



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Controlling Epidemic Spread: Reducing Economic Losses with Targeted Closures

John R. Birge, Ozan Candogan, Yiding Feng

To cite this article:

John R. Birge, Ozan Candogan, Yiding Feng (2022) Controlling Epidemic Spread: Reducing Economic Losses with Targeted Closures. *Management Science* 68(5):3175-3195. <https://doi.org/10.1287/mnsc.2022.4318>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Controlling Epidemic Spread: Reducing Economic Losses with Targeted Closures

John R. Birge,^a Ozan Candogan,^{a,*} Yiding Feng^b

^a Booth School of Business, University of Chicago, Chicago, Illinois 60637; ^b Microsoft Research New England, Cambridge, Massachusetts 02142

*Corresponding author

Contact: john.birge@chicagobooth.edu,  <https://orcid.org/0000-0002-7446-0953> (JRB); ozan.candogan@chicagobooth.edu,  <https://orcid.org/0000-0003-3920-402X> (OC); yidingfeng@microsoft.com,  <https://orcid.org/0000-0002-8258-6994> (YF)

Received: November 23, 2020

Revised: May 15, 2021; September 3, 2021

Accepted: September 22, 2021

Published Online in Articles in Advance: March 17, 2022

<https://doi.org/10.1287/mnsc.2022.4318>

Copyright: © 2022 INFORMS

Abstract. Data on population movements can be helpful in designing targeted policy responses to curb epidemic spread. However, it is not clear how to exactly leverage such data and how valuable they might be for the control of epidemics. To explore these questions, we study a spatial epidemic model that explicitly accounts for population movements and propose an optimization framework for obtaining targeted policies that restrict economic activity in different neighborhoods of a city at different levels. We focus on COVID-19 and calibrate our model using the mobile phone data that capture individuals' movements within New York City (NYC). We use these data to illustrate that targeting can allow for substantially higher employment levels than uniform (city-wide) policies when applied to reduce infections across a region of focus. In our NYC example (which focuses on the control of the disease in April 2020), our main model illustrates that appropriate targeting achieves a reduction in infections in all neighborhoods while resuming 23.1%–42.4% of the baseline nonteleworkable employment level. By contrast, uniform restriction policies that achieve the same policy goal permit 3.92–6.25 times less nonteleworkable employment. Our optimization framework demonstrates the potential of targeting to limit the economic costs of unemployment while curbing the spread of an epidemic.

History: Accepted by Carri Chan, healthcare management.

Supplemental Material: Data and the electronic companion are available at <https://doi.org/10.1287/mnsc.2022.4318>.

Keywords: COVID-19 • spatial epidemic models • network SEIR models • targeted interventions

1. Introduction

By November 2020, COVID-19 had infected more than 10 million individuals in the United States. To curb the spread of the epidemic, many cities instituted uniform (city-wide) suspension of economic activity to varying degrees. This contributed to unemployment, which reached a peak rate of 14.7% in the United States (CDC 2020, Congressional Research Service 2020). Similar effects were also observed elsewhere in the world.

The spread of the disease relies on human-to-human contact and has an inherent spatial nature: infected individuals potentially infect others in neighborhoods they visit. Thus, it may be possible to target some locations in a city with closures of public spaces and businesses so as to reduce the spread of the disease, while resuming economic activity elsewhere in the city. This could simultaneously reduce the economic losses and curb the spread of the disease. Determining how closures in targeted locations might achieve these goals depends on not only the health status and activity of residents in these locations but

also on their interactions with individuals throughout the region. Data on population movements can be useful in making these determinations. We provide a model that illustrates how such data can be leveraged and shed light on their value for controlling epidemic spread.

Targeted closure strategies have recently been employed by various cities/countries. Most existing strategies have not explicitly considered the economic consequences of closures. In addition, they have relied on local metrics but have not accounted for spillovers across locations and individual mobility as main drivers of disease spread. For instance, New York City (NYC) considered implementing targeted restrictions of economic activity in some zip codes and eventually settled on a more fine-grained microcluster-based strategy (New York State Governor's Office 2020b, c, d). Both strategies relied on metrics such as local positivity rate and targeted neighborhoods with high disease prevalence.¹ On the other hand, ignoring the economic consequences and the spillover effects due to population movements, as these strategies did,

can result in both higher rates of infection and levels of unemployment than can be achieved with a policy that recognizes these economic and health interactions.

To shed light on how data on population movements can help improve targeted interventions, we propose a spatial epidemic spread model that explicitly accounts for the spillovers of infections across different neighborhoods of a city. The model recognizes that individuals who reside in a neighborhood may spend some of their time in another neighborhood, for example, because of work or leisure. The aggregate amount of time that the members of each subpopulation spend in different neighborhoods determines (i) the amount of “mixing” between infected and susceptible agents in different neighborhoods and (ii) the rate with which the susceptible agents from a given neighborhood get infected. The model allows infected individuals from one neighborhood to infect residents of a different neighborhood whenever they spend time in a common neighborhood (possibly different from the neighborhood in which either group resides). We study the problem of a social planner who can restrict economic activity in different neighborhoods. In our model, reducing the permitted level of economic activity in a neighborhood triggers unemployment and decreases the number of individuals who visit that neighborhood. The latter effect reduces the infections among individuals who reside in that neighborhood as well as those who only visit it.

Our baseline model focuses on a static version of the problem, where the restrictions seek to ensure that in the short run the infections in all neighborhoods decrease. The social planner’s optimal policy is the one that achieves this policy goal while ensuring minimum unemployment. This model accounts for various practically relevant details, including differences in individuals’ trips (work versus leisure) and employment characteristics (teleworkable versus non-teleworkable employment). Yet it is simplified in at least three other practically relevant dimensions to clearly illustrate the key trade-offs resulting from closures. First, it is assumed that in response to the restrictions in one neighborhood, the individuals who would otherwise spend time there instead spend more time at home and effectively isolate. This is a simplification because these individuals could instead “substitute” their trips and visit other neighborhoods. Second, it may not always be feasible to target different neighborhoods in a city in a fine-grained way (e.g., at the zip-code level), and coarser targeting strategies may be needed. Third, dynamic policies are needed to control the spread of the disease over time. We first solve the baseline model and derive qualitative insights into the value and structure of targeted policies. These results demonstrate the significant potential that

suitably designed targeted policies have for reducing the economic costs of unemployment while controlling an epidemic.² Subsequently, we enrich the model to allow for different reactions to closures, coarser levels of targeting, and dynamics, using the solution of the baseline model as a building block to shed light on optimal targeted policies in these alternative settings.

We calibrate our baseline model and exemplify our approach using April 2020 infection levels of COVID-19 in NYC as well as cell phone data that are leveraged to represent population movements. We combine these data with disease-specific parameters for COVID-19 and obtain a representation of epidemic spread at the zip-code level. We also leverage data on workers’ abilities to work remotely and account for the impact of remote work on the spread of the epidemic. We then use our model to obtain an optimization framework that yields optimal targeting decisions. Within our framework assuming the aforementioned disease parameters and mobility choices, we show that appropriate targeting could achieve approximately a four- to six-fold increase in non-teleworkable employment relative to uniform policies that achieve the same infection-reduction goals. We also demonstrate that certain areas with high non-teleworkable employment and relatively low infection levels (e.g., Midtown Manhattan) may continue their economic activities while nearby regions with relatively less employment may face significant restrictions. Moreover, the optimal economic activity level in a target region heavily depends on the level of economic activity in the neighboring regions beyond the social planner’s control, highlighting the need for coordination among governing bodies (of different regions). In particular, in our example, the complete resumption of economic activity in neighboring counties can reduce the optimal activity level in NYC by 45.4%. We also explore variants of our baseline model that (i) require permitted economic activity levels in different neighborhoods to be similar, (ii) enable trading off between economic losses and faster reduction in infection rates, and (iii) allow for prioritizing low-risk businesses in each neighborhood. In all cases, we obtain similar insights and the gap between optimal targeted and uniform policies remains sizable.

We then examine alternatives to our baseline model along the three practically relevant dimensions highlighted earlier. Specifically, we (i) estimate a choice model that captures how trips may be substituted in response to closures, (ii) obtain coarse targeted strategies by clustering the solutions of our baseline model, and (iii) develop greedy dynamic heuristics for controlling the spread of the epidemic. For these alternatives, we once again illustrate that targeted restrictions in different locations can achieve the policy goal of reducing infections over time with substantially lower unemployment

than a policy with the same disease control goal that uses uniform restrictions throughout the region. The targeting structure in these extensions is also quite similar to that of the baseline solution.

Spatial spread of COVID-19 (and similar diseases) is a complex process, and there is still considerable uncertainty about it. Our objective is not to provide a definitive decision support tool for targeted interventions but rather to shed light on the value of targeting and to provide insights into its structure. We show that our findings are robust to various modeling assumptions that can impact the evolution of the epidemic, including (i) values of parameters governing epidemic spread, (ii) possibility of individuals' staying at home because of fear, and (iii) different individual mobility patterns and infection numbers. Moreover, we establish that even if the model is misspecified (which could be relevant, e.g., because of imperfect data), our targeted policies achieve substantially higher employment relative to their uniform counterparts without causing a significant increase in infections. We relegate our robustness checks to our electronic companion (EC) (see Section EC.3). In the EC, we also discuss further related literature (Section EC.1) beyond Section 1.1, present supplementary material on data sources and model calibration (Section EC.2), and present an alternative model that is appropriate for controlling a new emerging disease (Section EC.4).

1.1. Related Literature

Our paper contributes to the large literature on the control of epidemics. Since the first version of this paper, many other contemporaneous and subsequent works have studied interventions for curbing the spread of COVID-19. Here we review some of the works that are closest to ours. In Section EC.1, we include additional discussion on some other related directions.

1.1.1. Spatial Spread of Diseases. Spatial aspects of epidemic spread play an important role in prediction, estimation, and understanding of disease transmission (Anderson et al. 1992, McNeill and McNeill 1998, Allen et al. 2008, Martcheva 2015, Brauer et al. 2019). Different models have been proposed in epidemiology to analyze the spatial spread of diseases (Hethcote 1976, Hethcote and Van Ark 1987, Van den Driessche and Watmough 2002, Dingel and Neiman 2020). Spatial effects have received attention in the recent literature that studies COVID-19 as well. Using data on population flows (from Wuhan, China), Jia et al. (2020) propose a spatiotemporal "risk source" model to forecast confirmed cases and identify locales of high transmission risk. This model sheds light on the spread of COVID-19 and its growth pattern in China. Similarly, Chinazzi et al. (2020) use a metapopulation model and project the impact of travel limitations on the national and the international spread of the epidemic. We

use a spatial model of disease transmission similar to that of Post et al. (1983). Our main contribution to this literature is a novel framework that (i) allows for controlling the spread of an epidemic by spatially targeted restrictions and (ii) takes into account population movements and induced economic losses (through unemployment).

1.1.2. COVID-19 and Geographical Spillovers. Holtz et al. (2020) focus on the impact of shelter-in-place orders (impacting different states/counties) and conclude that these orders have a nontrivial impact on individual mobility. In addition, they show the importance of geographical spillovers and argue that interventions in one region cause a reduction in mobility in other regions. Carranza et al. (2020) provide evidence that reduction in mobility implies nontrivial reduction in infections. They also highlight that (targeted/localized) lockdowns lead to reduction in mobility.³ These findings are consistent with our model and results: we provide a framework for optimizing closures by taking into account individual mobility, a remote-work option, and spillovers across regions. Our model explicitly accounts for the impact of closures in one region on individuals in other regions (through mobility patterns) and, hence, provides a way of coordinating interventions in different locations. We also illustrate the need for coordination among neighboring counties and the impact of lack of coordination on the achievable employment levels. Ma et al. (2020) build on our framework and explore the question of targeted interventions in other geographical regions. In a subsequent work, Chang et al. (2020) also take into account population movements and confirm findings similar to ours in terms of the benefits of targeting. Pei and Shaman (2020) consider a metapopulation model at county resolution in the continental United States. In their model, they consider two types of movement: daily work commuting and random movement. They estimate the epidemiological parameters, use the calibrated model to project outbreaks in the continental United States, and evaluate the effects of social distancing and travel restrictions on the outbreaks. Gonzalez-Reiche et al. (2020) study the initial stages of the spread of COVID-19 in New York City. Unlike ours, these papers do not provide the optimal targeted decisions that take into account economic losses induced by closures.

1.1.3. Targeting Along Other Dimensions. Acemoglu et al. (2020) study a heterogeneous Susceptible-Infected-Recovered (SIR) model where different subpopulations (e.g., age groups) are susceptible to different risks (i.e., infection, hospitalization, and fatality rates vary between subpopulations). They consider the problem of a social planner who tries to control the spread of the disease by taking into account the induced economic losses as well as the disease-related deaths. They obtain

optimal age-based targeted policies and show that these policies significantly outperform uniform policies. Glover et al. (2020) consider a similar problem where the population is partitioned into age and sector groups, and a social planner chooses lockdown policies and how to redistribute income across the groups to optimize long-term social welfare. Evgeniou et al. (2020) propose an extension of the Susceptible-Exposed-Infected-Recovered (SEIR) model and study different exit policies using COVID-19 data and estimates for France as of early May 2020. Their study indicates that implementing an exit policy based on risk predictions could relax restrictions for millions of individuals in the lowest-risk population, whereas implementing the same exit policy without such predictions could exceed the hospital capacity or require long-term isolation for a substantial portion of the population for over a year in aggregate. Lipton and Lopez de Prado (2020) and Gershon et al. (2020) also focus on heterogeneous SEIR models with populations at different risk levels, and analyze how to manage the progression of the disease without exceeding the healthcare capacity or creating undue economic burden. Duque et al. (2020) propose an extension of the SEIR model and study how to design social-distancing policies for different age groups subject to a hospital capacity constraint. Chen et al. (2020) focus on an age-structured SAPHIRE (susceptible, exposed, presymptomatic infectious, ascertained infectious, unascertained infectious, isolated, removed) model and study how to allocate COVID-19 vaccines to different age groups subject to a limited supply constraint. Batlle et al. (2020) present a general framework for adaptive allocation of viral tests by taking into account the network of social contacts.

1.1.4. Economic Impacts on COVID-19. COVID-19 has already impacted economic activity around the globe immensely. Various economic impacts of the disease (e.g., on financial markets, fiscal policies, employment and wage, etc.) have been explored in the recent literature (Baker et al. 2020a, b, c; Coibion et al. 2020; Gormsen and Kojen 2020; Guerrieri et al. 2020; Hanson et al. 2020; He and Liu 2020; Mulligan 2021). Budish et al. (2020) discuss an economic policy response to the COVID-19 crisis and offer frameworks and guiding principles for dealing with the crisis. Budish (2020) proposes focusing on minimizing the social welfare loss while taking the appropriate steps to ensure that the disease is contained. The present paper adopts a similar approach and proposes targeted closures as a lever for achieving this goal.

2. Model

We first present our model of epidemic spread over n distinct subregions or neighborhoods $\mathcal{N} = [n] := \{1, \dots, n\}$ within a larger region, for example, a city. Each

neighborhood i has population N_i . We partition the population N_i into three subpopulations, given as N_i^A, N_i^B, N_i^C such that $N_i^A + N_i^B + N_i^C = N_i$, to capture different mobility effects from employment restrictions. Here N_i^A and N_i^B respectively denote the subpopulations of individuals with teleworkable and nonteleworkable employment and N_i^C denotes the subpopulation of individuals who are unemployed.

For $a \in \mathcal{A} = \{A, B, C\}$, we let $\{\tau_{ij}^a\}_{i,j \in \mathcal{N}}$ denote (unitless) parameters that capture the fraction of time (e.g., in a day) the individuals from subpopulation N_i^a spent in neighborhood j prior to the epidemic. Individuals can spend time in different neighborhoods either for work or for leisure. We distinguish between the two and let $\tau_{ij,W}^a$ and $\tau_{ij,L}^a$ respectively denote the work and leisure components of τ_{ij}^a , where $\tau_{ij}^a = \tau_{ij,W}^a + \tau_{ij,L}^a$. Note that $\tau_{ij,W}^a = 0$ for $i, j \in \mathcal{N}$, because individuals in subpopulations $\{N_i^C\}$ are unemployed.

To control the spread of the epidemic, we assume that economic activity (measured in our empirical simulations in terms of employment) in different neighborhoods can be restricted. Let $x_j \in [0, 1]$ denote the fraction of the baseline economic activity permitted⁴ at j . Individuals' trips to j depend on the permitted level of economic activity x_j there, and $\tau_{ij}^a(x_j)$ denotes the fraction of time individuals from subpopulation N_i^a spend in neighborhood j in aggregate (for leisure plus work) as a result of the restrictions. We assume that all individuals in subpopulations $\{N_i^A\}$ (who have teleworkable employment) will work from home during the epidemic.⁵ The leisure/work trips of the remaining subpopulations are scaled proportionally to the permitted level of economic activity in the target neighborhood. More precisely, we have

$$\tau_{ij}^a(x_j) = \begin{cases} x_j \tau_{ij,L}^a & \text{for } a = A, \\ x_j \tau_{ij,W}^a + x_j \tau_{ij,L}^a & \text{for } a = B, \\ x_j \tau_{ij,L}^a & \text{for } a = C. \end{cases} \quad (1)$$

We note that $\tau_{ij}^a(x_j) = \tilde{\tau}_{ij}^a x_j$, where $\tilde{\tau}_{ij}^A = \tau_{ij,L}^A$, $\tilde{\tau}_{ij}^B = \tau_{ij,L}^B + \tau_{ij,W}^B$, and $\tilde{\tau}_{ij}^C = \tau_{ij,L}^C$ for all $i, j \in \mathcal{N}$ and $a \in \mathcal{A}$. Using this notation the “effective population” of individuals who mix in neighborhood j and infect others/get infected there can be given as $\sum_{k \in \mathcal{N}} \sum_{a \in \mathcal{A}} N_k^a \tilde{\tau}_{kj}^a (x_j) = \sum_{k \in \mathcal{N}} \sum_{a \in \mathcal{A}} N_k^a \tilde{\tau}_{kj}^a x_j$.

We model the spread of the disease using a compartmental model. Specifically, we assume that each member of population N_i^a belongs to one of five “disease compartments” that correspond to populations $S_i^a, E_i^a, I_i^{c,a}, I_i^{sc,a}, R_i^a$. Here S_i^a denotes the population of susceptible individuals (who can get infected if they come in contact with an infected individual) and E_i^a denotes the population of individuals who are exposed to the disease but who are not infectious yet

(e.g., because the infection is in the incubation period). These individuals subsequently move to one of the two infected compartments. The first (second) of these compartments, population $I_i^{c,a}$ ($I_i^{sc,a}$), is referred to as the infected clinical (subclinical) population and captures the infected individuals who (do not) show symptoms. The individuals in both compartments are infectious but possibly with different infection rates. The final compartment, population R_i^a , captures the infected individuals who recover from the disease and become immune to it. They have no impact on the future progression of the disease. In this formulation, we do not model the disease-related deaths as they are a small enough fraction of the total population not to affect the rate of new infections. That said, our optimization framework below has the goal of reducing infections, which are proportional to deaths (where we assume a constant fatality to infection ratio).

In our model, susceptible individuals from a neighborhood i can become infected in neighborhood j because of contact with infectious individuals from another neighborhood k . The rate at which such infections take place is proportional to the fraction of time the susceptible individuals from i spend in j as well as the fraction of the effective population of j that is infected (see (2) and (3)). This model of disease evolution (without the restrictions that influence the time spent in different neighborhoods) is similar to the one in Post et al. (1983). The evolution of the epidemic is described by the following system of equations:⁶

$$S_i^{a'} = -\beta \sum_j S_i^a \underbrace{\tilde{\tau}_{ij}^a x_j}_{\substack{\text{Fraction of time} \\ \text{agents } \in S_i^a \text{ spend in } j}} \quad (2)$$

$$E_i^{a'} = \beta \sum_j S_i^a \tilde{\tau}_{ij}^a x_j, \left(\frac{\sum_k \sum_{b \in \mathcal{A}} I_k^{c,b} \tilde{\tau}_{kj}^b}{\sum_k \sum_{b \in \mathcal{A}} N_k^b \tilde{\tau}_{kj}^b} + \alpha \frac{\sum_k \sum_{b \in \mathcal{A}} I_k^{sc,b} \tilde{\tau}_{kj}^b}{\sum_k \sum_{b \in \mathcal{A}} N_k^b \tilde{\tau}_{kj}^b} \right) - \kappa E_i^a \quad (3)$$

$$I_i^{c,a'} = \rho \kappa E_i^a - \gamma I_i^{c,a} \quad (4)$$

$$I_i^{sc,a'} = (1 - \rho) \kappa E_i^a - \gamma I_i^{sc,a} \quad R_i^{a'} = \gamma I_i^{c,a} + \gamma I_i^{sc,a},$$

for $a \in \mathcal{A}$ and $i \in \mathcal{N}$. Here prime ($'$) denotes the time derivative and $\beta, \alpha, \kappa, \rho, \gamma$ are disease-specific parameters.⁷ The infectious rate κ captures how quickly the exposed individuals become infectious, and the recovery rate γ captures how quickly the infected individuals recover. Not all infections cause severe symptoms. The clinical rate ρ captures the fraction of infected individuals who show clinical symptoms. The subclinical cases have a lower transmission rate, captured by a discount

factor α . The transmission rate β captures how many effective contacts an individual has on a given day and how likely such a contact is to result in an infection, conditional on being with an infected individual.

In (2), the term $\tilde{\tau}_{ij}^a x_j$ captures the fraction of time individuals from S_i^a spend in j when the permitted economic activity level in j is x_j . The term within parentheses captures the fraction of the effective population of j that is infected (where we discount the subclinical populations by α to capture their lower infectiousness). The entire right-hand side of (2) is the rate at which the susceptible individuals in i move to the exposed compartment. The rest of the equations detail how these individuals subsequently move to the remaining compartments. Note that, because of the spatial structure of the epidemic spread, the restriction of economic activity in one neighborhood impacts the infections among individuals who reside there as well as among those who reside elsewhere but spend time in this neighborhood. As a result, in order to reduce the infections in neighborhood i , it may be critical to restrict activity in another neighborhood j where infections occur at a nontrivial rate and individuals from i spend substantial time.

We first focus on a static decision problem a social planner might face at any point in time, that of reducing the rate of infections, for example, protecting scarce healthcare resources from increased loads in all neighborhoods. Specifically, the planner chooses the permitted economic activity levels in neighborhoods $\mathcal{Z} \subset \mathcal{N}$. We assume that the permitted economic activity in neighborhoods outside the planner's control⁸ ($\mathcal{N} \setminus \mathcal{Z}$) is fixed at an exogenous level $x_i = y$. We further assume that the *policy goal* of the planner is to ensure a reduction in the total number of infected and exposed individuals in every neighborhood in \mathcal{Z} . More precisely, the rate of change in the total population of these individuals should be nonpositive *at the time the policy takes effect*,⁹ that is, $\sum_{a \in \mathcal{A}} E_i^{a'} + I_i^{c,a'} + I_i^{sc,a'} \leq 0$ for $i \in \mathcal{Z}$. Reducing the economic activity induces unemployment in \mathcal{Z} , denoted by $\sum_{i \in \mathcal{Z}} \tilde{c}_i (1 - x_i)$, where \tilde{c}_i represents the number of nonteleworkable jobs in neighborhood i . We characterize the targeted policies that maximize the number of individuals $\sum_{i \in \mathcal{Z}} \tilde{c}_i x_i$ who remain employed, while achieving the aforementioned policy goal. The optimal targeted restriction policy is obtained by solving the following linear program:

$$\begin{aligned} & \max_{\substack{x_i \in [0,1] \mid i \in \mathcal{Z} \\ x_i = y \mid i \in \mathcal{N} \setminus \mathcal{Z}}} \sum_{i \in \mathcal{Z}} \tilde{c}_i x_i \\ & \text{s.t. } \sum_j \sum_{a \in \mathcal{A}} S_i^a x_j \tilde{\tau}_{ij}^a \frac{\sum_k \sum_{b \in \mathcal{A}} (I_k^{c,b} + \alpha I_k^{sc,b}) \tilde{\tau}_{kj}^b}{\sum_k \sum_{b \in \mathcal{A}} N_k^b \tilde{\tau}_{kj}^b} \\ & \leq \sum_{a \in \mathcal{A}} \frac{\gamma}{\beta} (I_i^{c,a} + I_i^{sc,a}), \quad \forall i \in \mathcal{Z}, \end{aligned} \quad (P1)$$

where the constraint restates the policy goal $\sum_{a \in A} E_i^{a'} + I_i^{c,a'} + I_i^{sc,a'} \leq 0$ more explicitly using (2)–(4).

We explore the optimal policy as well as its *efficiency*, that is, the level of total economic activity it permits in neighborhoods in \mathcal{Z} relative to the baseline level (i.e., $\sum_{i \in \mathcal{Z}} \tilde{c}_i x_i / \sum_{i \in \mathcal{Z}} \tilde{c}_i$). This quantity is at most one, and it is close to one only if the economic restrictions are very mild. Efficiency is defined similarly in other variants of the baseline model where additional constraints are imposed on the feasible policies.

Remark. We focused on the policy goal of reducing infections in all neighborhoods. An alternative goal is to achieve reduction in *total infections* in all neighborhoods in \mathcal{Z} . An accompanying optimization formulation can easily be obtained by summing both sides of the constraints of (P1) over $i \in \mathcal{Z}$. The optimal policy for such a formulation permits economic activity aggressively in some neighborhoods (with a large number of jobs) to an extent that it increases infections there. The policy balances this by restricting activity elsewhere to achieve a net reduction in infections. Letting infections increase in some neighborhoods while suppressing them elsewhere could overwhelm healthcare resources in certain neighborhoods and could be viewed as unfair. In addition, the increasing infections in those neighborhoods can trigger future spread of infections, thereby leading to difficulties in controlling the infection in the long run. These observations motivated us to focus on the more stringent policy goal presented in this section. Policies that satisfy this goal (i) do not reduce cases in some neighborhoods at the expense of others and (ii) ensure that the loads on the healthcare resources in each neighborhood do not exceed their current levels. It is worthwhile to point out that (i) additional restrictions on $\{x_i\}$, for example, lower-bound constraints to ensure that at least some level of economic activity is permitted in each region (see Section EC.3.6), and (ii) additional infection-related costs capturing, for example, disease-related deaths (see Section EC.3.7), can readily be incorporated into (P1), while preserving its tractability.

3. Data Sources and Model Calibration

Before we apply our model to the study of COVID-19 in NYC, we go over relevant data sources and discuss how model primitives are calibrated. There are five main types of data that we use to determine the primitives of our model: (i) disease-specific parameters (e.g., transmission rate, clinical rate) for COVID-19, (ii) infection counts for NYC and neighboring counties (circa April 2020), (iii) populations of different neighborhoods, (iv) workplace and residence location information for individuals who reside in NYC and neighboring counties as well as information on whether they can work remotely, and (v) SafeGraph data¹⁰ that

record individuals' movements (SafeGraph 2020). We next discuss our data sources and detail how we construct the different primitives of our model using the available data. We also explain how we characterize interactions of populations in different neighborhoods while accounting for the effects of work and leisure trips as well as the potential for remote work.¹¹

The spatial nature of epidemic spread and our model admit a natural network representation in which nodes correspond to different locations/neighborhoods and (weighted) edges encode trips of individuals between these neighborhoods. In what follows, we also clarify how these nodes and edge weights are defined in our application to NYC. We sometimes use the terms nodes/neighborhoods interchangeably.

3.1. Disease-Specific Parameters

For our study of COVID-19, we calibrate the parameters $\beta, \alpha, \kappa, \rho, \gamma$, using the estimates provided in the recent literature. Specifically, the estimates of all parameters (or their inverses) other than the transmission rate β are given in Table 1 and are taken from Li et al. (2020).¹² Recall that the transmission rate β depends on the number of effective contacts an individual has on a given day. Because the number of effective contacts can be different for different cities/countries, in our numerical study we base this parameter on the NYC analysis provided in Fernández-Villaverde and Jones (2020). The latter paper estimates the basic reproduction number \mathcal{R}_0 of the disease in NYC before the lockdown to be 2.7 (in mid-March 2020), which in turn implies that $\beta = 0.78 \text{ day}^{-1}$. For our main analysis, we use the values of the disease-specific parameters reported here.

In Section EC.2, we provide confidence intervals for the parameter estimates we report here (also taken from the literature). We use these for our robustness analyses. Specifically, we fit distributions to the parameters that are consistent with the (median) estimates and confidence intervals. Then, we sample the parameters from these distributions and resolve our model. As we detail in Section EC.3.1, our takeaways still hold when the parameters are sampled from the aforementioned distributions and the insights obtained by focusing on the parameter estimates reported here turn out to be robust.

Table 1. Disease-Specific Parameters (Adapted from Li et al. 2020)

Parameter	Estimate (median)
$\kappa^{-1} \text{ day}$	3.69
ρ	0.14
α	0.55
$\gamma^{-1} \text{ day}$	3.47

3.2. Neighborhoods/Nodes (\mathcal{N})

NYC reports confirmed COVID-19 cases at the modified Zip Code Tabulation Area (MODZCTA) level (NYC Health 2020). MODZCTAs are obtained from Zip Code Tabulation areas (ZCTAs) by incorporating small ZCTAs into adjacent larger ones.¹³ ZCTAs in turn consist of census blocks (CBs).¹⁴ Motivated by the availability of the case data at the MODZCTA level, we define the neighborhoods/nodes in our analysis as the different MODZCTAs in the city.

We are mainly interested in understanding how NYC can reduce economic activity in different MODZCTAs so as to curb the spread of the epidemic. However, because of the spatial nature of the spread of the epidemic, infected individuals from outside NYC can infect susceptible individuals in NYC. To capture this, we define additional nodes that correspond to the counties neighboring NYC.¹⁵ Specifically, we use the SafeGraph data (discussed below) to obtain the 10 counties that have the largest inflow/outflow of individuals to/from NYC. These 10 counties together with the NYC MODZCTAs constitute the set of all nodes of our network. The list of these counties is provided in Table 2. In total, our network has 187 nodes.

3.3. Populations ($\{N_i^a\}$) and Baseline Economic Activity ($\{\tilde{c}_i\}$)

In order to define the total population of each node, we use the 2010 Census data that specify populations at the CB level (U.S. Census Bureau 2010). To capture the effect of population changes over years, we use the data in U.S. Census Bureau (2018) that report populations in 2018 at the county level. Using these data, we obtain a multiplier for each county that captures how much the population changed between 2010 and 2018. Then, we rescale the population of each CB reported in U.S. Census Bureau (2010) with the multiplier of the county to which it belongs. Finally, we use these to construct the population of each node in our network. These populations are relevant as they are the primitives ($\{N_i\}$) of our spatial epidemic model. We plot the spatial distribution of NYC's population in Section EC.2 (see Figure EC.1).

Next, we discuss how to construct subpopulations N_i^A, N_i^B, N_i^C . For each two-digit North American Industry Classification System (NAICS) industry (see Table EC.4 in Section EC.2 for a list of all NAICS industries), U.S. Census Bureau (2017) records (i) the number of jobs in each CB and (ii) the number of employees who reside in each CB and have occupations in that

industry. In addition, it considers three supersectors (each of which consists of different NAICS industries; see Table EC.3 in Section EC.2) and contains (iii) the total number of individuals in each supersector who reside in one CB and work in another. Aggregating (i), (ii), and (iii) at the MODZCTA level (or county level for neighboring counties), we construct the number of jobs $\{\bar{c}_i^{(k)}\}$ in (and employees $\{\bar{r}_i^{(k)}\}$) from neighborhood i for each NAICS industry k and the number of employees $W_{ij}^{(k)}$ who reside in neighborhood i and work in neighborhood j for each NAICS supersector K .

Define $\mathbb{S}^{(K)}$ as the set of NAICS industries contained in NAICS supersector K . Let $w_{ij}^{(k)}$ denote the number of employees who reside in neighborhood i , work in neighborhood j , and are employed in NAICS industry k . It follows from these definitions that $\{w_{ij}^{(k)}\}$ satisfy the following constraints:

$$\begin{aligned} \sum_{k \in \mathbb{S}^{(K)}} w_{ij}^{(k)} &= W_{ij}^{(K)} & \forall i, j, K, \\ \sum_j w_{ij}^{(k)} &\leq \bar{r}_i^{(k)} & \forall i, k, \\ \sum_i w_{ij}^{(k)} &\leq \bar{c}_j^{(k)} & \forall j, k, \\ w_{ij}^{(k)} &\geq 0 & \forall i, j, k, \end{aligned}$$

where the second and third constraints are inequalities because there is a small number of work-related trips between the locations in our network and other locations (e.g., counties other than those in Table 2). This set of constraints determines a polytope to which $\{w_{ij}^{(k)}\}$ belongs. For our analysis, we set $\{w_{ij}^{(k)}\}$ values to the centroid of this polytope.¹⁶ Let $r_i^{(k)} = \sum_j w_{ij}^{(k)}$ and $c_i^{(k)} = \sum_j w_{ji}^{(k)}$ denote the number of employees/jobs from/in each MODZCTA (or county) implied by this solution. We observe that the gap between $\bar{r}_i^{(k)}$ and $r_i^{(k)}$ (similarly $\bar{c}_i^{(k)}$ and $c_i^{(k)}$) is small.¹⁷

Dingel and Neiman (2020) estimate the *teleworkable ratio* (i.e., the fraction of jobs that can be conducted remotely) $\{\theta^{(k)}\}$ for each two-digit NAICS industry k (for completeness in Table EC.4 we provide the teleworkable ratios of different NAICS industries). Using these together with $\{c_i^{(k)}, r_i^{(k)}\}$ constructed earlier, we compute the number of jobs in each neighborhood i that are not amenable to remote work $\tilde{c}_i = \sum_k c_i^{(k)} (1 - \theta^{(k)})$ and the number of individuals from different neighborhoods who are employed in teleworkable and nonteleworkable jobs, that is, subpopulations

Table 2. Counties/Nodes Outside NYC

	Fairfield County, CT	Bergen County, NJ	Monmouth County, NJ
Essex County, NJ	Hudson County, NJ	Middlesex County, NJ	Westchester County, NY
Union County, NJ	Nassau County, NY	Suffolk County, NY	

$N_i^A = \sum_k r_i^{(k)} \theta^{(k)}$ and $N_i^B = \sum_k r_i^{(k)} (1 - \theta^{(k)})$. The rest of the population in each neighborhood is unemployed; hence, we have $N_i^C \triangleq N_i - N_i^A - N_i^B$.

According to the data in U.S. Census Bureau (2017), there are a total of 4.4M jobs in NYC. Using the estimation of $\{\tilde{c}_i\}$ described above, 58% of these, or 2.5M jobs, are not amenable to remote work. It is worth pointing out that the number of jobs $\{c_i = \sum_k c_i^{(k)}\}$ in each MODZCTA and the number of these jobs that are not amenable to remote work $\{\tilde{c}_i\}$ are strongly correlated and exhibit a Pearson correlation coefficient of 0.987. In other words, the distribution of all jobs and jobs that are not amenable to remote work in NYC show similar patterns. In Figure EC.3, we also illustrate this point through a plot of the spatial distribution of (nonteleworkable) jobs in NYC.

3.4. Sizes of Compartments ($\{S_i^a/E_i^a/I_i^{c,a}/I_i^{sc,a}/R_i^a\}$)

We focus on the state of the epidemic in NYC (and neighboring counties) on April 18, 2020, and solve the planner's decision problem. Note that, on the aforementioned date, the total number of infections in NYC was already nontrivial and the policy goal in (P1) (reducing infections in all neighborhoods) was natural. Recall that this formulation is based on the size of each of the disease compartments. Thus, we need to estimate the sizes of these components on April 18. We next describe how we do so. In what follows, we refer to April 18 as the target date and denote it by T_* . We also focus on April 21 and April 15 and denote these dates respectively by T_+ and T_- .

For the nodes of our network, we can extract the number of confirmed cases on a given date, using the data provided in NYC Health (2020) and New York Times (2020).¹⁸ On the other hand, recent studies suggest that there were a large number of individuals who had the disease but were not among confirmed cases (either because they did not show any symptoms or because of congestion in the healthcare system/diagnostic test shortages). In fact, antibody studies suggest that in NYC about 24.7% of the individuals may have experienced the disease (as of April 27, 2020) (see New York State Governor's Office 2020a), a much larger number than the number of confirmed cases.¹⁹

To bridge the gap between the confirmed cases and the total infections implied by the antibody studies, we define an identification rate. We first describe how this quantity and the compartment sizes are obtained for the MODZCTAs in NYC and then briefly discuss the extension of the approach to the neighboring counties.

We assume that on T_* 24.7% of NYC have had the disease.²⁰ We also compute the total number of confirmed infections in NYC by T_* (at all locations) using the data set in NYC Health (2020). Dividing the latter

number by the former, we obtain an identification rate, which represents the fraction of actual infections that become confirmed cases. We assume that the identification rate is the same in all neighborhoods.

The data in NYC Health (2020) yield the daily numbers of confirmed cases by MODZCTA for NYC. We analyze the data under the assumption that individuals spend three days in the exposed compartment and three days in the infected compartments (by rounding down the parameters reported above in our discussion on the disease-specific parameters). Observe that under this assumption, infected individuals on T_- leave the infected compartments by T_* . Hence, we focus on the cumulative number of confirmed cases on dates T_* and T_- in each MODZCTA and interpret the difference as the number of active identified infections on T_* . By dividing this quantity by our identification rate, we obtain the number of infections that are active in each MODZCTA of NYC on T_* . Note that these active infections may be clinical or subclinical; thus, the aforementioned quantity is $I_i^c + I_i^{sc}$ for MODZCTA i . We use the clinical rate ρ reported above to solve for I_i^c and I_i^{sc} (where $\rho = I_i^c / (I_i^c + I_i^{sc})$) on T_* .

Similarly, under our assumption on the time individuals spend in the exposed/infected compartments, the only infected individuals on T_+ are those who were in the exposed compartment on T_* . We focus on the cumulative number of confirmed cases in each MODZCTA on T_* and T_+ . Subtracting the former number from the latter one, we obtain the number of active identified infections on T_+ . Dividing this quantity by the identification rate, and repeating these steps for each MODZCTA, gives us the active infections in each MODZCTA on T_+ or equivalently the number of exposed individuals $\{E_i\}$ on T_* .

Finally, under our assumptions all of the confirmed cases by T_- move to the recovered compartment by T_* . Dividing the relevant number in each MODZCTA by the identification rate yields $\{R_i\}$ on T_* .

Consider node i and the compartment sizes $E_i/I_i^c/I_i^{sc}/R_i$ described above. We compute the difference between N_i and the sum of these quantities to obtain the size of the susceptible compartment, S_i , for all MODZCTAs i . Finally, we partition each compartment into subpopulations N_i^A, N_i^B, N_i^C proportional to their sizes, that is, we set $S_i^a = S_i N_i^a / N_i$ and similarly for the other compartments.

The compartment sizes for the 10 counties neighboring NYC are obtained following a similar approach. The only difference is that the antibody tests suggest lower prevalence in these counties than in NYC (New York State Governor's Office 2020a). Thus, we replace the number 24.7% with a county-specific number and repeat the same process (using the data from New York Times (2020) to compute the county-specific identification rates).

See Figure EC.5 and Figure EC.6 for illustrations of the spatial distribution of infections in NYC on T_* .

3.5. Edge Weights ($\{\tilde{\tau}_{ij}^a\}$)

SafeGraph uses data from a number of mobile devices and applications to derive anonymized location data that provide insights into locations visited by individuals as well as their movements across different locations. To preserve anonymity, the data are aggregated at the census block group (CBG) level. More specifically, the SafeGraph data set identifies a “home” CBG for each device/individual and reports the median home dwell time for each CBG. In addition, it reports the daily number of individuals who go from their home CBG to various destination CBGs. Similarly to MODZCTAs, each CBG consists of a number of CBs. But, in general, CBGs neither contain nor are contained by MODZCTAs. To get around this difficulty, we first disaggregate the SafeGraph data at the CB level. In particular, we focus on all “trips” that originate from a home CBG i and that go to a destination CBG j . We assume that the fraction of these trips that originate from a particular CB in CBG i is simply given by the ratio of the population of this CB (obtained from the data in U.S. Census Bureau 2010) to the total population of CBG i . Similarly, we assume that the fraction of these trips that end in a particular CB in CBG j is given by the ratio of the area of this CB to the area of CBG j .²¹ Under this assumption, we can break down the “flow” of individuals between home/destination CBGs to the flow between home/destination CBs. Then we aggregate these quantities at the MODZCTA level for NYC (at the county level for the nodes outside NYC) to define the flow of individuals from one node (MODZCTA in NYC or a neighboring county) to another. Similarly, we assume that all CBs in the same CBG share the same home dwell time; we define the home dwell time for a MODZCTA as the average of these for the CBs that are contained in the MODZCTA (and similarly for the neighboring counties).

To capture individuals’ movements prior to lockdowns in NYC, we focus on the SafeGraph data from January 1, 2020 to February 29, 2020 and for any home/destination neighborhood pair we aggregate all trips during this period.²² Our construction allows us to obtain a representative home dwell time for each node. Moreover, it allows us to capture the number of trips (in the data set) from one node to another. Given node i , we focus on the number of trips to a node j divided by the total number of trips leaving i . We use this quantity to define the fraction of time the individuals from i spend in j in aggregate, denoted by τ_{ij} . More precisely, if the home dwell time in node i is h_i hours, and the number of trips from i to ℓ is $k_{i\ell}$, we let $\tau_{ij} = (1 - (h_i/24)) \times (k_{ij}/\sum_\ell k_{i\ell})$.

We proceed by decomposing τ_{ij} into work/leisure-related components for each subpopulation. Recall that we obtained $\{w_{ij}^{(k)}\}$, that is, the number of employees who reside in neighborhood i and work in neighborhood j for NAICS industry k . We use these to construct the quantity $\zeta_{ij,W}^A = (\sum_k w_{ij}^{(k)} \theta^{(k)} / \sum_{l,k} w_{il}^{(k)} \theta^{(k)})$, which captures the pre-epidemic fraction of work-related trips of individuals from subpopulation N_i^A that end in node j . Similarly, we define $\zeta_{ij,W}^B = (\sum_k w_{ij}^{(k)} (1 - \theta^{(k)}) / \sum_{l,k} w_{il}^{(k)} (1 - \theta^{(k)}))$ and $\zeta_{ij,W}^C = 0$ for subpopulations N_i^B and N_i^C , respectively.

The American Time Use Survey (ATUS) reports the amount of time individuals spend on different activities (see U.S. Bureau of Labor Statistics 2020). We focus on three numbers derived from these surveys: (i) the average daily time spent at the workplace per individual with employment (h_W), (ii) the average daily time spent on outdoor leisure per individual with employment (h_L), (iii) the average daily time spent on outdoor leisure per individual without employment (h'_L). These quantities are respectively given as follows: $h_W = 4.82$ hours, $h_L = 1.28$ hours, $h'_L = 1.50$ hours.²³

We denote the leisure time employed (unemployed) individuals from i spend at location j by $h_L \zeta_{ij,L}$ ($h'_L \zeta_{ij,L}$). Here $\{\zeta_{ij,L}\}$ are leisure-trip factors that determine the portion of leisure time spent at different locations. Because we do not have more detailed data, we assume that subpopulations N_i^A, N_i^B, N_i^C share the same factors. On the other hand, it follows from the definitions of τ_{ij} , $\zeta_{ij,W}^a$, $\zeta_{ij,L}$ that these quantities satisfy the following relation:

$$24 N_i \tau_{ij} = h_W (N_i^A \zeta_{ij,W}^A + N_i^B \zeta_{ij,W}^B) + h_L (N_i^A + N_i^B) \zeta_{ij,L} + h'_L N_i^C \zeta_{ij,L}. \quad (5)$$

Here both sides of the equation reflect the total number of hours (in an average day) that individuals who reside in node i spend in node j . The left-hand side expresses this quantity directly by focusing on $\{\tau_{ij}\}$, that is, the fraction of time individuals who reside in i spend in j . The right-hand side expresses the same quantity by focusing on subpopulations $\{N_i^A, N_i^B, N_i^C\}$ as well as their work/leisure behavior. To see this more clearly, recall that a fraction $\zeta_{ij,W}^a$ of the work trips of subpopulation N_i^a end in j , and individuals who take these trips spend h_W hours a day on average working in node j . Similarly, as explained earlier, the leisure time that employed (unemployed) individuals who reside in i spend in j is given by $h_L \zeta_{ij,L}$ ($h'_L \zeta_{ij,L}$). Thus, the quantity on the right-hand side also gives the total number of hours that individuals who reside in i spend in j .

Using (5) and the quantities defined earlier, we solve for $\{\zeta_{ij,L}^a\}$ as well. Specifically, we express $\{\tau_{ij,W}^a, \tau_{ij,L}^a\}_{a \in \mathcal{A}}$ as follows:²⁴ $\tau_{ij,W}^A = (h_W/24)\zeta_{ij,W}^A$, $\tau_{ij,W}^B = (h_W/24)\zeta_{ij,W}^B$, $\tau_{ij,L}^A = \tau_{ij,L}^B = (h_L/24)\zeta_{ij,L}$, and $\tau_{ij,L}^C = (h_L'/24)\zeta_{ij,L}$. Finally, the $\{\tilde{\tau}_{ij}^a\}$ variables that appear in our optimization program (P1) can be computed in terms of the aforementioned quantities as $\tilde{\tau}_{ij}^A = \tau_{ij,L}^A$, $\tilde{\tau}_{ij}^B = \tau_{ij,L}^B + \tau_{ij,W}^B$, and $\tilde{\tau}_{ij}^C = \tau_{ij,L}^C$.

4. Optimal Targeted Closures

In this section, we use our model to illuminate critical issues for planners to consider in determining how to minimize the economic costs from unemployment while controlling an epidemic. We use the NYC data to demonstrate the potential advantages of targeted closures compared with uniform ones. The results show significant potential gains in employment (or significantly reduced economic costs from unemployment) from the use of targeted policies. In addition, they demonstrate that the implied policies differ substantially from naïve policies that do not consider mobility and only focus on localized infection numbers. They also highlight the significant impact that considering regions beyond the planner's jurisdiction can have on the planner's capabilities to control an epidemic.

We consider three different scenarios $y \in \{0, 0.5, 1\}$ to capture the permitted economic activity in the neighboring counties ($\mathcal{N} \setminus \mathcal{Z}$) and solve (P1) to obtain the optimal policy $\{x_i\}_{i \in \mathcal{Z}}$ for the MODZCTAs of NYC (\mathcal{Z}) in each scenario. The optimal policies are presented in Figure 1.²⁵

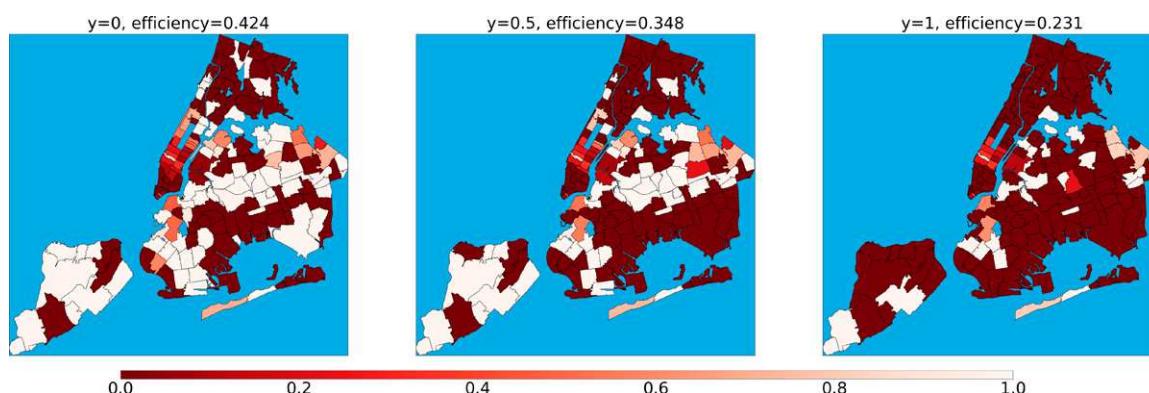
When the economic activity in the neighboring regions is completely suspended, the optimal targeted closure policy can guarantee 42.4% efficiency, while achieving the relevant policy goal. On the other hand, when the economic activity level in the neighboring counties is 0.5 or 1, the efficiency decreases to 34.8% and 23.1%, respectively. Hence, the change in the

economic activity level of neighboring counties from $y=0$ to $y=1$ reduces the achievable nonteleworkable employment level in NYC by 45.4%. This observation highlights the necessity of coordination between states or other governing bodies to ensure small economic losses.

In all cases, the optimal policy involves having some economic activity in Midtown Manhattan, while restricting economic activity substantially (and in some cases completely) elsewhere in Manhattan. The latter effect is more prominent when the economic activity in neighboring counties increases. The aforementioned structure of the optimal policy is partly because Midtown has the largest number of jobs in the city (both nonteleworkable and total; see Figure EC.3 in Section EC.2). It is worth noting that other neighborhoods (such as Lower Manhattan) also have a substantial number of jobs, yet the permitted economic activity at the optimal solution is (almost) zero. This is because if the permitted economic activity in those neighborhoods were increased, then because of the spatial nature of disease spread, in order to achieve the policy goals, the economic activity level in other locations (such as Midtown) would need to be reduced. The cost of the latter reduction outweighs the benefit of increasing economic activity in Lower Manhattan.

It is worth highlighting that the restrictions can change substantially even between adjacent neighborhoods (e.g., those in Staten Island). More interestingly, even among neighborhoods with a similar number of jobs, the ones with higher levels of permitted economic activity are not necessarily those with lower levels of infection. For instance, in Staten Island, some neighborhoods that are permitted to continue economic activity have higher infection rates than the adjacent neighborhoods that are completely shut down (see Figure EC.5 in Section EC.2). This is because, when a planner decides how much economic activity to allow, in addition to neighborhood-specific

Figure 1. (Color online) The Neighborhoods' Colors (as Summarized in the Color Bar) Indicate the Optimal $\{x_i\}$



metrics (such as the infection rate in a neighborhood), the spillovers across neighborhoods (i.e., imports/exports of the disease across neighborhoods) also matter.

The significant differences in permitted economic activity levels even in adjacent MODZCTAs may be undesirable because of fairness concerns. A conceptually more fair alternative is to limit the differences in the level of permitted economic activity in all neighborhoods. We study this alternative by including the “fairness constraint” $x_i - x_j \leq 1 - \theta$ for all $i, j \in \mathcal{Z}$ in (P1), where $\theta \geq 0$ characterizes the maximum allowed difference. In Figure 2(a), we focus on scenarios $y \in \{0, 0.5, 1\}$ and illustrate how the efficiency of the optimal solution varies with θ . In all scenarios, the fairness constraints result in a reduction in efficiency. The extreme where $\theta = 1$ captures *uniform policies* where the planner must permit the same level of economic activity everywhere.²⁶ Relative to the optimal targeted policy ($\theta = 0$), we see that the optimal uniform policy yields a 3.92–6.25-fold reduction in efficiency (or, equivalently, the permitted nonteleworkable employment level). In Section EC.3.6, we consider a similar variant where we impose the constraint $x_i - x_j \leq 1 - \theta$ on neighborhoods $i, j \in \mathcal{Z}$ that are adjacent.²⁷ We show that this variant yields similar conclusions in terms of the value of targeting (over uniform policies) as well as the targeting structure.

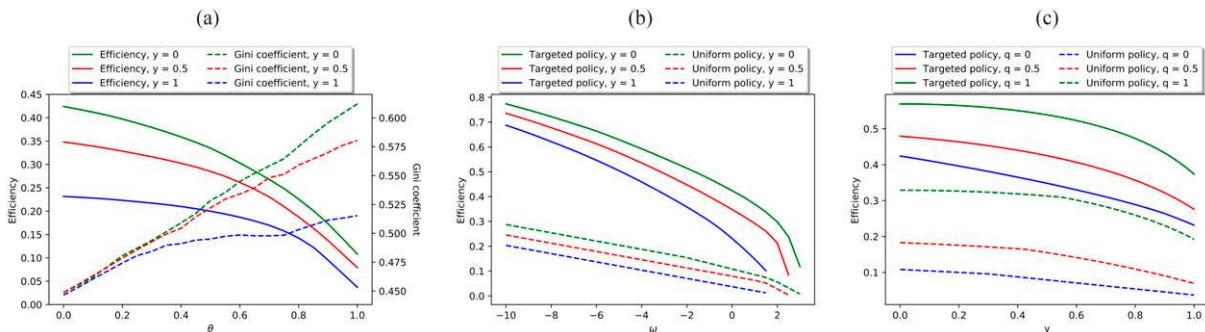
A subtle point about fairness constraints is worth highlighting: they ensure fairness in terms of restrictions of economic activity. On the other hand, given the heterogeneity in terms of the population movement patterns and infection levels across locations, such constraints may be unfair in terms of the induced infection rates. To illustrate this, at the optimal solution obtained for different θ , we compute the rate at which infections decrease ($z_i = |\sum_{a \in \mathcal{A}} E_i^{a'} + I_i^{c,a'} + I_i^{sc,a'}|$) in each neighborhood $i \in \mathcal{Z}$, that is, the slack in the relevant constraints of (P1). We then compute the Gini

coefficient $\sum_{i \in \mathcal{Z}} \sum_{j \in \mathcal{Z}} |z_i - z_j| / 2 \sum_{i \in \mathcal{Z}} \sum_{j \in \mathcal{Z}} z_i$, which measures the degree of inequality in these rates. Figure 2(a) illustrates that indeed more equal treatment in terms of closures could translate to less fair outcomes in terms of the spread of the infection.

Infections could induce additional societal losses (such as excess deaths); hence, it may be desirable to reduce cases more aggressively. We consider such a variant of the problem by replacing the constraint $\sum_{a \in \mathcal{A}} E_i^{a'} + I_i^{c,a'} + I_i^{sc,a'} \leq 0$ in (P1) with $\sum_{a \in \mathcal{A}} E_i^{a'} + I_i^{c,a'} + I_i^{sc,a'} \leq -(\omega/100) \sum_{a \in \mathcal{A}} (E_i^a + I_i^{c,a} + I_i^{sc,a})$. Here, ω can take positive or negative values and $\omega > 0$ encodes the requirement of reducing cases aggressively.²⁸ We compute the efficiency for different ω under the optimal targeted and uniform policies. The induced curves, referred to as Pareto frontiers, capture the best trade-offs achievable between ω and efficiency. Figure 2(b) presents Pareto frontiers for $y \in \{0, 0.5, 1\}$ and shows that the gap between the optimal targeted and uniform policies is large for any value of ω . For large ω , the social planner’s problem becomes infeasible; the lines are cut off at the ω -values where this happens. Note that larger values of ω guarantee faster reduction in cases at the expense of higher unemployment. To illustrate this for $y = 0.5$, we obtain the optimal targeted policies for different ω and compute the time it takes city-wide active infections to decrease to half of the initial level. We see that by setting $\omega = 0.5$, $\omega = 1.0$, $\omega = 2.0$, the social planner can guarantee that the infections will halve in 10, 9, and 8 days and achieve an efficiency of 32.1%, 29.2%, and 21.3%, respectively. We detail the trade-offs between halving times and efficiency that can be obtained through an appropriate choice of ω in Section EC.3.6.

Social planners could prioritize low-risk businesses over high-risk ones (e.g., bars and restaurants) by allowing the former to remain open while closing the latter. Our model is flexible enough to capture such prioritization. Specifically, suppose that if a fraction x_j

Figure 2. (Color online) Impact of Fairness Constraints (a), Pareto Frontiers (b), and Prioritizing Low-Risk Businesses (c) for $y \in \{0, 0.5, 1\}$



Note. (a) Efficiency and Gini coefficient for different θ ; (b) Pareto frontiers; (c) efficiency when low-risk businesses are prioritized.

of the businesses are open, then infections happen in j at a rate proportional to $x_j\eta(x_j)$. Here, η is a (weakly) increasing function that captures infection efficiency: as x_j increases, more risky businesses are allowed to operate, thereby leading to a disproportionately larger number of infections. We can reflect the impact of infection efficiency on the evolution of the epidemic and our optimization framework by replacing x_j with $x_j\eta(x_j)$ in (2) and (3) as well as the constraints of (P1).

Note that for increasing η , after this change, (P1) remains a convex program. To illustrate how the findings change under prioritization of low-risk businesses, we let $\eta(x_j) = (1 - q) + qx_j$. This parameterization of η ensures that when all businesses are closed ($x_j = 0$) and all of them are open ($x_j = 1$), $x_j\eta(x_j) = x_j$ and the rate of infection is the same as in the baseline model (because at these extremes, prioritization does not play a role). Here, q captures the value of prioritizing low-risk businesses: when $q = 1$ the marginal risk is arbitrarily smaller for $x_j \ll 1$ than for $x_j \approx 1$. Using the aforementioned optimization formulation, we obtain the optimal targeting and uniform policies. Figure 2(c) shows that for different values of q , the efficiency gap between these remains sizable. That is, similar to the baseline model ($q = 0$), targeting achieves the policy goal with substantially lower unemployment, even when low-risk businesses are prioritized (for all levels of y).

In the variants of our baseline model considered here, the optimal targeting structure is similar to Figure 1. For instance, when $y = 0.5$, the Pearson correlation coefficient between the baseline optimal solution and the variant in which (i) $\theta = 0.5$ is 0.71, (ii) $\omega = 2$ is 0.66, and (iii) $q = 0.5$ is 0.91. The differences in the solutions are mainly driven by neighborhoods outside Manhattan. In all cases, the optimal solutions exhibit similarities; for example, within Manhattan some economic activity is permitted in Midtown, whereas more stringent restrictions are imposed elsewhere (including Lower Manhattan). See Figure EC.20 in Section EC.3.6.

5. Trip Substitutions, Coarse Targeting, and Dynamic Policies

Whereas our baseline model applies directly to situations in which individuals prefer to stay at home (as opposed to traveling to other destinations) when a destination has restricted activity, the planner allows for fine-grained decisions on restrictions and the planner's focus is on the immediate containment of infections; alternative settings may be relevant. In this section, we incorporate three practically relevant features highlighted earlier into our baseline model and verify that our main observations (particularly in terms of the significant gains from targeted relative to uniform policies) still hold.

As a first alternative, we consider settings where trip substitutions play a role: individuals whose destination neighborhoods are targeted with restrictions may instead decide to visit other neighborhoods. Note that such trip substitutions could be relevant for leisure trips, but this is unlikely the case for work-related trips. This is because it is unlikely that many individuals would be able to switch jobs in a short amount of time in response to restrictions of economic activity in neighborhoods that contain their workplaces. Second, fine-grained (e.g., MODZCTA level) targeting may not always be feasible; and it may be necessary to develop coarser targeted policies. In addition, as opposed to specifying the policy goal in terms of instantaneous infection rates, it may be reasonable in practice to express it in terms of the cases induced over time. We first explain how we incorporate these features into our model and develop coarse targeted policies that control the spread of the epidemic. Using our trip substitution model and coarse-targeting approach as building blocks, we then turn our attention to controlling the disease using dynamic policies. Throughout this section, we discuss our findings by focusing on the $y = 0.5$ case for the employment level in regions outside the planner's control.

5.1. Leisure-Trip Substitutions

To capture leisure-trip substitutions, we posit a model similar to multinomial (MNL) choice models. Specifically, we assume that individuals from i can choose to spend their leisure time by (i) visiting neighborhood $j \neq i$, (ii) spending time in neighborhood i outside home, and (iii) spending time at home. We refer to these respectively as action j , action i , and action \emptyset and denote the set of possible actions by $\mathcal{L}_i = \{j \in \mathcal{N} \mid \sum_{a \in \mathcal{A}} N_i^a \times \tau_{ij,L}^a \geq 1\} \cup \{\emptyset\}$. As discussed in Section 2, $\tau_{ij,L}^a$ captures the fraction of time individuals from sub-population N_i^a spend in j for leisure and is constructed using the pre-epidemic mobility data. Set \mathcal{L}_i is obtained from \mathcal{N} by excluding neighborhoods to which almost no (more precisely, on average less than one) individual from i travels for leisure. We observe from data that leisure trips are concentrated locally, which suggests that for leisure trips the disutility increases quickly in the distance between the origin neighborhood and the destination neighborhood (see Figure EC.7 in Section EC.2 for an illustration).

Given the permitted economic activity levels $\{x_i\}_{i \in \mathcal{N}}$ ($\triangleq \mathbf{x}$), we denote the expected payoff from action $\ell \in \mathcal{L}_i$ by $u_{i,\ell}^a(\mathbf{x})$. We assume that

$$u_{i,\ell}^a(\mathbf{x}) = x_\ell (V(\ell) \mathbf{1}\{\ell \notin \{i, \emptyset\}\} + H_1(i) \mathbf{1}\{\ell = i\}) + H_2^a(i) \mathbf{1}\{\ell = \emptyset\} - K d_{i,\ell}^v \mathbf{1}\{\ell \notin \{i, \emptyset\}\}. \quad (6)$$

Here, $V(\ell)$ captures the value individuals derive from spending leisure time in neighborhood ℓ . Individuals

could derive a different value from visiting attractions that are local to them. We incorporate this into the model through the second term, which explicitly allows individuals who reside in i to have value $H_1(i)$ (possibly different from $V(i)$) for spending leisure time in i . The x_ℓ multiplier ensures that the payoff from these options decreases with restrictions in ℓ . Although the first two terms capture the value for spending time outside home, the third term captures the value from spending leisure time at home. We allow this term to be different for distinct subpopulations a to reflect the fact that some subpopulations (e.g., unemployed individuals) have more leisure time and may spend a different fraction of this time at home.²⁹ The last term captures the disutility due to traveling from neighborhood i to ℓ , which depends on the distance $d_{i,\ell}$ between them. As we mentioned above, it can be observed from data that leisure trips are concentrated locally, which suggests that for short distances the disutility increases quickly in $d_{i,\ell}$. To accommodate different impacts of distance on trip patterns, we allow the distance penalty to have an exponent ν and estimate our model with different ν . We observe that the best fit is achieved for $\nu = 0.1$ and use this value in our analysis.³⁰

We assume that the total utility an individual derives from action $\ell \in \mathcal{L}_i$ is given by $u_{i,\ell}^a(\mathbf{x}) + \epsilon_{i,\ell}^a$, where ϵ_i is the choice-specific shock drawn independent and identically distributed from Extreme Value Type I Distribution with location parameter 0 (Hotz and Miller 1993). Individuals first observe the realization of the choice-specific shock and then select the action to maximize their utility. For the given model of choice-specific shocks, the ex ante probability that action ℓ is chosen is proportional to $\exp(u_{i,\ell}^a(\mathbf{x}))$. In turn, this implies that the (average) fraction of time $\tau_{i,\ell,L}^a(\mathbf{x})$ an individual in subpopulation N_i^a spends on leisure by choosing action $\ell \in \mathcal{L}_i$ is proportional to $\exp(u_{i,\ell}^a(\mathbf{x}))$. Hence, after appropriate normalization, for each $i \in \mathcal{N}$, $a \in \mathcal{A}$ and $\ell \in \mathcal{L}_i$, we obtain

$$\tau_{i,\ell,L}^a(\mathbf{x}) = \begin{cases} \frac{\exp(u_{i,\ell}^a(\mathbf{x}))}{\sum_{\ell' \in \mathcal{L}_i} \exp(u_{i,\ell'}^a(\mathbf{x}))} \cdot \frac{\bar{h}_{i,L}^a}{24} & \ell \in \mathcal{L}_i, \\ 0 & \ell \notin \mathcal{L}_i, \end{cases} \quad (7)$$

where $\bar{h}_{i,L}^a$ is the average number of hours individuals in subpopulation N_i^a spend on leisure per day.

Note that the model readily captures trip substitutions. Specifically, if the economic activity in neighborhood ℓ is restricted, the payoff individuals derive from visiting ℓ decreases (see (6)). As a result, the portion of leisure time they spend in ℓ decreases and they start spending more time in neighborhood $j \neq \ell$ or at home. The rate at which trips are substituted depends on the values of different neighborhoods.³¹

5.2. Model Estimation

We estimate the parameters V , H_1 , H_2^a , and K , which capture the attractiveness of different neighborhoods (as well as the disutility due to traveling to more distant neighborhoods), by focusing on the period before lockdowns in NYC.

As mentioned earlier, U.S. Bureau of Labor Statistics (2020) reports (i) the average daily time spent on outdoor leisure per individual with employment (h_L) and (ii) the average daily time spent on outdoor leisure per individual without employment (h'_L). In addition, it reports (iii) the average daily time spent on leisure per individual (both outdoors and at home) with employment (\bar{h}_L) and (iv) the average daily time spent on leisure per individual (both outdoors and at home) without employment (\bar{h}'_L). These quantities are given as follows: $h_L = 1.28$ hours, $h'_L = 1.50$ hours, $\bar{h}_L = 4.18$ hours, and $\bar{h}'_L = 6.91$ hours. Thus, the ratio of total leisure time to outdoor leisure time is $\kappa_L = \bar{h}_L/h_L$ for employed individuals and $\kappa'_L = \bar{h}'_L/h'_L$ for unemployed ones. We assume that the ratio of total leisure time to outdoor leisure time depends only on the employment status and that these ratios are the same for all neighborhoods.

Note that $\bar{h}_{i,L}^a$ can be expressed in terms of these parameters as follows:

$$\begin{aligned} \bar{h}_{i,L}^A &= 24\kappa_L \sum_{j \in \mathcal{N}} \tau_{ij,L}^A(\mathbf{1}), & \bar{h}_{i,L}^B &= 24\kappa_L \sum_{j \in \mathcal{N}} \tau_{ij,L}^B(\mathbf{1}), \\ \bar{h}_{i,L}^C &= 24\kappa'_L \sum_{j \in \mathcal{N}} \tau_{ij,L}^C(\mathbf{1}). \end{aligned}$$

To see these relations, observe that $24 \sum_{j \in \mathcal{N}} \tau_{ij,L}^a(\mathbf{1})$ represents the time (in hours) individuals from subpopulation N_i^a spend on leisure outside home. Scaling this quantity by κ_L for subpopulations N_i^A, N_i^B (or κ'_L for subpopulation N_i^C) yields the average leisure time (both outdoors and at home) for that subpopulation. Similarly, $\tau_{i,\emptyset,L}^a(\mathbf{1})$ (i.e., the average fraction of time spent at home—in the absence of closures—for leisure by individuals in subpopulation N_i^a) satisfies

$$\tau_{i,\emptyset,L}^a(\mathbf{1}) = \frac{\bar{h}_{i,L}^a}{24} - \sum_{j \in \mathcal{N}} \tau_{ij,L}^a(\mathbf{1}),$$

because $24(\tau_{i,\emptyset,L}^a(\mathbf{1}) + \sum_{j \in \mathcal{N}} \tau_{ij,L}^a(\mathbf{1}))$ gives the average daily leisure time (in hours), $\bar{h}_{i,L}^a$. The parameters $\{\tau_{ij,L}^a(\mathbf{1}) = \tau_{ij,L}^a\}$ were obtained in Section 3; using them together with κ_L, κ'_L we derive $\{\bar{h}_{i,L}^a, \tau_{i,\emptyset,L}^a(\mathbf{1})\}$ as well.

Note that (7) implies that for each neighborhood i , subpopulation a , and action $\ell \in \mathcal{L}_i$, we have

$$\log(\tau_{i,\ell,L}^a(\mathbf{1})) - \log(\tau_{i,\emptyset,L}^a(\mathbf{1})) = u_{i,\ell}^a(\mathbf{1}) - u_{i,\emptyset}^a(\mathbf{1}). \quad (8)$$

Using (6), we can express the right-hand side as a linear function of V , H_1 , H_2^a , and K . The coefficients of

these linear terms are known and obtained from data, and the left-hand side of (8) is also derived from the data as described above.

Similar to MNL models, our model is identified up to a normalization constant; that is, if a constant is added to V, H_1, H_2^a , the choice probabilities/fractions remain unchanged. Thus, we first set one of the parameters $V(i^*)$ (where i^* corresponds to MODZCTA 11215) equal to zero and estimate the remaining parameters of our model by solving the following ordinary least squares problem:

$$\min_{\{\{V(i)\}_{i \neq i^*}, \{H_1(i)\}_i, \{H_2^a(i)\}_i, K\}} \sum_{i \in \mathcal{N}} \sum_{a \in \mathcal{A}} \sum_{\ell \in \mathcal{L}_i} \left(\log \left(\tau_{i\ell,L}^a(\mathbf{1}) \right) \right. \\ \left. - \log \left(\tau_{i\ell,L}^a(\mathbf{1}) \right) - u_{i,\ell}^a(\mathbf{1}) + u_{i,\ell}^a(\mathbf{1}) \right)^2$$

All parameters estimated in this linear regression are statistically significant (with all p -values less than 1%) and the adjusted R^2 is 0.751. The summary table for the regression results is included in Table 3.

We observe from the mobility data that the initial city-wide restrictions of economic activity reduced the aggregate leisure time spent outside home to roughly 10% of the total leisure time. We use this observation to choose the normalization constant such that if $x_i = 0$ for all i , the leisure time spent outside home matches this observation. We obtain our final estimates by adding this normalization constant to our estimates of V, H_1 , and H_2^a . We see that the resulting parameters are such that the $H_1(j)$ parameters are nonnegative for all but 14 neighborhoods; for the remaining neighborhoods, the aforementioned parameter is very close to zero. The parameters other than H_1 are also nonnegative. To ensure that $u_{ij}^a(x)$ is weakly increasing in x_j , in the remainder of our analysis we set the latter set of $H_1(j)$ parameters (which are close to zero) to zero.

The estimated parameters, together with (6) and (7), give us a model of trip substitutions in response to closures. We next analyze how the planner can use this model and obtain (coarse) targeted restrictions, in settings with trip substitutions.

5.3. Coarse Targeting

In the remainder of the paper, we focus on our model that allows trip substitutions. Under this model, the dependence of new infections on permitted activity levels is nonlinear. As a result, the constraint $\sum_{a \in \mathcal{A}} E_i^{a'}$

$+ I_i^{c,a'} + I_i^{sc,a'} \leq 0$ becomes nonconvex in the decision variables and the resulting variant of (P1) is less tractable. At the same time, this constraint captures only the instantaneous changes in infections; in practice, it may be more reasonable to focus on policy goals that ensure a reduction in active cases in a certain time frame (e.g., two weeks). Finally, targeted interventions at the fine-grained MODZCTA level may pose difficulties in implementation. We simultaneously address these issues and obtain targeted policies that take into account trip substitutions.

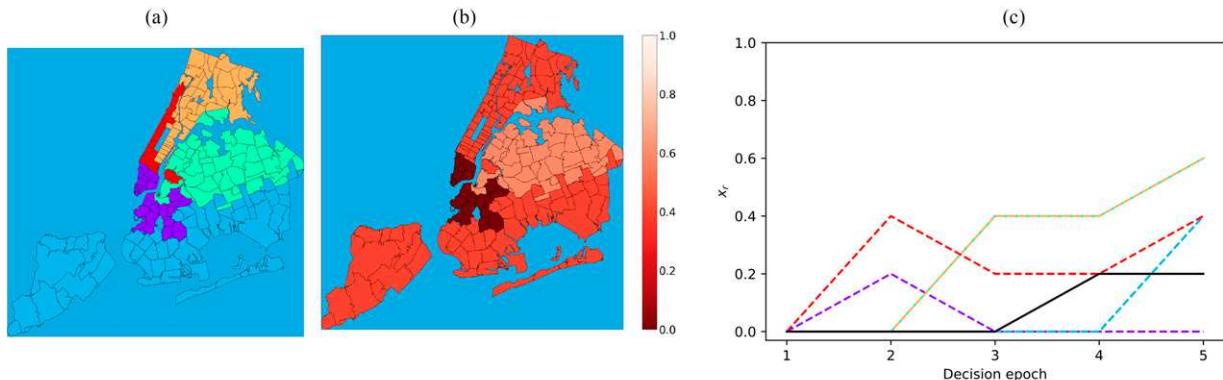
Though the targeting structure illustrated in Figure 1 is obtained ignoring trip substitutions, it still captures information on infections and employment characteristics in different neighborhoods as well as population movements. Thus, we consider coarse targeted policies by clustering MODZCTAs whose $\{x_i\}$ are similar, using a geographical clustering algorithm: max-p-regions (Duque et al. 2012). In the max-p-regions problem, the goal is to cluster n areas into an unknown maximum number of regions. Each area has a vector of attributes, and a distance function measures how different two vectors of attributes are. The heterogeneity of a region is the total distance between every pair of areas within that region. Moreover, there is an additional attribute, referred to as the spatially extensive attribute, where the feasible partitions of areas into regions satisfy (i) a spatial contiguity condition and (ii) a minimum threshold value condition that requires that the sum of spatially extensive attributes of the areas contained in each region is greater than a prespecified threshold. The latter condition can be used to ensure that no region is too small or large in size. The objective is to obtain a feasible partition that has the largest number of partition elements and ensures the smallest total heterogeneity among all feasible partitions with the same cardinality. Duque et al. (2012) present both an exact algorithm (using mixed-integer programming) and a heuristic algorithm (using local search) for the solution of the max-p-regions problem.

We formulate the max-p-regions problem to cluster MODZCTAs into several geographically connected³² regions (condition (i)) such that the number of nonteleworkable jobs in each region is at least 1/6 of the total number of nonteleworkable jobs in NYC (condition (ii)). We use $\{x_i\}$ in the solution of our baseline model as the attribute that is used to identify similar regions and define the attribute “distance” between two neighborhoods as the absolute difference of their attributes. Because of running-time considerations, we implement the heuristic algorithm of Duque et al. (2012) and get the clustering structure given in Figure 3(a). This partition \mathcal{R} has five regions (equal to the number of boroughs in NYC). Every region is geographically connected and accounts for at least 1/6 of the nonteleworkable employment in NYC. It can be seen that the

Table 3. Summary Table for Regression Results

Dep. variable:	y	R ²	0.755
Model:	OLS	Adj. R ² :	0.751
Method:	Least squares	F-statistic:	190.2
No. observations:	46,917	Prob(F-statistic):	0.00
Df residuals:	46,169	Log-likelihood:	-60,409
Df model:	747	AIC:	1.223×10^5
Covariance type:	Nonrobust	BIC:	1.289×10^5

Figure 3 (Color online) Clustering Structure, Optimal Static Solution with Trip Substitutions, and Dynamic Policies for $y = 0.5$



Notes. In (c), the solid curve is the optimal uniform policy; and the dashed curves define the greedy targeted policy. (a) Clustering structure. (b) Optimal solution w/substitutions. (c) Dynamic policies.

clustering structure largely follows the optimal solution in Figure 1, with some deviations due to geographical contiguity requirements.

We focus on coarse targeted policies where all neighborhoods in region $r \in \mathcal{R}$ have the same level of permitted activity, x_r . Motivated by practical implementation considerations, we restrict attention to a grid of permitted economic activity levels $\mathcal{G} = \{0\%, 20\%, \dots, 100\%\}$ for the targeted policy. Our policy goal is to ensure a reduction in active cases $\sum_{a \in \mathcal{A}} E_i^a + I_i^{c,a} + I_i^{sc,a}$ in all neighborhoods $i \in \mathcal{Z}$ in a time frame of two weeks (while keeping the permitted activity levels the same throughout). Specifically, for an arbitrary policy $\mathbf{x} = \{x_r\}$, let $A_{i,t}(\mathbf{x})$ denote the active infections in neighborhood i on day t (or more precisely t days after the restrictions are imposed, where, as before, we assume that the policy is employed starting from April 18, 2020). Our policy goal requires $A_{i,14}(\mathbf{x}) \leq A_{i,0}(\mathbf{x})$ for each MODZCTA $i \in \mathcal{Z}$. The optimal (coarse targeted) policy is the one that achieves this policy goal while yielding minimal nonteleworkable unemployment. This policy is found via exhaustive search (over five regions and six permitted economic activity levels) and is given in Figure 3(b).³³ As before, targeting yields substantially higher efficiency: in the scenario $y = 0.5$, the efficiency of the optimal coarse targeted solution is 34.0%, whereas the optimal uniform solution (obtained after we require $x_r = x_{r^*}$ for $r, r^* \in \mathcal{R}$) is 12.5%. The targeting structure exhibits similarities to the baseline model: no neighborhood in Manhattan is open at a higher level than Midtown, and there are neighborhoods outside Manhattan where even more economic activity is allowed. Because of coarse targeting, the permitted economic activity levels are more restricted and hence less extreme than before: they vary between 0 and 0.6 (as opposed to zero and one).

5.4. Dynamic Policies

Next, we focus on a setting where the planner decides how to control the spread of the epidemic over 10 weeks by adjusting his or her targeting decisions every two weeks (i.e., there are five decision epochs). We denote by $x_{r,t}$ the level of permitted activity in region $r \in \mathcal{R}$ on day t . We allow for trip substitutions and once again restrict attention to coarse targeted policies (such that $x_{r,t} \in \mathcal{G}$). For an arbitrary policy $\mathbf{x} = \{x_{r,t}\}$, as before, we denote by $A_{i,t}(\mathbf{x})$ the active infections in neighborhood i on day t .

We impose two types of constraints on our policies. The first type of constraint requires that the number of active infections be smaller at the end of each epoch than in the beginning for each MODZCTA $i \in \mathcal{Z}$, that is,

$$A_{i,t+14}(\mathbf{x}) \leq A_{i,t}(\mathbf{x}), \quad \forall i \in \mathcal{Z}, t \in \{0, 14, \dots, T-14\}, \quad (9)$$

where $T=70$ denotes the last day of the planning horizon. The second type of constraint requires that the city-wide active infections under dynamic policy \mathbf{x} be not much larger than those under the full closure policy (which sets $x_{r,t} = 0$ for all r and t), denoted by $\mathbf{0}$. More precisely, we require that

$$\sum_{i \in \mathcal{Z}} A_{i,t}(\mathbf{x}) \leq \sum_{i \in \mathcal{Z}} A_{i,t}(\mathbf{0}) + b \cdot \sum_{i \in \mathcal{Z}} A_{i,T/2}(\mathbf{0}), \quad \forall t \in \{0, \dots, T\}. \quad (10)$$

Here, $\sum_{i \in \mathcal{Z}} A_{i,t}(\mathbf{0})$ denotes the total active infections in the city under the full closure policy on day t . The variable b is a fixed constant, and $\sum_{i \in \mathcal{Z}} A_{i,T/2}(\mathbf{0})$ represents the active infections in the middle of the horizon under the full closure policy. When $b=0$, the only feasible policy is the full closure policy; $b=\infty$ yields policies that only enforce nonincreasingness in infections in each neighborhood. As b increases, infections

increase and unemployment decreases. By choosing $b > 0$ appropriately, different trade-offs between these can be obtained. For instance, to ensure swift reduction in infections (and in turn disease-induced deaths), small values of b can be adopted.

Finding the optimal dynamic targeted policy appears intractable (because of the high-dimensional state space and large number, 6^{25} , of different policies). We instead obtain the optimal dynamic uniform policy (which requires $x_{r,t} = x_{r^t,t}$ for $r, r^t \in \mathcal{R}$ and all t) via exhaustive search. In addition, we compute the greedy dynamic targeted policy, which at each decision epoch finds the targeting structure that satisfies the policy goal while inducing the lowest nonteleworkable unemployment level (at that epoch). Figure 3(c) shows the greedy dynamic targeted policy for $b = 10\%$ in the scenario $y=0.5$. Our greedy dynamic targeted policy yields much higher efficiency (averaged over time) than the optimal dynamic uniform policy: 18.0% versus 8.0%, respectively, highlighting the power of targeting even in the dynamic variants of the problem.³⁴

Despite its desirable performance against the optimal dynamic uniform policy, the greedy dynamic targeted policy is not necessarily the optimal dynamic targeted policy. Solving for the latter is computationally hard; thus, we also explore a more limited scenario in which only coarser targeted policies (with three regions over three epochs) are allowed and shed light on the efficiency gap between the greedy dynamic targeted policy and the optimal dynamic targeted policy. Specifically, following a similar approach to the one above (but requiring each region to have at least 1/4 of all the nonteleworkable employment), we obtain a coarser clustering structure that divides NYC into three regions (see Figure 4(a)). We consider dynamic policies over six weeks with three decision epochs (once again adjusting closure decisions every two weeks). We focus on the same types of constraints with both $b = 10\%$ and $b = 20\%$. The small dimension of this example enables us to compute the optimal dynamic targeted policy via exhaustive search. For $b = 10\%$ (Figure 4(b)), the greedy dynamic targeted policy is indeed the optimal dynamic targeted policy and yields higher efficiency than the optimal dynamic uniform policy: 10.1% versus 6.7%, respectively. For $b = 20\%$ (Figure 4(c)), both the optimal dynamic targeted policy and the greedy dynamic targeted policy yield relatively higher efficiency than the optimal dynamic uniform policy: 18.5%, 16.8% versus 13.3%, respectively. Moreover, in both cases the greedy dynamic policy is near optimal. Two additional comments about these results are worth highlighting. First, because this example considers very coarse targeting structures (i.e., three regions in NYC), the efficiency gap between the greedy dynamic targeted policy and the optimal dynamic uniform policy is smaller than the one in our five-region five-epoch

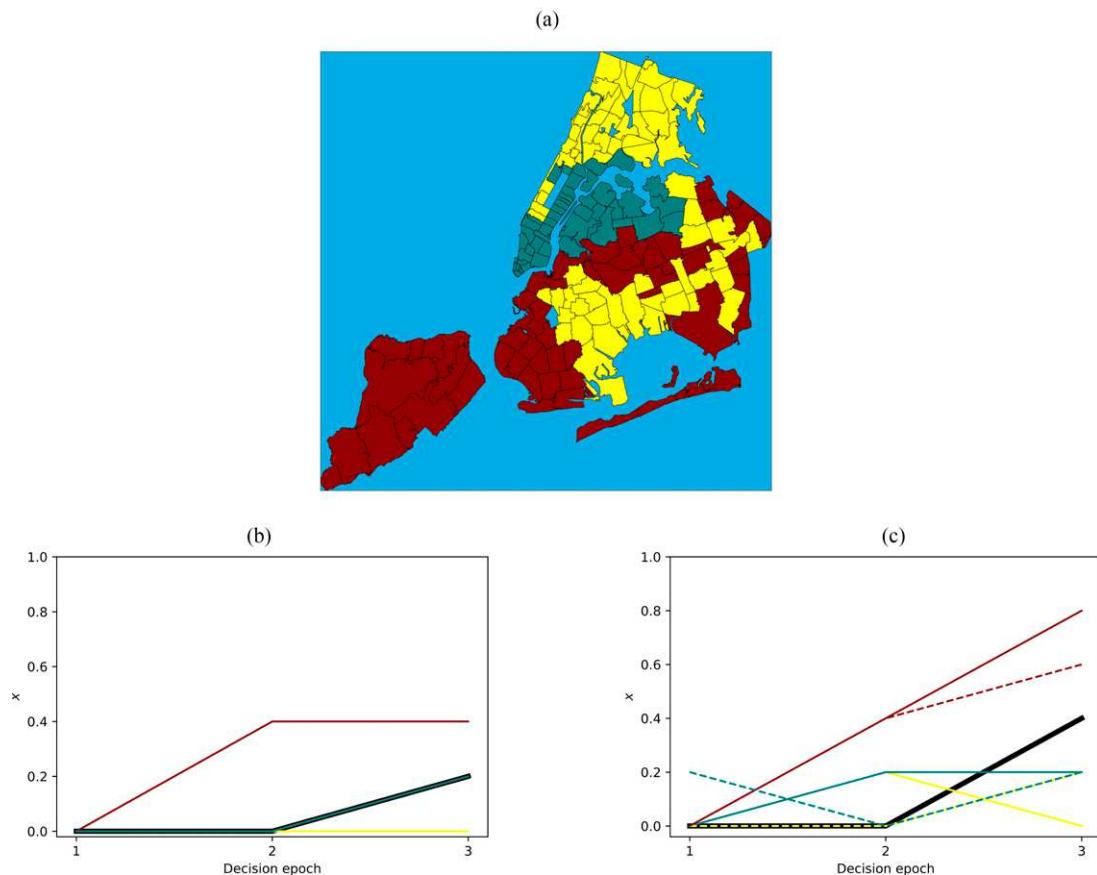
example. Second, interestingly, the optimal dynamic targeted policy also exhibits a nonmonotone structure (see, e.g., the solid yellow curve in Figure 4(c)).

We also analyze the performance of the greedy dynamic targeted policy and the optimal dynamic uniform policy under model misspecification. In this analysis, we capture the possibility that (i) the realized economic activity in each neighborhood may deviate from what is required by the planner and (ii) there is uncertainty in both the disease parameters and other data sources (i.e., initial infection numbers and mobility patterns). Specifically, we assume that in each neighborhood i , the realized economic activity is a noisy version of the activity level x_i set by the planner (using the greedy dynamic targeted policy or the optimal dynamic uniform policy), the actual values of the disease parameters differ from our assumptions (and are sampled from the confidence intervals reported in the literature), and only noisy information is available on the initial infection numbers/mobility patterns. We show that our greedy dynamic targeted policy (which does not know the actual realization of the model parameters and incorrectly assumes that the realized economic activity levels match the levels set by the planner) continues to achieve much higher efficiency than the uniform policy, while having a very similar “risk profile,” that is, inducing a very similar distribution of total infections. In other words, our findings seem to be robust with respect to model misspecification. We detail these results in Section EC.3.10.

6. Conclusions

This paper illustrates that targeted closure policies have the potential of reaching the same goals in terms of epidemic control as uniform closure policies but at substantially lower economic costs (due to unemployment) than uniform policies. In addition, the paper highlights that population movement within a region makes the optimal targeted policies quite different from policies built on local infection measures and that practices in adjoining regions may have a significant effect on the efficacy and design of epidemic control policies. This analysis opens up a number of interesting research directions. First, our observation of the importance of considering population movement patterns and costs associated with the control of the disease provide motivation for the control of infectious diseases in different settings from those considered here. Indeed, subsequent research has started exploring this direction (e.g., see Ma et al. 2020). Second, it is not a priori clear for which type of population movement patterns targeting becomes more valuable (relative to uniform policies); a theoretical analysis of this question is an interesting research direction. Third, although optimal dynamic control of infections appears intractable because of the large scale of the problem and

Figure 4. (Color online) Dynamic Policies with Three Epochs and Three Regions



Notes. (a) Three-region clustering structure. (b) Dynamic policies with $b = 10\%$. The thick curve is the optimal uniform policy, and the thin curves define the optimal targeted policy and the greedy targeted policy. (c) Dynamic policies with $b = 20\%$. The thick curve is the optimal uniform policy, and the thin dashed (resp. solid) curves define the greedy (resp. optimal) targeted policy.

the nontrivial evolution of the state, we proposed a greedy dynamic targeted policy and illustrated that it already provides a substantial improvement over the optimal dynamic uniform policy. Providing algorithms for obtaining better dynamic policies with provable performance guarantees remains an interesting research direction. Fourth, other stronger restrictions of individual activities (e.g., strict shelter-in-place orders or neighborhood quarantines) can be used to curb the spread of an epidemic. Providing an optimization framework for the design of such policies, exploring their performance, and shedding light on when such stronger controls provide substantial improvement over the less stringent ones we considered here are additional interesting future directions.

Our objective in this paper was not to provide a definitive decision support tool for targeted interventions but rather to shed light on the value of targeting and provide insights into its structure. A natural and important direction is to refine the model, assumptions, and optimization framework of this paper in order to provide a decision support tool for social

planners. To accomplish this, several practical considerations need to be addressed. For instance, implementing policies such as ours requires keeping track of the sizes of disease compartments (at each decision epoch), and relevant data may not always be available. Also, it may not be easy to restrict the activity in a neighborhood to a specified level. Finally, social planners may have other considerations beyond the trade-off between infection and unemployment rates. Exploring these and other practically relevant considerations, and adapting the targeted policies discussed in this paper to handle those issues, continue to be interesting future directions.

Acknowledgments

The authors thank SafeGraph for providing the mobility data.

Endnotes

¹ New York State Governor's Office (2020d) states that "instead of analyzing data by region, county, or even just ZIP, the micro-cluster strategy will use granular data to pinpoint the epicenters of viral

outbreaks in neighborhoods and smaller areas.” Notably, these policies showed promise in curbing the spread of the epidemic, at least locally: after keeping this targeted policy in place for two weeks, the positivity rate in targeted locations decreased substantially; for example, for Brooklyn and Rockland, the positivity rates before/after this intervention are respectively 5.88%/3.57% and 9.77%/3.21%; see New York State Governor’s Office (2020c).

² Although we frame our discussion in terms of the economic costs of lockdown-induced unemployment, it is worth noting that besides the economic costs, unemployment can also cause adverse health outcomes; see, for example, *The Hill* (2020).

³ Interestingly, they find that the change in mobility also depends on the socioeconomic characteristics of individuals. For both high- and low-income zones of the city, the reduction in mobility is substantial; but high-income zones exhibit more sizable reductions in mobility than lower-income ones. One explanation the authors provide for this is the fact that high-income individuals may have professions that are amenable to remote work. The authors also document that there are nontrivial spillovers, that is, lockdowns in a region impact other regions.

⁴ This parameter controls the aggregate capacity at which the businesses in a neighborhood operate. In practice, capacity reduction in a neighborhood can be achieved by, for example, (i) closing a fraction of businesses in a neighborhood (possibly by prioritizing essential businesses), (ii) limiting the number of people who can be present in businesses, and (iii) restricting the hours of operation.

⁵ The insights hold and the framework applies even if it is assumed that individuals cannot work remotely; see Section EC.3.5.

⁶ In practice, some infected individuals may not realize that they are sick or they may have only mild symptoms that they attribute to other causes. Motivated by this, we take a conservative approach and assume that they do not self-isolate in the absence of restrictions. It is straightforward to model self-isolation by scaling the populations of the infected compartments.

⁷ It can be seen that the closures in neighborhood j impact the infections among individuals who reside in other neighborhoods but visit j . In turn, provided that the latter neighborhoods do not experience full closures, this impacts the subsequent infections that take place there. With full closures in j (i.e., when $x_j=0$), our baseline model assumes that there will be no infections in j . In the case of full closures, in practice, one could still expect individuals (from all neighborhoods, including j) to spend a limited amount of time in j , which causes a small number of infections there. Our trip substitution model already accounts for this behavior; see Section 5.

⁸ In our NYC example, \mathcal{Z} corresponds to all modified Zip Code Tabulation Areas (MODZCTAs) in NYC and \mathcal{N} corresponds to all MODZCTAs in NYC and neighboring counties (see Table 2) that have the largest inflow/outflow of individuals to/from NYC. See Section 3 for further discussion.

⁹ The change in compartment sizes in a short amount of time can be approximately given by the right-hand sides of (2)–(4) times the duration. Thus, the aforementioned condition allows for controlling the evolution of the disease in the short term. Over time the susceptible population shrinks, and each infected individual can cause fewer infections. Thus, a (static) policy that initially satisfies this condition is also likely to achieve a reduction in infections in the long run. This observation suggests that our condition also offers a tractable proxy for the long run.

¹⁰ SafeGraph is a data company that aggregates anonymized location data from numerous applications in order to provide insights into physical places.

¹¹ In Section EC.2, we conduct a model validation exercise and verify that the model we detail here has reasonable predictions.

¹² Li et al. (2020) divide the infected population into two groups: (i) documented infected individuals with symptoms severe enough to be confirmed and (ii) the rest. They refer to the rate at which infected individuals belong to the first group as the reporting rate and estimate it to be 0.14. We assume that this quantity is equal to the clinical rate and use their estimate. The estimates reported here are also largely consistent with other available estimates in the literature; see, for example, Wu et al. (2020) and Ferretti et al. (2020).

¹³ The case data prior to June 2020 are also available at the ZCTA level. The aggregation of the data to the MODZCTA level allows more stable analysis and reporting of case numbers, and more recent data are reported only at the MODZCTA (or coarser) level (see NYC Health 2020). Motivated by this, we conduct our analysis by focusing on the MODZCTA level data. An earlier version of this paper was based on the ZCTA level data and obtained similar insights to those provided in this version; see Birge et al. (2020).

¹⁴ ZCTAs are closely related to zip codes, and in fact the ZCTA code of a block is the most common zip code contained in it; see U.S. Census Bureau (2020).

¹⁵ Here we aggregate data at the county level, as outside NYC the infection data were provided at the county level (New York Times 2020).

¹⁶ We also explore alternative specifications such as one where, for each i, j , the employees who live in i and work in j , are allocated equally to different NAICS industries. We verify that our insights remain unchanged.

¹⁷ For instance, $\{r_i^{(k)}\}$ account for 96% of all the employees in NYC, whereas the remaining employees live in NYC but work at locations outside NYC or the 10 counties mentioned above, and similarly $\{c_i^{(k)}\}$ account for 93% of all the jobs in NYC.

¹⁸ The data exhibit some day-of-the-week effect; for example, on certain days of the week fewer tests are done and cases reported. When constructing the initial sizes of infected populations, we focus on both the reported case counts and a version that is smoothed using a five-day window (to sidestep the aforementioned issue) and obtain similar results. In the remainder of the paper, we report the results using smoothed data.

¹⁹ Similarly, Hortaçsu et al. (2021) estimate that 4%–14% of the actual number of infections had been reported in the United States up to March 16.

²⁰ The antibodies appear with some delay after the infection. According to Wölfel et al. (2020), half of the infected mild cases developed antibodies within a week of the onset of symptoms, whereas the rest developed antibodies in the second week. Even though 24.7% was reported on April 27, because of the aforementioned delay as well as the time spent surveying the patients, it is reasonable to expect a similar infection rate earlier as well, consistent with our assumption on T_* .

²¹ We distribute the origins according to the populations of the CBs, whereas we distribute the destinations according to the areas. This is because the number of “homes” in a CB is naturally proportional to the population of that CB; hence, when allocating origins of trips to CBs, it is natural to focus on populations. On the other hand, the destinations of the trips can be to nonresidential areas and focusing on populations in that case could be misleading. Therefore, in the latter case, we focus on the areas of CBs. We also ran our model when the origins/destinations are both distributed according to areas or populations. We obtained similar results to our main model and verified the robustness of our findings.

²² We also do sensitivity analysis by considering different time windows and obtain similar findings; see Section EC.3.3.

²³ More precisely, the surveys project that on an average day (i) 112,700K individuals work (in the United States) and they spend

(on average) 7.62 hours on work and (ii) 92,274K individuals work at their workplace (at least partially) and they spend (on average) 7.86 hours working there. This implies that on an average (pre-epidemic), a total of $112700K \times 7.62$ hours of work is done and $92274K \times 7.86$ hours of work is done at the workplace. Hence, taking the ratio we conclude that approximately 84.4% of work is done at the workplace. We assume that workers in NYC work for 40 hours a week on average. Although not specific to NYC, similar numbers are reported by the ATUS: employed individuals work eight hours on an average weekday, whereas this number is 7.94 hours for single job holders (we also repeated our analysis using these numbers as well as 7.62 hours of work on an average work day mentioned earlier and obtained very similar results). This assumption together with our earlier observation implies that $40 \times 84.4\% = 33.76$ hours of work is done at the workplace in an average week by an employed individual. This in turn implies that on average an employed individual spends $33.76/7 = 4.82$ hours at the workplace per day. The surveys also directly report the time spent on different leisure activities by employment characteristics (see U.S. Bureau of Labor Statistics (2020), ATUS table 11A). The quantities h_L and h'_L are obtained by combining the time spent on various “outdoor” activities: (i) participating in sports, exercise, and recreation; (ii) socializing and communicating; and (iii) other leisure and sports activities, including travel.

²⁴ The aforementioned approach to obtaining $\{\zeta_{ij,L}\}$ yields a small number of i, j pairs for which $\zeta_{ij,L} < 0$. These i, j mainly correspond to locations between which only a small number of individuals travel in the SafeGraph data set. For such i, j , the left-hand side of (5) is smaller than the first term on the right-hand side. Denote by $\Xi(i)$ the set of j such that the (i, j) tuple exhibits this property. The quantity $\sum_{j \in \Xi(i)} (h_W(N_i^A \zeta_{ij,W}^A + N_i^B \zeta_{ij,W}^B) - 24N_i \tau_{ij})/N_i$, which is a measure of the aggregate inconsistency in node i , is bounded by 32 minutes for all i . In other words, the maximum inconsistency (over all nodes) is relatively small (2.2% of the daily window we use in our analysis). It can be attributed partly to the assumptions we made to deal with data unavailability (e.g., subpopulations N_i^A, N_i^B, N_i^C have identical leisure-trip factors $\{\zeta_{ij,L}\}_j$) and partly to the sampling errors in the SafeGraph data. In addition, U.S. Census Bureau (2017) data are a few years old and the employment patterns can change over time and have additional noise. Given that 2.2% of daily time is relatively small, we manually adjust those edges by setting $\zeta_{ij,L} = 0$ and $\zeta_{ij,W}^A = \zeta_{ij,W}^B = \tau_{ij} \times (24N_i/h_W(N_i^A + N_i^B))$ and derive $\{\tau_{ij,W}^A, \tau_{ij,L}^A\}$ accordingly.

²⁵ It is interesting to explore whether the optimal policy permits more economic activity in neighborhoods where individuals have higher incomes on average. We find no evidence for this. See Section EC.2 for details.

²⁶ In this and other variants of (P1), when the additional parameter is set equal to zero we recover the baseline model. In all variants we study, whenever we solve for the uniform policies we impose the constraint $x_i - x_j \leq 1 - \theta$ for $\theta = 1$ for $i, j \in \mathcal{Z}$.

²⁷ We say that MODZCTAs i and j are neighbors if (i) they share a border or (ii) they are separated by the East or the Harlem River and i is geographically the closest MODZCTA to j on the other side of the river (or vice versa) or (iii) they are connected by a bridge.

²⁸ This formulation is also useful to obtain robust solutions that ensure that small perturbations to the parameters do not violate the original policy goal.

²⁹ Though we verify that if this term takes the same value for all subpopulations, the change in the results is marginal.

³⁰ When i, ℓ both correspond to MODZCTAs or neighboring counties, we measure the distance by focusing on the distance between their centroids. When one of them is a MODZCTA and the other is a neighboring county, we measure the distance between the centroid of the MODZCTA and the closest point (at the border) of the neighboring county. We experiment with other utility specifications

and verify that the above functional form yields the best fit, and the results are not sensitive to the way the distances are measured (using centroids versus borders).

³¹ In light of the pandemic, there may have been other changes in individuals’ trip patterns, triggered, for example, by aversion to public transit. Although we do not explicitly model such changes here, we verified that the majority of leisure trips are of a local nature (even prior to the pandemic); hence, they are unlikely to be drastically impacted by attitudes toward public transit.

³² Here we use the same adjacency relation as the one in Endnote 29.

³³ Specifically, for each $\{x_r\}_{r \in \mathcal{R}} \in \mathcal{G}^{\mathcal{R}}$, we first numerically solve the system (2)–(4) and compute the active infections at the end of a two-week horizon. We then identify solutions that achieve a reduction in cases and among these find the solution that yields minimal nonteleworkable unemployment.

³⁴ Despite the fact that this policy is nonmonotone (in the sense that some regions may face harsher restrictions after initially opening up), a monotone variant can trivially be obtained by setting $\hat{x}_{r,t} = \min_{t' \geq t} x_{r,t'}$, where $x_{r,t}$ and $\hat{x}_{r,t}$ respectively denote the permitted economic activity in region r at epoch t under the greedy dynamic policy and its monotone variant. Interestingly, the efficiency of the monotone variant is also substantially larger than that of the optimal dynamic uniform policy: 16.2%.

References

Acemoglu D, Chernozhukov V, Werning I, Whinston MD (2020) A multi-risk SIR model with optimally targeted lockdown. NBER Working Paper No. 27102, National Bureau of Economic Research, Cambridge, MA.

Allen LJ, Brauer F, Van den Driessche P, Wu J (2008) *Mathematical Epidemiology*, vol. 1945 (Springer, Berlin).

Anderson RM, Anderson B, May RM (1992) *Infectious Diseases of Humans: Dynamics and Control* (Oxford University Press, Oxford, UK).

Baker SR, Bloom N, Davis SJ, Terry SJ (2020b) COVID-induced economic uncertainty. NBER Working Paper 26983, National Bureau of Economic Research, Cambridge, MA.

Baker SR, Farrokhnia RA, Meyer S, Pagel M, Yannelis C (2020c) How does household spending respond to an epidemic? Consumption during the 2020 COVID-19 pandemic. *Rev. Asset Pricing Stud.* 10(4):834–862.

Baker SR, Bloom N, Davis SJ, Kost K, Sammon M, Virayosin T (2020a) The unprecedented stock market reaction to COVID-19. *Rev. Asset Pricing Stud.* 10(4):742–758.

Batlle P, Bruna J, Fernandez-Granda C, Preciado VM (2020) Adaptive test allocation for outbreak detection and tracking in social contact networks. Preprint, submitted November 3, <https://arxiv.org/abs/2011.01998>.

Birge JR, Candogan O, Feng Y (2020) Controlling epidemic spread: Reducing economic losses with targeted closures. Working Paper No. 2020-57, Becker Friedman Institute for Economics, University of Chicago, Chicago.

Brauer F, Castillo-Chavez C, Feng Z (2019) *Mathematical Models in Epidemiology* (Springer, Berlin).

Budish EB (2020) R<1 as an economic constraint: Can we ‘expand the frontier’ in the fight against COVID-19? Working Paper No. 2020-31, Becker Friedman Institute for Economics, University of Chicago, Chicago.

Budish E, Kashyap A, Koijen R, Neiman B (2020) Three pillars of the economic policy response to the COVID-19 crisis. Working paper, Becker Friedman Institute for Economics, University of Chicago, Chicago.

Carranza A, Goic M, Lara E, Olivares M, Weintraub GY, Covarrubia J, Escobedo C, Jara N, Basso LJ (2020) The social divide of social distancing: Shelter-in-place behavior in Santiago during

the COVID-19 pandemic. Preprint, submitted September 14, <https://dx.doi.org/10.2139/ssrn>.

CDC (2020) CDC COVID data tracker. Accessed September 24, 2021, <https://covid.cdc.gov/covid-data-tracker>.

Chang S, Pierson E, Koh PW, Gerardin J, Redbird B, Grusky D, Leskovec J (2020) Mobility network models of COVID-19 explain inequities and inform reopening. *Nature* 589:82–87.

Chen X, Li M, Simchi-Levi D, Zhao T (2020) Allocation of COVID-19 vaccines under limited supply. Preprint, submitted August 31, <https://dx.doi.org/10.2139/ssrn.3678986>.

Chinazzi M, Davis JT, Ajelli M, Gioannini C, Litvinova M, Merler S, Piontti AP, et al. (2020) The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* 368(6489):395–400.

Coibion O, Gorodnichenko Y, Weber M (2020) Labor markets during the COVID-19 crisis: A preliminary view. Technical report, National Bureau of Economic Research, Cambridge, MA.

Congressional Research Service (2020) Unemployment rates during the COVID-19 pandemic: In brief. Accessed September 24, 2021, <https://fas.org/sgp/crs/misc/R46554.pdf>.

Dingel JI, Neiman B (2020) How many jobs can be done at home? *J. Public Econom.* 189:104235.

Duque JC, Anselin L, Rey SJ (2012) The max-p-regions problem. *J. Reg. Sci.* 52(3):397–419.

Duque D, Morton DP, Singh B, Du Z, Pasco R, Meyers LA (2020) Timing social distancing to avert unmanageable COVID-19 hospital surges. *Proc. Natl. Acad. Sci. USA* 117(33):19873–19878.

Evgeniou T, Fekom M, Ovchinnikov A, Porcher R, Pouchol C, Vayatis N (2020) Epidemic models for personalised COVID-19 isolation and exit policies using clinical risk predictions. INSEAD Working Paper No. 2020/22/DSC, INSEAD, Fontainebleau, France.

Fernández-Villaverde J, Jones CI (2020) Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. NBER Working Paper No. 27128, National Bureau of Economic Research, Cambridge, MA.

Ferretti L, Wymant C, Kendall M, Zhao L, Nurtay A, Abeler-Dörner L, Parker M, Bonsall D, Fraser C (2020) Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* 368(6491):eabb6936.

Gershon D, Lipton A, Levine H (2020) Managing COVID-19 pandemic without destroying the economy. Preprint, submitted April 21, <https://arxiv.org/abs/2004.10324>.

Glover A, Heathcote J, Krueger D J RR (2020) Health vs. wealth: On the distributional effects of controlling a pandemic. Working paper, Centre for Economic Policy Research, London.

Gonzalez-Reiche AS, Hernandez MM, Sullivan MJ, Ciferri B, Alshammary H, Obla A, Fabre S, et al. (2020) Introductions and early spread of SARS-CoV-2 in the New York City area. *Science* 369(6501):297–301.

Gormsen NJ, Koijen RS (2020) Coronavirus: Impact on stock prices and growth expectations. *Rev. Asset Pricing Stud.* 10(4): 574–597.

Guerrieri V, Lorenzoni G, Straub L, Werning I (2020) Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? NBER Working Paper No. 26918, National Bureau of Economic Research, Cambridge, MA.

Hanson SG, Stein JC, Sunderam A, Zwick E (2020) Business continuity insurance and loans: Keeping America's lights on during the pandemic. Working paper, Becker Friedman Institute for Economics, University of Chicago, Chicago.

He Z, Liu B (2020) *Dealing with a Liquidity Crisis: Economic and Financial Policies in China During the Coronavirus Outbreak. Impact of COVID-19 on Asian Economies and Policy Responses* (World Scientific, Singapore), 55–63.

Hethcote HW (1976) Qualitative analyses of communicable disease models. *Math. Biosci.* 28(3-4):335–356.

Hethcote HW, Van Ark JW (1987) Epidemiological models for heterogeneous populations: Proportionate mixing, parameter estimation, and immunization programs. *Math. Biosci.* 84(1):85–118.

Holtz D, Zhao M, Benzell SG, Cao CY, Rahimian MA, Yang J, Allen JNL, et al. (2020) Interdependence and the cost of uncoordinated responses to COVID-19. *Proc. Natl. Acad. Sci. USA* 117(33): 19837–19843.

Hortaçsu A, Liu J, Schwieg T (2021) Estimating the fraction of unreported infections in epidemics with a known epicenter: An application to COVID-19. *J. Econometrics* 220(1):106–129.

Hotz VJ, Miller RA (1993) Conditional choice probabilities and the estimation of dynamic models. *Rev. Econom. Stud.* 60(3): 497–529.

Jia JS, Lu X, Yuan Y, Xu G, Jia J, Christakis NA (2020) Population flow drives spatio-temporal distribution of COVID-19 in China. *Nature* 582:389–394.

Li R, Pei S, Chen B, Song Y, Zhang T, Yang W, Shaman J (2020) Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV2). *Science* 368(6490):489–493.

Lipton A, Lopez de Prado M (2020) Exit strategies for COVID-19: An application of the K-SIR model (presentation slides). Preprint, submitted April 22, <https://dx.doi.org/10.2139/ssrn.3579712>.

Ma Q, Liu YY, Olshevsky A (2020) Optimal lockdown for pandemic stabilization. Preprint, submitted October 24, <https://arxiv.org/abs/2010.12923>.

Martcheva M (2015) *An Introduction to Mathematical Epidemiology*, vol. 61 (Springer, Berlin).

McNeill WH, McNeill W (1998) *Plagues and Peoples* (Anchor, New York).

Mulligan CB (2021) Economic activity and the value of medical innovation during a pandemic. *J. Benefit-Cost Anal.* 12(3):420–440.

New York State Governor's Office (2020a) Amid ongoing COVID-19 pandemic, Governor Cuomo announces phase II results of antibody testing study show 14.9% of population has COVID-19 antibodies. Accessed September 24, 2021, <https://www.governor.ny.gov/news/amid-ongoing-covid-19-pandemic-governor-cuomo-announces-phase-ii-results-antibody-testing-study>.

New York State Governor's Office (2020b) Governor Cuomo announces new cluster action initiative. Accessed September 24, 2021, <https://www.governor.ny.gov/news/governor-cuomo-announces-new-cluster-action-initiative>.

New York State Governor's Office (2020c) Governor Cuomo announces Orange County Micro-cluster focus area meets metrics to exit "Red Zone." Accessed September 24, 2021, <https://www.governor.ny.gov/news/governor-cuomo-announces-orange-county-micro-cluster-focus-area-meets-metrics-exit-red-zone>.

New York State Governor's Office (2020d) New York "Micro-cluster" strategy. Accessed September 24, 2021, https://www.governor.ny.gov/sites/governor.ny.gov/files/atoms/files/MicroCluster_Metrics_10.21.20_FINAL.pdf.

New York Times (2020) Coronavirus (COVID-19) data in the United States. Accessed September 24, 2021, <https://github.com/nytimes/covid-19-data>.

NYC Health (2020) COVID-19: Data. Accessed September 24, 2021, <https://www1.nyc.gov/site/doh/covid/covid-19-data.page>.

Pei S, Shaman J (2020) Initial simulation of SARS-CoV2 spread and intervention effects in the continental US. Preprint, submitted March 27, <https://doi.org/10.1101/2020.03.21.20040303>.

Post W, DeAngelis D, Travis C (1983) Endemic disease in environments with spatially heterogeneous host populations. *Math. Biosci.* 63(2):289–302.

SafeGraph (2020) COVID-19 response datasets, social distancing metrics. Accessed September 24, 2021, <https://docs.safegraph.com/docs/social-distancing-metrics>.

The Hill (2020) The COVID-19 shutdown will cost Americans millions of years of life. Accessed September 24, 2021, <https://>

thehill.com/opinion/healthcare/499394-the-covid-19-shutdown-will-cost-americans-millions-of-years-of-life.

U.S. Bureau of Labor Statistics (2020) American time use survey (ATUS)-2019 Results. Accessed September 24, 2021, <https://www.bls.gov/tus/#tables>.

U.S. Census Bureau (2010) TIGER/line with selected demographic and economic data. Accessed September 24, 2021, <https://www.census.gov/geographies/mapping-files/2010/geo/tiger-data.html>.

U.S. Census Bureau (2017) Longitudinal employer-household dynamics (LODES). Accessed September 24, 2021, <https://lehd.ces.census.gov/data/>.

U.S. Census Bureau (2018) The American community survey. Accessed September 24, 2021, <https://www.census.gov/programs-surveys/acs/>.

U.S. Census Bureau (2020) ZIP code tabulation areas. Accessed September 24, 2021, <https://www.census.gov/programs-surveys/geography/guidance/geo-areas/zctas.html>.

Van den Driessche P, Watmough J (2002) Reproduction numbers and sub-threshold endemic equilibria for compartmental models of disease transmission. *Math. Biosci.* 180(1–2): 29–48.

Wölfel R, Corman VM, Guggemos W, Seilmaier M, Zange S, Müller MA, Niemeyer D, et al. (2020) Virological assessment of hospitalized patients with COVID-2019. *Nature* 581:465–469.

Wu JT, Leung K, Leung GM (2020) Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: A modelling study. *Lancet* 395(10225):689–697.