

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**Economics and Epidemics: Evidence from an Estimated Spatial
Econ-SIR Model**

Mark Bognanni, Doug Hanley, Daniel Kolliner, Kurt Mitman

2020-091

Please cite this paper as:

Bognanni, Mark, Doug Hanley, Daniel Kolliner, and Kurt Mitman (2020). "Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model," Finance and Economics Discussion Series 2020-091. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2020.091>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model*

Mark Bognanni[†] Doug Hanley[‡] Daniel Kolliner[§] Kurt Mitman[¶]

September 21, 2020

Abstract

Economic analysis of effective policies for managing epidemics requires an integrated economic and epidemiological approach. We develop and estimate a spatial, micro-founded model of the joint evolution of economic variables and the spread of an epidemic. We empirically discipline the model using new U.S. county-level data on health, mobility, employment outcomes, and non-pharmaceutical interventions (NPIs) at a daily frequency. Absent policy or medical interventions, the model predicts an initial period of exponential growth in new cases, followed by a protracted period of roughly constant case levels and reduced economic activity. Nevertheless, if vaccine development proved impossible, and suppression cannot entirely eradicate the disease, a utilitarian policymaker cannot improve significantly over the laissez-faire equilibrium by using lockdowns. Conversely, if a vaccine will arrive within two years, NPIs can improve upon the laissez-faire outcome by dramatically decreasing the number of infectious agents and keeping infections low until vaccine arrival. Mitigation measures that reduce viral transmission (e.g., mask-wearing) both reduce the virus’s spread and increase economic activity.

*We would like to thank our discussant, Kyle Herkenhoff, as well as Adrien Auclert, Timo Boppart, John Hassler, Per Krusell, Laura Pilossoph, Simon Mongey, and seminar participants at Bonn ECONtribute Virtual Macro Workshop, Federal Reserve Board, IIES, SafeGraph Consortium Seminar, and the Virtual Macro Seminar (VMACS) for stimulating discussions. We thank the IZA for generous research support through the Coronavirus Emergency Research Thrust. We thank SafeGraph for providing access to their data and the SafeGraph COVID-19 Data Consortium for helpful input. We thank Ray Sandza and Homebase for providing access to their payroll data. Mitman gratefully acknowledges support from the Ragnar Söderbergs stiftelse, and the European Research Council grant No. 759482 under the European Union’s Horizon 2020 research and innovation programme. Philipp Hochmuth and Fabian Sinn provided excellent research assistance. We presented earlier versions of this research project under the titles “Economic Activity and COVID-19 Transmission: Evidence from an Estimated Economic-Epidemiological Model,” and “Economic Activity and COVID-19 Transmission.”

[†]Board of Governors of the Federal Reserve System. Email: mark.j.bognanni@gmail.com. The views stated herein are those of the authors and are not necessarily those of the Board of Governors of the Federal Reserve System.

[‡]University of Pittsburgh. Email: doughanley@pitt.edu

[§]University of Maryland. Email: d.kolliner@gmail.com

[¶]IIES, Stockholm University, CEPR & IZA. Email: kurt.MIT.man@gmail.com

1 Introduction

The rapidly developing COVID-19 pandemic presents perhaps the most daunting challenge to economic policymakers since the Great Depression. By engaging in economic activity, individuals subject themselves to infection risk, while infected agents impose negative externalities on others by spreading the virus. In the absence of a vaccine or cure, policymakers can combat the virus’s spread only through “non-pharmaceutical interventions” (NPIs) (e.g., recommending social distancing, stay-at-home orders). While NPIs may effectively slow the rate of new infections, they do so primarily by reducing economic activity. Alternatively, allowing the virus to propagate unchecked may also entail economic costs, as rational agents will mitigate their exposure to life-threatening risks. Formulating a sound policy response to the pandemic thus requires quantifying the tradeoffs between health and economic activity, as well as accurate predictions of *both* economic and epidemiological variables in response to alternative policies. Managing an epidemic thus requires an integrated assessment model, combining economics and epidemiology and rigorously disciplined by data, analogous to the integrated climate-economic models for evaluating policies to combat climate change (Nordhaus, 1993).¹

This paper develops and estimates a parsimonious model of the joint determination of epidemiological and economic variables in an epidemic. We call our framework an “economic SIR” (Econ-SIR) model. The model features optimizing agents who endogenously respond to changes in both infection risk and NPIs. The framework nests the standard SIR model (Kermack and McKendrick, 1927) as the special case of a zero infection cost. When the infection cost is sufficiently high, the model endogenously gives rise to two phases of virus spread, with new infections initially growing exponentially before transitioning to a period of approximately linear growth. The virus’s effective replication rate in the quasi-linear regime is just below one, leading to protracted dynamics of decreased economic activity and roughly constant spread of the virus.

While others have concurrently and independently formulated similar models (e.g., Brotherhood et al. (2020); Farboodi et al. (2020); Keppo et al. (2019); Toxvaerd (2020)), this paper is the first to formally bring such a model to the data to identify the extent to which the real world demands the Econ-SIR’s divergence from pure epidemiological models. We develop a novel estimation approach to fit the model to a panel of daily economic and epidemiological data for approximately 1000 U.S. counties, covering over 80 percent of the U.S. population. Through revealed preference, we quantify both the disutility of contracting the disease and the economic consequences of mitigation policies—the two key trade-offs for policy analysis.

¹We thank Per Krusell for suggesting this analogy.

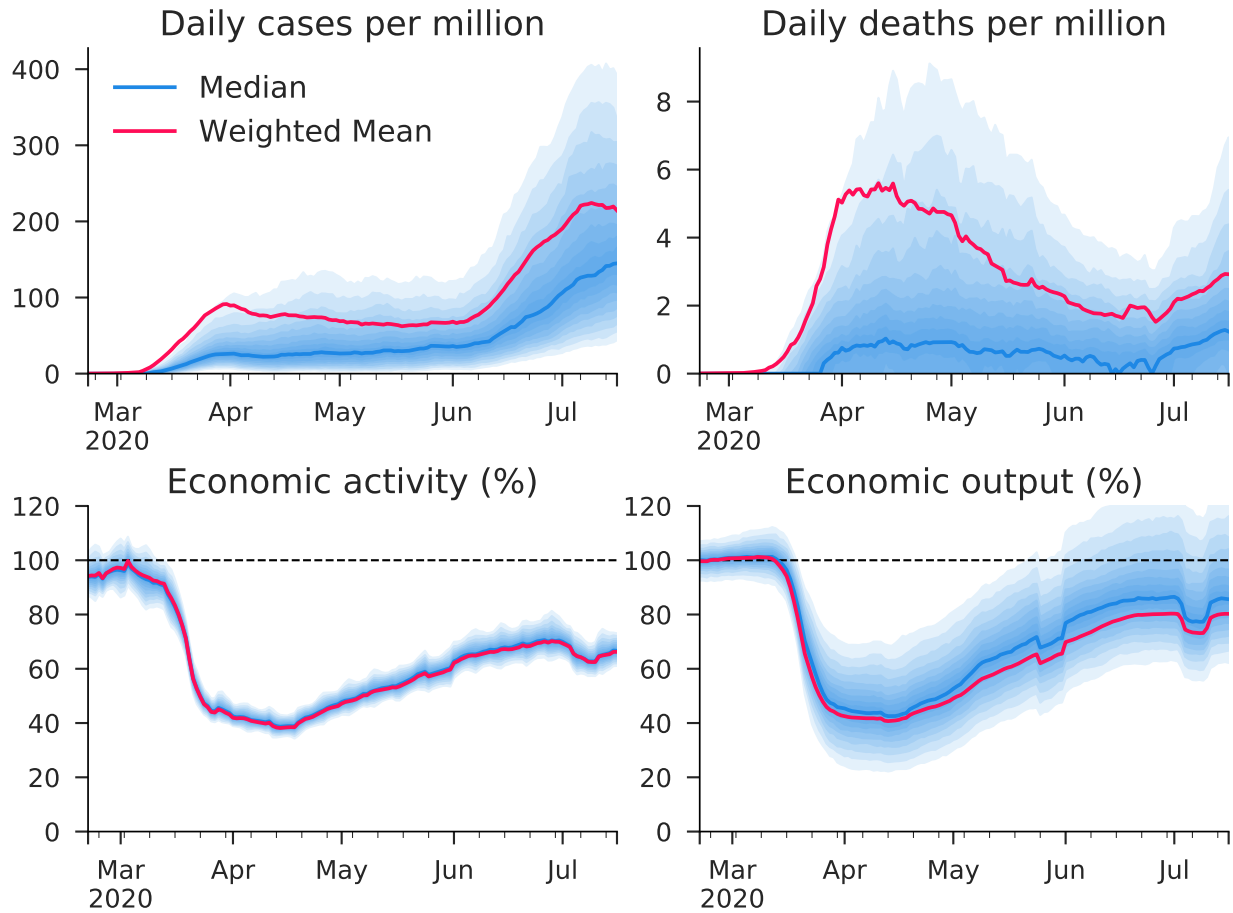


Figure 1: Clockwise from upper left: new reported COVID-19 cases, new reported COVID-19 deaths, hours worked and economic activity. In each panel, red lines are population-weighted means, the thick blue lines are the median county, and the light blue regions demarcate deciles of the distribution across counties (all of these statistics computed pointwise).

Our results highlight the importance of using an integrated assessment model for analyzing policy responses to epidemics, as the endogenous mitigation of private agents yields substantially different epidemiological dynamics than the standard SIR model. While our quantitative results focus on the COVID-19 epidemic, the framework and estimation method have broader applicability for studying policy responses to future epidemics.

The estimated model closely matches the joint evolution of economic and epidemiological variables throughout the pandemic, both in the aggregate time series and lower-level units of analysis, such as U. S. states. In particular, the model can capture both the beginning of the pandemic and the spatial heterogeneity in the virus’s “summer resurgence.” Looking beyond the data used directly for estimation, we find that our estimated model is consistent even with data on the seroprevalence of COVID-19 antibodies in numerous U.S. locations. At the estimated parameter values, individuals face a high infection cost, implying that the model’s dynamics depart significantly from the standard SIR model. We find that a combination of endogenous mitigation by agents and the rollout of NPIs explains the decline in economic activity and output at the onset of the epidemic. These results are consistent with identified effects of “stay-at-home” orders on activity from empirical studies examining the rollout and reversal of NPIs at the county level (Goolsbee and Syverson, 2020).² Furthermore, as in Goolsbee and Syverson (2020), we find that disentangling the endogenous mitigation of individuals from the effect of policy requires a county-level analysis (as opposed to state or national) as many counties enacted NPIs earlier than their corresponding states.

Confident in our model’s success in capturing the salient economic and epidemiological features of the epidemic, we proceed to use the framework as a laboratory to conduct counterfactual experiments. We characterize the implications of several policy scenarios for economic output and the virus’s human toll. The first experiment examines the *laissez-faire* outcome of the model. In particular, we simulate the model forward from August 1, 2020, assuming that all NPIs are reverted and no further NPIs are put into place. The model predicts that removing all restrictions on schools and restaurants would lead to an initial burst of economic activity and output, which yields a jump in the effective replication rate and generates surge in reported cases and a corresponding surge in daily deaths. Seeing the surge in infections, individuals then endogenously reduce economic activity, bringing the effective replication rate back below one by the early fall. Over the course of the first year of *laissez-faire*, the model predicts approximately 620,000 deaths from COVID-19, while output would remain more than 10% below the pre-pandemic level.

Considering *laissez-faire* over the longer run, and still assuming the development of neither

²See also earlier innovative empirical work on NPIs by Adda (2016) and recent cross-country evidence by Mendiola et al. (2020)

a vaccine nor a cure, the model reaches the “slow burn” regime sometime in 2021, where the behavior of the model is roughly linear on a day to day basis. In this phase, individuals reduce their economic activity until the effective reproductive factor (\mathcal{R}_0) is approximately 1. The pandemic then plays out over the next three to four years, with a roughly constant decline in daily deaths and cases and an increase in economic activity as the model approaches herd immunity. During this phase, COVID-19 causes roughly 300,000 U.S. deaths per year, and economic output recovers roughly linearly from about 90 percent of pre-pandemic levels. These protracted dynamics imply a final toll of roughly 1.25 million COVID-19 deaths over the next 4 years.

The PPF, while instructive about the tradeoffs a policymaker faces, is agnostic about how to value the tradeoff between output and lives. We proceed to compute the optimal adaptive lockdown policies (a Ramsey-like approach) of a utilitarian planner. The planner can only set the lockdown’s strength as a function of the current number of infected individuals. Under the assumption that a cure or vaccine never arrives, we find that the optimal adaptive policy closely approximates the laissez-faire outcome. The intuition behind this result is that neither lockdown policies nor endogenous social distancing by agents can change the herd immunity threshold in the model (recall that when infections are close to zero the model dynamics converge to those of the basic SIR model). The only role for policy in this scenario is to minimize the “overshoot” of the herd immunity threshold to prevent “unnecessary” deaths and infections. The endogenous mitigation under laissez-faire generates very little overshoot at our estimated costs of infection, hence prescribing a small role for policy.³

Given the stark predictions of the baseline model, we consider how the optimal adaptive policies change under two plausible extensions. First, we rectify one of the unpleasant (and unrealistic) features of the standard SIR model—namely, we allow for the eradication of the virus.⁴ One need not take eradication literally, but we consider this thought experiment also to approximate an aggressive test, trace, and quarantine (TTQ) that can effectively control any community-based transmission of the virus. Second, we consider the possibility of the arrival of a vaccine.⁵

³Thus far, we have assumed that our planner places the same value on life as do the economic agents. If the planner places a different value on life, e.g. from the notion of a statistical value of life, the optimal policy may be different.

⁴In the standard SIR model with Poisson transition rates between states, the share of infected agents never reaches zero in finite time (absent a cure). As such, there will still be some measure of infected agents alive even after curtailing economic and social activity for an extended period (e.g., 100 years). After reversing the lock-down, the system will immediately revert to exponential growth. We allow for the share of infected individuals to fall to zero in finite time if the stock of infected individuals falls below a pre-specified threshold (e.g., if the percentage of infected implies fewer than one infected person in a county).

⁵We have also considered the arrival of a cure. The properties of the optimal adaptive policy under this alternative scenario are qualitatively the same as to the arrivals of a vaccine.

With the possibility of eradicating the virus, a meaningful role for policy emerges. The endogenous social distancing by agents in the model is not strong enough to affect a quick reduction in infections to the eradication thresholds we consider.⁶ Now we find that a strict lockdown for two months followed by a graduate reopening over the course of a month can effectively bring infections down to the “eradication zone” whereby community transmission is eliminated or managed via TTQ. Note, given the spatial nature of the model, it is important that the lockdown be coordinated across locations, to prevent re-seeding via mobility. In this extension of the model, over the medium and long run, there is virtually no tradeoff between health and the economy. At one year (and longer horizons) one can maximize output (via temporary strict lockdowns followed by reopening) with minimal loss of life.

Finally, we consider the optimal adaptive policy when the arrival of a vaccine is near. We assume that a vaccine will arrive with certainty on August 1, 2021 and solve for the optimal adaptive policy until then. Under such a scenario, and given the current stock of infections, it is optimal to reduce activity to similar levels as under stay-at-home orders for two to four months. This significantly reduces the current stock of infections, after which the planner allows activity to increase until the effective replication rate is approximately 1. The planner achieves this by loosening NPIs to approximately the activity level estimated when restaurants are take-out only and schools are closed. NPIs are then rolled back once the vaccine arrives.

Our paper relates to the mushrooming literature studying the epidemiological and economic response to the COVID-19 pandemic (e.g. Alvarez et al. (2020); Baker et al. (2020); Berger et al. (2020); Brotherhood et al. (2020); Eichenbaum et al. (2020); Farboodi et al. (2020); Glover et al. (2020); Guerrieri et al. (2020); Hagedorn and Mitman (2020); Hur (2020); Jones et al. (2020); Kaplan et al. (2020); Keppo et al. (2019); Krueger et al. (2020); Lubik and Matthes (2020); Mitman et al. (2020); Toxvaerd (2020)).⁷ Given the real-time and rapidly changing nature of those papers we think that instead of providing a discussion of each it is more instructive and expedient to summarize where we have innovated relative to that body of work and what our main contribution is.

Another strand of papers has focused on forecasting purely the epidemiological dimension of the pandemic’s evolution in a panel of locations. See Liu et al. (2020) and Fernández-Villaverde and Jones (2020).

In addition to the now standard SIR model and its many variants, there is a large

⁶Our eradication zone thresholds are motivated by the experiences of Taiwan and S. Korea at effectively either eliminating community transmission in the case of the former, or managing it with TTQ in the latter.

⁷We apologize for any papers we have missed. Due to the rapidly changing landscape, this likely is not an exhaustive literature list.

epidemiological literature on the notion of “behavioral change” by agents in response to their environment and new information. Most of these studies do not explicitly incorporate economic factors or optimizing behavior into their analysis, though there are exceptions such as Morin et al. (2013) and Fenichel et al. (2011). See Verelst et al. (2016) for an extensive review of this literature.

Relative to standard epidemiological models, we have added economic behavior — what is sometimes called a behavioral SIR model.⁸ We have estimated the model at a granular level, using data on individual movements based on GPS and data on economic outcomes—namely, employment—to discipline the economic-epidemiological connection.

The rest of the paper is organized as follows. Section 2 presents the most parsimonious version of our Econ-SIR model. Section 3 focuses on the two distinct phases of virus spread to which the Econ-SIR model gives rise. Section 4 extends our Econ-SIR model to facilitate a more rigorous quantification of the COVID-19 pandemic. Section 5 describes our estimation strategy and documents the estimated model’s fit. Section 6 uses the estimated model to conduct and compare a variety of policy experiments. Section 7 concludes.

2 A Minimal Economic-SIR Model

We begin by describing the simplest version of our economic SIR model to focus on the principal mechanism by which economic behavior and virus transmission interact. In Section 4, we enrich the model with additional features particularly relevant to the COVID-19 pandemic.

On the economic side of the model, susceptible agents can choose whether or not to participate in the economic marketplace, facing the potential risk of being exposed to infected individuals. They optimally trade off the gains from market participation, with the cost of potentially becoming infected with the virus. The more widespread the infection is in the population, the higher the risk that market participants will become infected. Further, as more infected agents participate in the marketplace, more individuals become infected. The model thus features bi-directional interaction between economic activity and viral transmission.

We set our models in continuous time $t \in [0, \infty]$, and adopt the following notational conventions: constants and parameters are denoted with lowercase Greek letters, e.g., β ; stock variables, such as the measure of agents in a given epidemiological state, are denoted

⁸We have opted for the term “economic SIR” model, as our agents are fully rational, optimizing agents.

with capital English letters, e.g., $S(t)$; flow-variable counterparts to each stock variable use the same capital letter but add a “.” over the character, e.g., \dot{S} ; other endogenous variables are denoted with lowercase English letters, e.g., $n(t)$.

2.1 Epidemiological Structure

Epidemiological States. The model is populated with a time-invariant, unit measure of individuals. An individual can be in one of three epidemiological states: Susceptible (S), Infected (I), or Recovered (R). Susceptible agents can contract the virus, but have not yet done so. Infected agents have the virus and can spread it. Recovered individuals have previously had the virus, can no longer spread it, and possess immunity from reinfection. At a given time t , $S(t)$, $I(t)$, and $R(t)$ denote the measures of individuals in each epidemiological state.

Epidemiological state transitions. Given initial values S_0 , I_0 , and R_0 , the measures of agents in each state evolve according to the following system of equations,

$$\dot{S} = -n(t)S(t) \tag{1}$$

$$\dot{I} = n(t)S(t) - \delta I(t) \tag{2}$$

$$\dot{R} = \delta I(t), \tag{3}$$

in which Susceptibles contract the virus at rate $n(t)$ and Infecteds recover from the virus at rate δ . The number of new infections is given by $N(t) = n(t)S(t)$. If we declared $n(t)$ equal to $\beta I(t)$ and stopped here, then we would have the most standard SIR model. We next describe the economic environment that endogenizes the rate of infection, $n(t)$.

2.2 Economics and the Virus

We now introduce the novel economic dimension of the model. Agents face economic needs, which we model as a random variable with realizations denoted $z \in (0, \infty)$. Satisfying the economic need and receiving benefit equal to z requires making an excursion, while forgoing the excursion yields a benefit of only 0. For a Susceptible, an excursion exposes the agent to infection risk. For an Infected, an excursion exposes Susceptibles to increased infection risk. All else equal, we assume that agents prefer not to contract the virus and hence incur a cost ψ from infection.⁹ To capture the fact that not all economic needs are equally important, we

⁹As we discuss in 4, this cost captures all disutility associated with infection.

model z as drawn from a distribution $F(z)$.

Pausing at this level of generality, one can already see that the virus's economic implications will differ according to an agent's infection status. I and R agents face no infection risk, so they may undertake excursions without fear of infection to satisfy any z (though we will assume that infected agents partially self-quarantine). Susceptibles, however, face real infection risk and thus a tradeoff between the benefits of acquiring z with the potential cost of infection. This tradeoff makes Susceptibles the model's key economic decision-makers.

To z , or Not To z ? Conditional on taking an excursion, we model the rate of infection for Susceptibles as proportional to the economically active infected population share:

$$\tilde{I}(t) \equiv \beta I(t) a_I \quad (4)$$

where β is a time-invariant rate of virus transmission per unit of Infected and a_I is the rate at which infected agents take excursions. We refer to $\tilde{I}(t)$ as the “infection risk” (to an excursing Susceptible).

In light of their disutility from infection, Susceptibles optimally choose to satisfy a given economic need only if the benefit exceeds the expected cost of taking an excursion,

$$z(t) > \underbrace{\tilde{I}(t)\psi}_{\text{expected cost of excursion}} \equiv \bar{z}(t). \quad (5)$$

$\bar{z}(t)$ represents the reservation value of $z(t)$ for Susceptibles to excure, which varies endogenously over time through its dependence on $I(t)$.

Economic Activity's Endogenous Response to Infection Risk. Let $a_S(t)$ denote the rate of economic activity by a Susceptible—the rate at which the agent undertakes excursion. The optimizing decision rule in (5) implies that $a_S(t)$ is given by the rate at which Susceptibles face economic needs $z(t)$ greater than $\bar{z}(t)$. Denoting the complementary cumulative distribution function of z as $F^c(z) = 1 - F(z)$, the activity rate of Susceptibles is given by

$$a_S(t) = 1 - F(\bar{z}(t)) = F^c(\bar{z}(t)) = F^c(\tilde{I}(t)\psi). \quad (6)$$

With $\bar{z}(t)$ as given in (5), equation (6) reveals the the key channel by which the infectiousness of the environment effects the economic activity of Susceptibles: as the virus becomes more widespread ($I(t)$ increases), infection risk rises ($\tilde{I}(t) = \beta I(t) a_I$ increases), which causes Susceptibles to endogenously curtail their economic activity (as long as $\psi > 0$, $\bar{z}(t)$ increases

with $\tilde{I}(t)$ and $F^c(\bar{z}(t))$ falls). Noting that in the early stages of an epidemic the vast majority of agents belong to $S(t)$, this is the primary channel by which the virus affects aggregate activity. Thus, the model contains a channel by which the virus can suppress economic activity even in the absence of outside interventions.

Virus Transmission Rate’s Endogenous Response to Economic Activity. The rate of infection for Susceptibles, $n(t)$, is given by the product of their infection risk per excursion with the rate at which they take excursions. Using (4) and (6), the infection rate is thus

$$n(t) = \tilde{I}(t)a_S(t) = \tilde{I}(t)F^c(\tilde{I}(t)\psi) \quad (7)$$

The term $F^c(\tilde{I}(t)\psi)$, derived from the optimizing economic behavior of Susceptibles, introduces a nonlinear relationship between $I(t)$ and the new infection rate. Using (7), the flow rates \dot{S} and \dot{I} in Econ-SIR (see equations (1) and (2)) become

$$\dot{S} = -\tilde{I}(t)F^c(\tilde{I}(t)\psi)S(t) \quad , \quad \dot{I} = \tilde{I}(t)F^c(\tilde{I}(t)\psi)S(t) - \delta I(t) \quad (8)$$

The term $N(t) = n(t)S(t) = \tilde{I}(t)F^c(\tilde{I}(t)\psi)S(t)$ is the number of new infections. Equations (7) and (8) characterize the key channel by which the economic activity of Susceptibles affects the infectiousness of the environment: as Susceptibles reduce their economic activity ($a_S(t)$ falls), new infections decrease ($n(t)$ falls), which causes the number of infectious agents to fall (\dot{I} falls so $I(t)$).

Discussion. From equations (6) and (8), one can see that not only does virus prevalence have implications for economic activity, but economic activity has implications for virus prevalence. Lower economic activity by Susceptibles reduces the infection rate, which in turn reduces the number of new infections and the size of the Infected pool. When $I(t)$ increases, the endogenous reduction of activity by Susceptibles will tend to mitigate the rate of new infections.

2.3 Policy: Nonpharmaceutical Interventions in Econ-SIR

The primary policy lever exercised in response to COVID-19 has been the imposition of NPIs, such as lockdown orders. Particularly in the U.S., such orders have typically restricted the economic transactions that may take place (e.g. grocery shopping may be allowed out of necessity while tattoo parlors are required to close). We model such NPIs as a minimum $\bar{z}^P(t)$ that all agents internalize when deciding whether or not to undertake an excursion.




Figure 2: Structure of SIR and Economic SIR Models

Hence the *effective* reservation value of z for agents of type $\tau \in \{S, I, R\}$ becomes

$$\bar{z}_\tau^*(t) = \max \{ \bar{z}_\tau(t), \bar{z}^P(t) \}. \quad (9)$$

And hence activity by agents of type τ becomes

$$a_\tau(t) = F^c(\bar{z}_\tau^*(t)). \quad (10)$$

A few important aspects of how such a policy works are worth pointing out. First, for S agents, the endogenous $\bar{z}_S(t)$ forms an effective lower bound for \bar{z}^* : if a policymaker chooses a $\bar{z}^P(t)$ below agents' endogenous $\bar{z}(t)$, then the policy is inoperative and has no effect on agents of type τ . Policy options are thus notably asymmetric. Stated in real-world terms, policymakers can force people to stay home, but they cannot force people to go out.

Second, shutdowns are a blunt tool. Sufficiently restrictive shutdowns will reduce the activity of *all* agents. While such measures can mitigate or suppress the virus in the model, they are clearly a brute force way of doing so in light of the fact that, in principle, new infections could be driven to zero by driving just a_I to zero.

2.4 Economic Activity and Output

We will later treat economic activity and output as observables when estimating the model. However, in the data we cannot separately observe the activity of agents by epidemiological state, so the observables will correspond to the model's aggregate values. To fix ideas, it is helpful to first define these aggregate concepts within the basic Econ-SIR model.

Aggregate activity, denoted $a(t)$, is the measure-weighted mean of activity rates by each

group

$$a(t) = a_S(t)S(t) + a_I(t)I(t) + a_R(t)R(t) \quad (11)$$

In the absence of NPIs, $a_R(t) = 1$ and $a_I(t) = a_I = F^c(\bar{z}_I)$ and one can write $a(t)$ as¹⁰

$$a(t) = 1 - \underbrace{F(\bar{z}_S(t))}_{\text{rate of forgone excursions by } S} S(t) - \underbrace{F(\bar{z}_I)}_{\text{rate of forgone excursions by } I} I(t) \quad (12)$$

When $I(t)$ equals zero, the virus is behaviorally irrelevant and activity occurs at its full rate of 1.

Output is the rate at which value is accumulated from activity. The model differentiates between activity and output because whenever agents restrain their economic activity, they optimally sacrifice their least valuable activities first (forgo the lowest values of z , recall (5)). The optimizing behavior leads to a potentially non-linear relationship between activity and output, with the exact relationship determined by the form of F .

Letting $f(z) = F'(z)$ and $h(\bar{z})$ be the partial expectation of z , defined as

$$h(\bar{z}) = \int_{\bar{z}}^{\infty} z f(z) dz = \mathbb{E}[z \mid z > \bar{z}] F^c(\bar{z}) \quad (13)$$

the output rate is given by

$$y(t) = h(\bar{z}_S(t))S(t) + h(\bar{z}_I(t))I(t) + h(\bar{z}_R(t))R(t) \quad (14)$$

Equation (14) makes clear the differing output rates that result from the differing behavior between agents who are concerned with infection and agents who are not. In our estimated model, we use $F(z)$ such that $\log(z)$ is normally distributed with mean $-\frac{\sigma^2}{2}$ and variance σ^2 , in which case $\mathbb{E}[z] = 1$. Given this assumption, the variance parameter σ^2 summarizes the relationship between output and activity, as shown in Figure 3.

3 The Spread of the Virus in Econ-SIR: Exponential Growth and “The Slow Burn”

In this section we describe the dynamics of the Econ-SIR model, in particular the nature of the two phases of virus spread to which the model endogenously gives rise: an initial phase

¹⁰The following expression comes from substituting for $a_S(t)$ and using the fact that $S(t) + I(t) + R(t) = 1$.




Figure 3: Relationship between economic output and activity given various values of σ , the standard deviation of $\log(z)$.

of exponential growth following by a phase of approximately linear growth. When $\psi = 0$ Econ-SIR collapses to the standard SIR model whose dynamics are well-known from the epidemiological literature. As we describe below, increasing the cost of infection ψ will lead to a more muted region of exponential growth and a more protracted phase of approximately linear growth.

3.1 Phase 1: Exponential Growth of Infections

The baseline rate at which an Infected agent in the model reproduces the infection in the broader population, which we call the baseline reproductive rate \mathcal{R}_0 , is given by

$$\mathcal{R}_0 = \frac{\beta}{\delta}. \quad (15)$$

Namely, \mathcal{R}_0 is the ratio of the rate at which Infecteds transmit the virus per unit of Susceptibles to the rate at which the Infecteds recover and stop transmitting the virus altogether. The baseline reproductive rate will obtain when $S(t)$ is near 1.

Taking into account the transmission mitigation from reduced activity by Susceptibles, the *effective reproductive rate* in Econ-SIR is

$$\mathcal{R}_e = a_S(t)\mathcal{R}_0. \quad (16)$$

Using (2) and (7), the growth rate of infections in the model is given by

$$\frac{\dot{I}(t)}{I(t)} = \frac{\beta I(t) a_S(t) S(t) - \delta I(t)}{I(t)} = \beta a_S(t) S(t) - \delta \quad (17)$$

For very small values of $I(t)$ or ψ , the activity rate of Susceptibles will be $a_S(t) \approx 1$. Thus if initial infections I_0 are very small, to first-order dynamics of SIR and Econ-SIR are equivalent. Further, if $S(t) \approx 1$, Equation (17) then implies exponential growth of the infected population share according to rate

$$\frac{\dot{I}(t)}{I(t)} = \beta - \delta. \quad (18)$$

which is positive whenever $\mathcal{R}_0 > 1$.

3.2 Phase 2: The Slow Burn

As the virus spreads and $I(t)$ continues to increase, $a_S(t)$ endogenously declines and plays a more important role in mitigating the the pace of new infections. Here we develop intuition for the second phase of virus spread in which $I(t)$ is approximately constant, i.e. $\dot{I}(t) \approx 0$.

Recall that $\dot{I}(t)$ is given by

$$\dot{I}(t) = \beta I(t) a_S(t) S(t) - \delta I(t), \quad (19)$$

it follows that $\dot{I}(t) = 0$ when

$$\beta a_S(t) S(t) = \delta. \quad (20)$$

Considering the case where the infected share has been low enough that $R(t)$ is still negligible so that $S(t) = 1 - I(t) - R(t) \approx 1 - I(t)$, substituting for $a_S(t)$, and dividing through by β , one can write the condition in (20) as

$$[1 - F(\beta I(t) \psi)] (1 - I(t)) = \frac{\delta}{\beta} = \frac{1}{\mathcal{R}_0}. \quad (21)$$

Recalling that $\bar{z}(t) = \beta I(t) \psi$, one can express the previous relationship in terms of $\bar{z}(t)$ as

$$[1 - F(\bar{z})] \left(1 - \frac{\bar{z}(t)}{\beta \psi}\right) = \frac{1}{\mathcal{R}_0} \quad (22)$$

For large $\beta\psi$, equation (22) has the approximate solution for the endogenous object $\bar{z}(t)$ of

$$\bar{z}(t) \approx F^{-1} \left(1 - \frac{1}{\mathcal{R}_0} \right). \quad (23)$$

Again using $\bar{z}(t) = \beta I(t)\psi$, equation (22) implies that the pseudo-steady-state entails a constant fraction of agents in the infected state given by

$$I(t) = \frac{1}{\beta\psi} \cdot F^{-1}(1 - \mathcal{R}_0^{-1}) \quad (24)$$

Since new infections $N(t)$ must equal outflows, $\delta I(t)$, the constant levels of new infections are given by

$$N(t) = \frac{\delta}{\beta\psi} \cdot F^{-1}(1 - \mathcal{R}_0^{-1}) = \frac{1}{\psi\mathcal{R}_0} \cdot F^{-1}(1 - \mathcal{R}_0^{-1}). \quad (25)$$

The expression for $\bar{z}(t)$ in (23) immediately implies that the activity rate of Susceptibles satisfies $a_S(t) = 1 - F(F^{-1}(1 - \mathcal{R}_0^{-1})) = \mathcal{R}_0^{-1}$ and hence

$$\mathcal{R}_e = a_S(t)\mathcal{R}_0 = 1 \quad (26)$$

Thus, for sufficiently large $\beta\psi$, and while $R(t)$ remains relatively small, we observe an effective reproduction rate of approximately 1. Absent some other intervention, this “slow burn” regime will last until $R(t)$ becomes appreciably large, meaning “herd immunity” dynamics begin to have an impact.

3.3 Comparing SIR and Econ-SIR Epidemiological Dynamics

Here we highlight the key differences in the behavior of the SIR and Econ-SIR models. To understand the dynamics of both models it is instructive to understand the the curve in (S, I) space where $\dot{I}(t) = 0$, known as the nullcline of $I(t)$. Recall that $\dot{I}(t)$ is given by

$$\dot{I}(t) = \beta I(t)a_S(t)S(t) - \delta I(t), \quad (27)$$

it follows that the nullcline is characterized by

$$\beta a_S(t)S(t) = \delta. \quad (28)$$

In the SIR model, since $a_S(t) = 1 \forall t$, the nullcline of $I(t)$ is independent of $I(t)$ and is given by a vertical line at $S = 1/\mathcal{R}_0$, which is the herd immunity threshold (HIT).

Proposition 1 *When $\psi > 0$ and $F(0) = 0$, the nullcline of $I(t)$ is characterized by an upward sloping curve in (S, I) space, denoted $I^N(S)$, that originates at $(1/\mathcal{R}_0, 0)$. Furthermore, the slope of the nullcline is decreasing in ψ .*

Proof. Substituting (6) into (20) and rearranging, the nullcline is characterized by

$$I^N = \frac{1}{\beta\psi} F^{-1} \left(1 - \frac{1}{\mathcal{R}_0 S} \right)$$

Taking the limit as $S \rightarrow 1/\mathcal{R}_0 \Rightarrow I^N \rightarrow 0$. Differentiating I^N with respect to S yields:

$$\frac{\partial I^N}{\partial S} = \frac{1}{\beta\psi} \frac{\mathcal{R}_0 S^2}{f \left(F^{-1} \left(1 - \frac{1}{\mathcal{R}_0 S} \right) \right)} > 0$$

where the positivity follows since all parameters are positive and the pdf f is non-negative. The relationship is also decreasing in ψ . ■

Assuming initially that no agents have immunity to the virus $R_0 = 0$, if the initial seeding of the I_0 is such that $I_0 < I^N(1 - I_0)$, then the nullcline will describe the point that is the “peak” of the epidemic in terms of the stock of the number of infected agents and the effective replication rate is equal to unity, $\mathcal{R}_e = 1$. The differences between SIR and Econ-SIR are characterized graphically using the phase diagrams shown in Figure 4.

In this example, the parameters are chosen so that $\delta = \beta/2.5$, which yields $R_0 = 2.5$. In the SIR model, $I(t)$ then increases until $S(t)$ falls below 0.4, which is the value at which \dot{I} turns negative.¹¹ Reduction of $I(t)$ mass in Econ-SIR occurs well before reaching herd immunity, as agents try to avoid infection (as is visible from the bending of the nullcline from a vertical line). The strength of the mitigation increases with the cost of infection ψ . This results in both slower model dynamics — in SIR herd immunity is reached in 4-5 weeks, whereas in the rightmost example it takes 8-9 weeks. Further, the stronger mitigation results in smaller overshoot of the herd immunity threshold. Summing up, Econ-SIR generates more protracted dynamics and less overshoot relative to the SIR model. The extent of overshoot will become important later as this is the main externality that a policy maker is able to correct with lockdown policy.

¹¹Reduction of $I(t)$ mass in SIR does not occur until hitting “herd immunity.”




Figure 4: Phase diagrams for Econ-SIR models under alternative infection costs ψ . Blue lines represent sample time paths, dots along the time paths are spaced at one week intervals, black vertical dashed lines represent the herd immunity threshold, red dashed lines are nullclines of $I(t)$.

4 The Spatial Econ-SAIRD Model

This section extends our Econ-SIR model along multiple empirically-relevant dimensions towards rigorously quantifying the policy tradeoffs posed by the COVID-19 pandemic. We leave some of the mathematical specifics to Appendix B.

First, we add two features to the epidemiological environment that are particularly relevant for modeling COVID-19: asymptomatic virus carriers and mortality risk. It has been well documented that many COVID-19 carriers show few or no symptoms, sometimes only for a few days prior to becoming more seriously ill and sometimes agents recover with no increase in severity. As we now know, COVID-19 poses mortality risk uncommonly high for a virus so easily transmitted. Accordingly, when agents initially contract the virus in the extended model we assume that they move to an asymptomatic state (A). Asymptomatic agents in the model unknowingly transmit the virus. We model asymptomatic agents as typically unaware of their infectiousness, and hence they behave similarly to Susceptibles.¹² We should note that we keep ψ to denote the cost of being infected (as opposed to the mortality cost) to capture the fact that even non-fatal cases can entail significant medical costs to individuals (Kniesner and Sullivan, 2020). Asymptomatic agents recover directly from the A state at rate γ or become symptomatic and infectious, moving to I at rate λ . Infectious agents die (move to the D state) at rate κ . The other epidemiological states have

¹²We model this as an activity rate $a_A(t)$ that is a trailing moving average of the activity rate of Susceptibles.

the same interpretation as before. Individuals in the I state, being generally symptomatic, we model as having some awareness of their status and choosing to partially self-quarantine by undertaking actions no less valuable than some level \bar{z}_I . Activity among I agents is then given by $a_I = F^c(\bar{z}_I)$. We refer to the epidemiological state space of the model as SAIRD.

A second important feature of the real world is the presence of many separate locations, each with a distinct outbreak severity and virus transmission environment. Furthermore, real world agents travel across locations, potentially spreading the virus from one region to another. We add a spatial dimension to the model consisting of J distinct locations, each with its own measures of agents in the five epidemiological states (denoted with subscripts, as in $S_j(t)$) and its own values of β_j and ψ_j . All other parameters are the same in all locations. Denote the rate of economic activity by infectious agents in location j as

$$e_j(t) = a_{I,j}(t)I_j(t) + a_{A,j}(t)A_j(t). \quad (29)$$

If location j were an island, then the infection risk to an excursing Susceptible would be $\tilde{I}_j(t) = \beta_j e_j(t)$. However, allowing for movement across locations, we model the infection risk per $S_j(t)$ excursion as a weighted average of the local force of infection and the population-weighted national average force of infection,

$$\tilde{I}_j(t) = (1 - \alpha)\beta_j e_j(t) + \alpha \sum_k w_k \beta_k e_k(t), \quad (30)$$

where $\alpha \in [0, 1]$ and w_k is a weight proportional to location j 's relative population. We also allow for an information friction in which agents perceive their infection risk to be determined according to the form of (30), but with a value α^* that may differ from the true α . Denote an S_j agent's *perceived* infection risk as $\tilde{I}_j^*(t)$.

The economic decision agents face, “to z , or not to z ,” is similar to before. Susceptibles can take an excursion with (perceived) expected benefit of $z(t) - \tilde{I}_j^*(t)\psi_j$, or not. However, we enrich the environment to allow for a “no-contact consumption” option, by which agents can forgo an excursion and still obtain some of the value from a given z , namely $d \cdot z(t)$, where $d \in [0, 1]$. This model mechanism corresponds to the availability of delivery or takeout services for a variety of goods. With the delivery option, Susceptibles optimally take an excursion only if $z_j(t) > \tilde{I}_j(t)\psi_j/(1 - d) = \bar{z}_j(t)$.

We allow for the economic environment to evolve over time with trends in both the value of the delivery option ($d(t)$) and the safety of activity ($\beta(t)$). For $\beta(t)$, we represent the evolution via a common factor $\tilde{\beta}(t)$ that proportionally scales the values of β_j everywhere.

We model the trends as deterministic processes

$$d(t) = \bar{d} \cdot \Lambda(\lambda_d t) \quad \text{and} \quad \tilde{\beta}(t) = 1 + (\bar{\beta} - 1)\Lambda(\lambda_\beta t) \quad (31)$$

where \bar{d} and $\bar{\beta}$ represent the terminal values of the processes, $\lambda_d, \lambda_\beta \geq 0$ govern the speed of the transitions to the terminal values, and Λ is the “smoothstep” sigmoid function, which is increasing and for which $\Lambda(0) = 0$.¹³ Accounting for the common trend component, the value of the infection rate parameter in j at time t is given by $\beta_j(t) = \beta_j \cdot \tilde{\beta}_j(t)$.

NPI policy operates in fundamentally the same way in the extended model as in the Econ-SIR model (see Section 2.3). The activity of all agents in the model is potentially constrained by NPIs. This includes recovered R agents since NPIs have not been enforced differentially based on a person’s immunity status. When taking the model to data and estimating the severity of realized NPIs, we will allow for three different NPI severities (\bar{z}^{P_1} , \bar{z}^{P_2} , and \bar{z}^{P_3}), corresponding to school closures, restaurant take-out only orders, and stay-at-home orders (the most broadly used NPIs).

We model the distribution of economic needs as $\text{Lognormal}(\mu_z, \sigma_z)$.¹⁴ In addition to having support $(0, \infty)$, the Lognormal distribution allows for known expressions for partial and conditional expectations, which facilitates the calculation of quantities such as (14). When taking the model to the data there will be two additional considerations for the handling of z . Because we will normalize output in the data, we parameterize the distribution to have μ_z fixed and equal to $-\sigma_z^2/2$, which ensures $E[z] = 1$. Lastly, we will estimate the model from daily data, which necessitates some special considerations for the systematic fluctuations of activity and hours with day of the week. To accommodate the systematic day-of-week variation, we allow day-specific values of μ_z .

5 Estimating Spatial Econ-SAIRD’s Parameters

This section describes our approach to estimating the Econ-SAIRD model’s parameters and shows that the estimated model can generate the paths of economic and epidemiological variables that we observe over the course of the COVID-19 pandemic. Going forward, let θ denote the set of model parameters.

¹³The smoothstep function has a simple polynomial expression over the transition domain and is commonly used in computer graphics due to its ease of computability.

¹⁴With $z \sim \text{Lognormal}(\mu_z, \sigma_z)$, our parameters μ_z and σ_z represent the mean and standard deviation of the random variable $\log(z)$.

5.1 Data

Our estimation uses a panel of daily county-level observations over the period from February 15, 2020 to July 15, 2020 on NPIs, new COVID-19 cases, new COVID-19 deaths, point-of-interest visits (economic activity), and hours worked.¹⁵ After removing counties with insufficient data, the estimation uses 921 U.S. counties for inference, accounting for about 80 percent of the U.S. population.¹⁶ The economic variables are normalized county-by-county. For the next section, it will be useful to denote the vector of county i 's observables at time t as

$$\mathbf{x}_{i,t} = [Activity_{i,t}, Hours_{i,t}, NewCases_{i,t}, NewDeaths_{i,t}]' \quad (32)$$

and let $X \equiv \{\mathbf{x}_{i,t}\}_{i,t}$ denote the full panel of the observations.

The NPI data consists of three different types of NPIs: stay-at-home orders, restaurants and bars takeout-only orders, and school closures. While we observe the enactments of NPI orders, an NPI's implications for behavior must be inferred. Hence, an NPI in location i at time t of type k enters our data as a binary variable $b_{i,t,k}$, with a value of 0 for "off" and 1 for "on." The model represents the relevance of a type- k NPI through the value of the parameter \bar{z}^{P_k} , which we include in θ and estimate for each NPI type. For the next section, it will be useful to denote the full panel of NPIs as $B \equiv \{b_{i,t,k}\}_{i,t,k}$ and let $B_{0:t}$ be the panel of all NPIs up through date t .

5.2 Likelihood

Given θ and the panel of enacted NPIs, the SAIRD model generates a deterministic path of values for epidemiological and economic variables in each location.¹⁷ In particular, the model generates paths for newly infected and deceased agents, as well as economic activity and output. We take the deterministic structural model, along with measurement errors, to be the generating process for these four data series. The measurement errors create a probabilistic relationship between the data and model-generated paths that forms the basis of the likelihood function. We next describe our model for the data generating processes of the four variables of interest and then follow with a description of the resulting likelihood function.

¹⁵Appendix A contains additional details on the data.

¹⁶We use only counties with populations of at least 50,000. Because of missing or sparse observations in the Homebase hours data, we exclude a few additional counties that meet our minimum population threshold.

¹⁷The initial conditions for the epidemiological states are relevant as well, however we set them as a deterministic function of the elements of θ that we have already introduced. See Appendix C.1.

Economic Data Generating Process. We write $a_{i,t}^{\theta,B_{0:t}}$ to denote the model’s simulated value for economic activity in location i at period t , where the superscripts emphasize the simulated path’s dependence on the the particular values of the parameters and NPIs. For both of the economic variables we assume mean-zero additive Gaussian measurement errors around the model’s value. For economic activity, the data generating process is thus

$$Activity_{i,t} = a_{i,t}^{\theta,B_{0:t}} + \epsilon_{i,t,a} , \quad \epsilon_{i,t,a} \sim \text{Normal}(0, \sigma_a^2/w_i), \quad (33)$$

where w_i is a weight proportional to county i ’s relative population among counties in the sample. The model for $Output_{i,t}$ is analogous to (33), with the model object on the right-hand side instead being $y_{i,t}^{\theta,B_{0:t}}$. We estimate the series specific variances σ_a^2 and σ_y^2 .

Health Data Generating Process. Modeling the relationship between the model objects and health variables introduces some additional challenges. First, the data on new COVID-19 cases and deaths are discrete counts and their scale varies considerably with county populations. Meanwhile, the model’s predictions for new cases and deaths take the form of continuous population shares. Taking new cases as an example, one can straightforwardly map the model-implied flows into expected counts for a county with population pop_i as $pop_i \times N_i(t)$.¹⁸ The expected count can then discipline the *mean* of a discrete probability distribution over the positive integers.

Second, the data are subject to systematic measurement error of at least two types. The first type is reporting delays: the tendency of new infections or deaths occurring at time t to not show up in the data until $t + d$. We examined the empirical distribution of reporting delays in locations that separately report both day-of-death and “new deaths,” which lead us to incorporate a delay of $d = 10$ days for both cases and deaths (see Figure A.1 in the Appendix).

The second type of systematic measurement error in the health data is underreporting. We account for underreporting by assuming that asymptomatic cases are generally not reported (i.e. agents in $A(t)$ do not show up in new cases data) and that symptomatic cases (new entrants to $I(t)$) are systematically underreported, particularly in the early phase of the pandemic. We model the evolution of underreporting of new cases using a sequence of reporting factors $r_t \in [0, 1]$ that converge to 1 at an exponential rate. In the limit, cases are correctly reported.¹⁹ Accounting for the systematic delays and underreporting, we have the

¹⁸Note that the observed value $x_{i,t,j}$ satisfies $x_{i,t,j} = n_i p_{i,t,j}$ for a true population fraction $p_{i,t,j}$. The relevant model flows are given in equations (43) and (45) as $N_i(t) = \lambda A_i(t)$ and $\dot{D}_i(t) = \kappa I_i(t)$.

¹⁹While we build under-reporting of cases into the model and estimation, we do not do so for deaths. We estimate the model to best match the *reported* data for COVID-19 deaths (including the reporting delays).

following implication that will center the DGP for new cases,

$$E[NewCases_{i,t}] = r_{t-d} \cdot pop_i \cdot N_{i,t-d}^{\theta, B_{0:t-d}}. \quad (34)$$

Lastly, having accounted for the systematic dimensions of the measurement error, we model the idiosyncratic variation as arising according to a Conway-Maxwell-Poisson (CMP) distribution.²⁰ The $CMP(\lambda, \nu)$ distribution extends the Poisson distribution by allowing the center and dispersion to be specified separately. The parameter ν governs the distribution's dispersion, allowing us to model the difference in measurement accuracy between cases and deaths. Thus, we model the new cases data as

$$NewCases_{i,t} \sim CMP(\lambda_C, \nu_C) \quad (35)$$

where $\nu_C = 0.1$ and, conditional on ν_C , λ_C is chosen to fix the mean of the CMP distribution at the value prescribed by (34). We model the data on new deaths analogously, but with $\nu_D = 100$. The differing dispersion parameters reflect the view that the deaths data are considerably more reliable.

Location-specific parameters. All parameters except $\{\beta_i, \psi_i\}_i$ are common to all locations. We shrink the location-specific values of β_i and ψ_i towards common values by modeling them as random effects with $\beta_i \sim LogNormal(\bar{\beta}, \sigma_\beta^2)$ and $\psi_i \sim LogNormal(\bar{\psi}, \sigma_\psi^2)$.

Initial conditions. We initiate the virus for county j at the date when a positive number of confirmed cases first occurs in the data. At that date, we set $I(t) = NewCases/Pop$ and $A(t) = \mathbf{f}I(t)$, where \mathbf{f} is a function of other parameters (see Appendix C.1).²¹

Likelihood. Accounting for the DGPs and random effects, the likelihood takes the form

$$L(\theta|X, B) \propto p(X|\theta, B) \prod_i p(\beta_i|\bar{\beta}, \sigma_\beta^2) p(\psi_i|\bar{\psi}, \sigma_\psi^2). \quad (36)$$

A number of journalistic outlets have documented evidence of “excess deaths” in multiple locations, which exceed the number of reported COVID-19 deaths in those locations. To the extent that official death totals are meaningfully under-reported, the true mortality rate in our model (κ) should be higher and readers of this would want to accordingly adjust upwards our mortality predictions.

²⁰The CMP distribution of Conway and Maxwell (1962) is also known as the COM-Poisson distribution. See Sellers et al. (2012) and Daly and Gaunt (2016) for useful recent overviews of the distribution's features.

²¹At the initialization date, $R(t) = D(t) = 0$ and $S(t) = 1 - I(t) - A(t)$.

We assume the measurement errors are idiosyncratic in the sense of being independent across time, location, and data series, thus $p(X|\theta, B)$ decomposes as $p(X|\theta, B) = \prod_{i,t} p(\mathbf{x}_{i,t}|\theta, B_{0:t})$ where

$$p(\mathbf{x}_{i,t}|\theta, B_{0:t}) = \phi(\text{Activity}_{i,t}|a_{i,t}^{\theta, B_{0:t}}, \sigma_a^2/w_i) \cdot \phi(\text{Output}_{i,t}|y_{i,t}^{\theta, B_{0:t}}, \sigma_y^2/w_i) \cdot g(\text{NewCases}_{i,t}|\lambda_C, \nu_C) \cdot g(\text{NewDeaths}_{i,t}|\lambda_D, \nu_D) \quad (37)$$

for $\phi(\cdot|\mu, \varsigma)$ the pdf of a Normal distribution with mean μ and variance ς , $g(\cdot|\lambda, \nu)$ the probability mass function of a CMP distribution, and λ_C and λ_D set as described above. We estimate θ as the argmax of (the log of) L in (36).

Discussion. A few aspects of the estimator warrant further discussion. First, the simulated trajectories from the model might be interpreted as impulse responses to a “COVID-19 shock.” The entirety of the paths reflect a single COVID-19 infection shock that perturbs the measures of $S(t)$ and $I(t)$ in each county at the beginning of the pandemic and the policy paths in each county. In other words, the paths have no mechanism for correcting mid-course deviations from the data. In the language sometimes used when assessing DSGE models, one could say that our model fit results exclusively from *internal* propagation mechanisms following the infection impulses in early March.

Second, for the parameters common across locations, the estimator prioritizes fitting the dynamics of higher population areas. The economic variables are normalized county-by-county, but the weights w_i in (33) decrease the relative measurement error as county population increases. A similar relationship is implicitly induced by the CMP distribution for the health variables. One can show that increasing the mean of the CMP while holding ν fixed at a given value results in probability mass more concentrated around the distribution’s center. Thus the relative size of the idiosyncratic variation falls with higher expected counts, which, all else equal, occur in higher population areas. The greater “weight” on matching the dynamics in large population counties helps to ensure good fit to aggregates.

5.3 Fit

We show the estimated Econ-SAIRD model’s success at fitting the data along multiple dimensions, beginning with the time-series of aggregates and period-by-period cross section.

Time series. The two panels of Figure 5 compare the data and model predictions for the four variables of interest. In light of this fact, we think the model’s ability to generally

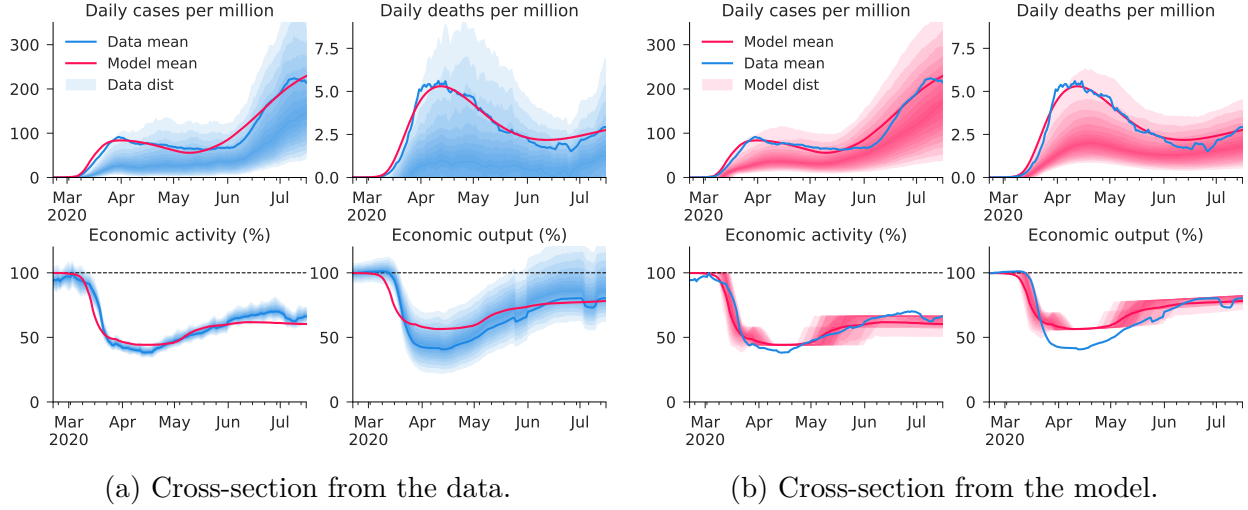


Figure 5: Data vs model: time series and cross sectional distributions.

match the paths of all 4 variables over a roughly 5 month period, including case resurgences in multiple states, based on only initial conditions is fairly remarkable.

Cross sectional units. While Figure 5 demonstrates the model’s fit to the aggregate time series, it is also important to show that the model accurately captures the time series evolution at lower levels of aggregation.²² Figure 6 compares the model and data for two states of particular interest: Florida and California. Appendix F contains analogous figures for each of the 44 states in our sample.²³

Seroprevalence. Since the model assumes underreporting of cases in the data, the model predicts that true cases always exceed the reported case data. Using seroprevalence studies we can quantitatively validate the model’s view of the number of unreported COVID-19 cases. This is an especially useful validation of the model mechanisms, as our estimation does not target these quantities at the few times and locations for which they are available.

The CDC reports the results from 10 local seroprevalence studies, each conducted at two different points in time.²⁴ Because our model provides daily county-level predictions for COVID-19 exposure, we can generally construct the model’s assessment of exposure in the same location, and at the same dates, as each seroprevalence study. The first panel of

²²We would not want to get the average right (in the form of aggregates) while getting all individual units very wrong.

²³Our data does not include any counties in six states because of lack of coverage in the hours worked data.

²⁴See <https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/commercial-labs-interactive-serology-dash.html>.




Figure 6: State-level fit. Solid black line indicates stay-at-home order, dashed black line indicates restaurant take-out-only order, dotted black line indicates school closure.

Figure 7a shows the model’s aggregate implications for true cases vs reported cases. The second panel shows the model’s implications for exposure in the locations and on the dates of the available seroprevalence studies.²⁵ We find this a particularly interesting dimension of the model to evaluate because our estimation procedure does not “target” or in any way use these features of the data. Further, given the uncertainty over what fraction of past infected people demonstrate antibodies, the fact that the model slightly over-estimates the data point estimates is reassuring.

6 Counterfactuals and Policy Experiments

In this section we examine the quantitative implications of the estimated Econ-SAIRD model for the trajectories of epidemiological and economic variables under different policy scenarios.

Eradication Zone. Under some policies, the model’s implications depend importantly on whether or not it is possible to eradicate the virus by driving the number of carriers sufficiently low. An artifact of standard SIR-based models is that it is not possible to eradicate a virus until the population reaches herd immunity. This is because some measure of agents

²⁵The values are taken from the CDC’s Commercial Laboratory Seroprevalence Survey Data. The CDC offers the following broad summary of the nature of the underlying data, “The survey includes people who had blood specimens tested for reasons unrelated to COVID-19, such as for a routine or sick visit during which blood was collected and tested by commercial laboratories in participating areas from each of the 10 sites.” As such the data collection should not reflect COVID-19 cases that are either patently symptomatic or have already resolved via mortality.

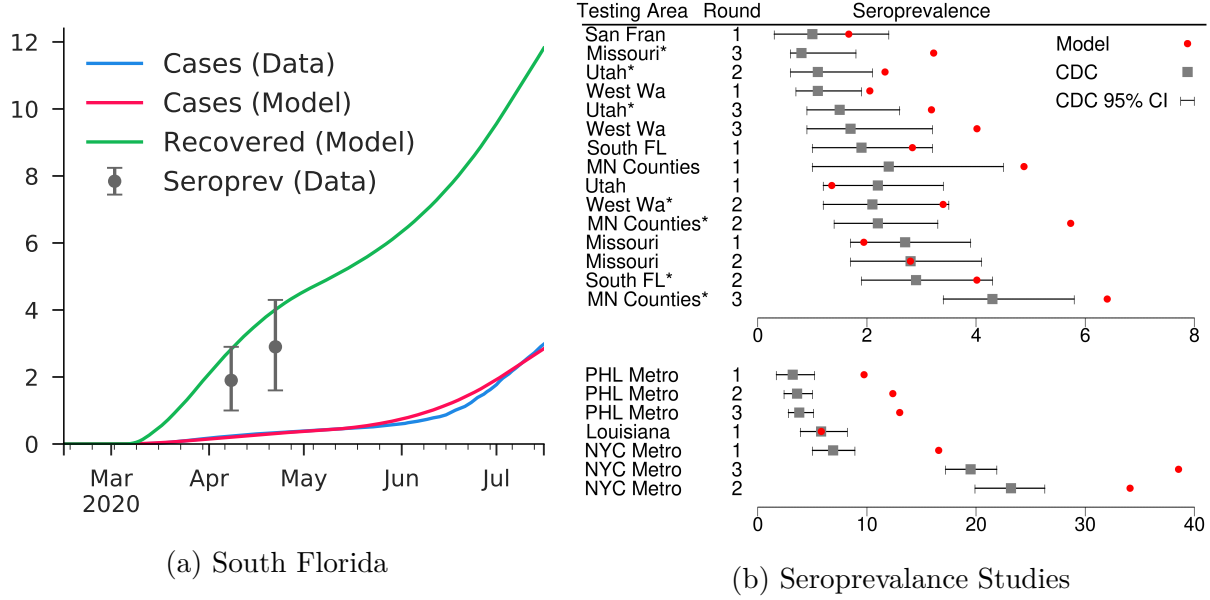


Figure 7: Seroprevalence in the model and data.

remain infected forever in this class of models and, as long as a large share of Susceptibles remains, the virus will propagate in the presence of any positive measure of infectious agents. We will consider a version of the model where $A(t)$ and $I(t)$ jump discontinuously to 0 if the total measure of infectious agents ($A(t) + I(t)$) falls below a particular threshold. At which point, we say that the virus has been eradicated.

When using this variant of the model, we entertain two sizes for this “Eradication Zone,” one at smaller (and hence weaker) one at 100 per million and a stronger one at 400 per million. Under reasonable asymptomatic transmission and case reporting assumptions, these $A + I$ numbers correspond very roughly to reported new case rates of 1 and 5 per million per day, respectively. Both of these zones are still fairly stringent, as daily reported case number in the US have been between 100 and 200 per million for some time. They are closer to numbers seen in S. Korea which are typically around 1 per million but spike up to 5-10 per million during small outbreaks.

6.1 Laissez-Faire

We first examine the scenario in which policymakers do nothing, i.e. laissez-faire policy. This scenario demonstrates the behavior that the model absent NPIs simply from the endogenous responses of Susceptibles with respect to the force of infection. As in other SIR-based models, the virus continues to spread until the population reaches herd immunity. As discussed in Section 2, infections peak earlier than the herd immunity threshold — at the point at which




Figure 8: **Laissez-Faire.** Red lines are population-weighted means, thick blue lines are for the median county, and the light blue shaded bands indicate percentiles of the cross sectional distribution of counties.

the inflow rate of new infections equals the outflow rate of infectious agents. In the Econ-SAIRD model achieving herd immunity in the full cross-section requires about 71 percent of the population having been infected. From a policy perspective, the important questions about the laissez-faire option are how long it takes to get to herd immunity, what are the costs of doing so, and, perhaps most importantly, what are the other options?

Figure 8 shows the long term trajectories of economic and epidemiological variables under laissez-faire policy in the Econ-SAIRD model. The critically different implication of our Econ-SIR-based model, relative to a standard SIR model, is the protractedness of the path back to normality. The endogenous virus avoidance of Susceptible agents, who reduce activity even in the absence of formal lock downs, generates a path of roughly 3 years to reach herd immunity, after which infections steadily die out over the subsequent 2 years. Along this path, activity and output remain significantly below their pre-pandemic levels. Our




Figure 9: Forward simulations under alternative policies.

results crystallize the notion that “reopening” is easier said than done. As long as people fear infection from the virus, they will continue to curtail their economic activity and output will remain below normal levels. As the model approaches herd immunity, economic activity increases and approaches its baseline level, ultimately yielding dynamics similar to those in models with a more simplistic infection mechanism, albeit occurring at a considerably later date in Econ-SIR. As this process plays out, cumulative mortality over the five years is about 0.38 percent of the population, which equates to roughly 1.25 million deaths in the United States.

6.2 Alternative Policies

With the model’s baseline behavior as context, we next consider the outcomes under various policies. In each case, we assume the policy is in effect from July 12, 2020 forward.

- **Status quo.** The status quo option entails the continuation of whatever NPI policies were in place at the end of our estimation period.

- **Stay-at-home indefinitely.** (Self explanatory).
- **Severe temporary lockdowns, followed by full reopening.** It has been suggested that a severe temporary lockdown would be an effective way to reduce U.S. infection levels to the point where test-and-trace strategies would be viable for containment. We consider the implications of such a severe shutdown if it were followed by simply a naive reopening.
- **Optimal adaptive local lockdowns.** We next consider optimal policy rules that enforce location-specific lockdowns of varying strength as a function of the local share of infectious agents.

Figure 9 shows the predicted trajectory of our benchmark variables under laissez-faire, status quo, and perpetual stay-at-home. From these trajectories one can see the main trade-offs that the model allows. While laissez-faire generates the highest output (even despite some degree of private mitigation by agents), it does so at the cost of vastly more infections than the other options over the near term. At the other end of the spectrum, our estimated stay-at-home orders would be sufficient to dramatically curtail the virus’s spread, they do at the cost of also dramatically curtailing economic output.

The long run implications of a temporary severe lockdown (if followed by a full reopening), depend critically on whether or not the eradication zone is “fair game.” Figure E.4 in the appendix shows the model outcome given a fairly stringent lockdown lasting for two months. In the short run, the lockdown is effective at reducing infection rates at the cost of reducing output by about 50 percent. If the eradication zone exists, then this is an extremely effective long run strategy for knocking out the virus and fully normal activity resumes within 4 months. However, with no eradication zone, infections come surging back after only a few months. On the economic side, there is a brief period of high output at the early phase of the virus resurgence, followed by a subsequent collapse due to renewed avoidance.

We next consider optimal policy rules that enforce location-specific lockdowns of varying strength as a function of the local share of infectious agents. We optimize a policy rule of the form

$$\bar{z}_{j,t} = \bar{z}^{\max} \cdot \Lambda(A_{j,t-1} + I_{j,t-1}) \quad (38)$$

where $\Lambda(s)$ is the smoothstep sigmoid function with left edge of 0 and right edge of \bar{s} . The parameters to optimize in (38) are the maximal lockdown intensity, \bar{z}^{\max} , and the value of \bar{s} of local virus prevalence that triggers the maximal lockdown. For each simulation, we

compute welfare via

$$W(t_0, T) = \sum_{t=t_0}^T \sum_{j=1}^J (y_{j,t} - \psi_j N_{j,t}). \quad (39)$$

and determine $(\bar{z}^{\max}, \bar{s})$ jointly as

$$(\bar{z}^{\max}, \bar{s}) = \operatorname{argmax}_{\bar{z}^{\max}, \bar{s}} W(t_0, T) \quad (40)$$

where t_0 is the end of the estimation sample, and T is 365 days after t_0 . For each location, the model-implied trajectories of output $y_{i,t_0:\bar{T}}$ and new infections $N_{i,t_0:\bar{T}}$ depend on the sequences $\bar{z}_{i,t_0:\bar{T}}$ generated by the policy rule. We will invoke the welfare calculation in (39) in additional experiments going forward. W is our preferred measure for comparing alternative policies, as it is the most internally-model-consistent summary of how agents value trajectories of the model's variables. Nevertheless, to see the model-implied tradeoffs in possibly more concrete dimensions, we also summarize the ramifications of alternative policies in terms of output and deaths. One can then form judgments about a given policy's appropriateness without fully accepting the primacy of W .

In the presence of a weak eradication zone (Figure 10), we estimate that the optimal policy takes the form of a lockdown level $\bar{z}^{\max} = 1.05$ and $\bar{s} = 212$ infectious agents per million. Meanwhile, with a strong eradication zone, we find the optimal policy to be $\bar{z}^{\max} = 1.08$ and $\bar{s} = 883$. To make the implications more concrete, such a policy (when active in a given location) would restrict economic activity to a level somewhat lower than that of our estimated stay-at-home order. The threshold for introduction is based on total active cases per million. Mapping this into daily reported cases per million, we can divide by about 4 to account for asymptomatic cases and under-reporting and divide by 20 ($\approx 1/\delta$) to go from a stock to a daily flow (so a total of 80). In this case, the thresholds map to 2.7 (for 212) and 11 (for 883).

Within the model, the nature of the tradeoff between deaths and output depends importantly on both the time horizon and the eradication zone. We follow others in the literature and refer to the set of achievable outcomes for deaths and output as the pandemic possibilities frontier, or PPF. Figure 11 plots the PPFs pertinent to 1 month and 1 year horizons under optimal adaptive lockdown with and without the eradication zone assumption. Without the eradication zone, (top row) we find classical tradeoffs between output and deaths. Note that the lower right most point in each PPF pertains to laissez-faire policy. If eradication is possible (bottom row), then the estimated model implies a PPF with an upward

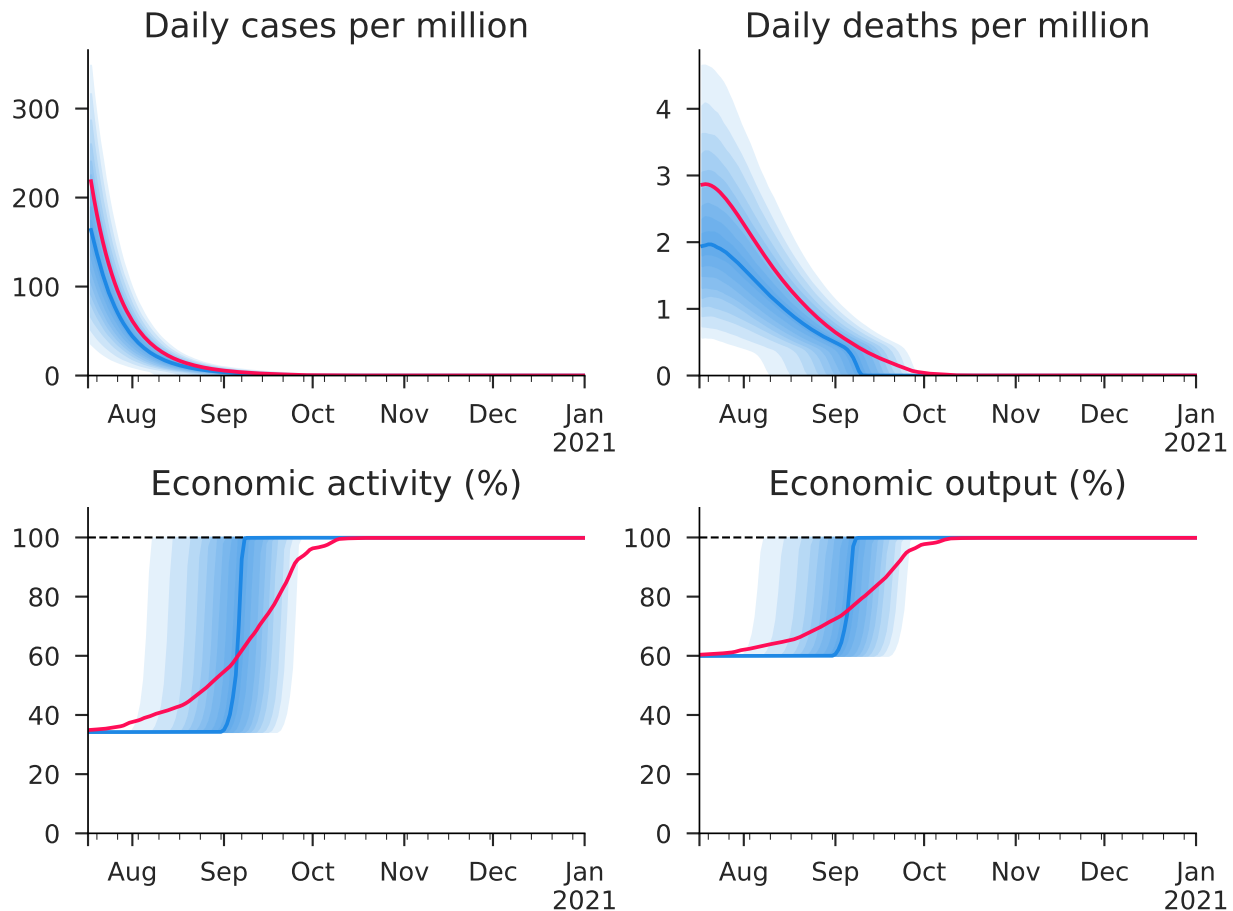


Figure 10: Optimal adaptive lockdown policy, with eradication zone.

sloping region at a one year horizon. This is achievable because once the virus is eradicated, output returns to normal levels, which can in principle outweigh sharp initial drops in output associated with a temporary stringent lockdown to eradicate the virus. It turns out, that such an approach yields gains in both dimensions, generating the upward sloping portion of the PPF.

6.3 Optimal Policy when a Vaccine is Near

The previous section's results underscore the fact that there is little scope for NPIs to decrease the virus's final toll if eradication is impossible and a vaccine never arrives. In this section we consider optimal policy when a vaccine is likely in the relatively near term — one year from now. The question then becomes, would any NPIs be worthwhile over such a time horizon?

One can develop the basic intuition for why NPIs could yield welfare gains under such a scenario by re-examining the laissez-faire path in Figure 8. Along the laissez-faire path, cases

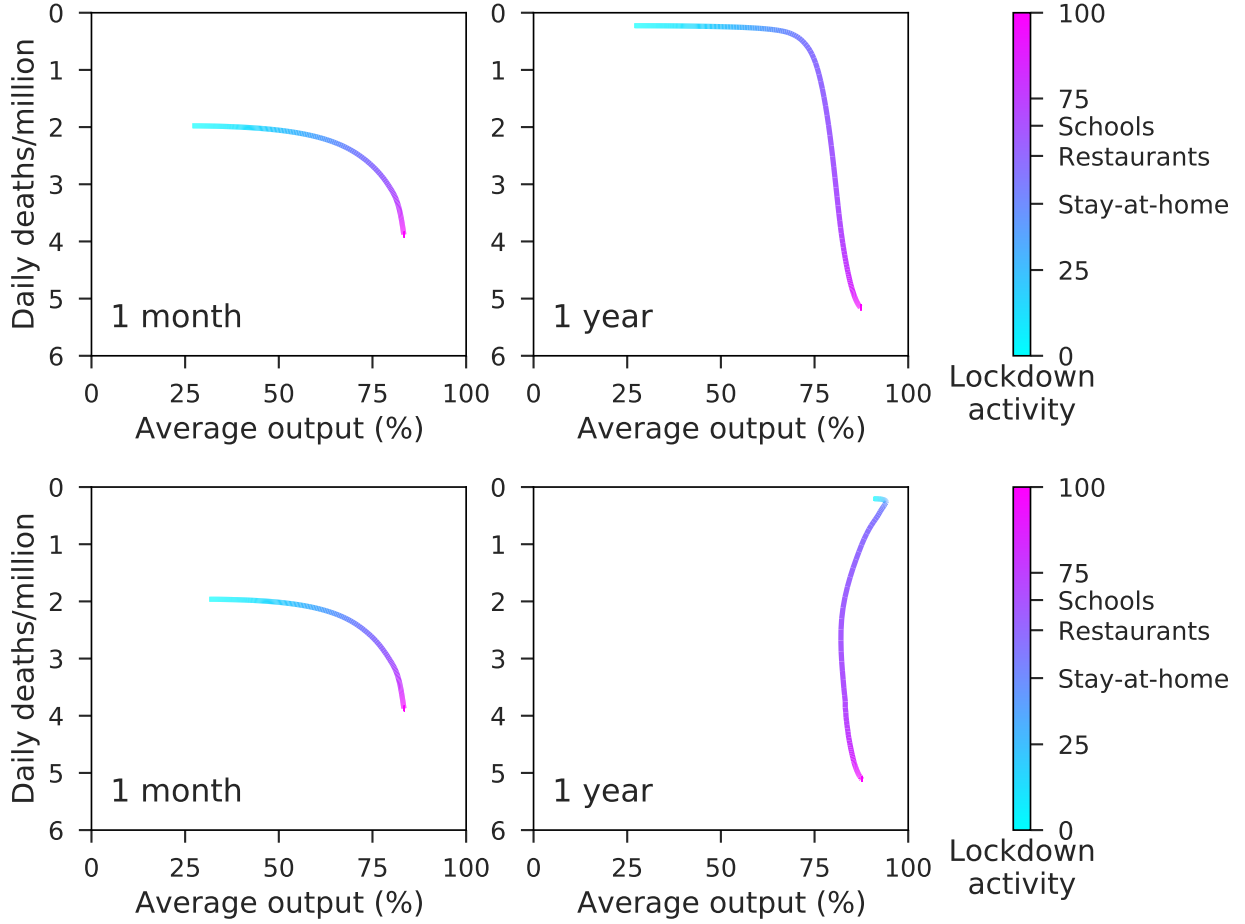


Figure 11: Forward looking PPF under adaptive lockdown policy. Top row of panels assumes no eradication zone exists. Bottom row of panels assumes a strong eradication zone.

and deaths are frontloaded. With a vaccine ready at time T_{vax} , we know that the virus's total cost beyond T_{vax} will be small regardless of how many infections have occurred by that date. In that case, it could conceivably be optimal to prevent the burst of infections and deaths that occur in the pandemic's early stages since the vaccine renders it unnecessary ever to incur them at all. One can contrast this with the scenario in which no vaccine could ever be developed, in which case simply shifting the timing of the costs (i.e. whether negatives are frontloaded or backloaded) would be largely irrelevant (for a sufficiently patient planner).

In particular we compute an optimal policy of the form of (40), tailored to the case of a vaccine 1 year away, and beginning from the observed state of the world in mid-July 2020. Figure 12 shows the trajectories of the four benchmark variables under the optimal adaptive NPI policy when a vaccine is one year away.²⁶ The resulting trajectories of cases and activity

²⁶We also considered the cases of 6 months and 2 years, which turned out to give results very similar to the 1 year case presented here. However, if it takes more than 4 years for a vaccine to arrive the optimal




Figure 12: Optimal adaptive lockdown with vaccine 1 year away.

are particularly informative for understanding the form of the policy.²⁷ In light of the initial stock of infectious agents, the policy begins by imposing a considerable reduction in activity to levels similar to our estimated effects of stay-at-home orders (activity slightly above 40 percent, output just below 70 percent). This phase remains in effect in most locations for 2 to 4 months, depending on the initial stock of infectious agents, during which time the activity reduction is great enough to dramatically reduce new infections and the stock of infectious agents. After this period, the optimal NPI loosens modestly, but is not removed entirely. The policy then prescribes a roughly 8 month period of reduced activity at levels similar to our estimated restaurants-takeout-only orders (activity around 60 percent, output around 80 percent). Having already reduced the stock of infectious agents, the somewhat higher activity level remains sufficient to keep the effective reproduction number around 1, and policy again resembles the laissez-faire equilibrium.

²⁷When the vaccine arrives, we model it as being administered to 2 million people per day, or 0.6 percent of the population. In the model, vaccinated agents are a mass of $S(t)$ shifted directly to $R(t)$.




Figure 13: Period by period costs of the pandemic under various policy scenarios.

thus keep cases from-reentering an exponential growth regime. Once the vaccine arrives and is distributed, the measure of Susceptibles begins to fall and NPIs can be safely progressively weakened while maintaining low levels of new infections. Finally, activity and output would return to normal in the fourth quarter of 2021.

6.4 The Costs of the Pandemic

Figure 13 shows the welfare cost of the pandemic, from July 2020 onward, under a variety of scenarios. Lower values in the figure indicate higher welfare. So that one might see the way different policies tradeoff the two welfare channels, we decompose the total cost along each path into the cost of infections and cost of lost output. Relative to the other policy options we consider, the key feature of the laissez-faire policy is that its relatively modest output losses are outweighed by its large disease cost, provided that a vaccine is on the horizon or eradication is possible. Conversely, the other options typically entail somewhat larger output costs, particularly early on, but with the benefit of large reductions in total infections.

7 Conclusion

In this paper we developed and estimated a spatial model of the joint evolution of economic variables and the spread of an epidemic across U.S. counties using high-frequency granular data. The model predicts a significant endogenous reduction in economic activity by agents in response to the spread of the virus — highlighting the importance of using an integrate-assessment model for evaluating policy responses to epidemics. Absent pharmaceutical advances to vaccinate against or treat the virus, our estimated model predicts a protracted march towards herd immunity. Infections and deaths continue to increase roughly

linearly for the next 3 to 5 years, with an ultimate death toll of around 1,250,000.

Lockdowns effectively lower the spread of the virus at the cost of lower economic output — giving rise to a pandemic possibility frontier. From a welfare perspective, however, lockdowns are desirable only under two conditions. First, if the policymaker expects a vaccine or cure to arrive within the next 2-3 years, lives saved from locking down outweigh the output costs. Second, suppose there is the possibility to curtail community transmission (either through eradication or management through testing, tracing, and quarantine). In that case, a strict lockdown for a few months can affect a quick return to pre-pandemic levels of activity. Policies that lower the rate of transmission (e.g., mask wearing), however, can successfully raise output and save lives.

One caveat to our conclusions is that we have assumed that individuals receive permanent immunity to the virus after recovering. Recent medical evidence suggests that reinfection is rare but possible (albeit in a milder form). While beyond the scope of the current analysis, temporary immunity would likely increase the welfare benefits of both finding a vaccine and curtailing community transmission through NPIs.

References

- Adda, J. (2016). Economic activity and the spread of viral diseases: Evidence from high frequency data. *The Quarterly Journal of Economics* 131(2), 891–941.
- Alvarez, F. E., D. Argente, and F. Lippi (2020). A simple planning problem for covid-19 lockdown. Technical report, National Bureau of Economic Research.
- Baker, S. R., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2020). How does household spending respond to an epidemic? consumption during the 2020 covid-19 pandemic. Technical report, National Bureau of Economic Research.
- Berger, D. W., K. F. Herkenhoff, and S. Mongey (2020). An seir infectious disease model with testing and conditional quarantine. Technical report, National Bureau of Economic Research.
- Brotherhood, L., P. Kircher, C. Santos, and M. Tertilt (2020). An economic model of the covid-19 epidemic: The importance of testing and age-specific policies.
- Conway, R. W. and W. L. Maxwell (1962). A queuing model with state dependent service rates. *Journal of Industrial Engineering* 12(2), 132–136.
- Daly, F. and R. E. Gaunt (2016). The conway-maxwell-poisson distribution: Distributional theory and approximation. *ALEA - Latin American Journal of Probability and Mathematical Statistics* 13, 635–658.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.
- Farboodi, M., G. Jarosch, and R. Shimer (2020). Internal and external effects of social distancing in a pandemic. Technical report, National Bureau of Economic Research.
- 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.
- Fernández-Villaverde, J. and C. I. Jones (2020, May). Estimating and simulating a sird model of covid-19 for many countries, states, and cities. Working Paper 27128, National Bureau of Economic Research.

- Glover, A., J. Heathcote, D. Krueger, and J.-V. Ríos-Rull (2020). Health versus wealth: On the distributional effects of controlling a pandemic. Technical report, National Bureau of Economic Research.
- Goolsbee, A. and C. Syverson (2020). Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. Technical report, National Bureau of Economic Research.
- Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? Technical report, National Bureau of Economic Research.
- Hagedorn, M. and K. Mitman (2020). Corona policy according to hank.
- Hur, S. (2020). The distributional effects of covid-19 and mitigation policies. *Globalization and Monetary Policy Institute Working Paper* (400).
- Jones, C. J., T. Philippon, and V. Venkateswaran (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical report, National Bureau of Economic Research.
- Kaplan, G., B. Moll, and G. Violante (2020). Pandemics according to hank. *Powerpoint presentation, LSE* 31.
- Keppo, J., E. Quercioli, and L. Smith (2019). The behavioral sir model, with applications to the swine flu epidemic. Technical report, University of Wisconsin.
- Kermack, W. and A. McKendrick (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society London A* 115(772), 700–721.
- Kniesner, T. J. and R. Sullivan (2020). The forgotten numbers: A closer look at covid-19 non-fatal valuations. *Available at SSRN* 3680348.
- Krueger, D., H. Uhlig, and T. Xie (2020). Macroeconomic dynamics and reallocation in an epidemic. Technical report, National Bureau of Economic Research.
- Kurmann, A., E. Lale, and L. Ta (2020). The impact of covid-19 on u.s. employment and hours: Real-time estimates with homebase data. Technical report.
- Liu, L., H. R. Moon, and F. Schorfheide (2020, May). Panel forecasts of country-level covid-19 infections. Working Paper 27248, National Bureau of Economic Research.

- Lubik, T. A. and C. Matthes (2020). How to go viral: A covid-19 model with endogenously time-varying parameters.
- Mendiola, S., O. Stavrunova, and O. Yerokhin (2020). Determinants of the community mobility during the covid-19 epidemic: the role of government regulations and information. Technical report.
- Mitman, K., S. Rabinovich, et al. (2020). Optimal unemployment benefits in the pandemic. Technical report, CEPR Discussion Papers.
- Morin, B. R., E. P. Fenichel, and C. Castillo-Chavez (2013). Sir dynamics with economically driven contact rates. *Natural resource modeling* 26(4), 505–525.
- Nordhaus, W. D. (1993). Rolling the “dice”: An optimal transition path for controlling greenhouse. *Resource and Energy Economics* 15, 27–50.
- Russell, T., J. Hellewell, C. Jarvis, K. Van-Zandvoort, S. Abbott, R. Ratnayake, S. Flasche, R. Eggo, A. Kucharski, C. nCov working group, et al. (2020). Estimating the infection and case fatality ratio for covid-19 using age-adjusted data from the outbreak on the diamond princess cruise ship. medrxiv. 2020 mar 5. DOI 10(2020.03), 05–20031773.
- Sellers, K. F., S. Borle, and G. Shmueli (2012). The com-poisson model for count data: a survey of methods and applications. *Applied Stochastic Models in Business and Industry* 28(2), 104–116.
- Squire, R. (2019). What about bias in your dataset? quantifying sampling bias in safe-graph patterns. Technical report.
- Toxvaerd, F. (2020). Equilibrium social distancing. *Cambridge-INET Working Paper Series* (2020/08).
- Verelst, F., L. Willem, and P. Beutels (2016). Behavioural change models for infectious disease transmission: a systematic review (2010–2015). *Journal of The Royal Society Interface* 13(125), 20160820.

A County-Level Data at Daily Frequency

For the empirical analysis and model estimation, we build a panel of daily county-level observations on two economic variables, two measures of virus spread, and NPIs. While the economic variables can be measured at even finer geographic granularity, we aggregate each of them to the county level to match the finest available granularity of the epidemiological data. We use daily observations from January 1, 2020 through May 24, 2020.

On the epidemiological side, we obtain daily county-level data on confirmed and probable COVID-19 cases and deaths from the *The New York Times*.²⁸ The *The New York Times* collects this data on a recurring basis from local public health authorities.²⁹ This is with one exception: the NYT data only records deaths for NYC as a whole, although the Bronx and Queens are distinct counties. We collect the county-level observations directly from New York City’s website and add them to our data set.

Our high-frequency economic variables consist of detailed data on both foot traffic and hours worked. We obtain data on foot traffic to over 3.6 million consumer “Points-of-Interest” (POIs), including grocery stores, restaurants, hospitals, and an array of other public and private establishments from SafeGraph Inc. SafeGraph collects the POI data as anonymized GPS location information from a large panel of more than 45 million smartphones in the U.S. We view this data as a highly accurate measure of a variable we will refer to throughout the paper as “economic activity” or just “activity.” Being derived from a panel of smartphones, the Safegraph data obviously does not represent a random sample of the U.S. population. However, Squire (2019) investigated potential sources of sampling bias in the SafeGraph data. The data are well representative at the county level (when compared to U.S. Census data from the American Community Survey) along a number of demographic dimensions, such as educational attainment and household income, although very low-income individuals (less than \$10,000 annual income) tend to be under-represented in the data.³⁰

We obtain hours worked from a large worker-firm matched data set provided by Homebase. Homebase is a scheduling and time clock software provider. The data consist of daily time-card records at the establishment level. The time-cards include information on hours worked and wages. The data are anonymized, but do specify the zip code and industry of the establishment. The data are provided in real-time when the time cards are reported

²⁸Data from *The New York Times*, based on reports from state and local health agencies. See <https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html>.

²⁹Note that it is unclear in the mortality data if the date corresponds to when the death occurred or when the death was officially reported. In our model estimation we account for this uncertainty by allowing for measurement error in recorded mortality outcomes.

³⁰The SafeGraph analysis does not address bias in age sampling, but we suspect that given that the panel is based on smartphone users that very old households would tend to be under-represented.




Figure A.1: Average Delay in COVID-19 Death Reporting

by the business that uses the software. As we detail in the appendix, small establishments and establishments in the food services and retail sectors tend to be overrepresented in the Homebase data. We refer the interested reader to Kurmann et al. (2020) for a more thorough discussion of the representativeness of the Homebase data and how it compares to the Quarterly Census of Employment and Wages conducted by the U.S. Census.

Finally, we assemble a detailed daily county panel for NPIs. To do so we use data from multiple sources, including data made available by other researchers as well as our own collection efforts using local and national news sources.

Reporting Delays. A few states publicly report the time series of assigned days of death, in addition to the more commonly available “new deaths.” We use the observed gap between the two series in some of these states to inform our choice of reporting delay during estimation. Figure A.1 shows how the average reporting delay has evolved over time in three such states.

NAICS	Firms HB	Firms QC	Wages HB	Wages QC	Weekly Wage HB	Weekly Wage QC
21		0.00		0.01		2,406
22		0.00		0.01		2,641
23		0.08		0.06		1,193
31		0.04		0.12		1,419
42		0.06		0.06		1,612
44	0.14	0.11	0.15	0.07	331	636
48	0.01	0.03	0.01	0.04	363	1,086
51		0.02		0.05		2,506
52		0.05		0.12		2,867
53		0.04		0.02		1,256
54	0.05	0.13	0.04	0.12	336	1,949
55		0.01		0.05		3,001
56		0.06		0.05		834
61	0.04	0.01	0.02	0.02	158	973
62	0.08	0.17	0.06	0.13	282	963
71	0.02	0.02	0.02	0.01	181	712
72	0.37	0.07	0.46	0.04	205	420
81	0.04	0.09	0.02	0.02	313	758
99	0.25	0.02	0.22	0.00	242	1,070

Table 1: Comparison of Homebase and Quarterly Census of Employment and Wages

B Econ-SAIRD Model Details

B.1 Epidemiological Environment

Epidemiological State Transitions. As in Econ-SIR, the S agents contract the virus at rate $n(t)$. The S category has no inflows from other states and hence the measure of Susceptibles decreases over time according to

$$\dot{S}(t) = -n(t)S(t). \quad (41)$$

We now assume that newly infected agents transition to the asymptomatic state. This feature of the model is consistent with reports that, when initially infected, agents are often asymptomatic but able to spread the virus for multiple days prior to developing symptoms. Asymptomatic agents transition to I at rate λ . Furthermore, since many virus carriers never exhibit symptoms at all, (Russell et al., 2020), we assume that agents in the A state can recover without ever developing symptoms, transitioning to R at rate γ . The measure $A(t)$

thus changes according to

$$\dot{A}(t) = n(t)S(t) - \lambda A(t) - \gamma A(t) \quad (42)$$

The measure $I(t)$ increases with the inflow of $\lambda A(t)$. As before, I agents recover at rate δ . We now also assume that I agents die at rate κ . $I(t)$ thus evolves according to

$$\dot{I}(t) = \lambda A(t) - \delta I(t) - \kappa I(t) \quad (43)$$

Lastly, the R and D states are absorbing states, with their measures only increasing over time, as they both have inflows but no outflows.³¹ With asymptomatic and infected agents recovering at rates γ and δ respectively, $R(t)$ increases at rate

$$\dot{R}(t) = \gamma A(t) + \delta I(t). \quad (44)$$

$D(t)$ grows with the deaths of I agents, and hence

$$\dot{D}(t) = \kappa I(t). \quad (45)$$

Given the parameters governing epidemiological state transition rates, λ, δ, κ , and γ , as well as the infection rate of Susceptibles, $n(t)$, equations (41)–(45) characterize the epidemiological dynamics in the economic-SAIRD model.

B.2 Activity and Output

Economic Activity. The rate of activity is dictated by the recovered agents and the active shares of Susceptibles, Asymptomatics, and Infecteds. One can then characterize the rate of activity either constructively, as in equation (46), or as a deviation from the full activity rate, as in equation (47):

$$a(t) = a_S(t)[S(t) + A(t)] + a_I I(t) + R(t) \quad (46)$$

$$= 1 - \underbrace{F(\bar{z}(t))}_{\text{rate of foregone excursions by } S \text{ \& } A} \cdot [S(t) + A(t)] - \underbrace{F(\bar{z}_I)I(t) - D(t)}_{\text{missing excursions by quarantined and deceased}} \quad (47)$$

³¹We are implicitly assuming that past infection to COVID-19 provides future immunity, at least over the time horizon relevant to the model estimation and usage.

Economic Output. Similarly, output can be equivalently characterized either constructively or in terms of shortfall from output when there are no infections:

$$\begin{aligned}
y(t) = & \overbrace{\mathbb{E}[z(t)]}^{\text{"no virus" output rate}} - [S(t) + A(t)] \cdot \overbrace{F(\bar{z}(t))}^{S(t) \text{ foregone excursion rate}} \cdot \overbrace{\mathbb{E}[z(t)|z(t) \leq \bar{z}(t)]}^{\text{value of foregone excursion}} \\
& - \underbrace{F(\bar{z}_I)I(t)\mathbb{E}[z(t)|z(t) \leq \bar{z}_I]}_{\text{foregone value by quarantined } I} - D(t) \cdot \mathbb{E}[z(t)]
\end{aligned} \tag{48}$$

B.3 Epidemiological Statistics

Once a person enters the Asymptomatic state, their transitions are stochastic but nonetheless exogenously determined by parameters. This makes it relatively straightforward to calculate standard epidemiological statistics. For instance, the case fatality rate (CFR), the number of deaths per confirmed cases, which in our model corresponds to the probability of dying conditional on reaching the Infected state is

$$CFR = \frac{\kappa}{\delta + \kappa} \tag{49}$$

Similarly, the infection fatality rate (IFR), the number of deaths per total cases, which (somewhat confusingly) corresponds to the probability of dying conditional on reaching the Asymptomatic state is

$$IFR = \frac{\lambda}{\gamma + \lambda} \cdot CFR = \frac{\lambda}{\gamma + \lambda} \cdot \frac{\kappa}{\delta + \kappa} \tag{50}$$

Next we turn to the various reproductive rates. Because there are many time varying quantities in our model (both from endogenous and exogenous sources), it is difficult to express in closed form how many additional individuals an infected person would expect to infect. However, we can answer that same question assuming conditions at that time prevailed in perpetuity. In this case we find

$$\mathcal{R}_e(t) = \mathcal{R}_0 \cdot e(t) \cdot a_S(t)S(t) \tag{51}$$

where we again have $\mathcal{R}_0 = \beta/\delta$. Notice that this figure reflects the fundamental reproductive number \mathcal{R}_0 as well as contributions from the excursions of Asymptomatic and Infected $e(t)$, the activity of Susceptibles $a_S(t)$, and the share of Susceptibles $S(t)$.

C Additional Estimation Details

C.1 Initial State

We do not observe the full underlying state of each county at the start of their outbreak. From the data, we know the initial number of confirmed cases (infected people) and deaths (generally zero), from which we can set I_0 and D_0 . However, we do not know the number of asymptomatic or recovered individuals. We assume that $R_0 = 0$ and that A_0 is proportional to I_0 , with $\mathbf{f}_0 \equiv A_0/I_0$ denoting the ratio between the two.

Identifying \mathbf{f}_0 from the data is challenging, so, rather than treat it as an additional free parameter, we derive a theory-driven estimate of it as a function of other model parameters. Letting $\mathbf{f}(t) \equiv A(t)/I(t)$, the law of motion for $\mathbf{f}(t)$ is

$$\frac{\dot{\mathbf{f}}(t)}{\mathbf{f}(t)} = n(t) \frac{S(t)}{A(t)} - \lambda - \gamma - \lambda \frac{A(t)}{I(t)} + \delta + \kappa \quad (52)$$

$$= \beta S(t) \left(a_A(t) + a_I \frac{1}{\mathbf{f}} \right) - \lambda - \gamma - \lambda \mathbf{f}(t) + \delta + \kappa \quad (53)$$

In the early stages of an outbreak, $S(t) \approx 1$ and $a_A(t) \approx a_S(t) \approx 1$. Further assuming $a_I \approx 0$ and imposing stationarity of $\mathbf{f}(t)$ ratio yields the estimate

$$\mathbf{f}_0 = \frac{\beta + \delta + \kappa - \lambda - \gamma}{\lambda}. \quad (54)$$

C.2 Parameter Estimates

Table 2 gives the estimates of the model parameters.

D A Historical Counterfactual: No May Reopening

The model makes a variety counterfactuals straightforward to compute.

In Figure D.2 we consider the possibility of continued lockdowns rather the reopenings implemented around the country in May.

Model Parameter	Value
$\mu[\beta]$	0.6167
$\sigma[\beta]$	0.3533
$\mu[\psi]$	354.5
$\sigma[\psi]$	224.1
λ	0.0527
γ	0.0827
δ	0.0460
κ	0.0006
σ	0.2893
f_0	10.024
z_i	1.7039
\bar{z}_1^L	0.8440
\bar{z}_2^L	0.9028
\bar{z}_3^L	0.9996
σ_a	0.1202
σ_o	0.2116

Table 2: Estimates of model parameters.




Figure D.2: **No May Reopening.** Counterfactual path in which lockdown continued indefinitely versus actual path in which reopening occurred.




Figure D.3: **Decomposing the Role of Technological Trends.**

D.1 The Roles of Mitigation and Activity Reduction

The estimated model allows for the possibility that agents have taken two types of precautions over time. First, agents have adopted measures to reduce virus transmission per unit of activity, modeled via declining β_t . The most salient real world example of such β -reducing measures is perhaps the wearing of face masks, but this mechanism in the model should be interpreted as also encompassing myriad other measures, such as the installation of plexiglass barriers at grocery stores, shifting restaurant dining outdoors, and simple six foot social distancing. Second, agents have increased their use of services that allow consumption with less infection risk, namely delivery options for a variety of goods, modeled via the z -substitute option. Here we consider the counterfactual in which these measures had not taken place.³² Figure D.3 shows the model's predicted evolution of the main variables in the absence of these trends.

³²This is straightforward to implement by setting $\beta_t = \beta_0$ for all t .

ONLINE APPENDIX

E Additional Figures




Figure E.4: Strong temporary lockdown with and without eradication zone.




Figure E.5: **Laissez-Faire until herd immunity—One year horizon.** Red lines are population-weighted means, thick blue lines are for the median county, and the blue shaded bands indicate percentiles of the cross sectional distribution of counties.




Figure E.6: Forward looking PPF under adaptive lockdown policy: weak eradication zone.

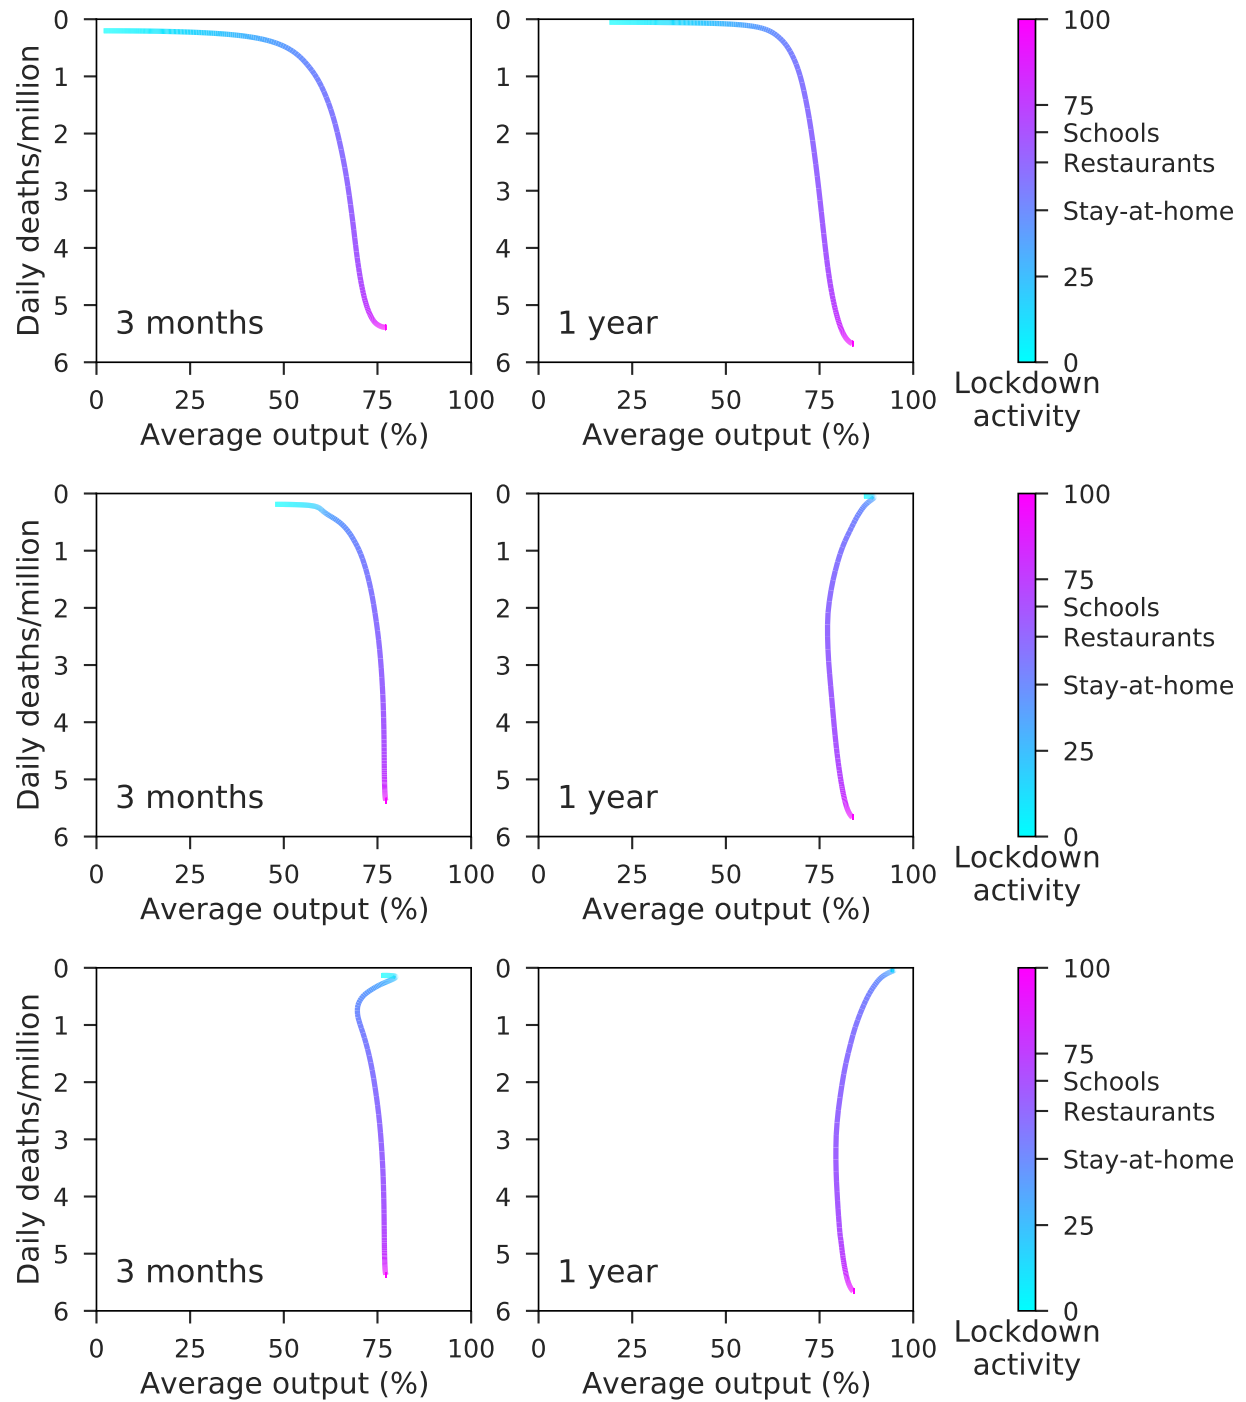




Figure E.7: Counterfactual PPF under adaptive lockdown policy: (a) without eradication, (b) with weak eradication zone, and (c) with strong eradication zone.

F State Fit


Alabama




Alaska




Arizona




Arkansas




California




Colorado




Delaware




District of Columbia




Florida




Georgia




Hawaii




Idaho




Illinois




Indiana




Iowa




Kansas




Kentucky




Louisiana




Maryland




Michigan




Minnesota




Mississippi




Missouri




Montana




Nebraska




Nevada




New Mexico




New York




North Carolina




North Dakota




Ohio




Oklahoma




Oregon




Pennsylvania




South Carolina




South Dakota




Tennessee




Texas




Utah




Virginia




Washington



West Virginia



Wisconsin



Wyoming

