicate that our primary treatment variable (city-level COVID-19 cases) is unrelated to attrition, which shows that attrition is exogenous to treatment (i.e., exposure). However, the analysis confirms that attrition is significantly associated with certain baseline characteristics, which suggests that attrition is non-random and warrants further investigation into potential selection bias.

In the presence of non-random attrition, a second standard procedure is to assess whether attrition is ignorable. To do so, we implement the BGLW (Becketti, Gould, Lilliard, & Welch, 1988) test which assesses whether attrition is statistically associated with our main dependent variables. The BGLW test involves regressing an outcome variable from the initial wave on a set of explanatory variables, an attrition dummy (capturing future attrition), and the attrition dummy interacted with the other explanatory variables. An F-test of the joint significance of the attrition dummy and the interaction variables can help to determine whether the explanatory variables differ systematically between non-attrited and attrited households. We implement the BGLW test for outcomes measured in both Waves 1 and 2, using the attrition dummy from the previous attrition test and its interaction with individual characteristics and city-level variables as the predictors. We reject the null hypothesis of no difference between attrited and non-attrited for only two of 12 outcomes, namely the Trust Game (Wave 1) and our measure of Altruism (Wave 2). Although we find no pervasive evidence that attrition is non-ignorable, differential attrition may still pose a threat to statistical inference from our analysis.

In an attempt to adjust for differential attrition, we use the inverse probability weighting (IPW) technique, following the procedure outlined in Wooldridge (2002, 2007). The key assumption of IPW methods is that by conditioning on a set of observed covariates, the complete-population density of an outcome of interest can be derived by weighting the conditional density by the inverse selection probabilities (Fitzgerald et al., 1998). We use the full set of individual and city-level characteristics observable at baseline (Wave 1), shown in column (3) of Table B4, to predict the probability (p_i) that an individual will be observed in all three survey waves. Each individual receives a weight equal to $1/p_i$, giving more weight to participants who are similar on baseline observables to those individuals who did not stay in the sample at Waves 2 or 3. We apply the IPW to all model estimates throughout the analysis.

Table B4: Attrition Probit

	(1) Attrition any Wave
*Ln*Cases	0.040
	(0.079)
Pay-off (Wave 1)	-0.005
	(0.003)
Cooperation	-0.000
	(0.000)
Trust	-0.001
	(0.002)
Risk Aversion	-0.004
	(0.017)
Risk Taking	0.000
	(0.000)
Age	0.102***
	(0.028)
Female	-0.537***
	(0.112)
Rural Hukou	0.013
	(0.141)
Only Child	0.177
	(0.123)
General Health	-0.001
	(0.063)
Depression Score	0.012
	(0.010)
Economics Major	0.281***
	(0.099)
Chronic Illness	0.094
	(0.176)
Perseverance	0.100
	(0.077)
Prosocial Trait	-0.122
	(0.121)
Competitiveness	-0.030**
	(0.013)
Immigration Rate (Wuhan)	0.067
	(0.144)
Hospitals (City)	0.002
	(0.003)
GDP per Capita (City)	-0.002
	(0.002)
Health Expenditure Share (City)	0.933
	(1.760)
Constant	-2.361***
	(0.788)
Observations	771

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C

Experimental Protocol & Questionnaire

Appendix C is hosted online at https://drive.google.com/file/d/1-JDd4r19m91zOsISqWJFPiEfJWVaKL 51/view?usp=sharing