

The (Structural) Gravity of Epidemics*

Alejandro Cuñat[†] Robert Zymek[‡]

May 2020

Abstract

Epidemiological models assume gravity-like interactions of individuals across space without microfoundations. We combine a simple epidemiological framework with a dynamic model of individual location choice. The model predicts that flows of people across space obey a structural gravity equation. By means of an application to data from Great Britain we show that our structural-gravity framework: provides a rationale for quarantines; offers a clear mapping from observed geography to the spread of a disease; and makes it possible to evaluate the welfare impact of (expected and unexpected) mobility restrictions in the face of a deadly epidemic.

JEL Classification codes: E65, F14, F17, J61, R23

Keywords: epidemics, Covid-19, gravity, regional mobility

*This work contains statistical data from the 2011 UK Census which is Crown Copyright. We are grateful to Philipp Kircher, Clara Santamaría, Yoto Yotov and webinar participants at the University of Edinburgh, the MacCaLM “Economics of Covid-19” workshop and the NBER Summer Institute for helpful comments and suggestions. Zymek gratefully acknowledges financial support from the UK Economic and Social Research Council (ESRC) under award ES/L009633/1.

[†]Department of Economics, University of Vienna and CESifo, Oskar Morgenstern Platz 1, Vienna 1090, Austria; alejandro.cunat@univie.ac.at.

[‡]School of Economics, University of Edinburgh and CESifo, 31 Buccleuch Place, Edinburgh, EH8 9JT, United Kingdom; robert.zymek@ed.ac.uk.

1 Introduction

The Covid-19 epidemic has thrust epidemiological models into the limelight. The “Imperial College study” (Ferguson et al., 2020) received widespread public attention, and is credited with having changed the UK Government’s stance on slowing the spread of the disease. Economists have started to integrate macroeconomic and epidemiological models in order to analyse jointly the economic and public-health impacts of different government interventions.¹

Disease transmission models such as the one used in Ferguson et al. (2020) assume that people interact across space in inverse proportion to (relative) distance.² The epidemiological literature explicitly refers to this as a “gravity” assumption.³ However, the functional forms assumed do not have a choice-theoretic microfoundation and are calibrated from available mobility data in an *ad hoc* manner. This precludes a formal welfare analysis of prominent interventions, such as mobility restrictions, on the basis of these models. Meanwhile, the first economic models of the epidemic have combined fully microfounded models of the macroeconomy with epidemiological frameworks by introducing somewhat arbitrary assumptions about how the disease transmission is affected by economic activity.⁴ This introduces new macro parameters that are difficult to calibrate with any degree of confidence.

In the present paper, we show that economists already possess a toolkit for improving on both approaches: structural-gravity modelling. Structural gravity models are now common in international trade, where they are used to study the observed pattern of economic interactions across space and to assess the impact of trade-policy changes. They have provided simple microfoundations to explain why certain types of data – such as trade, migration or commuting flows – exhibit “gravity” patterns. There exist well-understood empirical approaches for estimating the impact of geography on interactions consistently with these models. Moreover, such models share convenient properties that make it easy to analyse the welfare impact of barriers that restrict interactions across space.

By way of illustration, we combine a simple epidemiological framework – the SIR model (Kermack and McKendrick, 1927) – with a basic dynamic model of individual location choice. The model makes assumptions that ensure that flows of people across space obey a structural gravity equation. To demonstrate the uses of our structural-gravity SIR framework, we calibrate it to match regional mobility patterns from

¹See Atkeson (2020), Beenstock and Dai (2020), Eichenbaum et al. (2020), González-Eiras and Niepelt (2020).

²See the supplementary information of Ferguson et al. (2005) for a full description of the model, and Ferguson et al. (2006) for a discussion of its calibration to UK data.

³A detailed discussion can be found in Xia et al. (2004).

⁴For example, both Eichenbaum et al. (2020) and González-Eiras and Niepelt (2020) simply introduce the macro-level assumption that the infection rate of the disease is a positive function of economic activity.

British census data. We then use it to simulate the course of an epidemic, inspired by the properties of Covid-19, under different regional quarantine scenarios.

In the model, temporary mobility restrictions reduce welfare but also slow the spread of a disease. As a result, the model captures the welfare trade-off inherent in the imposition of quarantines in a microfounded fashion. Moreover, it highlights key parameters that govern this trade-off and could be estimated from micro data. A suggestive welfare analysis shows that quarantines are welfare-enhancing for reasonable parameter values. It also indicates that such quarantines may be more effective if imposed early, and if they are not anticipated by the public.

Our work borrows a number of insights from the international-trade literature. Anderson (1979) and Anderson and van Wincoop (2003) pioneered the use of structural-gravity models in international trade. We estimate our model using the Poisson Pseudo-Maximum Likelihood (PPML) estimator which was introduced by Santos-Silva and Tenreyro (2006) and ensures a straightforward, theory-consistent estimation of structural gravity models (Fally, 2015). Our structural mobility gravity equation is microfounded using the same choice-theoretic assumptions that underpin the trade gravity equation derived by Eaton and Kortum (2002). As a result, it shares the common welfare properties of this class of models, first pointed out by Arkolakis et al. (2012).

We are not the first to apply structural-gravity modelling in the context of regional mobility. McFadden’s (1974) classic study of urban travel demand exploits assumptions that are closely related to the microfoundation of gravity by Eaton and Kortum (2002). Anderson (2011) shows that structural-gravity models can be used in the context of migration flows. Several studies have modelled regional mobility and commuting in steady state using microfounded gravity.⁵ Our framework is more closely related to a small number of recent papers that employ dynamic structural gravity models to analyse the response of regional mobility to “macro” shocks at high frequencies.⁶ We show here that this approach is easy to integrate with existing epidemiological models and that doing so has the potential to deliver new insights about the impact of mobility restrictions in the face of an epidemic.

2 Model

2.1 Pre-infection Economy

We begin with a description of the model economy before the arrival of an epidemic.

⁵See Ahlfeld et al. (2015), Allen et al. (2016), Bryan and Morten (2018) and Monte et al. (2018).

⁶See Artuç et al. (2015), Caliendo et al. (2019) and Artuç et al. (2020).

2.1.1 Assumptions

There are $n = 1, \dots, N$ locations. Let L_{nt} denote the mass of people who spend t in location n . For simplicity, we will refer to it as the population of n in t . The total population $L = \sum_n L_{nt}$ is fixed, and there is no aggregate uncertainty. We will think of a period t as representing one day.

The decision problem of person in location n at the start of day t can be represented by the Bellman equation

$$V_t(n) = \max_{n' \in N} \left\{ \ln \left[\frac{u_{n'}}{\delta_{nn't}} z_{n't}(n) \right] + \beta E_t [V_{t+1}(n')] \right\}, \quad (1)$$

where $\beta \in (0, 1)$ denotes the discount factor; $u_{n'} > 0$ is a location-specific, constant flow of utility; $\delta_{nn't} \geq 1$ represents the cost of moving from location n to location n' , with $\delta_{nnn} = 1$ for all n ; and $z_{n't}(n)$ is a preference shock realised prior to an individual's location choice each t .

The location parameter u_n is a shortcut to the inherent characteristics that make a place attractive, such as its local labour market or the quality of local amenities.⁷ In our analysis, the moving cost $\delta_{nn't}$ will reflect bilateral travel costs that vary across places as a result of geography, but are normally constant across time. However, we will consider scenarios in which this moving cost becomes temporarily prohibitive as a result of expected or unexpected government interventions (“quarantines”).

The preference shock $z_{n't}(n)$ captures idiosyncratic reasons why an individual may want to move from n to n' on any given day. It is drawn from the Fréchet distribution

$$F_{nn'}(z) = e^{-\omega_{nn'} z^{-\theta}}, \quad (2)$$

where $\omega_{nn'} > 0$, $\omega_{nn} = 1$ for all n , and $\theta > 0$.

2.1.2 Bilateral Flows

The share of people in location n at the start of t who will find it optimal to move to n' is

$$m_{nn't} = \frac{(\tau_{nn't}/v_{n't})^{-\theta}}{\sum_{n'=1}^N (\tau_{nn't}/v_{n't})^{-\theta}}, \quad (3)$$

where $\tau_{nn't} \equiv \omega_{nn'}^{-\frac{1}{\theta}} \delta_{nn't}$ aggregates preferences and travel costs into an overall bilateral mobility barrier, with $\tau_{nnn} = 1$ for all n ;

$$v_{nt} \equiv u_n \left[e^\gamma \sum_{n'=1}^N (\tau_{nn't+1}/v_{n't+1})^{-\theta} \right]^{\frac{\beta}{\theta}}, \quad (4)$$

⁷For simplicity, we assume throughout that u_n is constant at short time horizons.

and γ is the Euler-Mascheroni constant. A proof is provided in the [online Appendix Section A.1.1](#).

The share $m_{nn't}$ depends negatively on the mobility barrier between n and n' (relative to all bilateral barriers), and positively on the “place value” of n' , $v_{n't}$ (relative to all place values). In turn, the place value of any n comprises the fundamental flow utility offered by n , u_n , as well as an index of connectivity, $\left[\sum_{n'} (\tau_{nn't+1}/v_{n't+1})^{-\theta}\right]^{1/\theta}$, reflecting the attractiveness of the locations to which n offers access going forward.

The functional form of $m_{nn't}$ implies that we can write the flow of people between locations n and n' on day t as

$$m_{nn't}L_{nt-1} = \left(\frac{\tau_{nn't}}{P_{n't}O_{nt}}\right)^{-\theta} L_{n't}L_{nt-1}, \quad (5)$$

where

$$P_{n't} \equiv \left[\sum_{n=1}^N \left(\frac{\tau_{nn't}}{O_{nt}}\right)^{-\theta} L_{nt-1}\right]^{-\frac{1}{\theta}}, \quad O_{nt} \equiv \left[\sum_{n'=1}^N \left(\frac{\tau_{nn't}}{P_{n't}}\right)^{-\theta} L_{n't}\right]^{-\frac{1}{\theta}}, \quad (6)$$

and P_{nt} and O_{nt} are the so-called inward and outward multilateral resistance terms (MRTs) of a location n , respectively.⁸

Equation (5) highlights that the flow of people between n and n' can be expressed as proportional to the product of the origin and destination populations, and inversely proportional to bilateral mobility barriers relative to the MRTs. The shape parameter of the Fréchet distribution emerges as the elasticity of bilateral flows with respect to mobility barriers.

2.1.3 Welfare

We assume that mobility barriers pre-infection are (expected to be) constant: $\tau_{nn't} = \tau_{nn'}$ for all t . In this case, $v_{nt} = v_n$ and $m_{nn't} = m_{nn'}$ for all t .

Equation (4) can be re-written as

$$v_n = u_n^{\frac{1}{1-\beta}} \left[(e^{-\gamma} m_{nn})^{-\frac{1}{\theta}} \right]^{\frac{\beta}{1-\beta}}. \quad (7)$$

Hence, variations in the connectivity index across locations in the pre-infection economy can be captured empirically by variations in the share of people who stay in their origin locations each period: places that provide easy access to many attractive locations will be characterised by a smaller share of “stayers”.

Let V_t/L denote average welfare. In the [online Appendix Section A.1.2](#) we show

⁸See Head and Mayer (2014) for a proof.

that

$$\frac{V_t}{L} = \frac{1}{1-\beta} \sum_n \frac{L_{nt-1}}{L} \left(-\frac{1}{\theta} \ln m_{nn} + \ln u_n + \frac{\gamma}{\theta} \right). \quad (8)$$

Suppose a government were to announce a permanent quarantine unexpectedly on day t , setting $\tau_{nnt'} \rightarrow \infty$ for all $t' \geq t$. This would imply $m_{nnt'} = 1$ for all $t' \geq t$. Based on equation (8), the welfare impact of such a scenario could be evaluated using only information on the distribution of the population across locations $\{L_{nt-1}/L\}_n$, the share of “stayers” in each location, $\{m_{nn}\}_n$, and the mobility elasticity, θ . This mirrors well-established results from the trade literature.⁹

2.2 An Epidemic Outbreak

We now model the course of an epidemic that arrives as an “MIT shock” in the form of some initial infections across locations. We let \tilde{I}_{n0} denote the mass of newly infected individuals on the initial day 0, with $\tilde{I}_{n0} \geq 0$ and $\tilde{I}_{n0} > 0$ for at least one n . Infections carry no inherent disutility, but for the duration of an infection, individuals face a probability π_d of death. For parsimony, the only private and social cost of death is the forgone utility of life.¹⁰

The dynamics of the epidemic follow a discrete-time version of the SIR model of Kermack and McKendrick (1927): individuals in location n during day t that have not yet contracted the disease are susceptible (S_{nt}); individuals in n that are currently infected (I_{nt}) create new infections among the susceptible in their location; and some infected probabilistically join the recovered (R_{nt}), whereupon they can no longer contract the disease.¹¹

2.2.1 New Assumptions

As a result of the probability of death arising from the epidemic, the aggregate population is no longer constant. The population of each location is now made up of the susceptible, infected and recovered, such that

$$L_t = \sum_n L_{nt} = \sum_n (S_{nt} + I_{nt} + R_{nt}). \quad (9)$$

Each day, the new sequence of events is as follows:

1. All survivors in t find themselves in their location n .
2. Preference shocks $\{z_{n't}(n)\}$ are realised.
3. Agents choose in which n' to spend t .

⁹See Arkolakis et al. (2012), Costinot and Rodríguez-Clare (2014), and Ossa (2015).

¹⁰This assumption is discussed more thoroughly in Hall and Jones (2007).

¹¹See Allen (1994) for an in-depth treatment of a one-location, discrete-time SIR model.

4. A mass $\tilde{I}_{n't} = \pi_s S_{n't} I_{n't} / L_{n't}$ of the susceptible in n' become newly infected.
5. The non-newly infected recover with probability π_r and die with probability π_d .

2.2.2 Bilateral Flows Revisited

Let $m_{nn't}(S)$, $m_{nn't}(I)$ and $m_{nn't}(R)$ denote the bilateral movement propensities of the susceptible, infected and recovered in the post-outbreak economy. In the [online Appendix](#) Section A.2.2, we show that if $\pi_d \rightarrow 0$,

$$m_{nn't}(S) \rightarrow m_{nn't}(I) \rightarrow m_{nn't}(R) = m_{nn't}. \quad (10)$$

Therefore, for small π_d , the susceptible, infected and recovered in location n post-outbreak behave (approximately) like the average person in location n of the pre-infection economy. From now on, we will restrict our attention to epidemic outbreaks characterised by $\pi_d \approx 0$. This allows us to make inferences about the (approximate) behaviour of agents in the wake of an epidemic outbreak from pre-infection mobility data in a straightforward way.

2.2.3 The Geographic Spread of the Epidemic

Assuming $\pi_d \approx 0$, we obtain

$$S_{n't} = \sum_n m_{nn't} (S_{nt-1} - \tilde{I}_{nt-1}), \quad (11)$$

$$I_{n't} = \sum_n m_{nn't} \left[(1 - \pi_r) I_{nt-1} + \tilde{I}_{nt-1} \right], \quad \tilde{I}_{nt-1} = \pi_s \frac{I_{nt-1} S_{nt-1}}{L_{nt-1}}, \quad (12)$$

$$R_{n't} = \sum_{n'} m_{nn't} (R_{nt-1} + \pi_r I_{nt-1}), \quad (13)$$

$$I_{n0} = R_{n0} = 0, \quad S_{n0} = L_{n0}, \quad \tilde{I}_{n0} \geq 0. \quad (14)$$

We consider two types of scenarios. In the first, people continue to expect that $\tau_{nn't} = \tau_{nn'}$ for all $t \geq 0$, as in the pre-infection economy. As long as they do, $m_{nn't} = m_{nn'}$. In the second scenario, people expect $\{\tau_{nn't}\}_{n' \neq n, t > 0}$ to vary as a result of government action. Given values for $\{u_n\}_n$ and $\{\tau_{n'n}\}_{n' \neq n}$ consistent with the observed pre-infection mobility patterns, $\{m_{n'n't}\}_{n, n', t \geq 0}$ in the post-infection economy can then be derived conditional on mobility-barrier expectations using equations (3) and (4).

2.3 Mobility Barriers and Disease Spread: Two Special Cases

We briefly explore two special cases of the model, characterising the spread of the disease under extreme assumptions about bilateral mobility barriers. These special cases offer some intuition about the impact of mobility barriers – due to geography or government policy – on the course of an epidemic. We impose $\pi_d = 0$ for the remainder of this section.

2.3.1 Perfect Mobility

Suppose there are no mobility barriers: $\tau_{nn't} = 1$ for all n, n', t . It is easy to show that in this case $m_{nn'} = L_{n'}/L$ for all n, n' .

It follows immediately from (12) that, irrespective of the distribution of initial infections across n on day 0, infections will be proportional to local populations from day 1 onwards. As a result, the behaviour of $S_t \equiv \sum_n S_{nt}$, $I_t \equiv \sum_n I_{nt}$ and $R_t \equiv \sum_n R_{nt}$ can be characterised completely independently of $\{\tilde{I}_{n0}\}_n$.¹²

2.3.2 No Mobility

Suppose now there are prohibitive mobility barriers: $\tau_{nn't} \rightarrow \infty$ for all $n \neq n', t$. In this case, the behaviour of S_t , I_t and R_t will reflect the weighted sum of S_{nt} , I_{nt} and R_{nt} across N autarkic “islands”. The spatial distribution of initial infections is now essential.

For different values of t , Figure 1 plots $(I_{nt} + R_{nt})/L_n$ – the share of the population in n that has contracted the disease by day t – against the share of initial infections in the population, \tilde{I}_{n0}/L_n . As can be seen from the figure, that relationship is concave for all t . Therefore, unless $\tilde{I}_{n0}/L_n = \sum_n \tilde{I}_{n0}/L$ for all n , the share of infections in the total population, I_t/L_t , will be smaller or equal than it would have been under perfect mobility. In the case in which $\tilde{I}_{n0}/L_n = 0$ for some n it will be strictly smaller forever. This illustrates the case for quarantines: mobility restrictions generally slow the overall spread of a disease in an economy, and they may even prevent some infections all together.¹³

¹²More details on this special case can be found in the [online Appendix](#) Section A.3.2.

¹³The concavity on display in Figure 1 is crucial to this argument. While the graph in the figure is drawn for particular disease parameters, we show in the [online Appendix](#) Section A.3.3 that the concave relationship is generic.

3 Data and Calibration

3.1 Data

3.1.1 Population, Migration and Commuting Flows from 2011 UK Census

To illustrate how our structural-gravity SIR framework can be put into action, we use it to analyse the spread of a disease across local authorities in Great Britain (England, Scotland and Wales). The properties of the disease are inspired by the Covid-19 pandemic. Crucially, they include a relatively small average daily probability of death for the infected (see Section 3.2.2).

We rely on information on population, migration and commuting from the latest UK Census, conducted in 2011.¹⁴ The data can be aggregated to the local-authority level using concordances provided by the UK’s Office for National Statistics (ONS). The boundaries of local authorities circumscribe areas administered by a single local government, typically a local council. After aggregation, we obtain data for 378 local authorities covering all of Great Britain. The median local-authority district had a population of 130,000 in 2018, with 90% of local-authority populations in the range of 60,000 to 360,000.

The 2011 Census reports information on all regular residents of an area in 2011 who lived at a different address one year prior. It also reports the location of an individual’s usual place of residence and place of work in 2011. After aggregation this allows us to compute, for any two local authorities A and B, what share of residents of A moved to B permanently in 2010-11 and what share of residents of A commuted for work to B in 2011. As we show in [online Appendix Section B.1.1](#), the corresponding bilateral flows of migrants and commuters exhibit strong gravity features: both “naive” and structural gravity regressions capture a large share of the variation in bilateral flows observed in the data.

3.1.2 Daily Bilateral Flows of People Between Local Authorities

We combine the Census migration and commuting data to calculate a daily flow of people between any two local authorities. For migrants, we divide the 2010-11 figures by 365 to obtain the daily flow. For commuters, we first “balance” flows to reflect that commuting represents gross flows that do not cause a net change in local populations. For example, if 60 people report commuting from A to B in 2011, and 40 people report commuting from B to A, we put the potential number of people from each place who could spend the day in the other at $(60+40)/2=50$. We then adjust for the fact that commuters will travel between A and B only for half of the average workday. In our example, this implies that on the average day $(5/7-34/365)\times 50/2$ people go for work

¹⁴See Office for National Statistics (2015).

from A to B, and from B to A, where we assume that the average work week is 5 days and the average annual number of holidays is 34.

Adding average daily migrant and commuter flows thus constructed, we obtain our final measure of the daily bilateral flow of people – including the shares of resident populations that tend to stay within their respective local authorities on the average day. This data is described in more detail in Section B.1.2 of the [online Appendix](#). Unsurprisingly, the bilateral daily flows inherit the gravity features of the underlying variables used to construct it. This can be seen in Table 1.

[Insert Table 1 here]

Column (1) of Table 1 reports the results of a “naive” gravity regression of bilateral flows on only size variables, distance measures and a constant term. Column (2) reports the results of a structural gravity regression of the form

$$m_{nn'}L_n = e^{\Pi_{n'} + \Omega_n + \phi_1 \ln dist_{nn'} + \phi_2 contig_{nn'}} \varepsilon_{nn'}, \quad (15)$$

where $\Pi_{n'}$ is a place- n' -as-destination fixed effect, Ω_n is a place- n -as origin fixed effect; $dist_{nn'}$ and $contig_{nn'}$ are measures of geographic distance; and $\varepsilon_{nn'}$ is an error term. Note that the sets of origin and destination fixed effects are not of full rank, so we impose the restriction $\Pi_N = 0$ for an arbitrary benchmark local authority N . Both the “naive” and structural gravity regressions are estimated in levels using Poisson Pseudo-Maximum Likelihood (PPML). This makes it possible to accommodate the fact that approximately 12% of daily bilateral flows between local authorities are zero. It also allows us to leverage some convenient properties of the PPML estimator in the context of structural gravity models, as discussed in subsection 3.2.1.

Columns (1) and (2) in Table 1 show that gravity-style regressions can account for a large share of the observed variation in bilateral flows of people between local authorities: both columns report very high values of R^2 . Moreover, column (2) reveals a distance elasticity of -1.9 and shows that daily bilateral flows between contiguous places are approximately ($e^{.971} =$) 2.6 times as large as between non-contiguous places.

3.2 Calibration

3.2.1 Bilateral Mobility Barriers and Relative Place Values

As shown in Fally (2015), the estimated fixed effects in a PPML specification of equation (15) are consistent with the definition of the inward and outward MRTs in

equation (5) and the equilibrium constraints that these need to satisfy.¹⁵ In turn, this implies that our structural gravity regression supplies us with two sets of parameter restrictions of the form

$$\frac{v_n^{-\theta}}{v_N^{-\theta}} = e^{-\hat{\Pi}_n}, \quad (16)$$

$$\tau_{nn'}^{-\theta} = (\text{dist}_{nn'})^{\hat{\phi}_1} (e^{\text{contig}_{nn'}})^{\hat{\phi}_2}. \quad (17)$$

Equation (16) pins down the place value of any local authority n relative to the benchmark local authority N , as a function of the destination fixed effects and the benchmark’s population, and up to the value of θ . Equation (17) yields the level of bilateral mobility barriers as a function of distance, the contiguity indicator, and the coefficient estimates $\hat{\phi}_1, \hat{\phi}_2$, up to the value of θ .

These parameter restrictions imply daily bilateral movement propensities

$$\hat{m}_{nn'} = \frac{(\text{dist}_{nn'})^{\hat{\phi}_1} (e^{\text{contig}_{nn'}})^{\hat{\phi}_2} e^{-\hat{\Pi}_{n'}}}{\sum_{n'=1}^N (\text{dist}_{nn'})^{\hat{\phi}_1} (e^{\text{contig}_{nn'}})^{\hat{\phi}_2} e^{-\hat{\Pi}_{n'}}}. \quad (18)$$

Note that $\hat{m}_{nn'}$ combines model-consistent estimates of the impact of geography on bilateral mobility with model-consistent estimates of the relative attractiveness of different destinations. This constitutes one of the key advantages of our structural-gravity SIR framework: it provides a clear mapping from geographic observables and regional mobility data into bilateral movement propensities that have choice-theoretic microfoundations and ultimately shape the spread of a disease across space. Using (7), (16) and (18), we can also back out the relative flow utilities associated with different places for a given values of β and θ .

Descriptive statistics for the relative place values, mobility barriers and relative flow utilities derived from our structural gravity regression can be found in the [online Appendix](#) Section B.2.1.

3.2.2 Disease Parameters, Initial Infections and Initial Populations

We base our calibration of the disease parameters on the “Imperial College Study” (Ferguson et al., 2020) that assessed the likely spread of the Covid-19 pandemic in the UK and the US in the absence of public intervention as of 16 March 2020. Ferguson et al. (2020) assume that the disease is characterised by a 6.5 day generation period, with an average probability of death of 0.9% among the infected. In line with this, we

¹⁵Crucially, in our setting the estimation of (15) by PPML with place-destination and place-origin fixed effects implies:

$$\sum_{n=1}^N \hat{m}_{nn'} L_{nt-1} = L_{n't}, \text{ and } \sum_{n'=1}^N \hat{m}_{nn'} L_{nt-1} = L_{nt-1}.$$

impose $\pi_d = .009/6.5$ and $\pi_r = .991/6.5$. In our model π_s/π_r represents the so-called “r zero” of the disease. Based on initial evidence from the spread of the pandemic in Wuhan, Ferguson et al. (2020) examine values of the “r zero” between 2.0 and 2.6. We choose a value close to the middle of this range, setting $\pi_s = 2.2\pi_r$.

Official statistics on Covid-19 cases in the UK have been released since 9 March 2020. For our simulations, we seed initial infections at the local-authority level consistent with the pattern of Covid-19 cases reported by the UK, Scottish and Welsh Governments on 10 March 2020. The sources of this data, and distribution of cases, are detailed in the [online Appendix](#) Section B.2.2. Ferguson et al. (2020) cite evidence from China and repatriation flights suggesting that 40-50% of infections are not identified as cases. To reflect this, and the relative initial scarcity of Covid-19 testing in the UK, we assume that the number of cases reported at the local-authority level on 10 March reflected 30% of actual infections, and set initial infection levels $\{\tilde{I}_{n0}\}_n$ accordingly.

For expositional convenience, we set initial local-authority populations $\{L_{n0}\}_n$ to equal the steady-state populations implied by $\{\hat{m}_{nn'}\}_{n,n'}$. However, these steady-state populations are almost perfectly correlated with 2018 mid-year population estimates for local authorities published by ONS.

3.2.3 Discount Factor, Mobility Elasticity and Value of Life

We set the discount factor to $\beta = .96^{1/365}$, implying an approximately 4% annual discount rate as in Eichenbaum et al. (2020). Since θ is a crucial parameter in our welfare analysis, we experiment with different values. However, in our baseline calibration, we impose $\theta = 3.3$ to reflect evidence on heterogeneity in location preferences from US commuting data (see Monte et al., 2018).

Finally, equations (7) and (16) only pin down place values and flow utilities in *relative* terms. We are thus free to select u_N to determine the absolute values of $\{v_n, u_n\}_n$. This “level” choice has no impact on individuals’ location decisions in the model, but it translates into the average daily utility received by agents. By assumption, the only utility consequence of an infection is the risk of death, and the only cost of death is the forgone utility of life. Therefore, u_N emerges as another crucial parameter for the welfare trade-off between mobility restrictions and disease control.

We perform our welfare analysis under two different calibrations of u_N . In the first, we make the conservative assumption that $u_{n'}z_{n't}(n)/\delta_{nn'}$ reflects only the real consumption of an agent from n who spends period t in n .¹⁶ We then set u_N such that

¹⁶This would be true under the following narrow interpretation of our model assumptions: all agents produce a homogenous, perfectly tradable good and choose locations to maximise their productivity in t , given by $u_{n'}z_{n't}(n)/\delta_{nn'}$.

average daily consumption is equal to \$126, which corresponds to 2018 UK daily GDP per capita in purchasing-power-adjusted US dollars. In the second calibration, we assume that $u_{n'} z_{n't}(n) / \delta_{nn'}$ reflects a broader notion of value of life, and set its average daily value across the pre-infection population to \$1040. Under the assumption of a 4% annual interest rate, the latter calibration translates into an average value of life of \$9.3m – the economic value of life used by US public authorities, such as the Environmental Protection Agency.

4 Simulations

4.1 Baseline: “Do Nothing”

We first simulate the course of the epidemic in the absence of any public intervention: a “do nothing” scenario. The resulting evolution of the shares of the susceptible, infected and recovered in the total population of Great Britain, as well as the number of deaths per day, are shown in Figure 2 (black lines). The share of the infected peaks at 18,500 per 100,000 population on day 67. The number of deaths per day peaks at 26 per 100,000 population on day 68. Over the entire course of the epidemic 494,000 people die, equivalent to 0.77% of the total population. 85% of the population ultimately contract the disease.

While our model is considerably simpler than the model used in Ferguson et al. (2020), it replicates the aggregate evolution of the Covid-19 epidemic envisaged in their baseline scenario fairly closely. In Ferguson et al. (2020), British infections and deaths in a “do nothing” scenario peak in late-May 2020, around day 70 in our model. The number of deaths per day peaks at 22 per 100,000 population, with a total number of deaths of 510,000 overall. In the long run, 81% of the British population contract the disease.¹⁷

There is also a short time window during which the model can be evaluated against actual data. After adopting a fairly light-touch approach to the containment of Covid-19 initially, the UK Government imposed a lockdown on 23 March 2020. In the [online Appendix](#) Section C.1, we compare the growth of infections reported across the 9 main administrative areas of the British National Healthcare System (NHS) during the 10-23 March period with the model-predicted growth of infections in this regions during days 0-13. We find that model-predicted and observed growth rates are strongly but by no means perfectly correlated.

¹⁷See Ferguson et al. (2020), pp. 6-8 and Figures 1 and 2.

4.2 Quarantine Scenarios

We now consider 3 regional-quarantine scenarios. In each scenario, the UK Government requires the public not to leave their local-authority districts for 120 days.¹⁸ We choose 120 days because this period is sufficiently long for new infections to have dropped to (almost) zero by its end in each simulation. In all cases, the government commits publicly and credibly to the length of the quarantine. The scenarios only differ as to the date in which the quarantine is introduced, and whether the introduction is expected by the public or not.

4.2.1 120-Day Regional Quarantine from $t = 0$

We first assume the quarantine is introduced as an “MIT shock” along with the original outbreak on day 0. The resulting dynamics of the epidemic are shown in Figure 2 (blue line). It is clear that the course of the epidemic is considerably milder. Infections and deaths peak around the same time as in the baseline scenario, but at much lower levels. In the long run, only 60% of the population contract the disease, and a quarter of the baseline-scenario deaths are avoided.

The relative effectiveness of the day-0 quarantine stems from the fact that a considerable number of local authorities are virus-free on day 0. With an instant quarantine, residents of these local authorities never contract the virus. This can be seen in Figure 3 which compares the regional spread of the disease on day 30 of our simulations across different scenarios. A comparison between panels A and B reveals the local authorities which are “spared” infection as a result of the quarantine. The long-run effect is a smaller share of the recovered and of deaths in the population.

4.2.2 Unexpected 120-Day Regional Quarantine from $t = 13$

We then assume that the quarantine is introduced as an “MIT shock” on day 13, roughly corresponding to the timing of the UK Government’s lockdown. Figure 2 (maroon line) illustrates that this delayed quarantine is considerably less effective than the day-0 quarantine: the long-run number of infections and deaths is the same as in the baseline scenario. However, even the delayed quarantine achieves some “flattening of the curve”: the peak of deaths and infections occurs later than in the baseline, and the peak number of infections and deaths are somewhat lower.

Unlike on day 0, by day 13 the virus has reached all local authority districts. However, the share of infections in local populations still varies considerably. As shown in Section 2.3.2 mobility restrictions can still slow the spread of a disease in such a setting. This gives rise to the “flattened” maroon curves in Figure 2. It is also

¹⁸However, people may continue to move freely within local authorities. This is much less restrictive than the lockdown actually imposed by the UK Government on 23 March 2020.

evident from a comparison of panels A and C of Figure 3: by day 30, the epidemic has spread more evenly across local authorities if the government does nothing than if a quarantine is unexpectedly introduced on day 13.

4.2.3 Expected 120-Day Regional Quarantine from $t = 13$

Finally, we assume that a quarantine is introduced on day 13, but that the public expects its introduction from day 0. The green line in Figure 2 traces the resulting course of the epidemic. However, it is obscured by the black line, since the course of the disease is virtually the same as in the “do nothing” baseline.

The anticipation of a 120-day regional quarantine on day 13 undoes all “flattening” benefits that would arise if the quarantine were introduced unexpectedly. The reason is that individuals re-optimize their locations during days 0-13. For the duration of the quarantine, in which all movements across local-authority boundaries are ruled out, locations with a low connectivity index in normal times gain in utility value relative to locations with a high index. Therefore, there is some reshuffling of populations from the latter to the former before day 13. With the initial conditions we impose in Section 3.2.2, the incidence of infections in period 0 is relatively high in high-connectivity places. Agents’ responses in anticipation of the quarantine thus spread the epidemic around the country more quickly.

Panel D of Figure 3 gives some sense of the impact: 30 days into the simulation, and 17 days after the imposition of a regional quarantine, the disease is spread more evenly around the country than in *any* of the other three scenarios.

4.3 Welfare Comparison

In the wake of an epidemic outbreak, the introduction of a quarantine presents a clear welfare trade-off: while mobility restrictions reduce social welfare in the short run, they may delay – or even prevent – infections and deaths. An advantage of our structural-gravity SIR framework is that it provides a tool to explore this trade off in a parsimonious manner, conditional on the values of a well-specified set of parameters.

Here we proceed with a suggestive welfare assessment of the three quarantine scenarios from Section 4.2 against the “do nothing” baseline. Needless to say, these scenarios do not span the full set of conceivable policy interventions. Moreover, the model we have outlined captures the fundamental trade-off between mobility and disease control in a bare-bones fashion. For these reasons, our welfare analysis should not be taken as the definitive assessment of plausible (or even likely) interventions in the face of a Covid-19-style epidemic. Instead, they only serve to illustrate that structural-gravity SIR can serve as a useful building block upon which to base such assessments.

Table 2 reports log changes in welfare relative to baseline from introducing the three quarantines described in Section 4.2. We report results for two different average values of life (values of u_N): \$126 and \$1040 per day, as discussed in Section 3.2.3. We also vary the value the mobility elasticity, θ , between 2 and 10 with a baseline value of 3.3. All other parameters are held constant.

Across different parametrisation, the table presents a consistent ranking of the alternative scenarios: the instant quarantine significantly improves welfare relative to “doing nothing”, while the delayed quarantine constitutes a marginal welfare improvement. The expected delayed quarantine *reduces* welfare, as it introduces temporary mobility restrictions without controlling the spread of the disease. If agents value mobility less (higher values of θ), the welfare improvements from the instant and delayed quarantines are larger. The same is true if the value of life is higher.

5 Conclusion

We set out to show that insights from structural-gravity modelling may prove useful in the emerging economics of epidemics. To this end, we build a bare-bones model in which flows of people across space are governed by a structural gravity equation, and contribute to the spread of a deadly disease. We demonstrate that the model can be readily applied to real-world data, and captures the fundamental welfare trade-off between mobility restrictions and disease control in a fully microfounded fashion.

Our simple framework could be generalised in a variety of ways to explore this welfare trade-off more thoroughly. Such generalisations may include the incorporation of a formally modelled production side (as in Eichenbaum et al. 2020), heterogeneous agents (as in Acemoglu et al., 2020), or broader welfare costs of infections and hospital-capacity constraints (as in Ferguson et al, 2020). With suitable modifications and the right data, it may also form the basis of an analysis of mobility restrictions and disease spread at the level of neighbourhoods and households, or at the level of countries.

References

- [1] Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston, 2020. “A Multi-Risk SIR Model with Optimally Targeted Lockdown,” NBER Working Paper No. 27102.
- [2] Ahlfeld, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf, 2015. “The Economics of Density: Evidence from the Berlin Wall,” *Econometrica*, 83, 6, pp. 2127-2189.
- [3] Allen, Linda J. S., 1994. “Some Discrete-Time SI, SIR and SIS Epidemic Models,” *Mathematical Biosciences*, 124, pp. 83-105.
- [4] Allen, Treb, Costas Arkolakis, and Xiangliang Li, 2016. “Optimal City Structure,” mimeo.
- [5] Anderson, James E., 1979. “A Theoretical Foundation for the Gravity Equation,” *American Economic Review*, 69, 1, pp. 106-116.
- [6] Anderson, James E., 2011. “The Gravity Model,” *Annual Review of Economics*, 3, 1, pp. 133-160.
- [7] Anderson, James E., and Eric van Wincoop, 2003. “Gravity with Gravitas: A Solution to the Border Puzzle,” *American Economic Review*, 93, 1, pp. 170-192.
- [8] Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, 2012. “New Trade Models, Same Old Gains?,” *American Economic Review*, 102, 1, pp. 94-130.
- [9] Artuç, Erhan, and John McLaren, 2015. “Trade Policy and Wage Inequality: A Structural Analysis with Occupational and Sectoral Mobility,” *Journal of International Economics*, 97, 2, pp. 278–294.
- [10] Artuç, Erhan, Paulo Bastos, and Eunhee Lee, 2020. “Trade, Jobs, and Worker Welfare,” mimeo.
- [11] Atkeson, Andrew G., 2020. “What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios,” manuscript, UCLA.
- [12] Beenstock, Michael, and Xieer Dai, 2020. “OLG as an Alternative to SIR,” *Covid Economics*, 10, pp. 87-115.
- [13] Bryan, Gharad, and Melanie Morten, 2018. “The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia,” *Journal of Political Economy*, 127, 5, pp. 2229-2268.

- [14] Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro, 2019. “Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock,” *Econometrica*, 87, 3, pp. 741-835.
- [15] Costinot, Arnaud, and Andrés Rodríguez-Clare, 2014. “Trade Theory with Numbers: Quantifying the Consequences of Globalization,” in Gita Gopinath, Elhanan Helpman and Kenneth S. Rogoff (eds.), *Handbook of International Economics*, 4, chapter 4, pp. 197-261.
- [16] Eaton, Jonathan, and Samuel Kortum, 2002. “Technology, Geography, and Trade,” *Econometrica*, 70, 5, pp. 1741-1779.
- [17] Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt, 2020. “The Macroeconomics of Epidemics,” NBER Working Paper No. 26882.
- [18] Fally, Thibault, 2015. “Structural Gravity and Fixed Effects,” *Journal of International Economics*, 97, 1, pp. 76-85.
- [19] Ferguson, Neil M., Derek A.T. Cummings, Simon Cauchemez, Christophe Fraser, Steven Riley, Aronrag Meeyai, Sophon Iamsirithaworn and Donald S. Burke, 2005. “Strategies for Containing an Emerging Influenza Pandemic in Southeast Asia,” *Nature*, 437, 8, pp. 209-214.
- [20] Ferguson, Neil M., Derek A.T. Cummings, Christophe Fraser, James C. Cajka, Philip C. Cooley, and Donald S. Burke, 2006. “Strategies for Mitigating an Influenza Pandemic,” *Nature Letters*, 442, 27, pp. 448-452.
- [21] Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunubá, Gina Cuomo-Dannenburg, Amy Dighe, Ilaria Dorigatti, Han Fu, Katy Gaythorpe, Will Green, Arran Hamlet, Wes Hinsley, Lucy C. Okell, Sabine van Elsland, Hayley Thompson, Robert Verity, Erik Volz, Haowei Wang, Yuanrong Wang, Patrick G.T. Walker, Caroline Walters, Peter Winskill, Charles Whitaker, Christl A. Donnelly, Steven Riley, and Azra C. Ghani, 2020. *Report 9: Impact of Non-pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand*.
- [22] González-Eiras, Martín and Dirk Niepelt, 2020. “On the Optimal ‘Lockdown’ During an Epidemic,” *Covid Economics*, 7, pp. 68-87.
- [23] Head, Keith, and Thierry Mayer, 2014. “Gravity Equations: Workhorse, Toolkit, and Cookbook,” in Gita Gopinath, Elhanan Helpman and Kenneth Rogoff (eds.), *Handbook of International Economics*, 2014, 4, pp. 131-195.

- [24] Kermack, William Ogilvy, and Anderson G. McKendrick, 1927. “A Contribution to the Mathematical Theory of Epidemics,” *Proceedings of the Royal Society of London, Series A*, 115, 772, pp. 700-721.
- [25] McFadden, Daniel, 1974. “The Measurement of Urban Travel Demand,” *Journal of Public Economics*, 3, 4, pp. 303-328.
- [26] Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg, 2018. “Commuting, Migration, and Local Employment Elasticities,” *American Economic Review*, 108, 12, pp. 3855-3890.
- [27] Office for National Statistics, 2015. *2011 Census: Flow Data* [data collection]. UK Data Service, SN: 7713.
- [28] Ossa, Ralph, 2015. “Why Trade Matters After All,” *Journal of International Economics*, 97, 2, pp. 266-277.
- [29] Santos Silva, João, and Silvana Tenreyro, 2006. “The Log of Gravity,” *Review of Economics and Statistics*, 88, 4, pp. 641-658.
- [30] Xia, Yingcun, Ottar N. Bjørnstad, and Bryan T. Grenfell, 2004. “Measles Metapopulation Dynamics: A Gravity Model for Epidemiological Coupling and Dynamics,” *The American Naturalist*, 164, 2, pp. 267-281.

Dep. variable:		
# persons from origin going to destination on the average day	(1)	(2)
Ln distance (km)	-1.925*** (0.015)	-1.907*** (0.013)
= 1 if contiguous	0.977*** (0.057)	0.971*** (0.053)
Ln population in origin	0.502*** (0.029)	
Ln population in destination	0.502*** (0.029)	
Observations	142,884	142,884
Places	378	378
Adjusted R^2	.997	.996
Fixed effects:		
- Constant term	Yes	No
- Place-origin	No	Yes
- Place-destination	No	Yes

* $p < .10$; ** $p < .05$; *** $p < .01$;

Table 1: Gravity regressions on daily bilateral flows of people between GB local authorities

Regressions estimated with Poisson Pseudo-Maximum Likelihood (PPML). Robust standard errors in parentheses. The dependent variable is the daily bilateral flow of people from the origin local authority to the destination local authority. This flow is calculated on the basis of UK 2011 Census data (see text for details). “Ln distance (km)” is the natural logarithm of the kilometre distance between the geographic mid-points of the origin and destination local authorities; “= 1 if contiguous” represents a dummy which takes value 1 if two local authorities share a common border, 0 otherwise; “Ln population in origin/destination” is the natural logarithm of the resident population in the origin/destination local authority.

Avg. value of life = \$126 per day				Avg. value of life = \$1040 per day			
θ	Quarantine in			θ	Quarantine in		
	$t = 0$	$t = 13$	$t = 13$		$t = 0$	$t = 13$	$t = 13$
		(unexp.)	(exp.)			(unexp.)	(exp.)
2	1.028	.021	-.017	2	1.489	.042	-.020
3.3	1.039	.032	-.013	3.3	1.500	.056	-.015
10	1.050	.043	-.008	10	1.51	.064	-.010

Table 2: Welfare assessments of different quarantine scenarios

The table shows the permanent log change in daily consumption the average agent would require from day 0 to be compensated for the “do nothing” baseline being adopted over a particular quarantine scenario. The required compensations are shown for different values of the mobility elasticity, θ , and different values of the reference local authority’s utility flow, u_N (resulting in different average values of life). For details on the calibration, see Section 3.2. For details on the different scenarios, see Sections 4.1-4.3.

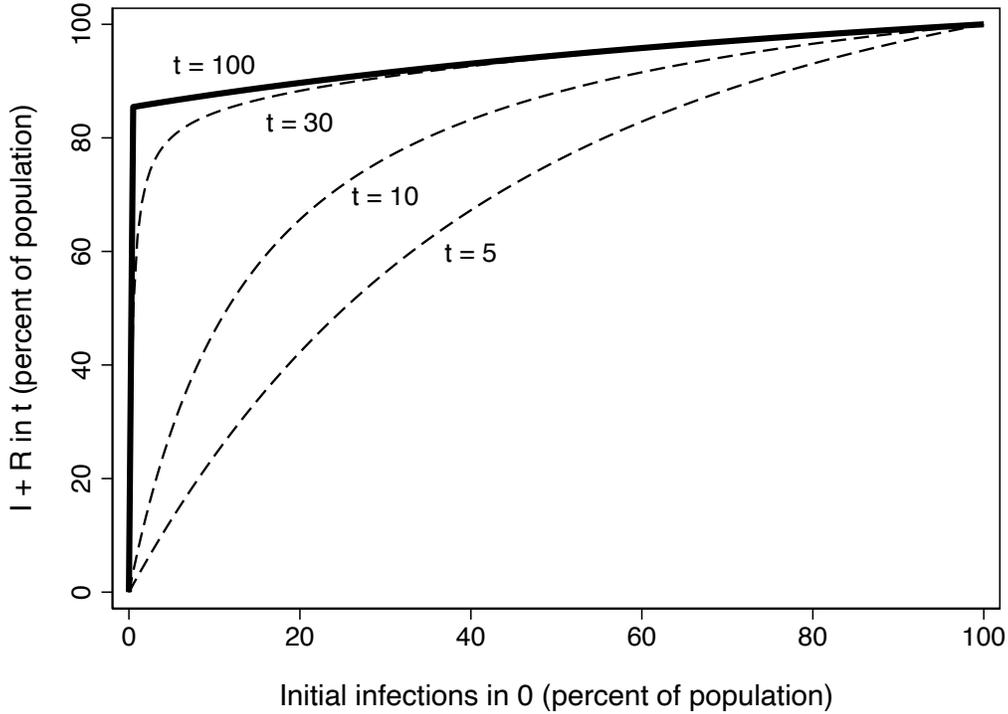


Figure 1: Share of population that has contracted disease in t against share of initial infections

The horizontal axis measures the infections seeded in a(n autarkic) location on day 0, as a percentage of the location's population. The vertical axis measures the number of people who have contracted the disease by day t in that location, $I_{nt} + R_{nt}$, as a percentage of the location's population. Graph drawn using equations (11)-(14) for $t = 5, 10, 30, 100$, imposing $\pi_d = 0$, $\pi_r = 1/6.5$, $\pi_s = 2.2\pi_r$.

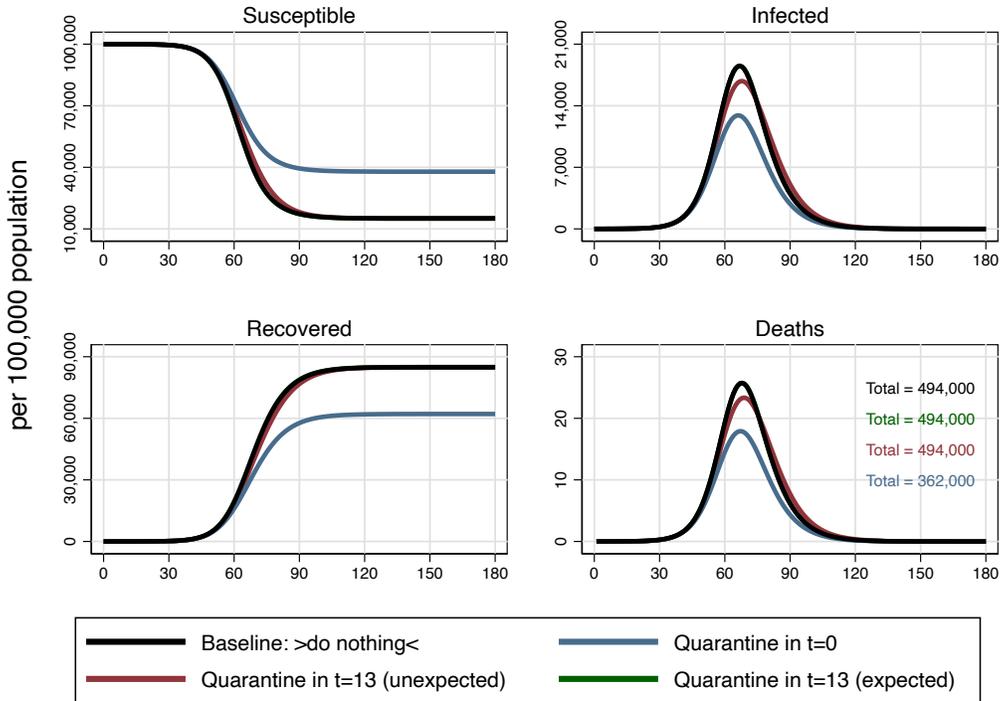


Figure 2: Susceptible, infected, recovered and deaths per day under different scenarios

The top-left panel plots the susceptible per 100,000 population against time in days. The top-right panel plots the infected per 100,000 population against time in days. The bottom-left panel plots the recovered per 100,000 population against time in days. The bottom-right panel plots deaths per day against time in days. All reported output represents Great Britain totals. Black line: baseline “do nothing” scenario (see Section 4.1). Blue line: quarantine in $t = 0$ (see Section 4.2.1). Maroon line: unexpected quarantine in $t = 13$ (see Section 4.2.2). Green line: expected quarantine in $t = 13$ (see Section 4.2.3).

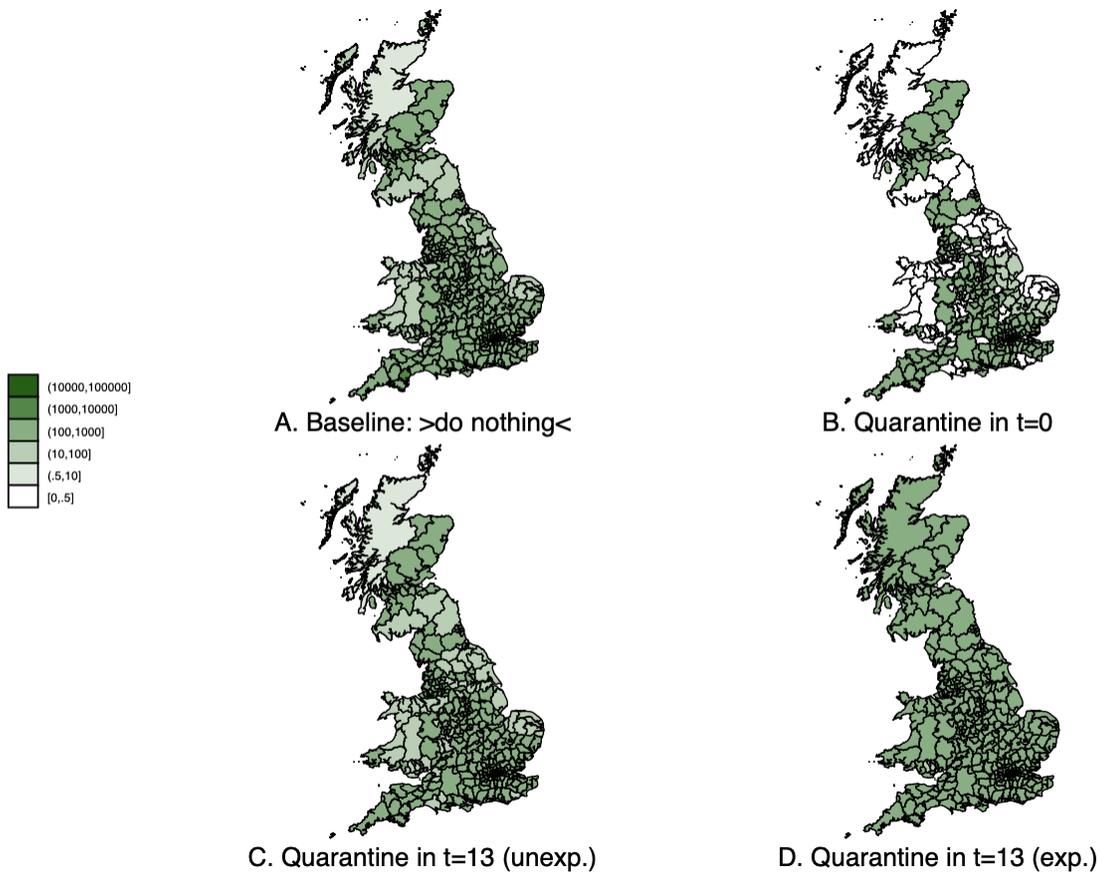


Figure 3: Infections per 100,000 population in $t = 30$ under different scenarios

Total number of infections at the local authority-level (I_{nt}) per 100,000 population in $t = 30$ for four different scenarios. Panel A: baseline “do nothing” scenario (see Section 4.1). Panel B: quarantine in $t = 0$ (see Section 4.2.1). Panel C: unexpected quarantine in $t = 13$ (see Section 4.2.2). Panel D: expected quarantine in $t = 13$ (see Section 4.2.3). Note: Shetland Islands excluded.