


ARTICLE

TIME SERIES MODELLING OF EPIDEMICS: LEADING INDICATORS, CONTROL GROUPS AND POLICY ASSESSMENT

Andrew Harvey* 

Faculty of Economics, University of Cambridge, Cambridge, United Kingdom

*Corresponding author. Email: ach34@cam.ac.uk

(Received 14 April 2021; revised 21 June 2021; accepted 01 July 2021)

This article shows how new time series models can be used to track the progress of an epidemic, forecast key variables and evaluate the effects of policies. The univariate framework of Harvey and Kattuman (2020, *Harvard Data Science Review*, Special Issue 1—COVID-19, <https://hdsr.mitpress.mit.edu/pub/ozgix0yn>) is extended to model the relationship between two or more series and the role of common trends is discussed. Data on daily deaths from COVID-19 in Italy and the UK provides an example of leading indicators when there is a balanced growth. When growth is not balanced, the model can be extended by including a non-stationary component in one of the series. The viability of this model is investigated by examining the relationship between new cases and deaths in the Florida second wave of summer 2020. The balanced growth framework is then used as the basis for policy evaluation by showing how some variables can serve as control groups for a target variable. This approach is used to investigate the consequences of Sweden's soft lockdown coronavirus policy in the spring of 2020.

Keywords: balanced growth; COVID-19; Gompertz curve; Kalman filter; stochastic trend.

JEL codes: C22; C32.

1. Introduction

The aim of this article is to show how time series models¹ can be used to track the progress of an epidemic, forecast key variables and evaluate the effects of policies. Developing effective techniques to accomplish these tasks is of some importance, because, as documented by Ioannidis *et al.* (2020), the performance of many of the methods used to forecast the current COVID-19 epidemic has not been impressive. The new models draw much of their inspiration from time series econometrics. However, the characteristics of time series for epidemics are different from those of most time series in economics and these differences need to be taken into account.

Harvey and Kattuman (2020)—hereafter HK—developed a class of univariate time series models for predicting future values of a variable which when cumulated is subject to an unknown saturation level. In these models, the logarithm of the growth rate of the cumulated series depends² on a time trend. Allowing this trend to be time-varying introduces flexibility which, in the context of an epidemic, enables the effects of changes in policy and population behaviour to be tracked. Nowcasts and forecasts of the variables of interest, such as the daily number of cases, its growth rate and the instantaneous

¹The application of classical time series methods to data on epidemics is relatively undeveloped. Most of the emphasis has been on building 'semi-mechanistic' models to simulate the path of an epidemic under different assumptions about behaviour and policies; see Avery *et al.* (2020).

²The models have been used as the basis for the NIESR COVID-19 tracker since the early part of 2021; see Harvey *et al.* (2021) and <https://www.niesr.ac.uk/latest-weekly-covid-19-tracker>

reproduction number, R_t , can be made. Estimation of the models is by maximum likelihood (ML) and goodness of fit can be assessed by standard statistical test procedures.

Time series models can also be used to address other questions by exploring relationships between different series. One application concerns how the time path of an epidemic in a country which suffers an outbreak before another can be used as a leading indicator. The rationale for modelling the logarithm of the growth rate (of the cumulated series) comes from the properties of a Gompertz growth curve and when two such curves follow the same time path, but one lags the other, the trends in the series on the logarithms of the growth rate are a constant distance apart. This suggests that when the trends are stochastic, the same will be true. This situation, known as a balanced growth, arises in macroeconomics and is a special case of what econometricians call co-integration; see, for example, Stock and Watson (1988). The situation is illustrated by showing how the time path of deaths in the UK in the first few months of the coronavirus epidemic follows the time path of deaths in Italy 2 weeks earlier.

The requirement that two series exhibit balanced growth, while highly desirable, is not necessary for one to be a good leading indicator of the other. The need for additional flexibility is explored with data from the 'second wave' of coronavirus in Florida in the early part of the summer of 2020 where it is shown how daily new cases can potentially offer improved forecasts of deaths in 2–3 weeks' time. The forecasts are based on a bivariate unobserved component time series model that combines the dynamic information in the two series by a common trend specified as an integrated random walk (IRW) but includes an independent random walk (RW) component for new cases.

Time series modelling of an intervention can be used to assess the impact of a policy. This was done in HK in connection with the UK lockdown of March 2020. Here, an attempt is made to answer the question 'What if lockdown had been imposed a week earlier?' The impact of lockdowns is explored further by developing the ideas associated with balanced growth to try to estimate the number of coronavirus deaths in Sweden had a more stringent lockdown been imposed. The methodology draws on the study of control groups in time series by Harvey and Thiele (2021). It is argued that the fact that death rates in Sweden were roughly 10 times those in neighbouring countries could be misleading; the growth paths of the UK and Italy provide more relevant information. A comparison is made with studies based on the synthetic control (SC) method of Abadie *et al.* (2010, 2015).

2. Growth curves and time series models

This section sets out the basic model in which the logarithm of the growth rate of the cumulated series consists of a stochastic trend plus an irregular term. It is then shown how the framework may be extended to model the relationship between two series.

2.1. Dynamic trend models

The observational model uses data on the time series of the cumulated total of confirmed cases or deaths, Y_t , $t = 0, 1, \dots, T$, and the daily change, $y_t = \Delta Y_t = Y_t - Y_{t-1}$. HK show how the theory of generalised logistic growth curves suggests models for $\ln y_t$ and $\ln g_t$, where $g_t = y_t / Y_{t-1}$ or $\Delta \ln Y_t$. For the special case of the Gompertz growth curve:

$$\ln y_t = \ln Y_{t-1} + \delta - \gamma t + \varepsilon_t, \quad \gamma > 0, \quad t = 1, \dots, T, \quad (1)$$

and

$$\ln g_t = \delta - \gamma t + \varepsilon_t, \quad t = 1, \dots, T, \quad (2)$$

where ε_t is a random disturbance term.

A stochastic, or time-varying, trend may be introduced into (2), to give the dynamic trend model:

$$\ln g_t = \delta_t + \varepsilon_t, \quad \varepsilon_t \sim NID(0, \sigma_\varepsilon^2), \quad t = 1, \dots, T, \quad (3)$$

where

$$\begin{aligned} \delta_t &= \delta_{t-1} - \gamma_{t-1} + \eta_t, & \eta_t &\sim NID(0, \sigma_\eta^2), \\ \gamma_t &= \gamma_{t-1} + \zeta_t, & \zeta_t &\sim NID(0, \sigma_\zeta^2), \end{aligned} \tag{4}$$

and the normally distributed irregular, level and slope disturbances, ε_t , η_t and ζ_t , respectively, are mutually independent. When σ_ζ^2 is positive, but $\sigma_\eta^2 = 0$, the trend is an IRW. HK found an IRW trend to be particularly useful for tracking an epidemic and it will be adopted in the applications here. The speed with which a trend adapts to a change depends on the signal-noise ratio, which for the IRW is $q_\zeta = \sigma_\zeta^2 / \sigma_\varepsilon^2$; when $q_\zeta = 0$ the trend is deterministic, as in (2).

Allowing γ_t to change over time means that the progress of the epidemic is no longer tied to the proportion of the population infected as it would be if Y_t followed a deterministic growth curve. Instead the model adapts to movements brought about by changes in behaviour and policies. If γ_t falls to zero, the growth in Y_t becomes exponential while a negative γ_t means that the growth rate is increasing. This flexibility also allows the model to deal with second waves, where infections start to increase sharply after having fallen to a relatively low level. The Florida example³ of Section 4.2 shows how the model deals successfully with a second wave and Harvey *et al.* (2021) report accurate forecast for the second UK wave of early 2021. A modification of the model, that is currently under investigation, is to re-initialize the cumulative total at the low point before a new wave begins. The way in which the cumulative total enters the model is important because a key feature of the dynamic Gompertz model is its ability to detect upcoming turning points and to make forecasts that show a downward movement even before a peak has been reached; see, for example, the forecasts made for Germany in HK.

Additional components, such as day of the week effects, can be added to (3). These may be deterministic or stochastic. Explanatory variables, including interventions, can also be included, as may stationary components. Thus (3) could become:

$$\ln g_t = \delta_t + \theta_t + \mu_t + \mathbf{x}'_t \boldsymbol{\beta} + \varepsilon_t, \quad t = 1, \dots, T, \tag{5}$$

where θ_t is a stochastic daily component, modelled as in Harvey (1989, pp. 43–4), μ_t is a stationary autoregressive process, \mathbf{x}_t is a vector of explanatory variables and $\boldsymbol{\beta}$ is a corresponding vector of parameters. Possible candidates for explanatory variables include stringency indices for governmental policies, as in Hale *et al.* (2021). All these models can be handled using techniques based on state space models and the Kalman filter (KF); see Durbin and Koopman (2012). Here, the STAMP package of Koopman *et al.* (2021) is used. Estimation of the unknown parameters is by ML. Diagnostic tests for normality and residual serial correlation are based on the one-step ahead prediction errors, $v_t = \ln g_t - \delta_{t|t-1}$, $t = 3, \dots, T$.

The KF outputs the estimates and forecasts of the state vector $(\delta_t, \gamma_t)'$. Estimates at time t conditional on information up to and including time t are denoted $(\delta_{t|t}, \gamma_{t|t})'$, while predictions j steps ahead are $(\delta_{t+j|t}, \gamma_{t+j|t})'$. The smoother, which estimates the state at time t based on all T observations in the series, is denoted $(\delta_{t|T}, \gamma_{t|T})'$.

Remark 1. *When the observations are small, a negative binomial distribution for y_t , conditional on past observations, may be appropriate. HK show how the model may be modified to deal with this possibility. However, the numbers in the applications here are big enough to allow y_t to be treated as conditionally lognormal and hence for the conditional distribution of $\ln g_t$ to be considered normal.*

2.2. Forecasts

The forecasts of the trend in future values of $\ln g_t$ in the dynamic Gompertz model are given by $\delta_{T+\ell|T} = \delta_{T|T} - \gamma_{T|T} \ell$, $\ell = 1, 2, \dots$, where $\delta_{T|T}$ and $\gamma_{T|T}$ are the KF estimates of δ_T and γ_T at the end of the sample.

³Appendix B illustrates the importance of the $\ln g_t$ transformation for new cases in Florida.

Forecasts of the trend in the daily observations are obtained from a recursion for the trend in their cumulative total, Y_t , namely,

$$\mu_{T+\ell|T} = \mu_{T+\ell-1|T} \left(1 + g_{T+\ell|T} \right), \quad \ell = 1, 2, \dots, \tag{6}$$

where $g_{T+\ell|T} = \exp \delta_{T+\ell|T}$ and $\mu_{1|T} = Y_T$. The trend in the daily figures is then,

$$\mu_{y,T+\ell|T} = g_{T+\ell|T} \mu_{T+\ell-1|T}, \quad \ell = 1, 2, \dots \tag{7}$$

Daily effects can be added to δ_t . In this case, forecasts of the observations themselves, that is $\hat{y}_{T+\ell|T}$ and $\hat{Y}_{T+\ell|T}$, are given by adding the filtered value of the daily component to the trend component, $\delta_{T+\ell|T}$.

Unlike most other forecasting methods, the dynamic Gompertz model yields prediction intervals. The way in which they are constructed is set out in Section 2.4 of Harvey *et al.* (2021) and examples are given.

2.3. Forecasting and nowcasting R

Harvey and Kattuman (2021) use filtered estimates of $g_{y,t}$, given by $g_{y,t|t} = g_{t|t} - \gamma_{t|t}$, to track the progress of an epidemic. A corresponding estimator of the instantaneous reproduction number, R_t , can be constructed in a number of ways, as in Wallinga and Lipsitch (2007). The most practical for COVID-19 are:

$$\tilde{R}_{t,\tau} = 1 + \tau g_{y,t|t} \quad \text{and} \quad \tilde{R}_{t,\tau}^e = \exp \left(\tau g_{y,t|t} \right), \tag{8}$$

where τ is the generation interval, which is the number of days that must elapse before an infected person can transmit the disease; setting $\tau = 4$ is a good choice. Harvey and Kattuman (2021) provide more details and show how forecasts of $R_{t,\tau}$, with associated prediction intervals. Harvey *et al.* (2021) illustrate the implementation of this approach in the NIESR tracker.

2.4. Panel data

The extended dynamic Gompertz model of (5) can be used as the basis for handling panel data. When there are N cross-sectional units,

$$\ln g_{it} = \delta_{it} + \mathbf{z}'_{it} \boldsymbol{\alpha}_i + \mathbf{x}'_{it} \boldsymbol{\beta} + \varepsilon_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \tag{9}$$

where δ'_{it} s are stochastic trend components and \mathbf{z}_{it} and \mathbf{x}_{it} are vectors of explanatory variables, with $\boldsymbol{\alpha}_i$, $i = 1, \dots, N$, and $\boldsymbol{\beta}$ denoting associated coefficients. It may be necessary to add autoregressive and day of the week components. Either way, we can pre-filter with the univariate filter, as in Harvey (1989, Sect. 3.4.2) and then, if the components are assumed to be mutually independent, use the transformed observation to estimate a standard panel data model. This procedure can be iterated to convergence.

Further generalisations would let the stochastic trends depend on common factors.

3. Comparing different growth curves

The Gompertz growth curve lies behind the notion of setting up time series models in which the logarithm of the growth rate of the cumulative total of a variable follows a trend. It is therefore able to provide insight on how to formulate and interpret models linking several series.

The Gompertz growth curve is:

$$\mu(t) = \bar{\mu} \exp(-\alpha e^{-\gamma t}), \quad \alpha, \gamma > 0, \quad -\infty < t < \infty, \tag{10}$$




Figure 1. (Colour online) Gompertz incidence curves, $\mu'(t)$, with $\gamma = 0.15$, $\alpha_1 = 20$ for the left hand curve and $\alpha_2 = 100$ for the right hand curves; the value of $\bar{\mu}$ in the upper curve is 1400 as opposed to 1000 as in the lower curve

where γ is a growth rate parameter, $\bar{\mu}$ is the upper bound or saturation level ($\mu(t) \rightarrow \bar{\mu}$ as $t \rightarrow \infty$) and α reflects initial conditions. The associated incidence curve is:

$$d\mu(t)/dt = \mu'(t) = \gamma\alpha\mu(t) \exp(-\gamma t),$$

with a peak at $t = \gamma^{-1} \ln \alpha$. Figure 1 shows an incidence curve with a peak at $t = 19.97$, together with the same curve shifted to the right so the peak is at 30.71. A curve above the right hand curve is also shown; this is higher because the value of $\bar{\mu}$ is 1400 rather than 1000 as it is for the other two curves. In all cases $\gamma = 0.15$, but for the left hand curve α is 20 whereas for the right hand curves it is 100.

Although the right hand curves in figure 1 clearly lag the left hand one, it is not immediately evident how to model the relationship. However, the logarithms of the growth rates of $\mu(t)$ are:

$$\ln g(t) = \delta - \gamma t, \quad t \geq 0, \tag{11}$$

where $\delta = \ln \alpha\gamma$; compare (2). Figure 2 shows the two lines for $\ln g(t)$ running in parallel. The distance between them depends on the intercepts, δ , which in turn depend on the initialization parameter, α . The height of the incidence curve, which depends on the saturation level, $\bar{\mu}$, is irrelevant; as a result, the lines corresponding to the two right-hand incidence curves in figure 1 are identical. This is important because it means that small populations can be compared with big ones: size does not matter.

When two lines are parallel, the upper line lags the lower one by:

$$k = \frac{\delta_2 - \delta_1}{\gamma} = \frac{\ln \alpha_2 - \ln \alpha_1}{\gamma}, \tag{12}$$

where δ_1 and δ_2 are the intercepts of the lower and upper lines, respectively and α_1 and α_2 are the corresponding initial conditions. In figure 2, the lag is $k = 10.73$. When the γ 's are different, the epidemic progresses at different speeds. The lines for $\ln g(t)$ are no longer parallel and the time lag is no longer constant.

4. A statistical model for leading indicators

Now consider observational models of the form (2) for two time series which are on the same growth path because $\gamma_1 = \gamma_2$ but the first series leads the second by k time periods. The observations run from




Figure 2. (Colour online) Logarithms of the growth rates for incidence curves in figure 1; $\gamma = 0.15$, $\alpha_1 = 20$ and $\alpha_2 = 100$ (upper line)

$t = 1$ to T but when the first series is lagged by k time periods, $\ln g_{1,t-k}$ runs from $t = k + 1$ to $T + k$. Subtracting the first series from the second gives:

$$\ln g_{2t} = \delta + \ln g_{1,t-k} + \varepsilon_t, \tag{13}$$

where $\delta = \ln(\alpha_2/\alpha_1)$ and the disturbance term is $\varepsilon_t = \varepsilon_{2t} - \varepsilon_{1,t-k}$. The equation takes the same form when the trends are stochastic, so long as there is balanced growth. The disturbance, ε_t , can be replaced by any stationary process.

When the two series are not on the same growth path, there is no longer a value of k for the contrast in (13) that makes it stationary. The stationarity test of Kwiatkowski *et al.* (1992)—the KPSS test—can be used to test for this possibility.

A bivariate time series model combines the dynamic information in the target series with that in the leading indicator. It is set up by lagging the observations on the leading indicator so that they are aligned with the target. Hence, defining $g_{1,t}^{(k)} = g_{1,t-k}$ for $t = k + 1, \dots, T + k$ gives:

$$\begin{aligned} \ln g_{1,t}^{(k)} &= \delta_t + \psi_t + \varepsilon_{1t}, & t = k + 1, \dots, T + k, \\ \ln g_{2t} &= \bar{\delta} + \delta_t + \varepsilon_{2t}, & t = k + 1, \dots, T. \end{aligned} \tag{14}$$

The k future values of $\ln g_{2,T+j}$, $j = 1, \dots, k$ are treated as missing observations.⁴ The trend, δ_t , is an IRW that is designed to capture the growth path of the target series. Its initial level has been (arbitrarily) assigned to the first series; hence the need for a constant term, $\bar{\delta}$, in the second. The role of the other stochastic component, ψ_t , is to allow for deviations of the leading indicator from the balanced growth path. A convenient specification for it is the first-order autoregressive process,

$$\psi_t = \phi\psi_{t-1} + \zeta_t, \quad \zeta_t \sim NID\left(0, \sigma_\zeta^2\right), \quad t = k + 1, \dots, T + k.$$

All disturbances in (14), including ε_{1t} and ε_{2t} , are Gaussian and assumed to be mutually as well as serially independent. Only a single lag is present. More lags could be included, but the aim is find a viable

⁴If the first k observations on the second series are reliable they could be used by treating the first k values of the first series as missing.

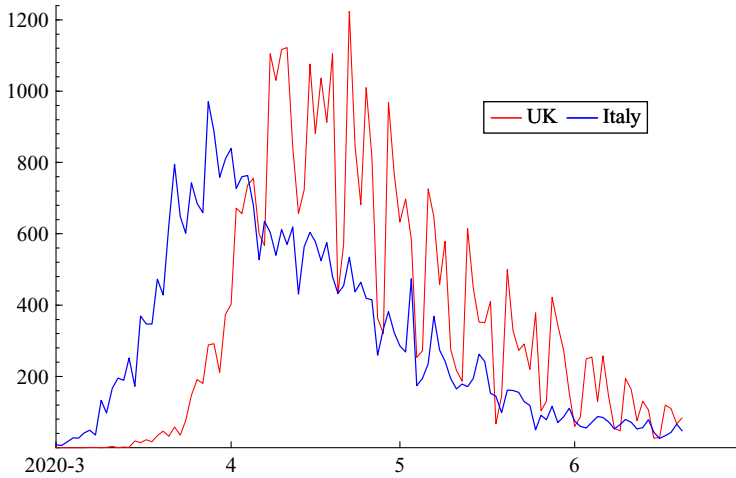


Figure 3. (Colour online) Daily deaths from COVID-19 in Italy and UK in 2020

leading indicator for movements in the trend rather than to estimate a distributed lag for the observations. Estimation is by state space methods. As new observations become available, nowcasts and forecasts are updated by the KF.

When $|\phi| < 1$, the series are co-integrated with balanced growth. In the absence of balanced growth, the suggestion is to let ψ_t be a RW, by setting $\phi = 1$. The value of k is then based on experimentation and prior information about what might constitute a reasonable lag. The hope is that the RW specification for ψ_t enables its movements to be separated from those in the IRW trend.

The filtered estimates, $g_{t|t}$ and $\gamma_{t|t}$, for the target series give the nowcast of $g_{y,t}$ at time $t = T$ and the forecast at $t = T + k$. Forecasts can also be made beyond $t = T + k$, but without the benefit of corresponding values of the leading indicator. The KF and smoother implicitly weights observations in both series in order to compute nowcasts and forecasts for the target.

4.1. Italy and the UK

Figure 3 shows the daily deaths in Italy and the UK from 2 March to 20 June 2020; after that the numbers for Italy start to become small. The figures are for when the deaths were recorded rather than when they occurred. Series based on date of death would not have the daily pattern but were difficult to obtain at that time. Data sources are given in Appendix A.

Italy clearly leads the UK but the relationship is captured more precisely in figure 4 which shows the logarithms of the growth rates of total deaths (LDL). The UK numbers are small at the beginning of March and so there are missing observations. A lag of 14 days is not inconsistent with prior information and it has the attraction of lining up the days of the week in the two countries. Other lags were tried but 14 minimised the gap between the two countries. Figure 5 shows the LDL series with Italy lagged by 14 days together with the contrast between the two countries obtained by subtracting Italy from the UK. The contrast series appears to be stationary with a mean close to zero; without the lag for Italy the values at the end of March and the beginning of April tend to be higher than the others, reflecting the later UK lockdown. Estimating a regression model with daily dummy variables removed most of the serial correlation and gave a mean of $\tilde{\delta} = -0.083$, with a SE on 0.035. The diagnostic statistics were⁵: $r(1) = -0.06$; $Q(14) = 13.40$, $r(1) = -0.06$, $BS = 1.85$ and $H = 1.24$.

⁵ $r(1)$ is the autocorrelation at lag one, $Q(P)$ is Box–Ljung statistic with P autocorrelations, BS is the Bowman–Shenton normality statistic and H is a heteroscedasticity statistic constructed as the ratio of the sum of squares in the last third of the sample to the sum of squares in the first third.

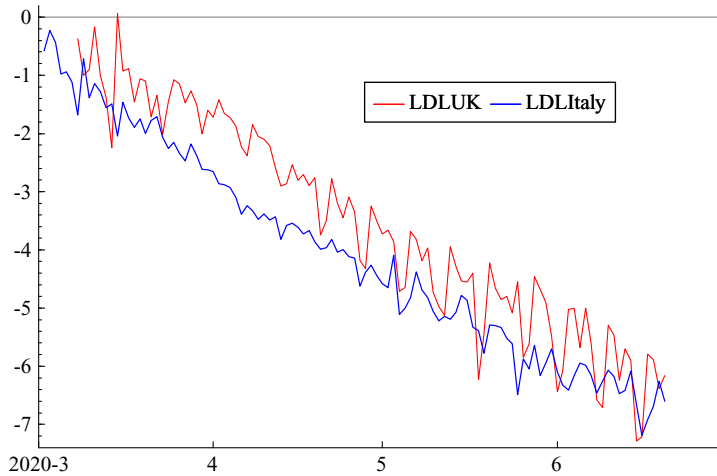


Figure 4. (Colour online) Logarithms of the growth rates (LDL) of total deaths in UK and Italy

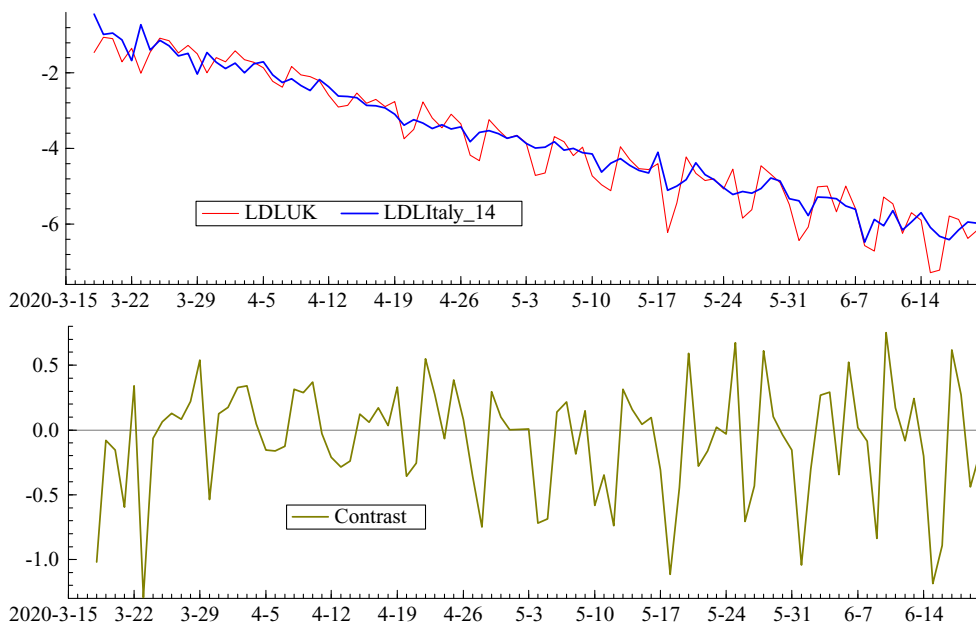


Figure 5. (Colour online) LDL series from 16 March to 20 June with Italy lagged by 14 days together with the contrast LDLUK-LDLItaly

Fitting a bivariate time series model of the form (14), starting on 16 March and finishing on 5 July, gave a slowly changing trend that was close to being deterministic. The δ_{1t} term was excluded but a daily component was included. The estimate of the daily growth rate of UK deaths 14 days beyond the final observation on 20 June was $g_{y,T+k|T} = -0.058$, giving a forecast of $\hat{R}_{T+k,4} = 0.77$.

4.2. Deaths and new cases in Florida

Daily cases of COVID-19 in the U.S. state of Florida peaked in early April. There was then a decline following a lockdown during April. After April restrictions were eased and there was a levelling out in

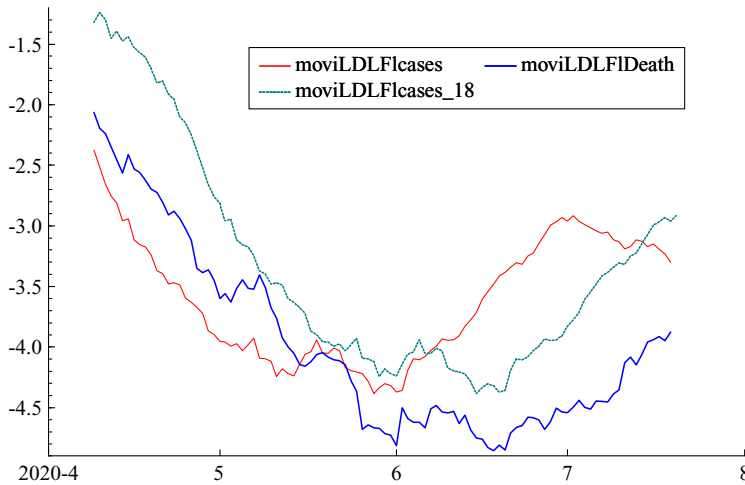


Figure 6. (Colour online) Seven-day moving average of LDL deaths in Florida, new cases and new cases lagged 18 days (dotted line) from 29 March till 19 July 2020

May, followed by a sharp rise in June. This second wave poses a challenge for a model in which new cases are used as a leading indicator for deaths. The model deals with the second wave by allowing $\gamma_{t|t}$ to become negative; estimates of R_t can still be obtained from $g_{y,t|t}$, as in (8).

Aside from the model having to deal with a situation where new cases and deaths rise and fall, there is the problem that the basis on which new cases are recorded changes over time. At the beginning of the pandemic, new cases in many countries were primarily hospital admissions, but over time testing became more widespread. In the case of Florida, there was an increase in testing in May, although the growth rate in tests was roughly constant from the end of May onwards. This suggests that the growth rate of confirmed new cases may still be a good indicator of the path of new infections; see Harvey and Kattuman (2021).

The observations, particularly deaths, have a strong weekly pattern. A clearer impression of the underlying trend is given by figure 6 which shows a 7-day moving average of the logarithms of the growth rates of total new cases and deaths from 29 March to 19 July 2020 inclusive. It is apparent that new cases are leading, but the relationship between deaths and new cases is not completely stable over time, partly because of an increase in the growth rate of testing in May and partly because of other factors, such as better hospital treatment. The inclusion of the ψ_t component in model (14) offers a way of dealing with the discrepancy. Nevertheless, despite the instability, it seems clear that new cases peak some time before the end of the sample whereas deaths appear to be at their peak, something confirmed by later observations.

The lag in (14) is chosen so as to get maximum benefit for new cases as a leading indicator. It is not trying to model the distribution of days from infection to death although the choice of k may be roughly aligned with the mean time to death. After some experimentation it was decided to fix the lag at 18. The LDL for new cases shifted in this way, and shown in figure 6, tracks deaths quite well.

The model, including day of the week variables, was fitted to the Florida data from 29 March till 19 July, with the new cases shifted forward 18 days so as to end on 6 August; thus $k = 18$. Specifying ψ_t as a first-order autoregressive process in (14) gave an estimated ϕ of 0.998, so a RW seems appropriate. The smoothed estimates of this RW component are shown in figure 7. The lower graph is the smoothed estimates of the day of the week component in deaths; note that the bivariate model is able to give estimates for the period after 19 July when there are no observations on deaths. The size and variability of the daily component in deaths is much bigger than it is for new cases, with the very high variation coinciding with the period when the numbers of deaths are relatively small. Similarly the prediction error

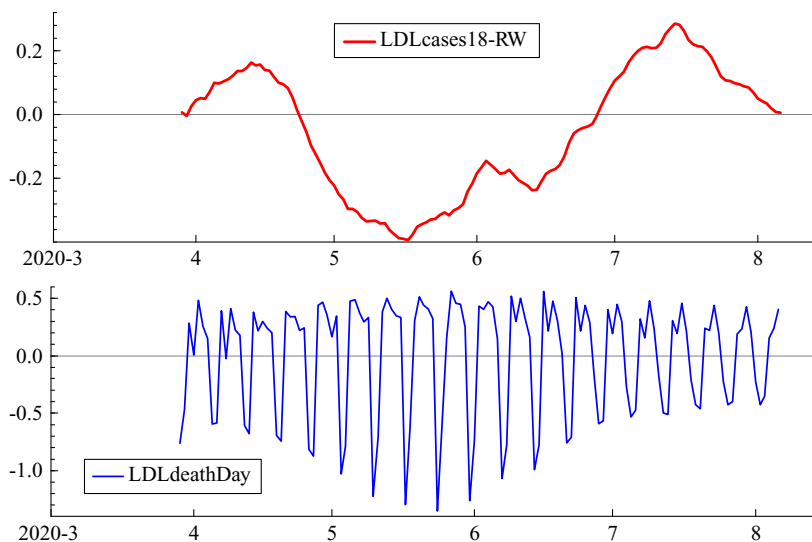


Figure 7. (Colour online) Smoothed estimates of the RW component in Florida new cases, shifted forward by 18 days and the associated daily component in deaths

variance of 0.115 for new cases was less than half the 0.253 obtained for deaths. Little serial correlation remained in the residuals for deaths: the Box–Ljung *Q*-statistic for the first 18 residual autocorrelations was 8.16, while the corresponding figure for cases was a little higher⁶ at 25.01. The signal-noise ratio was estimated as 0.00037, so the trend changes relatively slowly but is still able to adapt to changes in direction.

Figure 8 shows the forecasts of the logarithm of the growth rate deaths, obtained by using the leading indicator, together with the actual observations from 20 July to 4 August. The smooth dotted line is the trend in LDL deaths. As can be seen, the model foresees the turning point. By contrast, the growth rate of LDL deaths on 19 July is still positive, and estimating a univariate model up to this point gave forecasts continuing on an upward path, overshooting the actual observations.

5. Policy interventions and control groups

This section shows how the time series models can be used to assess the effects of policy. The first example uses univariate time series modelling to investigate the timing of the UK lockdown in the spring of 2020. The second illustrates how the balanced growth framework provides the basis for policy evaluation by showing how some variables can be used as control groups for a target variable. The consequences of Sweden’s soft lockdown coronavirus policy in the early part of 2020 are assessed and a comparison is made with studies based on the method of SC.

5.1. What if the March 2020 lockdown in the UK had been a week earlier?

The UK went into full lockdown on 23 March. Can we estimate how many deaths could have been saved if it had been a week earlier?

A slope intervention in (2) enables the effect of a policy which changes γ to be evaluated. Thus,

$$\ln g_t = \delta - \gamma t - \beta tw_t + \varepsilon_t, \quad t = 1, \dots, T, \tag{15}$$

⁶The suggestion in the STAMP manual is to test against a chi-square variable with allowance made for the loss in degrees of freedom due to estimated parameters which here is six. Thus, the chi-square may be taken to have 12 degrees of freedom.

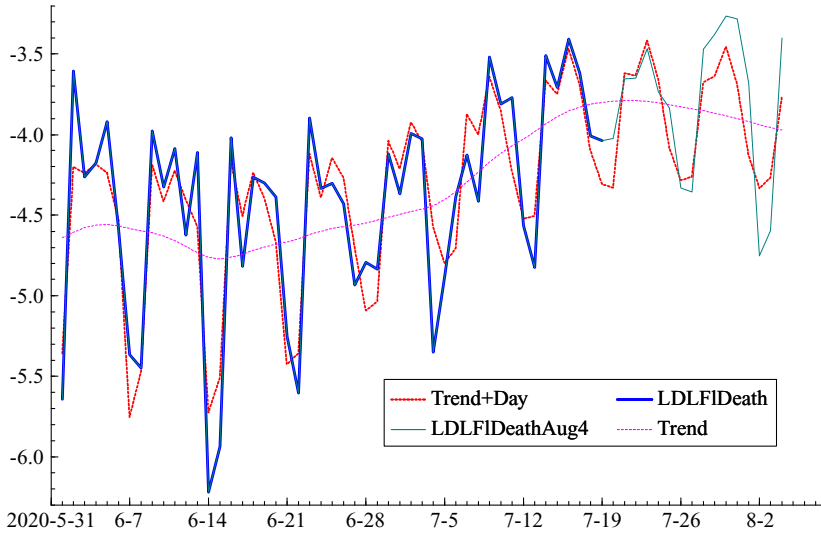


Figure 8. (Colour online) Forecasts (dots) and trend (smooth dots) of the logarithm of the growth rate of deaths, obtained by using the leading indicator, together with the actual observations from 20 July to 4 August; observations before 20 July (LDLFIDeath) shown by thick line

where w_t are intervention dummy variables. When the full effect is realised, the slope on the time trend will have moved from γ to $\gamma + \beta$. A positive β lowers the growth rate, g_t , the peak of the incidence curve and the final level. The intervention dummies can be constructed from a logistic cumulative distribution function, giving a response curve:

$$W(t) = 1 / (1 + \exp(-\gamma(t - t^{\text{med}}))), \quad -\infty < W(t) < \infty,$$

where t^{med} is the median. With t^L and t^U denoting the beginning and end of the time span during which the response to the intervention occurs, the dummy variables are defined as $w_t = 0$ for $t < t^L$, $w_t = W(t)$ for $t^L, t = t^L + 1, \dots, t^L, \dots, t^U$ and $w_t = 1$ for $t = t^U + 1, \dots, T$. HK fitted the static model in (15) to new cases in the UK, with an intervention starting on 26 March and ending on 12 April, using data from the beginning of March up to 29 April. The result was an estimate of β equal to 0.020 (0.004) and an estimate of γ also equal to 0.020. The overall effect⁷ is a new slope of 0.041. The trend, with the intervention included, is shown by the dashed line in figure 9.

The effect of implementing lockdown restrictions a week earlier can be estimated by shifting the intervention response forward by 1 week so it starts on 19 March, rather than on 26 March. The adjusted trend in the logarithm of the growth rate is then:

$$\ln g_t^* = \delta - \gamma t - \beta t w_{t+7}, \quad t = 1, \dots, T. \tag{16}$$

Once the effect of the intervention has worked itself through, the new slope is the same as before, as can be seen in the solid line in figure 9.

The predicted final total is:

$$\bar{\mu} \simeq \mu_T \exp((\exp \delta_{T|T}) / \gamma_{T|T}),$$

⁷When the slope was allowed to be stochastic, the estimate of β was reduced to 0.014 (0.006), but with such a small sample size, a stochastic slope risks some confounding with the intervention variable.

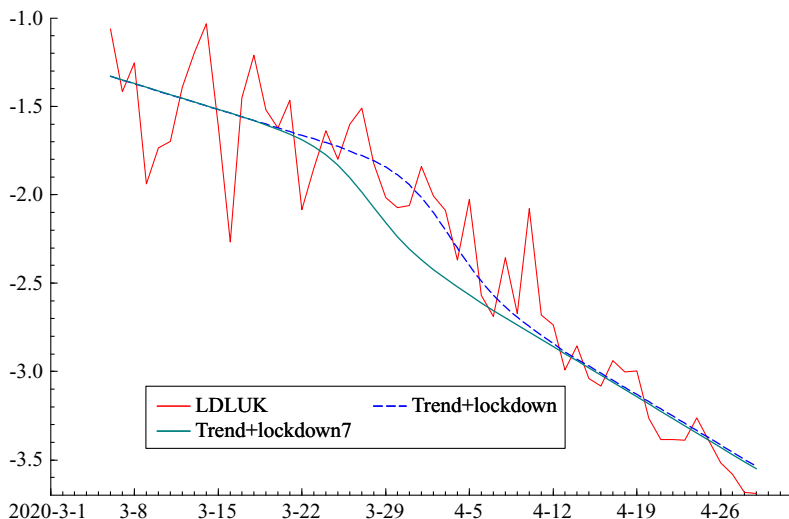


Figure 9. (Colour online) Estimates of logarithm of growth rate of total cases in UK with a logistic intervention and a daily effect

where T is 12 April. For the actual data, μ_T can be approximated by Y_T , but for the early lockdown scenario, μ_T will be smaller because the growth rate falls earlier. This implies that the level on 18 March is multiplied by $\exp(\sum g_t^*)$, where the summation is over the period from 19 March to 12 April. To ensure comparability, the actual level on 12 April is best estimated in the same way, rather than by Y_T . Thus an estimate of the ratio of the total number of cases for a hypothetical early lockdown to the actual total is given by:

$$\frac{\text{Hypothetical}}{\text{Actual}} = \frac{\exp(\sum g_j^*)}{\exp(\sum g_j)} = \frac{\exp(\sum \exp(\delta - \gamma t - \beta t w_{t+7}))}{\exp(\sum \exp(\delta - \gamma t - \beta t w_t))}$$

This ratio is 0.551 implying that the number of infections, as measured by data on daily coronavirus hospital admissions, could have been almost halved by an earlier lockdown. If a constant proportion of those admitted die, the implication is that deaths in the initial phase of the epidemic (up to the end of June) could have been almost halved by an earlier lockdown.⁸ This conclusion is not too different from ones obtained by other methods. For example, the BBC reported on 10 June that Professor Neil Ferguson of Imperial College told a committee of MPs: ‘Had we introduced lockdown measures a week earlier, we would have reduced the final death toll by at least a half’.

5.2. Fewer deaths in Sweden with a full lockdown?

Sweden did not opt for the full lockdown that other European countries imposed in March. Restrictions were minimal: the government recommended frequent handwashing, working from home, self-isolation for those who felt ill or were over 70 and social distancing⁹; see, for example, Kamerlin and Kasson (2020). Did this policy lead to the number of deaths being significantly higher than it might have been

⁸It should be stressed that these findings relate specifically to the effect of the full lockdown of March 2020. A full lockdown imposed now is unlikely to have the same impact because the environment is different in that social distancing restrictions are in place, behaviour has changed and the risk to care homes is better understood.

⁹Carl Bildt, a Former Prime Minister, was quoted as saying ‘Swedes, especially of the older generation, have a genetic disposition to social distancing anyway’.

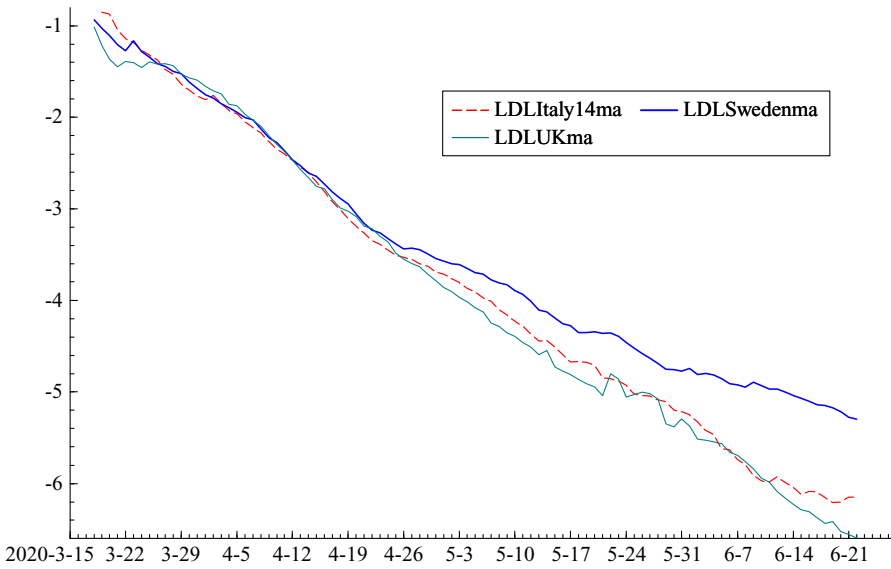


Figure 10. (Colour online) Seven day moving averages of the logarithms of the growth rate (LDL) from 18 March to 22 July

under a full lockdown? To answer this question we need to determine the growth path that Sweden would most likely have followed under a hard lockdown.

The analysis is based on daily deaths in Sweden, UK and Italy (lagged 14 days) from 18 March to 22 July; by the end of July numbers had become small. A comparison of actual and potential growth paths is best carried out with the logarithms of growth rates of the cumulative total for the reasons discussed earlier. Although Sweden is much smaller than the UK and Italy, there is no need to take deaths per 100,000 because it follows from the discussion in Section 2.3 that standardising in this way leaves the growth rate, g_t , unchanged. Because the day of the week effect is very strong, particularly in the UK, the logarithms of growth rates were smoothed with a 7-day moving average, centred on the fourth day. The graph in figure 10 shows that Sweden initially fell at the same rate as the UK and Italy but then started to diverge¹⁰ around 24 April, about a month after the UK lockdown began on 23 March.

If Sweden had kept on the same growth path as the UK and Italy there would have been fewer deaths. An estimate of the number of deaths under this alternative scenario is given by reference to the forecasting equations in Section 2.2. Let $t = m$ denote the date of divergence and let $\hat{\delta}_t$ denote the values of δ_t estimated for the lockdown growth path using the data on UK and Italy. Since the moving averages are quite smooth, $\hat{\delta}_t$ was constructed as a simple average of the two countries, rather than by restricted least squares (RLS) as¹¹ in Harvey and Thiele (2021). Then,

$$\hat{\mu}_{m+j} = \hat{\mu}_{m+j-1} \left(1 + \hat{g}_{m+j} \right) \approx \hat{\mu}_{m+j-1} \exp \hat{\delta}_{m+j}, \quad j = 1, 2, \dots, T - m. \tag{17}$$

The initial value is $\hat{\mu}_m = Y_m$, or a weighted average around that point. Solving the recursion gives:

¹⁰The growth path of deaths in the UK and Italy differs somewhat from the growth path of new cases. The growth rate of $\ln g_t$ for new cases, that is γ_t , drops significantly within a little over 2 weeks from the start of lockdown; HK estimate the UK fall by fitting intervention variables. A corresponding sharp drop in γ_t is less evident in the deaths data. The divergence of Sweden from Italy and the UK is more a consequence of the Swedish γ_t increasing, rather than the γ_t 's falling for the other countries.

¹¹The general methodology is to select a set of controls from a donor pool by using the KPSS test to determine which series are on a balanced growth path with the target. The control group weighting is then determined by RLS. The complication here is that, when there is an intervention, balanced growth may require lagging some of the series.

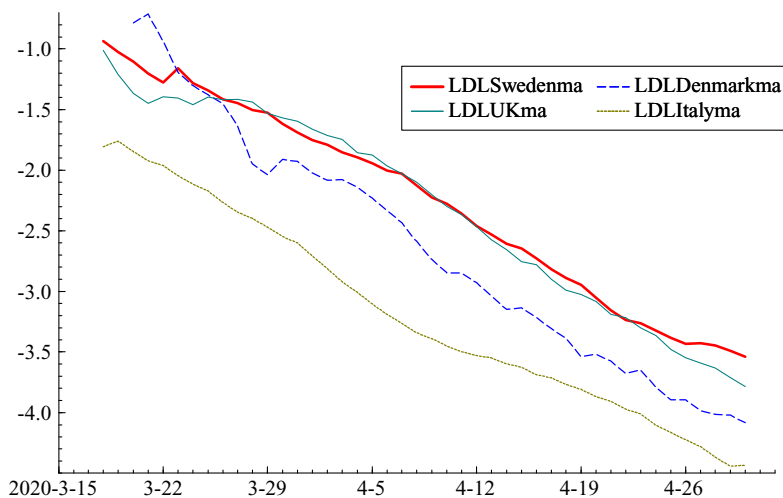


Figure 11. (Colour online) Seven-day moving averages of the logarithms of the growth rate from 18 March to 30 April

$$\hat{Y}_T = \hat{\mu}_T = Y_m \prod_{j=1}^{T-m} (1 + \hat{g}_{m+j}) \approx Y_m \exp \sum_{j=1}^{T-m} \hat{\delta}_{m+j}, \tag{18}$$

as the estimated total number of deaths, up to time T , under the lockdown scenario. The estimated number of deaths after time m is $\hat{Y}_T - Y_m$ while the actual is $Y_T - Y_m$. Here T is 22 July; the number of deaths after that is relatively small.

The total on 24 April was 2236 and using formula (18) gives an estimate of 4062 for 22 July as opposed to an actual figure of 5722, a difference of 1660. The sensitivity to the initial value can be gauged by noting that the estimates using the totals 2 days before and 2 days after 24 April are 3808 and 4378, respectively.

One way of reducing the dependence on the starting value is to estimate the underlying total for Sweden using formula (18) with \hat{g}_{m+j} replaced by the actual Swedish values. This gave a total of 5657. The ratio of \hat{Y}_T for the lockdown control group to that of Sweden is $1.816/2.530 = 0.718$. For $\hat{Y}_T - Y_m$ it is $0.816/1.530 = 0.533$. This implies that the actual increase from 24 April, which was 3486, could have been 1902. The first method gave $4062 - 2236 = 1826$. The overall conclusion is that, between 24 April and 22 July, there were perhaps 40–45 per cent more deaths than there might have been under a more stringent lockdown of the kind implemented in the UK and Italy.

It is worth noting that although Sweden may have had more deaths under its soft lockdown, this does not mean a higher death rate than countries which had a hard lockdown. On 4 September, the figures for deaths per one million for Sweden were 577 as against 611 for the UK and 587 for Italy. The rates for Denmark, Norway and Finland were 108, 49 and 61, respectively, but this should not lead one to infer that the soft Swedish lockdown resulted in a death rate of perhaps ten times what it might have been.

The number of deaths in Denmark is too small to allow a full analysis based on the logarithms of growth rates. The variability is high and after mid-May there are often days when no deaths occur. Numbers in Norway and Finland are lower still. However, up to the end of April the logarithm of the growth rate for Denmark is informative. Figure 11 shows the logarithms of the growth rates for Sweden, Italy, UK and Denmark. Denmark is on a similar growth path to that of the other countries but it is lower than the UK because coronavirus may have arrived earlier and lockdown was imposed on 13 March; the gap is consistent with Denmark leading the UK by about a week. During this period deaths in Denmark were much lower than in Sweden even though they were on the same growth path until close to the end of April. This difference therefore seems to be for reasons not directly connected to the policies of the two countries on lockdown.

On 30 April, 2714 deaths had been recorded in Sweden as against 443 in Denmark, a ratio of 6.13. On 24 April, the figures were 2236 and 394, a ratio of 5.68. (But bear in mind that the population of Sweden is 1.76 times that of Denmark so in per capita terms the ratio is closer to three.) On 22 July, the ratio of Swedish to Danish deaths had risen to 9.36. However, the ratio of the lockdown estimate of 4062 to the 611 Danish deaths is only 6.64 which is not far from the ratio at the end of April. Thus the estimate of the number of deaths obtained using the control group seems quite plausible. The conclusion is that for reasons unconnected with lockdown policy the death rate per head in Sweden was about three and a half times that in Denmark. The less stringent lockdown then raised this ratio to nearly five and a half.

5.3. Synthetic control

A number of researchers have analysed the Swedish experience using the method of SC. The recent paper by Cho (2020) is a careful and thoughtful analysis, containing a number of references to earlier papers on the topic. Cho uses daily infection case data per million people to construct a SC variable for Sweden using observations from 29 February to 24 March. The countries and their SC weights were: Finland (0.49), Greece (0.24), Norway (0.22), Denmark (0.03) and Estonia (0.02). The choice of these countries, with the exception of Greece, is not unexpected.¹² Cho concludes that, for the 75 days post-lockdown days, from 25 March until early June, synthetic Sweden is 75 per cent lower than actual Sweden. The SC method cannot be applied directly to deaths because, as noted above, the numbers for the key control group candidates are too small so Cho goes on to examine excess deaths by combining the analysis of new cases with weekly data on excess mortality. He concludes that excess deaths were about 25 per cent less in synthetic Sweden as compared with actual Sweden. What is striking is that in the balanced growth analysis the reduction in deaths is quite close, at 29 per cent, and converting to excess deaths might end up with a figure that is closer still.

Cho, in common with other SC researchers like Born *et al.* (2020), uses raw cases numbers, standardised for population. However, the logarithm of the growth rate could also be used and, since this yields better behaved time series, it would be interesting to see if it yields the same SC group.

Overall the balanced growth approach is simpler, more transparent and arguably more convincing. Harvey and Thiele (2021) reach the same conclusion in their analysis of the seminal SC applications of Abadie *et al.* (2010, 2015).

6. Conclusion

The aim of this article has been to provide a methodological framework for the statistical analysis of the relationship between time series of the kind that are relevant for tracking and forecasting epidemics and analysing the effects of policy. The examples illustrate how the methods may be applied in practice, although a degree of caution is needed in interpreting the results because of data revisions and different definitions of what constitutes a COVID-19 death.

The growth path of an epidemic is best captured by fitting a stochastic trend to the logarithm of the growth rate of the cumulated series. When two series are on a balanced growth path, the difference between them is stationary. The relationship between deaths from coronavirus in the UK and Italy in the first half of 2020 is a good example of balanced growth, with deaths in Italy 14 days earlier providing a leading indicator for deaths in the UK. A bivariate state space model takes full account of the dynamics in both series and, by extracting the common underlying trend, yields estimates of the daily growth rate of an epidemic and the associated value of R_t .

The balanced growth model was extended by including a RW component. This allows the growth path of the leading indicator to deviate from the growth path of the target series. A model of this kind

¹²In an earlier study, Born *et al.* (2020) selected a somewhat different group, namely the Netherlands (0.39), Denmark (0.26), Finland (0.19), Norway (0.15) and Portugal (0.01).

linking deaths to new cases in Florida was estimated for the period covering the second wave in early summer 2020. The forecasts made for deaths while they were still rising are remarkably successful in picking up the subsequent downward movement.

Policy evaluation can be carried out by using some series as control groups for others. A common trend or, better still, balanced growth is the key ingredient. The Swedish policy response to coronavirus provides an example of the methodology. It is shown that the average of the growth paths of deaths in the UK and Italy yields a suitable control group for deaths in Sweden. The Swedish growth path is initially the same as those of the UK and Italy but it begins to diverge towards the end of April. The difference in the growth paths then enables the implications of the Swedish soft lockdown policy to be assessed. The analysis suggests that, between 24 April and 22 July, there were perhaps 40–45 per cent more deaths than there might have been under a more stringent lockdown of the kind implemented in the UK and Italy.

Acknowledgements. Comments and suggestions from Leopoldo Catania, Stanley Cho, Jagjit Chadha, Radu Cristea, Paul Kattuman, Michael Höhle, Christopher Howe, Peter Kasson, Jonas Knecht, Rutger-Jan Lange, Daniel Mackay, Franco Peracchi, Jerome Simons, Craig Thamoheram, Herman van Dijk and a referee are gratefully acknowledged; of course they bear no responsibility for opinions expressed or mistakes made. Some of the ideas were presented at a Keynote talk at the (virtual) 40th International Symposium on Forecasting in October 2020 and at NIESR conferences in November 2020 and February 2021. The work was carried out as part of the University of Cambridge Keynes Fund project Persistence and Forecasting in Climate and Environmental Econometrics.

References

- Abadie, A., Diamond, A. and Hainmueller, J. (2010), 'Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program', *Journal of the American Statistical Association*, **105**, pp. 493–505.
- Abadie, A., Diamond, A. and Hainmueller, J. (2015), 'Comparative politics and the synthetic control method', *American Journal of Political Science*, **59**, pp. 495–510.
- Avery, C., Bossert, W., Clark, A., Ellison, G. and Ellison, S.F. (2020), 'An economist's guide to epidemiology models of infectious diseases', *Journal of Economic Perspectives*, **34**, pp. 79–104.
- Born, B., Dietrich, A.M. and Müller, G.J. (2020), 'Do lockdowns work? A counterfactual for Sweden', *Covid Economics*, **16**, pp. 1–22.
- Cho, S.-W. (2020), 'Quantifying the impact of nonpharmaceutical interventions during the COVID-19 outbreak: The case of Sweden', *Econometrics Journal*, **23**, pp. 323–44. <http://dx.doi.org/10.1093/ectj/utaa025>.
- Doornik, J.A., Castle, J.L. and Hendry, D.F. (2020), 'Short-term forecasting of the coronavirus pandemic', *International Journal of Forecasting*. <https://doi.org/10.1016/j.ijforecast.2020.09.003>.
- Durbin, J. and Koopman, S.J. (2012), *Time Series Analysis by State Space Methods*, Oxford: Oxford University Press.
- Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S. and Tatlow, H. (2021), 'A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)', *Nature Human Behaviour*, **5**, pp. 529–38.
- Harvey, A. and Kattuman, P. (2020), 'Time series models based on growth curves with applications to forecasting coronavirus', *Harvard Data Science Review*, Special Issue 1 - COVID-19, available online at <https://hdsr.mitpress.mit.edu/pub/ozgix0yn>.
- Harvey, A. and Kattuman, P. (2021), 'A farewell to R: Time series models for tracking and forecasting epidemics', *Journal of the Royal Society Interface*, forthcoming.
- Harvey, A., Kattuman, P. and Thamoheram, C. (2021), 'Tracking the mutant: Forecasting and nowcasting COVID-19 in the UK in 2021', *National Institute Economic Review*, **256**, pp. 110–26. <https://doi.org/10.1017/nie.2021.12>.
- Harvey, A.C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge: Cambridge University Press.
- Harvey, A.C. and Thiele, S. (2021), 'Co-integration and control: Assessing the impact of events using time series data', *Journal of Applied Econometrics*, **36**, pp. 71–85.
- Ioannidis, J.P.A., Cripps, S. and Tanner, M.A. (2020), 'Forecasting for COVID-19 has failed', *International Journal of Forecasting*. <https://doi.org/10.1016/j.ijforecast.2020.08.004>.
- Kamerlin, S.C.L. and Kasson, P.M. (2020), 'Managing coronavirus disease 2019 spread with voluntary public health measures: Sweden as a case study for pandemic control', *Clinical Infectious Diseases*, **71**, pp. 3174–81. <https://doi.org/10.1093/cid/ciaa864>.
- Koopman, S.J., Lit, R. and Harvey, A.C. (2021), *STAMP 9.00. Structural Time Series Analyser, Modeller and Predictor*, London: Timberlake Consultants.
- Kwiatkowski, D., Phillips, P.C.B., Schmidt P. and Shin, Y. (1992), 'Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?', *Journal of Econometrics*, **44**, pp. 159–78.

- Onder, G.** (2020). 'Case-fatality rate and characteristics of patients dying in relation to COVID-19 in Italy', *Journal of the American Medical Association*, **323**, pp. 1775–6.
- Stock, J. and Watson, M.** (1988), 'Testing for common trends', *Journal of the American Statistical Association*, **83**, pp. 1097–107.
- Wallinga, J. and Lipsitch, M.** (2007), 'How generation intervals shape the relationship between growth rates and reproductive numbers', *Proceedings of the Royal Society B*, **274**, pp. 599–604.

A. Appendix A. Data Sources

The data for European countries was obtained from the European Centre for Disease Prevention and Control (ECDC) website, <https://www.ecdc.europa.eu/en/publications-data/download-todays-data-geographic-distribution-covid-19-cases-worldwide>. For Florida the source was: <https://covidtracking.com/data>. The data were obtained at the end of August and the beginning of September. Data can be subject to revisions. For example, the UK definition of deaths was changed in August to include only people who had a laboratory-confirmed positive COVID-19 test and had died within 28 days of the date the test result was reported. Before that it included anybody who had ever tested positive for COVID-19 no matter how long before the actual death.

Case-fatality statistics in Italy are based on defining COVID-19-related deaths as those occurring in patients who test positive for SARS-CoV-2 via RTPCR, independently of pre-existing diseases that may have caused death. This method may have resulted in overestimation; see Onder (2020).

B. Appendix B. Transformations

The $\ln g_t$ transformation (LDL) is crucial in giving a series that stabilises the variability of the observations around the trend. The behaviour of $\ln y_t$ is similar when Y_t is large, but it may be quite different at the start of an epidemic. Other leading examples of statistical forecasting methods are based on different transformations. For example, Doornik *et al.* (2020) model g_t directly, but a comparison of g_t with its logarithm shows it to be much less stable in that its variability changes with the level. Figure A.1 shows these transformations for data on cases of coronavirus in Florida from 29 March to 19 July, as used in the leading indicator study in Section 4.2. Figure 4 reinforces the case for $\ln g_t$ by showing its downward trend and stability for UK and Italian deaths in the initial phase of the epidemic during the spring of 2020.

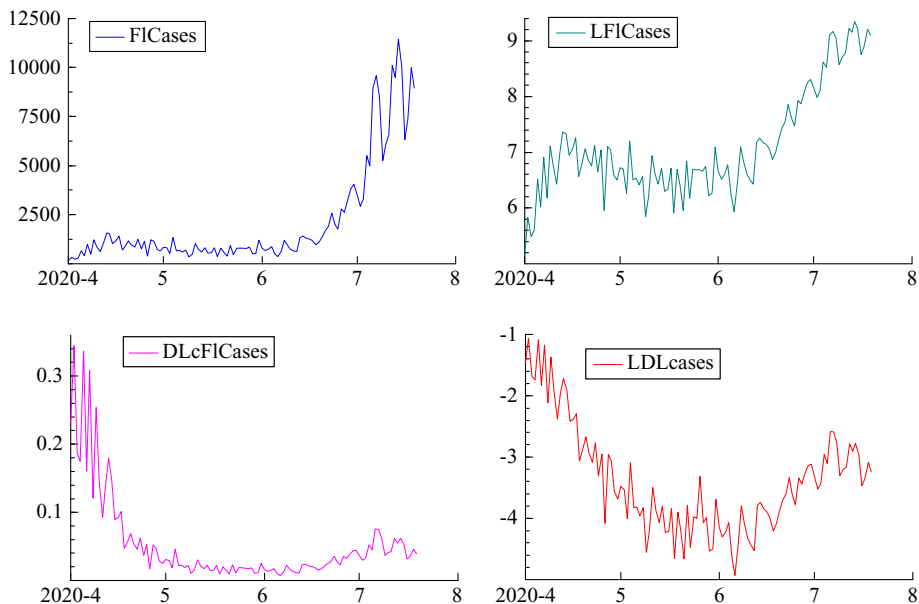


Figure A.1. (Colour online) Daily cases of coronavirus in Florida from 29 March to 19 July (top left hand graph), and its logarithm, $\ln y_t$, (top right hand), together with the growth rate of the cumulative total (lower left hand) and its logarithm, $\ln g_t$.