NBER WORKING PAPER SERIES

OPTIMAL LOCKDOWN IN A COMMUTING NETWORK

Pablo Fajgelbaum
Amit Khandelwal
Wookun Kim
Cristiano Mantovani
Edouard Schaal

Working Paper 27441
http://www.nber.org/papers/w27441

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2020

We thank Andy Atkeson for his comments. Schaal acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563) and the European Research Council Starting Grant 804095. We thank SafeGraph for making their data freely available. We thank Hyungmo Choi for providing excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Optimal Lockdown in a Commuting Network
Pablo Fajgelbaum, Amit Khandelwal, Wookun Kim, Cristiano Mantovani, and Edouard Schaal
NBER Working Paper No. 27441
June 2020
JEL No. C6,R38,R4

ABSTRACT

We study optimal dynamic lockdowns against Covid-19 within a commuting network. Our framework integrates canonical spatial epidemiology and trade models, and is applied to cities with varying initial viral spread: Seoul, Daegu and NYC-Metro. Spatial lockdowns achieve substantially smaller income losses than uniform lockdowns, and are not easily approximated by simple centrality-based rules. In NYM and Daegu—with large initial shocks—the optimal lockdown restricts inflows to central districts before gradual relaxation, while in Seoul it imposes low temporal but large spatial variation. Actual commuting responses were too weak in central locations in Daegu and NYM, and too strong across Seoul.

Pablo Fajgelbaum
Department of Economics
University of California, Los Angeles
Bunche Hall 8283
315 Portola Plaza
Los Angeles, CA 90095
and NBER
pfajgelbaum@gmail.com

Amit Khandelwal
Graduate School of Business
Columbia University
Uris Hall 606, 3022 Broadway
New York, NY 10027
and NBER
ak2796@columbia.edu

Wookun Kim
SMU Department Economics
3300 Dyer Street Suite 301
Dallas, TX 75205
wookunkim@smu.edu

Cristiano Mantovani
Universitat Pompeu Fabra
Ramon Trias Fargas, 23-25
08005 Barcelona
Spain
Barcelona
Spain
cristiano.mantovani@upf.edu

Edouard Schaal
Centre de Recerca en Economia Internacional
Universitat Pompeu Fabra
Ramon Trias Fargas, 23-25
08005 Barcelona
Spain
edouard.schaal@gmail.com

1 Introduction

Commuting networks are the backbone of cities, allowing interactions that are vital for economic growth. On a typical day, Manhattan receives as many commuters from neighboring counties as it has residents–about 1.6 million people. Two months after the onset of Covid-19, NYC metro commute flows were 49% below pre-pandemic levels. Weighing the economic costs against the benefits of stopping Covid-19, was this reduction too large or not large enough? To fight a highly infectious disease without a vaccine, public authorities must decide how to curtail movements. How should lockdown policies be set across locations and time? These questions have particular significance given the different patterns of lockdown implemented across locations connected via commuting and trade.[1]

In this paper we establish an efficient benchmark against which to measure the losses from uncoordinated or spatially uniform lockdown efforts. We study optimal dynamic lockdowns to fight pandemics in a commuting network using a framework that integrates standard spatial epidemiology and trade models.[2] In the model, a disease spreads through interactions of commuters at the workplace. Lockdown policies reduce the real income of workers who stay at home, and, indirectly, impact other locations as these workers consume fewer goods or services.[3] Our planning problem determines the fraction of each origin-destination commuting flow that is allowed to operate at each point in time, under a probability that a vaccine becomes available, to minimize the discounted economic costs and the loss of lives. Examples of how a planner could control the full commuting matrix include policies that close businesses, preclude commutes from specific areas, or disclose publicly the location of confirmed cases so residents take precautionary measures, as in the case of Seoul (see Argente et al., 2020).

We apply the model using real-time commuting data across districts in two South Korean cities, Seoul and Daegu, and cellphone mobility data across counties in the NYC Metro area (NYM). By studying different cities, we can compare optimal pandemic-fighting strategies across intensities of the initial virus shock and contrast them with the observed commuting responses. We analyze Korean cities because Korea has tested for Covid-19 at greater intensities than most countries, making the timeline of their case data more reliable.[4] Seoul is the largest city in Korea but experienced a very small caseload per capita, while Daegu (Korea's fourth largest city) experienced

[1] At the state level, lockdown mandates were announced fairly uniformly across bordering U.S. states, with a mean difference of 4 days and standard deviation of 3.5 days, although there has been variation in county-level policies as well. For example, New York, New Jersey, and Connecticut imposed almost simultaneous lockdown, while Illinois did so more than two weeks before Missouri (Raifman et al., 2020).

[2] The spatial SIR model that we formulate is closely related to the multi-city epidemic model in Arino and Van den Driessche (2003), in which the disease is transmitted from infected residents of location i to susceptible residents of location j when they meet in k. The quantitative trade model is in the style of Anderson and Van Wincoop (2003) and Eaton and Kortum (2002).

[3] Caliendo et al. (2018) and Monte et al. (2018) study diffusion of local shocks across and within cities in related gravity models.

[4] Korea had performed 0.878 tests per thousand people at the time of its 1000th patient compared to 0.086 in the U.S. Stock (2020), Manski and Molinari (2020), Korolev (2020), and Atkeson (2020a) discuss challenges that arise from infrequent testing of Covid-19.

the largest shock within Korea. We study NYM because of its economic importance and rapid spread.

We compute the optimal lockdown given the Covid-19 spread when lockdown policies were announced. The model matches pre-pandemic commuting flows and wages across locations (Korean districts and NYM counties). We estimate the transmission rate using data on the spatial distribution of new cases and commuting flows over time. To estimate the spatial frictions that determine the diffusion of lockdown through goods markets, we use confidential geocoded credit card expenditure data from Seoul.

Our first results show that the optimal lockdowns are not easily approximated by simple centrality-based rules. Instead, they depend on the geography of commute flows, real income, and the initial viral spread. In NYM and Daegu—where the virus initially spread very quickly—locations with high virus-diffusion potential are subject to a strict initial lockdown, eliminating 40% to 80% of pre-pandemic inflows depending on the city and location, which is relaxed over 3 to 6 months. In NYM many other locations are imposed an early lockdown, but only the top-3 central locations (Manhattan, Brooklyn, and Bronx) are still closed after 100 days, and remain so in expectation that a vaccine arrives. In contrast, in Seoul—where the initial spread of Covid-19 was much smaller—the planner initially locks down only a few locations of relatively high centrality. As the virus spreads, the lockdown intensifies and retains considerable spatial variation across locations.

Our second result shows large benefits from spatial targeting. Specifically, we find substantially lower real-income losses from allowing for spatial targeting compared to an optimal uniform lockdown (i.e., no spatial variation). Given the actual case count by April 30, spatial targeting would have led to 15%, 28%, and 50% lower economic costs in Daegu, Seoul, and NYM, respectively, than the optimal uniform lockdown.

Finally, we ask how the observed commuting reductions, resulting from a combination of government action and commuters' self-imposed precaution, compare with the optimal benchmark. On average across locations, commuting declines reached troughs of 79.3%, 37.6%, and 81.5% below pre-pandemic levels in Daegu, Seoul, and NYM before modestly reverting upward. In NYM and Daegu, these city-level declines are broadly consistent with the optimal benchmark. However, the most central (peripheral) locations exhibited a weaker (stronger) reduction in commuting than what would have been optimal. In Seoul, actual commuting reductions were too strong compared to the optimal across the city. As a result, across all three cities, the real income losses could have been much smaller through optimal spatial targeting, given the actual case count.

Studies of optimal control of epidemics in single-location economic models include Goldman and Lightwood (2002) and Rowthorn and Toxvaerd (2012) and, in the context of Covid-19, Atkeson (2020b), Alvarez et al. (2020), Jones et al. (2020), Piguillem and Shi (2020), Rowthorn (2020), and Rowthorn and Toxvaerd (2020), among others. Acemoglu et al. (2020), Baqaee et al. (2020), and Glover et al. (2020) among others study lockdown with heterogeneous agents.

Adda (2016) demonstrates the relevance of transportation networks for disease transmission

by exploiting exogenous variation in public-transport strikes, and Viboud et al. (2006) show that work-related flows correlate with the regional spread of influenza in the United States. Several studies document that travel across and within cities contributed to the spread of Covid-19.[5]

Early spatial SIR models were used to study influenza and measles (Rvachev and Longini Jr, 1985; Bolker and Grenfell, 1995). Several studies simulate specific policies in these frameworks, but empirical applications of optimal lockdown in realistic environments with human diseases are uncommon.[6] For Covid-19, Chinazzi et al. (2020) simulate the impact of travel restrictions on geographic diffusion of Covid-19, Birge et al. (2020) study location-specific but time-invariant lockdown, Giannone et al. (2020) simulate alternative location-specific patterns to reopening, and Argente et al. (2020) study targeted policies taking into account the benefits of case information disclosure.

Compared to this literature, our contribution is three-fold. First, we implement optimal lockdown over both time and space in a commuting network. Second, to evaluate the diffusion of economic costs through changes in spending we integrate a general-equilibrium trade framework. Third, we use real-time commuting data to estimate and compare the actual commuting responses over space with the optimal lockdowns. Our analysis demonstrates large advantages of optimal targeting in a commuting network.

2 Model

We use a standard spatial epidemiology model similar to Arino and Van den Driessche (2003). The general equilibrium corresponds to a standard quantitative gravity trade model as in Anderson and Van Wincoop (2003) and Eaton and Kortum (2002).

2.1 Spatial Diffusion

The economy consists of N locations in continuous time. Before the pandemic, in each location i there are $N_0(i)$ residents, of which a fraction $\lambda(i,j)$ commutes to j. We let $\mathbf{\Lambda}$ be the matrix of bilateral commuting flows such that $[\mathbf{\Lambda}]_{ij} = \lambda(i,j)$. Hence, $\sum_i \lambda(i,j) N_0(j)$ is the pre-pandemic number of jobs at j.[7]

At each time t, the surviving residents of location j are either susceptible, exposed, infected, or recovered from the disease in quantities $S(j,t)$, $E(j,t)$, $I(j,t)$, and $R(j,t)$, respectively. Susceptible agents become exposed after interacting with infected agents; exposed agents are latent

[5] Tian et al. (2020) argue that the Wuhan lockdown and suspending intra-city public transport delayed the spread of Covid-19 across China, and Fang et al. (2020) use real-time movement data to demonstrate that the lockdown reduced infection rates to cities across China. Kissler et al. (2020) show that differences in commuting across boroughs correlate with cases in New York.

[6] See Bussell et al. (2019) for a recent literature review. E.g., Germann et al. (2006), Eubank et al. (2004) and Drakopolous and Zheng (2017) study targeted policies against epidemics in spatial or network SIR models. Theoretical properties of optimal control problems in spatial SIR frameworks were studied by Rowthorn et al. (2009).

[7] These initial distributions could be the equilibrium of a spatial model as in Redding and Rossi-Hansberg (2017). Given the time frame of the analysis, we assume no job reallocation, so that the distribution of employer-employee matches changes only due to lockdown policies.

carriers who do not infect others and become infected at rate γ_I. Infected agents die at rate γ_D or recover and become immune at rate γ_R. The spatial distributions are collected in the (column) vectors $\mathbf{S}(t)$, $\mathbf{E}(t)$, $\mathbf{I}(t)$, and $\mathbf{R}(t)$.

The government can control the fraction $\chi(i,j,t)$ of commuting flows from i to j by imposing lockdown measures or incentives to commuters. Of the $1-\chi(i,j,t)$ workers not physically commuting, a fraction δ_u, $u=S,E,I,R$, telecommutes. Every agent is subject to the policy regardless of infection status. Among the infected, only the asymptomatic fraction ζ works. The lockdown policies are collected in the N by N matrix $\boldsymbol{\chi}(t)$.

The geographic spread of the disease depends on how infected and susceptible people interact in space via commuting.[8] Let $\tilde{S}(i,t) \equiv \sum_i \chi(j,i,t)\,\lambda(j,i)\,S(j,t)$ be the number of susceptible agents (from any origin) exposed in i, and $\tilde{I}(i,t) \equiv \zeta \sum_i \chi(j,i,t)\,\lambda(j,i)\,I(j,t)$ be the number of infected asymptomatic agents (from any origin) spreading the disease in i. In matrix form, $\tilde{\mathbf{S}}(t) \equiv \mathbf{H}_S(t)\,\mathbf{S}(t)$ and $\tilde{\mathbf{I}}(t) \equiv \mathbf{H}_I(t)\,\mathbf{I}(t)$, where $\mathbf{H}_S(t)$ and $\mathbf{H}_I(t)$ are spatial incidence matrices that depend on pre-pandemic commuting flows and current lockdown policies:

$$\mathbf{H}_u(\chi(t)) = \zeta_u\left(\boldsymbol{\Lambda}\cdot\boldsymbol{\chi}(t)\right)' \tag{1}$$

for $u=S,I$, where $\zeta_I=\zeta$ is the fraction of infected asymptomatic commuters, $\zeta_S=1$, and $\cdot$ is the element-wise product. The (ij) element of $\mathbf{H}_S(t)$ is the exposure of the $S(j,t)$ susceptible residents of location j to the $\tilde{I}(i,t)$ infected agents commuting to i. Similarly, the (ij) element of $\mathbf{H}_I(t)$ is the exposure of the $I(j,t)$ infected residents of location j to the $\tilde{S}(j,t)$ susceptible agents commuting to i.

The flow of new infections taking place in i is

$$M_i\left(\tilde{I}(i,t),\tilde{S}(i,t)\right), \tag{2}$$

where $M_i(\cdot)$ represents the matching process between infected and susceptible individuals meeting in i. The infections taking place in i are carried back by susceptible and infected agents to their residence. Of the infections taking place in i, a fraction $H_S(i,j,t)\,S(j,t)/\tilde{S}(i,t)$ corresponds to residents of j. Therefore, the flow of the new infections among location j's residents is:

$$\dot{S}(j,t) = -\sum_i \frac{H_S(i,j,t)\,S(j,t)}{\tilde{S}(i,t)} M_i\left(\tilde{I}(i,t),\tilde{S}(i,t)\right). \tag{3}$$

[8] We assume that susceptible and infected agents interact at the destination location. The infections could happen anywhere in that destination, e.g., at the workplace, train station, or restaurants. We could incorporate infections during trips to work (e.g. on the subway) by explicitly modeling links in the transport network where commuters traveling between different origin-destination pairs meet.

2.2 Real Income

The economic costs of lockdown enter through the distribution of real income,

$$U(i,t) = \frac{Y(i,t)}{P(i,t)}, \tag{4}$$

where $P(i,t)$ is the cost of living and $Y(i,t)$ is the nominal income of location-i residents:

$$Y(i,t) = \sum_{u=S,E,I,R} \sum_{j} N_u(i,j,t)\, w(j,t), \tag{5}$$

Here, $w(j,t)$ is the wage per efficiency unit in j at time t and $N_u(i,j,t)$ is the flow of efficiency units of type-u commuters from i to j:

$$N_u(i,j,t) = \zeta_u \left[\chi(i,j,t) + (1-\chi(i,j,t))\,\delta_u\right] \lambda(i,j)\, u(i,t), \tag{6}$$

for $u = S, E, I, R$, where $\zeta_I = \zeta$ is the fraction of asymptomatic infected and $\zeta_u = 1$ for $u \neq I$. The efficiency units flowing from i to j include those physically commuting, $\chi(i,j,t)$, and those not commuting scaled by the fraction of telecommuters, $(1-\chi(i,j,t))\,\delta_u$.

The jobs at j produce a good or service with time-invariant productivity $z(j)$. Consumers have a constant elasticity of substitution σ across the goods and services from different locations. Residents of j face trade costs $\tau(i,j) > 1$ when buying from i, which may include the cost of going to the store.[9] In equilibrium, all goods and service markets clear at all t:

$$w(i,t) \sum_{u=S,E,I,R} \sum_{j} N_u(j,i,t) = \sum_{j} s(i,j,t)\, Y(j,t) \text{ for all } i; \tag{7}$$

i.e., the total income of workers employed in i (in the left hand side) equals the aggregate expenditures in goods from i, where $s(i,j,t) \equiv \left(\frac{\tau(i,j)}{P(j,t)}\frac{w(i,t)}{z(i)}\right)^{1-\sigma}$ is the expenditures share of goods from i in location j, and where

$$P(j,t) = \left(\sum_{i} p(i,j,t)^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{8}$$

is the price index in j.

In equilibrium, $\{w(j,t), P(j,t)\}$ are such that (7) and (8) hold. The lockdown $\chi(i,j,t)$ affects the income of location-i residents and the supply of goods and services at j. Lower spending of location i diminishes purchases from locations with lower trade costs to i, and the lower supply of goods at j increases purchases from other locations. Empirically, these reallocations depend on the expenditure shares $s(i,j,t)$.

[9] We measure iceberg costs using pre-pandemic expenditure patterns within the city. We assume that these costs remain constant, although an increase through the pandemic would be straightforward to incorporate. Disease diffusion via shopping could also be incorporated at increased computational burden.

2.3 Planning Problem

A social planner chooses the lockdown matrix $\boldsymbol{\chi}(t)$ to maximize the present discounted value of the real income net of the cost ω per life lost. The previous economic equilibrium implies that the aggregate real income of location j defined in (4) depends on the spatial distribution of lockdown and residents by infection status:

$$U(j,t) \equiv U(j; \mathbf{S}(t), \mathbf{E}(t), \mathbf{I}(t), \mathbf{R}(t), \boldsymbol{\chi}(t)). \tag{9}$$

A vaccine and a cure become freely available with probability ν in every time period. If the cure occurs at time t, location j generates the real income $\bar{U}(j,t)$ forever.[10] The planning problem is:[11]

$$W = \max_{\boldsymbol{\chi}(t)} \int_0^\infty e^{-(r+\nu)t} \sum_j \left[U(j,t) + \frac{\nu}{r} \bar{U}(j,t) - \omega \gamma_D I(j,t) \right] dt \tag{10}$$

subject to

$$\dot{\mathbf{S}}(t) = -\mathbf{S}(t) \cdot \left[\mathbf{H}_S(\boldsymbol{\chi}(t))' \left(\mathbf{M}\left(\tilde{\mathbf{S}}(t), \tilde{\mathbf{I}}(t) \right) \cdot / \tilde{\mathbf{S}}(t) \right) \right], \tag{11}$$

$$\dot{\mathbf{E}}(t) = -\dot{\mathbf{S}}(t) - \gamma_I \mathbf{E}(t), \tag{12}$$

$$\dot{\mathbf{I}}(t) = \gamma_I \mathbf{E}(t) - (\gamma_R + \gamma_D) \mathbf{I}(t), \tag{13}$$

$$\dot{\mathbf{R}}(t) = \gamma_R \mathbf{I}(t), \tag{14}$$

where $\mathbf{M}\left(\tilde{\mathbf{S}}(t), \tilde{\mathbf{I}}(t) \right)$ is a vector with the new matches, and where $\tilde{\mathbf{S}}(t) = \mathbf{H}_S(\boldsymbol{\chi}(t)) \mathbf{S}(t)$ and $\tilde{\mathbf{I}}(t) = \mathbf{H}_I(\boldsymbol{\chi}(t)) \mathbf{I}(t)$ are the vectors of susceptible and infected agents present in each location.[12] The policy matrix $\boldsymbol{\chi}(t)$ controlled by the planner impacts aggregate real income $U(j,t)$ and the spatial incidence matrix $\mathbf{H}_u(t)$ through (9) and (1).

3 Data and Parametrization

The model parametrization uses case data and real-time commuting flows to estimate the virus transmission rate. The optimization is implemented starting at the lockdown announcement date using pre-pandemic data on bilateral commuting flows, wages, population, and spending. Our

[10] Hence, the real income at the recovery time is $\bar{U}(j,t) = U(j; 0, 0, 0, \mathbf{S}(t) + \mathbf{E}(t) + \mathbf{I}(t) + \mathbf{R}(t), \mathbf{1}_{J \times J})$ where $\mathbf{1}_{J \times J}$ is a matrix of ones. The framework could be readily extended to study the spatially optimal allocation of a vaccine that is in limited supply.

[11] The flow value for the planner's utility from location j is:

$$W_0(j,t) = (U(j,t) - \omega \gamma_D I(j,t)) dt + \nu dt e^{-r dt} \frac{\bar{U}(j,t)}{r} + e^{-r dt} (1 - \nu dt) W_0(j, t + dt).$$

This value includes the flow payoff $U(j,t) - \omega \gamma_D I(j,t)$, the probability νdt of transitioning to a vaccine and cure with associated present-discounted value $\frac{\bar{U}(j,t)}{r}$, and a probability νdt of continuing. Solving for the present discounted value for location j and adding up across locations yields (10).

[12] The notation $\cdot/$ stands for element-wise ratio.

units of analysis are the 25 districts in Seoul and 8 districts in Daegu. We define NYM to be 20 counties: 5 NYC boroughs, 5 counties in NY (Putnam, Rockland, Westchester, Nassau, Suffolk), 8 counties in New Jersey (Bergen, Essex, Hudson, Middlesex, Morris, Passaic, Somerset, Union), and 2 counties in Connecticut (Fairfield, New Haven).

3.1 Data

Covid-19 Data

The Seoul Metropolitan government released patient-level case data. We filed an Official Information Disclosure Act request to obtain patient-level data from the Daegu Metropolitan government. Using these individual records, we build a daily panel dataset with total confirmed cases in the Seoul and Daegu districts. County-level NYM cases come from Johns Hopkins University and the NY State Department of Health.[13]

There are limitations to using data on Covid-19 infections due to bias of the tested population and sensitivity of the tests. These limitations are more severe at the onset of the pandemic when testing intensity was low. Our estimation of the virus transmission rate uses data from the later periods when reporting and testing improved. These limitations are further mitigated in Korea where testing intensity has been high since the beginning of the pandemic.

Daily Commuter Data

For Seoul, we use district-to-district commute flows on the public transit system (subway and bus) through confidential individual trip-level data housed at the Seoul Big Data Campus. Passengers enter and exit public transit using a card with a unique identifier, which we use to identify the time, origin, and destination of each commute. We retain mornings on weekdays from 4am to 12pm (and 12pm to 8pm on weekends, since commutes start later) to capture the first commute leg. We aggregate over individual trips to bilateral commute flows across districts from January 2018 to April 2020. Pre-pandemic commute flows for Seoul are the 2019 averages.

For Daegu, we measure daily commuting using subway turnstile data from January 2018 to April 2020, made available by Daegu Metro Transit Corporation. We retain entries and exits using the same commute window as in Seoul. Stations' total entries (exits) plausibly capture the density of residents (jobs) if commuters enter (exit) the subway station closest to their residence (workplace). We aggregate the station-level data to the district level. Pre-pandemic commute flows for Daegu come from the 2015 Korean Population Census which records where people live and work.

For NYM, we measure county-to-county daily movements using cellphone data from SafeGraph (turnstile data within NYC is available but not for the suburban commuter rails). The data cover the sample period from January 1, 2020 to April 30, and are constructed from anonymized

[13]The Korean data allow us to exclude cases arriving from overseas travels that were stopped at the border and quarantined.

smartphone movement data collected daily at the Census block level.[14] We aggregate these data to the NYM counties. Pre-pandemic commute flows for NYM are the averages from January 1-20.

Appendix Table A.1 summarizes the commuter data and dates of key events described in the next subsection.

Wages and Population

Population and wage data for the Korean districts come from the 2019 resident registration database and the 2019 Statistical Yearbook of National Tax, respectively. NYM county wages are constructed from the 2017 LEHD Origin-Destination Employment Statistics, which reports the number of workers in a given wage bin by Census block.[15] NYM county population in 2019 are from the U.S. Census Bureau.

Credit Card Spending

We access confidential data on the universe of transactions using credit and debit cards issued by one of the top 3 banks in Korea, again housed at the Seoul Big Data Campus. The data contain transactions at brick-and-mortar shops within Seoul. We observe the addresses of cardholders and business, and the dollar values. We construct an average daily district-to-district spending matrix for 2019 among the 25 districts of Seoul, restricting the data to spending by Seoul residents. We use this data to estimate a pre-pandemic trade-to-distance elasticity and a same-district spending share to estimate the key parameters of the within-city trade block of the model in Section 3.3.

3.2 Commuting Responses

Figure 1 plots time fixed effects for commuting flows relative to pre-pandemic averages since January 2020 in each city.[16] In each figure the first vertical line denotes the first confirmed cases in the country, the middle line denotes the first case in the city, and the last line is the date of the lockdown announcement in each city. We overlay the daily counts of new Covid-19 scaled by population within each city.[17]

[14] It is possible that within-county flows increased as people stopped commuting and started working from home. SafeGraph provides a daily measure of the fraction of people who stay within their residence, which allows us to separate the within-county flows into those who stay at home versus those who commute to work within their county.

[15] Bins are defined by LEHD as lower than $1250, between $1250 and $3333, and above $3333 per month. Since only the number of workers per bin is available, we construct average wages by county as the weighted average of interval limits and the midpoint of intervals (i.e., $1250, $2291, $3333).

[16] For Seoul and NYM, where we observe bilateral commute flows, we estimate $\frac{N_{ijt}}{N_{ij,\tau(t)}} = \pi_t + \epsilon_{ijt}$, where the dependent variable is bilateral commute flows from i to j on date t relative to the average pre-pandemic flow and the $\tau(t)$ is a day-of-week and month dummy to control for seasonality and daily variation. For Daegu, since the raw data are from turnstiles, this figure reports $\frac{E_{it}}{E_{i,\tau(t)}} = \pi_t + \epsilon_{it}$, where the dependent variable is daily entries. The figure reports the π_t coefficients for each city. Standard errors are two-way clustered origin and destination in Seoul and NYM, and by origin in Daegu using the wild bootstrap to account for the small number of clusters (Cameron et al., 2008).

[17] Following Fang et al. (2020) we find a positive correlation between lags of commuting and new daily infections after controlling for location and date fixed effects. See Appendix B for details.

Before the first confirmed case within Korea on January 26, we observe no trends in commutes in both cities. That day, ridership fell by roughly 10% relative to pre-pandemic levels.[18] After the first confirmed case within Daegu on February 17, the virus spread quickly and commute responses declined steeply. In Seoul, there was no further change after the first confirmed case within the city on January 30. The right axes show that the spread of the disease was much larger in Daegu than in Seoul, which may explain the different commute responses between the two cities.

After virus spread throughout Korea during February, a national task force laid out a set of guidelines that included social distancing, working from home, canceling non-essential gatherings, and postponing the spring semester start date for all schools and universities. These guidelines were announced by the government on February 24. Following this announcement, public transit ridership continued to fall in both cities for roughly two weeks before trending back upwards. Overall, ridership fell 60.2% and 34.9% in Daegu and Seoul, respectively, relative to pre-pandemic levels. The standard errors suggest similar responses across districts within each city.

In contrast to the Korean cities, in NYM we do not observe declines in commuting before the first confirmed within-city case on March 3. New York State issued a lockdown order on March 22 that closed all non-essential businesses.[19] At the time of the lockdown announcement, commuting patterns were already trending downward and continued to fall until mid-April, at which point commuting was 70.1% below the pre-pandemic flows, before slightly returning back upwards.

3.3 Parametrization of the Model

To bring the model to the data the frequency is set to daily. Appendix Table A.2 summarizes the parameters and Appendix A describes the numerical resolution method.

Disease Dynamics

We calibrate parameters that determine transition rates across infection status. We show that the key qualitative features of the spatial-temporal patterns of optimal lockdown do not change for a range of values corresponding to existing estimates.

Following Ferguson et al. (2020) we set a value for γ_I consistent with an incubation period of 5.1 days and also consider results assuming 4.2 days as estimated by Sanche et al. (2020). Following the Wang et al. (2020), we set a recovery rate γ_R consistent with an average recovery time of 18 days and also show results assuming 10 days. Ferguson et al. (2020) obtain an infection fatality ratio of 0.9%, which we use as benchmark, and we also show results using a lower bound of 0.3% across studies reported by Hall et al. (2020). In a random sample of Iranian households, Alamian et al. (2019) estimate that 36% of cases infections are asymptomatic.[20] We use this number and also show that the key patterns are invariant to a 50% lower rate. For the fraction of telecommuters δ_u,

[18] Commuting falls one day before the first confirmed case due to a national holiday.

[19] New Jersey and Connecticut locked down a few days earlier but we assign the NY's lockdown date since Manhattan is the central node.

[20] This number equals the ratio between asymptomatic infections (12%) and infections (33%) in their sample.

Dingel and Neiman (2020) report that 46% of jobs in the U.S. could be done from home. According to a survey by Job Korea, 60% of workers are able to telecommute.[21] The probability of finding a vaccine is such that the expected time until a vaccine appears is 18 months, corresponding to the lower bound of existing estimates.[22] ρ matches an annual interest rate of 4%.

As benchmark we impose a value of life ω equal to 10 million USD on top of the expected present discounted value of wages,[23] but also implement the analysis for a range of values between 1/100 and 100 times the benchmark to trace a Pareto frontier.

Matching Function and Transmission rate

As is standard in the SIR model, we impose a multiplicative matching function. The total number of new infections taking place in location j is

$$M_j\left(\tilde{I}_j, \tilde{S}_j\right) = \beta_j \tilde{I}_j \tilde{S}_j. \tag{15}$$

β_j is a location-specific transmission rate capturing that some locations are more prone to contagion. We assume that, given the number of individuals interacting in a location, contagion is more prevalent in denser districts:[24]

$$\beta_j = \frac{\beta}{area_j}. \tag{16}$$

We set β to match the model-based infection dynamics to the Covid-19 case data. Using (3) and (15) the change in the number of susceptible agents is:

$$\Delta S(i,t) = -\beta\zeta\left[\sum_j \frac{1}{area_j}\chi(i,j,t)\,\lambda(i,j)\sum_{i'}\chi\left(i',j,t\right)\lambda\left(i',j\right)I\left(i',t\right)\right]S(i,t) + \varepsilon(i,t), \tag{17}$$

where $\lambda(i,j)$ is the pre-pandemic fraction of residents from i commuting to j at time t, $\chi(i,j,t)$ is the commuting from i to j at time t relative to pre-pandemic flows, and $\varepsilon(i,t)$ accounts for measurement error and other forces driving infections. We directly observe $\lambda(i,j,t)\,\chi(i,j,t)$ for Seoul and NYM, and for Daegu we apply the changes in entries and exits to the pre-pandemic flows to construct this variable. The susceptible and infected populations $S(i,t)$ and $I(i,t)$ are recovered using data on new infections, the calibrated transition rates $(\gamma_I, \gamma_E, \gamma_R)$ and the laws of motion (11) to (13). Given the asymptomatic rate ζ, we set β to minimize the sum of square errors $\sum_i \sum_t \varepsilon(i,t)^2$.

To mitigate concerns that the data on new cases is imprecise and driven by testing, we start

[21] See http://www.jobkorea.co.kr/GoodJob/Tip/View?News_No=16696.

[22] See https://www.nytimes.com/interactive/2020/04/30/opinion/coronavirus-covid-vaccine.html.

[23] This value is between 12 and 15 times the present discounted value of wages depending on the city.

[24] In the single-location SIR model, adjusting for area is not necessary. Here, this adjustment ensures that the aggregate infections are invariant to spatial aggregation. E.g., in an SIR model with S susceptible and I infected individuals equally divided among N locations, assuming away spatial interactions the aggregate number of infections is $\beta N\left(\frac{S}{N}\right)\left(\frac{I}{N}\right) = \frac{SI}{N}$, so that slicing a territory reduces infections. Normalizing the transmission rate by area yields instead SI, which is invariant to the number of locations.

the estimation 10 days after the peak in new cases in each city. This approach is consistent with the assumptions that the data on new cases became more precise in the latest periods. We show that the results are very similar if we start the estimation at the time of the peak of the caseload.

Appendix Figure A.1 shows the average number of new cases per district in the data and in the estimated model, both on average and for each location. The model replicates well the average number of new cases after the peak, and implies a fair amount of dispersion in the dynamics across locations. For the first week after patient zero, the estimation implies a city-level reproduction number (the number of new infections per infected individual) of 1.38 in Seoul, 1.32 in Daegu, and 2.27 in NYM. This reproduction number is in line with existing estimates using daily data in single-location SIR models. Specifically, Shim et al. (2020) estimate 1.5 in Korea during the first month of the pandemic and Fernández-Villaverde and Jones (2020) estimate 2.5 in New York and surrounding counties during the first days of the pandemic.[25]

Trade Model Parameters

We now parametrize the trade costs $\tau(i,j)$ and the productivities $z(i)$. We follow the approach in gravity models of specifying trade costs as a log-linear function of distance,

$$\tau(i,j) = \kappa_0 distance(i,j)^{\kappa_1} \tag{18}$$

for $i \neq j$ and $\tau(i,j) \equiv 1$. Adding an error term we obtain the gravity equation:

$$\ln X(i,j) = \psi(j) + \eta(i) - (\sigma-1)\kappa_1 \ln(distance(i,j)) + \varepsilon(i,j) \tag{19}$$

where $X(i,j)$ are district j's expenditures in goods and services produced by district i, and $\psi(j)$ and $\eta(i)$ are destination and origin fixed effects. Using the credit-card spending data from Seoul we estimate a coefficient on log distance $(\sigma-1)\kappa_1 = 1.48$ (0.034).[26] As in the Ramondo et al. (2016) estimation of within-country trade costs, we set $\sigma = 5$ to recover κ_1 and impose this parameter in Daegu and New York. Using the market-clearing condition (7) we recover, for each city, the distribution $z(i)$ and the scale parameter κ_0 of the trade cost to match the pre-pandemic data on wages $w(j)$ and the fraction of same-district expenditures in total expenditures, respectively. The latter is measured equal to 55% from the Seoul credit-card data, and assumed to be the same across cities.

[25] At each time t, the reproduction number is computed as the largest eigenvalue of the matrix $\text{diag}(\mathbf{S}(t))\mathbf{H}_S'\text{diag}(\boldsymbol{\beta})\mathbf{H}_I/(\gamma_D+\gamma_R)$, where $[\boldsymbol{\beta}]_j = \frac{\beta}{area_j}$. In a single-location SIR model, this number is $\beta/(\gamma_D+\gamma_R)$ but here it must be adjusted by the incidence matrix (Diekmann et al., 1990). Appendix Table A.2 reports the estimates of β. Unlike the city-level reproduction number, the scale of β is not comparable to existing estimates in single-location SIR models because of our spatial structure.

[26] Robust standard error is reported in parenthesis. The Commodity Flow Survey (CFS) in the U.S. is often used to estimate this parameter across regions in the U.S. E.g., Monte et al. (2018) estimate 1.29 using the bilateral trade flows between 123 CFS regions. Districts within Seoul are much more granular than the CFS regions.

4 Optimal Spatial Lockdowns

4.1 Centrality and Optimal Lockdown

We implement the model in each city. As initial condition, we use the spatial distribution of cases at the lockdown date. The left panel of Figure 2 shows a map with the eigenvector centrality of each location in the commuting network. This measure captures the potential of each location to diffuse Covid-19 if no lockdown is implemented.[27] The right panel shows the fraction of commuting inflows that is shut down relative to pre-pandemic flows under optimal lockdown.

We find broadly similar qualitative patterns in NYM and Daegu. The most central locations first experience a strong lockdown, of between 50%-80% in New York and 40%-60% in Daegu. In both cities the lockdown of central locations is relaxed over the course of 3 to 6 months, eventually affecting 5%-40% of inflows in the most central NYM locations and 10%-30% in Daegu. In NYM other locations also exhibit an early lockdown, but only the top 3 central counties (Manhattan, Brooklyn and Bronx) are closed after 100 days, and remain so in expectation that a vaccine arrives. In Daegu, a few peripheral locations exhibit strict lockdown for a long time.

These patterns contrast starkly with Seoul, where the spread of the virus was substantially smaller as shown in Figure 1. Despite the limited spread, the planner imposes a lockdown for a long period to restrain the growth of the disease. The planner first cracks down on a few locations of relatively high centrality for Covid diffusion, but maintains economic activity. It is only after the disease has spread that the lockdown intensifies across most locations.

The results demonstrate that no simple statistic guides the optimal lockdown. Instead, the optimal pandemic-fighting strategies over time and space depend on the full geography of commute patterns, the spatial distribution of real income, and on the spatial distribution and magnitude of the initial viral spread. When the spread is sufficiently large, the planner first places more weight on shutting down locations that are perceived as transmission hubs, even if they are the main sources of real income. When it is not, those locations are (relatively) spared. In either case, the policy maintains a considerable steady-state lockdown to avoid a re-emergence of the disease.[28]

The optimal lockdown paths are further visualized in Figure 3. For each city, the figure shows the lockdown every 30 days. The initial lockdown pattern radiates from geographically central locations and weakens over time.

4.2 Pareto Frontier: Uniform versus Spatially Optimal Lockdown

We compute a "Pareto" frontier describing the tradeoff between cases and economic costs.[29] We solve the optimal lockdown for values of life ω ranging between 1/100 and 100 times the benchmark.

[27] Specifically, it is defined as the eigenvector of the largest eigenvalue of the matrix governing Covid spread at time 0, $\text{diag}(\mathbf{S}(0))\mathbf{H}_S'\text{diag}(\boldsymbol{\beta})\mathbf{H}_I$.

[28] Figure A.2 shows that for a large shock affecting 1% of the population, or for a large value of life ω equal to 100 times the benchmark, the qualitative patterns in Seoul resemble those in Daegu and NYM. Figure 4 below shows the tradeoff between real income infections across a wide range of values of life.

[29] This is not a Pareto frontier in the usual sense. Instead, it traces the optimal allocation across marginal rates of substitution between value of life and real income.

To demonstrate the importance of spatially targeted policies, we also implement optimal uniform lockdown paths that are restricted to be constant over space. Figure 4 plots cumulative Covid-19 cases against cumulative economic cost at the last period of our data (April 30) across these values of ω, in both the spatial optimum and the uniform cases. We also show the economic costs and cases in the estimated model at the observed reductions in commuting flows.

We find large gains from implementing optimal spatial lockdown. Compared to uniform optimization, and given the actual number of cases, spatially targeted lockdown leads to 15%, 28%, and 50% lower economic costs in Daegu, Seoul, and NYM, respectively. In NYM and Daegu, the gap in economic cost between uniform and optimal policies grows for higher values of life, when the number of cumulative cases is low.

We also find that the economic costs and actual cumulative cases were far from the optimal. Under spatial targeting, the same number of cumulative cases could have been reached at 13%, 18%, and 33% lower economic costs in Daegu, Seoul, and NYM, respectively.

4.3 Optimal and Observed Commuting Reductions

We now compare observed reductions in commuting with the model's prediction of optimal flows. The left panel of Figure 5 shows the aggregate commuting flows relative to pre-pandemic values in the data and under optimal lockdown. Since the optimal policies are implemented at the time of lockdown, they are shown as a flat line until that time. The right panel shows inflows, with a breakdown between high and low centrality locations. In Daegu and NYM, the actual city-level reductions in commuting were close to the model-implied optimal benchmark. In New York, a 40% drop took place during the time leading up to the lockdown. In the model, the optimal lockdown in that period is very similar. We find a similar pattern in Daegu, where the average lockdown is relaxed over time, as observed in the data. However, we find that the most central (peripheral) locations of both NYM and Daegu exhibited a weaker (stronger) reduction in commuting than the optimal response. In Seoul, the actual reductions were stronger than in the data in all locations. These differences explain why the estimated economic costs from actual commuting responses were much larger than the spatially optimal ones, as observed in Figure 4.

4.4 Robustness of Optimal Lockdown Patterns

We implemented robustness with respect to key parameters, as described in Section 3.3. We consider a lower infection fatality ratio (0.03%); a faster recovery time (10 days); an estimation of the transmission rate starting at the peak of new cases; an asymptomatic rate of half the benchmark; a large shock such that 1% of the population is infected; twice the value of life of the benchmark; and a shorter incubation (4.1 days). Appendix Figure A.3 shows that the qualitative patterns of optimal lockdown from the benchmark are similar across these specifications, except for Seoul given a large shock, as previously mentioned. Doubling the value of life or introducing a large shock leads to stronger initial lockdown, in particular for central locations, while a lower death rate weakens it.

5 Conclusion

Our framework could be applied to other spatial scales, such as across cities, states, or countries. The model could be adapted to study the optimal spatial deployment of a vaccine in limited supply, or to account for the fact that shopping linkages, which we incorporated via a standard quantitative trade model, also spread the disease. Also, given the time frame and high frequency of the analysis, we have assumed that worker-job matches are kept constant throughout the pandemic. The model could be modified to account for job reallocations and for the fact that resuming work may take time once a lockdown ends.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. *National Bureau of Economic Research Working Paper* (27102).

Adda, J. (2016). Economic activity and the spread of viral diseases: Evidence from high frequency data. *The Quarterly Journal of Economics 131*(2), 891–941.

Alamian, A., S. Pourbakhsh, A. Shoushtari, and H. Keivanfar (2019). Seroprevalence investigation of newcastle disease in rural poultries of the northern provinces (golestan, gilan, and mazandaran) of iran. *Archives of Razi Institute 74*(4), 365–373.

Alvarez, F., D. Argente, and F. Lippi (2020). A simple planning problem for covid-19 lockdown. *National Bureau of Economic Research Working Paper* (26981).

Anderson, J. E. and E. Van Wincoop (2003). Gravity with gravitas: a solution to the border puzzle. *American Economic Review 93*(1), 170–192.

Argente, D. O., C.-T. Hsieh, and M. Lee (2020). The cost of privacy: Welfare effect of the disclosure of covid-19 cases. Technical report, National Bureau of Economic Research.

Arino, J. and P. Van den Driessche (2003). A multi-city epidemic model. *Mathematical Population Studies 10*(3), 175–193.

Atkeson, A. (2020a). How deadly is covid-19? understanding the difficulties with estimation of its fatality rate. Technical report.

Atkeson, A. (2020b). What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. *National Bureau of Economic Research Working Paper* (26867).

Baqaee, D. R., E. Fahri, M. Mina, and J. Stock (2020). Reopening scenarios. *National Bureau of Economic Research Working Paper* (27244).

Birge, J., O. Candogan, and Y. Feng (2020). Controlling epidemic spread: Reducing economic losses with targeted closure. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (57).

Bolker, B. and B. T. Grenfell (1995). Space, persistence and dynamics of measles epidemics. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences 348*(1325), 309–320.

Bussell, E. H., C. E. Dangerfield, C. A. Gilligan, and N. J. Cunniffe (2019). Applying optimal control theory to complex epidemiological models to inform real-world disease management. *Philosophical Transactions of the Royal Society B 374*(1776), 20180284.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2018). The impact of regional and sectoral productivity changes on the us economy. *The Review of economic studies 85*(4), 2042–2096.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008, aug). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics 90*(3), 414–427.

Chinazzi, M., J. T. Davis, M. Ajelli, C. Gioannini, M. Litvinova, S. Merler, A. P. y Piontti, K. Mu, L. Rossi, K. Sun, et al. (2020). The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak. *Science 368*(6489), 395–400.

Diekmann, O., J. A. P. Heesterbeek, and J. A. Metz (1990). On the definition and the computation of the basic reproduction ratio r 0 in models for infectious diseases in heterogeneous populations. *Journal of mathematical biology 28*(4), 365–382.

Dingel, J. I. and B. Neiman (2020). How many jobs can be done at home? Technical report, National Bureau of Economic Research.

Drakopolous, K. and F. Zheng (2017). Network effects in contagion processes: Identification and control. Technical report.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Eubank, S., H. Guclu, V. A. Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang (2004). Modelling disease outbreaks in realistic urban social networks. *Nature 429*(6988), 180–184.

Fang, H., L. Wang, and Y. Yang (2020, March). Human mobility restrictions and the spread of the novel coronavirus (2019-ncov) in china. Working Paper 26906, National Bureau of Economic Research.

Ferguson, N., D. Laydon, G. Nedjati Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunuba Perez, G. Cuomo-Dannenburg, et al. (2020). Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid19 mortality and healthcare demand.

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.

Germann, T. C., K. Kadau, I. M. Longini, and C. A. Macken (2006). Mitigation strategies for pandemic influenza in the united states. *Proceedings of the National Academy of Sciences 103*(15), 5935–5940.

Giannone, E., N. Paixao, and X. Pang (2020). Pandemic in an interregional model - staggered restart. Technical report, Working Paper.

Glover, A., J. Heathcote, D. Krueger, and J. V. Rios Rull (2020). Health versus wealth: On the distributional effects of controlling a pandemic.

Goldman, S. M. and J. Lightwood (2002). Cost optimization in the sis model of infectious disease with treatment. *Topics in Economic Analysis & Policy 2*(1).

Hall, R. E., C. I. Jones, and P. J. Klenow (2020). Trading off consumption and covid-19 deaths. Technical report, Working Paper.

Jones, C., T. Philippon, and V. Venkateswaran (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. *Working Paper*.

Kissler, S., N. Kishore, M. Prabhu, D. Goffman, Y. Beilin, R. Landau, C. Gyamfi-Bannerman, B. Bateman, D. Katz, J. Gal, et al. (2020). Reductions in commuting mobility predict geographic differences in sars-cov-2 prevalence in new york city.

Korolev, I. (2020). Quantifying social interactions using smartphone data. Working paper, Binghamton University.

Manski, C. F. and F. Molinari (2020, April). Estimating the covid-19 infection rate: Anatomy of an inference problem. Working Paper 27023, National Bureau of Economic Research.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, migration, and local employment elasticities. *American Economic Review 108*(12), 3855–90.

Piguillem, F. and L. Shi (2020). The optimal covid-19 quarantine and testing policies. *Working Paper*.

Raifman, J., K. Nocka, D. Jones, J. Bor, S. Lipson, J. Jay, and P. Chan (2020). Covid-19 us state policy database. Technical report.

Ramondo, N., A. Rodríguez-Clare, and M. Saborío-Rodríguez (2016). Trade, domestic frictions, and scale effects. *American Economic Review 106*(10), 3159–84.

Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics 9*, 21–58.

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

Rowthorn, R. (2020). A cost-benefit analysis of the covid-19 disease. *Working Paper*.

Rowthorn, R. and F. Toxvaerd (2020). The optimal control of infectious diseases via prevention and treatment. *Cambridge Working Papers in Economics 2027*.

Rowthorn, R. E., R. Laxminarayan, and C. A. Gilligan (2009). Optimal control of epidemics in metapopulations. *Journal of the Royal Society Interface 6*(41), 1135–1144.

Rvachev, L. A. and I. M. Longini Jr (1985). A mathematical model for the global spread of influenza. *Mathematical biosciences 75*(1), 3–22.

Sanche, S., Y. T. Lin, C. Xu, E. Romero-Severson, N. Hengartner, and R. Ke (2020). Early release-high contagiousness and rapid spread of severe acute respiratory syndrome coronavirus 2.

Shim, E., A. Tariq, W. Choi, Y. Lee, and G. Chowell (2020, apr). Transmission potential and severity of COVID-19 in south korea. *International Journal of Infectious Diseases 93*, 339–344.

Stock, J. H. (2020, March). Data gaps and the policy response to the novel coronavirus. Working Paper 26902, National Bureau of Economic Research.

Tian, H., Y. Liu, Y. Li, C.-H. Wu, B. Chen, M. U. Kraemer, B. Li, J. Cai, B. Xu, Q. Yang, et al. (2020). An investigation of transmission control measures during the first 50 days of the covid-19 epidemic in china. *Science 368*(6491), 638–642.

Viboud, C., O. N. Bjørnstad, D. L. Smith, L. Simonsen, M. A. Miller, and B. T. Grenfell (2006). Synchrony, waves, and spatial hierarchies in the spread of influenza. *science 312*(5772), 447–451.

Wang, H., Z. Wang, Y. Dong, R. Chang, C. Xu, X. Yu, S. Zhang, L. Tsamlag, M. Shang, J. Huang, et al. (2020). Phase-adjusted estimation of the number of coronavirus disease 2019 cases in wuhan, china. *Cell discovery 6*(1), 1–8.

Figure 1: Commute Response and Disease Spread

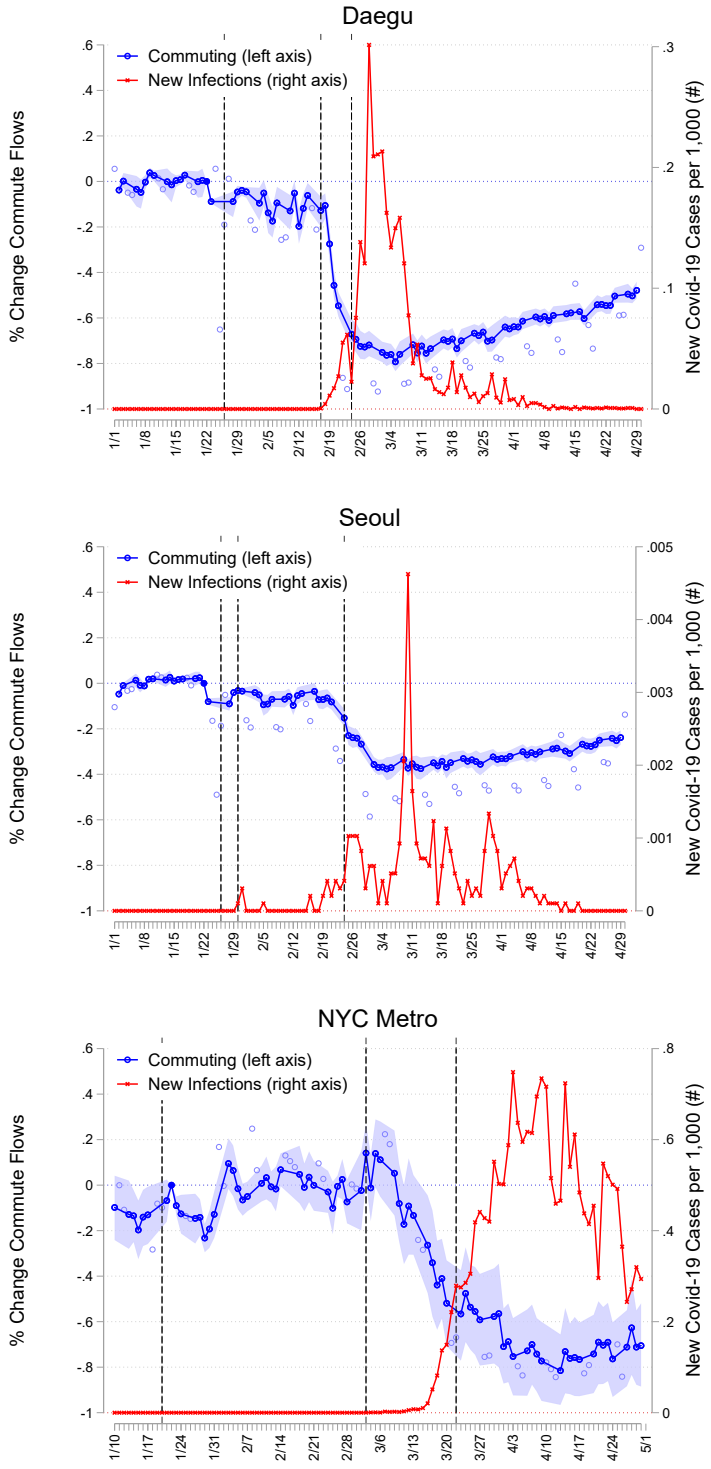

Note: Figure reports the average daily changes in commute flows relative to the pre-pandemic levels, corresponding to the time fixed effects from the equation in Footnote 16. The time fixed effects are normalized to 1 on January 22, 2020 and reported in logs. Weekdays (weekends/holidays) are denoted in darker (lighter) blue circles. The sample size is 968 for Daegu, 75,625 for Seoul, and 46,325 for NYM. Each observation is a district by date pair for Daegu and a district/county-pair by date tuple for Seoul and NYM from January 1, 2020 to April 30. The first vertical line denotes the date of the first case in Korea (top/middle) and the U.S. (bottom). The middle and the last vertical lines denote the date of the first case and the lockdown announcement in each city, respectively. Wild bootstrap standard errors are clustered by district for Daegu, and twoway-clustered by origin and destination for Seoul and NYM. Error bars show 95% confidence intervals. The right axis reports the average (across locations) daily new Covid-19 cases in Daegu (top), Seoul (middle) and NYM (bottom).

Figure 2: Centrality of Commuting Locations and Optimal Policies

(a) Daegu

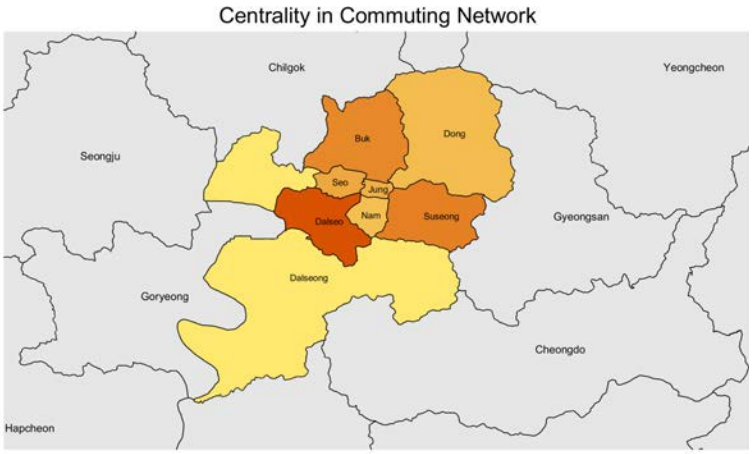

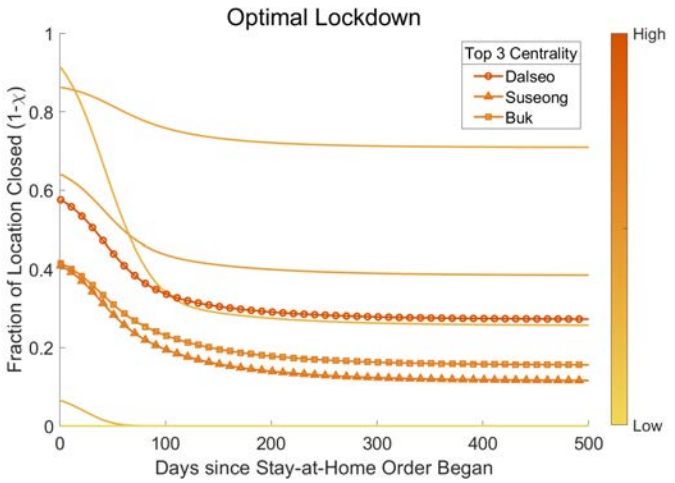

(b) Seoul

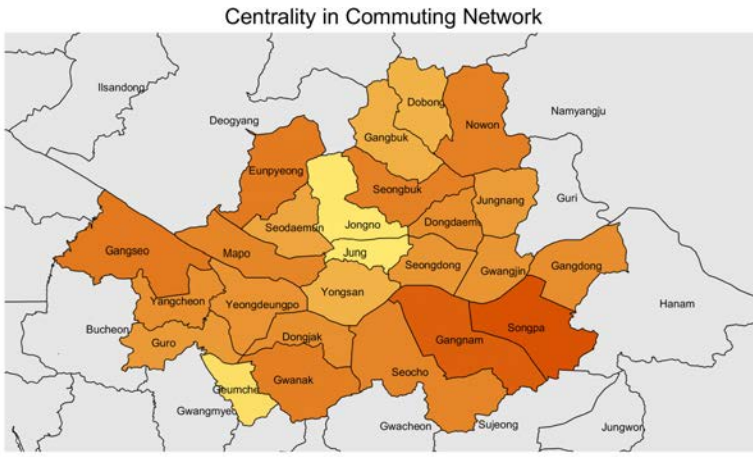

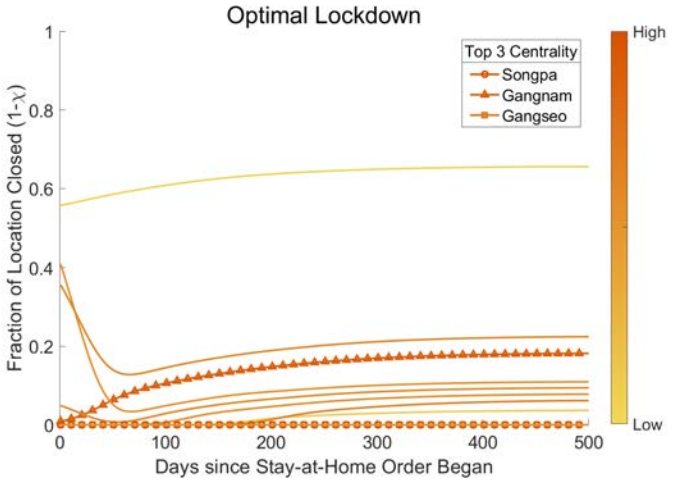

(c) NYC Metro

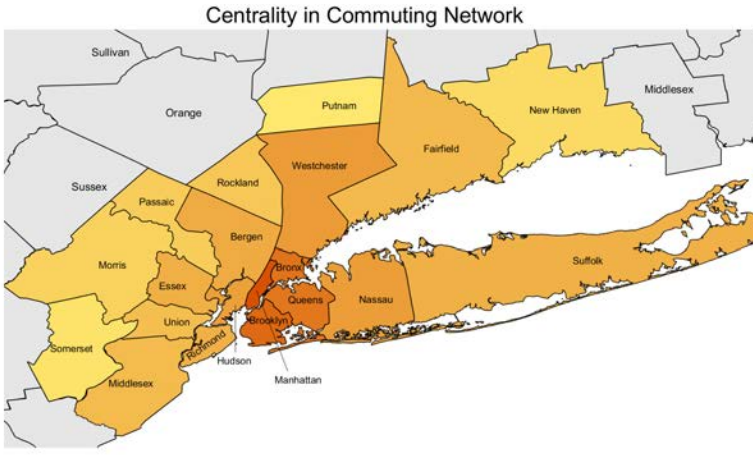

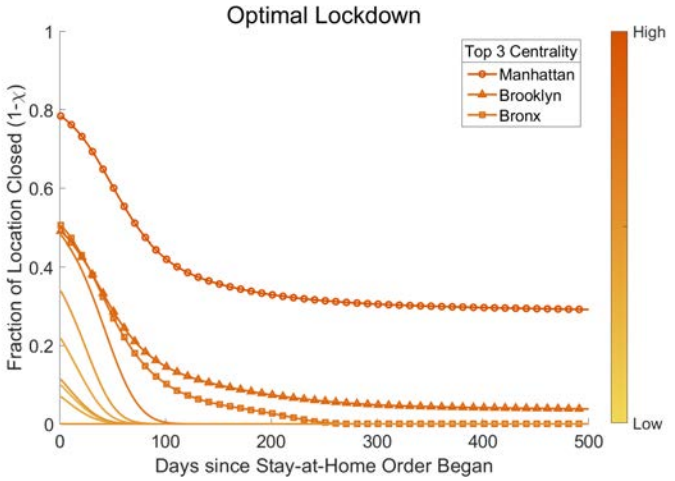

Note: The left panel denotes the (log) centrality of a location (see footnote 27), normalized so that the most central location is 1. The right panel plots the optimal policies over time for each location in the network. The color of the line represents the centrality of the location in the network. The three most central locations in the network are indicated in the legend.

Figure 3: Optimal Lockdown over Districts and Time

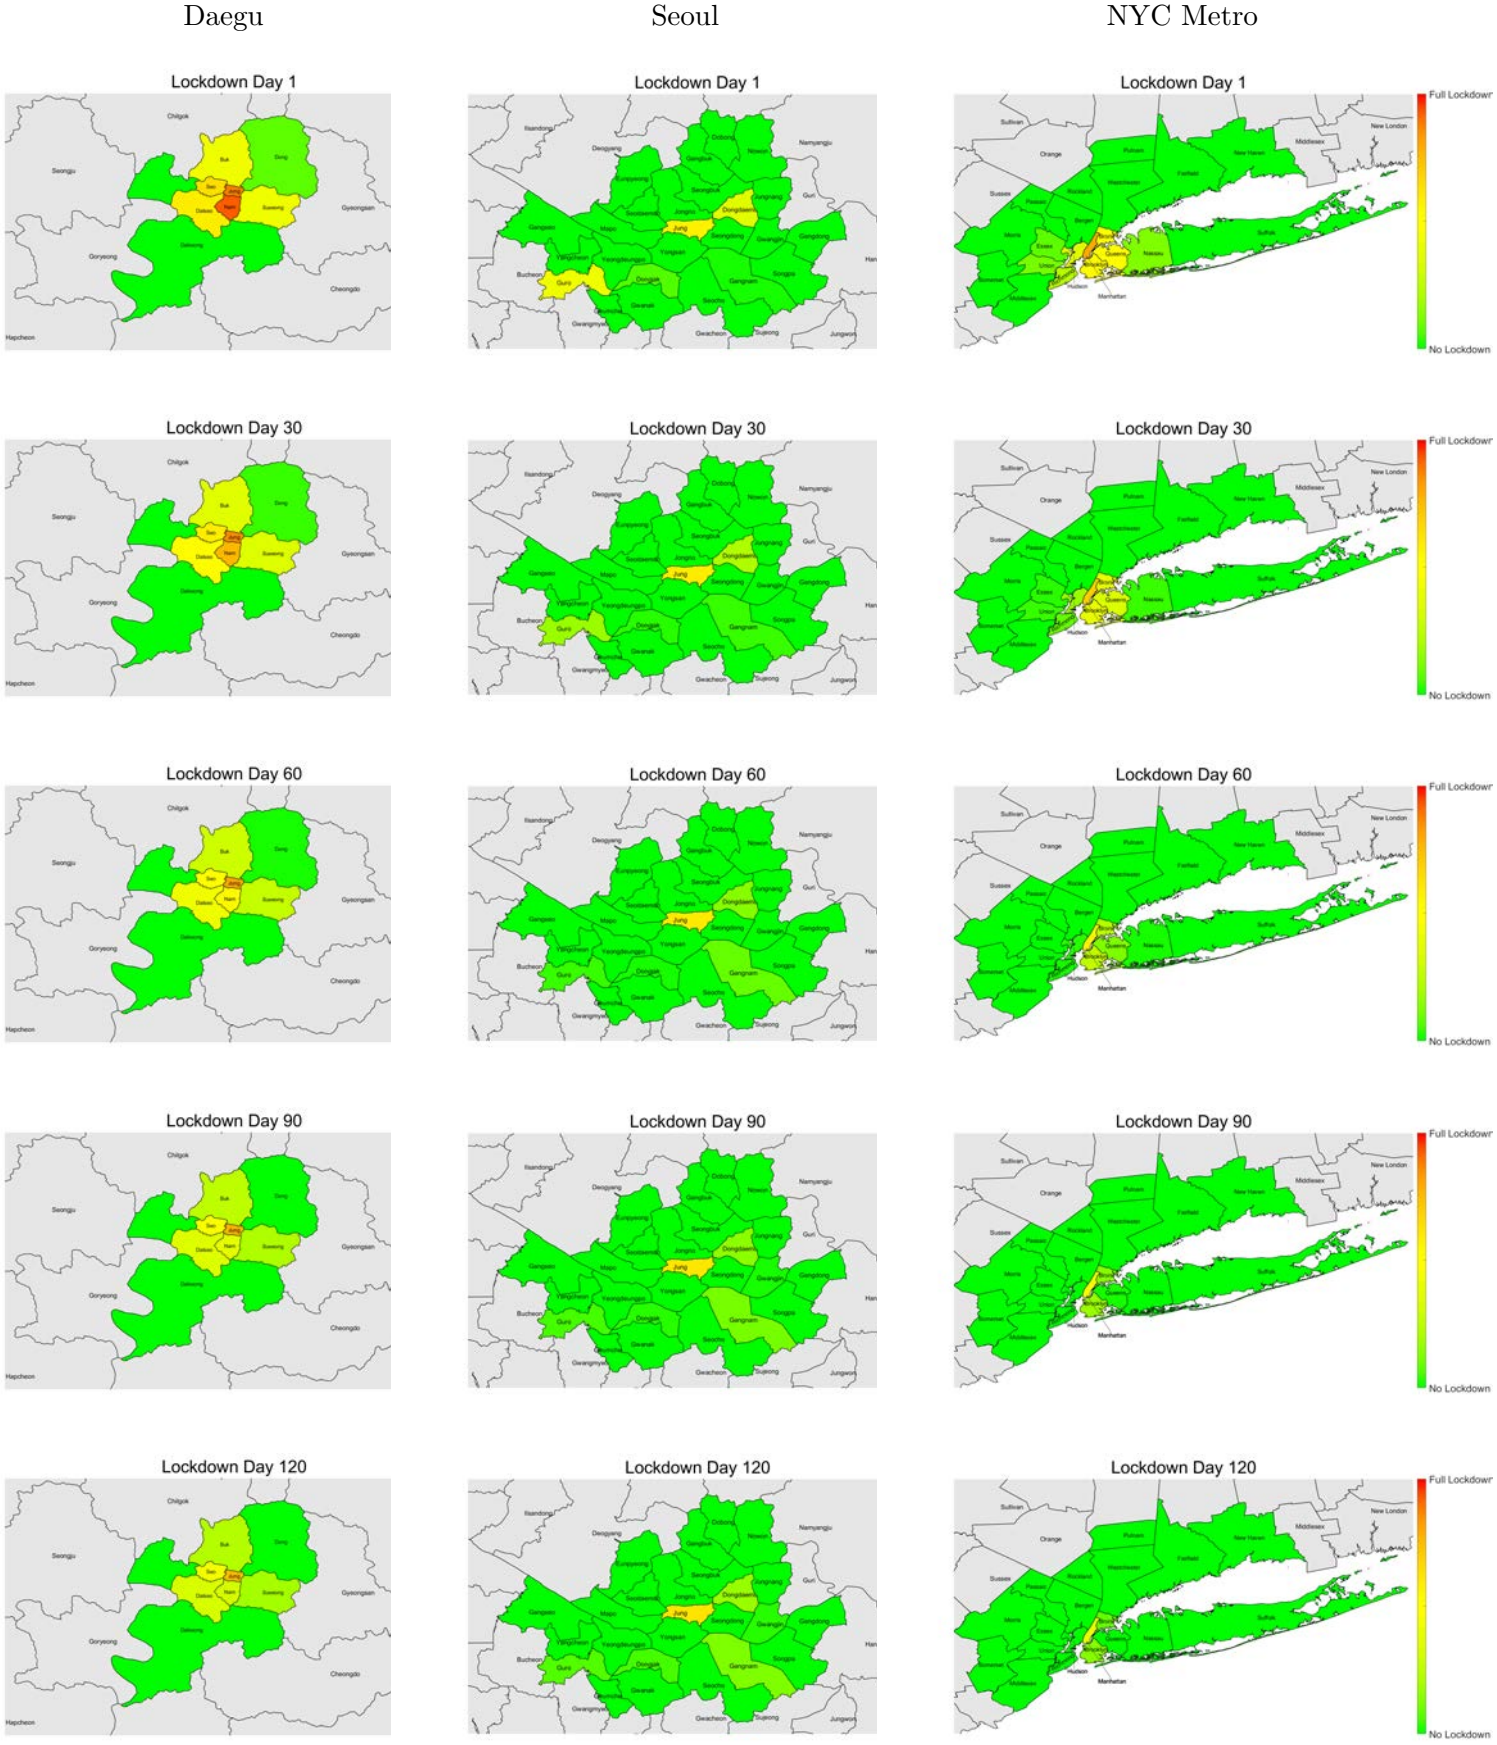

Note: The figure plots the optimal policy in the commuting area at different points in time. Redder colors denote more stringent lockdowns.

Figure 4: Pareto Frontiers

(a) Daegu

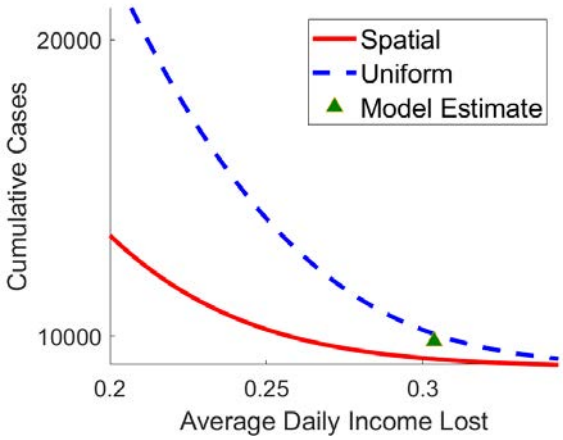

(b) Seoul

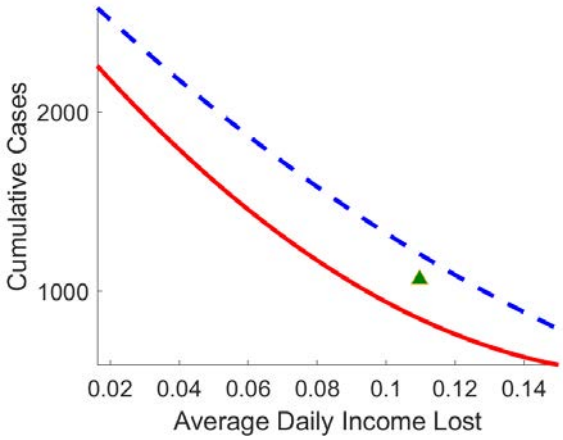

(c) NYC Metro

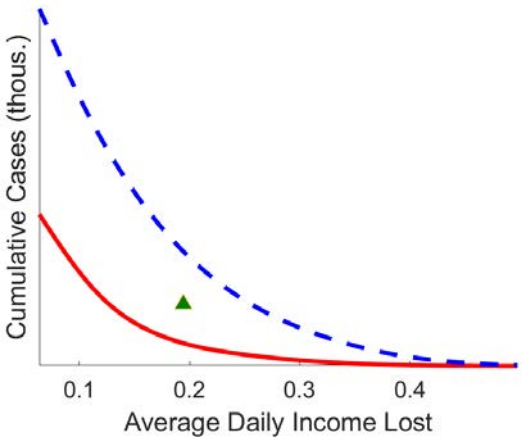

Note: The figures plot the cumulative number of new cases (y-axis, log scale) and the average real income lost per day between the date of the first confirmed case (see Appendix Table A.1) and April 30, 2020 for parametrizations of the value of life (ω) ranging from 1/100 to 100 times the benchmark, in both the optimal lockdown with space and time variation ("Spatial") and in the spatially uniform optimal solution with time variation only ("Uniform", i.e., the same lockdown across all locations). The green triangle shows the case count and real income lost implied by the estimated model on April 30, 2020.

Figure 5: Changes in Commuting Flows: Optimal and Observed

(a) Daegu

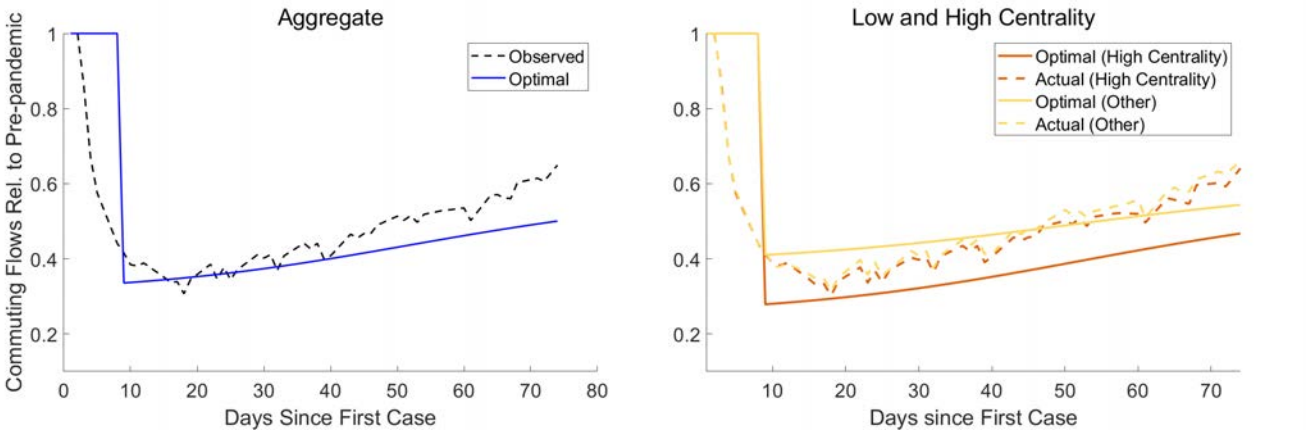

(b) Seoul

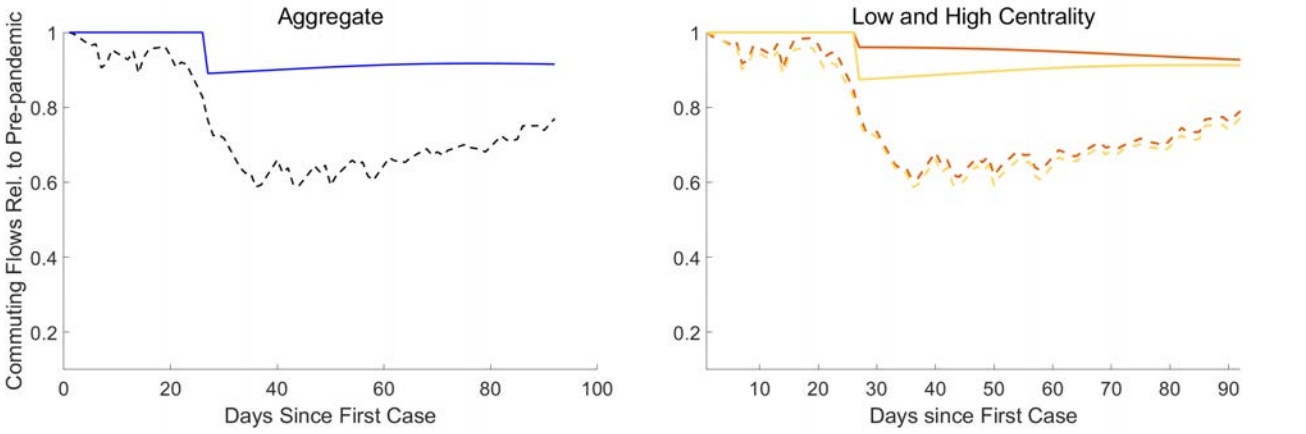

(c) NYC Metro

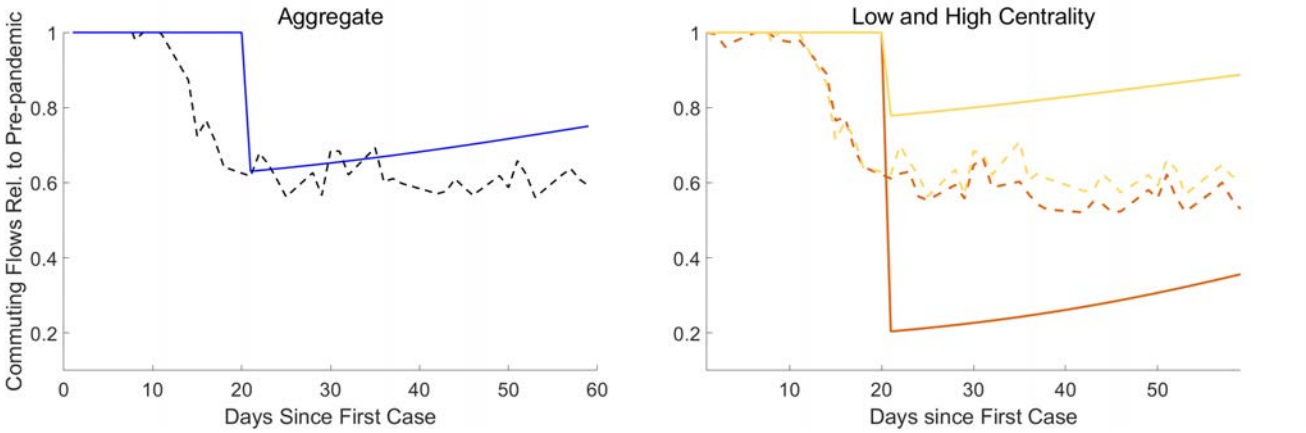

Note: In the left panel, the dashed black line shows the aggregate commuting flows in each city starting from the date of the first confirmed case in each city. The solid and circled blue lines show the aggregate commuting flows implied by the optimal spatial policy. In the right panel, optimal and observed commuting responses are divided by top-3 centrality locations (darker shade) and the other locations.

Online Appendix

A Numerical Implementation

This section describes how the optimal planning problem is numerically implemented.

A.1 Optimal control problem

Assuming that the matching function is $M_j\left(\tilde{I},\tilde{S}\right)=\beta_j\tilde{I}\tilde{S}$, the optimal control problem (10) simplifies to

$$W=\max_{\boldsymbol{\chi}(t)}\int_0^\infty e^{-(r+\nu)t}\sum_j\left[U\left(j,t\right)+\frac{\nu}{r}\overline{U}\left(j,t\right)-\omega\gamma_D I\left(j,t\right)\right]dt$$

subject to

$$\dot{\mathbf{S}}\left(t\right)=-\mathbf{S}\left(t\right)\cdot\left[\mathbf{H}_S\left(t\right)'\operatorname{diag}\left(\boldsymbol{\beta}\right)\mathbf{H}_I\left(t\right)\mathbf{I}\left(t\right)\right] \tag{A.1}$$

$$\dot{\mathbf{E}}\left(t\right)=-\dot{\mathbf{S}}\left(t\right)-\gamma_I\mathbf{E}\left(t\right) \tag{A.2}$$

$$\dot{\mathbf{I}}\left(t\right)=\gamma_I\mathbf{E}\left(t\right)-\left(\gamma_R+\gamma_D\right)\mathbf{I}\left(t\right) \tag{A.3}$$

$$\dot{\mathbf{R}}\left(t\right)=\gamma_R\mathbf{I}\left(t\right) \tag{A.4}$$

and

$$\begin{aligned}U\left(j,t\right)&=U\left(j;\mathbf{S}\left(t\right),\mathbf{E}\left(t\right),\mathbf{I}\left(t\right),\mathbf{R}\left(t\right),\boldsymbol{\chi}\left(t\right)\right)\\ \overline{U}\left(j,t\right)&=U\left(j;0,0,0,\mathbf{S}\left(t\right)+\mathbf{E}\left(t\right)+\mathbf{I}\left(t\right)+\mathbf{R}\left(t\right)\right).\end{aligned}$$

The present-value hamiltonian can be written

$$\begin{aligned}H\left(t\right)=\sum_j\Bigg[&U\left(j,t\right)+\frac{\nu}{r}\overline{U}\left(j;t\right)-\omega_D\gamma_D I\left(j,t\right)\\&+\mu_S\left(j,t\right)\dot{\mathbf{S}}\left(j,t\right)+\mu_E\left(j,t\right)\dot{\mathbf{E}}\left(j,t\right)+\mu_I\left(j,t\right)\dot{\mathbf{I}}\left(j,t\right)+\mu_R\left(j,t\right)\gamma_R\mathbf{I}\left(j,t\right)\Bigg],\end{aligned}$$

where $\boldsymbol{\mu}_u$, $u=S,E,I,R$, are $J\times1$ vectors of costate variables associated to each sickness status.

The first-order conditions of the problem are

$$\begin{aligned}
[\mathbf{S}\left(t\right)]\quad&\left(D_{\mathbf{S}}\mathbf{U}\left(t\right)+\frac{\nu}{r}D_{\mathbf{R}}\overline{\mathbf{U}}\left(t\right)\right)'\mathbf{1}_{J\times1}+\operatorname{diag}\left(\mathbf{H}_S\left(t\right)'\operatorname{diag}(\boldsymbol{\beta})\mathbf{H}_I\left(t\right)\mathbf{I}\left(t\right)\right)\left(\boldsymbol{\mu}_E\left(t\right)-\boldsymbol{\mu}_S\left(t\right)\right)\\
&\quad=-\dot{\boldsymbol{\mu}}_S\left(t\right)+\left(r+\nu\right)\boldsymbol{\mu}_S\left(t\right) &\text{(A.5)}\\
[\mathbf{E}\left(t\right)]\quad&\left(D_{\mathbf{E}}\mathbf{U}\left(t\right)+\frac{\nu}{r}D_{\mathbf{R}}\overline{\mathbf{U}}\left(t\right)\right)'\mathbf{1}_{J\times1}+\gamma_I\left(\boldsymbol{\mu}_I\left(t\right)-\boldsymbol{\mu}_E\left(t\right)\right)=-\dot{\boldsymbol{\mu}}_E\left(t\right)+\left(r+\nu\right)\boldsymbol{\mu}_E\left(t\right) &\text{(A.6)}\\
[\mathbf{I}\left(t\right)]\quad&\left(D_{\mathbf{I}}\mathbf{U}\left(t\right)+\frac{\nu}{r}D_{\mathbf{R}}\overline{\mathbf{U}}\left(t\right)\right)'\mathbf{1}_{J\times1}-\omega_D\gamma_D\mathbf{1}_{J\times1}\\
&+\mathbf{H}_I\left(t\right)'\operatorname{diag}\left(\boldsymbol{\beta}\right)\mathbf{H}_S\left(t\right)\operatorname{diag}\left(\mathbf{S}\left(t\right)\right)\left(\boldsymbol{\mu}_E\left(t\right)-\boldsymbol{\mu}_S\left(t\right)\right)\\
&\quad-\left(\gamma_R+\gamma_D\right)\boldsymbol{\mu}_I\left(t\right)+\gamma_R\boldsymbol{\mu}_R\left(t\right)=-\dot{\boldsymbol{\mu}}_I\left(t\right)+\left(r+\nu\right)\boldsymbol{\mu}_I\left(t\right) &\text{(A.7)}\\
[\mathbf{R}\left(t\right)]\quad&\left(D_{\mathbf{R}}\mathbf{U}\left(t\right)+\frac{\nu}{r}D_{\mathbf{R}}\overline{\mathbf{U}}\left(t\right)\right)'\mathbf{1}_{J\times1}=-\dot{\boldsymbol{\mu}}_R\left(t\right)+\left(r+\nu\right)\boldsymbol{\mu}_R\left(t\right) &\text{(A.8)}\\
[\boldsymbol{\chi}\left(t\right)]\quad&D_{\boldsymbol{\chi}}H\left(t\right)=0 &\text{(A.9)}
\end{aligned}$$

A.2 Algorithm

We solve the optimal control problem using the following steps. Set the terminal period T to be a large number. Given some initial condition $\left\{\mathbf{S}^{(n)}(0), \mathbf{E}^{(n)}(0), \mathbf{I}^{(n)}(0), \mathbf{R}^{(n)}(0), \mathbf{D}^{(n)}(0)\right\}$,

1. Initialize $n := 1$. Guess the policy $\boldsymbol{\chi}^{(1)}(t)$ for $t = 0 \dots T$ at the first iteration.
2. Using $\boldsymbol{\chi}^{(n)}(t)$, solve the partial differential equations (A.1)-(A.4) forward using the Euler method to recover $\left\{\mathbf{S}^{(n)}(t), \mathbf{E}^{(n)}(t), \mathbf{I}^{(n)}(t), \mathbf{R}^{(n)}(t), \mathbf{D}^{(n)}(t)\right\}$ for $t = 1 \dots T$. Solve for the economic allocation and the corresponding Jacobian in each t as described in the subsection A.3.
3. Using $\boldsymbol{\chi}^{(n)}(t)$ and the disease states, solve the partial differential equations (A.5)-(A.8) of the costates $\{\boldsymbol{\mu}_S(t), \boldsymbol{\mu}_E(t), \boldsymbol{\mu}_I(t), \boldsymbol{\mu}_R(t)\}_{t=0}^{T-1}$ backward using the Euler method with terminal condition
$$\{\boldsymbol{\mu}_S(T), \boldsymbol{\mu}_E(T), \boldsymbol{\mu}_I(T), \boldsymbol{\mu}_R(T)\} = 0.$$
4. Compute $\boldsymbol{\chi}^*(t) = \text{argmax}_{\boldsymbol{\chi}} H^{(n)}(t; \boldsymbol{\chi})$ using a numerical optimizer. This step uses the analytical gradient for the trade model described in the next section.
5. Stop if $\left\|\boldsymbol{\chi}^{(n)} - \boldsymbol{\chi}^*\right\| < \varepsilon$. Otherwise, set $\boldsymbol{\chi}^{(n+1)} = \lambda\boldsymbol{\chi}^* + (1-\lambda)\boldsymbol{\chi}^{(n)}$ where $0 < \lambda < 1$, set $n := n + 1$ and return to step (2).

Computing the maximizer $\boldsymbol{\chi}^*(t)$ for all t is the most computationally expensive step. Computation times can be improved by sampling a smaller number of dates t_1,t_2...t_N and interpolating the policy between those dates.

A.3 Solving the General Equilibrium Trade Model

We use two methods to compute the solution of the general equilibrium trade model at different stages of the numerical optimization:

Exact Solution

Wen solving the SEIR model forward, we compute the exact general equilibrium solution of the trade model by iterating over $w(j,t)$ for each t on the goods market equilibrium equation given a distribution of the state variables. Specifically, combining (4) and (7) we obtain:

$$w(j,t) = \frac{\sum_k Y(k,t)\left(\frac{\tau(j,k)}{P(k,t)}\frac{w(j,t)}{z(j)}\right)^{1-\sigma}}{\sum_{u=S,E,I,R}\sum_k N_u(k,j,t)}. \tag{A.10}$$

Further using (5) and (8) we obtain a system for wages at time t of the form

$$w(j,t) = H_j(w(1,t), .., w(J,t), t) \tag{A.11}$$

where the operator $H_j(w_1, .., w_J)$ takes the form

$$H_j(w_1, .., w_J) = \left(\frac{1}{\sum_{u=S,E,I,R}\sum_{k'} N_u(k',j,t)}\sum_k \frac{\left(\frac{\tau(j,k)}{z(j)}\right)^{1-\sigma}}{\sum_i\left(\frac{\tau(i,k)}{z(i)}w_i\right)^{1-\sigma}}\sum_{u=S,E,I,R}\sum_{i'} N_u(k,i',t)\, w_k\right)^{\frac{1}{\sigma}} \tag{A.12}$$

where $N_u(i,j,t)$ is given by (6).

Gradients

When evaluating the Jacobians $D_u\mathbf{U}(t)$ and $D_u\overline{\mathbf{U}}(t)$ for $u = S, E, I, R$ or when maximizing the Hamiltonian, we linearize the trade model around a nonlinear solution (the equilibrium under the current $\boldsymbol{\chi}$ to evaluate the Jacobian

and the current equilibrium with $\boldsymbol{\chi} = 1$ for the Hamiltonian maximization).

Solving for the linearized equilibrium solely requires inverting a matrix for which we have an analytical expression. Specifically, totally differentiating the equilibrium conditions given shocks to the bilateral flows $d \ln N_u(j,i)$ of type-u workers, and dropping the time subscript to save notation, we obtain the following linear system

$$\begin{aligned}
d \ln Y &= \sum_u \left(s_R(u) \cdot d \ln N_u \right) 1_{J1} + \sum_u s_R(u) \, d \ln w \\
d \ln w &= s_X d \ln Y - \left(s_X \cdot d \ln s_M \right) 1_{J1} - \sum_u \left(s_W(u) \cdot d \ln N_u \right)' 1_{J1} \\
d \ln P &= s_M' d \ln w \\
d \ln s &= (1-\sigma) \left(d \ln w \right) 1_{1J} - 1_{J1} \left(d \ln P \right)'
\end{aligned}$$

where the first line is the total differential of (5), the second line corresponds to (7), the third line differentiates the price index (8) and the last line is the changes in the expenditure share, and where we are using vector notation such that $[d \ln N_u]_{i,j} = N_u(i,j)$, $[d \ln Y]_j = d \ln Y(j)$, $[d \ln P]_j = d \ln P(j)$, $[d \ln w]_j = d \ln w(j)$, and $[s]_{ij} = s(i,j)$. In these expressions, we have also defined $[s_X]_{ij} = s_X(i,j)$, $[s_W(u)]_{ij} = s_W(u,i,j)$ and $[s_R(u)]_{ij} = s_R(u,i,j)$ such that $s_X(i,j) \equiv \frac{Y(j) s(i,j)}{w(i) \sum_{u=S,E,I,R} \sum_j N_u(j,i)}$ is location j share in i's sales, $s_W(u,i,j) = \frac{N_u(i,j)}{\sum_{u=S,E,I,R} \sum_k N_u(k,j)}$ is the fraction of j's efficiency units corresponding to type u commuters from i, and $s_R(u,i,j) = \frac{N_u(j,i) w(i)}{Y_R(j)}$ is the fraction of resident of j's income corresponding to commuters to i in type u.

We can summarize the expressions as a solution for wages as a function of expenditure shares and labor flows shares:

$$d \ln w = \Omega_w^{-1} \left[\sum_u s_X \left(s_R(u) \cdot d \ln N_u \right) 1_{J1} - \sum_u \left(s_W(u) \cdot d \ln N_u \right)' 1_{J1} \right] \tag{A.13}$$

where

$$\Omega_w = I_J - \sum_u s_X s_R(u) - (1-\sigma) \operatorname{diag}(s_X 1_{J1}) + (1-\sigma) s_X s'. \tag{A.14}$$

The gradients with respect to χ, S, E, I and R then follow using the definition of $N_u(i,j)$ in (6).

B Commuting and Infections Relationship

We follow Fang et al. (2020) to assess the relationship between new daily infections and observed commuting in a reduced form specification. We correlate new daily infections in district i on date t in response to 21-day lags of commuting:

$$\ln(1 + \text{new cases}_{it}) = \alpha_i + \gamma_{city(i),t} + \sum_{k=0}^{21} \beta_k \ln(flow_{i,t-k}) + \epsilon_{it} \tag{A.15}$$

where α_i is a district fixed effect and $\gamma_{city,t}$ is a fixed effect that varies by date and city. This specification flexibly controls for city-level trends due to forces other than commuting. It identifies the impact of flows from cross-sectional variation by exploiting a district's flows above or below its average. The variable $flow_{it}$ is either the total number of people who leave from district i (outflows) or arrive into district i (inflows); we try both. The specification pools over the three cities since number of districts in each city is small (20 for NYM, 25 for Seoul, and 8 for Daegu), and we weight the regression so that each city contributes equally. Standard errors are clustered by i using the block bootstrap to account for a small number of clusters (Cameron et al., 2008).

We do not have an instrument for commuting that varies across space, so the coefficients must

be interpreted with caution. For example, changes in commuting may be correlated with behavioral changes. Figure A.4 reports the coefficients. It shows an inverted-U shape peaking between 8-15 days when looking at inflows. The p-value of the joint test $\beta_k = 0$ for $k = \{0, ..., 7\}$ is 0.479, for $k = \{8, ...15\}$ is 0.045, and for $k = \{16, ..., 21\}$ is 0.452. The figure for outflows is noisier. This finding seems consistent with an incubation period after which people showing symptoms get tested or come to the hospital. These moments are not targeted by the structural model.

C Additional Tables and Figures

Table A.1: Commuter Data Summary Statistics

	Daegu	Seoul	NYC Metro
Population	2,438,031	9,729,107	19,467,622
# Districts	8	25	20
Sample Period	Jan 1, 2018–Apr 30, 2020	Jan 1, 2018–Apr 30, 2020	Jan 1, 2020–Apr 30, 2020
Data Source	Subway ridership	Subway/bus ridership	Mobile phones
Flow Type	Turnstile	Bilateral	Bilateral
First Case	Feb 17, 2020	Jan 30, 2020	Mar 3, 2020
Lockdown Date	Feb 24, 2020	Feb 24, 2020	Mar 22, 2020
# Cumulative Cases	6,778	354	389,603

Notes: Table reports summary statistics for the Daegu, Seoul, and NYC Metro data. Administrative units within the two Korean cities are called districts with an average population of 368,701 and an average land area of 45 km^2. Administrative units within NYC Metro are counties with an average population of 1,232,768 and an average land area of 690 km^2. Cumulative Covid-19 cases are as of April 30, 2020.

Table A.2: Summary of Parameter Values

Parameter	Definition	Value	Source
Disease Dynamics			
γ_I	Exposed to Infected Rate	{1/5.1, 1/4.2}	Ferguson et al. (2020), Sanche et al. (2020)
γ_R	Infected to Recovered Rate	{1/18, 1/10}	Wang et al. (2020)
γ_D	Infected to Death Rate	{0.0005, 0.0002} (see Table note)	Ferguson et al. (2020), Hall et al. (2020)
ζ_I	% asymptomatic	{0.545, 0.272}	Alamian et al. (2019)
Matching Function			
β	Transmission Rate	Daegu: 0.58 Seoul: 1.58 NYM: 0.16	Case Data and Commuting (see Section 3.3)
Trade Model			
κ_1	Distance-Trade Cost Elasticity	0.37	
κ_0	Scale of Trade Costs	Daegu: 0.69 Seoul: 1.23 NYM: 0.62	Credit Card Expenditures (see Section 3.3)
σ	Demand Elasticity	5	Ramondo et al. (2016)
Other Parameters			
δ_I	Telecommuting Rate	Korea: 0.62 NYM: 0.46	http://www.jobkorea.co.kr/GoodJob/Tip/View?News_No=16696 Dingel and Neiman (2020)
v	Probability of Vaccine	1/(365*1.5)	Expected time of 1.5 years until vaccine
ω	Value of Life	{1/100,..,100}*10 Million USD	
ρ	Discount rate	0.04/365	

Note: The table reports the parameters in the benchmark calibration and robustness exercises. The rate γ_D is chosen such that the infection fatality ratio 0.009 equals $\frac{\gamma_D}{\gamma_D+\gamma_R}$, and the robustness uses 0.003 instead. See Section 3.3 for more details.

Figure A.1: New Cases per District: Data and Estimation

(a) Daegu (b) Seoul (c) NYC Metro

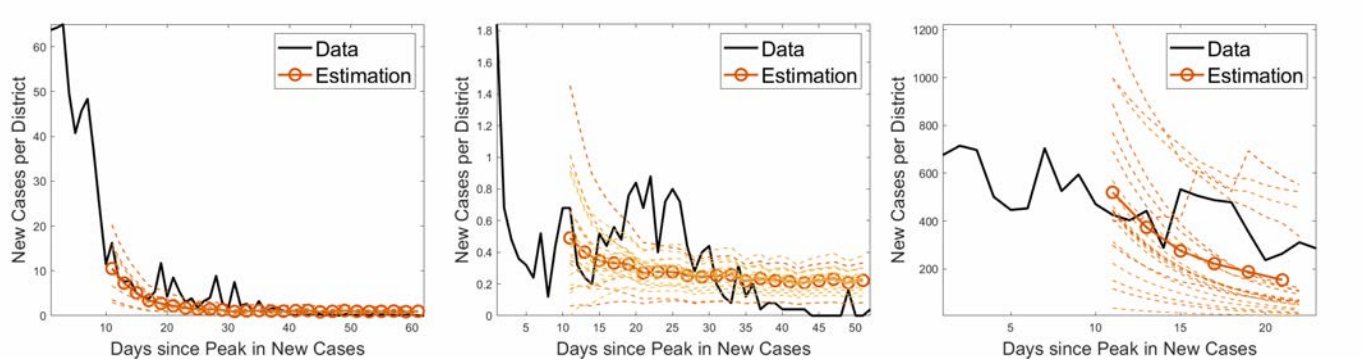

Note: The solid lines show the average number of new cases per district over time since the peak in new cases in the data. The dashed lines are the number of new cases by district in the estimated model corresponding to equation (17), assuming that commuting changed as observed in the data (the shade of the lines represent the share of commuter inflows, darker representing more inflows). The calibration is implemented using case data starting 10 days after the peak in new cases in each city. The solid line with circle markers is the total case number per district in the estimated model.

Figure A.2: Seoul: Optimal Lockdown in Baseline and Alternative Scenarios

(a) Baseline (b) Large Shock (c) High Value of Life

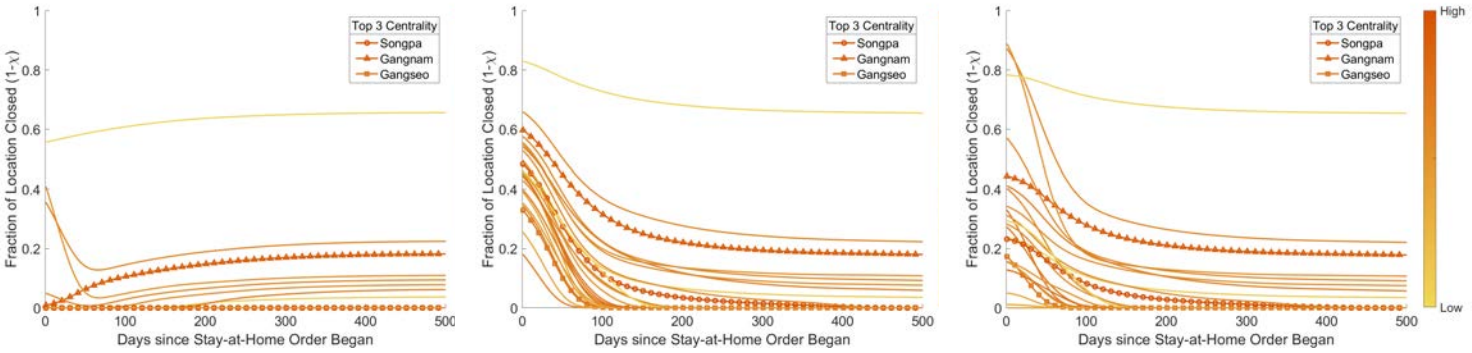

Note: The three panels show the results for Seoul under the baseline calibration (left panel), a large shock infecting 1% of the population (middle panel) and a value of life that is 100 times the benchmark (right panel).

Figure A.3: Robustness

(a) Daegu

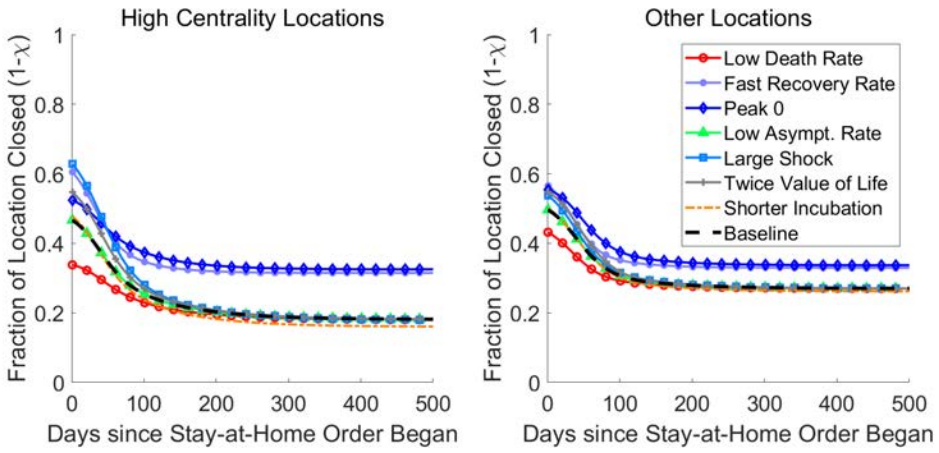

(b) Seoul

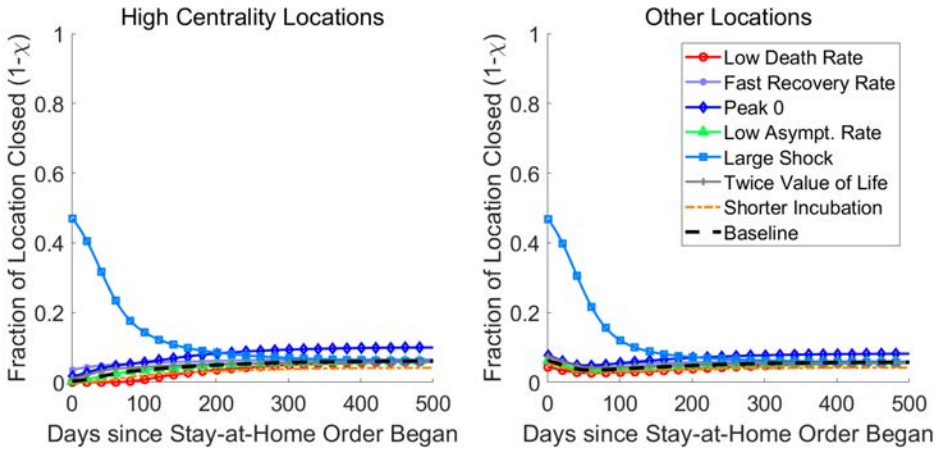

(c) NYC Metro

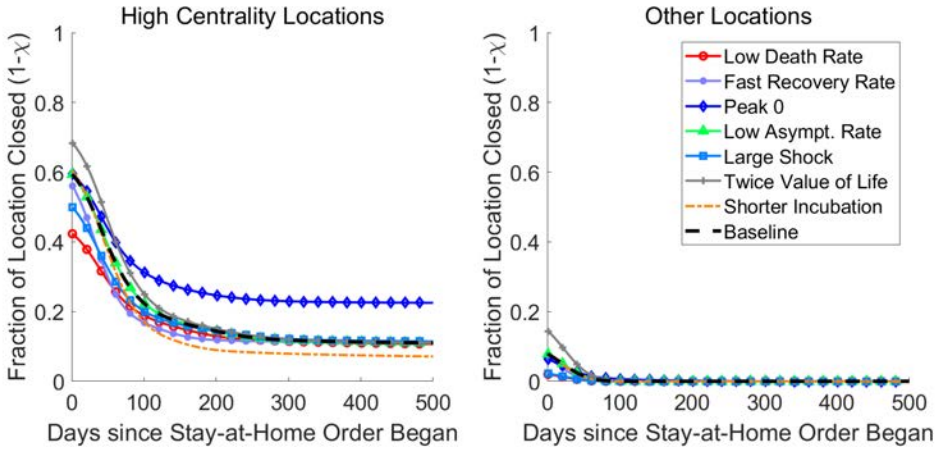

Note: Plotted optimal policies are defined as mean policy for high centrality vs other locations for each city. The different cases correspond to the alternative parametrizations described in Section 3.3 and discussed in Section 4.4.

Figure A.4: Commuting and New Daily Cases

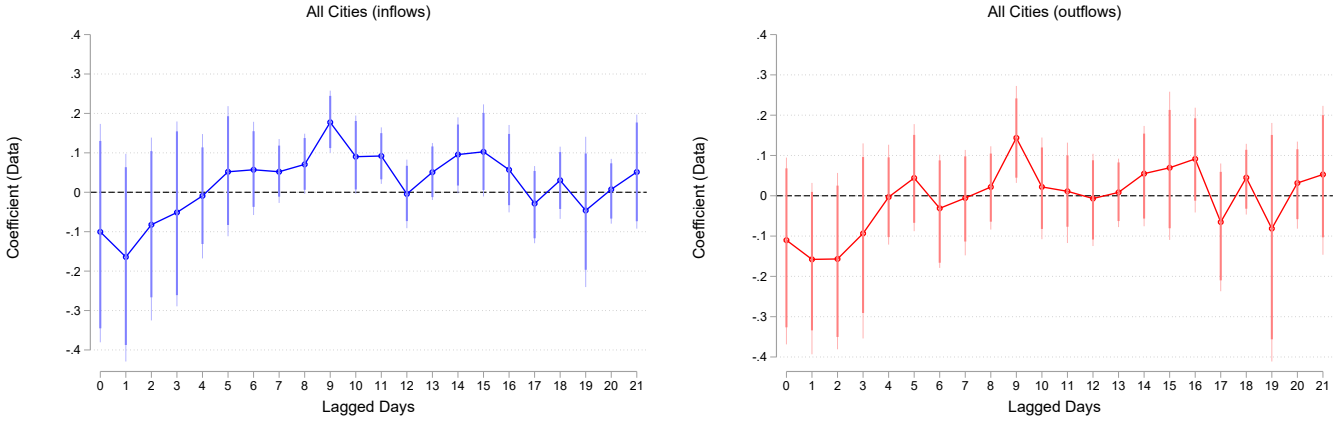

Note: The figure plots the coeffients from A.15. The left panel reports results using inflows as the independent variable. The right panel reports outflows as the independent variable. The regression pools over the three cities and applies weights so that each city contributes equally. The regression is run using data since January 22, 2020. Error bars show 90 percent (thick) and 95 percent (thin) confidence intervals. Standard errors are clustered by using the block bootstrap to account for a small number of clusters.