Department of Economics Discussion Paper Series

The Cost of Staying Open: Voluntary Social Distancing and Lockdowns in the US

Adam Brzezinski, Valentin Kecht, David Van Dijcke

Number 910
September, 2020

Department of Economics
Manor Road Building, Manor Road
Oxford, OX1 3UQ

ISSN 1471-0498

The Cost of Staying Open: Voluntary Social Distancing and Lockdowns in the US*

Adam Brzezinski[1], Valentin Kecht[2], David Van Dijcke*[1,3]

[1] Department of Economics, University of Michigan

[2] Department of Economics, Bocconi University

[3] Institute for New Economic Thinking at the Oxford Martin School, University of Oxford

2020-09-30

Abstract

In combating the spread of COVID-19, some governments have been reluctant to adopt lockdown policies due to their perceived economic costs. Such costs can, however, arise even in the absence of restrictive policies, if individuals' independent reaction to the virus slows down the economy. This paper finds that imposing lockdowns leads to lower overall costs to the economy than staying open. We combine detailed location trace data from 40 million mobile devices with difference-in-differences estimations and a modification of the epidemiological SIR model that allows for societal and political response to the virus. In that way, we show that voluntary reaction incurs substantial economic costs, while the additional economic costs arising from lockdown policies are small compared to their large benefits in terms of reduced medical costs. Our results hold for practically all realistic estimates of lockdown efficiency and voluntary response strength. We quantify the counterfactual costs of voluntary social distancing for various US states that implemented lockdowns. For the US average, we estimate that lockdowns reduce the costs of the pandemic by 0.8% of annual GDP per capita, compared to purely voluntary responses.

Keywords: COVID-19, difference-in-differences, SIR model, social distancing, lockdown, big data.

JEL-Classification: I12, I18, H12, D04, C33, H51.

*We thank SafeGraph for providing their data and support. We also thank Austin L. Wright, Rick Van der Ploeg, Nicola Gennaioli, R. Maria del Rio-Chanona, Carlo Altomonte, Sebastian Hager, Gustavo Mellior, Xiaowen Lei, Pietro Buri, Doyne Farmer and the INET Complexity Economics Group for support and helpful discussions. Correspondence to: David Van Dijcke, 611 Tappan Ave, Ann Arbor, MI 48109, United States. Email: `dvdijcke@umich.edu`. Declarations of interest: Van Dijcke acknowledges support from the Belgian American Educational Foundation.

1 Introduction

There is an emerging consensus that social distancing is effective at diminishing the spread of COVID-19 by reducing interpersonal transmission (Anderson et al., 2020; Bai et al., 2020; Fowler et al., 2020; Viner et al., 2020). Yet, the way in which political leaders aim to flatten the infection curve differs vastly across countries. Many governments successfully decreased contagion by mandating lockdown policies (Hsiang et al., 2020), while others relied on the voluntary response of their population, under the argument that lockdown policies may be associated with significant economic costs. It is therefore a major policy concern to understand how high levels of social distancing can be reached (Briscese et al., 2020) at minimal economic cost.

In this paper, we leverage high-resolution location trace data across US counties and combine them with microeconometric methods to disentangle the voluntary and mandatory social distancing response during the pandemic. For the purpose of calculating the voluntary and lockdown responses, we concentrate on the time period from the start of the pandemic (January 1, 2020) until the last day before the first easing of a lockdown takes place (April 23, 2020). Using a controlled SIR model (Gros et al., 2020), we translate these social distancing response estimates into assessments of medical and economic costs associated with COVID. We thus arrive at two sets of cost estimates: first, those that would arise in the absence of a lockdown; and second, those that would arise under a state-wide lockdown that is not eased before the end of the pandemic.

We find that the percentage of people who stay at home voluntarily increases on average by 5.1% in response to the occurrence of the first local cases of COVID. The effect following the implementation of a state-wide shelter-in-place mandate is of a similar size. Combining these estimates with the SIR model based on Gros et al. (2020), we estimate the pandemic to cost the US around 16.1% of annual GDP per capita under a non-lockdown scenario compared to 15.2% if a lockdown is imposed. When taking into account the statistical value of life, these values increase to 20.9% and 18.1%, respectively.

The main implications of our paper are threefold. First, independent of government policies, movement decreases markedly in response to a local outbreak of the virus. Thus, economic costs are inevitable even in the absence of lockdown policies. Second, voluntary and lockdown responses are of similar magnitude. Governments that rely on their citizens to respond voluntarily need to take into account that the associated change in social distancing behavior will be significantly smaller without a lockdown policy in place. Third, our augmented version of the SIR model indicates that not imposing a lockdown barely improves economic performance, while it drastically increases medical costs—both in terms of lives lost and in terms of hospitalization costs.

Our paper contributes to several strands of literature. First, we support the finding that the additional economic costs of lockdown measures compared to non-lockdown scenarios are relatively small, while their medical benefits are large. Low economic costs have also been estimated by Baek et al. (2020), who find that only one fourth of the unemployment claims in the early phases of the crisis can be attributed to shelter-in-place policies. Hence, our results stand in opposition to the view that voluntary responses are a better strategy than lockdown policies as proposed in Krueger et al. (2020), among others. An often cited example for the efficacy of relying on voluntary social distancing is Sweden. Current data on GDP growth and forecasts suggest, however, that the Swedish economy may contract in a similar magnitude as in other European countries or in the US (EC, 2020).[1] Our results also contrast with those of Chudik et al. (2020), who find that voluntary responses have a small impact on the spread of the pandemic compared to lockdowns, but lead to large costs in terms of employment losses. Second, this paper speaks to a growing literature that jointly models the dynamics of the economic and health threats arising from COVID (Acemoglu et al., 2020; Allcott et al., 2020a; Atkeson, 2020; Barro et al., 2020; Coibion et al., 2020; Eichenbaum et al., 2020; Jones et al., 2020; Kaplan et al., 2020). Third, a number of papers has used cellphone data to study movement patterns during the COVID pandemic. They show that patterns in social distancing depend on partisanship (Allcott et al., 2020b; Grossman et al., 2020; Kushner Gadarian et al., 2020; Painter and Qiu, 2020), political polarization (Cornelson and Miloucheva, 2020), poverty and economic dislocation (Wright et al., 2020), belief in science (Brzezinski et al., 2020b), risk perception (Allcott et al., 2020b; Barrios and Hochberg, 2020; Engle et al., 2020) and civic capital (Barrios et al., 2020; Durante et al., 2020). Our central contribution is to combine data from mobile devices with a controlled SIR model to estimate the medical and economic costs under different policy scenarios.

2 Data

We construct a dataset at the county-day level spanning the period between January 1, 2020 and April 23, 2020. The data contains measures of social distancing and the lockdown policies enacted by counties.We only consider data up to April 23 because Georgia was the first state to partially ease its lockdown on April 24, but we are interested in the effect of imposing lockdowns, not easing them.

Social distancing. Our outcome variables are based on location trace data from SafeGraph, obtained by tracking GPS pings from up to 40 million devices across the United States. Our analysis is based on a new data product SafeGraph developed to allow the tracking of social distancing in response to the virus, *Social Distancing Metrics*. The data comes from an underlying panel of up to 40 million mobile devices with home addresses in all 200,000+ census

[1]Sweden only saw an estimated 4 p.p. smaller decrease in aggregate consumer spending than neighboring Denmark, which did impose a lockdown and experienced a drop in spending of around 29 percent (Andersen, 2020).

block groups (CBG) across the United States. It has been aggregated in an exhaustive 6-step process designed to guarantee reliability, granularity and anonymity.[2] As our main outcome variable we calculate the *percentage of devices that stayed home all day* by summing, at the county level, the number of the devices that exclusively emitted GPS pings from their home location over the course of a day, and dividing it by the total number of devices observed in each county on that same day. Variations of this variable have previously been exploited in Brzezinski et al. (2020b), Cotti et al. (2020), Holtz et al. (2020), Lasry et al. (2020), and Simonov et al. (2020). A device's home is determined as its common nighttime location over the course of 6 weeks, down to a Geohash-7 (153m x 153m) granularity. Note that, due to the limited frequency of GPS pings, our outcome variable is likely to be downward biased for more densely populated areas, where short trips outside of the home might not be registered. We take this into account by controlling for county-specific factors in our main regression specifications.

The panel exhibits limited bias along several dimensions. When it comes to geographic bias, the absolute difference between the panel's density and the true population density as measured by the US census never exceeds 1% at the county level. The correlation between both variables is 0.97. In addition, the panel also has a low degree of demographic sampling bias. Although device-level demographics are not collected for privacy reasons, average demographic patterns can be studied using panel-weighted, CBG-level Census data. Here again, the frequency of salient demographic and income groups in the panel closely tracks the same frequency in the Census. This supports the representativity of the sample with respect to the whole population, because cellphone use in the US is common across a wide range of demographic groups.[3]

Lockdown Policies. Data on the lockdown measures implemented to combat the spread of COVID have been retrieved from several sources: the National Association of Counties (NACO)[4]; the National Governors' Association (NGA)[5]; Education Week[6] ; and The New York Times.[7] The data incorporates information on shelter-in-place policies as well as school and business closures. Business closure orders require all non-essential businesses to shut down, while shelter-in-place policies call for all citizens to stay at home. Essential needs, such as grocery shopping, exercise and medical emergencies, are the only exceptions to shelter-in-place orders. Nonetheless, such orders were generally not strictly enforced in the US, allowing for varying degrees of compliance. People working in essential businesses were still allowed to go to work. Additionally, all states implemented school closures. We codify the date of a policy's

[2]CBGs with less than 5 devices are excluded from the sample. To further enhance privacy preservation, SafeGraph collects data not directly from cellphones, but only from secondary sources. Thus, the data products and maps derived from the mobility patterns are aggregate results that do not allow the re-identification of individuals.

[3]See, for example, `https://www.pewresearch.org/internet/fact-sheet/mobile/`.

[4]For details , see `https://ce.naco.org/?dset=COVID-19&ind=State%20Declaration%20Types`.

[5]`https://www.nga.org/coronavirus/#states`

[6]`https://www.edweek.org/ew/section/multimedia/map-coronavirus-and-school-closures.html?override=web`

[7]`https://www.nytimes.com/interactive/2020/us/states-reopen-map-coronavirus.html`.

implementation as the official date it went into effect if it did so before 12pm, or one day later if it did so after 12pm.

COVID Statistics. County-level statistics on COVID cases and deaths in the United States are retrieved from the COVID-19 data portal provided by the New York Times.[8] We also collected data on the state-day effective reproduction number R_t, which is calculated according to a modified version of the model described in Bettencourt and Ribeiro (2008).[9]

Unemployment. Data on weekly unemployment claims and rates by state was retrieved from the United States Department of Labor, Employment and Training Administration.[10]

3 Quantifying Voluntary and Lockdown Responses

In order to quantify and compare voluntary and lockdown-induced social distancing responses, we pursue a set of difference-in-differences (DiD) strategies. First, we estimate the state-specific voluntary social distancing response by analyzing how movement patterns vary around the appearance date of the first local COVID cases controlling for whether a lockdown policy is in place. Second, we isolate the social distancing response to state-wide shelter-in-place policies, where we control for the trajectory of voluntary behavior caused by the spread of the virus. We also explicitly control for spillovers caused by policies of connected states and counties using the approach of Holtz et al. (2020). Following Goodman-Bacon and Marcus (2020), we complement our baseline DiD estimates with an event-study approach to check the robustness of our results.

The inherent differences in our two types of events of interest, i.e. the appearance of county-level local cases and the implementation of state-wide lockdown policies, as well as their staggered nature, implies that we need to pursue two separate DiD strategies to obtain estimates of the voluntary and lockdown responses. To estimate the voluntary response, we can leverage the fact that COVID cases appear at different periods across counties in a given US state, allowing us to estimate the voluntary response separately for each state. We cannot pursue the same strategy for estimating the effect of state-level lockdowns, since these are, by definition, implemented simultaneously in all counties of a state. Thus, we implement separate DiD strategies for the estimation of voluntary and lockdown responses, for which we employ different sets of control variables, as will be described in more detail further below.[11]

Figure 1 shows the evolution in the growth rate in the percentage of people who stay at home, alongside the share of states that have experienced their first COVID cases and deaths as well

[8] https://developer.nytimes.com/covid

[9] https://rt.live/

[10] https://oui.doleta.gov/unemploy/claims.asp

[11] For instance, to capture increases in social distancing that are really voluntary, we need to account for whether county- or state-level policies have been implemented. On the other hand, when estimating the lockdown response, we require additional controls that account for the voluntary increase in social distancing, since lockdowns are enacted for the most part after the appearance of first cases.

as the share that implemented state-wide shelter-in-place policies. In most states, the first case occurred during the first half of March, while first deaths and lockdown policies were lagging behind. Social distancing gained ground from March 8 onward, the date of the first death in the US, until the end of the month and remained stable afterwards. To disentangle voluntary and lockdown-induced social distancing, we adopt a DiD approach where the counterfactual comprises counties that have not (yet) been exposed to a first COVID case or a shelter-in-place order, respectively. Hence, our identification strategy hinges on a parallel trend assumption in the outcome variable relative to the day of the treatment. In this case, the approach yields estimates that can be interpreted as causal average treatment effects.

— [insert Figure 1 here] —

Voluntary response. To estimate the voluntary response, we pursue the following DiD approach:

$$pct_{i,t} = \alpha_i + \delta_t + \zeta Lock_{j,t} + \beta First_{i,t} + \gamma\left(Lock_{j,t} \times First_{i,t}\right) + \mathbf{\Omega y_{i,t}} + \epsilon_{i,t}, \tag{1}$$

where $pct_{i,t}$ denotes the percent of devices that stay home all day in county i and on day t. α_i and δ_t refer to county and day fixed effects. $Lock_{j,t}$ takes value 1 if the shelter-in-place policy has been implemented at or before time t in state j. Similarly, $First_{i,t}$ is a dummy variable that is equal to 1 if the first case has already occurred, and 0 otherwise. Lastly, the vector $\mathbf{y_{i,t}}$ includes controls for state-wide cases and deaths, state-wide school closures, county-wide business closures and state- and county-wide emergency declarations.[12]

Since we control for the interaction of $First_{i,t}$ and the policy variables, the coefficient β captures the social distancing response when no policy is enacted, in other words, the voluntary response. For each state, we use a balanced sample of its counties that goes up to two weeks after the state's lockdown. For the US-wide stacked DiD, we use the full balanced sample. Further robustness checks suggest that our estimates are not strongly affected by bias due to either time or unit heterogeneity (see section D).

Figure 2 shows the results from regression Equation 1, estimated separately for each state. Blue (red) lines correspond to Democratic (Republican) states as per the 2016 presidential election, dashed lines indicate parameters that are insignificant, and the horizontal line plots the point estimate for the US as a whole. The main US-wide regression results for the estimation are reported in Table 1, Column 1.[13] Our state-specific estimates range from showing no significant effect to an effect of 6.2 percentage points, with a US-wide estimate of 1.28 percentage points.

— [insert Figure 2 here] —

[12]We do not control for state-wide business closures as these generally coincide with shelter-in-place orders. Note also that our specification assumes that the effect of a county-level policy is constant over time, and that the effect of a state-level lockdown only depends on whether an outbreak has already taken place. This seems a reasonable approximation: the estimates from our event-study approach for the state-level lockdown show that the coefficients are relatively stable across the first days after a lockdown (see Figure D.1).

[13]Full results in Appendix Table A.3.

Lockdown response. We pursue a similar stacked DiD approach to estimate the lockdown response:

$$pct_{i,t} = \alpha_i + \delta_t + dayssince_{i,t} + \theta Lock_{j,t} + \zeta D^{GEO}_{i,t} + \mathbf{\Psi y_{i,t}} + u_{i,t}, \quad (2)$$

where all variables are defined as above, with the exception that we include the additional days-since-first-case fixed effects $dayssince_{i,t}$. Additionally, we include $D^{GEO}_{it} = \sum_k w_{i,k} \times D_{k,t}$, which is a vector containing a geographic-adjacency-weighted average of all other counties' policies.[14] The weights $w_{i,k}$ measure the normalized posterior probability of travel to county k for someone who lives in county i, and D_{it} is a vector of dummies for county- and state-level shelter-in-place orders, business closures and school closures.[15] This term allows us to control for spillovers of policies implemented in geographic alter counties. We incorporate this variable to control for the trajectory of voluntary responses coming from the progression of the disease, before any lockdown has occurred. As previously shown, individuals in affected areas exhibit voluntary social distancing over and above any country-wide behavioral trends. Thus, omitting $dayssince_{i,t}$ would result in an upward bias in the estimate for θ. The event-study approach in section D confirms that there is no significant pre-trend in our analysis.

As before, we balance the US-wide sample by county and date. We exclude counties from the analysis that have implemented county-level policies before the introduction of state-level policies to reduce the threat of anticipation effects. In contrast to the voluntary response, we are unable to obtain state-specific lockdown results, since, by definition, state-level lockdowns are enacted in all counties of a state simultaneously.

Column 2 in Table 1 shows the estimate for θ from Equation 2 for the US. We find that the lockdown increases the percentage of people who stay at home by 1.32 percentage points after the policy is implemented. The full regression results presented in Appendix Table A.3 confirm that the control variables have the expected signs or are insignificant.

— [insert Figure 1 here] —

4 The Costs of Voluntary and Imposed Social Distancing

In this section, we provide calculations of the medical and economic costs incurred by voluntary and policy-induced social distancing. We do so by combining the estimates for the two forms of social distancing with the 'controlled SIR model', a modified version of the workhorse SIR

[14]Note that for the estimation of the voluntary response, we cannot include similar geographic-adjacency-weighted average measures of all other counties' first-case appearance. The reason is that the appearance of COVID cases in one county is highly collinear with the appearance of such cases in neighbouring counties. We also do not include geographic-adjacency-weighted average measures of other counties' policies for the estimation of voluntary responses, since the first cases appear generally much earlier than the implementation of policy measures.

[15]See Holtz et al. (2020, S2.3) for further information on the geographic adjacency matrix.

model, that allows the disease reproduction number to endogenously depend on the political and societal response intensity to the virus (Gros et al., 2020).

As the model admits an analytical solution, the total response intensity can be estimated from the data. We interpret this response intensity as the combination of voluntary and lockdown-induced social distancing, which we can then decompose based on our counterfactual econometric estimates from the preceding section. Armed with this decomposition, we re-simulate the model for several US states so as to quantify the costs associated with the observed total and estimated voluntary reductions in movements. We thus combine econometric estimates based on detailed micro-data with a modeling approach firmly embedded in epidemiology to conduct a counterfactual policy simulation of the costs of staying open. The theoretical model we use is parsimonious and transparent, a feature we believe makes it preferable to heavily parametrized economic models, as it reduces the uncertainty associated with estimating a large number of parameters, and facilitates discussion of the underlying model assumptions. Moreover, while agent heterogeneity is likely to be an important factor in the spread of the virus, our model assumes that governments are either not capable of or not willing to target lockdowns towards specific population groups.

Identification of the model parameters requires that a peak of the pandemic is clearly identifiable and has already passed. As a result, we only include states in the analysis where a clear peak is identifiable before the easing of a state-level lockdown. The calculation of the estimates of the costs of lockdowns requires the further assumption that the government lockdowns would have been kept in place until case levels are low enough to be contained by more targeted measures. This allows us to calculate two types of counterfactual medical cost estimates for each state: those that would arise under no lockdown, and those that would arise under a lockdown that was kept in place until the end of the pandemic. Since we only look at states that saw a clear initial peak during lockdown, the parameters estimated on the basis of our specification are well-identified even if these states see resurgences in cases after the lockdown is lifted.

Controlled SIR model. At any time $t > 0$, a population is composed of three types of individuals: susceptible, infected and recovered; with associated quantities $S = S(t)$, $I = I(t)$ and $R = R(t)$.[16] Population size is constant and normalized such that $S + I + R = 1$. The following set of ordinary differential equations describes an isolated epidemic outbreak:

$$\tau\frac{dS}{dt} = -gSI, \qquad \tau\frac{dI}{dt} = gSI - I, \qquad \tau\frac{dR}{dt} = I, \tag{3}$$

where g is the reproduction number and $\tau > 0$ is the time scale. According to the first equation, the growth rate of S decreases in g and in the number of contacts between susceptible and infected individuals SI. The change in infections in the second equation equals this reduction

[16]This paragraph closely follows the exposition in Gros et al. (2020).

in susceptible individuals gSI. Finally, the third equation posits that all infected individuals recover or die. Note that, in the case of COVID, this invites an interpretation of the model's time scale as the duration of the disease's infectiousness, which is estimated to be 2 weeks (WHO, 2020).

In contrast to the standard SIR model, societal and political response to the spread of the virus is assumed to push g below its intrinsic number $g_0 > 0$, which depends on the virological characteristics of the disease. This notion can be operationalized by

$$g = \frac{g_0}{1 + \alpha X}, \qquad X = 1 - S. \tag{4}$$

The parameter $\alpha \geq 0$ describes the overall reaction strength of the population to the spread of the virus, captured by the cumulative number of cases X. Equation 4 implies that social distancing has decreasing returns in reducing the reproduction number. This is in line with epidemiological evidence that the distribution of individual infectiousness is highly right-skewed, with a small number of 'superspreaders' infecting many individuals at large social events (Lloyd-Smith et al., 2005).[17] Given g_0, the parameter α determines the level to which the reproduction number converges as the disease progresses, allowing us to interpret α as the degree of social distancing. As such, a business-as-usual scenario amounts to $\alpha = 0$ and $g = g_0$. Using our estimates from section 3, we can decompose $\alpha = \alpha_v + \alpha_l$, where $\alpha_v \geq 0$ captures the degree of voluntary social distancing and $\alpha_l \geq 0$ the additional response induced by lockdown policies. Note that our DiD design allows for such an additive interpretation of both responses, as the estimated lockdown response quantifies the additional social distancing compared to the counterfactual of no lockdown—i.e. of only voluntary social distancing.[18]

Under the assumption that individuals and policymakers react to X rather than I—what Gros et al. (2020) call long-term and short-term control, respectively—, the controlled SIR model described by Equations 3 and 4 admits an analytical solution. Since $X = 1 - S$ by definition, this implies that the reaction depends directly on the size of the susceptible population. From a theoretical point of view, this is more intuitive than to assume that the government only reacts to the current rate of infections, regardless of the progression of the virus and the amount of people at risk of being infected. Moreover, the assumption that the government reacts to X, and the implied relationship between X and the reproduction rate, holds true empirically (see

[17]This dispersion in individual infectiousness is commonly denoted as k. Early estimates suggest the k for COVID is similar (Riou and Althaus, 2020) or even lower (Endo et al., 2020) than it was for SARS and MERS, two other corona viruses that exhibited strong clustering.

[18]For both the voluntary and the lockdown response estimations, the absence of significant pre-trends in the event study estimates presented in D.1 provides evidence in favor of this interpretation. With the use of suitable control variables, we ensure that our estimate of α_v only captures the voluntary response and of α_l only the lockdown response (see section 3). For this purpose, in the case of the voluntary response we include controls for county- and state-level policies. For the estimation of lockdown responses, in turn, we use a rich set of controls for the progression of the disease, and thus account for the voluntary response pre-lockdown. Therefore, the estimated lockdown response quantifies the additional social distancing compared to a no-lockdown scenario.

the discussion of Figure C.1 in the Appendix). The model then permits the estimation of α and g_0 from the data using the derived 'XI representation'

$$I = \frac{\alpha + g_0}{g_0} X + \frac{1+\alpha}{g_0} \log(1-X), \tag{5}$$

which is independent of τ. This relation holds remarkably well in the data across highly dissimilar regions, validating the model and supporting its robustness (Gros et al., 2020, Fig 1). Figure B.1 in Appendix shows a similarly close model fit for 23 different US states. What is more, the model is robust to differences in testing as long as I and X are mismeasured by the same ratio.

PARAMETRIZATION. To estimate the model parameters, we take I_t and X_t from our daily data as the number of newly confirmed COVID cases and the cumulative number of confirmed cases.[19] We smooth the number of cases with a symmetric moving average of length 5, as in Gros et al. (2020), and normalize by dividing I_t and X_t by the total population size of the geographical unit under consideration. To parametrize Equation 5, we use weighted least squares to estimate, for state j at time t,

$$I_{jt} = \beta_1 X_{jt} + \beta_2 \log(1 - X_{jt}) + \epsilon_{jt}, \tag{6}$$

where the weights equal I_{jt} to account for the fact that the case data is crowded at low levels (Gros et al., 2020, p.9). For each state, we use a sample that goes until the lockdown is lifted. One can then obtain $\hat{g}_0$ and $\hat{\alpha}$ using $\beta_1 = (\alpha + g_0)/g_0$ and $\beta_2 = (1+\alpha)/g_0$. The estimated parameters are reported for selected US states in Table A.1. Based on $\hat{\alpha}$, we obtain estimates for the voluntary response $\hat{\alpha}_v$ by using the estimates from Section 3 and the relationship $\alpha = \alpha_v + \alpha_l$. In particular, we use the state-specific voluntary response estimates (V) shown in Figure 2 and the US-wide lockdown response (L) from Table 1 to estimate the relative size of $\hat{\alpha}_v$ by scaling α by the ratio

$$\frac{V}{V+L}, \tag{7}$$

for each state.[20]

[19]Note that this is a valid measure of I as long as the infectiousness period does not change in the sample. This seems reasonable since no probably effective drug treatment had yet been put to wide use during our sample period.

[20]Note that our decomposition of α constitutes a conservative way of estimating the fraction of the total measured response made up by the voluntary response. The reason for this is that there are likely other 'weaker' forms of voluntary social distancing that our estimate does not take into account, which are, for example, captured in the country-wide time trend (Farboodi et al., 2020). Our definition of voluntary social distancing is 'strict' in that we only consider people's response to the local spread of the virus, which is a feature we should always expect to be associated with a pandemic outbreak. The country-wide time trend, on the other hand, likely captures the response to more contingent features of the outbreak, such as national emergency declarations. Our 'strict' definition is thus conservative in the sense that it underestimates what percentage of the total response is due to voluntary social distancing, and hence underestimates the economic costs associated with voluntary distancing.

Because we are interested in decomposing the effect of lockdown, we only consider states that implemented such a policy in our sample. Additionally, we drop all states where daily case counts did not clearly reach a first peak before the lockdown was eased.

With these parameter estimates in hand, one can simulate the discrete-time version of the model. Denoting the discrete-time equivalents of g and g_0 by ρ and ρ_0, respectively, it holds that $\hat{\rho}_0 = \exp(\hat{g}_0)$, and the discretized model can be written as

$$I_{t+1} = \hat{\rho}_t \, I_t \, (1 - X_t), \qquad X_t = \sum_{k=0}^{\infty} I_{t-k}, \tag{8}$$

where ρ_t is described by the discrete version of Equation 4.[21] Note that the estimated baseline reproduction numbers ρ_0 in Table A.1 are closely in line with those from the epidemiological literature (Li et al., 2020; Liu et al., 2020).

Cost estimates. As shown in Equations 3 and 4, the framework explicitly models the positive feedback loop between infections and the reproduction number. Political and societal factors can lead to variation in g and therefore influence the costs associated with the outbreak. We use the findings from section 3 to estimate the costs related to voluntary and policy measures. We distinguish four types of costs, expressed in terms of GDP per capita: (1) the production loss due to infected workers going on sick leave, (2) the medical expenses associated with infections, (3) the loss of human lives, and (4) the costs of social distancing (Gros et al., 2020, p.12).

Since the first three types are related to health and medical costs, they are directly proportional to X_{tot}, the cumulative number of cases at the end of the pandemic.[22] Hence, these costs can be approximated as $C^{\text{medical}} = kX_{\text{tot}}$, where Gros et al. (2020) suggest a value of $k \approx 0.305$, or $\tilde{k} \approx 0.14$ if the value of human life is not considered. These estimates, based on data from the Diamond Princess cruise ship, include appraisals of the costs induced by the lost work time of sick people, as well as conservative estimates of hospitalization costs in Europe, which are based on early estimates of the hospitalization rate of COVID (CDC, 2020).[23] Under a laissez-faire scenario, $X_{\text{tot}} = 0.94$[24] and the upper bound for the medical cost estimate is roughly 29% (or 13% if not accounting for the value of life) of annual GDP per capita.

The fourth cost factor is a function of voluntary and mandatory social distancing. Denote the function that maps social distancing into economic costs by $f(\alpha_v, \alpha_l, S)$, which depends positively on voluntary social distancing α_v and lockdown-induced social distancing α_v, and

[21] We initialize $I_0 = 2 \times 10^{-5}$ and end the simulation at $I_t = 10^{-5}$, after which it is assumed that new cases are fully reduced to zero by test and trace strategies (Gros et al., 2020).

[22] We assume that $X_{\text{tot}} \approx X$ at time $t = T$, with T the final time period.

[23] Note that these are truly conservative estimates insofar as hospitalization costs in the US are generally higher than in Europe.

[24] We stress that this value is not to be confused with herd immunity, the point where transmission of the disease starts to decline. In fact, consistent with the literature (e.g. Kwok et al., 2020), plugging our estimates in the standard formula $1 - 1/\rho_0$ (Fine et al., 2011) yields an estimated herd immunity point of around 70% for the US as a whole.

negatively on the size of the susceptible population S. By Equation 4, this cost can be captured by the decrease in ρ/ρ_0:

$$C^{\text{econ}} = f(\alpha_v, \alpha_l, S) = \textstyle\sum_{I>I_{\min}} \; m\left[1 - \dfrac{\rho}{\rho_0}\right]\dfrac{2}{52}, \tag{9}$$

where Equation 9 sums over all points in time when the pandemic is in a large-scale containment stage, that is, when individual testing and tracing is not possible. The parameter m maps the social distancing efforts into their economic costs. We estimate m based on the observed relation between the weekly average of the state-level reproduction rate and the weekly insured unemployment rate across US states, as described in section C.[25] Our estimate of 0.28 is very close to the value of 0.25 calculated by Gros et al. (2020) relying on data from China.

Overall, the model reflects the trade-off faced by governments between medical and economic costs. Less restrictive policies lower the direct economic burden of the pandemic due to an increase in g at the expense of an increase in medical costs, which enters the estimated costs through a higher level of X_{tot}. Clearly, our estimates of the economic and medical costs of the pandemic are based on a highly stylized model. For instance, it does not take into account the possibility of a collapse of the health care system, which would lead to a discontinuous surge in the death toll ("The benefits and costs of flattening the curve for COVID-19"). Moreover, the model also abstracts from externalities arising from social distancing, such as increases in domestic violence or mental health issues. We address some of these issues by checking the robustness of our results for a wide range of different parameter estimates further below.

Another potential limitation is that increases in voluntary responses and in lockdown responses affect our cost estimates in the same way. In other words, we do not allow for the possibility that, on the margin, an increase in the voluntary response yields different medical and economic implications than an equivalent increase in the lockdown response of the population. However, this line of reasoning would require that the heterogeneity in responses differs systematically across demographic and income groups. The available evidence, in turn, suggests otherwise: For instance, people from lower income groups exhibit both a weaker voluntary (Brzezinski et al., 2020a) and a weaker lockdown response (Wright et al., 2020). As a result, similar heterogeneities for the voluntary and mandatory social distancing map into similar marginal costs, implying that the additive decomposition is justified at the aggregate level.[26]

Results. Figure 3 plots our results. It shows the total costs, in terms of GDP per capita, of the total containment response (T) and the counterfactual voluntary response (V) by US state. As explained above, the total response cost is estimated using the α measured from the data,

[25] We exclude 5 outlier states, though the estimated m is very similar (13.74) when we retain them.

[26] Moreover, this issue would only pose a problem for our main results if, on the margin, an increase in the lockdown response would yield systematically lower health benefits and higher economic costs than an equivalent increase in the voluntary response.

while the voluntary response is a counterfactual based on our estimates from section 3. Also shown are the total and voluntary costs for the GDP-weighted average of all included states (2019Q4, annualized). Both costs are broken down into economic (red) and medical (blue) costs. Note that the state-level costs should be interpreted as relative to the GDP per capita of the state in question, not of the US as a whole.

Panel a) in the figure depicts a stark result. For the overall average, we estimate COVID to generate costs of around 16.1% of annualized GDP per capita under a laissez-faire scenario compared to 15.2% if a lockdown is imposed. Taking into account the statistical value of a life, these estimates increase to 20.9% and 18.1% (see Figure B.3). For all states considered, reducing the measured containment response to levels consistent with voluntary social distancing would only marginally decrease the economic costs, while at the same time drastically increasing the medical costs. This finding is more or less pronounced depending on both the estimated total containment response α and the estimated voluntary response for each state from Figure 2. States with a combination of a high estimated containment efficiency and a high estimated voluntary response, such as Montana and Idaho, barely see a change in either medical or economic costs when changing to voluntary containment, as they remain close to the respective asymptotic minimum and maximum of both cost cost categories.[27] On the other hand, states with a combination of low estimated containment efficiency and low estimated voluntary response, such as New York and Massachusetts, see a stronger change in both cost categories when moving to voluntary response, with the increase in medical costs outweighing the decrease in economic costs.

The reason why the model predicts that voluntary social distancing will lead to higher overall costs becomes evident by considering Figure B.2 in the Appendix, which shows how total costs change with an increase in control strength α. The model implies a trade-off between medical and economic costs as α gets higher. While the former fall with a stronger reaction strength, the latter increase. Initially, for low values of α, an increase in its value rapidly generates higher economic costs, causing larger estimates of total costs. However, once α is sufficiently high, further increases in the control strength will generate smaller marginal increases in economic costs, which are easily offset by the medical benefits. This means that the economic costs are approximately concave in α. Following this logic, for a range of values for the control strength there are two strategies that yield the same costs: either a weak response that implies small economic and large medical costs, or a strong response that generates small medical and larger economic costs. For all the states in our sample, we estimate that the strong responses result in lower overall costs.

Robustness. The previous section has concluded that lockdown strategies yield lower economic costs than strategies that rely on voluntary responses. We corroborate this finding by performing

[27] Note that the Figure actually shows a slight decrease in total costs for ID, MT and NV, which is due to numerical inaccuracies and discontinuities in the cost function.

a series of robustness checks, showing that our conclusion holds for a large set of different model parametrizations and estimation approaches. These robustness checks are described in detail in section D. First, we consider the robustness of our baseline DiD results by pursuing a staggered event-study approach. As pointed out in Abraham and Sun (2018) and Goodman-Bacon (2018), heterogeneous treatment effects over time may induce bias in the 'stacked' DiD approach in section 3. Yet, the results presented in section D in the Appendix indicate that such heterogeneity is rather limited and there is no significant pre-trend prior to the treatments of interest. In addition, the estimated coefficients are very close to our baseline results.

Second, we show that our main conclusion would still hold for estimates that differ considerably from our baseline, by considering the sufficient conditions under which lockdowns are the most cost-effective strategy. For this purpose, we take the model calibration of our SIR model and show the range of alternative estimates of the combined response strength and of the relative magnitude of the voluntary response for which our conclusions would still hold.

Third, we assess the robustness of our results to different estimates of the cost parameter m. Importantly, higher estimates for m increase the economic costs for any given level of social distancing. As higher economic costs render no-lockdown strategies more favorable, differences in m may modify our main conclusion. However, Figure D.2 supports the robustness of our results with respect to variations in m: starting from a baseline value of 0.28, a small increase in the parameter would lead the first two states to prefer a no-lockdown strategy, while even under unreasonably high elasticities, it would still be efficient for most states to impose a lockdown. Lastly, we emphasize that our results are not merely artifacts of the model. As we show in Figure D.3, different combinations of α and m can indeed generate lower overall costs under the voluntary scenario. However, our estimates paired with any realistic parametrization of the C-SIR model support the view that lockdowns have been the efficient strategy during the COVID pandemic.

— [insert Figure 3 here] —

5 Conclusion

Lockdown policies are an important tool for policymakers to combat the spread of COVID-19. Nonetheless, several governments across the world have been reluctant to adopt such policies, fearing that their economic costs outweigh their medical benefits. In contrast to this view, this paper argues that the costs of staying open outweigh the benefits. This is because individuals engage in substantial voluntary social distancing even in the absence of lockdowns, once the virus takes hold in their area. Hence, substantial economic costs are unavoidable, even when not locking down. At the same time, we show that lockdowns lead to a significant uptake in social distancing over and above any voluntary response. While large economic costs materialize in any event, such additional social distancing still plays an important role in further reducing medical costs by flattening the curve. Indeed, for our estimates of the voluntary and mandated social distancing responses, all US states that imposed a lockdown would have incurred larger overall costs had they stayed open. This suggests that the observed reluctance of some governments to lock down was unwarranted insofar as it was guided by a concern over the economic costs of such a policy.

Our study is subject to several caveats. First, the cost estimates presented are simulated for the trajectory of a large-scale pandemic for which test and trace strategies are not feasible. This implies that our results may not apply to other diseases and should not be interpreted as a justification for lockdowns in all contexts. Second, note that we only allow for variations in the strength of the response—and not, for example, in its speed—while assuming that containment remains operative until the number of infectious people drops below a certain threshold. Third, we do not specifically take into account possible externalities on outcomes such as mental health (Brodeur et al., 2020; Rossi et al., 2020; Twenge and Joiner, 2020) or inequality (Adams-Prassl et al., 2020; Alon et al., 2020; Bell et al., 2020; Galasso, 2020). Fourth, while we hope our results can guide the response to possible future waves of COVID-19, such extrapolation should be done with caution and with consideration of the fraction of the population infected in the first wave.[28] Finally, our model does not preclude the possibility that well-targeted partial lockdowns that take into account regional and demographic heterogeneity can achieve similar reductions in medical costs, at lower economic costs.[29] Our central contribution is to show that, even in the extreme case of an across-the-board lockdown, such a measure in all likelihood fares better than a laissez-fare approach.

[28]Indeed, non-linear dynamic systems such as the C-SIR model are highly sensitive to initial conditions.

[29]Due to the skewed distribution of virus transmission (Endo et al., 2020; Riou and Althaus, 2020; Sneppen and Simonsen, 2020) and the importance of 'superspreader' events, local and quick measures may indeed be best suited for mitigating the spread of COVID (see e.g. Bonardi et al., 2020 or here.

Figure 1: Contagion, Lockdown Policies and Social Distancing , Feb-Apr 2020

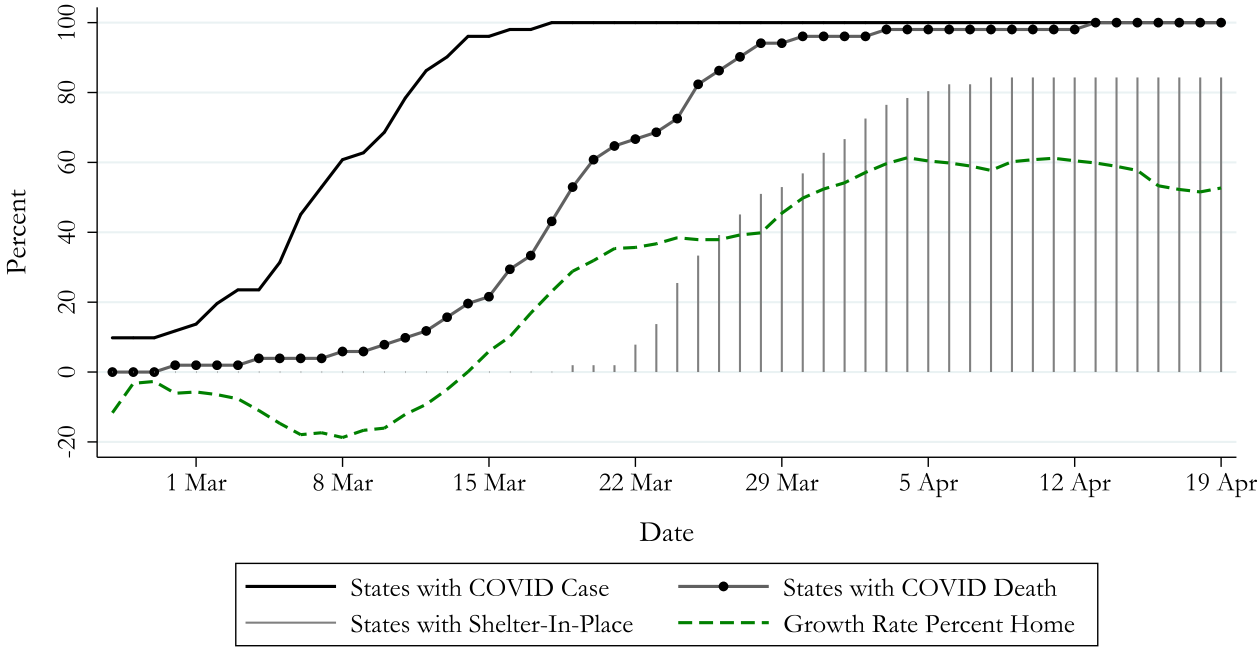

Note: The black solid (connected) line shows the percentage of states with at least one confirmed case (death). The spikes indicate the percentage of states that have shelter-in-place orders. The green dashed line depicts the growth rate of the median percentage of devices that stayed home over all counties, smoothed with a moving-average of length 7 to eliminate weekly patterns. Baseline period is February 2020.

Figure 2: Voluntary Response, by State

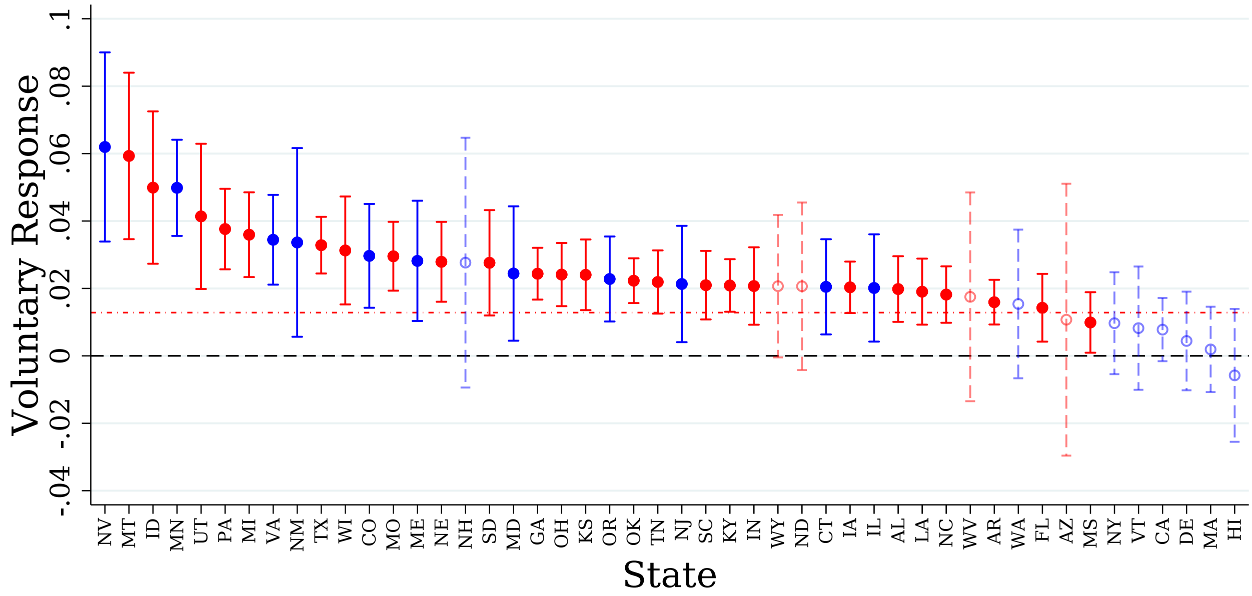

Note: The figure plots the cumulative change in the percentage of devices that stay completely at home following the first case in a county (parameter β from Equation 1). Parameter estimates and confidence intervals come from separate regressions for each state, while the overall effect (horizontal line) comes from a country-wide regression. Blue (red) lines correspond to Democratic (Republican) states as per the 2016 presidential election; dashed lines indicate parameters that are insignificant at $p < 0.05$. The sample is balanced on county and date, where we cut the sample for each state two weeks after the state-wide shelter-in-place. 95% confidence intervals are based on standard errors double-clustered by county and date, as recommended in Brzezinski et al. (2020c).

Table 1: US-Wide Voluntary and Lockdown Responses

	Dependent Variable: Percent at Home	
	Voluntary Response	Lockdown Response
Coefficient	0.0128***	0.0132***
	(0.00428)	(0.00488)
Date FE	X	X
County FE	X	X
County-Level Controls	Eq. 1	Eq. 2
Observations	352,560	300,690
R-squared	0.782	0.841

Note: The table shows US-wide regression results for β and θ from Equations 1 and 2 in Columns 1 and 2, respectively. Standard errors are double-clustered by county and date, as recommended in Brzezinski et al. (2020c). Full results shown in Appendix Table A.3. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 3: Results From Controlled SIR Model

(a) Costs of Voluntary and Total Social Distancing by State, Without Value of Life

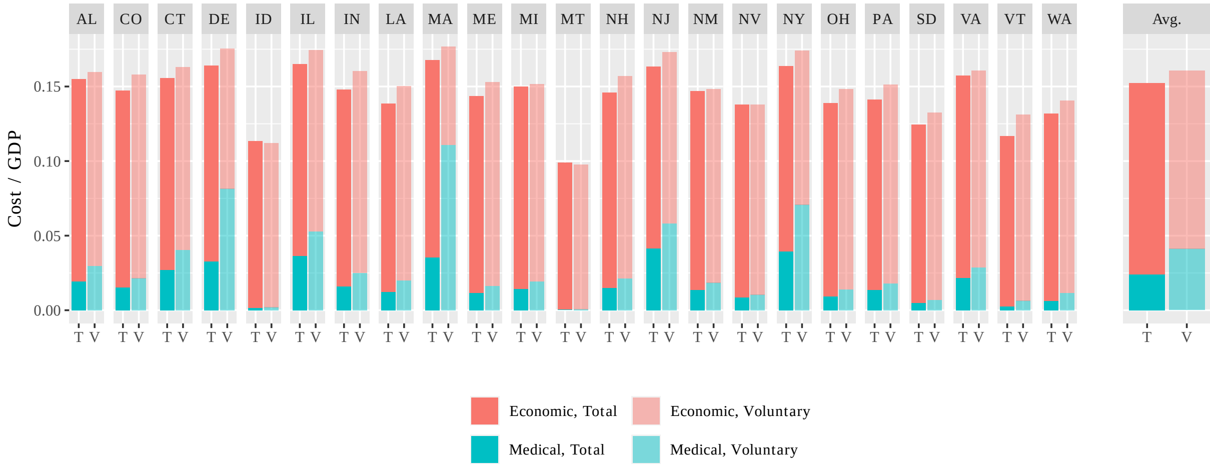

(b) Intensities of Voluntary Social Distancing That Incur Same Cost as Lockdown, Plus Costs for Each, for Various Lockdown Strengths

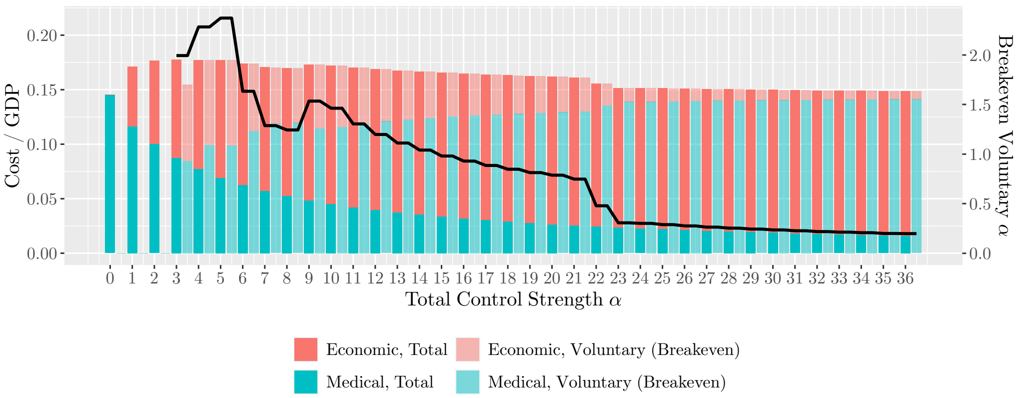

[1] **Panel a)**: estimated costs under voluntary social distancing **(V)** compared to total social distancing **(T)**—which includes voluntary (V) and lockdown-induced reductions (L) in movement. Estimated costs for (T) are based on simulations of the discretized C-SIR model using the α and ρ_0 estimates reported in Table A.2. Costs for (V) are re-estimated with α scaled by the estimated ratio $V/(V+L) = V/T$. Estimates of responses V and L are those reported in Figure 2 and Table 1. Costs are in terms of GDP per capita. Costs are broken down into economic (red) and medical (blue, without value of life), see Equation 9. Avg. serves as a proxy for the US overall and refers to the weighted average of the states in the sample, weighted by state-level GDP in Q4:2019 measured in current prices. NB the small decrease in total costs for ID, MT, NV is due to numerical inaccuracies and discontinuities in the cost function.

[2] **Panel b)**: *Black line*: gives, for each control strength α on the x-axis, the corresponding lower α on the y-axis that would incur the same costs, i.e. the maximal corresponding level of voluntary social distancing that could obtain without incurring higher costs than full lockdown. Estimates are from C-SIR model simulations with ρ_0=3. *Bars*: give the total (saturated color) and voluntary (transparent color) costs as % of GDP p.c. incurred by the total control strength α given on the x-axis and the corresponding (given by black line) voluntary response strength α on the right y-axis. Costs are broken down into economic (red) and medical (blue).

References

Abraham, Sarah and Liyang Sun (2018). "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects". *Available at SSRN 3158747*.

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston (2020). "A multi-risk sir model with optimally targeted lockdown". *NBER Working Paper* (w27102).

Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh (2020). "Inequality in the impact of the coronavirus shock: evidence from real time surveys". *Journal of Public Economics*, Forthcoming.

Allcott, Hunt et al. (2020a). "Economic and health impacts of social distancing policies during the coronavirus pandemic". *Available at SSRN 3610422*.

Allcott, Hunt et al. (2020b). "Polarization and public health: partisan differences in social distancing during the coronavirus pandemic". *NBER Working Paper* (w26946).

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt (2020). "The impact of the coronavirus pandemic on gender equality". *Covid Economics 4*, 62–85.

Andersen, Martin (2020). "Early evidence on social distancing in response to covid-19 in the united states". *Available at SSRN 3569368*.

Anderson, Roy M, Hans Heesterbeek, Don Klinkenberg, and T Déirdre Hollingsworth (2020). "How will country-based mitigation measures influence the course of the covid-19 epidemic?" *The Lancet 395* (10228), 931–934.

Atkeson, Andrew (2020). "What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios". *NBER Working Paper* (w26867).

Baek, ChaeWon, Peter B McCrory, Todd Messer, and Preston Mui (2020). "Unemployment effects of stay-at-home orders: evidence from high frequency claims data". *Institute for Research on labor and employment working paper*, 101–20.

Bai, Yan et al. (2020). "Presumed asymptomatic carrier transmission of covid-19". *Jama 323* (14), 1406–1407.

Ball, Laurence M, Daniel Leigh, and Prakash Loungani (2013). "Okun's law: fit at fifty?" *NBER Working Paper* (w18668).

Barrios, John M., Efraim Benmelech, Yael V. Hochberg, Paola Sapienza, and Luigi Zingales (2020). "Civic capital and social distancing during the covid-19 pandemic". *BFI Working Paper 2020-74*.

Barrios, John Manuel and Yael V Hochberg (2020). "Risk perception through the lens of politics in the time of the covid-19 pandemic". *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-32).

Barro, Robert J, José F Ursúa, and Joanna Weng (2020). "The coronavirus and the great influenza pandemic: lessons from the spanish flu for the coronavirusâs potential effects on mortality and economic activity". *NBER Working Paper* (w26866).

Bell, Brian, Nicholas Bloom, Jack Blundell, and Luigi Pistaferri (2020). "Prepare for large wage cuts if you are younger and work in a small firm". *VoxEU.org*.

Bettencourt, Luis MA and Ruy M Ribeiro (2008). "Real time bayesian estimation of the epidemic potential of emerging infectious diseases". *PLoS One 3* (5).

Bonardi, Jean-Philippe, Quentin Gallea, Dimtrija Kalanoski, and Rafael Lalive (2020). "Fast and local: lockdown policies affect the spread and severity of covid-19". *Covid Economics 23*, 325–251.

Borusyak, Kirill and Xavier Jaravel (2017). "Revisiting event study designs". *Available at SSRN 2826228*.

Briscese, Guglielmo, Nicola Lacetera, Mario Macis, and Mirco Tonin (2020). "Compliance with covid-19 social-distancing measures in italy: the role of expectations and duration". *NBER Working Paper* (w26916).

Brodeur, Abel, Andrew E Clark, Sarah Fleche, and Nattavudh Powdthavee (2020). "Assessing the impact of the coronavirus lockdown on unhappiness, loneliness, and boredom using google trends". *arXiv preprint arXiv:2004.12129*.

Brzezinski, Adam, Guido Deiana, Valentin Kecht, and David Van Dijcke (2020a). "The covid-19 pandemic: government vs. community action across the united states". *Covid Economics: Vetted and Real-Time Papers 7*, 115–156.

Brzezinski, Adam, Valentin Kecht, David Van Dijcke, and Austin L Wright (2020b). "Belief in science influences physical distancing in response to covid-19 lockdown policies". *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-56).

——— (2020c). "Hierarchical clustering in staggered difference-in-differences designs: a randomization inference approach".

CDC (2020). "Severe outcomes among patients with coronavirus disease 2019 (covid-19)âunited states, february 12–march 16, 2020". *Morbidity and Mortality Weekly Report (MMWR) 69* (12), 343–346.

Chudik, Alexander, M Hashem Pesaran, and Alessandro Rebucci (2020). "Voluntary and mandatory social distancing: evidence on covid-19 exposure rates from chinese provinces and selected countries". *NBER Working Paper* (w27039).

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (2020). "The cost of the covid-19 crisis: lockdowns, macroeconomic expectations, and consumer spending". *NBER Working Paper* (w27141).

Cornelson, Kirsten and Boriana Miloucheva (2020). "Political polarization, social fragmentation, and cooperation during a pandemic".

Cotti, Chad D, Bryan Engelhardt, Joshua Foster, Erik Nesson, and Paul Niekamp (2020). "The relationship between in-person voting, consolidated polling locations, and absentee voting on covid-19: evidence from the wisconsin primary". *NBER Working Paper* (w27187).

Durante, Ruben, Luigi Guiso, and Giorgio Gulino (2020). "Civic capital and social distancing: evidence from italiansâ response to covid-19". *VoxEU Column*.

EC (2020). “European economic forecast-summer 2020 – statistical annex”. *Institutional Paper 132*.

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt (2020). “The macroeconomics of epidemics”. *NBER Working Paper* (w26882).

Endo, Akira, Sam Abbott, Adam J Kucharski, Sebastian Funk, et al. (2020). “Estimating the overdispersion in covid-19 transmission using outbreak sizes outside china”. *Wellcome Open Research 5* (67), 67.

Engle, Sam, John Stromme, and Anson Zhou (2020). “Staying at home: mobility effects of covid-19”. *Covid Economics: Vetted and Real-Time Papers 4*, 86–102.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer (2020). “Internal and external effects of social distancing in a pandemic”. Tech. rep. w27059.

Fine, Paul, Ken Eames, and David L Heymann (2011). “Âherd immunityâ: a rough guide”. *Clinical infectious diseases 52* (7), 911–916.

Fowler, James H, Seth J Hill, Nick Obradovich, and Remy Levin (2020). “The effect of stay-at-home orders on covid-19 cases and fatalities in the united states”. *medRxiv*.

Galasso, Vincenzo (2020). “Covid: not a great equaliser”. *Covid Economics 19*, 241–265.

Goodman-Bacon, Andrew (2018). “Difference-in-differences with variation in treatment timing”. *NBER Working Paper* (25018).

Goodman-Bacon, Andrew and Jan Marcus (2020). “Using difference-in-differences to identify causal effects of covid-19 policies”.

Gros, Claudius, Roser Valenti, Lukas Schneider, Kilian Valenti, and Daniel Gros (2020). “Containment efficiency and control strategies for the corona pandemic costs”. *arXiv:2004.00493v2*.

Grossman, Guy, Soojong Kim, Jonah Rexer, and Harsha Thirumurthy (2020). “Political partisanship influences behavioral responses to governors’ recommendations for covid-19 prevention in the united states”. *Proceedings of the National Academy of Sciences*. eprint: `https://www.pnas.org/content/early/2020/09/14/2007835117.full.pdf`.

Holtz, David et al. (2020). “Interdependence and the cost of uncoordinated responses to covid-19”. *Proceedings of the National Academy of Sciences 117* (33), 19837–19843.

Hsiang, Solomon et al. (2020). “The effect of large-scale anti-contagion policies on the covid-19 pandemic”. *Nature 584* (7820), 262–267.

Jones, Callum J, Thomas Philippon, and Venky Venkateswaran (2020). “Optimal mitigation policies in a pandemic: social distancing and working from home”. *NBER Working Paper* (w26984).

Kaplan, Greg, Benjamin Moll, and Gianluca Violante (2020). “The great lockdown and the big stimulus: tracing the pandemic possibility frontier for the u.s.” *NBER Working Paper* (w27794).

Krueger, Dirk, Harald Uhlig, and Taojun Xia (2020). “Macroeconomic dynamics and reallocation in an epidemic”. *Covid Economics: Vetted and Real-Time Papers 5*, 21–55.

Kushner Gadarian, Shana, Sara Wallace Goodman, and Thomas B Pepinsky (2020). “Partisanship, health behavior, and policy attitudes in the early stages of the covid-19 pandemic”. *Health Behavior, and Policy Attitudes in the Early Stages of the COVID-19 Pandemic (March 27, 2020)*.

Kwok, Kin On, Florence Lai, Wan In Wei, Samuel Yeung Shan Wong, and Julian WT Tang (2020). “Herd immunity–estimating the level required to halt the covid-19 epidemics in affected countries”. *Journal of Infection 80* (6), e32–e33.

Lasry, Arielle et al. (2020). “Timing of community mitigation and changes in reported covid-19 and community mobilityâfour us metropolitan areas, february 26–april 1, 2020”. *Morbidity and Mortality Weekly Report (MMWR) 69* (15), 451–457.

Li, Qun et al. (2020). “Early transmission dynamics in wuhan, china, of novel coronavirus–infected pneumonia”. *New England Journal of Medicine*.

Liu, Ying, Albert A Gayle, Annelies Wilder-Smith, and Joacim Rocklöv (2020). “The reproductive number of covid-19 is higher compared to sars coronavirus”. *Journal of travel medicine*.

Lloyd-Smith, James O, Sebastian J Schreiber, P Ekkehard Kopp, and Wayne M Getz (2005). “Superspreading and the effect of individual variation on disease emergence”. *Nature 438* (7066), 355–359.

Painter, Marcus O. and Tian Qiu (2020). “Political belief affect compliance with covid-19 social distancing orders”. *Covid Economics: Vetted and Real-Time Papers 4*, 103–123.

Riou, Julien and Christian L Althaus (2020). “Pattern of early human-to-human transmission of wuhan 2019 novel coronavirus (2019-ncov), december 2019 to january 2020”. *Eurosurveillance 25* (4), 2000058.

Rossi, Rodolfo et al. (2020). “Covid-19 pandemic and lockdown measures impact on mental health among the general population in italy. an n= 18147 web-based survey.” *medRxiv*.

Simonov, Andrey, Szymon Sacher, Jean-Pierre Dubé, and Shirsho Biswas (2020). “The persuasive effect of fox news: non-compliance with social distancing during the covid-19 pandemic”. *NBER Working Paper* (w27237).

Sneppen, Kim and Lone Simonsen (2020). “Impact of superspreaders on dissemination and mitigation of covid-19”. *medRxiv*.

Thunstrom, Linda, Stephen Newbold, David Finnoff, Madison Ashworth, and Jason F Shogren. “The benefits and costs of flattening the curve for covid-19”. *Journal of Benefit-Cost Analysis*, Forthcoming.

Twenge, Jean and Thomas Joiner (2020). “Mental distress among us adults during the covid-19 pandemic”.

Viner, Russell M et al. (2020). “School closure and management practices during coronavirus outbreaks including covid-19: a rapid systematic review”. *The Lancet Child & Adolescent Health*.

WHO (2020). *Report of the who-china joint mission on coronavirus disease 2019 (covid-19)*.

Wright, Austin L, Konstantin Sonin, Jesse Driscoll, and Jarnickae Wilson (2020). "Poverty and economic dislocation reduce compliance with covid-19 shelter-in-place protocols". *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-40).

A Tables

Table A.1: Model Parameters

Parameter	Description	Value
k	Proportionality factor linking X_{tot} and medical costs $C_{medical}$; including cost of value of lifes	0.305
$\tilde{k}$	Proportionality factor linking X_{tot} and medical costs $C_{medical}$; excluding cost of value of lifes	0.14
m	Proportionality factor linking social distancing and social costs C_{social}	0.28
I_0	Starting value for share of infected I	2×10^{-5}
$I_{\min}$	Share of infected in a population I that can be controlled by testing and tracing without the need for imposing large-scale lockdown policies	10^{-5}

Note: The table provides a summary of the parameters used for the model in section 4, including the descriptions and the values obtained from Gros et al. (2020).

Table A.2: Controlled SIR Estimates

State	α	g_0	ρ_0	CEI
NJ	12.2	1.114	3.047	0.916
NY	13.2	1.129	3.094	0.921
IL	13.5	1.079	2.943	0.926
MA	14.8	1.109	3.031	0.93
DE	15.8	1.089	2.971	0.935
CT	20.7	1.118	3.058	0.949
VA	25.1	1.077	2.937	0.959
AL	28.1	1.062	2.892	0.964
IN	35.9	1.091	2.976	0.971
CO	37.2	1.085	2.96	0.972
NH	37.7	1.074	2.926	0.972
MI	42.1	1.114	3.047	0.974

State	α	g_0	ρ_0	CEI
NM	42.6	1.085	2.96	0.975
PA	44.7	1.117	3.055	0.976
LA	56.1	1.193	3.296	0.979
ME	48.7	1.055	2.872	0.979
OH	65.9	1.097	2.996	0.984
NV	69.5	1.087	2.967	0.985
WA	99.5	1.114	3.047	0.989
SD	164.2	1.241	3.458	0.992
US	**168.5**	**1.202**	**3.326**	**0.993**
VT	307	1.221	3.391	0.996
ID	530.1	1.282	3.604	0.998
MT	1225.6	1.251	3.495	0.999

Note: The table shows the parameter estimates for the controlled SIR model by state. Parameters α and g_0 are obtained from Equation 6, estimated by weighted least squares with weights I_t. Data obtained from the New York Times, where X_t is the cumulative confirmed cases and I_t the daily new cases. We also calculate the discretized bi-weekly reproduction number $\rho_0 = exp(g_0)$ and the containment efficiency index CEI $= \alpha/(g_0 + \alpha)$, CEI $\in [0, 1]$, following Gros et al. (2020).

Table A.3: Full Regression Results, With Spillovers

	Dependent Variable: Percent at Home	
	Voluntary Response	Lockdown Response
Policies		
First Case	0.0128***	
	(0.00428)	
State Lockdown	-0.00374	0.0132***
	(0.00638)	(0.00488)
State Lockdown × First Case	0.0126***	
	(0.00469)	
County Business Closure	0.0140	-0.00106
	(0.0147)	(0.00990)
School Closure	0.000575	-0.00162
	(0.00221)	(0.00404)
SOE	-0.0539***	0.00126
	(0.0139)	(0.00132)
County SOE	0.00614***	0.00287
	(0.00215)	(0.00183)
Spillovers		
Geo Adj Business		0.0134***
		(0.00213)
Geo Adj School		0.00103
		(0.00560)
Geo Adj Shltr		0.00445
		(0.00708)
COVID-19 Spread		
State Cases	7.40e-07***	4.79e-07***
	(1.70e-07)	(1.63e-07)
State Deaths	-6.82e-06***	-4.92e-06**
	(2.15e-06)	(2.15e-06)
Date FE	X	X
County FE	X	X
Observations	346,560	78,588
R-squared	0.793	0.858

Note: The table shows US-wide regression results from Equations 1 and 2 in Columns 1 and 2, respectively. The policy variables refer to dummy variables equal to 1 if the policy is in place or if the first case has occurred, respectively. Spillovers are modelled by using the geographic weighting approch described in section 3. Covid-19 cases and deaths are included in levels. Both regressions also include date and country fixed effects. Standard errors double-clustered by county and date as recommended in Brzezinski et al. (2020c).
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Figures

Figure B.1: Data Collapse for XI Representation, US States

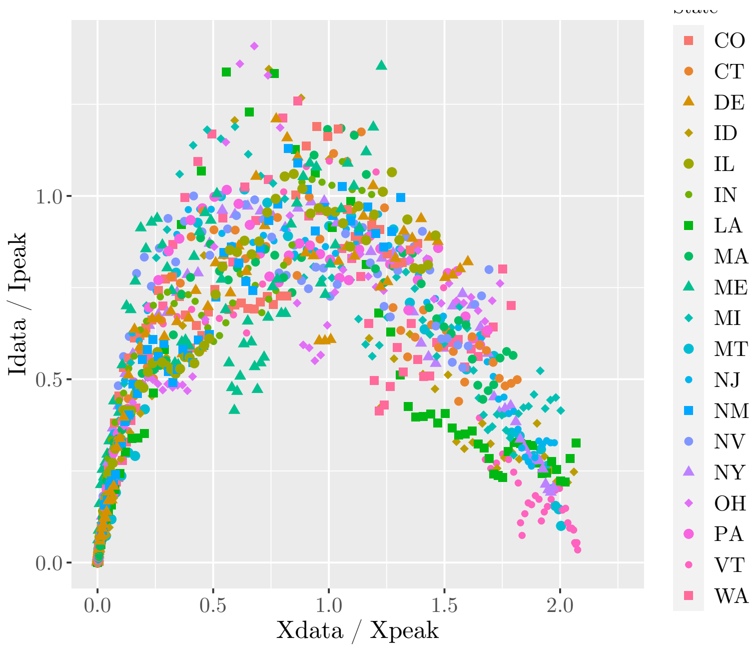

Note: Figures shows data collapse of Equation 5 to near-universal inverted parabola, obtained by rescaling by X_{peak} and I_{peak}, as calculated from $\hat{g}_0$ and $\hat{\alpha}$ for each state. As in Gros et al. (2020, Fig.1), but for 23 US states.

Figure B.2: Costs As A Function of Control Strength α

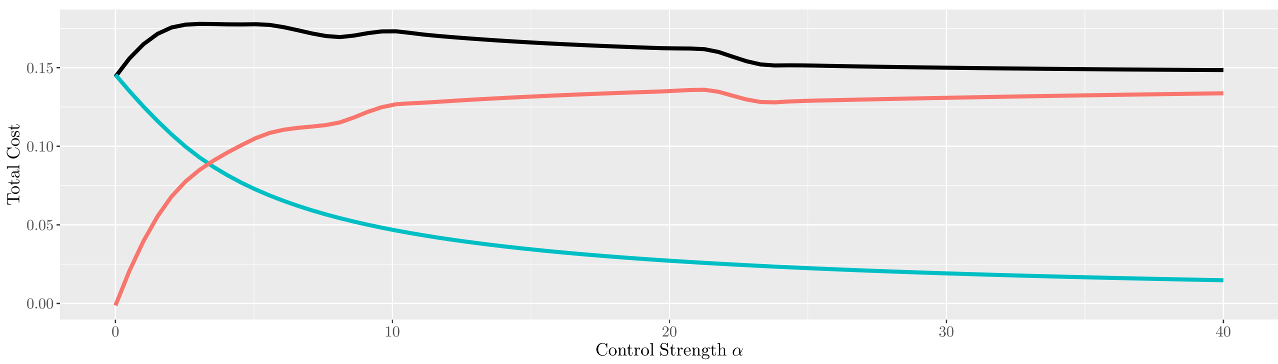

Note: Reproduction of Gros et al. (2020, Fig.5) Costs are in terms of GDP per capita. Total costs (black line) are defined as $C^{\text{total}} = C^{\text{econ}} + C^{\text{medical}}$. Economic costs (red line) are calculated as $C^{\text{econ}} = \Sigma_{I>I_{\text{min}}}\ 0.25[1 - \rho/\rho_0] \times 2/52$, while medical costs (blue line) are calculated as $C^{\text{medical}} = 0.14X_{\text{tot}}$, which excludes the estimated statistical value of a life.

Figure B.3: Costs of Voluntary and Total Social Distancing by State, With Value of Life Included

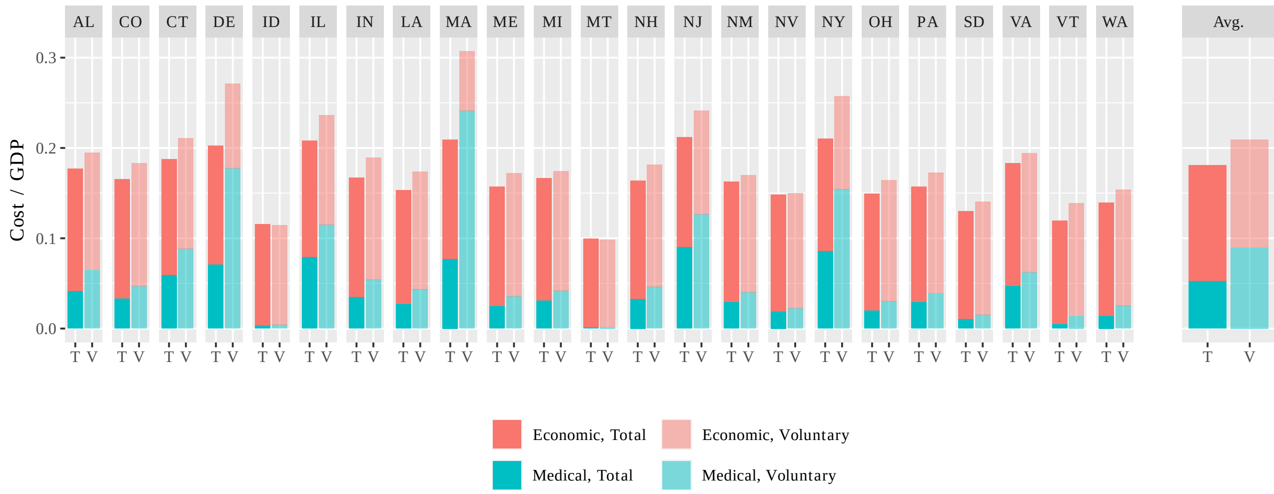

Note: The figure shows the estimated costs under the voluntary social distancing scenario **(V)** compared to total social distancing **(T)**, which includes voluntary (V) and lockdown-induced reductions (L) in movement. The estimated total costs for (T) are based on simulations of the discretized C-SIR model using the α and ρ_0 estimates reported in Table A.2. The costs for (V) are re-estimated with α scaled by the estimated ratio $V/(V+L) = V/T$. Estimates of the social distancing response V and L are those reported in Figure 2 and Table 1. Costs are in terms of GDP per capita. Costs for both (V) and (T) are broken down into their economic and medical components. Economic costs are calculated as $C^{\text{econ}} = \Sigma_{I>I_{\text{min}}}\ 0.2804[1-\rho/\rho_0] \times 2/52$, while medical costs are calculated as $C^{\text{medical}} = 0.305X_{\text{tot}}$, which is based on Gros et al. (2020) and includes the estimated statistical value of a life. Avg. serves as a proxy for the US overall and refers to the weighted average of the states in the sample, weighted by state-level GDP in Q4:2019 measured in current prices.

C Specification of the C-SIR Model

Short vs. long-term control. The controlled SIR model offers two different variants: In the first one, long-term control, policymakers and individuals react to the cumulative number of cases, while in the second one, short-term control, their actions only depend on the current number of infections. For the case of long-term control, starting from equation 4 and taking logs yields

$$\ln(g) = \ln(g_0) - \ln(1 + \alpha X), \tag{C.1}$$

and equivalently in the case of short-term control:

$$\ln(g) = \ln(g_0) - \ln(1 + \alpha I). \tag{C.2}$$

For each state, g_0 is constant by definition and α is assumed to remain constant over the course of the pandemic. Therefore, one may expect a linear relationship between $\ln(g)$ and $\ln(X)$ or $\ln(X)$, respectively. We check this assumption by plotting the residuals of these variables from regressions on state fixed effects in figure C.1. Both the cumulative and the daily case numbers exhibit a clear negative correlation with the reproduction rate, with R^2 of 0.91 and 0.83. This is intuitive as both variables are very closely related during the first wave of the pandemic, on which we focus in this paper. Due to the better tractability of the model under long-term control and the higher explanatory power of X, we opt for using the cumulative case numbers.

Figure C.1: Daily Cases I, Cumulative Cases X and R_t

(a) Short-Term Control **(b)** Long-Term Control

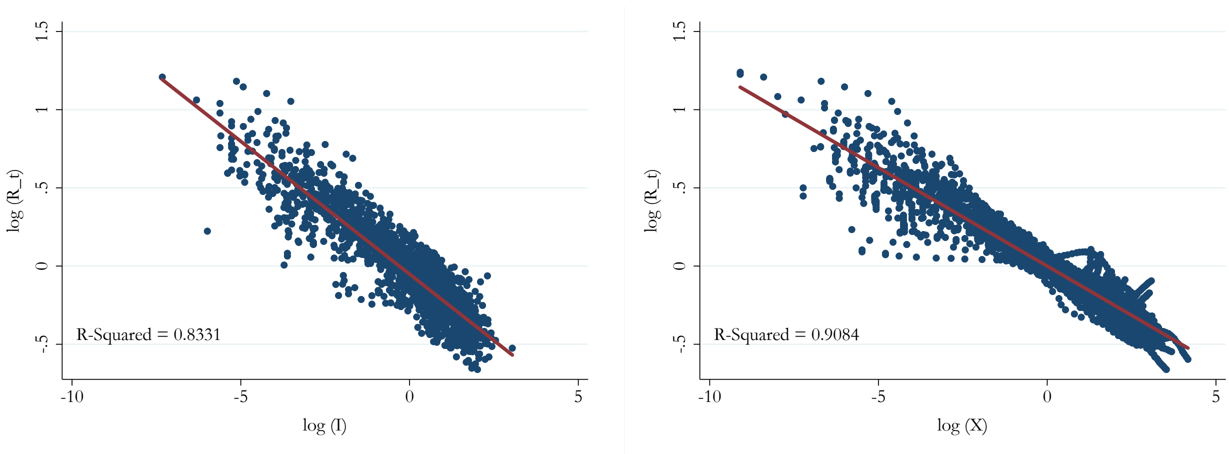

Note: The figure plots the reproduction rate R_t against the number of daily cases I (Panel (a)) and the cumulative number of cases X (Panel (b)). All three variables refer to residuals from regressions on state fixed-effects. The sample starts with the occurrence of the first case in each state, as R_t cannot be computed before this point, and ends on April 23.

Economic cost parameter m. The cost parameter m links social distancing and economic costs. We estimate m from the weekly average unemployment rate across US states and the reduction state-level reproduction rate. The unemployment rate has been retrieved from the

United States Department of Labor, while we estimate the reproduction rate using a modified version of the algorithm of Bettencourt and Ribeiro (2008), and express it with respect to the basic reproduction number $R_0 = 3$ taken from the literature (Liu et al., 2020).[30] This gives an estimate of the relationship between the change in the unemployment rate and the change in the reproduction rate of $m_u = 0.1262$. From there, we obtain the relationship between regional GDP and the reproduction rate as $m = m_u \times 2.22 = 0.2804$, using an estimate of Okun's coefficient for the US (Ball et al., 2013). In the next section, we perform a series of robustness checks across different values of m to validate our results.

Our baseline estimate of 0.28 is close to the estimate Gros et al. (2020) use based on China's experience (0.25). One would expect that when a larger share of GDP is consumed domestically (low openness to trade), a drop in domestic demand corresponding to a decrease in R_t leads to a higher drop in GDP. In that light, it makes sense that the US's estimate is close to, but slightly larger than, China's. Note that the economic costs estimated in this way will include local general equilibrium as well as partial equilibrium effects. The local unemployment rate might, in addition, be affected by cross-state and international general equilibrium effects. Yet, insofar as these should not be systematically related with the local reproduction rate, such effects should only act by shifting the intercept, while not affecting the slope of the line in Figure C.2. Lastly, note that since the US pandemic was in effect marked by several semi-distinct regional outbreaks, it makes sense to study the relation between R_t and economic output at a regional level.

[30] As time progresses and more economic data becomes available, it will be possible to estimate m more precisely. Already, however, our robustness exercises in Section D suggests that our results hold for a wide range of credible parametrizations of m.

Figure C.2: ΔR_t Vs. Δ Unemployment, Feb-April 2020, US States

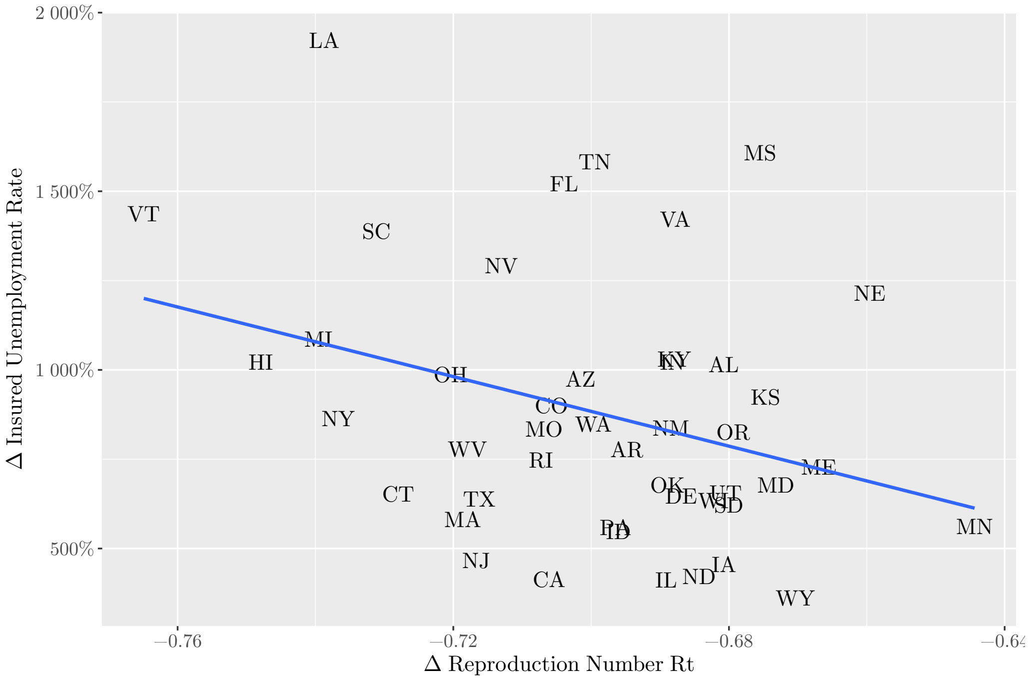

Note: % change in insured unemployment rate and the effective reproduction number R_t, as measured in the week of April 18, 2020. Δ unemployment rate is the change with respect to February 2020, as measured by the US Department of Labor. ΔR_t is the change in R_t, as estimated using a modified version of Bettencourt and Ribeiro (2008), with respect to the baseline reproduction number $R_0 = 3$ (Liu et al., 2020).

D Robustness

Robustness of DiD estimates for voluntary responses. The estimates presented in Figure 2 are based on a so-called 'stacked' DiD design. As shown in Goodman-Bacon (2018) and Abraham and Sun (2018), this estimand averages over heterogeneous treatment effects and is biased if these effects vary over time. To assess the robustness of our results we therefore apply an event-study approach where we estimate the effect separately for each day after the first case or policy implementation date. This strategy is recommended in Goodman-Bacon and Marcus (2020) in order to check for potential biases that may arise from the dynamic evolution of effects over time. Lockdown responses have been estimated in this way by a number of papers in the literature (Brzezinski et al., 2020a; Painter and Qiu, 2020; Wright et al., 2020). For our baseline specification, however, we do not opt for this approach since the cost calculations presented in section 4 require time-invariant estimates.

To study the dynamic response of social distancing to first local cases, we estimate the following equation:

$$pct_{i,j,t} = \alpha_i + \delta_t + \tilde{\zeta} Lock_{j,t} + \sum_{s=-9}^{10} (\beta_s C_{i,t_0+s} + \gamma_s C_{i,t_0+s} \times Lock_{j,t_0+s}) + \tilde{\mathbf{\Omega}} \mathbf{y_{i,t}} + \tilde{\epsilon}_{i,t}, \quad \text{with } s \neq -1 \tag{D.1}$$

where all the variables are as in Equation 1, except that the indicator $First_{i,t}$ is replaced by the set of dummies $C_{i,t+k}$. These are centered around the time period t_0 that refers to the first county-level COVID case, and are equal to 1 at time $t_0 + k$ if the first case is k days away and 0 otherwise. For all t where it holds that $t < t_0 - 10$ and $t > t_0 + 10$, we include two binning dummies that equal 1 when these respective conditions hold and 0 otherwise. β_{-9} to β_{-2} serve as placebo checks for the common trends assumption. As pointed out by Borusyak and Jaravel (2017), fully dynamic event studies with two-way fixed effects suffer from an underidentification problem, in that at least two restrictions are required to pin down a constant and a linear term. Thus, we follow the authors in omitting the two most disparate placebo checks as reference periods, namely -10 and -1, to improve the stability of the results.

Panel a) in Figure D.1 shows the resulting estimates for the voluntary respons, which are very close to the baseline estimates presented in Table 1. One day after the first case, there is a 0.9 percentage point increase in the number of devices that stay completely at home, which increases slightly during the following days. Compared to a base-level of 25% in February, this amounts to a 3.6% increase. This estimate is very similar to the one in the main specification. Moreover, across almost all dates after the treatment date considered, our baseline estimates in section 3 lie within the 95% confidence interval of the event-study estimates. Concerning the pre-trends, we do not detect any significant chances prior to the first local cases, indicating that we sufficiently control for the previous dynamics in order to attribute a causal interpretation to the result. Finally, the limited time heterogeneity in the treatment effects suggests that the stacked DiD estimates are not strongly affected by this potential source of bias.

Note that our definition of voluntary responses as the event-study estimates around the occurrence of first local cases is very narrow. Alternatively, one could for instance view the country-wide trend in absence of lockdowns a behavioral change which is not induced by policies and hence voluntary. These estimates, presented in Farboodi et al. (2020), among others, are a much larger than the results shown here, strongly reinforcing the overall conclusion of this paper.

Robustness of DiD estimates for lockdown response. For the case of the lockdown response, we pursue an equivalent event-study strategy to evaluate the bias arising from heterogeneous treatment effects over time. This approach closely follows a number of papers in the literature (Brzezinski et al., 2020a; Painter and Qiu, 2020; Wright et al., 2020). In particular, we estimate the following equation:

$$pct_{i,j,t} = \alpha_i + \delta_t + dayssince_{i,t} + \sum_{s=-9}^{10} \theta_s P_{j,t_0+s} + \tilde{\zeta} D_{i,t}^{GEO} + \tilde{\mathbf{\Psi}} \mathbf{y_{i,t}} + \tilde{u}_{i,t}, \quad \text{with } s \neq -1 \tag{D.2}$$

where all the variables are defined as in the main specification, except that P_{j,t_0+k} is a dummy variable centered around the state-wide shelter-in-place policy implementation date t_0 which is equal to 1 at time t_0+k and 0 otherwise. As before, k does not take the value -1, and we bin periods larger than t_0-10 and t_0+10.

Panel b) in Figure D.1 shows the results. The introduction of shelter-in-place policies leads to a 2.2 percentage points increase in the devices that stay completely at home one day after the implementation. As before, the estimated coefficients are reasonably close to those from the baseline specification, while time heterogeneity is fairly limited. The pre-trend is not significantly different from zero, providing evidence that we sufficiently control for increased social distancing before the enactment of the shelter-in-place policy.

Figure D.1: Staggered DiD estimates

(a) Voluntary Response, Event-Study Design

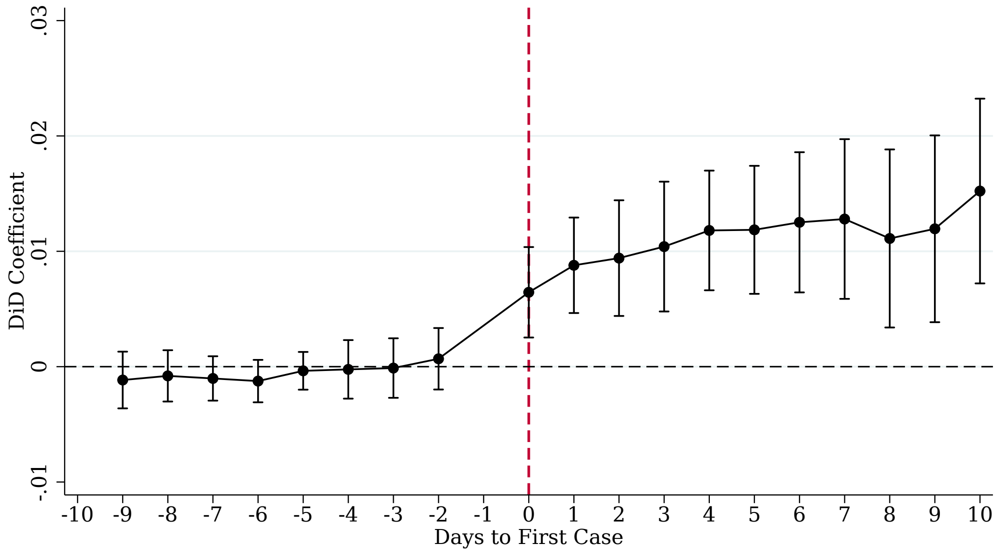

(b) Lockdown Response, Event-Study Design

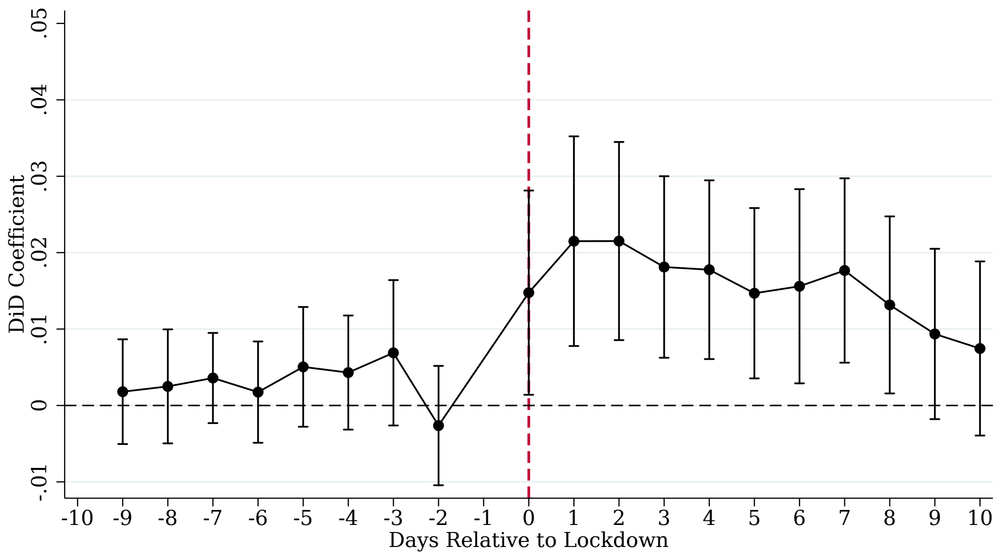

[1] **Panel a)**: Event-study approach for the Voluntary Response showing estimates from 10 days before to 10 days after the first local case. Reference periods are -10 and -1. Periods > 10 and < -10 are absorbed by 'bin' dummies. Vertical lines indicate 95% confidence intervals.

[2] **Panel b)**: Event-study approach for the Lockdown Response showing estimates from 10 days before to 10 days after a lockdown. Reference periods are -10 and -1. Periods > 10 and < -10 are absorbed by 'bin' dummies. Vertical lines indicate 95% confidence intervals.

Robustness to different response estimates ($\alpha, \alpha_v, \alpha_l$). To consider the robustness of our results to different estimates of the voluntary versus lockdown responses obtained in sections 3-4, panel b) of Figure 3 further elaborates on the trade-off between higher economic costs and lower medical costs. The black line shows, for each response strength $\alpha \in$ (x-axis), the corresponding lower $\alpha \in$ (right y-axis) that generates the same amount of total costs (as indicated in Figure B.2).[31] Thus, the black line links the two strategies that yield the same economic costs: a *lockdown strategy* that relies on diminishing medical costs at the expense of higher economic costs (x-axis), and the corresponding *voluntary strategy* that would yield the same costs by alleviating pressures on the economy at the cost of higher medical costs (y-axis). Due to the (approximate) concavity of total costs in α in this range, the more intensive the lockdown strategy, the less intensive the corresponding voluntary strategy required to generate the same level of costs. The reason for this is further elaborated by the bars in the figure, which show the cost decomposition for the lockdown strategy (saturated color) and voluntary strategy (transparent color) for each α: the former relies on low medical costs and the latter on low economic costs.

Thus, there are two necessary conditions for a voluntary strategy to yield lower costs than a lockdown strategy: First, α needs to be lower than around 36, as a value of above 36 yields lower costs that cannot be replicated by any α below. Second, the voluntary response needs to be sufficiently weak such that the economic costs are indeed kept at low levels when no lockdown is imposed. The latter condition implies that the counterfactual control strength α that would apply under no lockdown has to be at or below the black line of Panel b) in Figure 3.

These two conditions are very unlikely to be jointly satisfied. First, around two thirds of the states lie above the critical threshold of 36, implying that for these states no voluntary strategy can yield lower total costs—regardless of the precise ratio of voluntary to lockdown response. In this case, lockdown measures are always the most cost-effective solution. Second, even if some estimates are below 36, it is unlikely that the counterfactual voluntary responses would be sufficiently muted to generate lower costs than a lockdown strategy. For the most part, this would require very low levels of voluntary social distancing that do not come close to our results from section 3, combined with unrealistically low levels of α that are far from our lower bound estimates. For instance, even an α of 10—substantially lower than our lowest estimate—would require for the voluntary response to be at most at $\alpha = 1.5$, or less than 15% of the total control strength. This is unlikely according to our results in section 3: Even for the lowest significant point estimate for the voluntary effect obtained from our sample (Massachusetts), the size of the voluntary response is as large as 50% of the combined response, as the voluntary and lockdown response are approximately equal.

Two conclusions about the robustness of our results emerge from this discussion. First, since most of our estimates of the total control strength all lie above 36 (see Table A.1), the benefits of a lockdown always outweigh the costs for most of the states we consider. Thus, our main finding is not dependent on the precise magnitude of the estimates in section 3. Second, we show that, even for those states that have a very low estimated lockdown effectiveness, it is highly unlikely that an approach that relies on voluntary social distancing would yield lower total costs than a lockdown strategy. This is because one would require implausibly low counterfactual voluntary

[31] Note that the relevant range for the figure is $\alpha \in [3, 36]$, since below 3 and above 36 there is no lower α that would lead to the same total costs.

responses, combined with a very ineffective lockdown response, for the voluntary scenario to be more efficient.

Robustness to different economic cost estimates (m). We now consider how changing the baseline economic cost estimates m would affect our conclusion that lockdown responses generate lower total costs. Higher estimates of m could in theory change our conclusion, as they would change the trade-off between economic and medical costs: for any given level of social distancing, the economic costs would be more substantial, making a voluntary strategy more attractive.

Figure D.2 shows how many of the states we consider would experience lower total costs from a voluntary response than from a lockdown, for different estimates of m. As discussed in section 4, for our baseline estimate of 0.28, no state would fare better under a voluntary strategy. Yet, under slightly larger value of m the first two states would perform better under a voluntary strategy. Note that even at an estimate of 0.7, we would estimate only twelve states to perform better under a voluntary response. This estimate, however, is unreasonably high: it would imply that a change in the reproduction number of a factor of 10 would decrease GDP per capita by 70%. Since the reproduction numbers empirically changed by a factor between 8-10, this would thus imply an economic cost of more than 50% of GDP. Even at such unreasonable estimates, however, the vast majority of states in our sample would be better off with a lockdown strategy, given the magnitudes of their estimated voluntary response.

Figure D.2: Number of States for which Estimated Voluntary Costs are Lower than Lockdown Costs, by m

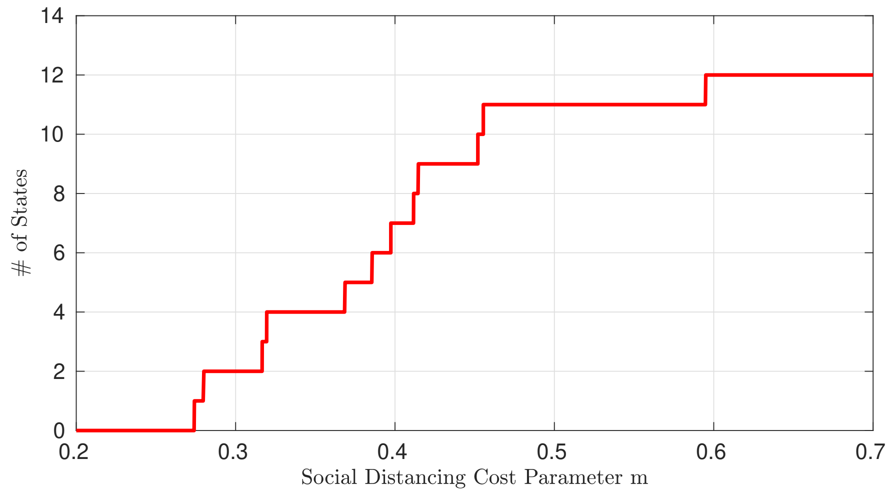

Note: The figure shows, for increasing cost parameter m, how many states could have 'empirically' attained a break-even between voluntary and lockdown response for their estimated control strength α. The empirical break-even point here corresponds to the estimated state-specific voluntary responses from Figure 2.

It is important to note that the model very much allows for the possibility that, in theory, the voluntary response can lead to smaller costs. Thus, the finding that lockdowns generate smaller costs is not simply a foregone conclusion that trivially follows from our SIR specification. Figure D.3 further elaborates on this point by looking at the necessary estimates of m under which it would be theoretically possible that the voluntary response generates lower costs, given the overall α parameters estimated for our sample of states. In other words, the figure considers the counterfactual situation where the voluntary response would have been close to zero for different cost parameter estimates. In general, already at a parameter of $m = 0.3$ it could be possible for half of the states to have lower costs under a voluntary regime, if people barely responded independently to the spread of the virus. However, our estimates of the relative strength of the voluntary responses and lockdown responses found in section 3 combined with any reasonable parametrization of our SIR model suggest that lockdown responses most likely generate lower total costs than voluntary responses across all US states in our sample.

Figure D.3: Number of States That Could Theoretically Break Even, by m

(a) Number of States

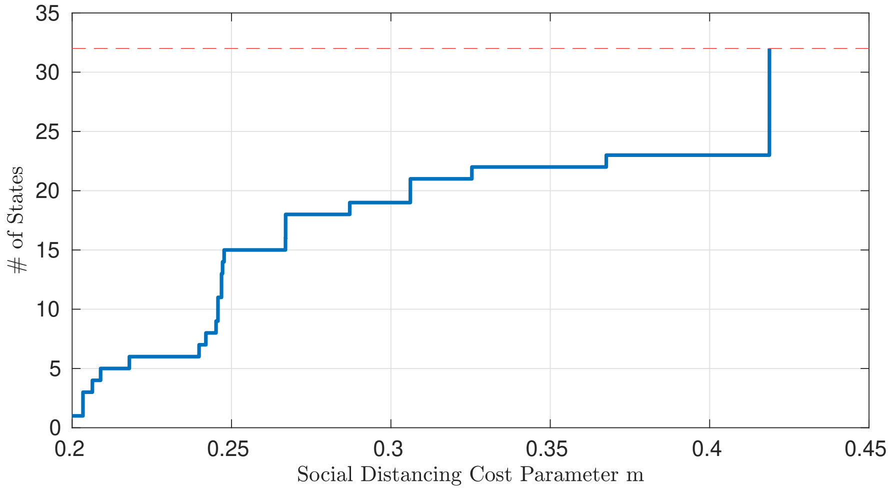

(b) Corresponding 'Threshold Control Strength' α

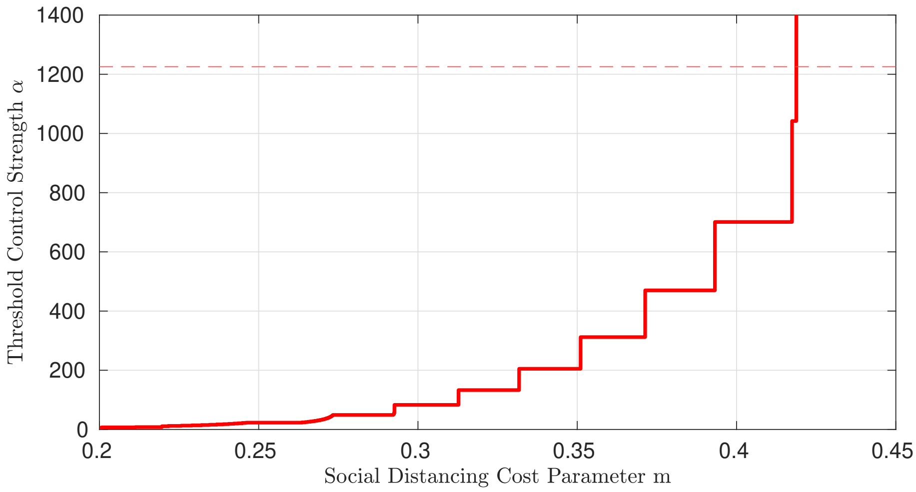

[1] **Panel a)**: Shows, for increasing cost parameter m, how many states could have 'theoretically' attained a break-even between voluntary and lockdown response for their estimated lockdown control strength α. The theoretical break-even point here corresponds to the absence of any voluntary response, i.e. $\alpha = 0$. Obtained from simulations of C-SIR model, as in Figure 3.

[2] **Panel b)**: Plots corresponding 'threshold control strength' α. Lockdown responses equivalent to any α' higher than this threshold are always better than no response at the given cost parameter m.