The Spread of COVID-19 in London: Network Effects and Optimal Lockdowns*

Christian Julliard † Ran Shi ‡ Kathy Yuan §

October, 2020

Abstract

We generalise a stochastic version of the workhorse SIR (Susceptible-Infectious-Removed) epidemiological model to account for spatial dynamics generated by network interactions. Using the London metropolitan area as a salient case study, we show that commuter network externalities account for about 42% of the propagation of COVID-19. We find that the UK lockdown measure reduced total propagation by 57%, with more than one third of the effect coming from the reduction in network externalities. Counterfactual analyses suggest that: *i*) the lockdown was somehow late, but further delay would have had more extreme consequences; *ii*) a targeted lockdown of a small number of highly connected geographic regions would have been equally effective, arguably with significantly lower economic costs; *iii*) targeted lockdowns based on threshold number of cases are not effective, since they fail to account for network externalities.

Keywords: COVID-19, networks, key players, spatial modelling, SIR model.

JEL codes: I12, I18, C31, C51, D85.

*We thank seminar participants at the London School of Economics and London Business School.

†Department of Finance, FMG, and SRC, London School of Economics, and CEPR; c.julliard@lse.ac.uk

‡Department of Finance, London School of Economics, r.shi1@lse.ac.uk

§Department of Finance, FMG, London School of Economics, and CEPR; k.yuan@lse.ac.uk

Introduction

Facing the challenges of the COVID-19 pandemic, governments worldwide have resorted to large scale lockdown policies to contain the infectious disease. However, large scale lockdown policies have high economic costs. For example, according to the Office for National Statistics (ONS), UK gross domestic product (GDP) fell in the second quarter of 2020 – the period when the UK was in national lockdown – by 20.4% compared with the previous three months. This is the biggest quarterly decline since comparable records began in 1955. In this context, it is worthwhile to explore whether it is possible to minimise the (expected) size of infected population by optimally choosing the lockdown areas. Our estimation of the COVID-19 epidemic dynamic in London allows us to investigate whether targeted lockdowns might have achieved the same outcome, in containing the epidemic spreading, as the full-scale lockdown. Therefore, this paper aims at providing information for the design of optimal lockdown policies in order to lower their negative economic impact.

To capture the COVID-19 epidemic dynamic in London, we generalise the single-population *deterministic* susceptible-infectious-recovered (SIR) model developed by Kermack and McKendrick (1927).[1] We incorporate *stochastic* transmission within and between the heterogeneous subpopulations of the London boroughs that are connected via the *commuting network*. Our model features multiple groups, allows disease spread dynamics in each of them to be different (*autoregressive effects*), and connects them through commuter network links (*network effects*). In short, we name it the *Network-SIR* model. We allow the model to account for potential confounding variables such as the time-varying testing capacity and positivity ratio – the *endemic effects*. To quantify the impacts of lockdown policies, we parameterise the model to

[1]The core ingredients of SIR models were first formulated by Lowell Reed and Wade Frost in the 1920s, but they were not published.

reflect the changes induced by the nationwide lockdown in the UK on March 23, 2020.

We estimate the model for daily COVID-19 surveillance data across 32 London boroughs. In particular, we decompose the contributing factors of the pandemic spread into three components: 1) local dynamic (intra-location) effects (people spreading the disease over time to others living in the same borough); 2) the network (inter-location) effect (people spreading the disease to other boroughs via their commute to, and interactions at, work locations); 3) the endemic effects (the confounding effects such as e.g., testing policy regimes, as well as borough specific characteristics and unobserved heterogeneity). Our estimation results confirm the essential role of network spillover effects via the commuting network, with existing cases in one London borough transmit the disease to residents in other boroughs. The estimated magnitude of these network effects is large: they contribute to over 42% of all COVID-19 cases in London. In comparison, local within-borough dynamics account for less than 35% of all COVID-19 cases.

According to our estimates, the March 23 nationwide lockdown is effective in containing the spread of COVID-19 cases in London. The policy reduces the number of cases generated by local transmission by around 75% of its pre-lockdown level, lowers the spillovers from workplaces to residential areas by about 12%, and decreases the home-to-work transmissions by as much as 80%. We also find that a reduction in the network transmission channel has occurred about one week before the official lockdown, indicating self-imposed changes in behaviour by the London population in an effort to avoid contracting the disease.

From our estimated stochastic Network-SIR model, a tight and salient upper-bound can be computed for the so-called disease R_0 (basic reproduction number). This number represents the expected number of new cases that are *directly* infected by one existing case within a given time span (one week in our benchmark). The point

estimate is around 1.4 before the nationwide lockdown, and is reduced to about 0.8 afterwards. The reduction is large, and is especially encouraging as the upper bound for R_0 drops well below one after the policy change.

To identify the pivotal boroughs for the spreading of the COVID-19 pandemic, we also derive and estimate the network impulse response functions (NIRFs) motivated by Denbee, Julliard, Li, and Yuan (2020). We find that the Westminster/City of London location, due to its centrality in the commuting network, dominates all other London boroughs in terms of externalities generated: before the lockdown, one additional case at this location generates three new expected cases in the whole greater London area within one week (despite Westminster/City of London ranking only 26th out of 32 boroughs in terms of total local cases before the lockdown).

Finally, we simulate alternative targeted lockdown policies based on the estimated dynamics of the Network-SIR model. When considering a lockdown limited to only one borough, we find that isolating Westminster/City of London – the borough with the largest NIRF, but very few local cases – minimises the total expected number of infected cases. This result indicates that optimal lockdown should not only target areas with large number of cases, but also the locations in the network that are key for the COVID-19 propagation dynamics. Similarly, when constraining the optimal lockdown to two boroughs only, we find that the optimal target areas are Westminster/City of London – again, the borough with the largest NIRF – and Southwark – the borough with the highest number of cases. Therefore, optimal lockdown policies should be based on both the number of cases and the network centrality in transmitting the disease. Furthermore, our simulations suggest that a lockdown of just these two boroughs would have achieved the same outcome, in terms of total infections in the Greater London area, as the actual national lockdown. This finding questions the optimality of the full-scale lockdown and calls for a careful redesign of the lockdown

policies.

Related Literature. There is a heated ongoing discussion on the trade-off between sustaining economic output and saving lives in the context of the COVID-19 pandemic. A cornerstone of this literature is a well-specified disease transmission dynamic. Avery, Bossert, Clark, Ellison, and Ellison (2020) provide an overview of various models for the spread of COVID-19. The most dominant paradigm that emerges from these models is still the SIR framework, and most existing works extend the classical Kermack-McKendrick SIR model to account for richer and finer population categorisations, transmission dynamics, and policy responses. Embedding them into economic models featuring costs of lives or negative impacts on productivity (due to infection) and, sometimes, calibrating these models to epidemiological data, are key characteristics of this literature. Along this line of work, Rowthorn and Toxvaerd (2012) analyse (non-pharmaceutical) prevention and treatment policy responses and show that they are history-dependent. Alvarez, Argente, and Lippi (2020) formulate and solve the planner's dynamic control problem. Garriga, Manuelli, and Sanghi (2020) identify value of death as the major policy determinant. Eichenbaum, Rebelo, and Trabandt (2020) illustrate the equilibrium interactions between economic decision and epidemic dynamics. Jones, Philippon, and Venkateswaran (2020) introduce a healthcare congestion externality in addition to the classic infection externality to the planner's problem. Farboodi, Jarosch, and Shimer (2020) derive an optimal policy which features nonrestrictive and discontinuous social distancing. Acemoglu, Chernozhukov, Werning, and Whinston (2020) introduce heterogeneity in social interaction and productivity among social groups by categorising the whole population by their different risk profiles after contracting the disease (which is mainly driven by age). Acemoglu, Chernozhukov, Werning, and Whinston (2020) calibrate their model and demonstrate the benefits from lockdown policies that are tailored for each group separately. Beyond calibration, there are also estimation exercises for the COVID-19

epidemics. For example, Fernández-Villaverde and Jones (2020) estimate the deterministic version of single-population SIR across multiple countries. Our paper departs from this by estimating the network element in a stochastic environment and demonstrating its relevance and policy potential.

Our paper is also related to the network literature on epidemiology (see e.g. Jackson (2008)[2] and Easley and Kleinberg (2010)[3]). Our optimal lockdown policy targeting either the level and NIRF key players is motivated by Ballester, Calvó-Armengol, and Zenou (2006) and Denbee, Julliard, Li, and Yuan (2020), respectively.

1 A Network SIR Model

This section introduces our network SIR model. We first extend the deterministic single-population SIR model to its stochastic counterpart and derive necessary distributional results for mapping the model to disease surveillance data. Then, we extend the statistical model to multiple subpopulations, with a special account for network effects, formalising a network SIR framework. Finally, we enrich the framework by introducing a parametrisation scheme tailored for analysing the spreading of COVID-19 in London.

1.1 The deterministic SIR dynamics

We begin with notation. For a given population of fixed size N, the triplet $\{S_t, I_t, R_t\}$ represents, respectively, the *cumulative* numbers of susceptible (S), infectious (I), and recovered (R) individuals at time t. Susceptible individuals get infected through mix-

[2]Chapter 7, Sections 7.1,7.2.
[3]Chapter 21.

ing with infectious individuals, featuring the dynamic

$$\dot{S}_t = -\theta_I \frac{I_t}{N} S_t, \tag{1}$$

in which the overdot notation represents derivatives with regard to time; I_t/N is the level of disease prevalence in the current population; θ_I is a parameter measuring the contact rate (times the probability of infection per contact). Infected individuals recover at a rate θ_R, which implies

$$\dot{R}_t = \theta_R I_t. \tag{2}$$

The sum $S_t + I_t + R_t$ equals the total population N, constant by assumption. Therefore, $\dot{S}_t + \dot{I}_t + \dot{R}_t = 0$. Plugging in (1) and (2) and rearranging terms,

$$\dot{I}_t = \left(\theta_I \frac{S_t}{N} - \theta_R \right) I_t. \tag{3}$$

Our analysis will focus on the dynamics of the infected population I_t. We make a simplifying assumption that $S_t/N \approx 1$. For COVID-19 in London, this ratio is greater than 99.7% in our sample period. Treating each London borough separately, this ratio ranges from 99.6% (Brent, four cases per thousand population) to 99.8% (Islington, two cases per thousand population). Based on this assumption, a constant, namely α, defined as $\alpha \triangleq \theta_I - \theta_R \approx \theta_I S_t/N - \theta_R$, is sufficient to capture the dynamics of the infectious and infected subpopulations in our investigation sample. The framework can be generalised to allow, as we do below, time variation in α.

1.2 The stochastic SIR model

To map the deterministic SIR model to the data, we need to introduce probabilistic "error" terms. However, arbitrarily introduced errors such as Gaussian errors may reduce statistical power due to misspecified likelihood functions. The issue is more pronounced when the number of disease incidents is relatively small (compared with the total population size), which is the case in our data (less than three detected cases per thousand people in London by the end of the lockdown period). Thus, we aim to derive this distribution coherently from a well-defined probabilistic analogue to the deterministic SIR dynamics. The natural stochastic extension to the differential equation $\mathrm{d}I_t/\mathrm{d}t \approx \alpha I_t$ is a continuous-time Markov chain of the form:

$$\mathbb{P}\left[y \text{ new infections in } (t,\, t+\mathrm{d}t) \mid I_t = x\right] = \begin{cases} \alpha x \mathrm{d}t, & y = 1 \\ o(\mathrm{d}t), & y \geq 2, \\ 1 - \alpha x \mathrm{d}t - o(\mathrm{d}t), & y = 0 \end{cases}, \tag{4}$$

where $o(\mathrm{d}t)$ satisfies $o(\mathrm{d}t)/dt \to 0$ as $\mathrm{d}t \to 0$. Interpreting this probabilistic statement is straightforward. Given the current number of infective individuals x, within an infinitesimally short time interval, one additional person can contract the disease with probability $\alpha x \mathrm{d}t$. This random process is a simple birth process (also known as the Yule-Furry process, see for example, Grimmett and Stirzaker (2001, p. 250)).

When working with surveillance data of infectious diseases, we only observe numbers of new cases within discrete time intervals (say, one day or one week). Based on the specification of (4), we can solve for the distribution of new case counts within a time interval of length h (instead of length $\mathrm{d}t$ which goes to zero), denoted by

$$p_h(y \mid x) = \mathbb{P}\left[y \text{ new infections in } (t,\, t+h) \mid I_t = x\right].$$

Solving the implied Kolmogorov forward equation (see Grimmett and Stirzaker (2001, p. 250)), the simple birth process in (4) yields an analytical expression of the probability mass function $p_h(y \mid x)$:

$$p_h(y \mid x) = \frac{\Gamma(y+x)}{\Gamma(x)\Gamma(y+1)} \left(e^{-\alpha h}\right)^x \left(1 - e^{-\alpha h}\right)^y, \tag{5}$$

where $\Gamma(\cdot)$ represents the gamma function. The above expression describes a negative Binomial distribution – a mixture of Poisson distributions with mixing of the Poisson rate driven by a gamma distribution. It provides the chance of y successes after exactly x failures in a sequence of independent Bernoulli trials, each having a probability of success $p = 1 - e^{-\alpha h}$. The negative Binomial is an appropriate representation for discrete arrival data over an unbounded positive range whose sample variance exceeds the sample mean. In such cases, the observations are overdispersed with respect to a Poisson distribution (for which the mean is equal to the variance). Since the negative binomial distribution has one more parameter than the Poisson distribution, the second parameter can be used to adjust the variance separately from the mean. Furthermore, it implies that first and second conditional moments are positively correlated – a feature consistent with epidemic dynamics. The probability $p(= 1 - e^{-\alpha h})$ also has a clear interpretation in our context of disease transmission. Let $y = 0$ and $x = 1$ in (5), then p is the probability that an existing disease case infects *at least* one person within a period of length h (since $e^{-\alpha h} = 1 - p = p_h(0 \mid 1)$, implying that $p = 1 - p_h(0 \mid 1) = \sum_{y \geq 1} p_h(y \mid 1)$).

Denote the daily count of new infected individuals as y_t, $t = 1, \ldots, T$. Normalising the length of the time interval h to one day and denoting with $\mathcal{F}_t$ the information available up to time t, we have

$$y_t \mid \mathcal{F}_{t-1} \sim \text{NegBinom}\left(p,\, x_{t-1}\right), \tag{6}$$

where $p = 1 - e^{-\alpha}$. Conditional expectations of new disease cases can be computed directly as

$$\mathbb{E}\left[y_t \mid \mathcal{F}_{t-1}\right] = \mu_t = a x_{t-1} \tag{7}$$

where $a = e^{\alpha} - 1 = p/(1-p)$ can be understood as the odds ratio comparing the probability that an infected person does transmit the disease against that he does not. This simple linear relationship between the conditional expectations of new case numbers and existing infectious individuals is the key modelling assumption that we adopt throughout our analysis. Equation (7) also implies that

$$p = \frac{\mu_t}{x_{t-1} + \mu_t} \tag{8}$$

which offers a simple mapping from the conditional expectations μ_t to the probability parameter p in the negative binomial distribution (5).

From the negative binomial distribution, we can also calculate the conditional variance of new case counts as

$$\operatorname{var}\left[y_t \mid \mathcal{F}_{t-1}\right] = (a+1) a x_{t-1}. \tag{9}$$

If $a + 1 = 1/(1-p) \to 1$, that is $p \to 0$, which means that the disease transmission rate is extremely small, the conditional variance of y_t equals its conditional mean. This corresponds to a Poisson distribution specification for y_t as $y_t \sim \text{Pois}(a x_{t-1})$. Equations (7) and (9) imply a positive correlation between the first and second moments for the number of new infected individuals – a feature consistent with the data under analysis.

The remaining issue is how to determine the number of *actively* infectious cases x_t. This issue arises because people who have been infected might recover (or they may die), as we have initially discussed in the SIR model. We adopt the following

formulation for x_t:

$$x_t = \sum_{\ell=0}^{L-1} \nu(\ell) y_{t-\ell}, \tag{10}$$

which assumes that infected individuals can transmit the disease for L periods. The decay function $\nu(\ell) \in (0,1)$ characterises 'the 'rate of infectivity," in the language of Kermack and McKendrick (1927).[4] That is, for a person who has been infected for ℓ periods, the chance of transmitting the disease to another person is reduced to $100 \times \nu(\ell)$ percent of the initial level. Another way of interpreting this specification is that $100[1 - \nu(\ell)]$ percent of the infected individuals are no more infectious ℓ periods after having the contracted the disease. Taken together, equations (7) and (10) imply an autoregressive dynamic for the number of new infections.

1.3 A stochastic network SIR model

Now, we extend the stochastic model introduced above to account for multiple subpopulations connected in a network. In the context of the COVID-19 spread in London, we treat London boroughs as subpopulations of constant sizes N_i, $i = 1, \ldots, n$. Time-t conditional expectations of new case counts are concatenated into a vector $\boldsymbol{\mu}_t = [\mu_{1t}, \ldots, \mu_{nt}]^\top$, where $\mu_{it} = \mathbb{E}[y_{it} \mid \mathcal{F}_{t-1}]$ is the expected number of new cases in borough i at time t. The distribution of newly infected individuals y_{it} follows the negative Binomial structure in (6)–(8), specified as follows:

$$y_{it} \sim \text{NegBinom}\left(p_{it},\, x_{i,t-1}\right) \quad \text{where} \quad p_{it} = \frac{\mu_{it}}{x_{i,t-1} + \mu_{it}}. \tag{11}$$

To extend the conditional mean equation (7) to the multivariate case, we specify

[4]The simple SIR, especially that with constant rate of recovery as in (2), is a tractable special case considered in Kermack and McKendrick (1927). In the general setup of their model, constant rate of recovered is replaced with a specification similar as in (10).

$\boldsymbol{\mu}_t$ as

$$\boldsymbol{\mu}_t = \boldsymbol{A}\boldsymbol{x}_{t-1} + \boldsymbol{\mu}_t^{EN} = \underbrace{\boldsymbol{\mu}_t^{AR}}_{\text{diag}(\boldsymbol{A})\boldsymbol{x}_{t-1}} + \underbrace{\boldsymbol{\mu}_t^{NE}}_{(\boldsymbol{A}-\text{diag}(\boldsymbol{A}))\boldsymbol{x}_{t-1}} + \boldsymbol{\mu}_t^{EN} \tag{12}$$

where $\boldsymbol{x}_t = [x_{1t}, \dots x_{nt}]^\top = \sum_{\ell=0}^{L-1} \nu(\ell)\boldsymbol{y}_{t-\ell}$ is the vector of infectious individuals in the London boroughs and $\boldsymbol{A}$ is an $n \times n$ matrix of coefficients. The first term in (12), $\boldsymbol{A}\boldsymbol{x}_{t-1}$, and the definition of $\boldsymbol{x}$, imply *vector* autoregressive dynamics in the disease propagation. The additional term $\boldsymbol{\mu}_t^{EN}$, which we call the *endemic* term, aims to capture variations in disease dynamics that are not explained by the *epidemic* component $\boldsymbol{A}\boldsymbol{x}_{t-1}$. Such endemic forces aim to accommodate seasonality, behavioural responses, or transmission dynamics induced by external forces.[5] This endemic/epidemic decomposition is commonly adopted in the empirical analysis of epidemiological surveillance data (see, e.g. Finkenstädt and Grenfell (2000), Held, Höhle, and Hofmann (2005), Lawson (2013)).

By further separating the endemic effect in equation (12) into its intra- and inter-borough elements, we have a natural decomposition of the conditional expectation of new infections into three components: the local (intra-borough) autoregressive effects $\boldsymbol{\mu}_t^{AR}$, the network (inter-boroughs) effects $\boldsymbol{\mu}_t^{NE}$, and the endemic effects $\boldsymbol{\mu}_t^{EN}$.

Autoregressive effects The local autoregressive effects capture disease dynamics as if each subpopulation were in isolation – i.e., as if new cases were driven only by infectious residents within the same borough. These are equivalent to assigning the single-population dynamic in equation (7) independently for each subpopulation. We adopt the simplest specification for the autoregressive effects by treating them as homogeneous among subpopulations. This implies that $\text{diag}(\boldsymbol{A}) = \gamma \boldsymbol{I}$, where the coefficient γ can be understood as the odds ratio parameter a in equation (7).[6]

[5] This is to be partially distinguished from the concept of "endemic diseases", which command stationary transmission dynamics due to the combination of slow propagation rate and temporary immunity.

[6] The specification is easily generalisable to location specific local dynamics by setting $\text{diag}(\boldsymbol{A}) =$

Network effects Moving from one homogeneous population to multiple subpopulations, a key new component is the network effects μ_t^{NE}, which is the main focus of our analysis. These effects are driven by a network (or multiple networks) of connections between subpopulations. Individuals contact and transmit the disease along predefined network links; the intensity of transmission is captured by the "strength" of links (i.e., varying edge values of the network). These network effects of disease transmission have long been acknowledged and were brought to light by Newman (2002). The network effects can also be interpreted as matching intensities between individuals from different subpopulations as in Acemoglu, Chernozhukov, Werning, and Whinston (2020). In particular, given an observable adjacency matrix $W = \{w_{ij}\}_{1\leq i,j\leq n}$ capturing the connections among London boroughs, we model the network effects μ_t^{NE} in (12) as

$$\mu_{it}^{NE} = \phi \left(\sum_{j \neq i} w_{ij} x_{j,t-1} \right). \tag{13}$$

Under this formulation, for residents of borough i, the chance of being infected by residents from borough j is proportional to the strength of connections between the two boroughs, as describe by the network topology W. The constant ϕ determines the overall strength of these network effects. The specification in (13) can be extended to allow for multiple networks, with adjacency matrices given by $W^{(g)} = \left\{ w_{ij}^{(g)} \right\}_{1\leq i,j\leq n}$, $g = 1, \ldots, G$, as follows:

$$\mu_{i,t}^{NE} = \sum_{g=1}^{G} \phi^{(g)} \left(\sum_{j \neq i} w_{ij}^{(g)} x_{j,t-1} \right). \tag{14}$$

Here, the topology specific coefficients $\phi^{(g)}$ capture the strength of the disease propagation though the various sets of links considered.

The above assumptions about the autoregressive and network effects directly trans-

---

$\text{diag}\left(\{\gamma_j\}_{j=1}^n \right)$.

late into a parametric specification for matrix $\boldsymbol{A}$ in (12):

$$\boldsymbol{A} = \gamma \boldsymbol{I} + \sum_{g=1}^{G} \phi^{(g)} \boldsymbol{W}^{(g)}. \tag{15}$$

Hence, our formulation implies that the conditional expectation of new infections is driven by both local autoregressive, as well as spatial, dynamics.

The set of parameters $\left\{\gamma, \phi^{(1)}, \ldots, \phi^{(G)}\right\}$, jointly with the adjacency matrices $\boldsymbol{W}^{(g)}$, fully determine matrix $\boldsymbol{A}$, which is essential for understanding the disease transmission dynamics. We expect these parameters to be non-negative for the data-generating process to be well defined, and impose this restriction by working with exponential forms. Furthermore, as the lockdown policy may affect these fundamental parameters, we model them as time varying. That is:

$$\gamma \equiv \gamma_t = \exp\left(\gamma_0 + D_t \gamma\right), \tag{16}$$

$$\phi^{(g)} \equiv \phi_t^{(g)} = \exp\left(\phi_0^{(g)} + D_t \phi^{(g)}\right), \quad g = 1, \ldots, G, \tag{17}$$

where $D_t \in \{0, 1\}$ is a dummy variable equal to zero before the lockdown and one afterwards. The impact of the policy change is therefore quantified as $\exp(\gamma)$ for autoregressive effects and $\exp\left(\phi^{(g)}\right)$ for network effects, multiplicatively. For example, if $\exp(\gamma) = 0.70$, then the policy reduces the local autoregressive effects by $100 \times (1 - 0.7)\% = 30\%$.

Endemic effects The endemic effects in (12) are specified as

$$\mu_{i,t}^{EN} = \exp\left(z_t^\top \boldsymbol{\beta} + \eta_i\right) N_i, \qquad i = 1, \ldots, n \tag{18}$$

which is proportional to the size of subpopulations N_i. The vector z_t contains deterministic time trends such as polynomial and trigonometrical functions of time t.

Additional time-varying terms for control, such as time-varying testing intensity and positive-to-test ratio, are also included in this vector. Location-specific fixed effects η_i, are added in the endemic terms to account for static demographic and socioeconomic heterogeneities among subpopulations. In particular, control variables in the endemic term are specified as

$$z_t = \Big[t,\ t^2,\ t \times D_t,\ t^2 \times D_t,\ \text{positive-to-test ratio},\ \text{number of tests}\Big].$$

Linear and quadratic terms of time are motivated by the hump-shaped disease spread trend (see, e.g., Li and Linton (2020)). These deterministic trends are allowed to change after the lockdown policy, with two interaction terms, $t \times D_t$ and $t^2 \times D_t$, added in. The two testing-related variables are both smoothed. Numbers of tests are scaled to per one million population, and are seven-day moving averages. The positive-to-test ratio is, by definition, the number of positive cases divided by the number of tests, both are also rolling seven-day statistics. Since we do not have granular data regarding the number of tests at the level of local authorities, these two controls are all based on nationwide statistics. We add these two controls to account for common variations across time due to limited testing capacity, especially in the early stage of the disease outbreak in UK.

Estimation of the model parameters is performed via maximum likelihood. The tuning parameter ρ of the exponential decaying kernel $\nu(\ell) = \exp(-\rho\ell)$ is selected via (maximising) the profile likelihood of each model. Standard errors are all calculated based on the outer product of the score vectors à la Berndt, Hall, Hall, and Hausman (1974).

We focus on performing inferences about parameters that determine the epidemic (autoregressive and network) effects. For parameters governing the baseline (before-policy) disease spread dynamics, we test whether their exponentials are significantly

greater than zero, that is, $\exp(\gamma_0) > 0$ and $\exp\left(\phi_0^{(g)}\right) > 0$ for all $g \in \mathcal{G}$ are our null hypotheses. These tests evaluate if there are substantial autoregressive or network effects during the COVID-19 pandemic in London. For parameters quantifying the policy impacts, we ask if their exponentials are significantly smaller than one, that is, if our null hypotheses are $\exp(\gamma) < 1$ and $\exp\left(\phi^{(g)}\right) < 1$. This is because (the exponentials of) these parameters all measure (gross) percentage changes from baseline estimates after the lockdown policy is implemented.

2 Data and Estimation

This section surveys the multiple data sources, including networks of London boroughs we use in this study, and presents estimation results and their implications.

2.1 COVID-19 data and demographic information

Disease surveillance data are provided by the UK government.[7] This database reports the number of daily new cases found in each local authority of the UK. We focus on local authorities of London, consisting of 32 London boroughs.[8] Our sample period is from March 1 to June 4, 2020. We begin our analysis from March because the number of reported cases is small from January to February in London.[9]

We illustrate in Panel (A) of Figure 2 the total number of reported COVID-19 cases in each borough until June 4, 2020. In general, the number of cases of the pandemic in each borough correlates with subpopulation sizes, which is shown in Panel (B) of the same figure. The correlation is 0.78 throughout our sample period.

[7] `https://coronavirus.data.gov.uk/`

[8] The city of London, strictly speaking, is not an official London borough (but, still a local authority), and in our dataset it is merged with Westminster. It is noteworthy that the city of London has an extremely small number of residents (less than ten thousand) and only thirteen disease cases in total (till June 4, 2020).

[9] There are only fifteen cases in total, with Southwark borough having the largest number, which is only three.

This association corroborates our specification for the endemic terms in equation (18), which is proportional to the subpopulation sizes.

There are three major policy dates for COVID-19 in the UK. They are March 16 when citizens were recommended to avoid traveling and stay home, March 20 when schools and pubs were closed, and March 23 from which full nationwide lockdown measures became effective. We choose March 23 as our policy date and evaluate the impacts of nationwide lockdown measures on the autoregressive and network effects using the specifications in (16) and (17).

Additionally, for subpopulation (borough) sizes, we use records from a housing-led population projection conducted by the Greater London Authority (GLA).[10] The GLA demography team constructs these projections based on trends in fertility, mortality and migration, and housing development on an annual basis. The starting point of these projections is dwelling records from the 2011 UK census.[11]

We use the UK nationwide testing data from the GitHub repository of Our World in Data (OWID).[12] Since May 23, 2020, the UK government has stopped publishing data regarding the number of "people tested". OWID uses instead the official data for the number of "tests performed" since then. For these numbers, only official swab tests count and all serology tests are disregarded.[13] Available data for both series, the numbers of people tested and tests performed, are plotted in Figure 1. Within their overlapping time window (April 26 to May 22, 2020), the two series agree well.

[10] Link: `https://data.london.gov.uk/dataset/housing-led-population-projections`

[11] Population census across all four UK countries is taken every ten years. The 2011 census is the most recent. We also use the census data for constructing our networks, as we discuss below.

[12] Link: `https://github.com/owid/covid-19-data/tree/master/public`

[13] *Official* swab tests are those conducted in Public Health England (PHE) labs and National Health Service (NHS) hospitals, as well as those processed *in-person* under government guidance for a wider population. Swab tests for surveillance purpose undertaken by PHE, Office for National Statistics (ONS), Biobank, universities, and other partners do not count.

2.2 Network construction

We construct networks connecting residents of different boroughs through commuting links. We use data from the 2011 UK Census to create a *directed* and *weighted* graph, of which 32 nodes denote the London boroughs[14]. Edges of this graph quantify employment links between boroughs. For example, an edge of value $10,000$ from Camden to Southwark means that there are $10,000$ individuals living in Camden who go to work in Southwark. We ignore all self-pointing edges, meaning that we drop the numbers of people working in the same boroughs where they live, since the local effects are already meant to be captured by the autoregressive component μ_t^{AR}. In other words, the adjacency matrix of this graph has a diagonal of constant zeros.

This graph is visualised in Figure 3. The size of a node in this figure is proportional to its in-degree (total number of people coming to work in this borough). For clarity, an edge is drawn only if its value exceeds the 80% percentile of all edge values. Widths of these plotted edges are also proportional to their values. A clear pattern of Figure 3 is that Westminster/City of London attracts a disproportional amount of London workforce with a total in-degree that is significantly larger than any other London borough. In addition, Camden and Tower Hamlets also attract a relatively large amount of workforce.

We denote by $K = \{k_{ij}\}_{1\leq i,j\leq n}$ the adjacency matrix of this graph, where i indexes home and j indexes work. We have constructed other networks based on this K matrix. Specifically, we consider three networks, $\boldsymbol{W}^{(1)} = K$, $\boldsymbol{W}^{(2)} = K^\top$, and $\boldsymbol{W}^{(3)} = KK^\top$, defined respectively through adjacency matrices.[15] These three adjacency matrices

[14]The 2011 census covers detailed employment information which includes office locations. Thus, for any local authority, the number of its people working in other local authorities can be calculated. Aggregate data regarding these statistics are available from `https://data.london.gov.uk/dataset/place-residence-place-work-local-authority`.

[15]Notations here are to provide basic ideas. In detail, when estimating the network SIR models, we always divide K and $K^\top$ by the largest singular value of K to rescale their spectral norms to one. This operation improves numerical stability (of nonlinear maximum likelihood estimation) and ensures

capture transmission from different contact networks: 1) “work-to-home" transmission, which measures the spreading of disease from residents in commuter’s work borough to residents in commuter’s home borough; 2) “home-to-work" transmission, which is from the opposite direction; and 3) “home-to-home" transmission, which is between different places of residence via a common workplace.

Generally speaking, left-multiplying the vector of active cases x_t by the matrix $W^{(1)}$ as in equation (14) features the transmission of COVID-19 from workplaces to residential areas. To be clear, for borough i, the vector $\left[w_{ij}^{(1)}\right]_{j\neq i}$ will *overweight* boroughs where more of borough i’s residents go to work. For example, for many London boroughs, their residents are more likely to go to work at Westminster/City of London. Thus, the propensity for contact and infection *from* Westminster/City of London is set to be higher by this network. Intuitively, larger nodes of the graph shown in Figure 3 are more emphasised in the transmission dynamics.

Similarly, left multiplying x_t by $W^{(2)}$ characterises the disease spread from residential areas to workplaces. Panel (A) of Figure 4 shows the graph defined by $W^{(2)}$. Edges are again thresholded at the 80% for presentation. Node sizes are also proportional to their total in-degrees. As we can see from Panel (A) of Figure 4, network effects commanded by this graph highlight the role of boroughs such as Wandsworth, Lambeth, Lewisham, and Southwark. These boroughs are those with large numbers of residents going to work in other boroughs.

Interpreting the impact of our third network on the disease spread dynamics – the home-to-home network – is also straightforward. The adjacency matrix of this network is symmetric with element (i,j) and (j,i) defined by $\sum_{\ell=1}^{n} k_{i\ell}k_{\ell j}$. It quantifies

comparability across network parameters. For computing the adjacency matrix of the third network, namely $W^{(3)}$, we first normalise each row of K to sum one, then calculate the product of it with its transpose. Diagonal elements of this product are also replaced with zero. The resulting matrix is then divided by its largest singular value (which is also its largest eigenvalue because this matrix is symmetric).

the propensity of residents of borough i and j to contact and infect each other with the disease via common workplaces. Specifically, if work and home locations were independent, the entries of this particular adjacency matrix would be proportional to the probability of individuals from the two locations meeting each other. Panel (B) of Figure (4) visualises the graph of this network. Nodes with large in-degrees such as Westminster/City of London, Camden, and Tower Hamlets should be influential in driving the disease spread dynamics mediated through this network. These boroughs are also featured in the work-to-home network graphed in Figure 3. In addition, Wandsworth, Lambeth, and Southwark, which emerge as pivotal nodes in the home-to-work graph, are also highlighted here, meaning that boroughs with large numbers of residents working outside those boroughs also have strong network effects. An exception to this appears to be Lewisham, which does not appear to be influential in the graph.

3 Empirical Findings

The main estimation results are presented in Table 1. Five specifications are included: three specifications with one of the networks ($\mathcal{G} = (1), (2), (3)$) each; one specification containing all three networks ($\mathcal{G} = (1, 2, 3)$); and one specification incorporating the first two networks ($\mathcal{G} = (1, 2)$). The first three specifications serve the purpose of univariate analysis. Parameter estimates and standard errors are reported in the table, as well as the pseudo-R^2 defined as

$$\text{pseudo-}R^2 = 1 - \frac{\sum_{i=1}^{n}\sum_{t=1}^{T}(y_{it} - \widehat{\mu}_{it})^2}{\sum_{i=1}^{n}\sum_{t=1}^{T}(y_{it} - \overline{y})^2},$$

where $\widehat{\mu}_{it}$ is the conditional expectation of daily new COVID-19 cases calculated from the model using estimated parameters; $\overline{y}$ stands for simple average across time and boroughs.

Before-policy estimations for the autoregressive effects, that is, estimates of $\exp(\gamma_0)$, are always significantly greater than zero. Across the five specifications, the smallest estimate is 0.252 with a standard error of 0.041. As this parameter quantifies within-borough disease transmission dynamics, this finding confirms that serious community-acquired infection has occurred during the initial outbreak of COVID-19 in London in early 2020. The nationwide lockdown policy on March 23 has reduced the infection significantly: except for the second specification (which is 5%) the magnitude ranges from 24% to 38%. The reduction is economically and statistically significant as $\exp(\gamma)$ is significantly smaller than one (again with model (2) as an exception). The analysis shows that the lockdown policy is effective in cutting down within-community transmission rates.

Estimation results from the first specification (column 1 of Table 1) show that the impact of the work-to-home network is large in magnitude: the before-policy coefficient $\exp\left(\phi_0^{(1)}\right)$ is around 0.5, which implies strong directional spillover effects from people living in major workplaces such as Westminster/City of London to residents in other boroughs via the work-to-home contact network. We also find that the work-to-home disease spread is partially reduced after the lockdown. The magnitude of reduction is around 12% for the first specification and around 18% after controlling for other networks, while being not statistically different from one. The effectiveness of the lockdown policy on blocking cross-borough transmission is questionable at this level. It seems that transmitting from pivotal places such as Westminster/City of London, Camden, and Tower Hamlets to other boroughs still remains active even after the nationwide lockdown.

We find that the impact of the home-to-work network on the epidemics of COVID-19 in London is relatively smaller (column 2 of Table 1). The baseline estimation $\exp\left(\phi_0^{(2)}\right)$ is less than 0.12.[16] This parameter is around 0.06 after controlling for other

[16]Our parameter estimates for different networks are comparable because all graph adjacency ma-

contact networks. Though smaller in magnitude, it is significantly greater than zero except when including the home-to-home network. We find that the lockdown policy has a stronger impact on this home-to-work spreading. After the lockdown policy is enforced, the spreading through home-to-work network is reduced by over 80% after controlling for other network effects.

We find a strong home-to-home network effect when estimating it alone (column 3 of Table 1), but it is insignificant after controlling for the work-to-home and home-to-work networks (column 4 of Table 1). Thus, we choose to use the last specification including both work-to-home and home-to-work networks as the benchmark specification (column 5 of Table 1). The maximised log-likelihood function for the fourth and the fifth specifications are -7335.229 and -7335.153, respectively. A likelihood ratio test simply fails to offer any definitive evidence on including the home-to-home network. Therefore, we report further analyses of the network effects based on this benchmark specification for expositional clarity in the rest of the paper.

Table 1 also reports results for testing-related variables in vector z_t that drives the endemic effects. Regression coefficients for the number of tests (i.e. lag test) are consistently positive across specifications, being around 1.8 with standard errors around 0.3. That is, we find that greater testing capacity leads to larger endemic effects and predicts a larger number of new cases. There are two potential explanations for this result. Firstly, testing efforts are motivated by the past disease propagation trajectory, which also determines the number of new cases. The association between testing numbers and new case counts is thus due to the fact that they both relate to the current epidemic severity. The second explanation points toward under testing. If there is a serious under testing problem, increased testing capacity will artificially pick up cases that have been omitted early on, thus increasing the expected number of cases. This explanation is partly supported by the significantly negative coeffi-

trices have been rescaled by their largest singular value before feeding into our models.

cients for positive-to-test ratios in Table 1. Across models, this regression coefficient is at least -1.09 with a standard error being at most 0.49. Higher positive rates predict weaker future endemic effects. Intuitively, high positive rates are more likely to be a result of serious under testing (assuming testing kit technologies remain unchanged), which artificially discounts the expected number of cases. This view has been discussed in the media, for example,

> *"The thinking is that higher positive test rates equate to more missed cases"*
>
> (Bloomberg Opinion, July 1, 2020)

Our findings that *both* larger testing capacity and smaller positive rates are associated with more future cases are suggestive of under testing, which may cause reported cases deviating from the actual number. Adding these two test-related variables into endemic effects serves as an adjustment to this issue.

3.1 Decomposition of expected daily new cases

We now plug the estimated parameters into expected new case numbers μ_t as defined in (12) to evaluate and compare the contribution of autoregressive, network, and endemic components in explaining the observed data. We rely on parameter estimates from the benchmark specification incorporating the work-to-home and home-to-work networks (column 5 in Table 1) to perform this decomposition in this section.

Panel (A) of Figure 5 graphs the decomposition of the contribution of the autoregressive, the network, and the endemic components to the total cases aggregating across the 32 London boroughs in the sample period. That is, we plot the time-series of the three ratios $\left(\mathbf{1}^\top \widehat{\boldsymbol{\mu}}_t^{AR}\right) / \left(\mathbf{1}^\top \widehat{\boldsymbol{\mu}}_t\right)$, $\left(\mathbf{1}^\top \widehat{\boldsymbol{\mu}}_t^{NE}\right) / \left(\mathbf{1}^\top \widehat{\boldsymbol{\mu}}_t\right)$, and $\left(\mathbf{1}^\top \widehat{\boldsymbol{\mu}}_t^{EN}\right) / \left(\mathbf{1}^\top \widehat{\boldsymbol{\mu}}_t\right)$ as defined in Section 1.3. The graph demonstrates a substantial and persistent network effect, accounting for 42% of the expected total daily new cases in London. The au-

toregressive and endemic effects contribute to 35% and 23%, respectively.

In Panels (B) and (C) of the same figure, we present the decomposition result for each London borough. We plot the time series of the fractions of conditional means explained by each of the three components, $\widehat{\mu}_{it}^{AR}/\widehat{\mu}_{it}$, $\widehat{\mu}_{it}^{NE}/\widehat{\mu}_{it}$, $\widehat{\mu}_{it}^{EN}/\widehat{\mu}_{it}$, for each borough. Light-coloured bands in these plots show 10% to 90% percentiles across all boroughs, delineating cross-sectional heterogeneity. Solid lines represent the median levels.

The left graph in Panel (B) shows community transmission (autoregressive effects) rises quickly rand remains a strong driver of the disease spread since March. The middle graph in Panel (B) shows a similar pattern for network effects but with an additional peak in mid-March. The right graph in Panel (B) shows that at the very early stage, endemic effect is a dominating component, which we interpret as infections being brought in by people coming from outside of London. Endemic effects keep dropping in early March and remain stable for a while. They are diminished to a low level in late May. Comparing all three graphs in Panel (b), we find that the fading endemic impact is largely subsumed by the network effects.

In Panel (C) of Figure 5 we further decompose the network effects by examining the contributions of the work-to-home and the home-to-work network to the total network effect across time in our sample period, respectively. Based on plots in this panel, we observe that almost all network effects in our sample period can be attributed to the work-to-home network, highlighting the importance of boroughs where many people go to work in transmitting the disease. This finding results from parameter estimates in Table 1, where the network coefficient for the second specification $\exp\left(\phi_0^{(2)}\right)$ is almost ten times smaller than its counterpart for the first specification $\exp\left(\phi_0^{(1)}\right)$. Although weaker in terms of magnitude, we also observe that the home-to-work network effects have undergone much stronger reduction after the lockdown policy (i.e. $\exp\left(\phi^{(2)}\right) < \exp\left(\phi^{(1)}\right)$).

In the graphs in all three panels, we highlight March 23, the lockdown date. Panels (B) and (C) show that the reductions in autoregressive and (work-to-home) network effect have occurred at least one week before March 23, indicating that a significant number of London residents have started observing social distancing within their communities and working from home before the official policy announcement date.

3.2 Disease R_0

Our model offers guidance on the basic reproduction number R_0, which quantifies the expected number of new cases *directly* generated by one existing case. To be more specific, since our model features borough-level heterogeneity, the expected number of new cases varies with regard to the residence of the existing infected cases. Thus, if the "one" existing case comes from borough i, we have a basic reproduction number $R_0^{(i)}$. What we aim to compute here is an estimated upper bound on the maximum $\max_{i\in\{1,\ldots,n\}} R_0^{(i)}$.

To proceed, we first calculate a simple plugging-in estimator of matrix A in equation (12), denoted by $\widehat{A}$. For one additional case in borough i, denoted by vector e_i in which only the i-th element equals one and all other elements equal zero. This case remains contagious for L periods by assumption. The total (expected) number of new cases created directly by e_i can be estimated as $\sum_{\ell=0}^{L-1} \exp\left(-\widehat{\rho}\ell\right) \widehat{A}e_i$. This quantity above is uniformly bounded by $\sum_{\ell=0}^{L} \exp\left(-\widehat{\rho}\ell\right) \sigma_{\max}\left(\widehat{A}\right)$ because $\|e_i\| = 1$, in which $\sigma_{\max}(\cdot)$ represents the function that computes the largest singular value of a matrix. Thus, we have an upper bound for the estimates of R_0 as follows:

$$\widehat{R}_0 \leq \max_{i\in\{1,\ldots,n\}} \left\{ \sum_{\ell=0}^{L-1} \exp\left(-\widehat{\rho}\ell\right) \widehat{A}e_i \right\} \leq \sum_{\ell=0}^{L} \exp\left(-\widehat{\rho}\ell\right) \sigma_{\max}\left(\widehat{A}\right).$$

Estimates of (upper bounds of) R_0 are presented in Table 2. According to our specification in Section 1.3, we have separate estimates of $\widehat{A}$ before and after the lockdown policy. Thus, the table shows the largest singular values and R_0s both before and after the policy date. The disease R_0 is around 1.4 before the lockdown policy and is around 0.8 afterwards based on our benchmark specification of incorporating work-to-home and home-to-work networks (last two columns of Table 2). The magnitude of reduction is large though not statistically significant. This reduction is due to the impacts of lockdown policies on both the autoregressive and network effects.

3.3 The network impulse response functions

To further understand how innovations in daily new COVID-19 cases propagate through networks, we define and calculate the network impulse response function (NIRF) of our model motivated by the analysis in Denbee, Julliard, Li, and Yuan (2020). For a unitary shock (or change in levels) of disease incidents in borough i, its impact on the expected total number of cases across *all* locations τ-period ahead is measured by

$$\text{NIRF}_i(\tau) = \sum_{j=1}^{n} \frac{\partial \mathbb{E}\left[y_{j,t+\tau} \mid \mathcal{F}_t\right]}{\partial y_{it}}. \tag{19}$$

The empirical model we work with allows an analytical formula for the NIRF, as detailed in the Appendix.

Plots in Figure 6 illustrate NIRFs across each borough for the time horizon of one week, that is $\tau = 7$. Panel (A) shows the impulse responses before the lockdown. The Westminster/City of London subpopulation strongly dominates all other London boroughs. For one additional case that emerges in this area, three more cases are expected to occur in the whole Greater London area, even after one week. This identifies the Westminster/City of London area as a "key player" for shock propagations

through the network in the language of Denbee, Julliard, Li, and Yuan (2020). Camden, Tower Hamlets, Southwark, as well as Lambeth are among the other key players that appear to show strong network impulse responses, but the magnitude is much weaker than for Westminster/City of London. Panel (B) presents results after the lockdown. It offers another angle for us to understand the effectiveness of the lockdown policy, as there is a distinct reduction in the NIRF measure for the key areas such as Westminster/City of London and other boroughs.

Analysing network impulse responses is valuable for designing "smart" partial lockdown policies that selectively lockdown a few regions instead of deploying a full-scale lockdown. When prescribing a partial lockdown plan, the conventional wisdom is to shutdown areas that have witnessed the largest number of existing cases and are undergoing rapid growth in new cases.[17] Our network impulse response analysis offers another perspective. In addition to focusing on regions that have reported severe outbreaks, lockdown policies should also target areas that are key to the disease transmission. Isolating subpopulations that are key players in the network can forestall rapid spread among the whole population, even if few cases have occurred in these areas. Optimal (partial) lockdown policies should combine both perspectives, as we demonstrate in our counterfactual analysis in the next session.

4 Counterfactual Simulations

This section presents simulations to evaluate counterfactual outcomes from alternative policy interventions. We start by investigating the impacts of earlier or later nation-wide lockdown measures, and compare them with actual numbers. We then compute optimal partial lockdown arrangements and their potential outcomes. We emphasise on comparing optimal policies with "naive" policies that only target areas

[17] See, for example, the lockdown of Hubei province in China and Lombardy region in Italy, both of which were the epicentres of COVID-19 outbreaks when the policies came out.

with the largest number of existing cases. Our goal when making these comparisons is to illustrate the importance of shutting down pivotal nodes of the network as a preventive measure against disease transmission.

4.1 Alternative dates for nationwide lockdown

The UK government implemented a nationwide lockdown on March 23. The timing of this policy has been under intense public scrutiny. Dr. Neil Ferguson who, with his research group's "Imperial College" model, has facilitated the lockdown decision of the UK government said:

> *"Had we introduced lockdown measures a week earlier, we would have reduced the final death toll by at least a half."* (BBC News, June 10, 2020)

Similar arguments have been made by scientists such as Dr. James Annan and Dr. Kit Yates.[18] In the meantime, Dr. Yates has acknowledged that

> *"There had been an 'overreliance' on certain models when determining how fast the epidemic had been doubling; ...that some of the modelling groups had more influence over the consensus decisions than others."*
>
> (BBC News, June 10, 2020)

We attempt the same inspection using our model, although within a limited scope, by only focusing on the case of London. With our estimates, we change the policy indicator D_t by allowing for different policy intervention dates and simulate the model

[18]These proponents have been broadly covered by the media. James Annan's conclusion was drawn upon his calculation made public on May 12, 2020 through a blog post, which is available at https://bskiesresearch.wordpress.com/2020/05/12/the-human-cost-of-delaying-lockdown/. Citing James Annan's calculation, Kit Yates wrote "locking-down a week earlier translates to beginning lockdown with roughly a quarter of the total cases..." in an essay to HuffPost on May 22, 2020. Details can be found at https://www.huffingtonpost.co.uk/entry/lockdown-uk-deaths_uk_5ec6efd8c5b68038a74a50ad?utm_hp_ref=uk-opinion.

outcomes. Specifically, we base simulation exercises on parameter estimates from the benchmark specification (work-to-home and home-to-work networks) reported in Table 1. We consider policy dates two weeks or one week, both before and after March 23. We simulate each of the resulting models $10,000$ times and then average the $10,000$ paths of daily case counts across all London boroughs as our counterfactual outcomes. Throughout our simulations, we *fix* the endemic terms, although this element of the model, by definition, is also affected by the specific date of the policy due to variables such as $t \times D_t$ and $t^2 \times D_t$. Under this simulation design, we are indeed treating the endemic term as a deterministic force in the model. The endemic effects are introduced only to isolate the autoregressive and network effects through controlling for variations that are not epidemic. By doing so, we are also ignoring the impacts of alternative policies on the nationwide positive-to-test ratio.[19]

Simulation results are presented in Figure 7. Plots in this figure show the cumulative number of all London cases from March 1 to June 4. The plot on the left panel compares the simulated outcomes from locking-down one or two weeks earlier with the actual outcome. Based on our results, locking down two weeks earlier translates into a reduction of 12% total cases ($3,283$ cases on the absolute level) in London during the period under study. The number is 9% ($2,553$ on the absolute level) if lockdown would have been one week earlier. These numbers indicate that, at least for London, positive action earlier than the March 23 lockdown dates would have yielded some reduction in total cases but the magnitude is not large.

The right panel plots alternative situations in which lockdown takes place one or two weeks later. Delaying the March 23 lockdown would cause serious consequences in terms of a large increase in the number of people infected. If the lockdown would have happened one week later, the total number of infected cases up till June 4 would

[19] In practice, each simulation creates new time-series of London cases, which should count as part of the total positive cases in Britain, thus changing the numerator of this ratio.

have increased by 17% (4,514 on the absolute level). Postponing the lockdown for two weeks would have led to an even larger increase of 56%, that is, over fifteen thousand more people in London would have been infected.

4.2 Optimal partial lockdowns

This section discusses optimal lockdown policies. With the parameter estimates for our network SIR model, we are able to map out the dynamics of COVID-19 spread across London boroughs. This knowledge makes it possible for us to answer the following question: at any given time, what is the optimal borough-specific lockdown scheme that can minimise the number of infected cases in the future? For clarification, our definition of optimality is a constraint on the number of boroughs being locked down. That is, we are searching for optimal solutions to problems such as: "If only three London boroughs are to be shut down for controlling the spread, which ones are they?"

To address these types of questions, we simulate the dynamics of COVID-19 transmission in London based on parameter estimates of the benchmark in Table 1. We then consider policies that shut down different combinations of London boroughs on the exact same date of the actual nationwide lockdown, March 23. A total of 10,000 paths of disease case counts is generated from simulations. For each path, we evaluate potential outcomes of all possible partial lockdown policies. For example, if two out of the 32 London boroughs are allowed to be locked down, we then have $\binom{32}{2} = 496$ different lockdown designs. The outcomes we focus on are the total number of COVID-19 infections in London, averaged across all simulations.

The effects of locking down a particular borough are quantified in two ways. The first is more "optimistic" in the sense that boroughs being locked down will no longer

interact with other boroughs through the networks at all. That is, corresponding columns and rows of matrix $W^{(1)}$ and $W^{(2)}$ are shifted to zero. The second borrows information from our estimates for the impact of nationwide lockdown policies. Under this setting, for boroughs that are locked down, corresponding columns and rows of matrix $W^{(1)}$ and $W^{(2)}$ are downscaled to 82.7% and 16.6% of their original values, respectively.[20] Both ways of evaluating lockdown plans ignore the impact on the autoregressive and endemic effects. Hence, the resulting outcomes are conservative when compared with the actual numbers (which can be regarded as direct outcomes due to the March 23 nationwide lockdown that have changed the dynamics of all three effects).

We start with the simplest case of lockdown policy: only one borough is placed on lockdown. Simulation results are presented in Figure 8. The left panel corresponds to the case where, when shutting down a borough, all in-flow and out-flow through the network via this borough are blocked. The right panel uses parameters from our estimation results as described above. In either one of the two panels, the red line represents the cumulative case numbers across time if there are no lockdown policies implemented at all. The blue line shows the real number. Outcomes from different lockdown plans are marked in grey. We highlight the policy of shutting down only one borough with the highest cumulative number of infected cases (the borough of Lambeth) on the policy date (March 23) in orange. This type of case-targeting policy reflects the conventional wisdom. For comparison, we characterise the optimal policy in pink. This policy is optimal in the sense that it minimises the total (expected) number of infected cases. In our analysis, the optimal single-borough lockdown plan is to lockdown Westminster/City of London, which features the largest network impulse responses as shown in Figure 6. The results confirm that lockdown

[20]The relevant quantities here are $\exp\left(\phi^{(1)}\right) = 0.827$ and $\exp\left(\phi^{(2)}\right) = 0.166$ for the benchmark specification in Table 1 (column 5).

plans which only pay attention to the current number of infected cases are far from optimal (by comparing the orange and pink lines). Shutting down the region that is a pivotal node for the network propagation of shocks brings major improvements: a reduction of around $15,000$ total cases in London.

We proceed to consider the case of locking down two or three boroughs. Results are shown in Figures 9 and 10, respectively. The optimal arrangement of locking down two boroughs is to pick Westminster/City of London and Southwark, outcomes of which are shown in pink lines in Figure 9. Up till the date of March 23, the top two boroughs that have witnessed the most severe outbreak are Lambeth and Southwark. Locking down these two boroughs gives rise to outcomes shown in orange in the same figure. We can observe again that the optimal lockdown policy of targeting Westminster/City of London and Southwark leads to major improvements over the policy that only targets regions with the largest existing number of cases. It is noteworthy mentioning that the two boroughs with the highest NIRFs are Westminster/City of London and Camden. Therefore, we find that the optimal lockdown policy is a combination of two targets: highest number of COVID-19 cases (the inclusion of Southwark) and largest NIRF (the inclusion of Westminster/City of London). Both targets are essential for the design of optimal lockdown policies.

Similar findings appear in Figure 10 when investigating polices that lockdown three boroughs down. The optimal policy dictates shutting down Westminster/City of London, Southwark, and Lambeth. By comparison, the top three boroughs that have been most harshly hit by the pandemic up till March 23 are Lambeth, Southwark, and Brent; while boroughs with the largest NIRFs are Westminster/City of London, Camden, and Tower Hamlets. Two boroughs out of the highest number of COVID-19 cases category are included to the optimal policy and the most pivotal node with the largest NIRF, Westminster/City of London, is also incorporated.

Our simulations show that the partial lockdown plans are powerful substitutes of the full-scale lockdown. Suppose that locked-down regions are completely removed from the network of London boroughs, an optimistic assumption, the optimal two-borough lockdown plan (locking down Westminster/City of London and Southwark) achieves almost the same outcome as the actual nationwide plan (see the left panel of Figure 9). The optimal three-borough lockdown plan (locking down Westminster/City of London, Southwark, and Lambeth) surpasses the nationwide plan and lowers the number of total cases in London by around $2,000$ (see the left panel of Figure 10). This optimistic assumption is not an exaggeration, in that locking down on a smaller scale in targeted areas can be much better administered. If we are more cautious in regarding the implementation of partial lockdown and quantifying its impact using the numbers we estimated for the nationwide lockdown, the optimal three-borough lockdown plan still leads to similar outcomes compared with the actual nationwide lockdown, increasing the total number of cases in London by less than $3,000$. We would like to point out that all comparisons above are highly conservative in terms of favouring partial lockdown plans, because we totally ignore the potential changes these plans can bring to the autoregressive and endemic components of epidermic transmission.

5 Conclusions

In this paper, we present and estimate a network-SIR model of the spreading of COVID-19 disease in London. Our estimates show that network play a major role in transmitting COVID-19 disease and they cannot be ignored. Based on the estimated epidermic dynamics, we investigate whether a certain target lockdown policy could contain the spread of COVID-19 disease as much as the full scale lockdown and, hence, have a lower economic cost. Our simulations show that an optimal lockdown should target areas that not only have the highest number of existing case, but also

those that play a key role in transmitting disease in the contact network among the population. In our case, the contact network corresponds to the commuting network in London. In designing a lockdown policy, our finding calls for special attention to be focused on the network role of the COVID-19 transmission. These network could be train or flight networks, or any other traffic networks at the national or international level including the migration network (from hard-hit COVID-19 hotspots) identified by Coven, Gupta, and Yao (2020). As networks potentially connect regions with different jurisdictional governments, our finding indicates that coordinated regional quarantine and lockdown policies are essential in containing the spread of the COVID-19 pandemic, a conclusion echoed by the theoretical work in Chandrasekhar, Goldsmith-Pinkham, Jackson, and Thau (2020).

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. Technical report, National Bureau of Economic Research.

Alvarez, F. E., D. Argente, and F. Lippi (2020). A simple planning problem for covid-19 lockdown. Technical report, National Bureau of Economic Research.

Avery, C., W. Bossert, A. Clark, G. Ellison, and S. F. Ellison (2020). Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. Technical report, National Bureau of Economic Research.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who's who in networks. wanted: The key player. *Econometrica 74*(5), 1403–1417.

Berndt, E. R., B. H. Hall, R. E. Hall, and J. A. Hausman (1974). Estimation and inference in nonlinear structural models. *Annals of Economic and Social Measurement 3*(4), 653–665.

Chandrasekhar, A. G., P. Goldsmith-Pinkham, M. O. Jackson, and S. Thau (2020). Interacting regional policies in containing a disease. Technical report, Stanford University.

Coven, J., A. Gupta, and I. Yao (2020). Urban flight seeded the covid-19 pandemic across the united states. Technical report, NYU Stern School of Business.

Denbee, E., C. Julliard, Y. Li, and K. Yuan (2020). Network risk and key players: A structural analysis of interbank liquidity. *Journal of Financial Economics forthcoming*.

Easley, D. and J. Kleinberg (2010). *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.

Farboodi, M., G. Jarosch, and R. Shimer (2020). Internal and external effects of social distancing in a pandemic. Technical report, National Bureau of Economic Research.

Fernández-Villaverde, J. and C. I. Jones (2020). Estimating and simulating a sird model of covid-19 for many countries, states, and cities. Technical report, National Bureau of Economic Research.

Finkenstädt, B. F. and B. T. Grenfell (2000). Time series modelling of childhood diseases: a dynamical systems approach. *Journal of the Royal Statistical Society: Series C 49*(2), 187–205.

Garriga, C., R. Manuelli, and S. Sanghi (2020). Optimal management of an epidemic: Lockdown, vaccine and value of life. Technical report.

Grimmett, G. R. and D. R. Stirzaker (2001). *Probability and Random Processes (Third Edition)*. Oxford University Press.

Held, L., M. Höhle, and M. Hofmann (2005). A statistical framework for the analysis of multivariate infectious disease surveillance counts. *Statistical Modelling 5*(3), 187–199.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.

Jones, C. J., T. Philippon, and V. Venkateswaran (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical report, National Bureau of Economic Research.

Kermack, W. O. and A. G. McKendrick (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A 115*(772), 700–721.

Lawson, A. B. (2013). *Bayesian Disease Mapping: Hierarchical Modeling in Spatial Epidemiology*. CRC press.

Li, S. and O. Linton (2020). When will the covid-19 pandemic peak? *Journal of Econometrics*.

Newman, M. E. (2002). Spread of epidemic disease on networks. *Physical review E 66*(1), 016128.

Rowthorn, B. R. and F. Toxvaerd (2012). The optimal control of infectious diseases via prevention and treatment.

Table 1: Estimation results of the Network-SIR model of Section 1.3. Results in columns (1)-(3) correspond, respectively, to specifications with only one of the following transmission networks: "work-to-home;" "home-to-work;" "home-to-home". Column (1,2,3) considers the three networks jointly while column (1,2) uses only the first two networks.

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Autoregressive effect										
$\exp(\gamma_0)$	0.311	0.035	0.286	0.040	0.286	0.037	0.252	0.041	0.253	0.041
$\exp(\gamma)$	0.617	0.077	0.958	0.136	0.750	0.102	0.760	0.129	0.754	0.127
Network effect										
$W^{(1)} = K$:										
$\exp\left(\phi_0^{(1)}\right)$	0.477	0.070					0.503	0.144	0.510	0.073
$\exp\left(\phi^{(1)}\right)$	0.885	0.136					0.821	0.169	0.827	0.127
$W^{(2)} = K^\top$:										
$\exp\left(\phi_0^{(2)}\right)$			0.115	0.022			0.062	0.036	0.064	0.023
$\exp\left(\phi^{(2)}\right)$			0.374	0.080			0.148	0.218	0.166	0.181
$W^{(3)} = KK^\top$:										
$\exp\left(\phi_0^{(3)}\right)$					0.261	0.035	0.000	0.085		
$\exp\left(\phi^{(3)}\right)$					0.570	0.082	0.284	0.000		
Testing-related endemic effect										
pos-to-test	-1.37	0.47	-1.09	0.28	-1.46	0.41	-1.46	0.48	-1.46	0.49
lag test	1.85	0.31	1.77	0.23	1.88	0.28	1.91	0.32	1.91	0.32
pseudo-R^2	82.09%		82.00%		82.02%		82.12%		82.12%	
# obs.	3008		3008		3008		3008		3008	

Table 2: Estimation-implied upper bound for the basic reproduction number R_0 of COVID-19 in London. $\sigma_{\max}$ denotes the largest singular value of the estimated matrix A in equation (15).

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Before Lockdown Policy: $D_t = 0$										
$\sigma_{\max}$	0.840	0.590	0.511	0.106	0.639	0.287	0.853	0.533	0.766	0.537
R_0	1.467	1.030	1.577	0.327	1.531	0.688	1.391	0.876	1.369	0.953
After Lockdown Policy: $D_t = 1$										
$\sigma_{\max}$	0.510	0.369	0.270	0.038	0.301	0.123	0.501	0.398	0.422	0.392
R_0	0.891	0.644	0.832	0.118	0.722	0.296	0.818	0.649	0.754	0.701

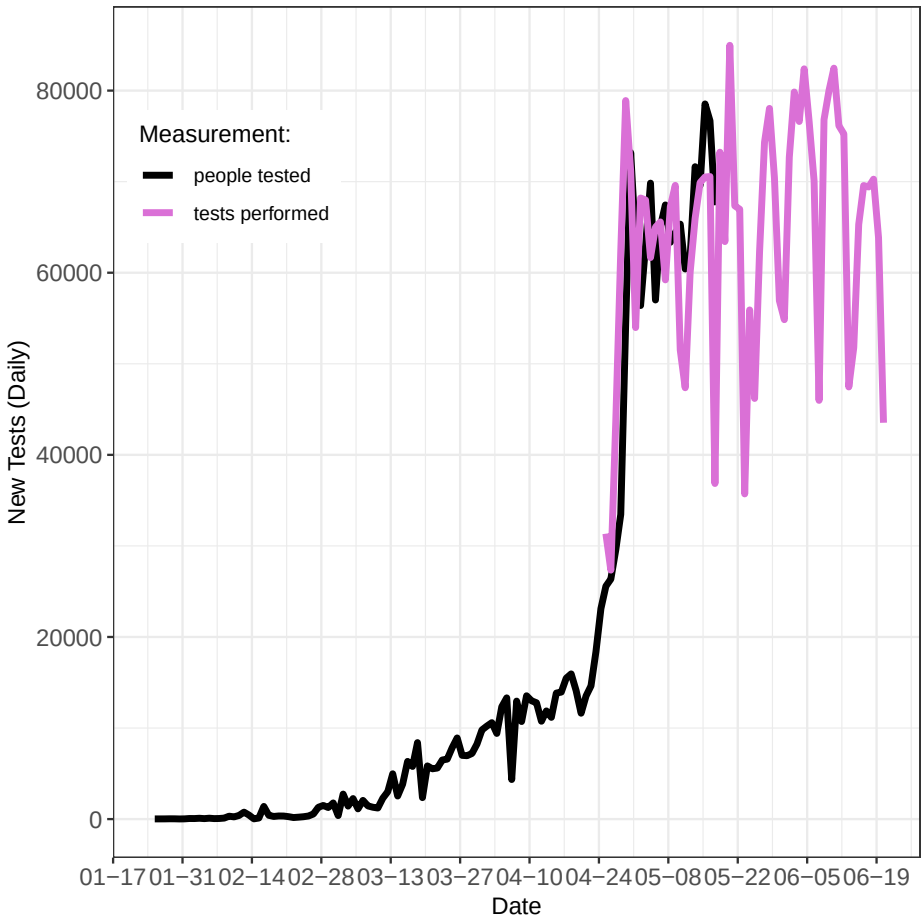

Figure 1: UK nationwide COVID-19 tests: people tested and tests performed

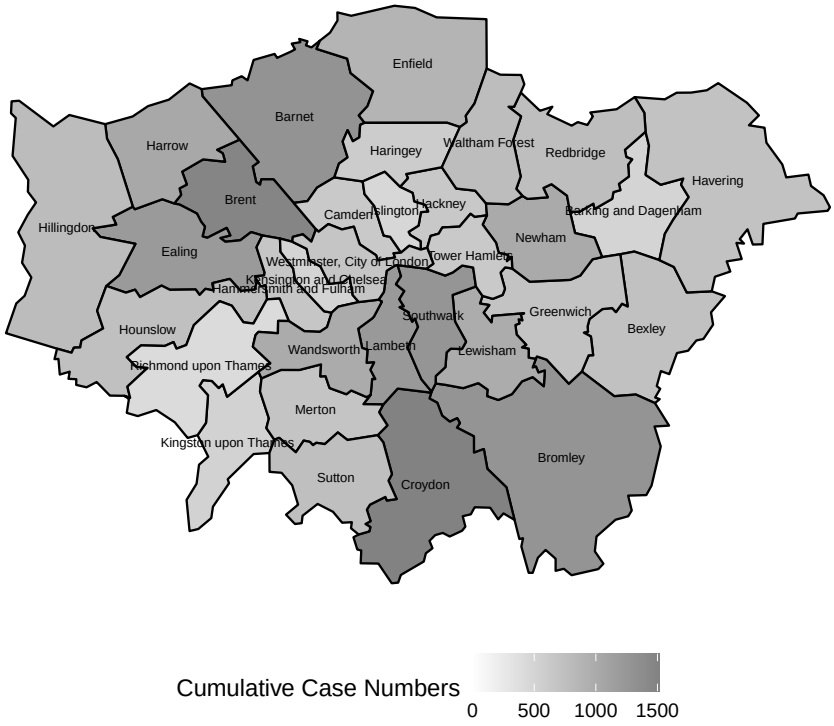

(A) COVID-19 cases for boroughs in London (up till June 4, 2020)

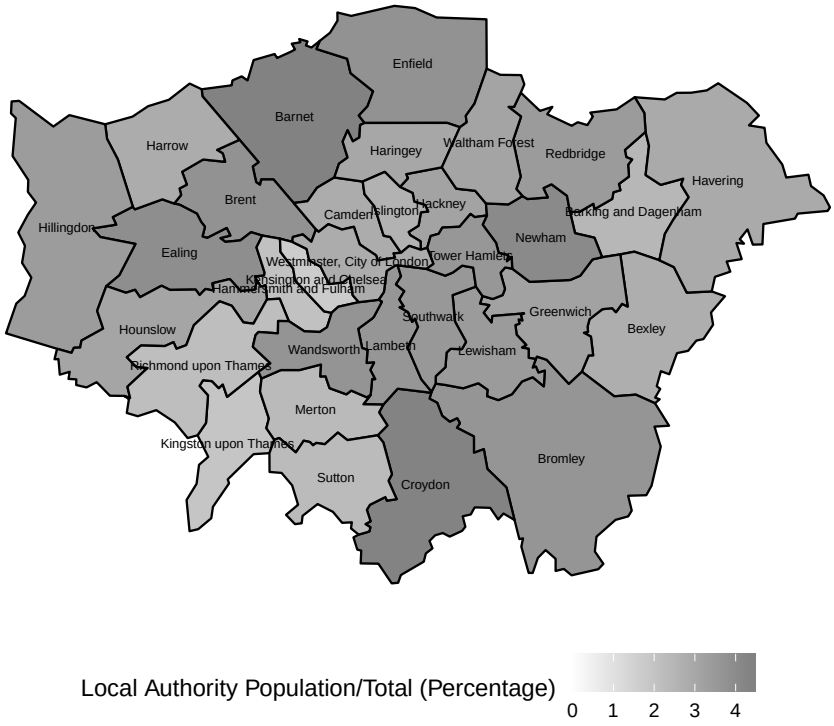

(B) Population sizes for boroughs in London (year 2020)

Figure 2: Summary statistics: cumulative number of cases (panel A) and subpopulation sizes (panel (B)).

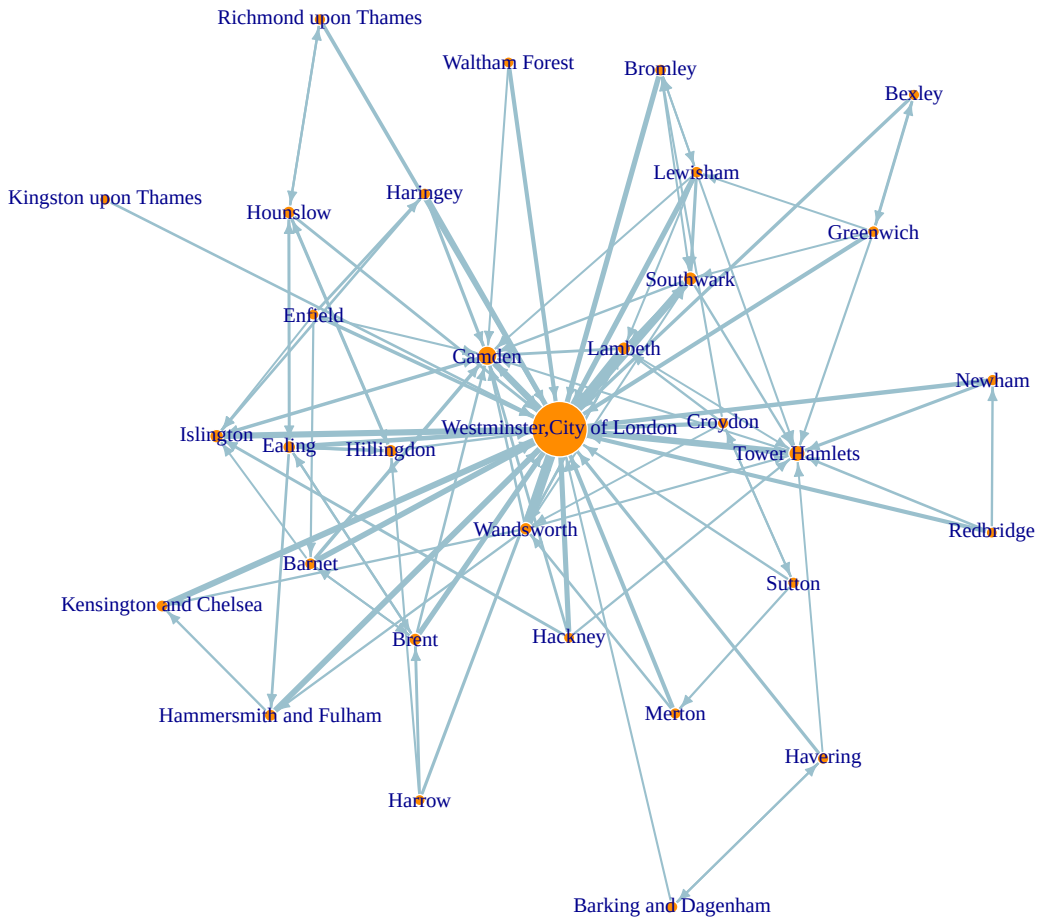

Figure 3: The graph of network $W^{(1)} = K$. This network is constructed using 2011 UK census data. Nodes represent London boroughs. Node sizes are proportional to their total in-degrees. Directed edges represent the amount of people going to work from one borough to another. Edge widths are proportional to their respective values in the graph. Only edges which exceed the 80% percentile are drawn.

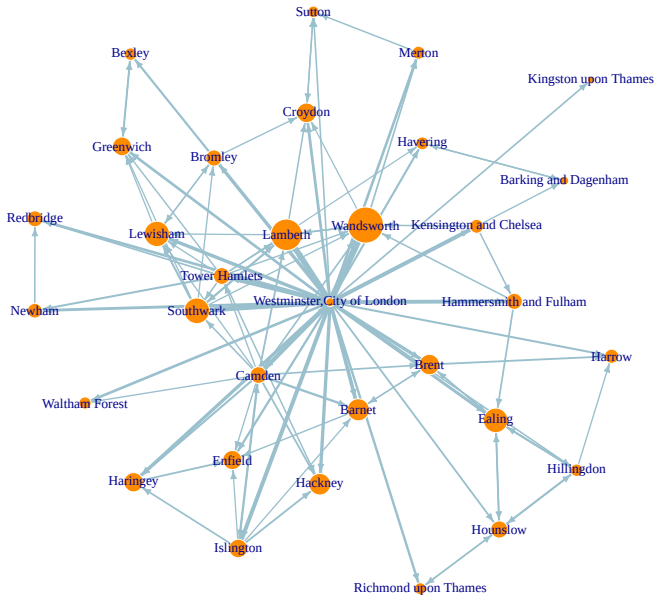

(A) Graph of $W^{(2)}$: defined by adjacency matrix $K^\top$

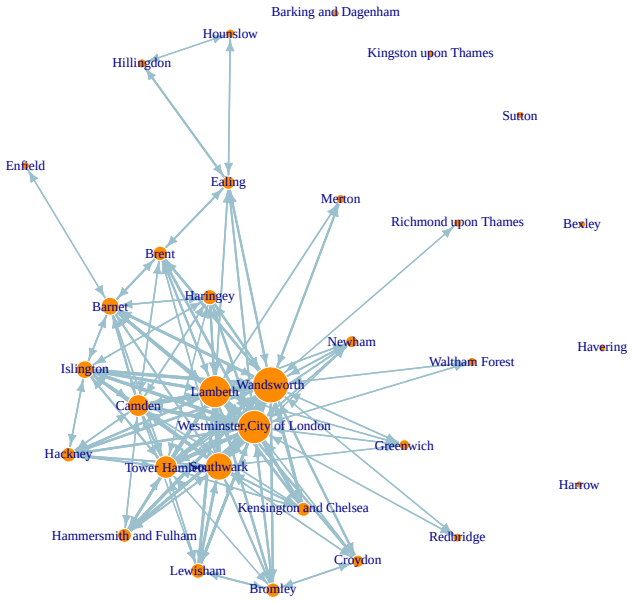

(B) Graph of $W^{(3)}$: defined by adjacency matrix $KK^\top$

Figure 4: Graphs of network $W^{(2)}$ and $W^{(2)}$. Both networks are constructed based on network $W^{(1)} = K$ shown in Figure 3. Nodes represent London boroughs. Node sizes are proportional to their total in-degrees. Directed edges represent the amount of people going to work from one borough to another. Edge widths are proportional to their respective values in the graph. Only edges which exceed the 80% percentile are drawn.

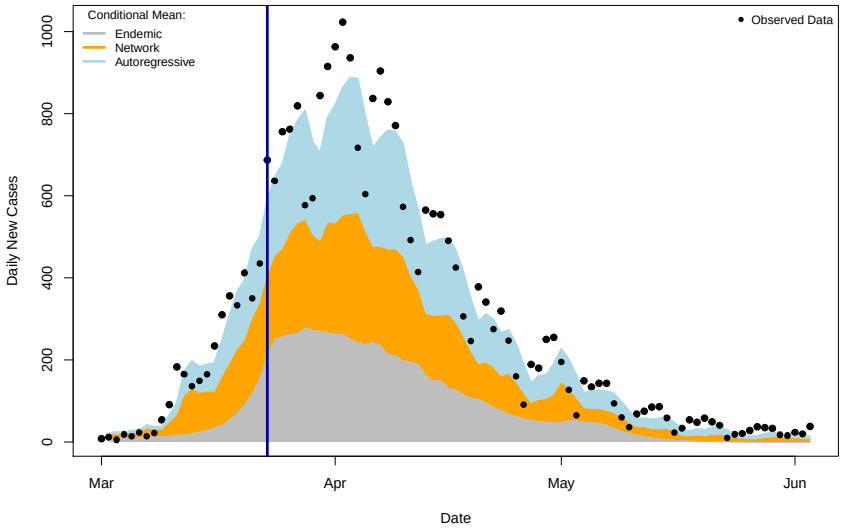

(A) Decomposition of total (expected) daily new cases in London across time

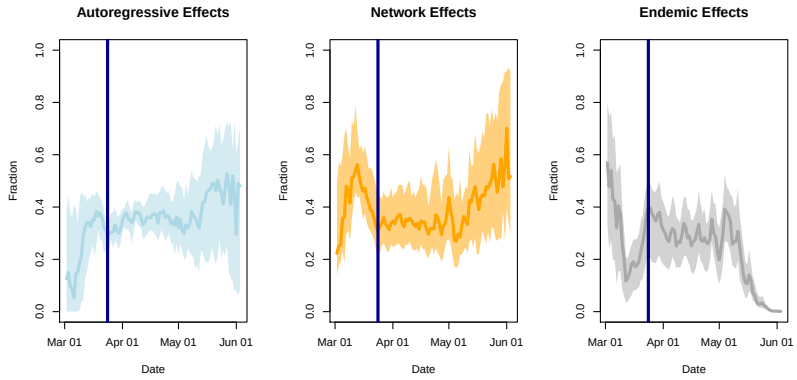

(B) The autoregressive, (total) network, and endemic effects as fractions: μ_{it}^{AR}/μ_{it}, μ_{it}^{NE}/μ_{it}, μ_{it}^{EN}/μ_{it}

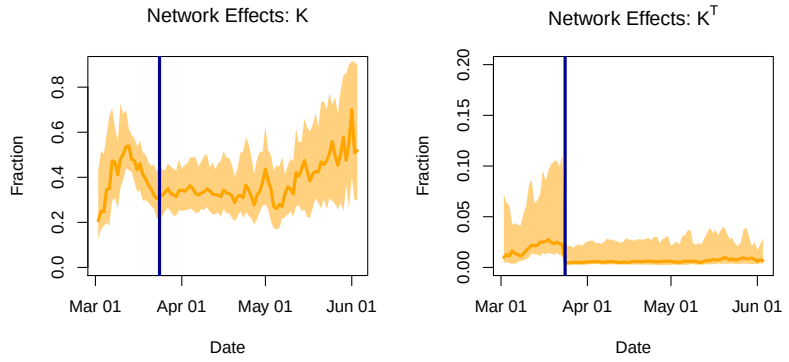

(C) The two network effects as fractions: $\mu_{it}^{NE,(g)}/\mu_{it}$, $g \in \{1,2\}$

Figure 5: Decomposition of the conditional mean: total daily new cases in London. Effects in Panel (A) are based on summing across all London boroughs. Bands in Panel (B) are 10% to 90% percentiles across all London boroughs. Solid lines are the median level.

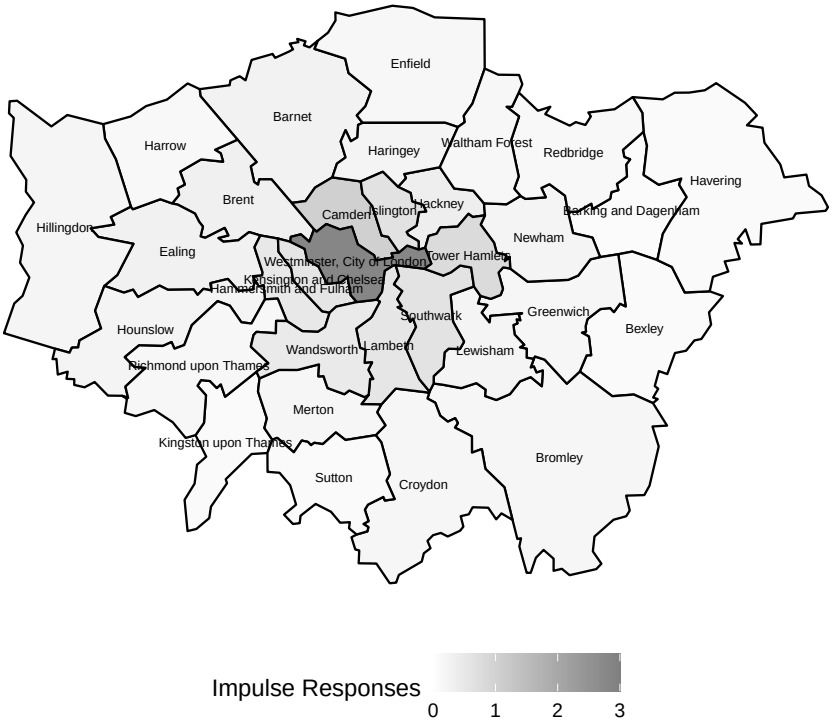

(A) Before lockdown

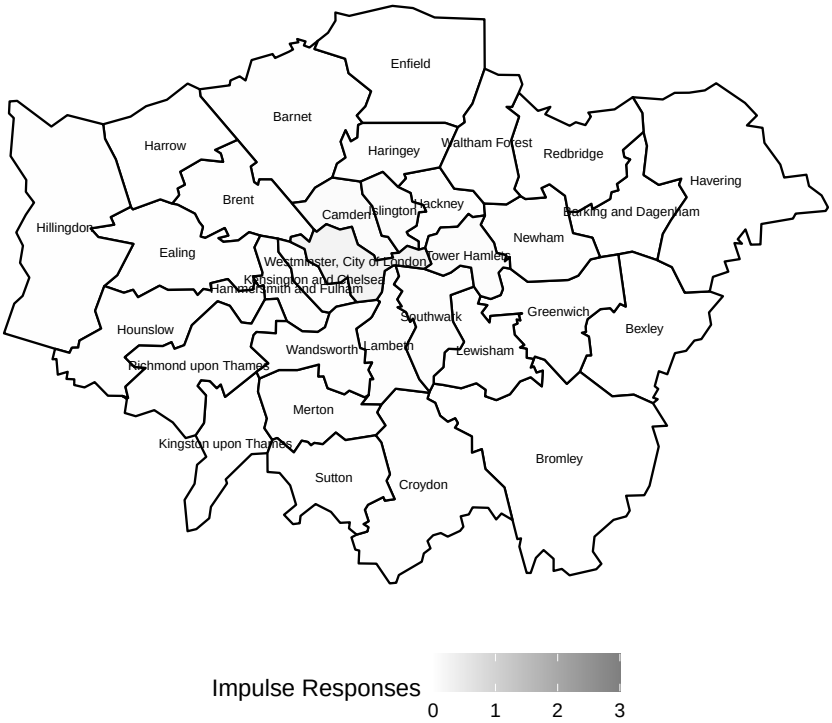

(B) After lockdown

Figure 6: Network impulse response functions (one-week horizon), from equation (19), before (panel (A)), and after (panel (B)), lockdown.

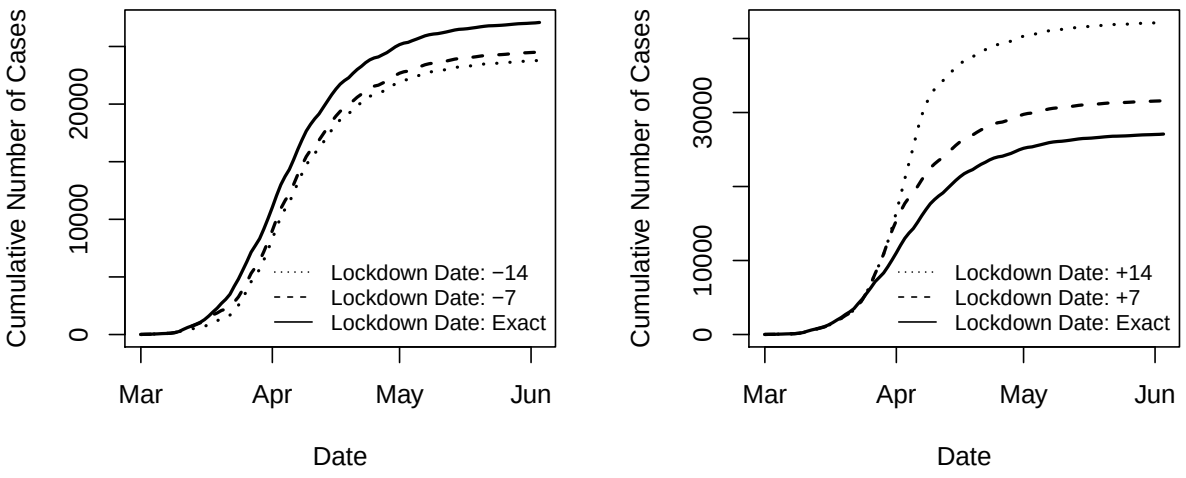

Figure 7: Cumulated number of cases over time (continuous black line) and counterfactual cumulated number of cases with earlier (left figure) and later (right figure) lockdown time.

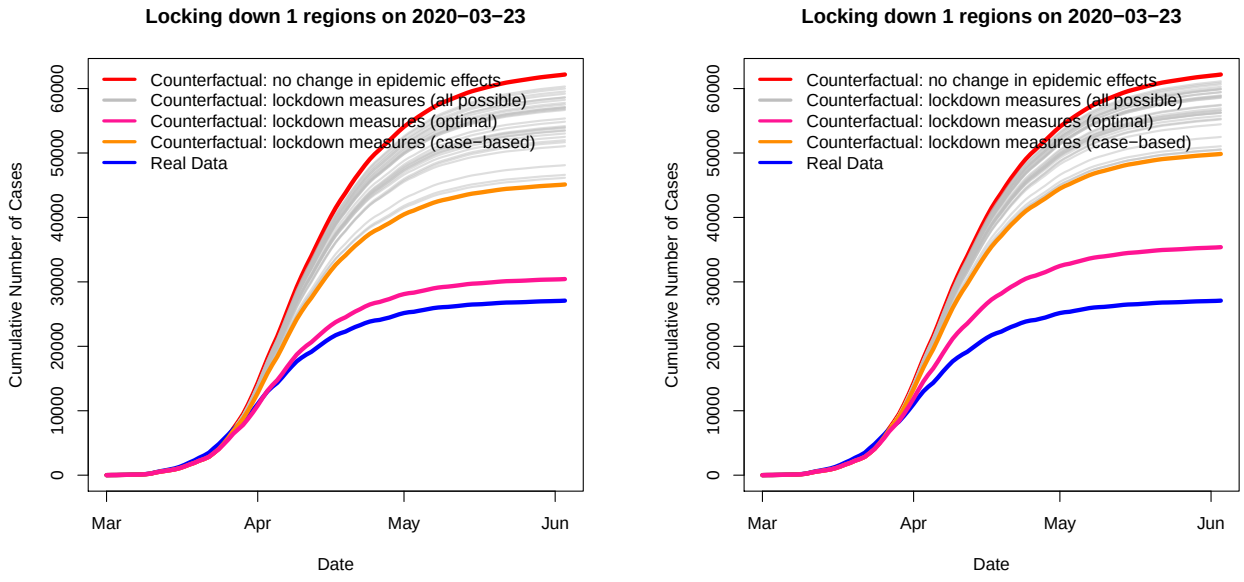

Figure 8: Cumulated number of cases over time: actual (blue line); counterfactual without lockdown (red line); counterfactual with lockdown of only one borough (all cases, grey); counterfactual with lockdown of only the borough with the most cases (orange line); counterfactual with optimally chosen borough (purple line). In the left figure, the counterfactual lockdown effects are computed by severing all linkages to and from those assumed to be in lockdown, while in the right figure, instead the corresponding columns and rows of matrix W(1) and W(2) are downscaled, respectively, to 82.7% and 16.6% of their original values.

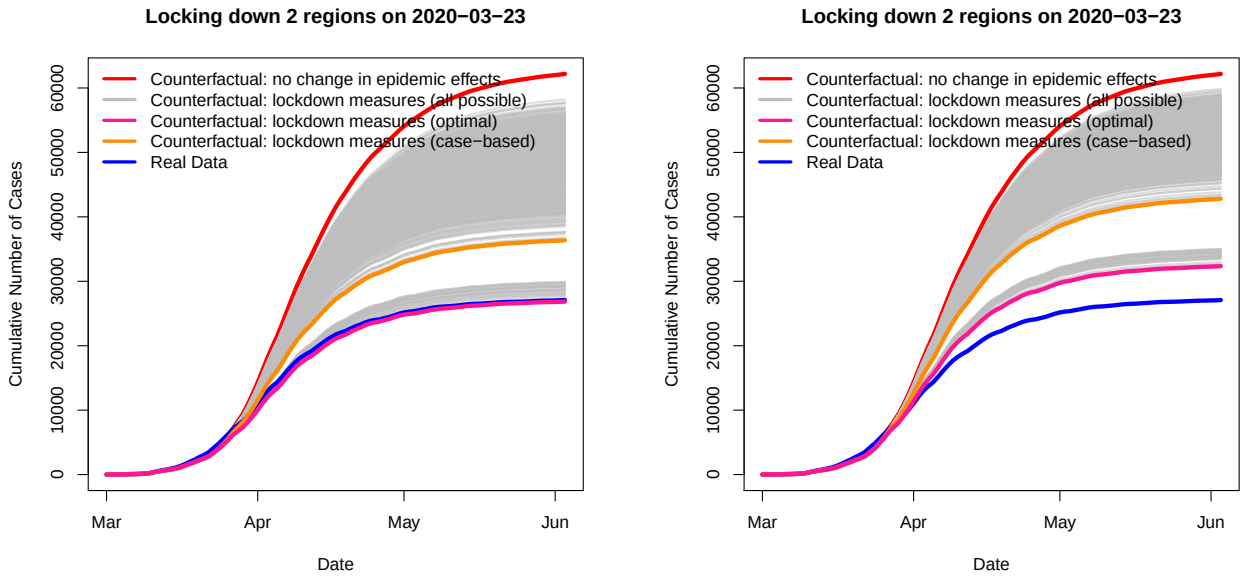

Figure 9: Cumulated number of cases over time: actual (blue line); counterfactual without lockdown (red line); counterfactual with lockdown of only two boroughs (all cases, grey); counterfactual with lockdown of only the two boroughs with most cases (orange line); counterfactual with optimally chosen two boroughs (purple line). In the left figure, the counterfactual lockdown effects are computed by severing all linkages to and from those assumed to be in lockdown, while in the right figure, instead the corresponding columns and rows of matrix W(1) and W(2) are downscaled, respectively, to 82.7% and 16.6% of their original values.

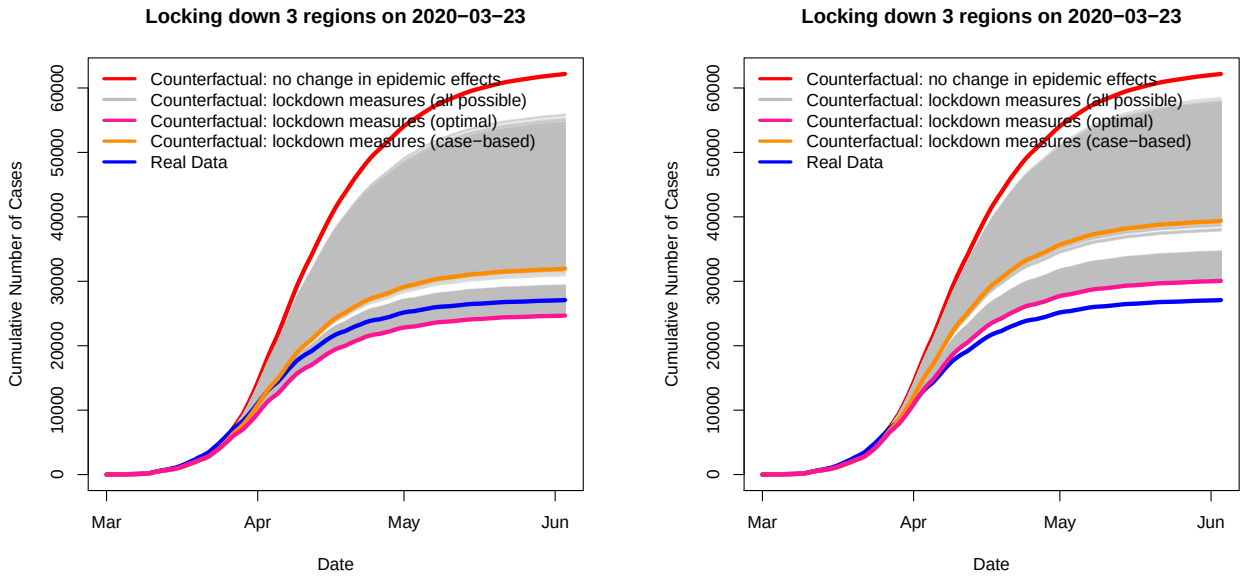

Figure 10: Cumulated number of cases over time: actual (blue line); counterfactual without lockdown (red line); counterfactual with lockdown of only two boroughs (all cases, grey); counterfactual with lockdown of only the three boroughs with most cases (orange line); counterfactual with optimally chosen three boroughs (purple line). In the left figure, the counterfactual lockdown effects are computed by severing all linkages to and from those assumed to be in lockdown, while in the right figure, instead the corresponding columns and rows of matrix W(1) and W(2) are downscaled, respectively, to 82.7% and 16.6% of their original values.

Appendix for "COVID-19 Spread in London: Network Effects and Optimal Lockdowns?"

A MLE details

Under a model $\mathcal{G}$, the log-likelihood function given the observed panel of daily new case numbers $\boldsymbol{Y}$ for parameters $\Theta(\mathcal{G}) = \left\{\gamma_0, \gamma, \left\{\phi_0^{(g)}\right\}_{g\in\mathcal{G}}, \left\{\phi^{(g)}\right\}_{g\in\mathcal{G}}, \boldsymbol{\beta}, \{\eta_i\}_{i=1}^n\right\}$ can be written as

$$\begin{aligned}
&\ell(\Theta \mid \boldsymbol{Y}, \mathcal{G}) \\
&= \sum_{i=1}^{n}\sum_{t=1}^{T}\Bigg\{ \log\Gamma\left(y_{it} + x_{i,t-1}\right) - \log\Gamma\left(x_{i,t-1}\right) - \log\Gamma(y_{it}+1) \\
&\qquad - \left(x_{i,t-1} + y_{it}\right)\log\left(1 + \mu_{it}x_{i,t-1}^{-1}\right) + y_{it}\left(\log\mu_{it} + x_{i,t-1}^{-1}\right)\Bigg\}
\end{aligned}$$

where $\Gamma(\cdot)$ is the standard gamma function; μ_{it} is the conditional mean of daily new cases defined as

$$\mu_{it} = \underbrace{\exp\left(\gamma_0 + D_t\gamma\right)x_{i,t-1}}_{\mu_{it}^{AR}} + \sum_{g\in\mathcal{G}}\underbrace{\exp\left(\phi_0^{(g)} + D_t\phi^{(g)}\right)\left(\sum_{j\neq i} w_{ij}^{(g)}x_{j,t-1}\right)}_{\mu_{it}^{NE,(g)}} + \underbrace{\exp\left(\boldsymbol{z}_t^\top\boldsymbol{\beta} + \eta_i\right)N_i}_{\mu_{it}^{EN}}.$$

Taking derivative, we have the following set of score functions:

$$\boldsymbol{S}(\Theta) = \left\{S_{\gamma_0}, S_\gamma, \left\{S_{\phi_0^{(g)}}\right\}_{g\in\mathcal{G}}, \left\{S_{\phi^{(g)}}\right\}_{g\in\mathcal{G}}, \boldsymbol{S}_{\boldsymbol{\beta}}, \{S_{\eta_i}\}_{i=1}^n\right\}(\Theta)$$

defined as

$$S_{\gamma_0}(\Theta) = \frac{\partial\ell}{\partial\gamma_0} = \sum_{i=1}^{n}\sum_{t=1}^{T}\frac{y_{it}-\mu_{it}}{1+x_{i,t-1}^{-1}\mu_{it}}\frac{\mu_{it}^{AR}}{\mu_{it}}, \quad S_\gamma(\Theta) = \frac{\partial\ell}{\partial\gamma} = \sum_{i=1}^{n}\sum_{t=1}^{T}\frac{y_{it}-\mu_{it}}{1+x_{i,t-1}^{-1}\mu_{it}}\frac{\mu_{it}^{AR}}{\mu_{it}}D_t,$$

$$S_{\phi_0^{(g)}}(\Theta) = \frac{\partial\ell}{\partial\gamma} = \sum_{i=1}^{n}\sum_{t=1}^{T}\frac{y_{it}-\mu_{it}}{1+x_{i,t-1}^{-1}\mu_{it}}\frac{\mu_{it}^{NE,(g)}}{\mu_{it}}, \quad S_{\phi^{(g)}}(\Theta) = \frac{\partial\ell}{\partial\gamma} = \sum_{i=1}^{n}\sum_{t=1}^{T}\frac{y_{it}-\mu_{it}}{1+x_{i,t-1}^{-1}\mu_{it}}\frac{\mu_{it}^{NE,(g)}}{\mu_{it}}D_t$$

for any $g\in\mathcal{G}$, and

$$\boldsymbol{S}_{\boldsymbol{\beta}}(\Theta) = \frac{\partial\ell}{\partial\boldsymbol{\beta}} = \sum_{i=1}^{n}\sum_{t=1}^{T}\frac{y_{it}-\mu_{it}}{1+x_{i,t-1}^{-1}\mu_{it}}\frac{\mu_{it}^{EN}}{\mu_{it}}\boldsymbol{z}_t,$$

$$S_{\eta_i}(\Theta) = \frac{\partial\ell}{\partial\eta_i} = \sum_{t=1}^{T}\frac{y_{it}-\mu_{it}}{1+x_{i,t-1}^{-1}\mu_{it}}\frac{\mu_{it}^{EN}}{\mu_{it}}$$

for each $i = 1,\ldots,n$.

Our MLE estimator $\widehat{\Theta}$ solves the system of score equations, that is $\boldsymbol{S}(\widehat{\Theta}) = \boldsymbol{0}$. Standard errors are then computed as $\sqrt{\text{diag}\left(\left\{\boldsymbol{S}(\widehat{\Theta})\boldsymbol{S}(\widehat{\Theta})^\top\right\}^{-1}\right)}$, based on the outer product of score vectors à la Berndt,

Hall, Hall, and Hausman (1974).

B Network impulse response function

To compute the NIRFs, we begin from the following fully vectorized representation. Define an $(n \times L)$-dimension vector $\widetilde{\boldsymbol{y}}_t^\top = \left[\boldsymbol{y}_t^\top, \boldsymbol{y}_{t-1}^\top, \ldots, \boldsymbol{y}_{t-L+1}^\top\right]^\top$ which concatenates the current and lagged-$(L-1)$ observations. Then

$$
\begin{aligned}
\mathbb{E}\left[\widetilde{\boldsymbol{y}}_{t+1} \mid \mathcal{F}_t\right] &= \begin{bmatrix} \boldsymbol{A}\boldsymbol{x}_t + \boldsymbol{\mu}_t^{EN} \\ \boldsymbol{y}_t \\ \boldsymbol{y}_{t-1} \\ \vdots \\ \boldsymbol{y}_{t-L+2} \end{bmatrix} \\
&= \begin{bmatrix} \boldsymbol{A}\sum_{\ell=0}^{L-1} \nu(\ell)\boldsymbol{y}_{t-\ell} \\ \boldsymbol{y}_t \\ \boldsymbol{y}_{t-1} \\ \vdots \\ \boldsymbol{y}_{t-L+2} \end{bmatrix} + \begin{bmatrix} \boldsymbol{\mu}_t^{EN} \\ \boldsymbol{0} \\ \boldsymbol{0} \\ \vdots \\ \boldsymbol{0} \end{bmatrix} \\
&= \underbrace{\begin{bmatrix} \boldsymbol{A}\nu(0) & \boldsymbol{A}\nu(1) & \cdots & \boldsymbol{A}\nu(L-2) & \boldsymbol{A}\nu(L-1) \\ \boldsymbol{I} & \boldsymbol{0} & \cdots & \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{I} & \cdots & \boldsymbol{0} & \boldsymbol{0} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{0} & \boldsymbol{0} & \cdots & \boldsymbol{I} & \boldsymbol{0} \end{bmatrix}}_{\widetilde{\boldsymbol{A}}(L)} \widetilde{\boldsymbol{y}}_t + \begin{bmatrix} \boldsymbol{\mu}_t^{EN} \\ \boldsymbol{0} \\ \boldsymbol{0} \\ \vdots \\ \boldsymbol{0} \end{bmatrix}.
\end{aligned}
$$

As a result, $\mathbb{E}\left[\widetilde{\boldsymbol{y}}_{t+\tau} \mid \mathcal{F}_t\right] = \left[\widetilde{\boldsymbol{A}}(L)\right]^\tau \widetilde{\boldsymbol{y}}_t + \text{const.}$, which implies that

$$
\text{NIRF}_{it}(\tau) = \sum_{j=1}^{n} \frac{\partial \mathbb{E}\left[y_{j,t+\tau} \mid \mathcal{F}_t\right]}{\partial y_{it}} = \sum_{j=1}^{n} \left[\widetilde{\boldsymbol{A}}(L)\right]_{ji}^{\tau},
$$

where the subscript (ji) of a matrix denotes its element on the j-th row and the i-th column. Estimations of the NIRF's are then calculated through plugging parameter estimates into the expression for matrix $\boldsymbol{A}$.

Table A1: Estimation results of the Network-SIR model of Section 1.3 without controlling for test-related variables. Results in columns (1)-(3) correspond, respectively, to specifications with only one of the following transmission networks: "work-to-home;" "home-to-work;" "home-to-home". Column (1,2,3) considers the three networks jointly while column (1,2) uses only the first two networks.

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Autoregressive effect										
$\exp(\gamma_0)$	0.381	0.033	0.456	0.034	0.424	0.032	0.365	0.036	0.344	0.038
$\exp(\gamma)$	0.472	0.050	0.519	0.045	0.446	0.042	0.492	0.057	0.522	0.066
Network effect										
$W^{(1)} = K$:										
$\exp\left(\phi_0^{(1)}\right)$	0.563	0.066					0.496	0.091	0.592	0.071
$\exp\left(\phi^{(1)}\right)$	0.694	0.098					0.594	0.189	0.637	0.102
$W^{(2)} = K^\top$:										
$\exp\left(\phi_0^{(2)}\right)$			0.076	0.021			0.000	0.000	0.041	0.021
$\exp\left(\phi^{(2)}\right)$			0.598	0.175			0.000	0.000	0.203	0.424
$W^{(3)} = KK^\top$:										
$\exp\left(\phi_0^{(3)}\right)$					0.280	0.035	0.056	0.048		
$\exp\left(\phi^{(3)}\right)$					0.518	0.073	0.696	0.779		
pseudo-R^2	81.86%		80.90%		81.58%		81.89%		81.89%	
# obs.	3008		3008		3008		3008		3008	

Table A2: Estimation results of the Network-SIR model of Section 1.3 adding day fixed effects. Results in columns (1)-(3) correspond, respectively, to specifications with only one of the following transmission networks: "work-to-home;" "home-to-work;" "home-to-home". Column (1,2,3) considers the three networks jointly while column (1,2) uses only the first two networks.

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Autoregressive effect										
$\exp(\gamma_0)$	0.301	0.024	0.295	0.018	0.283	0.025	0.261	0.026	0.261	0.026
$\exp(\gamma)$	0.455	0.049	0.468	0.032	0.483	0.050	0.523	0.064	0.523	0.064
Network effect										
$W^{(1)} = K$										
$\exp\left(\phi_0^{(1)}\right)$	0.271	0.041					0.266	0.043	0.266	0.043
$\exp\left(\phi^{(1)}\right)$	0.278	0.108					0.271	0.110	0.271	0.110
$W^{(2)} = K^\top$										
$\exp\left(\phi_0^{(2)}\right)$			0.061	0.020			0.053	0.020	0.053	0.020
$\exp\left(\phi^{(2)}\right)$			0.000	0.000			0.000	0.000	0.000	0.000
$W^{(3)} = KK^\top$										
$\exp\left(\phi_0^{(3)}\right)$					0.132	0.025	0.000	0.000		
$\exp\left(\phi^{(3)}\right)$					0.000	0.000	1.840	0.000		
pseudo-R^2	86.32%		86.36%		86.32%		86.34%		86.34%	
# obs.	3008		3008		3008		3008		3008	

Table A3: Estimation results of the Network-SIR model of Section 1.3 adding week fixed effects. Results in columns (1)-(3) correspond, respectively, to specifications with only one of the following transmission networks: "work-to-home;" "home-to-work;" "home-to-home". Column (1,2,3) considers the three networks jointly while column (1,2) uses only the first two networks.

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Autoregressive effect										
$\exp(\gamma_0)$	0.302	0.033	0.323	0.020	0.274	0.032	0.259	0.036	0.260	0.035
$\exp(\gamma)$	0.621	0.079	0.660	0.077	0.718	0.109	0.716	0.094	0.716	0.114
Network effect										
$W^{(1)} = K$										
$\exp\left(\phi_0^{(1)}\right)$	0.351	0.062					0.362	0.065	0.355	0.064
$\exp\left(\phi^{(1)}\right)$	0.810	0.174					0.666	0.219	0.797	0.175
$W^{(2)} = K^\top$										
$\exp\left(\phi_0^{(2)}\right)$			0.082	0.035			0.070	0.026	0.064	0.024
$\exp\left(\phi^{(2)}\right)$			0.102	0.213			0.000	0.000	0.000	0.000
$W^{(3)} = KK^\top$										
$\exp\left(\phi_0^{(3)}\right)$					0.216	0.038	0.000	0.000		
$\exp\left(\phi^{(3)}\right)$					0.420	0.101	1.562	0.000		
pseudo-R^2	81.25%		80.12%		81.96%		81.42%		81.27%	
# obs.	3008		3008		3008		3008		3008	

Table A4: Sensitivity analysis of the Network-SIR model of Section 1.3. Results in columns (1)-(3) correspond, respectively, to specifications with only one of the following transmission networks: "work-to-home;" "home-to-work;" "home-to-home". Column (1,2,3) considers the three networks jointly while column (1,2) uses only the first two networks. The policy date is changed to March 16.

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Autoregressive effect										
$\exp(\gamma_0)$	0.355	0.046	0.285	0.049	0.294	0.052	0.248	0.051	0.251	0.051
$\exp(\gamma)$	0.567	0.078	0.995	0.175	0.767	0.139	0.812	0.172	0.805	0.168
Network effect										
$\boldsymbol{W}^{(1)} = K$:										
$\exp\left(\phi_0^{(1)}\right)$	0.266	0.083					0.264	0.389	0.308	0.097
$\exp\left(\phi^{(1)}\right)$	1.523	0.475					1.528	2.255	1.307	0.413
$\boldsymbol{W}^{(2)} = K^\top$:										
$\exp\left(\phi_0^{(2)}\right)$			0.197	0.042			0.117	0.095	0.130	0.043
$\exp\left(\phi^{(2)}\right)$			0.233	0.055			0.000	0.000	0.005	0.084
$\boldsymbol{W}^{(3)} = KK^\top$:										
$\exp\left(\phi_0^{(3)}\right)$					0.310	0.066	0.048	0.294		
$\exp\left(\phi^{(3)}\right)$					0.436	0.102	0.000	0.000		
Testing-related endemic effect										
pos-to-test	-2.30	0.49	-1.50	0.33	-1.70	0.41	-2.19	0.48	-2.34	0.49
lag test	2.04	0.31	2.07	0.23	2.02	0.27	2.04	0.30	2.05	0.31
pseudo-R^2	81.98%		81.84%		81.88%		82.00%		82.00%	
# obs.	3008		3008		3008		3008		3008	

Table A5: Sensitivity analysis of the Network-SIR model of Section 1.3. Results in columns (1)-(3) correspond, respectively, to specifications with only one of the following transmission networks: "work-to-home;" "home-to-work;" "home-to-home". Column (1,2,3) considers the three networks jointly while column (1,2) uses only the first two networks. The policy date is changed to March 20.

Model $\mathcal{G}$	(1)		(2)		(3)		(1,2,3)		(1,2)	
Value	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*	*est.*	*se.*
Autoregressive effect										
$\exp(\gamma_0)$	0.333	0.040	0.301	0.044	0.314	0.042	0.282	0.046	0.282	0.046
$\exp(\gamma)$	0.596	0.078	0.946	0.140	0.703	0.098	0.709	0.122	0.708	0.120
Network effect										
$W^{(1)} = K$:										
$\exp\left(\phi_0^{(1)}\right)$	0.362	0.078					0.377	0.230	0.384	0.085
$\exp\left(\phi^{(1)}\right)$	1.134	0.247					1.073	0.664	1.059	0.234
$W^{(2)} = K^\top$:										
$\exp\left(\phi_0^{(2)}\right)$			0.120	0.026			0.057	0.053	0.057	0.028
$\exp\left(\phi^{(2)}\right)$			0.374	0.090			0.040	0.290	0.025	0.190
$W^{(3)} = KK^\top$:										
$\exp\left(\phi_0^{(3)}\right)$					0.245	0.044	0.004	0.143		
$\exp\left(\phi^{(3)}\right)$					0.621	0.116	0.104	9.481		
Testing-related endemic effect										
pos-to-test	-2.45	0.56	-1.79	0.35	-2.44	0.50	-2.58	0.57	-2.60	0.57
lag test	2.06	0.32	2.05	0.24	2.09	0.29	2.08	0.32	2.08	0.32
pseudo-R^2	81.98%		81.86%		81.87%		81.99%		81.99%	
# obs.	3008		3008		3008		3008		3008	