

Social distancing in networks: A web-based interactive experiment

Edoardo Gallo*, Darija Barak, Alastair Langtry

July 2023

Abstract

The outbreak of the COVID-19 pandemic led to an unprecedented rise in the use of social distancing as a means to curb the spread of infection. We examine the effectiveness of fines and informational messages (nudges) in promoting social distancing in a framed web-based interactive experiment conducted during the first wave of the pandemic on a near-representative sample of the US population. Fines promote distancing, while nudges have a smaller and less robust impact. Individuals do more social distancing when they are aware they are a superspreader. Using an instrumental variable approach, we show suggestive evidence that progressives are more likely to practice distancing, and they are somewhat more responsive to fines.

JEL: C99, D85, D91, I12

Keywords: social distancing, online experiment, nudge, superspreader, political ideology

Declaration of interest: none. Authors affiliation: University of Cambridge, Faculty of Economics, Sidgwick Avenue, Cambridge, UK. Edoardo Gallo and Darija Barak contributed equally to this work. Informed consent was obtained from all participants, and the research was approved by the Faculty of Economics Ethical Committee at the University of Cambridge (No: UCAM-FoE-20-01). The authors are grateful to Matthew Elliott, Freya Jephcott, Frederic Moisan, Gabriella Santangelo, Flavio Toxvaerd and two anonymous referees for helpful comments and suggestions, and to Ilia Shumailov for technical assistance. Thanks to seminar participants at the Large Scale Experiments Webinar and Cambridge-INET Networks Reading Group. Edoardo Gallo acknowledges support from the Cambridge Endowment for Research in Finance, and the Cambridge Humanities Research Grant Scheme. Darija Barak acknowledges support from the Economic and Social Research Council (ESRC) and the Royal Economics Society. Alastair Langtry acknowledges support from the ESRC. This work is the sole responsibility of the authors, and does not necessarily represent the official views of any of the funding agencies.

* Corresponding author: edo@econ.ac.uk (E. Gallo)

In 2020, the COVID-19 pandemic caused one of the most significant global disruptions since World War II (Chudik et al., 2020; Eichenbaum et al., 2020; Kissler et al., 2020; Polyakova et al., 2020), and paralyzed the global economy for over a year. Governments worldwide turned to a range of Non-Pharmaceutical Interventions (NPIs) to prevent healthcare systems from being overwhelmed and reduce the number of deaths (Brauner et al., 2021; Murray, 2020; Verelst et al., 2020). In particular, regulations to minimize social interactions by limiting exposure to others (social distancing) and/or self-isolating emerged as a primary policy tool to manage infection levels in the short and medium-term (Hale et al., 2020; IMF, 2020).

Widespread adoption of, and compliance with, social distancing guidelines was unprecedented. Throughout the pandemic, governments used a variety of policy tools to promote social distancing behavior. Some countries imposed heavy fines on anyone found breaching social distancing measures. For example, during 2020, Singapore punished first-time offenders with a fine of up to 10,000 SGD (approx. 7,000 USD), and repeated violators could be jailed (Singapore Ministry of Health, 2020). At the other end of the spectrum, India and the UK relied heavily on informational messages (nudges hereafter) early on, but then introduced fines after violations turned out to be common (Hunter, 2020; Sibony, 2020).

Interestingly, the roll-out of various COVID-19 vaccines in December 2020 did not put an immediate end to the use of NPIs. The SARS-CoV-2 virus continued mutating, leading to the emergence of new variants of the disease which were more resistant to vaccines (Andrews et al., 2022; Vogel, 2022). Moreover, some countries, such as China, continued to adhere to strict COVID-19 policies including mandatory lockdowns of entire cities up until January 2023 when the zero-COVID-19 policy officially ended. In practice, social distancing policies were widely used worldwide for nearly two years. However, even today, governments lack robust scientific evidence on the relative effectiveness of different NPIs aimed at encouraging social distancing practices.

We conduct a framed interactive online experiment to investigate the effectiveness of fines and nudges as policy interventions in promoting social distancing, and the role of network position and the contagiousness of the disease. In the fine treatment, subjects pay a fine if

they do not practice social distancing. In the nudge treatment, subjects watch a 3-minute video that explains the harm to others of not distancing. To investigate the role of network position, we assign participants to either a 5-node complete network (where everyone is connected to everyone else) or a 5-node star network (where participants are only connected to a central “superspreader”). Finally, we vary the contagiousness of the disease by assigning participants to either a high or a low contagion treatment. The subject pool, recruited through Amazon Mechanical Turk (henceforth, MTurk), is a near-representative sample of the US population in terms of age, gender, and geographical location. All sessions took place in May 2020 at the height of the first wave of the COVID-19 pandemic.

Our main finding is that fines significantly increase the level of social distancing by participants, while the impact of nudges is marginal. In particular, the estimated effect of the nudge in our experiment is about half of the impact of the fine for non-compliance with social distancing. Moreover, the effect of the nudge is less robust in statistical analysis, and overall appears more short-lived. This supports the increased use of financial penalties by many governments, rather than just relying on informational messages as in the early stages of the pandemic (Hunter, 2020; Murray, 2020).

Second, network position matters. Superspreaders practice distancing more than both poorly connected peripheral individuals and individuals with the same number of interactions in a homogeneous group. Superspreaders played a pivotal role in the COVID-19 pandemic, but, to our knowledge, research has been largely limited to the biological dimension (Adam et al., 2020; Lau et al., 2020). There is, however, also a social dimension to being a superspreader – individuals with many social interactions are more likely to spread the virus. Individuals are aware of being a superspreader in terms of social interactions, and our findings show that this leads to behavioral responses that are absent in the biological realm.

Third, political conservatives are less likely to practice distancing. The magnitude of the effect is significant and comparable to the impact of introducing a fine. We use an instrumental variable approach to provide suggestive evidence that this relationship is causal. This is consistent with the well-documented partisan divisions that characterize the COVID-19 response in

the US (Allcott et al., 2020; Gollwitzer et al., 2020). Note that, since this result is not backed by random assignment to treatments, it clearly relies on much stronger assumptions than the first two key results.

This study has immediate methodological implications for policymakers considering NPIs to contain the spread of contagious diseases such as COVID-19. It also contributes to both economics and epidemiology literatures. Below, we briefly summarize our contributions to these areas.

Our paper proposes a game that is novel in the experimental economics literature to investigate social distancing decisions. It is related to the standard public good game because it explores a trade-off between a costless self-interested decision and a costly one that benefits everyone in the group (Ledyard, 1994). The increase in social distancing in the fine intervention is analogous to the increase in contributions when there is a punishment mechanism in a public good game (Fehr and Gächter, 2000). Financial punishment for noncompliance with a policy has been studied in the behavioral and experimental economics literature in a variety of contexts. Importantly, a strand of this literature has documented that fines can backfire in some circumstances, primarily by crowding out intrinsic incentives (Gneezy et al., 2011; Gneezy and Rustichini, 2000a,b). For this reason, it is important to collect evidence on the likely effects of such fines before their deployment. Our paper provides evidence on the effectiveness of fines to ensure compliance with social distancing in the context of COVID-19.

When it comes to policy interventions and their effect on individual behavior, a growing literature investigates the effectiveness of nudges. The attractiveness of nudges, especially in the form of informational messages, is that they are a relatively inexpensive tool that can be implemented and scaled easily. Yet, the evidence on the effectiveness of nudges from existing research is mixed. For example, a meta-analysis by Hummel and Maedche (2019) finds that in empirical nudge studies only 62% of nudging treatments are statistically significant. Hence, testing the effectiveness of nudges may be useful before rolling them out in the population. When it comes to the COVID-19 pandemic, there is some evidence that informational messages *do* have an impact in certain settings. For example Banerjee et al. (2020) test an informational

message in West Bengal, India and find a significant impact on behavior – i.e. reporting of health symptoms to community health workers and travel beyond one’s village. Our paper contributes evidence of the efficacy of such messages in the US.

A further contribution of our work is to show how awareness of being a superspreader affects behavioral responses. The role of superspreaders has received a lot of attention in the epidemiological literature, but studies are limited to the biological component where the individual is unaware of being a superspreader (Adam et al., 2020; Lau et al., 2020) and/or empirical studies that correlate network position with health outcomes (Chen et al., 2021). In our work, instead, we show that position in the network *causes* changes in social distancing behavior, and we can investigate how these behavioral responses vary with policy interventions. Importantly, we find that subjects in the superspreader position practice more social distancing relative to the case of a complete network (where everyone has the same number of connections).

In political economy, recent studies have investigated the impact of political ideology on social distancing decisions. For instance, both Gollwitzer et al. (2020) and Allcott et al. (2020) use sociomobility data from smartphones to measure physical distancing. Looking at county-level data, they find that Republicans do significantly less social distancing than Democrats. In contrast, we look at decision-making at the individual level, which allows us to relate individuals’ political leanings to their behavior. Our findings are strongly complementary to the existing literature – we also find that conservatives do significantly less distancing.

Aside from the above key contributions, we also contribute to theoretical economics. A crucial new component of the social distancing game in our study is that decisions are affected by stochastic factors in the environment – the contagion process and the partial effectiveness of social distancing. This relates our paper to experimental work on public goods with stochastic elements. In particular, Fudenberg et al. (2012) look at repeated play in a prisoner’s dilemma where intended actions are implemented with noise. Further, Charness et al. (2014) investigate how uncertainty about network structure affects subjects’ ability to coordinate on efficient outcomes in games of strategic substitutes and complements. Unlike these works, in our social distancing game stochasticity originates from a probabilistic contagious process through which

the disease spreads across the network. While using a contagious process is not entirely new to experiments on self-protection decisions (see, for example, Böhm et al. (2016); Chen et al. (2013)), our study is, to the best of our knowledge, the first to incorporate an explicit network structure into the setup.¹ As discussed above, this allows us to causally investigate the effects of the architecture of social connections in the group on individual behavior and aggregate outcomes.

Field experiments have become a prominent methodology to test policies because they provide clean evidence of causality in real settings (Duflo, 2020). The COVID-19 pandemic has largely deprived policymakers of the ability to deploy field experiments due to movement restrictions including self-isolation, social distancing, and travel bans. In the last decade, several studies have run interactive web-based experiments in other disciplines (Rand et al., 2011; Shirado and Christakis, 2017) and, in a more limited fashion, within economics (Gallo and Yan, 2022; Jackson and Xing, 2014). Importantly, to the best of our knowledge, our work is the first to exploit the diversity of online subject pools to achieve representativeness along some sociodemographic dimensions (i.e. gender and geographic place of residence). Recent studies have shown that representative samples can differ from student populations with respect to fundamental behaviors (see, for example, Chapman et al. (2022) on loss-aversion). Using a near-representative sample allows us to get a more informed perspective on the behaviors in the US population. Another contribution of our work is to show how framed interactive web-based experiments can be used to test policies, complement existing methods, and provide distinct advantages. The final section of this paper contains further discussion on how framed interactive web-based experiments can complement and enhance studies based on sociomobility data and surveys.

The epidemiology literature about social distancing focuses on its macro-level impacts, i.e. how it might affect the evolution of the pandemic. Consequently, studies are mostly simulation-based. For example, Fenichel et al. (2011) explicitly incorporate optimizing behavior by a representative agent into the classic SIR model and use simulations to examine how this affects

¹See Choi et al. (2016) for a comprehensive review of the experimental literature on network games.

disease dynamics. More recently, there has been a profusion of studies focused on the COVID-19 pandemic. Kissler et al. (2020) and Chang et al. (2020) both simulate disease dynamics for COVID-19, and explore the effects of a variety of NPIs. In contrast, we elicit actual behavioral responses to counterfactual policies. Aside from their direct relevance, our findings can also be used to refine simulation-based epidemiological models by informing the choice of behavioral rules.

The remainder of this paper is structured as follows. Section 1 describes our experimental design. Section 2 summarizes our methods of data collection and presents the resulting sample. Section 3 presents the main results from our experiment, while Section 4 discusses covariates of social distancing decisions. Section 5 interprets the results from the previous two sections. Finally, Section 6 contains a discussion of our results and the role of framed interactive online experiments in informing policy.

1 Experimental Design

This section describes the game, the workflow of the experiment, and treatments.

Game. Figure 1 illustrates the social distancing game which corresponds to a round of the experiment. At the beginning of a round, participants are randomly allocated to one of the positions in an unweighted and undirected network (top left). The structure of the network is common knowledge. Subjects must simultaneously decide whether to practice social distancing at a known cost $c > 0$ (top center). After the decisions are made, one and only one participant – *patient zero* – is randomly picked to contract COVID-19 directly. In the example in Figure 1, patient zero is the participant highlighted in red in the top right panel.

There are two benefits to practicing social distancing. First, if the individual is randomly picked to be patient zero, then she becomes infected with 50% probability, rather than for sure. Second, any individual who practices distancing cannot transmit COVID-19 to others or contract it from infected individuals. In the example in Figure 1, patient zero decided to not practice social distancing, hence she becomes infected.

COVID-19 then spreads through the network from infected to healthy individuals who do not practice social distancing through contagion with a known probability $\alpha \in (0, 1)$. The bottom panels of Figure 1 show a possible spread in the example with three individuals infected and two healthy ones at the end of the contagion process.² At the end of the round, healthy individuals receive a benefit of $b = 100$ points, whereas those infected receive zero benefit. Any individual who chose to practice social distancing pays the cost $c = 35$ points, irrespective of their final infection status.

Assuming self-interested fully rational agents, the model predicts that the pure strategy Nash equilibrium outcomes will be inefficient for a wide range of parameter values. In fact, given an arbitrary network, it is possible to fully characterize both equilibrium and efficient outcomes.³ In particular, we show that, depending on the network structure, the equilibrium number of agents practicing social distancing can be below the efficient number. In this model, social distancing is essentially a merit good so it exhibits positive externalities – welfare loss may therefore occur in the absence of intervention.

Workflow. In the experiment participants first play 20 rounds of the baseline game above in fixed groups of five subjects, with positions on the network being randomly reallocated in each round. Notice that the participants are not informed about the decisions and outcomes of other members of their group at any point during or after the experiment. After the initial 20 rounds, their group is treated with either the fine or the behavioral nudge policy intervention. In the fine treatment, whenever a participant decides not to practice social distancing she receives a fine of $f = 15$ points, independent of her infection status at the end of the round. In the nudge treatment, every participant must watch a 3-minute video explaining how failing to practice social distancing may harm others. The same group of 5 participants then play 20 rounds of the game under one of the policy interventions. Hence we investigate the impact of each policy on social distancing behavior using a within-subjects design and the relative effectiveness of the

²The algorithm that we use is iterative and relies on the idea that each of the infected agents interacts with other agents exactly once. As a result, in a star network there may be at most two stages of contagion, whereas a complete network with five nodes permits at most four stages (for further details see Online Appendix A.1).

³See Online Appendix for the full characterization.

two policies with a between-subjects design. Once participants have completed the interactive part of the experiment, they proceed to the post-experimental questionnaire, with basic demographics questions and a set of knowledge and attitudes questions on a range of topics including the COVID-19 pandemic, social distancing, religion, global warming, ideology, and political affiliation. Finally, participants complete an incentivized Bomb Risk Elicitation Task (BRET) (Crosetto and Filippin, 2013) to elicit risk preferences. Subjects are explicitly primed to think about COVID-19 in the instructions by naming the disease and describing the main symptoms according to the guidelines by the Centers for Disease Control and Prevention.

Treatments. The focus of our experiment is on three treatment dimensions. First, we consider two 5-node network architectures for the structure of interactions among participants: (1) a complete network, where everyone is connected to everyone else, and (2) a star network, where one node is connected to all other nodes which are not connected between themselves. Next, we vary the level of contagiousness of COVID-19 to be either low ($\alpha = 15\%$) or high ($\alpha = 65\%$). Finally, we focus on two intervention methods: (1) a fine $f = 15$ points for not practicing social distancing, or (2) a behavioral nudge in the form of an informational video highlighting the harm caused to others by not practicing social distancing. Consequently, we have a $2 \times 2 \times 2$ full-factorial design, with a total of 8 treatments.

We choose to study the effect of an informational video for two main reasons. First, as a (relatively) cheap, easy to deploy, and non-invasive intervention, it is a natural benchmark against which to compare other policies. To justify more expensive and onerous measures – such as fines – policymakers ought to want those measures to be more effective than a nudge, not just more effective than doing nothing. Second, governments across the globe used informational messages extensively during the height of the COVID-19 pandemic.⁴ Understanding whether these interventions are likely to work is therefore important in its own right. Additionally, while a pure informational message has no impact in the standard rational agent framework (which our model assumes for simplicity), a range of studies find that they can affect behavior (Banerjee

⁴For example, the UK Government spent over £240 million on COVID-19 advertising in 2020 alone (More About Advertising, 2021).

et al., 2020; Barak et al., 2022).

The complete and star networks we use as treatments are stark, but they are a useful starting point for investigating the role of the network. They present clearly differentiated structures, making them useful for identifying causal effects.⁵ Additionally, they allow tractable analytic solutions to our model and hence sharper theoretical predictions (see Online Appendix A.1).

2 Data

In this section, we summarize our data collection methods and present the resulting dataset. The section also includes a description of the key demographics of the final sample and a convergence analysis of the decision data from the experiment.

Following standard practices in interactive online experiments (Gallo and Yan, 2022; Suri and Watts, 2011), we first recruit a standing panel of subjects using a short survey on MTurk. During recruitment, we use the 2018 US Census data to generate a representative panel of the adult US population in terms of age, gender, and geographical location (U.S. Census Bureau, 2019). As part of recruitment, we also collect subjects’ social preferences using the Social Value Orientation (SVO) task (Murphy et al., 2011). The incentivised SVO task classifies individuals into four categories – prosocial, individualistic, competitive, and altruistic.

For each of the experimental sessions, we invite a random representative sample of subjects from our standing panel.⁶ We obtain a near-representative sample of $n = 400$ participants completing the experiment. To check that our sample is balanced when it comes to treatment assignment, we run a chi-squared test on participants’ assignment to treatment (a categorical variable with 8 options) with respect to gender, age category, and geographical location variables. Results indicate that subjects in all eight experimental treatments are not significantly

⁵They are also common choices for networks experiments (see, for example, Choi and Lee (2014); Frey et al. (2012); Rosenkranz and Weitzel (2012)). However, if responsibility for protecting others is more diffused in less stark networks, it is possible that network structure plays a lesser role there. Falk et al. (2020) presents evidence suggesting that people are more willing to take actions seen as immoral if their role in decisions is not pivotal.

⁶We restrict participation to US-resident participants who have completed at least 500 tasks and have an overall ranking of at least 96%. We also ensure that no two subjects with the same IP address participate in the experiment. Finally, in line with recent contributions (Gallo and Yan, 2015; Rand et al., 2011), multiple participation is not allowed even by participants who have only seen the instructions.

different from each other (two-sided χ^2 , $p > 0.1$, see Table A4 for details).

The experiment took place between May 4th and 29th, 2020.⁷ Each session had between one and five independent groups of five subjects playing simultaneously, with the exact numbers depending on turnout. The experiment lasted on average 33 minutes, and subjects earned an average of 5.81 USD (including a fixed participation fee of 1 USD). Subjects remained completely anonymous throughout the experiment, and informed consent was obtained from subjects before participation.

Prior to the experiment, we screened Mturkers in two ways. First, we only recruited Mturkers with 500 completed Human Intelligence Tasks (HITs) on MTurk *and* an approval rate of at least 96% into our standing panel. Second, we tracked IP addresses, and did not permit the same IP address to take part more than once.⁸ Within the experiment, we screened participants in two main ways. First, we used a 3-question quiz to check their understanding of the instructions. Participants who failed three times were excluded before the experiment started. Second, we gave subjects 20 seconds to make each social distancing decision, and we disqualified them if they missed three consecutive decisions. Throughout the experiment, 6 subjects were disqualified for this reason. We dropped all data from their respective groups from our analysis – but adding it back in has no material impact on any of our results. A.3 of the Online Appendix provides further detail on these, and other steps we took to ensure data quality.

All 400 subjects in our sample are residents of the US, 47% are female, and the mean age is just under 44 years. The average subject in our sample has 15.3 years of education and 70% are employed (either full or part-time). In terms of religious attitudes, 51% of our participants consider themselves religious, with 39% reporting Christianity as their primary religion. The average subject in the sample is moderately risk-averse, with a CRRA score of 0.53 – 0.55. With respect to social values, our sample mostly consists of prosocials (58%) and individualists

⁷We used Qualtrics (www.qualtrics.com) for recruitment, and the main experiment was programmed in o-Tree (www.otree.org) (Chen et al., 2016) with a server deployed on Heroku (www.heroku.com).

⁸As well as preventing a given person participating more than once, this also seeks to stop more than one person from the same household from participating. For these purposes, even viewing the experimental instructions counts as taking part.

(41.8%), and only one subject is classified as competitive.⁹

Overall, the final dataset contains 40 decisions for each of the 400 subjects for a total of 16,000 observations. As subjects played simultaneously in groups of five, there are 80 distinct clusters (i.e. groups). Figure 2A shows the evolution of average individual propensity to social distance aggregated over contagiousness and network treatments. The figure shows a downward trend in social distancing behavior in the first rounds. Such behavior is typical for public goods experiments whereby contributions to the public good decline as the session progresses (Ledyard, 1994). In this paper, we are interested in the limiting outcomes of this convergence process, and therefore we would like to restrict our attention to those rounds where the majority of subjects have settled on some stable strategy.

In particular, we define that a participant has converged to a stable strategy by a certain round if she used this strategy for the previous four rounds, and in all subsequent rounds the number of consecutive deviations from the chosen strategy does not exceed two.¹⁰ Given this convergence definition, we find that at least 80% of participants converge to a stable strategy in every treatment after the first 10 rounds. Our main analysis, therefore, focuses on the last 10 rounds of the baseline and intervention in order to control for the time trend due to learning and experimentation by subjects. The results are robust to including all of the data.

3 Determinants of social distancing decisions

This section investigates the determinants of social distancing decisions, using both non-parametric and parametric methods. The analysis in this section is based entirely on our treatment variables – policy intervention (fine or nudge), network architecture and the rate of contagiousness of the disease.

The non-parametric analysis uses the Mann-Whitney U-test (MW hereafter) (Mann and

⁹See Online Appendix for a more detailed description of the demographics of the sample.

¹⁰We consider three types of convergence strategies. In both networks, we look at the strategy where the subject always chooses the same action. For the star network, we also consider two extra strategies. In one strategy the participant always chooses the same action when she is in the superspreader position and the complement action when she is peripheral. In the other strategy, she always chooses the same action when she is the superspreader and alternates between the actions when peripheral.

Whitney, 1947) for unmatched samples and the Wilcoxon Signed-Rank test (WSR hereafter) (Wilcoxon, 1945) for matched samples. Samples are aggregated at the group level and matched when they are generated by the same groups, and they are unmatched when making comparisons across groups from different treatments.¹¹

Our parametric approach relies on the Random Effects Logit (henceforth, REL) to model the binary choice of whether to practice social distancing. Formally, it models a random utility model with individual-specific random effects:

$$y_{it} = 1 \iff x_{it}\beta + v_i + \epsilon_{it} > 0 \quad (1)$$

where y_{it} is the decision to practice social distancing, x_{it} is the set of controls, v_i are individual specific random-effect and ϵ_{it} is an error term with a logistic distribution. It allows individual-specific propensity to practice social distancing, but assumes that the impact on utility of a variable x is not individual-specific. We cluster errors at the group-level. Our main specification (column 5 in Table 1) includes the following five categories of independent variables x_{it} : (1) dummy variables for the experimental treatments (fine intervention, nudge intervention, contagion level, and node types), (2) demographics controls, (3) location and institutional controls, (4) preference controls, and (5) ideology controls. P-values throughout this section are presented for a single specification, but, as Table 1 shows, results are robust across all specifications.¹²

A preliminary question is whether the policy intervention has an impact on the aggregate level of social distancing. We conduct a Wald test to identify a structural break. This test assumes that there is at most one structural break in the data, but is agnostic as to when it occurs and whether it occurs at all. In the fine intervention, a change in the evolution of social distancing occurs at round 21, immediately after the implementation of the fine (Wald test, $p < 0.001$ for all data and subdivided by the other treatments). The same break at round 21

¹¹All results using these non-parametric tests are robust to using their parametric analogs, i.e. t-test for unpaired and paired samples (see appendix A.4 of the Online Appendix).

¹²Further, the results are robust to using a Linear Probability Model or a Probit model. We discuss this further in the Online Appendix.

occurs in the nudge intervention in the aggregate data (Wald, $p < 0.001$), but the result is not robust once we subdivide by the other treatments. See Table A5 in the Online Appendix for a full breakdown. The introduction of the fine therefore has an immediate impact on social distancing behavior, while the impact of the nudge is less clear-cut.

Further, we find that fines increase the level of social distancing, and the effect is significant both in the non-parametric (WSR, $p = 0.0008$) and in the regression (Specification 1, S1 hereafter, $p = 0.001$) analyses. The increase in social distancing behavior is present both in low (WSR, $p = 0.04$) and high (WSR, $p = 0.003$) contagion environments. In the last 10 rounds of the fine intervention, there is an 8% increase in social distancing compared to the last 10 rounds of the pooled baseline so the effect of the fine is sizable and long-lasting.

In contrast, the impact of the nudge intervention is not robust to all specifications (WSR, $p = 0.008$; S1, $p = 0.03$). Interestingly, narrowing the focus to a specific contagion environment, the nudge significantly increases distancing with high contagion (WSR, $p = 0.0003$) but it has no effect with low contagion (WSR, $p = 0.5$). Comparing the two policies, the impact of the fine is marginally higher than the nudge (MW, $p = 0.08$). Overall, the increase in social distancing in the last 10 rounds due to the nudge intervention is only 2% against the pooled baseline.

A crucial driver in the spread of COVID-19 is the presence of superspreaders – individuals who go on to infect a much larger number of others than the average (Adam et al., 2020; Lau et al., 2020). A primary determinant of being a superspreader is biological – something that cannot be varied experimentally and is outside of the scope of this study. However, there is also a social element driven by the wide variations in the number of social interactions across individuals (Jackson and Rogers, 2007). An important difference between biological and social drivers of being a superspreader is that an individual is aware of the latter, and therefore the tendency to social distance may vary with position in the network.

Figure 2B illustrates the three positions in our experiment drawn with the node size proportional to the amount of social distancing in that position for each policy treatment. On the left side, there is the complete network in which all participants are in what we dub a ‘close-knit’ position. On the right side, there is the star network with one participant – the superspreader –

interacting with all others in what we dub a ‘peripheral’ position. Superspreaders practice more social distancing than peripheral participants both in the baseline and policy interventions (MW and S1, $p < 0.0001$ for all). Moreover, superspreaders also practice social distancing more than close-knit participants (MW, $p = 0.004$; S1, $p = 0.03$) despite having the same number of interactions. One potential driver of this behavior is that they are aware of their central role and want to protect the group. Another potential driver is that they realize peripheral participants are less likely to distance due to their isolated position. Our experimental design does not allow us to investigate the relative importance of these two drivers, but this is an interesting avenue for future research. High contagiousness leads to significantly more social distancing (MW, $p = 0.01$, S1, $p < 0.0001$), and the size of the effect is similar to the difference between being in the close-knit and peripheral positions. Social distancing is 10% and 15% higher in the high contagion setting than in the low contagion setting in baseline and intervention respectively.

4 Covariates of social distancing decisions

This section describes the association between social distancing decisions and a number of individual-specific variables including political ideology, risk attitudes, social preferences, and a range of socio-demographic characteristics. Section 4.1 focuses solely on political ideology where we use an Instrumental Variables (IV) approach to argue that the relationship may be causal. It also discusses the strong identifying assumptions we make in the analysis. Section 4.2 describes other correlates of social distancing decisions, where we make no claim of causal effects.

4.1 Political ideology

We find that self-reported Democrats are significantly more likely to practice distancing than Republicans (S4, $p = 0.0002$). However, this binary classification is a very coarse measure. To obtain a more fine-grained picture of participants’ political leanings, we construct an ideology index, with a score for each subject based on responses to questions about (1) support for

President Trump’s handling of the COVID-19 pandemic, (2) support for universal healthcare, and (3) belief that social distancing measures impose unjustified economic costs. Questions are on a 5-point Likert scale, so we obtain a 0-12 index with 0 and 12 indicating extremely progressive and extremely conservative participants respectively.

Figure 3A shows that the probability of practicing distancing decreases the more conservative a participant is – an increase in ideology index from 1 to 5 (25th to 75th percentile among our subjects) corresponds to a 14.6% decrease in the probability of distancing. The ideology index is a highly significant correlate of social distancing decisions (S5, $p < 0.0001$), in line with recent sociomobility and survey-based studies in the US (Allcott et al., 2020; Barrios and Hochberg, 2020; Simonov et al., 2020).

Our experimental design exogenously varies policy intervention, contagiousness, and the network, allowing us to make causal inferences on these dimensions. The same is obviously not the case with ideology. We investigate the causal effect of ideology with an instrumental variable approach, using a measure of participants’ skepticism of global warming as the instrument. As a partisan issue in the US, attitudes toward global warming are strongly correlated with political ideology (Hornsey et al., 2016; Pew Research Center, 2019). The exclusion restriction our IV identification relies on is that global warming attitudes do not affect social distancing decisions separately from their association with ideology.

However, it is likely that global warming attitudes at least partly reflect geography and local institutions – both of which might affect social distancing decisions. Failure to control for these would lead to violation of the exclusion restriction. We therefore add a set of proxy variables: federal regions indicator (see Online Appendix A.2 for details), population density, total number of COVID-19 cases to date, daily number of COVID-19 deaths, and an indicator of whether a stay-at-home order was in place.¹³ While these are necessarily imperfect proxies, it suggests that our instrument is not just capturing geography and local institutions.¹⁴

¹³Population density is at the county level. COVID-19 variables are at the state level. Cumulative number of COVID-19 cases, new daily COVID-19 deaths, and stay-at-home order indicator are for the date immediately before the day subject participated in the experiment. We do not report coefficients for them in Table 1 (we do so in Table A7 in the Online Appendix).

¹⁴Following a comment from an anonymous referee, we further checked that ideological composition of the

Further, it may be possible that our skepticism of global warming index is picking up subjects' general attitudes towards science and trust in information from experts. In turn, lack of trust in science may have an effect on subjects' social distancing behavior. While such concerns cannot be completely eliminated, existing research suggests that whereas political conservatism is strongly associated with skepticism towards global warming, understanding of science (or lack thereof) is not a climate change acceptance predictor (Rutjens et al., 2018).

Specification 6 (S6) in Table 1 is identical to our main specification (S5), except that it uses an instrument for ideology. The instrument is an index of climate change attitudes – constructed from subjects' beliefs that global warming is (1) happening, (2) caused mostly by humans, and (3) affecting weather in the USA.¹⁵ Political ideology appears to be a significant causal determinant of social distancing, with conservatives less likely to practice distancing (S6, $p = 0.04$). Figure 3B illustrates an interesting further interaction between political ideology and the type of policy intervention. While the nudge intervention has no differential impact on subjects with different ideologies, conservative participants are marginally less responsive to the fine (S6, $p = 0.08$). Note that using the ideology index itself (as in specification 5) instead of the IV approach (as in specification 6) has very little effect on the results.

4.2 Other covariates

Social preferences play an important role in the decision to social distance. First, an analysis of equilibrium decisions with self-interested individuals predicts levels of social distancing that are significantly lower than participants' decisions in all treatments (MW, $p < 0.0001$). Second, using the SVO task, we classify subjects into individualistic and prosocial. Prosocial participants are about 20% more likely to choose distancing (S2, $p < 0.0001$).

Next, we elicit participants' risk preferences using BRET. As expected, risk-seeking individuals are less likely to practice social distancing (S2, $p = 0.001$), and the effect is about half

state of residence – as captured by the political party affiliation of the governor as of May 2020 – and the level of religiosity – measured by the share of people who report at least weekly worship attendance and those who pray daily (Pew Research, 2016) – play no effect in determining distancing behavior. All results reported in this paper are robust to these checks.

¹⁵We used the questions from Howe et al. (2015).

the size of being prosocial in the SVO task. The post-experimental survey includes a question asking why subjects choose to stay at home – participants who list “protecting others” as one of the stated reasons are more likely to social distance (S2, $p = 0.007$), and the effect is similar in size to being risk-averse.

In terms of sociodemographics, females (S2, $p = 0.004$) and older individuals (S2, $p < 0.0001$) are more likely to distance. From the data, being female translates to a 12.6% increase in the probability of distancing, and there are no significant interactions between gender and the effect of each policy intervention. Social distancing is increasing with age – an average 60-year-old is 40.1% more likely to distance than an otherwise identical 20-year-old. Interestingly, older subjects are relatively more responsive to the nudge (S3, $p = 0.02$) and relatively less responsive to the fine (S3, $p < 0.0001$). Whites are less likely to distance (S2, $p = 0.04$), but the effect is not robust. Education, religion, and employment status are not associated with distancing behavior. These associations between sociodemographic characteristics and distancing are broadly consistent with evidence from sociomobility (Baradaran Motie and Biolsi, 2020) and survey-based data (Galasso et al., 2020; Masters et al., 2020; Papageorge et al., 2020; Pedersen and Favero, 2020).

5 Interpreting the results

The nonlinearity of the Random Effects Logit model makes interpreting the point estimates in Table 1 difficult. The partial effect of any given variable depends on both the initial value for that variable and on the values of all variables. For a given individual, the partial effect of changing a variable from x to x' is not the same as that of changing it from x' to x'' . As a consequence, the estimated partial effect is different for every person.

To meaningfully compare the effect of different factors, we compute the Average Partial Effects (APE) calibrated on the characteristics and/or decisions of the subjects in our experiment. To do this, we calculate the partial effect for each participant in the experiment, and then take the simple mean over these individual-specific partial effects. For the individual-specific partial

effect, we compare the estimated probability of distancing for each subject for two different values of the variable of interest while keeping all other variables at their observed value for that individual. For binary variables, we compare 0 and 1. For continuous variables, we compare the 25th and 75th percentiles (based on the experimental data).

Figure 4 shows the mean partial effects (APEs) for all variables that are statistically significant in our main specification, as well as their 25th, 50th, and 75th percentiles. Blue bars indicate the impact of determinants of social distancing decisions (i.e. treatment variables). A high contagion setting and being in the peripheral network position have the largest impact with 18% and -22% respectively. The fine intervention and being in the superspreader position have a moderate impact of about 8%. The smallest impact is the nudge intervention with only a 4% effect.

The rest of the bars in Figure 4 refer to the covariates of social distancing decisions. Green bars indicate sociodemographics correlates. Increasing age from the 25th to the 75th percentile in our data has a 23% effect on distancing. Gender and race have a moderate effect of around 10%. Orange bars display the effect of correlates related to preferences. Being prosocial is associated with a large 20% effect on distancing, while risk preferences and the desire to protect others have a moderate impact of about 10%. Finally, the red bar shows that political leanings have a moderate 12% effect, which we argue is causal in our analysis.

Omitted from Figure 4 are a set of demographic, institutional, and geographic factors that are not statistically significant. A subject's level of education, religion, and labor force participation are all unrelated to their social distancing decisions. Also insignificant are their geographic location (as measured by Federal Region), the number of COVID-19 cases and deaths in their state, and population density. Finally, Figure 4 clearly shows how the partial effect of any given variable is highly heterogeneous across subjects. In fact, the range in partial effects from the 25th to 75th percentiles *exceeds* the mean in many cases.

6 Concluding remarks

The sudden disruption brought about by the COVID-19 pandemic demanded a quick response from policymakers who had to implement novel social distancing policies whose success determined the fate of tens of thousands of lives. Our work provides experimental evidence on the efficacy of these policies and how it varies across social and individual characteristics. Additionally, it shows how the methodology of framed interactive web-based experiments can play an important role to inform policymakers, and complement data obtained using sociomobility and survey-based studies.

Our main finding is that fines are effective at increasing social distancing, while the estimated impact of nudges is about half in magnitude and not robust to all specifications.¹⁶ A limitation of our design is that we examine one specific form of nudge – an informational video that emphasizes the harms to others of failing to practice social distancing. Moreover, the ineffectiveness of nudges in our experiment seems to be at odds with some other evidence in the literature. For instance, Banerjee et al. (2020) show that a similar nudge – in the form of 2.5-minute information videos – was effective in encouraging social distancing in the Indian context. However, the efficacy of nudges is clearly sensitive to a multitude of factors, and whatever works in one context may not perform well in another. For example, as we show in Section 4, demographic characteristics, such as age, could contribute to the effectiveness of nudges. Further, in a related study conducted on a Chinese sample, we actually find a significant and robust effect of the same informational nudge at increasing compliance with social distancing (Barak et al., 2022). This suggests that the effectiveness of nudges may crucially depend on national and/or cultural contexts.

There is also a growing literature on the effectiveness of nudges in a range of other contexts. For example, de Ridder et al. (2022) show the importance of personal preferences in moderating the effects of nudges. Moreover, they show that the factors driving the effectiveness of nudges

¹⁶This result shows an important benefit of fines. Governments should also take into account the costs of using fines when setting policy. These costs could include non-economic factors, such as *per se* objections to intrusion on individuals' lives.

are rarely identified even in studies where nudges are effective. Indeed, Szaszi et al. (2022) and Maier et al. (2022) show that, overall, nudge effects tend to be barely significant after properly adjusting for publication bias. Taken together, this suggests that more research should be directed at systematizing our understanding of nudge-type interventions in the context of promoting health enhancing as well as other desirable behaviors.

Another important finding is that subjects practice more social distancing when they are in a superspreader position and are aware of their central position. In our setup, individuals in the superspreader position are more likely to social distance compared to participants in the close-knit network, despite both nodes having the same number of links. This suggests that awareness of being central in a social network may play an important role in shaping behavior, and maps a potential direction for future research in the contagion literature.

Finally, we find that conservative-leaning individuals are less likely to practice social distancing in the US. While this tendency has been documented in other empirical studies (Allcott et al., 2020; Gollwitzer et al., 2020), we provide suggestive evidence of a causal relationship using an IV approach. An open question is whether this relationship is peculiar to the US context or it extends to other countries with a less stark partisan divide.

In terms of methodology, this study highlights the important role framed interactive web-based experiments can play in providing timely information to policymakers to test the effectiveness of policies. Experiments are the gold standard to test the causal impact of policy interventions, but lab and/or field experiments were severely constrained during the COVID-19 pandemic, depriving researchers of the standard tools to identify causal effects (Haushofer and Metcalf, 2020). Web-based experiments, however, were unaffected. Moreover, they can be deployed quickly, scaled up at minimal cost, and easily reach a diverse sample and/or specific populations of particular interest.

Evidence regarding behavioral responses to social distancing policies has so far been mostly based on either sociomobility data or survey studies based on self-reported claims about hypothetical behavior. Aside from a clean identification of causality, web-based interactive experiments complement these methodologies and present some further distinct advantages.

Mobility data provides detailed information on behavior in real-life settings that can verify the external validity of web-based experiments. For instance, recent studies validate our finding that conservative-leaning individuals practice less social distancing (Allcott et al., 2020; Barrios and Hochberg, 2020). However, it is usually stripped of any personal information, including sociodemographics, so that inferences about the effects of these variables can only be made for geographic units rather than individuals. Moreover, sociomobility data is unable to shed light on the effectiveness of counterfactual policies, while policymakers frequently need information in advance of implementation.

Surveys are a standard methodology for fast, large-scale data collection, but our study raises some question marks about their accuracy in gauging actual behavioral responses. In our post-experimental survey, we ask participants to estimate how effective fines and nudges are in promoting social distancing. Participants significantly overestimate the efficacy of nudges compared to their actual choices in the experiment independently on whether they are in the fine (MW, $p < 0.0001$) or nudge (WSR, $p = 0.001$) treatment.¹⁷ This difference between survey-based self-reported intentions and actual behavior is consistent with recent evidence on the short-lived efficacy of nudges (Brandon et al., 2017). In other words, picking a policy based on self-reported measures may lead policymakers to opt for the wrong solution. These results caution against excessive reliance on survey-based self-reported data to gauge behavioral responses.

¹⁷For further details of this analysis see Section A.4.5 of the Online Appendix.




Figure 1: Flow of a typical round of baseline and intervention. In the experiment, we use the following parameterization: $f = 0$ or 15 points (fine intervention only); $c = 35$ and $b = 100$ points; $\alpha = 0.15$ (low contagion) or 0.65 (high contagion). Final payoffs for the round are a combination of individual social distancing choice and infection status. For example, a participant who practices social distancing and is healthy, receives $(-c + b) = 65$ points. In the figure, the chosen network architecture is the star. In the experiment, half of the treatments had the star network architecture and the other half had the complete.




Figure 2: Policy and network position. (a) Evolution of aggregate probability of distancing split by policy intervention. (b) Probability of distancing split by network position and policy intervention. Node size is proportional to distancing probability. Width of dotted band around the circumference of each circle indicates 95% confidence intervals.




Figure 3: Ideology and social preferences. (a) Probability of distancing and (b) Percentage change in probability of distancing split by policy intervention and quintile of ideology index. Vertical bars indicate 95% confidence intervals.




Figure 4: Average Partial Effects (APE) for variables that are significant in the REL main specification. Bar boundaries indicate 25th and 75th percentiles, vertical line inside bars indicate 50th percentile, and black dots show the mean. Blue bars indicate treatment variables. Green and orange bars indicate effects from demographics and social preferences respectively. Red bar is the ideology variable. Variables not included because they are not significant include: education, religion, labor force participation, case numbers, and geography controls.

Table 1: Logit regression analysis of the experimental data.

Variables	Decision to practice distancing (binary)					
	(S1)	(S2)	(S3)	(S4)	(S5) Main	(S6) IV
Treatment	Fine	0.755*** (0.231)	0.761*** (0.242)	3.023*** (0.515)	0.942*** (0.287)	1.146*** (0.352)
	Nudge	0.347** (0.157)	0.343** (0.155)	-0.928 (0.582)	0.353 (0.202)	0.438* (0.261)
	High contagion	1.330*** (0.336)	1.537*** (0.323)	1.561*** (0.326)	1.238*** (0.387)	1.547*** (0.323)
	Superspreader	0.975** (0.438)	0.830** (0.421)	0.845** (0.425)	0.870* (0.484)	0.833** (0.424)
	Peripheral	-1.608*** (0.322)	-1.780*** (0.333)	-1.819*** (0.338)	-1.760*** (0.424)	-1.785*** (0.333)
Preferences	Bomb risk score		-0.0296*** (0.00889)	-0.0297*** (0.00894)	-0.0226** (0.0107)	-0.0296*** (0.00896)
	Prosocial values		1.485*** (0.334)	1.478*** (0.337)	1.570*** (0.365)	1.495*** (0.335)
	Protect others		0.918*** (0.342)	0.984*** (0.342)	0.975** (0.381)	0.926*** (0.346)
Demographic	Age (years)		0.0793*** (0.0137)	0.0867*** (0.0142)	0.0873*** (0.0164)	0.0796*** (0.0139)
	Female = 1		1.069*** (0.372)	1.070*** (0.370)	1.122*** (0.428)	1.072*** (0.372)
	Race = White		-0.960** (0.469)	-1.079** (0.468)	-0.941 (0.680)	-0.955** (0.468)
Ideology	Conservative index		-0.282*** (0.0565)	-0.278*** (0.0567)		-0.250*** (0.0591)
	Republican dummy				-1.529*** (0.403)	-0.207** (0.0995)
Interactions	Fine-ideology interaction				-0.110* (0.0618)	-0.110* (0.0618)
	Nudge-ideology interaction				-0.0266 (0.0447)	-0.0265 (0.0448)
	Fine-age interaction			-0.0505*** (0.0112)		
	Nudge-age interaction			0.0309** (0.0130)		
Other Controls ^a		No	Yes	Yes	Yes	Yes
Constant		1.456*** (0.312)	-2.051 (1.670)	-2.537 (1.721)	-3.062 (1.928)	-2.213 (1.666)
Observations		8,000	7,880	7,880	5,520 ^b	7,880
Clusters		80	80	80	80	80

Robust standard errors in parentheses. Note that 6 subjects did not complete the SVO task.

^a Dummies for: religion, out of labor force, unemployed, federal regions, stay at home order in place. Continuous variables for: population density, cumulative COVID-19 cases in state, daily COVID-19 deaths in state, years of schooling, and number of quiz failures. Specification 5 also includes interactions between the Conservative Index and fine and nudge treatment variables.

^b 276 subjects used in this regression. 124 subjects are excluded because they self-identified as “Independent” or “Other” with respect to political party support.

References

- Adam, D., P. Wu, J. Wong, E. Lau, T. Tsang, S. Cauchemez, G. Leung, and B. Cowling (2020). Clustering and superspreading potential of severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) infections in Hong Kong. *Nature Medicine* 26, 1714–1719.
- Allcott, H., L. Boxell, J. Conway, M. Gentzkow, M. Thaler, and D. Y. Yang (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. Working Paper 26946, National Bureau of Economic Research.
- Andrews, N., J. Stowe, F. Kirsebom, S. Toffa, T. Rikeard, E. Gallagher, C. Gower, M. Kall, N. Groves, A.-M. O’Connell, et al. (2022). COVID-19 vaccine effectiveness against the Omicron (B.1.1.529) variant. *New England Journal of Medicine* 386(16), 1532–1546.
- Banerjee, A., M. Alsan, E. Breza, A. G. Chandrasekhar, A. Chowdhury, E. Duflo, P. Goldsmith-Pinkham, and B. A. Olken (2020). Messages on COVID-19 prevention in India increased symptoms reporting and adherence to preventive behaviors among 25 million recipients with similar effects on non-recipient members of their communities. Working Paper 27496, National Bureau of Economic Research.
- Baradaran Motie, G. and C. Biolsi (2020). County-level determinants of social distancing (or lack thereof) during the COVID-19 pandemic. *Contemporary Economic Policy* 39, 264–279.
- Barak, D., E. Gallo, K. Rong, K. Tang, and W. Du (2022). Experience of the COVID-19 pandemic in Wuhan leads to a lasting increase in social distancing. *Scientific Reports* 12(1), 18457.
- Barrios, J. M. and Y. Hochberg (2020). Risk perception through the lens of politics in the time of the COVID-19 pandemic. Working Paper 27008, National Bureau of Economic Research.
- Böhm, R., C. Betsch, and L. Korn (2016). Selfish-rational non-vaccination: Experimental evidence from an interactive vaccination game. *Journal of Economic Behavior & Organization* 131, 183–195.
- Brandon, A., P. J. Ferraro, J. A. List, R. D. Metcalfe, M. K. Price, and F. Rundhammer (2017). Do the effects of social nudges persist? Theory and evidence from 38 natural field experiments. Working Paper 23277, National Bureau of Economic Research.
- Brauner, J. M., S. Mindermann, M. Sharma, D. Johnston, J. Salvatier, T. Gavenčiak, A. B. Stephenson, G. Leech, G. Altman, V. Mikulik, A. J. Norman, J. T. Monrad, T. Besiroglu, H. Ge, M. A. Hartwick, Y. W. Teh, L. Chindelevitch, Y. Gal, and J. Kulveit (2021). Inferring the effectiveness of government interventions against COVID-19. *Science* 371, 1–8.
- Chang, S. L., N. Harding, C. Zachreson, O. M. Cliff, and M. Prokopenko (2020). Modelling transmission and control of the COVID-19 pandemic in Australia. *Nature Communications* 11(1), 1–13.
- Chapman, J., E. Snowberg, S. W. Wang, and C. Camerer (2022). Looming large or seeming small? Attitudes towards losses in a representative sample. Working Paper 30243, National Bureau of Economic Research.

- Charness, G., F. Feri, M. A. Meléndez-Jiménez, and M. Sutter (2014). Experimental games on networks: Underpinnings of behavior and equilibrium selection. *Econometrica* 82(5), 1615–1670.
- Chen, F., A. Griffith, A. Cottrell, and Y.-L. Wong (2013). Behavioral responses to epidemics in an online experiment: using virtual diseases to study human behavior. *PloS one* 8(1), e52814.
- Chen, M. K., J. A. Chevalier, and E. F. Long (2021). Nursing home staff networks and COVID-19. *Proceedings of the National Academy of Sciences* 118(1), 1–7.
- Choi, S., E. Gallo, and S. Kariv (2016). Networks in the laboratory. In Y. Bramoullé, A. Galeotti, and B. Rogers (Eds.), *The Oxford Handbook of the Economics of Networks*, pp. 440–475. Oxford University Press.
- Choi, S. and J. Lee (2014). Communication, coordination, and networks. *Journal of the European Economic Association* 12(1), 223–247.
- Chudik, A., K. Mohaddes, M. H. Pesaran, M. Raissi, and A. Rebucci (2020). A counterfactual economic analysis of COVID-19 using a threshold augmented multi-country model. Working Paper 27855, National Bureau of Economic Research.
- Crosetto, P. and A. Filippin (2013). The “bomb” risk elicitation task. *Journal of Risk and Uncertainty* 47(1), 31–65.
- de Ridder, D., F. Kroese, and L. van Gestel (2022). Nudgeability: Mapping conditions of susceptibility to nudge influence. *Perspectives on Psychological Science* 17(2), 346–359.
- Duflo, E. (2020). Field experiments and the practice of policy. *American Economic Review* 110(7), 1952–73.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020). The macroeconomics of epidemics. Working Paper 26882, National Bureau of Economic Research.
- Falk, A., T. Neuber, and N. Szech (2020). Diffusion of being pivotal and immoral outcomes. *The Review of Economic Studies* 87(5), 2205–2229.
- Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review* 90(4), 980–994.
- Fenichel, E. P., C. Castillo-Chavez, M. G. Ceddia, G. Chowell, P. A. G. Parra, G. J. Hickling, G. Holloway, R. Horan, B. Morin, C. Perrings, et al. (2011). Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences* 108(15), 6306–6311.
- Frey, V., R. Corten, and V. Buskens (2012). Equilibrium selection in network coordination games: an experimental study. *Review of Network Economics* 11(3).
- Fudenberg, D., D. G. Rand, and A. Dreber (2012, April). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review* 102(2), 720–49.

- Galasso, V., V. Pons, P. Profeta, M. Becher, S. Brouard, and M. Foucault (2020). Gender differences in COVID-19 related attitudes and behavior: Evidence from a panel survey in eight OECD countries. Working Paper 27359, National Bureau of Economic Research.
- Gallo, E. and C. Yan (2015). The effects of reputational and social knowledge on cooperation. *Proceedings of the National Academy of Sciences* 112(12), 3647–3652.
- Gallo, E. and C. Yan (2022). Efficiency and equilibrium in network games: An experiment. *forthcoming in The Review of Economics and Statistics*.
- Gneezy, U., S. Meier, and P. Rey-Biel (2011). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives* 25(4), 191–210.
- Gneezy, U. and A. Rustichini (2000a). A fine is a price. *The Journal of Legal Studies* 29(1), 1–17.
- Gneezy, U. and A. Rustichini (2000b). Pay enough or don't pay at all. *The Quarterly Journal of Economics* 115(3), 791–810.
- Gollwitzer, A., C. Martel, W. J. Brady, P. Pärnamets, I. G. Freedman, E. D. Knowles, and J. J. Van Bavel (2020). Partisan differences in physical distancing are linked to health outcomes during the COVID-19 pandemic. *Nature Human Behaviour* 4, 1186–1197.
- Hale, T., N. Angrist, E. Cameron-Blake, L. Hallas, B. Kira, S. Majumdar, A. Petherick, T. Phillips, H. Tatlow, and S. Webster (2020). Oxford COVID-19 government response tracker.
- Haushofer, J. and C. J. E. Metcalf (2020). Which interventions work best in a pandemic? *Science* 368(6495), 1063–1065.
- Hornsey, M. J., E. A. Harris, P. G. Bain, and K. S. Fielding (2016). Meta-analyses of the determinants and outcomes of belief in climate change. *Nature Climate Change* 6(6), 622–626.
- Howe, P. D., M. Mildenerger, J. R. Marlon, and A. Leiserowitz (2015). Geographic variation in opinions on climate change at state and local scales in the USA. *Nature Climate Change* 5(6), 596–603.
- Hummel, D. and A. Maedche (2019). How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioral and Experimental Economics* 80, 47–58.
- Hunter, D. J. (2020). COVID-19 and the stiff upper lip – the pandemic response in the United Kingdom. *New England Journal of Medicine* 382(16), e31.
- International Monetary Fund (IMF) (2020). Policy responses to COVID-19. Retrieved from: <https://www.imf.org/en/Topics/imf-and-covid19/Policy-Responses-to-COVID-19>.
- Jackson, M. O. and B. W. Rogers (2007). Meeting strangers and friends of friends: How are social networks? *American Economic Review* 97(3), 890–915.

- Jackson, M. O. and Y. Xing (2014). Culture-dependent strategies in coordination games. *Proceedings of the National Academy of Sciences* 111(Supplement 3), 10889–10896.
- Kissler, S. M., C. Tedijanto, E. Goldstein, Y. H. Grad, and M. Lipsitch (2020). Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. *Science* 368(6493), 860–868.
- Lau, M. S., B. Grenfell, M. Thomas, M. Bryan, K. Nelson, and B. Lopman (2020). Characterizing superspreading events and age-specific infectiousness of SARS-CoV-2 transmission in Georgia, USA. *Proceedings of the National Academy of Sciences* 117(36), 22430–22435.
- Ledyard, J. O. (1994). *Public goods: A survey of experimental research*. California Institute of Technology.
- Maier, M., F. Bartoš, T. D. Stanley, D. R. Shanks, A. J. L. Harris, and E.-J. Wagenmakers (2022). No evidence for nudging after adjusting for publication bias. *Proceedings of the National Academy of Sciences* 119(31), e2200300119.
- Mann, H. B. and D. R. Whitney (1947). On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics* 18(1), 50–60.
- Masters, N. B., S. F. Shih, A. Bukoff, K. B. Akel, L. C. Kobayashi, A. L. Miller, H. Harapan, Y. Lu, and A. L. Wagner (2020). Social distancing in response to the novel coronavirus (COVID-19) in the United States. *Plos one* 15(9), e0239025.
- More About Advertising (2021). UK government pandemic advertising tops £240m in 2020. Retrieved from: <https://www.moreaboutadvertising.com/2021/03/uk-government-pandemic-advertising-tops-240m-in-2020/>.
- Murphy, R. O., K. A. Ackermann, and M. Handgraaf (2011). Measuring social value orientation. *Judgment and Decision Making* 6(8), 771–781.
- Murray, C. J. (2020). Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months. *medRxiv*.
- Papageorge, N. W., M. V. Zahn, M. Belot, E. van den Broek-Altenburg, S. Choi, J. C. Jamison, E. Tripodi, et al. (2020). Socio-demographic factors associated with self-protecting behavior during the COVID-19 pandemic. Technical report, Institute of Labor Economics (IZA).
- Pedersen, M. J. and N. Favero (2020). Social distancing during the COVID-19 pandemic: Who are the present and future non-compliers? *Public Administration Review* 50(5), 805–814.
- Pew Research (2016). How religious is your state? <https://www.pewresearch.org/fact-tank/2016/02/29/how-religious-is-your-state/?state=alabama>.
- Pew Research Center (2019). U.S. public views on climate and energy.
- Polyakova, M., G. Kocks, V. Udalova, and A. Finkelstein (2020). Initial economic damage from the COVID-19 pandemic in the United States is more widespread across ages and geographies than initial mortality impacts. *Proceedings of the National Academy of Sciences* 117(45), 27934–27939.

- Rand, D. G., S. Arbesman, and N. A. Christakis (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences* 108(48), 19193–19198.
- Rosenkranz, S. and U. Weitzel (2012). Network structure and strategic investments: An experimental analysis. *Games and Economic Behavior* 75(2), 898–920.
- Rutjens, B. T., R. M. Sutton, and R. van der Lee (2018). Not all skepticism is equal: Exploring the ideological antecedents of science acceptance and rejection. *Personality and Social Psychology Bulletin* 44(3), 384–405.
- Shirado, H. and N. A. Christakis (2017). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature* 545(7654), 370–374.
- Sibony, A. L. (2020). The UK COVID-19 response: A behavioural irony? *European Journal of Risk Regulation* 11(Special Issue 2), 350–357.
- Simonov, A., S. K. Sacher, J. P. H. Dubé, and S. Biswas (2020). The persuasive effect of Fox News: Non-compliance with social distancing during the COVID-19 pandemic. Working Paper 27237, National Bureau of Economic Research.
- Singapore Ministry of Health (2020). Infectious diseases (COVID-19 – stay orders) regulations 2020. Retrieved from: <https://www.imf.org/en/Topics/imf-and-covid19/Policy-Responses-to-COVID-19>.
- Suri, S. and D. J. Watts (2011). Cooperation and contagion in web-based, networked public goods experiments. *Plos one* 6(3), e16836.
- Szaszi, B., A. Higney, A. Charlton, A. Gelman, I. Ziano, B. Aczel, D. G. Goldstein, D. S. Yeager, and E. Tipton (2022). No reason to expect large and consistent effects of nudge interventions. *Proceedings of the National Academy of Sciences* 119(31), e2200732119.
- U.S. Census Bureau (2019). Single Year of Age and Sex Population Estimates: April 1, 2010 to July 1, 2019. Retrieved from: <https://www.census.gov/data/tables/time-series/demo/popest/2010s-state-detail.html>.
- Verelst, F., E. Kuylen, and P. Beutels (2020). Indications for healthcare surge capacity in European countries facing an exponential increase in coronavirus disease (COVID-19) cases, March 2020. *Eurosurveillance* 25(13), 2000323.
- Vogel, G. (2022). New subvariants are masters of immune evasion. *Science (New York, NY)* 376(6594), 679–680.
- Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics* 1(6), 80–83.

A Online Appendix

This Appendix contains additional information. Section A.1 presents the theoretical framework and derives the hypotheses that we investigate in the experiment. Section A.2 describes our dataset. Section A.3 summarizes our data quality protocols. Section A.4 contains details of the analysis. Section A.5 presents sample instructions for the experiment.

A.1 Theory

In this section, we set out the theoretical model with greater generality. The model is similar to those found in the theoretical economics literature on self-protection against a contagious process on fixed networks (Acemoglu et al., 2016; Cerdeiro et al., 2017).

Consider a set of $N = 1, 2, \dots, n$ risk-neutral agents on an unweighted and undirected network $\mathbf{G}$. If a link between agents i and j is present, then $G_{ij} = G_{ji} = 1$. Otherwise $G_{ij} = G_{ji} = 0$. Each agent i simultaneously chooses whether to practice social distancing, at a cost $c > 0$ to herself.

After agents make this choice, exactly one agent is chosen uniformly at random to be exposed to COVID-19. Call this agent *patient zero*. If patient zero is practicing social distancing, then she becomes infected with probability $\gamma < 1$, and cannot pass COVID-19 on to anyone else. Otherwise, she becomes infected with probability 1 and passes it on to each of her neighbors who are also not practicing social distancing uniformly at random with probability $\alpha \in [0, 1]$. Any other agent j who becomes infected with COVID-19 passes it on to each of her own neighbors who are not practicing social distancing with probability $\alpha \in [0, 1]$, again uniformly at random. An agent does not pass it on to any neighbor who is practicing social distancing. Therefore, an agent who practices social distancing may only become infected if she is patient zero.

At the end of the game, an agent receives a benefit $b > c$ if she is not infected, and 0 otherwise. Additionally, an agent pays a fine, $f \geq 0$ if she did not practice social distancing – regardless of whether she infected anyone else or became infected herself. Note that setting $f = 0$ means no fine is present. When a subset of agents $S \subseteq N$ practice social distancing, the probability that agent i becomes infected is $p_{i|S_i}$. This depends on the network structure and on who practices social distancing. Notice that if $i \in S$ then $p_{i|S_i} = \gamma/n$, since she can only become infected if she is patient zero.

We assume that all agents are self-interested – they care only about themselves and so ignore any effects their choices have on others. Then agent i receives an expected payoff, π_i , of:

$$\pi_i = \begin{cases} (1 - \frac{\gamma}{n})b - c, & \text{if } i \in S \\ (1 - p_{i|S_i}) \cdot b - f, & \text{otherwise.} \end{cases} \quad (2)$$

Experimental algorithm. The contagion process described above is implemented in the experiment as follows. Suppose agent i is patient-zero who does not practice social distancing, and so is infected for sure. Then, i interacts once with her first-order neighbors who do not practice social distancing. We record all neighbors of i that become infected as a result of this interaction, and add them to a set $\text{Inf}(1)$. Next, each agent in set $\text{Inf}(1)$ interacts with her healthy first-order neighbors who do not practice social distancing. Any newly infected agents are added to a set $\text{Inf}(2)$, and in turn agents from this set interact with their healthy first-order

neighbors who do not practice distancing in the second round of contagion. The process ends when no more agents are infected as a result of the last contagion step. After this, payoffs are calculated based on the decisions and health outcomes of individual players.

Parameterization. We now set out some of the parameter values and the networks we use in the experiment. Doing so allows us to state the hypotheses clearly. We have two different networks: a 5-node complete network, and a 5-node star network. In the complete network, each agent is connected to every other agent. In the star network, one agent is connected to every other agent, and there is no other link in the network. We call agents in the complete network “close-knit”, agents at the center of the star “superspreaders”, and agents on the arms of the star “peripheral”. Denote these agents C , S , and P respectively.

We set the cost c of practicing social distancing equal to 35 points, the benefit b of avoiding infection equal to 100 points, and the probability γ of patient zero becoming infected if she is practicing social distancing equal to 0.5. This allows us to investigate how the Nash equilibrium predictions and efficiency vary as the rate of contagion α is varied on $[0, 1]$. Figure A1 shows the equilibrium and socially efficient sets in the complete network and the star network for all values of α . A set is socially efficient if it maximizes the sum of the expected payoffs. In the experiment, we focus on two different levels of contagion: $\alpha = 0.15$ (low contagion), and $\alpha = 0.65$ (high contagion).

Below, we formulate a set of hypotheses that are the focus of our analysis. First, fines make distancing relatively more attractive and so we expect that they weakly increase the amount of social distancing in society, while nudges have no impact on payoffs. So trivially fines have a weakly greater effect than nudges. Note that due to the non-uniqueness of equilibria, it is possible that *different* agents practice social distancing when fines are present.

Hypothesis 1. *A fine $f > 0$ for agents who do not practice social distancing weakly increases the number of agents practicing social distancing.*

Hypothesis 2. *A behavioral ‘nudge’, in the form of an informational video, has no impact on the social distancing decisions of any agent.*

Hypothesis 3. *A fine $f > 0$ increases the amount of social distancing weakly more than a behavioral ‘nudge’.*

The model predicts that the network position and contagiousness play important roles in determining social distancing decisions. Peripheral agents have few links and so a low incentive to practice social distancing. This means superspreaders ought to practice more distancing than close-knit agents, as superspreaders cannot rely on protection from others. All else equal, higher contagiousness makes infection more likely, increasing the incentive to practice distancing.

Hypothesis 4. *Superspreaders practice more social distancing than close-knit agents, and in turn, close-knit agents practice more social distancing than peripheral agents.*

Hypothesis 5. *Conditional on her position in the network, an agent practices more social distancing in the high contagion setting than in the low contagion setting.*


Regardless of their position in the network and the level of contagion, self-interested agents do not account for the benefits they provide to others – in the form of reduced infection risk –

Complete network

0	.11	.14	.2	.19	.24	.28	3	.49	3	.62	4
0	0	0	0	1	1		2		3		3

Star network

0	.13	S	.31	S	.57	S	.63	S	.72	S
0		0		S		S or 2P		S		S or 3P



Notes. Parameterisation: $b = 100$, $c = 35$, $\gamma = 0.5$, $N = 5$. For the complete network, digits indicate the number of close-knit agents practising social distancing in an efficient / equilibrium outcome. For the star network, S stands for the superspreader, and #P indicates the number of peripheral agents practising social distancing in an efficient / equilibrium outcome. Faded regions are where the equilibrium outcome is inefficient.

Figure A1: Equilibrium and socially efficient outcomes for the rate of contagion $\alpha \in [0, 1]$. The other parameters are set at: $b = 100$, $c = 35$, $\gamma = 0.5$, $N = 5$. For the complete network, digits indicate the number of close-knit agents practicing social distancing in an equilibrium/efficient outcome. For the star network, S stands for the superspreader, and #P indicates the number of peripheral agents practicing social distancing in an equilibrium/efficient outcome.

when they practice social distancing. Therefore, we expect that agents practice less social distancing than would be optimal. Hypothesis 6 confirms this intuition for the complete network, while Hypothesis 7 shows that there can be too little or the right amount of social distancing (relative to the social optimum) in the star. Moreover, for certain other values of α , *too much* social distancing is also possible, due to the presence of multiple equilibria in the star network.

Hypothesis 6. *In the complete network, the Nash equilibrium involves fewer agents practicing social distancing than the social optimum.*

Hypothesis 7. *In the star network: (1) in the low contagiousness environment, the Nash equilibrium involves fewer agents practicing social distancing than the social optimum; (2) in the high contagiousness environment, the Nash equilibrium and the social optimum coincide (with the hub agent practicing social distancing).*

Practicing social distancing reduces the variability of an agent's own payoffs (as she is definitely healthy unless randomly selected to be patient zero) and also has external benefits for everyone else in the network (as it reduces the probability that any other agent becomes infected). Therefore more risk-averse agents and more altruistic preferences weakly increase the amount of social distancing an agent does. This does not depend on how we choose to model risk preferences and altruistic preferences.

Table A2: Summary statistics the main controls.

Variable	$\bar{X}$	s.d.	Comments
<u>Demographic controls</u>			
Age	43.91	14.13	Measured in years
Gender	0.47	0.02	Female = 1
Race	0.84	0.02	White = 1
Education	15.34	2.04	Measured in years
Out of labor force	0.19	0.02	Yes = 1
Unemployed	0.11	0.02	Yes = 1
Christian religion	0.39	0.02	Yes = 1
Other religion	0.12	0.02	Yes = 1
<u>Preference controls</u>			
Protect others	0.57	0.02	Yes = 1, see text for details
Ideology score	3.08	3.06	$\in [0, 12]$, see text for construction details
Global warming skepticism score	1.99	3.10	$\in [0, 12]$, see text for construction details
BRET score	34.73	18.23	$\in [0, 100]$
SVO type	0.58	0.02	Prosocial = 1
<u>Location-based controls</u>			
Population density	2.35	6.32	Measured in 1,000's at county level
Cumulative cases	63.43	86.33	Measured in 1,000's at state level
Daily deaths	54.81	64.23	Measured at state level
Stay-at-home-orders	0.62	0.02	Yes = 1, measured at state level

Sample size is 400, except for SVO task which 6 subjects did not complete; $\bar{X}$ – mean value, or proportion in case of binary variables; s.d. – standard deviation.

Hypothesis 8. *The amount of social distancing an agent practices is weakly increasing in her level of risk aversion.*

Hypothesis 9. *An agent with other-regarding (i.e. altruistic) preferences practices weakly more social distancing than in the Nash equilibrium.*

A.2 Dataset

This section supplements Section 2 and provides additional details on our dataset. Apart from data on subjects' decisions in the experiment, we collect data on a set of variables, which can be broadly categorized as follows: demographic controls, preference controls, and location-based controls. Table A2 presents summary statistics for some of these controls.

Demographic controls. All participants are residents of the US, 47% are female, and the mean age is just under 44 years. 84% of our subjects are white, and 70% are employed (either full- or part-time). Further, we estimate the number of years of education received based on educational attainment.¹⁸ Based on this, the average subject in our sample has 15.3 years of

¹⁸To do this we assume that all subjects took the standard number of years to complete each qualification, and undertook no education that did not lead to a qualification.




Figure A2: Bomb Risk Elicitation Task (BRET) and Social Value Orientation (SVO) Scale distributions for the sample. We draw a vertical line on the subplots for each subject whose score in BRET/SVO is of the corresponding value. More intense line color indicates that more subjects are concentrated at that value.

education. Finally, 51% of our participants consider themselves religious, with 39% reporting Christianity as their religion.

Preference controls. A total of 92.5% of subjects reported that they try to stay at home as much as possible because of the COVID-19 pandemic. Of these, 61.6% (or 57% of the full sample) report that the desire to ‘protect others’ is one of the main reasons behind this decision.

To capture subjects’ political ideology, we construct an index from three questions in the post-experimental survey. These questions ask about: 1) subjects’ support for President Donald Trump’s handling of the COVID-19 pandemic, 2) their support for universal healthcare, and 3) their belief that social distancing measures impose unjustified economic costs. All are on a 5-point Likert scale (Likert, 1932), and so yield a 13-point index (0-12). Responses that indicate support for President Trump, opposition to universal healthcare, and belief in the unjustified economic costs of social distancing measures are scored positively. Higher scores, therefore, indicate a more Conservative ideology. In our sample, the mean ideology score is 3.1 (s.d. 3.1), and over 75% of subjects have a score of 5 or less.

We also collect information on participants’ attitudes to climate change. Specifically, our post-experimental survey includes three questions asking whether subjects believe that global warming is (1) happening, (2) caused mostly by human activity, and (3) affecting weather in the United States (Howe et al., 2015). The resulting climate change attitudes index is on a 0-12 scale, with higher scores indicating greater skepticism towards climate change. The mean score in our sample is 1.99.

Further, we collect information on subjects’ social value (SVO) and risk (BRET) preferences. Figure A2 presents the distributions of these preferences for our sample. The average subject in the sample is moderately risk-averse, with a BRET score of 34.73 boxes. When it comes to social values, the majority of our subjects (58%) are classified as prosocial (with SVO angle $\geq 22.45^\circ$ and $< 57.15^\circ$). The second-largest category in the sample are individualists (the angle is $\geq 12.04^\circ$ and $< 22.45^\circ$). Notice that we only have 1 subject who is classified as competitive (angle $< -12.04^\circ$), and no subjects classified as altruistic (angle $\geq 57.1^\circ$).

Location-based controls. As part of both recruitment and the experiment, we collect subjects’ IP-addresses. We use these to infer their geographic location, down to the county-level.

Table A3: Age-gender-location distribution of the US adult population.

Region	18-25		25-30		30-35		35-45		45-55		55+	
	M	F	M	F	M	F	M	F	M	F	M	F
I	0.29	0.29	0.20	0.19	0.19	0.19	0.34	0.35	0.38	0.41	0.85	1.00
II	0.51	0.50	0.41	0.40	0.39	0.39	0.70	0.71	0.73	0.77	1.51	1.83
III	0.56	0.54	0.43	0.42	0.41	0.41	0.74	0.76	0.78	0.81	1.70	2.01
IV	1.21	1.15	0.92	0.91	0.83	0.84	1.58	1.65	1.66	1.74	3.64	4.33
V	0.99	0.95	0.73	0.70	0.67	0.66	1.27	1.27	1.31	1.33	2.88	3.33
VI	0.83	0.78	0.63	0.61	0.60	0.58	1.10	1.10	1.01	1.03	1.97	2.29
VII	0.28	0.26	0.19	0.18	0.18	0.18	0.34	0.34	0.33	0.33	0.77	0.89
VIII	0.25	0.23	0.19	0.18	0.18	0.17	0.33	0.31	0.28	0.27	0.60	0.65
IX	0.98	0.92	0.81	0.76	0.75	0.71	1.34	1.31	1.27	1.28	2.54	2.92
X	0.26	0.24	0.22	0.2	0.21	0.20	0.38	0.36	0.34	0.34	0.76	0.85
Total:											100%	

All numbers are in percentage terms; ‘18-25’ etc. – age categories; ‘M’ – males, ‘F’ – females; we divide states into 10 Standard Federal Regions (Office of Management and Budget) defined as follows: I – CT, ME, MA, NH, RI, VT; II – NJ, NY; III – DE, DC, MD, PA, VA, WV; IV – AL, FL, GA, KY, MS, NC, SC, TN; V – IL, IN, MI, MN, OH, WI; VI – AR, LA, NM, OK, TX; VII – IA, KS, MO, NE; VIII – CO, MT, ND, SD, UT, WY; IX – AZ, CA, HI, NV; X – AK, ID, OR, WA. Note that we exclude US territories.




Figure A3: Location distribution, created with www.mapchart.net. States are divided into 10 Standard Federal Regions, using the same definition as Table A3. X/Y indicate participants counts, where X indicates the actual count in our sample, and Y – the target representative count.

Figure A3 shows the region-level distribution of our sample.

Further, using location information, we create a set of location-based control variables. First, for each subject, we record population density at the county level. The average density in the sample is 2.4k per square mile (U.S. Census Bureau, 2011, 2019). Second, we use location data together with data on the development of the COVID-19 pandemic in the US. In particular, for each subject, we record the number of cumulative COVID-19 cases (average 63.4k) and new COVID-19 related deaths (average 59) in the state on the day of the experiment. We also record whether stay-at-home orders were in place in the participant’s state on the day of the experiment. In our sample, 62% of the subjects had stay-at-home orders in place on the day they participated in our experiment.

Sample representativeness. We check the representativeness of our sample along the three dimensions – i.e. age, gender, location. Table A3 presents the distribution of the US adult population. The gender distribution of our sample is not statistically different from that of the adult US population (two-sided t-test, $p = 0.11$). Further, Figure A3 presents a map of the US divided into the 10 Standard Federal Region. For each region, we indicate the target number of people and the realized count in our sample. The difference between the observed and target location distributions is not significant (two-sided χ^2 , $p = 0.08$). When it comes to age, our distribution is unfortunately not quite representative of the adult US distribution (two-sided χ^2 , $p < 0.0001$). In particular, the 35-45 and 45-55 age categories are over-represented, while the 55+ is under-represented.

Sample balance. We check that our sample is balanced with respect to treatment assignment along the three dimensions of interest – gender, age category and geographical region of residence. We run a Pearson chi-squared test on participants’ treatment assignment (an 8-category variable) against each of these three variables separately. All tests are not significant at the 10% level. So we fail to reject the null hypothesis of no difference between our 8 treatments with respect to the three demographic characteristics. Table A4 reports the results of these tests.

Table A4: Pearson Chi-squared tests for balance of treatment assignment.

Variable	χ^2	p-val	df
Gender (Female = 1)	6.58	0.4735	7
Age (6 categories as in Table A3)	44.03	0.1409	35
Region (10 categories as in Table A3)	52.73	0.8185	63

χ^2 – value of the test statistic; p-val – corresponding p-value, df – degrees of freedom (number of distinct categories minus the number of parameters).

A.3 Data quality protocols

Ensuring data quality is important, particularly in online experiments where experimenters have relatively less control of the environment. This section explains our protocols for ensuring data quality.

Recruitment. We only make our survey visible to US residents who have completed at least 500 Human Intelligence Tasks (HITs) on MTurk *and* have an approval rate of at least 96%.¹⁹

¹⁹This means that the task-setters judged that they completed the tasks as required in at least 96% of submitted tasks.

We keep track of IP addresses of participants and ensure that no two persons with the same IP address take part in the experiment. Multiple participation is not allowed even by participants who have only seen the instructions.

Our recruitment survey also includes an attention check, which was failed by about 3% of our standing panel. Existing research suggests that excluding subjects based on failed attention checks may not necessarily be appropriate (Goodman et al., 2013). Therefore, we do not exclude subjects based on this attention check and rather use it as a robustness check for our analysis. In practice, 13 out of 400 subjects (a rate similar to that of the overall standing panel) in our dataset failed this check and excluding their data from the analysis does not affect the results reported in this paper.

Experiment. There are two main elements of our data quality control protocol at the experimental stage – comprehension checks and the use of timers. First, subjects are quizzed on their understanding of instructions (both Baseline and Intervention). For the Baseline, they need to answer three questions concerning their understanding of the contagious mechanism and the way points for their actions are calculated. Subjects have three attempts and need to answer all three questions correctly to participate in the study. Each failed attempt is followed by an explanation of correct answers, and questions are different between attempts. After the instructions for the Intervention part of the experiment, subjects need to answer one question – either asking how the fine works or what the video is about – and have three attempts. These procedures are standard for all experiments in economics, both lab and online. We use the number of attempts subjects took to complete the instructions as a control. This variable is not significant at all conventional levels (see Table A7).

Next, throughout the experiment – both Baseline and Intervention parts – in each round subjects are given 20 seconds to make their social distancing decision. If a subject fails to submit the decision in time, she receives a penalty of 200 points and her decision is automatically recorded as a “No”. A participant who fails to submit her decisions in three consecutive rounds is disqualified from the experiment and does not receive any compensation. A subject who is disqualified is replaced by a bot with pre-specified behavior (we use 4 different types of behavior) which allows the other members of her group to complete the experiment. Note that if the dropout occurs in Baseline, then the remaining subjects do not participate in the second part of the experiment, and just proceed to the post-experimental survey.

In practice, 6 subjects failed to submit their choices in three consecutive rounds of the Baseline. We had no instances of disqualification of this type in the Intervention part. We do not use data from the corresponding 6 groups in our analysis, but incorporating it back does not affect the results in robustness checks. The dataset analyzed in the paper contains a total of 13 missed choices across the 400 participants (0.08% of the total number of decisions), which has no material impact on the results reported in this paper.

Finally, another potential stage where subjects can lose attention and drop out is between the Baseline and Intervention stages of our experiment. We cap waiting time after completing the quiz for the Intervention part of instructions to 10 minutes. That is, if the longest waiting subject in a group has been waiting for 10 minutes or more, the whole group skips the Intervention part and proceeds to post-experimental survey to avoid excessive waiting. In practice, we discarded 2 groups as a result of this ‘bottleneck’ – one in the fine treatment and another in the nudge. Again, the data from these groups is not part of our main analysis, but we perform a robustness check to ensure that our results are robust to using it.

Post-experimental survey. We include an attention check into our post-experimental ques-

tionnaire which was failed by 4 of the 400 subjects in our dataset. Again, we do not discard those subject from our main analysis, but dropping their data as part of robustness checks does not alter our results.

A.4 Statistical analysis

This section presents a detailed statistical analysis. Section A.4.1 tests for a structural break in the data between Baseline and Intervention. Section A.4.2 then presents analysis at the aggregate (group) level, and Section A.4.3 focuses on the individual level. Next, Section A.4.4 contains convergence analysis for individual behavior, while Section A.4.5 reports on the difference between attitudes to interventions in the sample and observed behavior.

A.4.1 Structural break identification

We use a Wald test to look for a single structural break in social distancing decisions over time. This involves performing a simple linear regression of social distancing decisions (left-hand side) on the round number (right-hand side), and then testing whether the coefficient on the time variable exhibits a structural break. The test assumes there is at most one structural break but is agnostic as to at what point it may occur.

We perform this test on the data at three levels of aggregation. First, we consider all observations together. Second, we split by the rate of contagion and the policy intervention. This gives four equal-sized group: (1) 15% contagion, fine; (2) 15% contagion, nudge; (3) 65% contagion, fine; and (4) 65% contagion, nudge. Finally, we subdivide each of the four groups above by network type (i.e. complete vs star).

If the policy intervention is effective then we expect the test would find a structural break at round 21. With all observations together, the Wald test finds the structural break at round 21 ($p < 0.001$). When considering the dis-aggregated data, the Wald test finds a structural break at round 21 for each of the subsets, except for complete networks with the nudge treatment.

Table A5: Results from Structural Break analysis using the Wald Test

Low contagion: $\alpha = 0.15$				High contagion $\alpha = 0.65$			
Treatment	Network	Break	p-value	Treatment	Network	Break	p-value
fine	all	21	< 0.001	fine	all	21	< 0.001
fine	complete	21	< 0.001	fine	complete	21	< 0.001
fine	star	21	< 0.001	fine	star	21	< 0.001
nudge	all	21	< 0.001	nudge	all	21	< 0.001
nudge	complete	7	< 0.001	nudge	complete	16	< 0.001
nudge	star	21	< 0.001	nudge	star	21	0.009

A.4.2 Aggregate level

In this section, we focus on the aggregate analysis of the decision data and conduct a set of nonparametric tests. Since individual observations are correlated at the group level, our unit of observation is a group of 5 participants. Recall that we have 80 groups in total, equally

distributed across 8 treatments. For each group, we calculate the average proportion of participants that practice social distancing in the last 10 rounds of both parts of the experiment. Below, we refer to this average as ‘distancing levels’. We discard the first 10 rounds of both parts of the experiment because our participants display clear convergence behavior. In particular, by round 11, in both Baseline and Intervention, at least 80% of participants converge to a particular strategy in all treatments. Further details on convergence behavior is in Section A.4.4. All key results below are also robust to including all 20 rounds of both parts of the experiment.

Table A6 presents the results of the non-parametric tests. All of the findings of this section are robust to using parametric analogs of these non-parametric tests. In our aggregate analysis, we test the set of hypotheses in Section A.1. Note that we cannot test Hypotheses 8 and 9 here, which are instead covered in Section A.4.3.

Part 1 of Table A6 shows that the data from the Baseline from fine and nudge treatments can be pulled together as these samples are statistically identical at all conventional significance levels. This result is robust to various specifications.

In Part 2 of Table A6, we examine the effect of introducing a fine in Intervention. From the table, it is evident that a fine has a positive and statistically significant effect on distancing levels compared to the Baseline with no fine in all specifications considered. Most results in this part are statistically significant at least at the 5% level. The effect of the fine is also sizable – it results in an increase in mean distancing levels of 5% to 9% depending on the specification. The following result summarizes the above observations. Notice that the numbering of all results in this section corresponds to the numbering of Hypotheses in Section A.1.

Result A-1. *The introduction of the fine increases the level of social distancing. The effect is observed for all three positions as well as for the two rates of contagion. The effect is both statistically significant (WSR, $p < 0.05$ in all but one specification where $p = 0.09$) and sizable in magnitude.*

In Part 3 of Table A6 we repeat the analysis from Part 2, but now looking at the nudge Intervention rather than the fine. The tests find a statistically significant effect of the nudge in only 4 of the 6 specifications. Generally, we can see that the nudge has a smaller effect which is not as robust as that of the fine. In particular, the effect of the nudge is very significant and large in magnitude (7%) under high contagion, but disappears completely both statistically and in terms of economic magnitude under low contagion. Also, the nudge seems to have no effect on the superspreader. Note also that the effects of the nudge are not robust to using data from all rounds of Baseline and Intervention.

Result A-2. *The introduction of the nudge generally increases the levels of social distancing. The effect, however, is not robust to alternative specifications and is smaller in magnitude than that of the fine. In particular, the effect is statistically significant (WSR, $p = 0.0003$) and sizable in magnitude under high contagion, but not low contagion (WSR, $p = 0.5$). Further, a statistically significant effect is observed for the close-knit ($p = 0.05$) and peripheral (WSR, $p = 0.06$) participants but not for superspreaders.*

Part 4 of Table A6 shows that the data from the Intervention in the second part of the experiment is not statistically different for the fine and nudge treatments. In particular, even though the mean level of distancing in Intervention in fine treatments is higher than that in nudge treatments, our non-parametric testing fails to find any difference between them in 4 out of 6

Table A6: Non-parametric analysis at the aggregate level, last 10 rounds used.

Null Hypothesis	Rate of contagion	Position	Alt	Test	n	$\Delta\bar{X}$	$\Delta\tilde{X}$	p-val	Sig
Part 1. Baseline									
<i>Fine vs nudge</i>	all	all	2-sided	MW	40	.03	.05	.33	
	all	close-knit			20	.01	0	.81	
	all	superspreader			20	.02	.15	.47	
	all	peripheral			20	.07	.1	.20	
	15%	all			20	.02	-.01	.73	
	65%	all			20	.05	.01	.34	
Part 2. Intervention vs Baseline. Hypothesis 1 (fine has an effect)									
<i>Fine: Intervention vs Baseline</i>	all	all	1-sided	WSR	40	.07	.03	.0008	***
	all	close-knit			20	.07	.07	.01	**
	all	superspreader			20	.09	0	.09	*
	all	peripheral			20	.07	0	.04	**
	15%	all			20	.05	.03	.04	**
	65%	all			20	.09	.1	.003	***
Part 3. Intervention vs Baseline. Hypothesis 2 (nudge has an effect)									
<i>Nudge: Intervention vs Baseline</i>	all	all	1-sided	WSR	40	.04	.04	.008	***
	all	close-knit			20	.04	.07	.05	**
	all	superspreader			20	.01	.05	.50	*
	all	peripheral			20	.04	.06	.06	
	15%	all			20	0.0	-0.1	.50	***
	65%	all			20	.07	.08	.0003	
Part 4. Intervention. Hypothesis 3 (fine is more effective than nudge)									
<i>Fine vs nudge</i>	all	all	1-sided	MW	40	.07	.04	.08	*
	all	close-knit			20	.04	0	.24	*
	all	superspreader			20	.1	.1	.06	
	all	peripheral			20	.09	.04	.14	
	15%	all			20	.07	.03	.15	
	65%	all			20	.07	.03	.12	
Part 5. Baseline. Hypothesis 4 (superspreader > close-knit > peripheral)									
<i>Superspreader vs close-knit</i>	all	—	1-sided	MW	40	.1	.19	.004	***
	15%	—			20	.1	.18	.03	**
	65%	—			20	.1	.21	.03	**
<i>Close-knit vs peripheral</i>	all	—	1-sided	MW	40	.16	.11	.0001	***
	15%	—			20	.17	.20	.006	***
	65%	—			20	.15	.01	.003	***
<i>Superspreader vs peripheral</i>	all	—	1-sided	WSR	40	.26	.30	<.0001	***
	15%	—			20	.27	.38	.0007	***
	65%	—			20	.25	.31	.0001	***

Continued on next page

Continued from previous page

Null Hypothesis	Rate of contagion	Position	Alt	Test	n	$\Delta\bar{X}$	$\Delta\tilde{X}$	p-val	Sig
Part 6. Intervention. Hypothesis 4 (superspreader > close-knit > peripheral)									
<i>Super-spreader vs close-knit</i>	all	—	1-sided	MW	40	.1	.17	.002	***
	15%	—			20	.12	.16	.01	**
	65%	—			20	.08	.12	.01	**
<i>Close-knit vs peripheral</i>	all	—	1-sided	MW	40	.16	.16	.0002	***
	15%	—			20	.21	.18	.0009	***
	65%	—			20	.11	.16	.02	**
<i>Superspreader vs peripheral</i>	all	—	1-sided	WSR	40	.25	.33	<.0001	***
	15%	—			20	.32	.34	.0001	***
	65%	—			20	.19	.28	.0007	***
Part 7. Baseline. Hypothesis 5 (more distancing in high contagion compared to low contagion)									
<i>65% vs 15% rate of contagion</i>	all	all	1-sided	MW	40	.1	.06	.01	**
	all	close-knit			20	.09	.07	.09	*
	all	superspreader			20	.09	.1	.16	
	all	peripheral			20	.11	.16	.03	**
Part 8. Intervention. Hypothesis 5 (more distancing in high contagion compared to low contagion)									
<i>65% vs 15% rate of contagion</i>	all	all	1-sided	MW	40	.15	.14	.0002	***
	all	close-knit			20	.12	.14	.009	***
	all	superspreader			20	.08	.1	.05	**
	all	peripheral			20	.22	.16	.002	***
Part 9. Baseline. Hypothesis 6 (actual versus equilibrium/efficient outcomes in complete network)									
<i>Actual vs equilibrium</i>	15%	all	1-sided	MW	20	.63	.62	<.0001	***
	65%	all			20	.11	.09	.004	***
<i>Actual vs efficient</i>	15%	all	1-sided	MW	20	.23	.22	<.0001	***
	65%	all			20	-.09	-.11	.01	**
Part 10. Baseline. Hypothesis 7 (actual versus equilibrium/efficient outcomes in star network)									
<i>Actual vs equilibrium</i>	15%	superspreader	1-sided	MW	20	.73	.8	<.0001	***
		peripheral			20	.46	.42	<.0001	***
	65%	superspreader			20	-.18	-.1	<.0001	***
		peripheral			20	.57	.59	<.0001	***
<i>Actual vs efficient</i>	15%	superspreader	1-sided	MW	20	-.27	-.2	<.0001	***
		peripheral			20	.46	.42	<.0001	***

Null Hypothesis: for each hypothesis, sample #1 and sample #2 appear consecutively *in italics*; Alt: either 2-sided or 1-sided alternative, a 1-sided alternative is always s.t. median(sample #1) > median(sample #2); Test: Mann-Whitney U-test (MW) for unmatched samples and Wilcoxon Signed-Rank test (WSR) for matched samples; n: the number of observations per sample, where one observation is one group of 5 subjects; $\Delta\bar{X}$: difference in means between the two samples; $\Delta\tilde{X}$: difference in medians between the two samples; p-val: p-value for the test of the Null Hypothesis against the Alternative; Sig: significance of the test where * – 10%, ** – 5%, *** – 1%.

specifications. We do, however, find evidence that the fine is more effective than the nudge (1) when data is pulled across treatments, and (2) for the superspreader. The effect is substantial in magnitude (7-10%) and significant at 10% level.

Result A-3. *There is limited evidence that the fine is more effective than the nudge, and the effect is not robust. In particular, the fine increases distancing levels by more than the nudge (1) when all data is pulled together (MW, $p = 0.08$), and (2) separately for the superspreader (MW, $p = 0.06$). The effect, however, is not statistically significant for other specifications.*

In Part 5 of Table A6, we compare distancing levels in the three positions – close-knit, superspreader, and peripheral – in Baseline. From the table, we can see that (1) distancing levels are higher in the superspreader position relative to both the close-knit and the peripheral, and (2) higher in the close-knit relative to the peripheral. We run hypothesis tests both (a) aggregating over the rate of contagion, and (b) separately by two levels of the rate of contagion. All results in this part of the table are significant at the 5% level, with most also being significant at 1%. The differences between mean distancing levels in different positions are also large in magnitude – between 10% and 27% depending on the specification.

Part 6 of Table A6 repeats the analysis in Part 5, but now focusing on Intervention, rather than Baseline. The conclusions here are the same as above, with all results being significant at the 5% and most also at the 1% level.

Result A-4. *Superspreaders practice more social distancing than close-knit participants, who, in turn, practice more distancing than peripheral participants. The differences are both statistically significant (MW, $p < 0.05$) and large in magnitude in all specifications.*

In Parts 7 and 8 of Table A6, we investigate the effects on the rate of contagion on distancing behavior, separately for Baseline and Intervention. We find that there is generally significantly more social distancing under 65% rate of contagion relative to 15%. The test is not statistically significant only for the superspreader in Baseline. The effect is also large in magnitude – an average of 8-22% depending on the specification. Note that the effect on superspreaders is smallest, but is probably explained, at least in part, by ceiling effects. In particular, in all treatments, the mean distancing levels in the last 10 rounds in the superspreader position are at least 70%, so there is limited room for a further increase.

Result A-5. *Subjects in the experiment generally practice more social distancing in high contagion environments relative to low contagion environments. The effect is statistically significant for close-knit (MW, $p = 0.9$ in Baseline and $p = 0.009$ in Intervention) and peripheral (MW, $p = 0.03$ and $p = 0.002$ in Baseline and Intervention respectively) subjects. For superspreaders, the effect is relatively smaller, and less robust to specifications (MW, $p = 0.16$ in Baseline and $p = 0.05$ in Intervention).*

Finally, in Parts 9 and 10 of the table, we focus on testing Hypotheses 6 and 7. Specifically, we compare the outcomes in the complete and the star networks to (1) predictions of the Nash equilibrium, and (2) efficient outcomes.

We find that the levels of social distancing are not in line with theoretical predictions, with all tests being statistically significant at least at the 5% level. In particular, distancing levels in the complete network are well above those predicted by Nash equilibrium. The difference is particularly large in the low contagion environment, where the equilibrium prediction is no social distancing, but the actual average level of social distancing in the last 10 rounds of Baseline

stands at 62%. Distancing levels are also higher than those predicted by efficiency under low contagion. On the other hand, in the high contagion environment, the actual distancing levels are below efficiency requirements.

Result A-6. *Prior to intervention, the observed levels of social distancing in the complete differ from equilibrium and efficiency predictions. In particular, more social distancing is observed than predicted by Nash equilibrium. When it comes to efficiency, actual levels of distancing are above efficient under low contagion, but below efficient under high contagion. All results are statistically significant (MW, $p \leq 0.01$).*

In the star network, the amount of social distancing done by the peripheral participants is well above both equilibrium predictions and efficiency requirements. On the other hand, the levels of social distancing observed for the superspreaders is above the equilibrium but below the efficient level for low contagion, and below both the equilibrium and efficient levels for high contagion.

Result A-7. *Prior to intervention, the observed levels of social distancing in the star network differ from equilibrium and efficiency predictions. The result is true for both positions and both rates of contagion. In low contagion, participants in both positions practice more distancing than predicted by equilibrium analysis, while in high contagion superspreaders practice less and peripheral participants practice more than predicted. When it comes to efficiency, peripheral agents practice more social distancing and superspreaders practice less social distancing in both high and low contagion. All results are statistically significant (MW, $p < 0.0001$).*

A.4.3 Individual level

We now turn to individual behavior, and examine the determinants of individual social distancing decisions. To do this, we use a Random Effects Logit model. This implicitly models the following Random Utility framework for binary choice:

$$y_{it} = 1 \iff x_{it}\beta + v_i + \epsilon_{it} > 0. \quad (3)$$

According to this model, subjects choose to practice social distancing if and only if they receive greater utility from doing so than from not doing so. Note that (implicitly) we have normalized the utility from not distancing to zero – this is without loss of generality. This is a flexible and tractable way to model binary decisions. The model assumes that the subject-specific random effect, v_i , is normally distributed and that ϵ_{it} follows the logistic distribution.

In addition, we only use data from the final 10 rounds of each part of the experiment, because it takes time for subjects' behavior to converge. We have five categories of controls that collectively cover a wide variety of factors. First, and most obviously, we control for the experimental treatments: the fine, nudge, rate of contagion, and network position (model 1 of table A7). As we randomly assign subjects to these treatments, we can be confident that their effects are causal.

Next, we add controls for social demographics (in model 2), geographic and institutional factors (in model 3), social and risk preferences (in model 4), and finally, ideology (in model 5). To complete our model, we also control for interactions between ideology and the policy interventions (fine/nudge). This is because we find that subjects' responsiveness to the fine depends

on their ideology – more conservative subjects are *less responsive* to the fine.

Results. Ex ante, it seems plausible that each of the controls could be related to social distancing decisions. However, as we can see in model 6, only some of them are.²⁰ All experimental treatments have a significant effect – both statistically and in terms of economic relevance – and their direction is intuitive. Statistical tests are t-test on coefficients in the Random Effects Logit regression (REL hereafter), or t-test on coefficients in the instrumental variables Random Effects Logit (IV, hereafter).

First, the fine significantly increases social distancing (Hypothesis 1), while the nudge marginally increases social distancing (Hypothesis 2). Qualitatively, the fine is more effective (the statistical significance of the nudge is also not robust), but in our preferred model, the difference between the two effects is not statistically significant (Hypothesis 3). Note that the numbering of the results corresponds exactly to the number of the hypotheses in Section A.1, and the “I” prefix denotes individual-level results.

Result I-1. *The fine increases the probability that an individual practices social distancing (REL, $p = 0.001$).*

Result I-2. *The nudge marginally increases the probability that an individual practices social distancing (REL, $p = 0.09$). The effect is approximately half that of the fine and is not robust to changes in the regression specification.*

Result I-3. *The effect of the fine is larger than the effect of the nudge, but the difference is not statistically significant (REL, $p = 0.11$).*

Second, superspreaders distance more than ‘close-knit’ agents, even though they have the same number of links in the network (Hypothesis 4). This suggests that when subjects know they are relatively highly connected, they take some action to compensate. Note that our experimental design cannot identify *why* they do this – whether to protect themselves or to protect others. Peripheral agents distance significantly less – being less exposed to community contagion in the first place reduces the likelihood that agents take protective action. Further, subjects do more social distancing in a high contagion environment than in a low contagion environment, irrespective of their network position (Hypothesis 5).

Result I-4. *Superspreaders practice social distancing with a greater probability than close-knits subjects (REL, $p = 0.05$), who in turn, practice social distancing with a greater probability than peripheral subjects (REL, $p < 0.0001$).*

Result I-5. *Subjects practice social distancing with a greater probability in the high contagion setting than in the low contagion setting (REL, $p < 0.0001$).*

Beyond the experimental treatments, we find significant effects for age, gender, race, social and risk preferences, and ideology. Older, female, and non-white subjects all practice more social distancing. More risk-averse agents and prosocial agents (as classified by the SVO task) also practice more social distancing (Hypothesis 8 and Hypothesis 9).

Result I-8. *More risk-averse subjects practice social distancing with a higher probability (REL, $p = 0.001$).*

²⁰The p-values in all results relate to model 6.

Result I-9. *Subjects who indicate greater concern for others' well-being practice social distancing with a higher probability (REL, $p < 0.0001$).*

Perhaps the most interesting non-treatment effect is the subjects' political ideology. Subjects with stronger Conservative ideology practice less social distancing, and are marginally *less responsive* to fines. Recall that our ideology index is constructed from subjects' responses to three questions from the post-experiment questionnaire. They ask about subjects' support for President Donald Trump's handling of the COVID-19 pandemic, their support for universal healthcare, and their belief that social distancing measures impose unjustified economic costs.

Result I-10. *More conservative subjects both practice distancing with a lower probability (iv, $p = 0.04$) and are marginally less responsive to the fine (IV, $p = 0.08$).*

It is important to bear in mind that subjects were not randomly allocated to their ideology (as they were to the experimental treatments) – clearly, that would not be feasible. As such, the association we have found between ideology and social distancing decisions may not be causal – it is possible that ideology is *endogeneous*. That is, it could be correlated with some unobserved factor that affects social distancing decisions.

Therefore, we use an instrumental variables approach to deal with possible endogeneity. We use a measure of subjects' skepticism of global warming as the instrument. As a partisan issue in the United States, this is strongly correlated with political ideology (McCright and Dunlap, 2011). However, as it is unrelated to COVID-19 and the types of decisions that we ask subjects to make in this experiment, it should not have a *separate* effect on social distancing decisions.

Mathematically, this means that after we have controlled for ideology, then skepticism of global warming is uncorrelated with the error term, ϵ , in the regression. This property is required for the instrument to be *valid*, and so to give consistent estimates. Note that it is not possible to test this – it is an assumption.²¹

As we are using a Logit specification, it is not possible to do standard 2-Stage Least Squares. Instead we use a Control Function method (Train, 2009). The Control Function method takes predicted residuals from the first stage regression, and adds them into the Logit regression. This is in contrast to 2-Stage Least Squares, which takes predicted values from the first stage, and uses them instead of the (potentially) endogenous variable in the second stage.

Model 7 of table A7 reports the preferred regression specification with the instrumental variables method. Using the instrument does not have a large impact on the point estimate for the ideology variable. This suggests that there may in fact be a causal relationship between ideology and social distancing decisions – conservatives practice less social distancing.

Partial Effects. Using our preferred Logit specification – model 6 of table A7 – we can predict the probability that an agent chooses to practice social distancing. With estimated regression coefficients $\hat{\beta}$ and a set of subject characteristics x , then:

$$Pr(y_{it} = 1|x_{it}) = \frac{e^{\hat{\beta}'x}}{1 + e^{\hat{\beta}'x}}. \quad (4)$$

However, calculating partial effects with a Logit model is far from straightforward. In a non-linear model, the partial effect of one variable depends on the full set of an individual's characteristics, and so is highly heterogeneous. For example, the estimated partial effect of age

²¹See, for example, Wooldridge (2010).

Table A7: Main logit regression results.

Variables	Decision to practice distancing (binary)						
	(1)	(2)	(3)	(4)	(5)	(6) Main	(7) IV
<u>Treatment</u>							
Fine	0.755*** (0.231)	0.752*** (0.234)	0.752*** (0.234)	0.762*** (0.242)	0.761*** (0.242)	1.146*** (0.352)	1.147*** (0.352)
Nudge	0.347** (0.157)	0.353** (0.154)	0.353** (0.154)	0.344** (0.155)	0.343** (0.155)	0.438* (0.261)	0.437* (0.261)
High contagion	1.330*** (0.336)	1.477*** (0.334)	1.404*** (0.318)	1.529*** (0.330)	1.537*** (0.323)	1.547*** (0.323)	1.543*** (0.322)
Superspreader	0.975** (0.438)	1.234*** (0.450)	1.184*** (0.423)	0.927** (0.427)	0.830** (0.421)	0.833** (0.424)	0.846** (0.429)
Peripheral	-1.608*** (0.322)	-1.364*** (0.332)	-1.416*** (0.319)	-1.684*** (0.334)	-1.780*** (0.333)	-1.785*** (0.333)	-1.771*** (0.340)
<u>Demographic controls</u>							
Age (years)		0.0859*** (0.0129)	0.0858*** (0.0135)	0.0813*** (0.0144)	0.0793*** (0.0137)	0.0796*** (0.0139)	0.0796*** (0.0139)
Female		1.237*** (0.352)	1.193*** (0.367)	0.909** (0.374)	1.069*** (0.372)	1.072*** (0.372)	1.054*** (0.380)
Race = white		-1.310*** (0.462)	-1.553*** (0.480)	-1.149** (0.463)	-0.960** (0.469)	-0.955** (0.468)	-0.986** (0.464)
Education (years)		0.106 (0.0908)	0.0490 (0.0916)	0.0104 (0.0958)	0.00341 (0.0984)	0.00423 (0.0984)	0.00713 (0.0977)
Religion = Christian		-0.517 (0.409)	-0.443 (0.410)	-0.287 (0.413)	0.0632 (0.436)	0.0616 (0.437)	0.00188 (0.439)
Religion = other		-0.166 (0.543)	-0.0445 (0.560)	0.278 (0.542)	0.384 (0.525)	0.370 (0.525)	0.347 (0.523)
Out of labor force		0.0943 (0.455)	0.164 (0.484)	-0.0715 (0.486)	0.0233 (0.481)	0.00933 (0.484)	-0.00632 (0.488)
Unemployed		0.752 (0.592)	0.611 (0.571)	0.273 (0.568)	0.301 (0.538)	0.301 (0.542)	0.299 (0.543)
<u>Location-based controls</u>							
Population density			0.0320 (0.0334)	0.0409 (0.0300)	0.0429 (0.0293)	0.0435 (0.0294)	0.0430 (0.0293)
Cumulative cases			0.00572 (0.00475)	0.00570 (0.00442)	0.00444 (0.00418)	0.00450 (0.00418)	0.00470 (0.00423)
Daily deaths			-0.001131 (3.626)	-0.004514 (3.576)	-0.003625 (3.627)	-0.00366 (3.632)	-0.00382 (0.00371)
Stay-at-home order			-0.355 (0.501)	-0.343 (0.492)	-0.446 (0.452)	-0.451 (0.456)	-0.433 (0.466)
Region controls	No	No	Yes	Yes	Yes	Yes	Yes

Continued on next page

Continued from previous page

Variables	Decision to practice distancing (binary)						
	(1)	(2)	(3)	(4)	(5)	(6) Main	(7) IV
Preference controls							
Bomb risk score				-0.0339*** (0.00934)	-0.0296*** (0.00889)	-0.0296*** (0.00896)	-0.0300*** (0.00885)
Prosocial values				1.564*** (0.338)	1.485*** (0.334)	1.495*** (0.335)	1.503*** (0.334)
Protect others				1.512*** (0.332)	0.918*** (0.342)	0.926*** (0.346)	1.017** (0.395)
Quiz attempts					0.189 (0.170)	0.197 (0.171)	0.197 (0.171)
Ideology							
Conservative ideology					-0.282*** (0.0565)	-0.250*** (0.0591)	-0.207** (0.0995)
Fine–ideology interaction						-0.110* (0.0618)	-0.110* (0.0618)
Nudge–ideology interaction						-0.0266 (0.0447)	-0.0265 (0.0448)
Residual							-0.0663 (0.120)
Constant	1.456*** (0.312)	-3.547** (1.611)	-2.753* (1.627)	-2.927* (1.579)	-2.051 (1.670)	-2.213 (1.666)	-2.415 (1.700)
Observations	8,000	8,000	8,000	7,880	7,880	7,880	7,880
No of subjects	400	400	400	394	394	394	394

The variable “Quiz attempts” measure the number of attempts subjects required to pass the quiz prior to the main experiment. It is a proxy for a subject’s sophistication.

Population density = 1000’s of people per square mile. Cumulative cases = 1000’s of confirmed cases in state.

Significance levels: * – 10%, ** – 5%, *** – 1%.

depends on an agent’s gender, race, religion, degree of risk aversion – and *all other* variables in the model. Note that this even includes variables that are not statistically significant.

Given this, we calculate a variant on Average Partial Effects (APE). This is more useful than a Partial Effect at the Average; especially given our extensive use of binary variables – a Partial Effect at the Average would give us the partial effect for a subject who is 47% female, 85% white, and at a node position that simply does not exist (50% ‘close-knit’, 40% peripheral, and 10% superspreader).

To calculate our variant on an APE for a variable, for example, gender, we predicted the probability of social distancing for each of our 400 subjects, first assuming that they are all male, and then second assuming that they are all female. This gives us an individual partial effect of gender for each subject – taking an average across all subjects yields the variant on an APE. We assume that subjects are not exposed to a policy intervention (fine/nudge) when calculating APEs for all (other) variables. When the variable is continuous, we first assume that

all subjects are at the 25th percentile of that variable (based on the actual distribution in the experiment), and second that they are at the 75th percentile.

We use this variant – only looking at the 400 subjects in our experiment – due to data constraints. While data is readily available on the population-wide distributions of most of the individual variables, they are only available as marginal distributions, not as joint distributions. That is, one can easily find an age distribution, a gender distribution, and an education distribution for the US population; but they are only available separately. A joint distribution – especially over *all* of the variables in our preferred specification is not available. Therefore, our variant allows us to calculate partial effects for individuals who actually exist, and so makes our APE meaningful. Figure 4 in the main text shows a box plot Partial Effects for variables that are statistically significant.

A.4.4 Convergence

The analysis above focuses on the last 10 rounds of Baseline and Intervention. Here, we validate this choice by showing that the majority of participants exhibit clear convergence behavior after the first 10 rounds in both parts of the experiment. We define individual convergence as follows:

Definition 1. *A participant converges to a strategy s by round n if (i) she used this strategy for the last k rounds (including n), and (ii) in all subsequent rounds $[n + 1, 20]$ the number of consecutive deviations from the chosen strategy does not exceed a .*

We consider three types of convergence strategies as follows. In both networks, we look at the strategy where the subject always chooses the same action. For the star network, we also consider two extra strategies. In one strategy the participant always chooses the same action when she is in the superspreader position and the complement action when she is peripheral. In the other strategy, she always chooses the same action when she is the superspreader and alternates between the actions when peripheral.

We set $k = 4$ and $a = 2$. This means that the earliest a participant can be considered to converge to a particular strategy is by round 4 if she has not deviated from this strategy in the last four rounds, and in all subsequent rounds, she never performs more than two consecutive deviations.


Using the above definition, by round 11 at least 80% of our participants have converged to a certain strategy in all parameterizations.

As a robustness check, we also consider $a = 1$ and $a = 3$, allowing for one and three consecutive deviations respectively. Figure A4 plots the share of converged participants for each round separately by parts for $a \in [1, 3]$. We can see that the share of converged subjects does not change much when we allow for a more/less conservative definition. In particular, with $a = 1$ the share of subjects who converge by round 11 in Baseline drops to 79.8% while in Intervention it reaches 83.5%. With $a = 3$ the share of subjects who converge by round 11 in Baseline and Intervention stands at 86.3% and 90% respectively.

The above analysis suggests that it is reasonable to claim that the absolute majority of subjects converge to a particular strategy by round 11 in both parts of the experiment.

A.4.5 Perception and Behavior

As part of the post-experimental questionnaire, subjects are asked whether they agree that fines and nudges are effective in promoting social distancing. These two questions are both on a



Notes. We exclude the first 3 rounds in both parts since $k = 4$.

Figure A4: Evolution of the share of converged participants throughout the experiment separately for Baseline and Intervention, $a \in [1, 3]$. Note: we exclude the first 3 rounds in both parts since $k = 4$.

5-point Likert scale, with higher values indicating greater disagreement with the statements. Table A8 summarizes individual self-reported perceptions of fines and nudges in our sample. The majority of subjects believe that nudges are effective in promoting social distancing, while attitudes to fines appear to be polarized.

Table A8: Self-reported perception of fines and nudges, $n = 400$.

	strongly agree	agree	neither agree nor disagree	disagree	strongly disagree
fines are effective	11.75%	35.5%	21.25%	20.25%	9.25%
nudges are effective	31%	50.5%	13.75%	3.5%	1.25%

We construct a group-level index for the perception of fines, equal to one if at least three subjects in the group believe that fines are effective, i.e. they (strongly) agree with the corresponding statement, and an identical index for the perception of nudges. We then construct another index that captures changes in the observed behavior. This index is equal to one if the average level of social distancing in the last 10 rounds of Intervention is higher than in the last 10 rounds of Baseline. We then compare the perception indices to the behavior index.

Figure A5 summarizes how group-level perception of fines and nudges compares to behavior. Two patterns emerge. First, subjects seem to underestimate the effectiveness of fines: only 42% of groups in the fine treatment (55% in the nudge treatment) believe in the effectiveness of fines, whereas 70% of groups subjected to fines show an increase in average distancing levels. Second, participants seem to overestimate the effectiveness of nudges – 95% of the groups in the nudge treatment (98% in the fine treatment) believe that nudges are effective, while only 65% of the groups actually increase their average distancing levels when subjected to nudges.

To test this formally, we use non-parametric analysis with the three group-level indices. We find that irrespective of treatment they were placed in, subjects' perception of fines matches their (fines') measured effectiveness (WSR and MW, $p = 0.99$ and $p > 0.92$ respectively). On the other hand, participants overestimate the effectiveness of nudges (WSR and MW, $p = 0.001$ and $p < 0.0001$ respectively).




Figure A5: Perception of effectiveness of fines and nudges versus observed effectiveness.

A.5 Experimental instructions

Appendix A.5.1 presents instructions and interface for the experiment. Analogously, Appendices A.5.2 and A.5.3 present instructions and interface for the social preferences (SVO) and risk preferences (BRET) elicitation tasks that we employ.

A.5.1 Main Experiment

This section contains instructions for the Baseline and Intervention parts of the experiment. For the Baseline (Part 1), the rate of contagion is set at 65%. Note that the type of network and intervention do not feature in this part of instructions. For the Intervention (Part 2), we show instructions for the fine. Instructions for the nudge are similar, except that instead of explaining how the fine is implemented, participants are asked to watch a 3-minute video. The video can be accessed online at <https://youtu.be/tyf6EpSMegs>. The section concludes with Figure A6, which shows the decision and results interfaces from the main experiment.

Part 1: Instructions (page 1/6)

Welcome to this interactive experiment!

This experiment consists of two parts. You will be paid a **fixed reward of \$1** for completing all parts of the experiment. Additionally, you can earn points for your choices in Parts 1 and 2, which will be converted into \$ at the end of this experiment. There may also be a Bonus Task at the end of the experiment.

In Parts 1 and 2 of the experiment, you will **interact with a group of other real people recruited through MTurk**. Recall that you will never learn the identities of other people and no one will learn about your identity.

The expected duration of the experiment is **30-40 minutes** and your average expected total earnings will be **\$3-6 excluding the Bonus Task**. Note that you can earn less or more than this amount depending on your choices and the choices of others in your group.

At the beginning of Part 1 of the experiment, you will be randomly allocated to a **group of 5** and you will **remain in this group** for the duration of Parts 1 and 2 of the experiment. Note that you might have to wait while we are matching you with 4 other people, but we will compensate you for the time you wait.

Since this experiment is interactive, it is important that you **remain continuously attentive**, otherwise you may slow down others and may even be disqualified from the experiment.

In Part 1 of the experiment, you will be asked to play a game with the other members of your group. The instructions on the next 5 pages explain the rules of the game. You will receive more information about Part 2 after you complete Part 1.

All participants are given the **same** instructions. It is important that you read these instructions carefully. Note that there is no deception in this experiment.

Once you read the instructions, you will be required to pass a short understanding Quiz. **If you fail the Quiz, you will not be allowed to take part in the experiment and will not receive the fixed reward.**

To continue to the instructions, press the Next button below.

Part 1: Instructions (page 2/6)


In Part 1 of the experiment you are asked to play a game with the other members of your group. In what follows, you and the other members of your group are referred to as participants.

At the start of the game, you are presented with a diagram with 5 circles labeled by capital letters (P, E, C, M, Q) and lines between them.

Each circle represents a **position** – one for each participant. At the start of the game, each participant is **randomly allocated** to one of these positions in the diagram. Your position is the one colored in blue.

The lines between positions indicate the structure of interactions between participants in these positions. These lines **indicate which participants interact** with one another in the game.

An example is the diagram below. Here, you are in position M and directly interact with participants in positions P and E, but you do not interact directly with C and Q.



The next page of the instructions explains the choice you need to make in the game.

To continue to the next page, press the Next button below. To go to the previous page, press the Back button.

Part 1: Instructions (page 3/6)

In the game, you and the other participants face a **risk of getting infected with COVID-19** (commonly known as coronavirus).

The main symptoms of COVID-19 are shortness of breath, a high fever and a new, continuous cough. Most patients experience mild symptoms and recover in 1-2 weeks, but cases can progress to pneumonia and organ failure in the most vulnerable individuals.

After you learn your position in the diagram, your need to **choose whether to practice social distancing**. Below you can see what the buttons to make your choice look like.

Do you want to practice social distancing?

Yes

No

You have **20 seconds** to make your choice and the timer is displayed at the top of the interface throughout the experiment. If you do not make a choice within the allowed time, your choice is automatically recorded as a 'No'.

After everyone in your group has made their social distancing choice, the computer randomly chooses **one and only one participant to contract COVID-19 directly**. If a participant does not practice social distancing and is randomly picked by the computer then s/he gets infected for sure. In other words, **if you do not practice social distancing there is a 20% chance you get infected with COVID-19 directly**.

Social distancing gives you some level of protection against COVID-19. If you practice social distancing and you are the participant randomly chosen by the computer to be infected then the computer flips a fair coin. If the coin flip is Head you are infected with COVID-19. If the coin flip is Tail then you are not infected with COVID-19. In other words, **if you practice social distancing there is a 10% chance you get infected with COVID-19 directly**.

A participant who practices social distancing cannot pass COVID-19 to other participants and cannot be infected with COVID-19 by another participant.

On the other hand, an infected participant who does not practice social distancing may infect other participants through contagion. In particular, other participants who do not practice social distancing face a risk of getting infected because COVID-19 may spread through **interactions** between the participants.

The next page of the instructions explains how COVID-19 spreads through interactions between participants.


To continue to the next page, press the Next button below. To go to the previous page, press the Back button.

Part 1: Instructions (page 4/6)

A healthy participant who does not practice social distancing may get infected with COVID-19 through **contagion** by interacting with infected participants who do not practice social distancing. The probability a participant contracts COVID-19 through interaction with another participant is referred to as the **rate of contagiousness** of COVID-19.

Throughout Part 1 of the experiment the rate of contagiousness of COVID-19 is fixed at 65%.

Consider again the example diagram of interactions and, as an example, suppose that:



- You (participant M) do not practice social distancing.
- Participants in position E and Q practice social distancing, while participants in positions P and C do not.

In this example, suppose C is randomly picked by the computer to contract COVID-19 directly.


First, C has chosen not to practice social distancing so s/he gets infected for sure. Moreover, C can pass COVID-19 to other participants because s/he does not practice social distancing.

Second, E and Q cannot get infected with COVID-19 through contagion because they practice social distancing.

Next, P may get infected through contagion because s/he does not practice social distancing and interacts with C. This can happen with probability 65% – the rate of contagiousness.

Finally, you may also become infected by contagion through your interaction with P. Specifically, there is a 65% probability you might get infected through your interaction with P if s/he becomes infected. However, if P remains healthy you will also remain healthy for the duration of the game.

It follows that in this example if 1) you do not practice social distancing, 2) E and Q practice social distancing whereas P and C do not, and 3) C contracts COVID-19 directly, then you may become infected with probability 42.25%. The diagram below shows how this percentage is computed.



Note that you will not be informed of the choices of other participants at any point during the experiment.

The next page of the instructions explains how you earn points in Part 1 of the experiment.

To continue to the next page, press the Next button below. To go to the previous page, press the Back button.

Part 1: Instructions (page 5/6)

At the end of the game you earn the following points depending on your social distancing choice and infection status:


- **100 points:** if you did not practice social distancing and did not get infected;
- **65 points:** if you practiced social distancing and did not get infected (100 points for being healthy minus 35 points cost of social distancing);
- **0 points:** if you did not practice social distancing and got infected;
- **-35 points:** if you practiced social distancing and got infected (0 points for being infected minus 35 points cost of social distancing).

Note that **if you fail to submit** your social distancing choice then you will receive a **penalty of 200 points**.

The information about the points you can earn, the rate of contagiousness of COVID-19, the structure of interactions between participants and your position are always displayed on the screen when you make your social distancing choice.

Below you can see an example of how this part of the interface looks like. The diagram of interactions is on the right, while the textual information on the left reminds you of the rate of contagiousness and the possible outcomes.


Rate of contagiousness	65%
Points:	
- not practiced distancing, healthy	100
- practiced distancing, healthy	65
- not practiced distancing, infected	0
- practiced distancing, infected	-35



At the end of the game, you are reminded of the structure of interactions between participants, your position within that structure, and your social distancing choice. You are also informed of your infection status and the number of points earned.

Below you can see an example of how this part of the interface looks like.

Your distancing choice	practiced distancing
Your infection status	healthy
Points earned	65



You have 15 seconds to review this information and the timer is always displayed at the top of the interface.

The next page of the instructions explains how points are converted into your earnings in \$.

To continue to the next page, press the Next button below. To go to the previous page, press the Back button.

Part 1: Instructions (page 6/6)

Part 1 of the experiment has **20 separate games** as described in these instructions. **The choice you make in one game has no effect on other games.**

The only variation between games is the random reassignment of the positions of all the participants (including you) in the diagram of interactions. The participants assigned to your group, the structure of interactions between positions, the probability of contracting COVID-19 directly, the rate of contagiousness of COVID-19 and the number of points earned depending on social distancing choice and infection status remain unchanged.

At the end of each game, you can review the history of your choices and outcomes for the last 5 games. The table below shows an example of how your history after 9 games might look like.

Your History

Game	Position	Distancing Choice	Infection Status	Points Earned
9	Q	practiced distancing	healthy	65
8	E	practiced distancing	healthy	65
7	P	not practiced distancing	healthy	100
6	M	practiced distancing	infected	-35
5	C	not practiced distancing	infected	0

It is important that you make a choice in every game. If you fail to make a choice for **3 consecutive games**, you will be **disqualified from the experiment**. In this case, you will not receive any payment for this experiment.

At the end of this experiment, the computer randomly picks **4 out of 20 games** to determine your earnings for Part 1 of the experiment.

The points are converted to \$ at a rate of 115 points per \$1.

Suppose that you earn 260 points in the 4 randomly drawn games. Then, your total earnings for Part 1 are \$2.26.

To continue to the short Quiz, press the Next button below. To go to the previous page of the instructions, press the Back button.

Part 2: Instructions

You have completed Part 1 of the experiment and will now proceed to Part 2.

Below are the instructions for Part 2 of the experiment. It is important that you read these instructions carefully.

This part of the experiment also has 20 games, and you are assigned to the same group of 5 people as in Part 1. The structure of interactions between participants, the probability of contracting COVID-19 directly and the rate of contagiousness of COVID-19 are the same as in Part 1.

The single difference from Part 1 is that in Part 2 of the experiment you will receive a fine of 15 points in any game in which you do not practice social distancing.

Hence, in Part 2 of the experiment the points you earn at the end of the game are:

- **85 points:** if you did not practice social distancing and did not get infected (100 points for being healthy minus 15 points fine);
- **65 points:** if you practised social distancing and did not get infected (100 points for being healthy minus 35 points cost of social distancing);
- **-15 points:** if you did not practice social distancing and got infected (0 points for being infected minus 15 points fine);
- **-35 points:** if you practised social distancing and got infected (0 points for being infected minus 35 points cost of social distancing).

Note that if you fail to submit your social distancing choice then you will receive a penalty of 200 points.

Your earnings for Part 2 are computed in the same way as in Part 1. At the end of this experiment, the computer randomly picks 4 out of 20 games to determine your earnings for Part 2 of the experiment.

As in Part 1, the points are converted to \$ at a rate of **115 points per \$1**.

Suppose that you earn 300 points in the 4 randomly drawn games. Then, your total earnings for Part 2 are \$2.61.

Before you can start Part 2 of the experiment, you must answer a **Quiz question on the instructions above**. If you fail to answer the question correctly, you will not be able to continue to Part 2 of the experiment.

To continue to the short Quiz, press the Next button below.

The experiment: Interface

Figure A6 shows the decision and results interface from the main experiment. We focus on game 5 of the Intervention part of the experiment, where the rate of contagion is set at 65%, the network is the star, and fine is the intervention. In this game, our example participant is assigned to position P. She decides to practice social distancing, and does not get infected.

A.5.2 Social Value Orientation (SVO) Slider Measure

This section contains instructions from and interface of the Social value Orientation (SVO) task which participants complete as part of the recruitment survey.

SVO: Instructions

You have answered the qualifying questions correctly and are now in the Bonus Task.

In this Bonus Task you will be making a series of decisions about allocating money between you and another anonymous Turker (hereafter, Other). All of your decisions will be completely confidential.

There is a total of 6 decisions to make which are independent of each other. For each decision, you are asked to pick the distribution of money between yourself and Other that you prefer, all values are stated in cents. After you have made your decision, select the resulting distribution of money by clicking on the button below your choice. As you will see, your choices will influence both the amount of money you receive as well as the amount of money the Other receives.

There are no right or wrong answers, this is all about personal preferences.

Every time 50 Turkers complete the task, we will randomly pick two of them and pay them for the Bonus Task as follows. We will pick one of the 6 decisions, and randomly implement the decision of one of the two chosen Turkers.

For example, suppose we randomly chose Turkers X and Y out of those 50 Turkers, and that we further randomly chose to implement decision 3 of Turker Y. Suppose that in decision 3 Turker Y allocated 150 cents to themselves, and 140 cents to Other. Therefore, Turkers X and Y will be paid 140 and 150 cents respectively.

SVO: Interface

Figure A7 presents part of the interface of the SVO task. We focus on one of the six decisions. In this example, a participant decided to allocate 288 cents to self and 188 cents to other.

A.5.3 Bomb Risk Elicitation Task (BRET)

This section contains instructions from and interface of the Bomb Risk Elicitation Task (BRET) which participants complete after the main experiment.

BRET: Instructions

Thank you very much for taking part in the experiment!

You are now in the Bonus Task in which you have an opportunity to earn an extra payment.

On the next page, you will see 100 boxes. As soon as you start the task by pressing the ‘**Start**’ button, one box is collected per second starting from the top left corner. Once collected, the box is marked by a tick symbol.

For each box collected, you earn 2 cents.

One of the 100 boxes contains a **bomb** that destroys all of your earnings. **You do not know where the bomb is located.** You only know that the bomb can be in any box with equal probability.

Your task is to choose when to stop collecting the boxes and open those you have collected. You stop collecting boxes by pressing ‘**Stop**’ at any time. After that you open the boxes you have collected by pressing the ‘**Open**’ button. Note that once you press ‘Stop’ you cannot restart collecting boxes.

A dollar or a bomb symbol will be shown on each of the boxes you have collected.

If the bomb symbol does not appear, that means that you have **not collected the box with the bomb**. In this case, you **earn the amount accumulated by the boxes you have collected**.

If the bomb symbol appears, that means that you **have collected the box with the bomb**. In this case, you **earn zero** for the Bonus Task.

To proceed to the Bonus Task, press the button below.

BRET: Interface

Figure A8 presents the interface of the BRET. Figure A8a shows an example of the interface after a participant has stopped collecting boxes. We can see that she opened 56 boxes. Figure A8b shows the interface after she opened the boxes. We can see that the bomb is among the 56 collected boxes. Consequently, in this example the participant earns zero for the BRET.


Part 2 Game 5 of 20: Your Social Distancing Decision

Time left to complete this page: 0:10

Rate of contagiousness 65%

Points:

- not practiced distancing, fine, healthy 85
- practiced distancing, healthy 65
- not practiced distancing, fine, infected -15
- practiced distancing, infected -35



Do you want to practice social distancing?

Yes

No


By pressing one of the buttons you will submit your decision.

(a) Decision interface

Part 2 Game 5 of 20: Your Results

Time left to complete this page: 0:05

Your distancing choice practiced distancing
 Your infection status healthy
 Points earned 65



To proceed to the next game, press the button below.

Next

Your History

Game	Position	Distancing Choice	Infection Status	Points Earned
5	P	practiced distancing	healthy	65
4	C	practiced distancing	healthy	65
3	P	not practiced distancing	infected	-15
2	C	not practiced distancing	healthy	85
1	C	practiced distancing	healthy	65

(b) Results interface

Figure A6: Game 5 of the Intervention: interface





Figure A7: Social Value Orientation (SVO) Slider Measure: Decision interface



(a) Decision interface

(b) Results interface

Figure A8: Bomb Risk Elicitation Task (BRET) interface

Appendix References

- Acemoglu, D., A. Malekian, and A. Ozdaglar (2016). Network security and contagion. *Journal of Economic Theory* 166, 536–585.
- Cerdeiro, D. A., M. Dziubiński, and S. Goyal (2017). Individual security, contagion, and network design. *Journal of Economic Theory* 170, 182–226.
- Chen, D. L., M. Schonger, and C. Wickens (2016). oTree – an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9, 88–97.
- Goodman, J. K., C. E. Cryder, and A. Cheema (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making* 26(3), 213–224.
- Howe, P. D., M. Mildenerger, J. R. Marlon, and A. Leiserowitz (2015). Geographic variation in opinions on climate change at state and local scales in the USA. *Nature Climate Change* 5(6), 596–603.
- Likert, R. (1932). A technique for the measurement of attitudes. *Archives of Psychology* 140(22), 5–55.
- McCright, A. M. and R. E. Dunlap (2011). The politicization of climate change and polarization in the American public’s views of global warming, 2001-2010. *The Sociological Quarterly* 52(2), 155–194.
- Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.
- U.S. Census Bureau (2011). Land Area. Retrieved from <https://www.census.gov/library/publications/2011/compendia/usa-counties-2011.html>.
- U.S. Census Bureau (2019). County Population by Characteristics: 2010-2019. Retrieved from <https://www.census.gov/data/datasets/time-series/demo/popest/2010s-counties-detail.html>.
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.