

Meenagh, David; Minford, Patrick

Working Paper

A structural model of coronavirus behaviour for testing on data behaviour

Cardiff Economics Working Papers, No. E2020/4

Provided in Cooperation with:

Cardiff Business School, Cardiff University

Suggested Citation: Meenagh, David; Minford, Patrick (2021) : A structural model of coronavirus behaviour for testing on data behaviour, Cardiff Economics Working Papers, No. E2020/4, Cardiff University, Cardiff Business School, Cardiff

This Version is available at:

<https://hdl.handle.net/10419/250314>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Cardiff Economics Working Papers



Working Paper No. E2020/4

A structural model of coronavirus behaviour for testing
on data behaviour

David Meenagh and Patrick Minford

January 2021

ISSN 1749-6010

Cardiff Business School
Cardiff University
Colum Drive
Cardiff CF10 3EU
United Kingdom
t: +44 (0)29 2087 4000
f: +44 (0)29 2087 4419
business.cardiff.ac.uk

This working paper is produced for discussion purpose only. These working papers are expected to be published in due course, in revised form, and should not be quoted or cited without the author's written permission. Cardiff Economics Working Papers are available online from:
<http://econpapers.repec.org/paper/cdfwpaper/> and
business.cardiff.ac.uk/research/academic-sections/economics/working-papers
Enquiries: EconWP@cardiff.ac.uk

A structural model of coronavirus behaviour for testing on data behaviour

David Meenagh*

(Cardiff Business School, Cardiff University)

Patrick Minford

(Cardiff Business School, Cardiff University, and CEPR)

January 2021

Abstract

We fit the logistic function, the reduced form of epidemic behaviour, to the data for deaths from Covid-19, for a wide variety of countries, with a view to estimating a causal model of the covid virus' progression. We then set out a structural model of the Covid virus behaviour based on evolutionary biology and social household behaviour; we estimated and tested this by indirect inference, matching its simulated logistic behaviour to that found in the data. In our model the virus' progression depends on the interaction of strategies by household agents, the government and the virus itself as programmed by evolution. Within these interactions, it turns out that there is substitution between government topdown direction (such as lockdown) and social reaction to available information on the virus' behaviour. We also looked at experience of second waves, where we found that countries successfully limited second waves when they had had longer first waves and followed policies of localised reaction in the second.

Keywords: coronavirus, Covid-19, evolution, optimisation, indirect inference, lockdown

JEL Classification: C54, I12

*Corresponding Author: Cardiff Business School, Colum Drive, Cardiff, CF10 3EU, UK. Email: MeenaghD@cardiff.ac.uk

1 Introduction

In this paper we propose a causal model of coronavirus behaviour and estimate it and test it empirically by its ability to match the behaviour of infections and deaths observed in the UK and in other parts of the world. We then use the model to evaluate the effectiveness of the contrasting government policies in different countries for dealing with the virus. It might be asked why there is a need for such a model, when there are virological models of the corona virus disease. We suggest that these models do not take account of the reactions by agents, the virus itself and governments to the shocks created by the arrival of the pandemic in a population. These models essentially assume a deterministic process of susceptibility, infection and recovery; yet, as we shall show, there are many points in the process where there are both choices and stochastic processes; we build on both economic theory and evolutionary biology for our relatively simple utility-maximising causal model. There is now a large body of data which can be explored for the purposes of modelling these processes; our main contribution is to estimate this causal model by the powerful estimation and testing mechanism of indirect inference which uses both the logistic reduced form estimates for the disease and the causal model predictions to pin down the causal parameters.

We begin in section 2 by describing the logistic model's description of the data behaviour produced by the virus' progress. In section 3, we set out our structural ("causal") model of the virus' and households' behaviour in response to shocks, including from government policy. In section 4 we explain the estimation method and how we use it to evaluate our structural model by its ability to match the data behaviour revealed by our descriptive logistic model, termed the 'auxiliary model' in this method. We then, in section 5, show the results of our estimation and testing of the structural model. Next we comment on some key estimates from the model and discuss their implications for policy, before concluding in the final section.

To anticipate our main results, we use our models to contrast the experiences of the UK and Sweden, and also those of 28 countries experiencing second waves. We find from the Sweden/UK comparison that there is substitution between government topdown direction (such as lockdown) and social reaction to available information on the virus' behaviour; and that the government's optimal role is in providing information and coordinating social reaction at a decentralised level, so avoiding topdown disruption to the economy. From the experience of second waves we found

that countries successfully limited second waves when they had had longer first waves and followed policies of localised reaction in the second.

2 Describing the data — the logistic curve




Figure 1: Confirmed Deaths by Day Since Total Passed 20

The chart of the progress of deaths shown in Figure 1¹ — plotted on a log (“ proportional) scale — show a common and coherent pattern, which comes from an underlying ‘logistic stock-flow’ model of the virus; such a model is widely used to project how innovations spread through a population — whether it is new ideas, new technologies, or as here infections. Imagine that you have a population free of the virus, ranged from those with easy infectability at the one end to some at the other with great immunity. Enter the virus, with a mechanism of transmission from person to person via coughing, touching etc. In the initial slow stage, the virus will take time to infect a substantial group. In the second rapid stage, there will be a high speed of infection as the susceptible will quickly catch it and pass it on to other susceptible people of whom many are available. At this point the virus’ reproductive rate (R_0) will be high, with each infection leading to several others in a short time. The progress will look ‘exponential’ (an exponential curve grows without limit) but it is not, because there is a further stage.

¹Source of data: Johns Hopkins University Center for Systems Science and Engineering

As the stock of infected people accumulates, the virus needs to spread to people with greater natural immunity. The rate of infection (the flow of new infections) and that R_0 rate will slow. As the stock of infected people reaches the last tranche of people with the highest immunity, the rate will gradually fall to a stop. In the end the whole infectable population will have the virus or have had it.

These three stages — initial infection, rapid spread through widely available cases, and finally slowing in the face of saturation — must occur regardless of the epidemiological details. These details show up in the estimated parameters of the describing logistic curve, which therefore is a reduced form representation in the data for all epidemics and similar population-penetrating processes, whatever their structural details. The reduced form parameters are: the maximum penetration, the rate of infection and the point of inflection where saturation starts to set in. The problem for epidemiological models is that so little is known about this virus. But with the logistic curve we can observe for many countries what these estimated parameters, that reflect this unknown virus' character, are. From this diverse experience we can estimate the progression process in the UK, our primary focus, and also the effects of lockdown, the policy now being fiercely debated across the world. Batista (2019) and Golinski and Spencer (2020) have estimated logistic models for various countries.

But we can do more than find the best logistic curve description by building the structural (“causal”) model of virus behaviour that underlies this logistic reduced form. This structural model, if empirically reliable, can give us an understanding of how policy interventions affect the virus' progress. However, we need a means to establish the model's empirical reliability. For this we use the method of indirect inference where we check the model's capacity to generate the reduced form logistic behaviour we observe in the data. As we will show below when we discuss this method, most familiar in its form as the method of simulated moments, it gives us substantial power to discriminate against inaccurate or misspecified models.

Hence our account of the virus' logistic progress is not intended to replace the careful modelling of the detailed causal processes driving the virus epidemic; rather it is intended to describe the data behaviour of the virus' progress. A structural model of the virus' behaviour, which we develop below, can guide us on the effects of policy interventions such as lockdowns. Medical interventions, such as drugs and vaccines, require specifically medical research, which is being energetically pur-

sued by clinical companies in search of a vaccine and effective drug treatments. But so far none have been found or used except experimentally. Apart from financing and encouraging this pursuit, governments have intervened in two main ways: first by attempted denial of entry of the virus into uninfected populations, through testing, tracing and quarantining and second by lockdown of infected populations. The first has been used by Singapore and South Korea rather effectively. Other countries tried it for a time, the UK among them, but ineffectively, with general popular interaction releasing the virus into general circulation in spite of their efforts. The second intervention of lockdown then has had a plainly visible impact, namely in slowing the early rate of infection and delaying the point of inflection in time. Against this background, structural model estimates can give us practical guidance on what will happen from what has happened so far. This guidance can help to assess orders of magnitude for future cases and deaths which is important when one major clinical group, at Imperial College London, have predicted that deaths would have reached half a million had lockdown not occurred and will reach nearly 50,000 even with the lockdown in place since late March.

3 The rationale of a causal model

We now develop a structural model of the coronavirus' behaviour. Our intention is to test and estimate this model by indirect inference, in which we compare the model's simulated behavior with actual data behaviour and evaluate the match statistically, in a way we explain below in detail. We will fit it to data for the UK and Sweden, with the aim of identifying differential policy effects between the two countries, both in terms of lockdown and general public health protection; in both, policies differed starkly enough for us to identify the effects with moderate precision. In future work these methods could be extended to other countries to evaluate the effects of the wide variety of policies they all followed.

In our structural model there are two (representative) agents: the coronavirus and the household. We treat the coronavirus as having an optimised strategy for infecting a population it has been donated by chance to infect. We can think of this optimisation as having been crafted by natural selection over a long period of evolution; in other words today's virus has evolved to survive because its strategy has been optimised for survival. These ideas belong partly to evolutionary

biology (Nesse et al, 2010) and partly to recent DSGE modelling in macroeconomics (Le et al, 2011) where agents are treated as if they are optimising strategic decision-makers; here the virus is treated as an optimising agent, whose strategy has been selected by mutation and evolution. We think of the virus as having mutated by natural selection over previous episodes of contact with populations. However, we are currently modelling a particular episode’s population that constitutes a new environment, with differences from the previous ones. We divide this environment into elements the virus cannot control but must simply react to, due to the ‘surprises’ in the current population: these include the death rate, which will reflect the particular make-up of the population (e.g. more or fewer old and unhealthy people), and detailed shocks introduced for example by other diseases present and policies adopted by governments. The virus adopts reactions to these elements that reflect behaviour that has proved optimal for evolution to maximise surviving viruses: this maximand is its ‘utility’.

It may seem puzzling that a virus, lacking consciousness, can ‘respond’. However, this ‘response’ is simply the result of evolution in the behaviour of surviving mutations. Any given virus at one time will consist of many surviving strains, or mutated versions, each infecting in a different way. For example, we know that some versions spread quickly via asymptomatic ‘superspreaders’, who even after quarantine may be infectious. On the other hand other strains that hospitalise people tend to die out, as people either recover with strong antibodies, that kill the virus, or die. When people self-isolate, the virus stops spreading in the blocked channels but continues to spread via channels still open, such as superspreader chains. This is pre-programmed reactivity from the virus, picked up in our model as optimising behaviour.

Furthermore, we include in the model household agents, who also act strategically to avoid the costs the virus generates. In this element, underlined as a key one by Cochrane (2020), our work links with a large earlier literature on agents’ behavioural responses within epidemics, largely related to the AIDS virus- examples are Geoffard and Philipson (1996, 1997), Philipson and Posner (1993) and Kremer (1996). In this work, the authors could draw on surveys of individual behaviour as well as a wide range of cross-country and timeseries data, besides a good clinical understanding of the virus. Unfortunately such rich data is not yet available on the Covid-19 virus. Hence our agents here are treated as homogeneous in utility; though it is already known that the young and healthy are much less at risk than the old and unhealthy, we assume that all care about not getting

the disease, even the robust group if only because they could pass it onto to others less robust. Our 'representative agent' is a population-weighted average of the different types.

We must first go through the biology of susceptibility, infection and recovery, which is used in S-I-R models (Atkeson, 2020) usually with fixed parameters that define a mechanical progress of the virus. S people are those Susceptible to being infected. If infected, they become I people. Having been infected, they then after some time either die or develop powerful enough antibodies to kill the virus; or finally they may recover without killing the virus, so that the virus continues in them in a coexisting state, and they remain susceptible to further infection; those who die or recover and kill the virus are denoted R (Recovered) people. We then obtain the following relationships: $\Delta S_t = -\Delta R_t$; $\Delta R_t = \rho \Delta I_{t-d}$; $\Delta I_t = \psi_t S_t$. The moving parameter ψ_t is the rate at which susceptible people get infected on day t by the virus; ρ is the share of those infected that either die or recover, killing the virus.

Hence the virus' utility rises with the expected number infected who have not either died or killed off the virus in recovery. These represent all living clusters of the virus; so we assume it is aiming for as many living virus clusters as possible at any future point of time. As it is infinitely lived, with time preference and risk-aversion, it gives value to all these future clusters, discounted by its time preference and in logs, reflecting its risk-aversion (diminishing marginal utility of its 'consumption'). It plans on an infinite life, surviving to infect a future population that may be donated to it. We assume there is some cost of the speed of infection, r_t ; we think of this as due to increasing infection 'effort' which in turn represents the rising risk of policy resistance by the population the faster the infection rate, e.g. the faster development of vaccine or drugs, which will kill the virus. The biology of the actual infection speed implies that the higher the infected proportion of the population, the slower it is, and we add a term, ξ_t , which represents social reactions, and a policy intervention response to the infection, such as lockdowns. These interventions and the existing rate of infection increase the cost to the virus of achieving infection.

The usual epidemiological model treats infection rates as exogenous to the virus. It then introduces population characteristics, and calculates the interaction of the infection rates with these characteristics in an essentially mechanical way (Atkeson, 2020, surveys these S-I-R group models). In these models, the key parameter is the rate, ψ_t at which the infected I group who have not recovered or died (the R group) pass the virus on to the uninfected susceptible group, S; this para-

meter can be directly controlled by lockdown and other measures controlling people’s interactions. However, this is to treat the virus as unresponsive to circumstances, which would plainly endanger its survival chances. The optimising framework we use here endogenises infection rates, allowing the virus to respond in the best way for its ultimate success in surviving. Beenstock and Xieer (2020) point out there are large variations in contagion rates across countries and over time.

As we will see, in our model here, the contagion rate is affected by both known and unknown factors, responding to these as stochastic elements. Our approach allows us to estimate a complete structural model of virus behaviour, and test it powerfully against a reduced form of the data behaviour which we know to be a logistic curve process. By estimating model parameters and the exogenous shocks, we can identify, from different countries’ estimated behaviour, policy effects on death rates, and on the parameters of the virus’ response to the environment. This allows us to estimate the effects of a range of policy interventions — such as the huge variety adopted across many different countries — rather than simply those directly controlling people’s interactive behaviour.

We will focus on the growth rate of infection as the instrument chosen by the virus. The growth identity is $I_t = r_t I_{t-1}$, where I is the number of infected people and r is the gross rate of infection. This implies that $\Delta \ln I_t = \ln r_t$; since from the SIR relationships $\Delta \ln I_t = \psi_t S_t / I_{t-1}$, it follows that choosing $\ln r_t$ implies a simultaneous choice of ψ_t . Our most reliable data is on deaths, as reported cases suffer from downward reporting bias on actual. We assume the death rate, δ , is constant, with a constant lag of d days; so deaths $D_t = \delta I_{t-d}$, just as $R_t = \rho I_{t-d}$. We will use the model to predict deaths, the most reliable data we have, and compare the model predictions with the data behaviour of deaths (normally those that get hospitalised), using the indirect inference simulation method of testing and estimation. We will also estimate a reporting ratio of reported cases, C_t , to actual cases, I_t : $C_t = \Psi I_t$. We will also estimate Ψ so that reported cases can also be predicted.

Let the virus’ utility at the start of the infection be given by:

$$U_V = \sum_{t=0}^{\infty} \beta^t \left(\ln[I_t - R_t] - A r_t^\gamma \left(\frac{I_{t-1}}{POP} \right)^\mu \xi_t \right)$$

The second term in the utility function is the cost to the virus both of a higher infection rate

and of a rising lagged population share of infected people; both of these require the virus to work harder in terms of finding more cases to infect. ξ_t is a variable reflecting the varying infection-countering behaviour found in different sections of the population — which we will model shortly; $\left(\frac{I_{t-1}}{POP}\right)^\mu$ reflects special measures of protection taken by the government (lockdown etc.), as well as resistance rising with the expanding population share of infected people. t is days.

The virus maximises this utility subject to the infection state model above, viz $I_t = r_t I_{t-1}$, or in logs, $\ln I_t = \ln r_t + \ln S_{t-1}$. The number of days over which the epidemic lasts is infinite because its progress is asymptotic, never reaching full infection of the population. As noted above, the virus needs to survive and so there must be infected people carrying the live virus for ever.

The first order conditions are simply found by creating the Lagrangean, while substituting the model of Recovered into the utility function and noting that $\ln(I_t - R_t) = \frac{1}{1-\rho} \ln I_t - \frac{\rho}{1-\rho} [\ln \rho + \ln I_{t-d}]$; the virus maximises, with respect to R_t and I_t the Lagrangean

$$L = \sum_{t=0}^{\infty} E_0 \left\{ \beta^t \left[\frac{1}{1-\rho} \ln I_t - \frac{\rho}{1-\rho} [\ln \rho + \ln I_{t-d}] - A[r_t]^\gamma \xi_t \left(\frac{I_{t-1}}{POP} \right)^\mu \right] + \lambda_t [\ln I_t - \ln r_t - \ln I_{t-1}] \right\}$$

to yield:

$$0 = \frac{dL}{dr_t} = -\beta^t A \gamma r_t^{\gamma-1} \xi_t \left(\frac{I_{t-1}}{POP} \right)^\mu - \lambda_t r_t^{-1}$$

whence:

$$-\beta^t A \gamma r_t^{\gamma-1} \xi_t \left(\frac{I_{t-1}}{POP} \right)^\mu = \lambda_t$$

Secondly, the first order condition w.r.t I_t yields:

$$0 = \frac{dL}{dI_t} = \beta^t \frac{1}{1-\rho} I_t^{-1} - \beta^{t+d} \frac{\rho}{1-\rho} I_t^{-1} - \beta^{t+1} A [r_{t+1}]^\gamma \xi_{t+1} \mu \left(\frac{I_t}{POP} \right)^\mu I_t^{-1} + (\lambda_t - \lambda_{t+1}) I_t^{-1}$$

and so:

$$\begin{aligned}
0 &= \frac{dL}{dI_t} = \beta^t \frac{1}{1-\rho} - \beta^{t+d} \frac{\rho}{1-\rho} - \beta^{t+1} [r_{t+1}]^\gamma \xi_{t+1}^\mu \left(\frac{I_t}{POP} \right)^\mu A + (\lambda_t - \lambda_{t+1}) \\
&= \beta^t \frac{1}{1-\rho} - \beta^{t+d} \frac{\rho}{1-\rho} + \beta \frac{\mu}{\gamma} \lambda_{t+1} + (\lambda_t - \lambda_{t+1})
\end{aligned}$$

It follows that:

$$\left(1 - \left(1 - \beta \frac{\mu}{\gamma} B^{-1} \right) \right) \lambda_t = - \left[\beta^t \frac{1}{1-\rho} - \beta^{t+d} \frac{\rho}{1-\rho} \right]$$

where B^{-1} is the forward expectations operator leading the variable and keeping the expectations date given.

Hence

$$\lambda_t = - \left[\beta^t \frac{1}{1-\rho} - \beta^{t+d} \frac{\rho}{1-\rho} \right] / \left(1 - \left(1 - \beta \frac{\mu}{\gamma} \right) \right)$$

Now note that

$$\lambda_t = -\beta^t \gamma r_t^\gamma \xi_t \left(\frac{I_{t-1}}{POP} \right)^\mu A$$

then we have:

$$0 = -\beta^t \gamma r_t^\gamma \xi_t \left(\frac{I_{t-1}}{POP} \right)^\mu A = - \left[\beta^t \frac{1}{1-\rho} - \beta^{t+d} \frac{\rho}{1-\rho} \right] / \left(1 - \left(1 - \beta \frac{\mu}{\gamma} \right) \right)$$

or

$$\gamma r_t^\gamma \xi_t \left(\frac{I_{t-1}}{POP} \right)^\mu A = \frac{1}{(1-\rho) \left[1 - \left(1 - \beta \frac{\mu}{\gamma} \right) \right]} [1 - \rho \beta^d]$$

Finally in logs we obtain:

$$\ln r_t = \frac{1}{\gamma} \left\{ \ln \left[\frac{[1 - \beta^d \rho]}{(1-\rho) \left(1 - \beta \left[1 - \frac{\mu}{\gamma} \right] \right)} \right] - \ln \gamma - \ln \xi_t - \mu \ln \left(\frac{I_{t-1}}{POP} \right) - \ln A \right\} \quad (1)$$

and so using $\ln I_t = \ln r_t + \ln I_{t-1}$

$$\ln I_t = \frac{1}{\gamma} \left\{ \ln \left[\frac{[1 - \beta^d \rho]}{(1-\rho) \left(1 - \beta \left[1 - \frac{\mu}{\gamma} \right] \right)} \right] - \ln \gamma - \ln \xi_t + \mu \ln POP - \ln A \right\} + \left(1 - \frac{\mu}{\gamma} \right) \ln I_{t-1} \quad (2)$$

$$D_t = \delta I_{t-d}; R_t = \rho I_{t-d} \quad (3)$$

$$C_t = \Psi I_t \quad (4)$$

The model tells us that the daily infection rate responds inversely to the current self-isolation efforts of the population, ξ_t , and the existing (lagged) share of infected population, offsetting these in order to keep the costs of infection smooth over time, while still ensuring that the population gets steadily infected, ensuring new infections indefinitely.

We now insert household behaviour into the model. We will assume that household utility is reduced by infection but also by the personal inconvenience of avoiding infection by self-isolation activity, ξ_t . As this increases, the personal costs of not participating socially and economically rise directly with the extent of isolation, and rise indirectly the more uninfected people there are, as this lowers the personal risk of infection from participating, which raises the net costs of self-isolating (the economic costs net of the gain in lower infection risk). There is also a preference error, ϵ_t :

So

$$U_H = \sum_{t=0}^{\infty} \beta^t \left(-\ln I_t - \left[\xi_t \left(\frac{POP}{I_t} \right)^\phi \epsilon_t \right] \right)$$

Households maximise this utility with respect to ξ_t subject to the virus' behaviour set out above.

Hence its Lagrangean is

$$L_{H0} = \sum_{t=0}^{\infty} E_0 \left(-\nu_t \left[\ln I_t - \frac{1}{\gamma} \left\{ \ln \left[\frac{\beta^t \left(-\ln I_t - \left[\xi_t \left(\frac{POP}{I_t} \right)^\phi \epsilon_t \right]}{(1-\rho)(1-\beta[1-\frac{\mu}{\gamma}])} \right] - \ln \gamma - \ln \xi_t + \mu \ln POP - \ln A \right\} - \left(1 - \frac{\mu}{\gamma} \right) \ln I_{t-1} \right] \right) \right)$$

Going through analogous Lagrangean steps to find the first order conditions yields from $0 = \frac{dL}{d\xi_t}$:

$$\nu_t = -\beta^t \gamma \xi_t \left(\frac{POP}{I_t} \right)^\phi \epsilon_t$$

and from

$$\begin{aligned}
0 &= \frac{dL}{dI_t} = -\beta^t I_t^{-1} + \beta^t \phi \xi_t \left(\frac{POP}{I_t}\right)^\phi \epsilon_t I_t^{-1} - \nu_t I_t^{-1} + \left[1 - \frac{\mu}{\gamma}\right] \nu_{t+1} I_t^{-1} \\
&= -\beta^t + \beta^t \phi \xi_t \left(\frac{POP}{I_t}\right)^\phi \epsilon_t - \nu_t + \left[1 - \frac{\mu}{\gamma}\right] \nu_{t+1} \\
&= -\beta^t - (\phi/\gamma) \nu_t - \nu_t + \left[1 - \frac{\mu}{\gamma}\right] \nu_{t+1} \\
&= -\beta^t + \left[1 + (\phi/\gamma) - \left[1 - \frac{\mu}{\gamma}\right] B^{-1}\right] (-\nu_t) \\
&= -\beta^t + \left[1 + \phi/\gamma - \beta \left[1 - \frac{\mu}{\gamma}\right] B^{-1}\right] \{\beta^t \gamma \xi_t \left(\frac{POP}{I_t}\right)^\phi \epsilon_t\} \\
&= -1 + \left[1 + \phi/\gamma - \beta \left[1 - \frac{\mu}{\gamma}\right] B^{-1}\right] \{\gamma \xi_t \left(\frac{POP}{I_t}\right)^\phi \epsilon_t\}
\end{aligned}$$

so that: $\gamma \xi_t \left(\frac{POP}{I_t}\right)^\phi \epsilon_t = 1 / \left\{1 + \phi/\gamma + \beta \left[1 - \frac{\mu}{\gamma}\right]\right\}$

$$\ln \xi_t = -\ln\left\{1 + \phi/\gamma + \beta \left[1 - \frac{\mu}{\gamma}\right]\right\} - \ln \gamma - \ln \epsilon_t - \phi \ln POP + \phi \ln I_t$$

When this is substituted into the infections equation we obtain:

$$\begin{aligned}
\ln I_t &= \frac{1}{\gamma + \phi} \left\{ \ln \left[\frac{[1 - \beta^d \rho]}{(1 - \rho) \left(1 - \beta \left[1 - \frac{\mu}{\gamma}\right]\right)} \right] + \ln\left\{1 + \phi/\gamma + \beta \left[1 - \frac{\mu}{\gamma}\right]\right\} + (\mu + \phi) \ln POP - \ln A + \ln \epsilon_t \right\} \\
&\quad + \left(\frac{\gamma - \mu}{\gamma + \phi}\right) \ln I_{t-1}
\end{aligned}$$

The model is fitted to deaths, D . Unfortunately, we do not have data on the actual infections, I , because tests have not been good enough to estimate these reliably. However, the model gives us estimates of total infection rates, the death rate, infection growth rates and the reporting ratio that are consistent with the actual data. We report these in section 5.

Equations (1)-(4) constitute the structural model of the virus' behaviour, from the optimal first order conditions and the state variables' evolution. The intuition is that as the infected population share gets higher, infection becomes harder and the infection rate drops.

For simulation we can extract the random variable ϵ_t from the data and equation (2). In practice we will use deaths in place of infections, applying a lag on deaths, D_t , of 21 days from infection; so $D_t = \delta I_{t-d}$ where $d = 21$, and δ , the death rate, must be estimated. We will also estimate the

reporting rate, Ψ , of infections, so that reported cases, $C_t = \Psi I_t$.

4 Estimation Methodology and data applications

Survey data is available on the numbers in total infected by the virus in the UK and Sweden, according to antibody tests which check whether people were infected two to three weeks before, this being the period to antibody production. This data combined with data on deaths gives us a strong estimate of the IFR, a key parameter of the model.

One widely-held hope among virologists opposed to lockdown, such as Prof Carl Heneghan at Oxford and Anders Tegnell the state epidemiologist in Sweden was that a majority of the population had contracted the virus without getting more than weak symptoms. This would imply that there was close to herd immunity. This hope seemed to have been dashed by available surveys of specific Covid-19 antibody prevalence in several countries, which turns out to be low, in the range of 5%–7%, in the UK, in Sweden and in Spain, with big cities like London, Stockholm and Madrid reaching 20% or less. Outside big cities large numbers of small areas have had prevalence close to zero. However, the latest medical research finds that only seriously infected people develop antibodies² and that another 40-60% of the population already have general immunity to coronaviruses³ and so may have repelled weak infections. In the UK, in addition to these ONS surveys of those currently infected (about 2%) and of those with antibodies specific to Covid-19 (about 7%), an online survey running since end March at King’s College London, and known as ZOE (the name meaning ‘life’ in Greek, ZOE is a public company backed by King’s College researchers) has recorded those with symptoms of Covid-19; to date it has found a cumulative total of 60% of the population reporting symptoms, mostly fairly weak ones. Total UK deaths had reached (at the time of writing, summer 2020) around 41500, 620 per million — mostly in hospitals and care homes. Recent medical research, as just noted, finds that only those with serious infections develop specific antibodies, while there is general immunity in the population against coronaviruses which can protect against weak infections.

We also find in the latest medical research that those mildly infected, as well as their uninfected family members, develop T-cell responses which are used by the immune system to fight the virus. Thus they have significant immunity to Covid-19, even though they do not develop specific

²<https://www.biorxiv.org/content/10.1101/2020.05.21.108308v1.full>

³<https://www.cell.com/action/showPdf?pii=S0092-8674%2820%2930610-3>

antibodies; furthermore this T-cell immunity is long lasting, over several years, as it seems T-cells are kept in the immune system's memory for long unlike antibodies. At the time of writing, the details of this are not clear, awaiting more research. However it now seems that the 60% of the UK population identified by ZOE as having had the infection, have also probably acquired T-cell based immunity against a second wave. If one adds this 60% to the 7% badly infected who acquired antibodies (in addition presumably to T-cells), it suggests that the UK may well now have nearly 70% who have been infected, and who also have long term immunity (i.e. against a second wave), implying herd immunity of the basic sort, namely that not requiring any social responses.

How should we take account of this extra data? It would seem that the official PHE etc. reported data very greatly underestimates the true extent of infection. Furthermore, the antibody-based infection estimates greatly underestimate those with immunity, especially with long-term immunity (i.e. against a second wave). By implication the IFR is much overestimated by the official figures. On the latest figures of deaths vs total so far infected it is 0.1%.

Before the results of the antibody surveys were announced, figures like this were our best estimates based on our modelling and other scraps of data. The antibody tests seemed to destroy them. However, this later data and research suggests they were roughly correct after all. It would seem we are dealing with a population with widely differing pre-existing immunity and resistance to the Covid-19 disease. Many, the vast majority, rapidly developed defences so that they were only weakly symptomatic or even asymptomatic or nearly so. Others, a minority, were badly infected and a proportion of these died.

The data from ZOE records those reporting symptoms of Covid-19 and the ZOE team's estimates of the percent infection rate in the population based on their sample of reports and swab tests taken by reporting people. From these estimates it is possible to estimate daily new cases from the identity $\Delta Infections_t = NewCases_t - RecoveringCases_t$. To obtain the total who have ever been infected at date t we can sum all $NewCases$ up to and including t . Due to recent developments and having access to more data ZOE has recently changed its $Infections$ series from the 11th June. In Figure 2 this change can be seen as a slight drop in the number of people infected. Accumulating the number of people infected we get 40 million as the current number who have been infected — approximately 60% of the population. Adding the 7% estimated to be seriously infected brings the total infected to nearly 70%.




Figure 2: ZOE: number of people calculated to have COVID symptoms

Accordingly, we estimate our main model for the UK, benchmarked to this estimate of the total who have ever been infected; we find a good fit of the model with about 60% or more of the population being predicted to be infected in the long run. We report below the results for the UK on this basis, as ‘benchmarking to ZOE data’. This model represents our best estimate of what is going on in the UK with all Covid infections, recoveries and deaths.

In addition, we have estimated a model for serious infections alone, together with deaths stemming from these (of course all deaths come from serious infections). As we have seen above, serious infections in the UK have accumulated to around 7% of the population; and this is also the case in Sweden. In this model the benchmarked IFR implied relative to (serious) infections is of course higher than in our model benchmarked to the ZOE data of all infections. If we take a figure of 7% for end-May in both the UK and Sweden for total serious accumulated infections, this would imply an IFR of 0.0054 in Sweden, and one of 0.0083 for the UK. We calibrate our models with these two rates, and search for estimates in the region of these IFRs. We report these results below for the two countries as ‘benchmarking to serious infections only’. These comparative results allow us to draw some policy lessons from the comparison.

In a final exercise we model the latest evidence on second waves. On the basis of recent data, 28 countries have experienced second waves of infection, following the dying off of a first wave.

The countries fall into two groups, those with a small second wave, ‘small’ group, and those with a large one, the ‘large’ group. We have used our methods here to estimate logistic curves for these waves for the average of each country group, small and large; we also estimate a structural model for each. Our aim is to compare the behaviour of each group for clues to the policies conducive to a small second wave.

5 The Indirect Inference estimation and testing method

To test and estimate the model we use the method of Indirect Inference (for a detailed overview, see Le et al, 2016). The method involves measuring how close the simulated model is to the actual data, as in the familiar method of simulated moments. To do this it uses a descriptive model to capture the behaviour of the data, one possible such description being moments; this can also be, as here, the reduced form of the model, whose coefficients can then be used as the descriptors, the ‘auxiliary model’. We then measure how close the coefficients found in the actual data are to the mean of the coefficients found when estimated from the model simulations. If those in the actual data are within the 95% bounds of the distribution of those estimated on the simulated data then the model is not rejected and can be considered a good causal account of the data. When estimating the model we vary the structural parameters to find that set with the highest p-value, which is the closest fit to the data. In this paper we use the logistic function to describe the data. The logistic function is of the form:

$$f(x, a, b, c) = \frac{c}{1 + e^{-(x-b)/a}}$$

where x is time, and the three parameters are:

a is the infection speed

b is the day when the maximum number of new infections occurred

c is the total number of recorded infected people at the end of the infection

It is these three parameters that we will try to match in the Indirect Inference estimation procedure. A detailed description of the Indirect Inference method is in Appendix 1.

To evaluate the reliability of this indirect inference test, we can use a Monte Carlo experiment in which we check the power of the test as the structural model parameters deviate in accuracy

from the true ones. In the experiment we assume a model of the type here to be true and generate a large number of simulated samples from it. We then deliberately falsify the model parameters by some percentage, x , and see what percentage of these samples would reject the model. This gives us a measure of the test power — i.e how frequently the test will reject a model of given falsity.

For our virus model we have found the results shown in the following Table 6. The rejection rate rises steadily as the model parameters are falsified. It reaches 100% as $x\%$ reaches 25%. This implies our model is likely to lie not too far from the truth and so the test has reasonable power.

Falseness	Rejection Rate
0	5.00
3	5.84
5	10.14
7	15.68
9	25.12
15	73.60
20	97.49
25	99.88

Table 1: Monte Carlo: Power Check

6 Estimated results for the UK and Sweden for model benchmarked to serious infections only

To understand how the model works, consider the hurdles faced by the virus, all of which are reflected in its utility function. First, there is the death rate, inherited from its evolution through whatever species it has inhabited. Second, there is γ , the measure of how far speed of infection provokes increasing resistance from people with increasing immunity. This parameter is largely set by the population structure — the proportions of types, such as by age, fitness and existing health — since the faster the infection rate, the higher the proportion of people with immunity that the virus will be attempting to infect. Third, there is μ , measuring how far the proportion of uninfected people in the population stimulates the rate of spread. This is policy-related, in that targeting or lock down arrests the spread to the uninfected. Fourth, there is the household parameter, ϕ , measuring how households react to the risk of getting the virus by self-isolating, social distancing, hygiene etc. Finally, there is the constant, which, when divided by $\mu + \phi$, measures

the population proportion that will eventually be infected. This is partly related to population structure, partly to government policy and household reactions in stopping the spread via lockdown, track/trace/isolate, and self-isolation. We should note that $\gamma, \mu,$ and ϕ are identified separately by entering different utility functions and being attached to differing variables.

These factors determine the speed with which the virus spreads and also the extent to which it will spread in the end. The model is matched to the logistic data behaviour of deaths for the UK, Sweden and a global average of 25 countries with a large number of deaths. The model estimated parameters are shown in Table 2.

	UK	Sweden	Global
δ	0.0084	0.0052	0.0015
μ	4.11	0.151	2.55
γ	59.53	40.59	79.02
ϕ	0.17	2.95	0.62
$(\mu + \phi)/(\gamma + \phi)$	0.07	0.07	0.04
Constant	-11.94	-8.605	-10.36
% Population Infected to Date	7	7	7
% Population Infected Long Term	7	7	7
Reported/Actual Infections (inverse)	0.0499(20)	0.0442(23)	0.0337(30)
P-value	0.93	0.82	0.70

Table 2: Structural Model Parameter Estimates

The viral rate of spread depends directly on how many are uninfected by the joint parameter, $(\mu + \phi)/(\gamma + \phi)$, which measures the stimulus of the uninfected population share (reducing lockdown and reactivity) relative to the resistance from population immunity and reactivity as infection increases. The higher this measure, then when many in the population are uninfected, the spread is faster. The measure is similar in both countries. Hence in both the virus spread fast, and has by now infected about 7% of the population according to the model. Effectively lockdown and social resistance are close substitutes in their effect on virus prevalence.

The death rate is lower in Sweden, at 0.0052 vs 0.0084 in the UK. This lower death rate is presumably associated with general public health policies that were more effective in protecting vulnerable groups with the high death rates; the UK's problems with personal protective equipment supplies in hospitals and with care home conditions have been well publicised.

Both the models predict the number of (serious) infections at around 7% of the population long

term — low prevalence in line with the latest surveys of antibody presence. Our results fit well statistically as can be seen. The match of the model to the logistic estimates is good with p-values (the probability that the data does not reject the model); 0.82 for Sweden and 0.93 for the UK.

Nevertheless there is randomness and uncertainty at work. The error term, ϵ_t , measures the variability in the model’s behaviour, which comes from biological and other (mostly policy) shocks to the rate of infection. The consequence of these shocks for the behaviour of deaths can be seen in the estimated shocks, ϵ_t , and the resulting illustrative simulated histories below for the UK, which vary substantially and are far from the smooth progressions imputed by the logistic curve. Faced with spikes like these, it is not surprising that governments were driven to use drastic lockdowns to make sure of suppression.




Figure 3: Innovations for the ϵ_t process and Resulting Illustrative Simulations of Deaths

	Actual	Lower 2.5%	Upper 2.5%	Mean
a	10.2775	2.2953	20.8867	8.3010
b	47.4228	30.8611	77.5785	46.2417
c	38738.1847	18601.3551	89002.1526	40782.6827

Table 3: The auxiliary model estimates and bounds are for the logistic curve, as fitted for UK data.

Fitting a logistic function to the deaths data results in the parameters shown in Table 3. The bounds shown come from the simulated variation from the structural model — not from the logistic estimates on the data, which are rather tightly estimated, as listed in section 2. They are indicating that a wide variety of logistic models could emerge from the structural model with some probability. The logistic model estimated on the UK deaths data is highly probable according to the UK structural model with a p-value of 0.93, as we have seen.

We now go on to consider the implications for policy of the causal model estimates.

6.1 Policy implications of the comparative model estimates benchmarked to serious infections

We can now discuss the experience of the two countries and the different estimates we get from them for these factors. From this we can learn a fair amount about the effectiveness and costs of different government policies. Our main focus in the policy discussion that follows is on the UK, using Sweden as the main identifying benchmark, for outcomes of alternative policies, of no lockdown but instead information and advice on social distancing, together with other general public health policies.

Comparing the UK and Sweden we find that the parameter of natural resistance to the virus' rate of progression (γ) is much the same; but the Swedish IFR is substantially lower. This will be related to the effectiveness of controlling the access of the disease to vulnerable groups, like the ill and elderly; the better the protection against infection within hospitals and care homes, the less this access. In the UK, problems with PPE in the NHS and care homes have been well publicised.

Also the Swedish ϕ estimate, reflecting social reaction, is much higher than in the UK, where it is close to zero, while the μ parameter reflecting government-imposed policies like lockdown is around zero in Sweden, much lower than in the UK. These two parameters are of course close substitutes, since the social reaction compensates for lack of policy reaction.

Our interest in policy lies particularly in the effect of the UK lockdown. According to our model, this is found in the policy-reaction parameter μ . However, as we have seen, the higher this parameter the lower the social reaction parameter ϕ ; there is strong substitution. It is the two together that determine the equilibrium progress of the virus, both its end infected share of the population and its rate of spread. The model suggests that there is no difference in the behaviour of the virus between the two economies. We can illustrate this from the almost identical paths of actual deaths shown in Figure 4.

The implication is that lockdown achieved nothing extra compared with what a decentralised social reaction strategy, as pursued in Sweden, would have achieved. We could assume that the Swedish and UK relative policy costs are reflected in the relative Consensus Forecasts made in September 2020 for the fall in their 2020 GDP: for Sweden this is about 4.3%. but for the UK it is 10.1% in 2020, 5.8% more, or about £116 billion. According to the model the UK lockdown saved

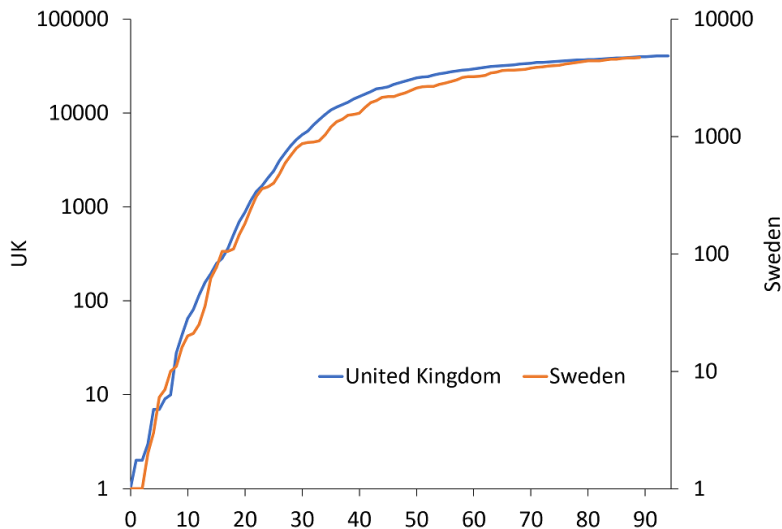


Figure 4: Cumulative Confirmed Deaths (log scale)

no deaths but cost the UK economy a GDP loss of over £100 billion. Plainly relying on social responses as in Sweden would have been far more cost-effective than lockdown, for much the same outcome in deaths. We do not need to appeal to any cost per life saved as in transport policy where typically £11 million per life saved⁴ is used as a benchmark; the point is that lockdown has cost a lot for no lives saved at all.

However, the Swedish experience suggests other policies there, of a general public health nature, succeeded in reducing deaths by lowering the death rate. Had Swedish health policies been applied in the UK, the population infected would have been the same but the death rate 0.3% lower at 0.52%. UK deaths would have been about 24000, some 16000 less than the total at the time of writing.

UK policy in the first wave aimed to lift the lockdown quickly but introduce stringent test/trace/isolate policies of localised lockdown, stopping localised outbreaks fast. However, the Swedish experience suggests that decentralised social reaction will do the same job without this heavy-handed government action. All that the government needs to contribute is any information it can provide, such as from surveys and local hospital reports: the people will do the rest, including sheltering

⁴The J-value (cost per life saved by safety measures) used by the US Dept of Transportation is between £4 and £10 million per life saved through road safety measures, with the UK value being around £9 million. <https://www.bristol.ac.uk/media-library/sites/policybristol/PolicyBristol-Report-April-2018-value-human-life.pdf>

the most vulnerable. Given that there are random shocks to the model, we can think of these as random starts of mini-waves; however, they provoke social reactions which bring the case and deaths outcomes back to the equilibrium path.

7 Estimation results for the UK based on ZOE benchmarking

We now turn to estimates for the UK based on the ZOE benchmark for total infections. The main difference this creates compared with our previous benchmarking on serious infections only, is that the implied IFR is considerably lower, at around 0.1%. Again, we fit our structural model to deaths, searching for the best estimates with an IFR close to this. Not surprisingly, since the model is matching the deaths logistic curves as before, the estimated parameters are close to the previous ones. The main difference lies in the constant and the IFR, given that we are benchmarking the model to generate total infections in line with ZOE and also total deaths in line with those reported as can be seen in Table 4

	UK
δ	0.000975
μ	5.14
γ	50.9
ϕ	0.14
$(\mu + \phi)/(\gamma + \phi)$	0.103
Constant	-2.80
% population infected long term	65
Reported/Actual Infections	0.0069
P-value	0.22

Table 4: Structural Model Parameter Estimates

8 Estimation results: modelling the dynamics of second waves

There is by now substantial evidence about second waves, since 28 countries have experienced one. We can compare the logistic curve estimates for the two waves, and attempt to find the policy implications for containment of the second.

In these second waves, the same virus has attacked the same population, having already completed a first attack. Evolutionary biology tells us that two main things could have changed between

the two waves. First, all organisms come in numerous copies or mutations, from which natural selection weeds out the least fit to survive; hence the surviving virus mutations in the second wave will exclude those that died in the first, whether by killing their host or by being killed by their host recovering from a nasty bout of covid-19 which produced killer antibodies. Second, the population being attacked in the second wave may have more or less immunity than the one attacked in the first; probably more since it will include those who survived from the first wave, while it will no longer sadly include those who did not. So, to sum up it is likely that the virus has ‘weakened’ and the population has ‘strengthened’.

We divide the 28 countries with second waves into two groups: 12 with a big second wave relative to the first (the ‘Large’ group), and 16 with a small one (the ‘Small’ group). The Large includes the US, Japan and Vietnam; the Small Germany, Netherlands and Singapore. What distinguishes the second group is that it had a longer first wave (typically four months against three for the other), and seems to have had a good localised test-trace-isolate system operating second time around, if not also in the first wave. What all 28 countries have in common is a rapidly falling death rate per reported case. This supports the idea of a weaker virus meeting a stronger population. It could also be due to better treatment, but this had already occurred by the end of the first wave, where the cumulative death rate came down impressively to 10% from a peak of 40%; the fall in the second wave has however been proportionately much bigger, to about 3%. This can be clearly seen in Figure 5, showing the average cumulative death rate for all 28 countries in each wave.

If we use this data to estimate the underlying causes at work, we estimate a much larger ‘social reaction’ response for the small second wave group than for the other, confirming that they had better policies for ‘whack-a-mole’. These we now know involve local area officials providing good local guidance for local behaviour.

The differences between the Small and Large second wave groups is shown in chart form in Appendix 2. What the first group share is a long-lasting first wave and a policy of localised test-trace-and-isolate after it. All their logistic curves lie below that of the first wave; they all have a lower c , so they reach a lower cumulative total. Mostly they also have a lower b implying that infections peak sooner in the wave: and a lower a , so they fall faster. The second group of twelve tend to the opposite. All their logistic curves lie above that of the first wave; they all have a higher c , so they reach a higher cumulative total. Mostly they have a larger a , so infections fall off more

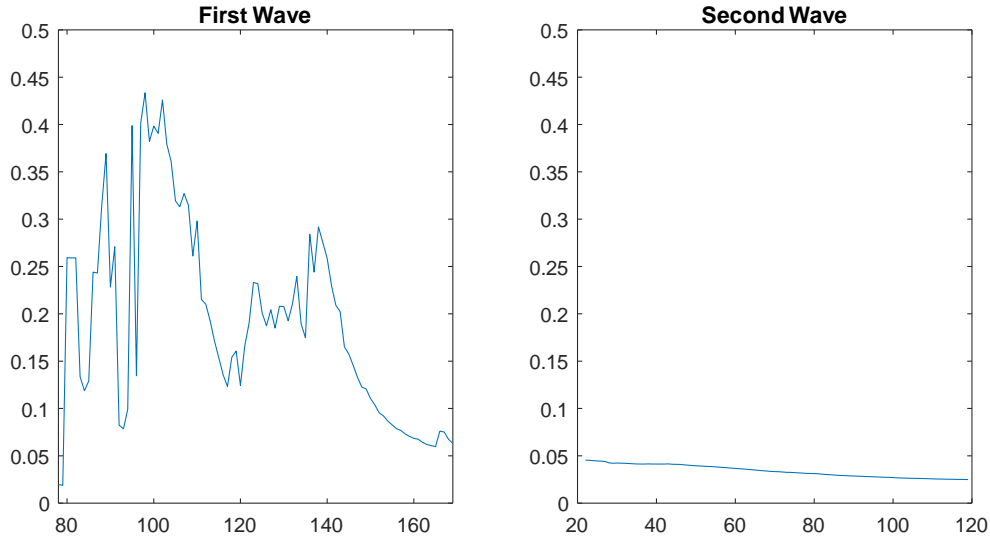


Figure 5: Cumulative Deaths/Cumulative Cases 21 days before, since start of wave. Simple average for 28 countries with second waves

slowly. They also tend to have a lower b implying that infections peak sooner in the second wave.

To discover more of what propels these two different second wave trajectories, we estimated a full structural model for the average Small second wave group and the average Large group country, using their average logistic curve behaviour as the matching criterion — the auxiliary model. We estimated their (‘augmented’) behaviour over the two waves of deaths combined, since the second wave alone furnishes too little data. The data behaviour reveals that the augmented curve for the Small group resembles closely that for its first wave alone (Figure 6) — since only limited extra cases occur in the second wave. However, for the Large second wave group, the augmented behaviour is heavily dominated by the large second wave ‘tail’ (Figure 7). Accordingly the model estimates for the Large group, compared with the Small group as shown in Table 5, show weaker government (μ) and social reaction (ϕ) parameters, together with weaker general immunity (γ). Furthermore, the Small group has a longer first wave than the Large group: its length of first wave is around 4 months versus around 3 months for the Large group. In sum, the results suggest that the source of the better outcomes for the Small group second wave is a combination of more immunity in the population (perhaps from more active sheltering of those more at risk), more government reaction and more social reaction, together with a longer exposure to the first wave (perhaps giving more

immunity in the second wave).




Figure 6: The Augmented Small Second Wave Deaths




Figure 7: The augmented large second wave deaths

9 Conclusions

In this paper we have fitted the logistic function, the reduced form of epidemic behaviour, to the data for deaths from Covid-19, for a wide variety of countries, with a view to estimating a causal model of the covid virus' progression that can match this logistic behaviour. We then set out a structural model of the Covid virus behaviour based on evolutionary biology and social household behaviour; we estimated and tested this by indirect inference, matching its simulated logistic behaviour to that found in the data. In our model the virus' progression depends on the interaction of strategies by household agents, the government and the virus itself (as programmed by evolution). Within these

	Small second wave-augmented	Large second wave-augmented
δ	0.000194	0.00107
μ	8.5	3.62
γ	82.7	49.0
ϕ	31.47	17.15
$(\mu + \phi)/(\gamma + \phi)$	0.35	0.31
Constant	-11.18	-10.27
% population infected long term	75	38
P-value	0.96	0.17

Table 5: Structural Model Parameter Estimates For the Second Wave

interactions, it turns out that there is substitution between government topdown direction (such as lockdown) and social reaction to available information on the virus' behaviour. We examined this substitution particularly in the case of the UK and Sweden, which we chose because they followed different policies, especially on lockdown. Our basic policy finding was that the general public health policies pursued in Sweden with no lockdown were more effective in reducing deaths than UK public health policies plus lockdown; and that the UK lockdown was no more effective in reducing deaths than the Swedish reliance on voluntary socially aware behaviour, whereas the economic cost of the UK policy was enormously bigger. We also looked at experience of second waves, where we found that countries successfully limited second waves when they had had longer first waves and followed policies of localised reaction in the second.

References

- [1] Atkeson, A., (2020) 'What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios', NBER Working Paper No. 26867, March 2020.
- [2] Batista, M. (2020) 'Estimation of the final size of the coronavirus epidemic by the logistic model'. Working Paper.
- [3] Beenstock, M. and Xieer, D. (2020) 'The natural and unnatural histories of Covid-19 contagion', Covid Economics: Vetted and Real-Time Papers, Issue 10, pp. 87–115. CEPR.

- [4] Cochrane, J. (2020) 'An SIR model with behaviour', <https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html>
- [5] Geoffard, P., & Philipson, T. (1996). Rational Epidemics and Their Public Control. *International Economic Review*, 37(3), 603-624. doi:10.2307/2527443
- [6] Geoffard, P., & Philipson, T. (1997). Disease Eradication: Private versus Public Vaccination. *The American Economic Review*, 87(1), 222-230. Retrieved from www.jstor.org/stable/2950864
- [7] Golinski, A. and Spencer, P. (2020) 'Coronometrics: The UK turns the corner', University of York Discussion Papers in Economics No. 20/04.
- [8] Kremer, M. (1996). "Integrating Behavioral Choice into Epidemiological Models of AIDS," *The Quarterly Journal of Economics*, Oxford University Press, vol. 111(2), pages 549-573.
- [9] Le, V. P. M., Meenagh, D., Minford, P. and Wickens, M. (2011) 'How Much Nominal Rigidity is there in the US Economy? Testing a New Keynesian DSGE Model Using Indirect Inference', *Journal of Economic Dynamics & Control*, Volume 35(12), December 2011, pp 2078-2104.
- [10] Le, V. P. M., Meenagh, D., Minford, P., Wickens, M. and Xu, Y. (2016). 'Testing Macro Models by Indirect Inference: A Survey for Users', *Open Economies Review*, Volume 27(1), pp. 1-38.
- [11] Nesse, R. M., Bergstrom, C. T., Ellison, P. T., Flier, J. S., Gluckman, P., Govindaraju, D. R., Niethammer, D., Omen, G. S., Perlman, R. L., Schwartz, M. D., Thoman, M. G., Stearns, S. C., Valle, D. (2010) 'Making evolutionary biology a basic science for medicine', *Proceedings of the National Academy of Sciences* Jan 2010, 107 (suppl 1) 1800-1807; DOI: 10.1073/pnas.0906224106
- [12] Philipson, T. and Posner, R. (1993) *Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective*. Harvard University Press

10 Appendix 1: How indirect inference is carried out

Suppose we have hit on a way to describe the data: this description is the auxiliary model. Here it is the known model describing the data behaviour of an epidemic: the logistic curve model set out earlier. Remember that we may want to come back to this choice if we find that it implies too much or too little power; we will want to do a Monte Carlo experiment with our model and our data to check this — as we will explain below.

Thus we have our sample of data and our auxiliary model describing its behaviour succinctly. We also have our structural model which we propose to test. This model contains error terms, ‘shocks’ (here ϵ_t) which follow some simple univariate time-series model, we will assume. We begin by specifying the model with a set of numerical parameter values which could be taken from micro-studies or otherwise ‘calibrated’ (meaning that their values reflect some previous idea we or others have had about them). Sometimes when no such values can be found the model can be estimated in some way, e.g. by Bayesian means, to find a set of such values to start with. Once these structural parameter values have been set, we have the model in a testable quantitative form. The final step is to deduce the errors, given the data and the parameters. Where there are no expectations terms in the structural model equations, as in this model, the error is simply whatever makes the equation add up properly, given the data for the variables in the equation.

We now have a model with some estimates of errors and by simple estimation also their univariate time-series processes. The model to be tested now is fully specified quantitatively and is, through its errors, consistent with the data. Notice that after the errors’ time-series processes have been estimated we have errors consisting of lagged effects (from past errors) and ‘innovations’ or ‘shocks’ (current events assumed to be i.i.d.). This model is a description of a causal world in which shocks affect agents’ behaviour given the existing optimised rules of behaviour these agents are following; and these shocks generate ‘impulses’ to variables in the economy. The implication is that any sample behaviour is the result of shocks that occurred randomly. In other words history would have been different had different shocks occurred.

The empirical indirect inference test can be thought of as a way of rewriting history many different times for this sample period. Our model is the history-generating construct — the causal engine of history. We can ask what histories it could have produced under randomly varied shock

combinations. We know that by construction one of these histories is the actual one that happened and that we have as our sample. But the question is: how probable would it have been had this model been the causal engine?

We can discover this probability by repeatedly drawing sets of shocks to create many different histories, each of which implies a different data behaviour — a process known as ‘bootstrapping’ to obtain ‘simulated samples’ (each sample being one of our histories). Going to our auxiliary model we estimate it on each of these histories and record the values of the auxiliary model coefficients. This gives us a distribution of possible coefficient values.

Take a simple example of an auxiliary model of the economy which records two descriptive coefficients only, the autocorrelation of interest rates and that of inflation. Figure 8 illustrates the joint distribution of these two autocorrelations coming from our structural model simulations.




Figure 8: Bivariate Normal Distributions with correlation of 0 and 0.9. Two possible data points shown: $x=0.1, 0.9$ and $y=0.0$

One can see by inspection that the top distribution is one where the two autocorrelations are quite unconnected across different sample histories, creating a rounded mountain. This means that when, for example, inflation is very persistent over a history, interest rates have no tendency to be more or less persistent than normal. Now in the bottom distribution one can see by contrast that the two autocorrelations are closely correlated, creating a ‘ridge-like’ mountain.

Both these mountains are possible in principle from different structural models. However, consider a typical causal macro model of the economy: any such model will have an equation

linking interest rates to inflation because the nominal interest rate is equal to the expected real interest rate plus expected inflation by definition. For simplicity assume expected real interest rates vary little. Now consider two samples: one where inflation is highly persistent and one where it is not persistent at all. In the first expected inflation and inflation will be highly correlated, so inflation will also be persistent and so therefore will interest rates due to its expected inflation component. In the second, inflation and expected inflation will not be correlated; inflation will have no persistence and expected inflation will always be zero. Interest rates will not vary due to expected inflation, only due to real interest rate variation which is minimal; hence interest rates will have no tendency towards persistence due to inflation persistence. We see from these two samples that the autocorrelations of inflation and interest rates will be linked; when inflation is persistent so will interest rates be, when it is not, so too will they not be.

This example illustrates a general finding in causal models: their equations create tight linkages between variables so that these variables' behaviour is also closely linked. Thus most model distributions of pairs of auxiliary model coefficients look like the bottom panel of Figure 8, with a ridge-like shape indicating a high covariance across different samples. In our virus model here there is link between the speed of infection, which shows up in a , and the date of maximum infection speed, b , which will come sooner (a lower b) the faster the infection rate. In a sample with a high a , b will tend to be low..

Now look at the two dots, red and blue, on Figure 8. These are two examples of data from the actual historical sample that we might have. The red dot shows two autocorrelations that are both low (zero in fact). The blue dot shows one that is very high and one that is very low.

The rising height of the mountain shows the rising probability of encountering the autocorrelation combinations. One can see that on the top mountain the blue dot has some probability, the red dot has none. But on the second mountain, the red dot has some probability and the blue dot has none. This is because of the covariance. When it is low, the distance between the two autocorrelations, which is high on the blue dot, does not matter so it does well, and the red dot badly (as on average it is a long way from the mean); when the covariance is high, the distance matters and the low distance red dot does well as a result.

Since autocorrelations are averages over a whole sample, the Central Limit Theorem implies that their distribution will be jointly normal. The likelihood, L , of any combination, α_s , is given

by the height of the mountain and its formula is:

$$L(\alpha_s) = \frac{1}{(2\pi)^k \left| \Omega(\alpha[\hat{\theta}]) \right|} e^{-0.5[a_s - \alpha_S(\hat{\theta})]' \{ \Omega(\alpha[\hat{\theta}]) \}^{-1} [a_s - \alpha_S(\hat{\theta})]}.$$

where k is the number of simulated auxiliary model coefficients, α_s (two in this illustration). Ω is the variance-covariance matrix of these coefficients as simulated by the DSGE model (with coefficients θ). The exponent is (-0.5 times) the Wald statistic (IIW) based on the bootstrap distribution (implied by the assumed DSGE model coefficients $\hat{\theta}$) of a_S around their bootstrap means, $\alpha_S(\hat{\theta})$: $[a_s - \alpha_S(\hat{\theta})]' \{ \Omega(\alpha[\hat{\theta}]) \}^{-1} [a_s - \alpha_S(\hat{\theta})]$. This has an approximate Chi-squared distribution (with k degrees of freedom). $\alpha_S(\hat{\theta})$ is the mean of α_S across all the histories, which are all functions of θ , the parameters of the structural model (including the error time-series parameters). It follows that the lower the Wald statistic for a given data combination α_T given by Wald

$$(\alpha_T) = [a_T - \alpha_S(\hat{\theta})]' \{ \Omega(\alpha[\hat{\theta}]) \}^{-1} [a_T - \alpha_S(\hat{\theta})]$$

the higher the likelihood of the combination,

$$L(\alpha_T) = \frac{1}{(2\pi)^k \left| \Omega(\alpha[\hat{\theta}]) \right|} e^{-0.5[a_T - \alpha_S(\hat{\theta})]' \{ \Omega(\alpha[\hat{\theta}]) \}^{-1} [a_T - \alpha_S(\hat{\theta})]}$$

One can see that the Wald and the likelihood are related to the distance between the data combination and the model-simulated mean coefficient.

From here on the Indirect Inference (II) procedure is simple in principle but in practice requires a lot of computer time. One records the Wald statistic for the originally calibrated DSGE model. Probably it will be rejected due to the bias in the estimation procedures that have been used (such as FIML) or due to there being no estimation at all. At this point the researcher needs to define the limits theory places on the DSGE parameters and search for a set of parameters within those limits that minimises the distance above. At that point the model can be tested properly allowing for estimation so that mere numerical approximation does not create rejection. It turns out that this II estimator has very low small sample bias.

In what follows we will go over some of the features of the testing procedure.

To evaluate the reliability of our indirect inference test we can use a Monte Carlo experiment in which we check the power of the test as the structural model parameters deviate in accuracy from the true ones. In the experiment we assume a model of the type here to be true and generate a large number of simulated samples from it. We then deliberately falsify the model parameters by some percentage, x , each alternately plus or minus and see what percentage of these samples would reject the model. This gives us a measure of the test power — i.e. how frequently the test will reject a model of given falsity.

For our virus model we have found the results shown in the following Table 6. The rejection rate rises steadily as the model parameters are falsified. It reaches 100% as $x\%$ reaches 25%. This implies our model is likely to lie not too far from the truth and so the test has reasonable power.

Falseness	Rejection Rate
0	5.00
3	5.84
5	10.14
7	15.68
9	25.12
15	73.60
20	97.49
25	99.88

Table 6: Monte Carlo: Power Check

11 Appendix 2: data on first and second waves for 28 countries

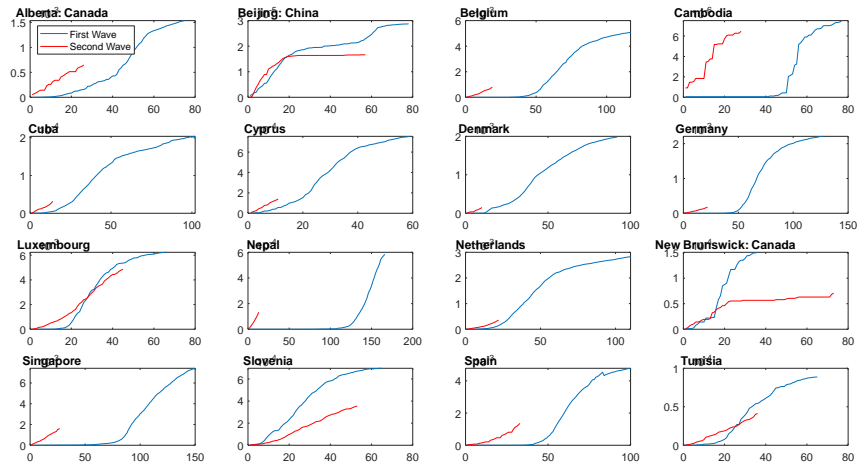


Figure 9: Cumulative Cases for Countries with Small Second Wave

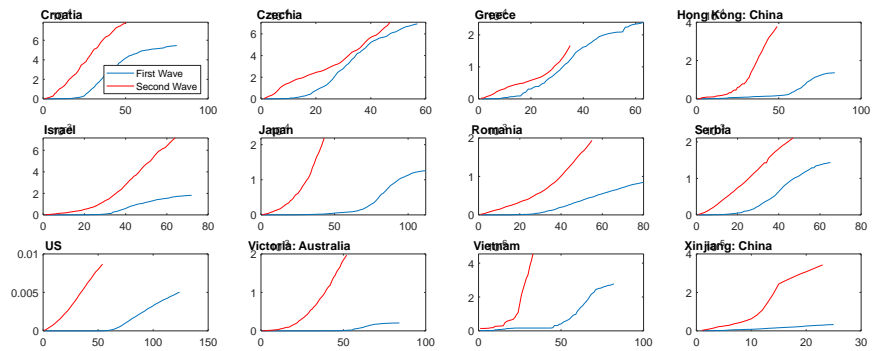


Figure 10: Cumulative Cases for Countries with Large Second Wave