



Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Sparse HP filter: Finding kinks in the COVID-19 contact rate[☆]

Sokbae Lee^{a,b}, Yuan Liao^c, Myung Hwan Seo^{d,*}, Youngki Shin^e

^a Department of Economics, Columbia University, 420 West 118th Street, New York, NY 10027, USA

^b Centre for Microdata Methods and Practice, Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E 7AE, UK

^c Department of Economics, Rutgers University, 75 Hamilton St., New Brunswick, NJ 08901, USA

^d Department of Economics, Seoul National University, 1 Gwanak-ro, Gwanak-gu, Seoul 08826, Republic of Korea

^e Department of Economics, McMaster University, 1280 Main St. W., Hamilton, ON L8S 4L8, Canada



ARTICLE INFO

Article history:

Received 18 June 2020

Received in revised form 18 June 2020

Accepted 10 August 2020

Available online 26 September 2020

JEL classification:

C51

C52

C22

Keywords:

COVID-19

Trend filtering

Knots

Piecewise linear fitting

Hodrick–Prescott filter

ABSTRACT

In this paper, we estimate the time-varying COVID-19 contact rate of a Susceptible–Infected–Recovered (SIR) model. Our measurement of the contact rate is constructed using data on actively infected, recovered and deceased cases. We propose a new trend filtering method that is a variant of the Hodrick–Prescott (HP) filter, constrained by the number of possible kinks. We term it the *sparse HP filter* and apply it to daily data from five countries: Canada, China, South Korea, the UK and the US. Our new method yields the kinks that are well aligned with actual events in each country. We find that the sparse HP filter provides a fewer kinks than the ℓ_1 trend filter, while both methods fitting data equally well. Theoretically, we establish risk consistency of both the sparse HP and ℓ_1 trend filters. Ultimately, we propose to use time-varying *contact growth rates* to document and monitor outbreaks of COVID-19.

© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Since March 2020, there has been a meteoric rise in economic research on COVID-19. New research outputs have been appearing on the daily and weekly basis at an unprecedented level.¹ To sample a few, [Ludvigson et al. \(2020\)](#) quantified the macroeconomic impact of COVID-19 by using data on costly and deadly disasters in recent US history; [Manski and Molinari \(2020\)](#) and [Manski \(2020\)](#) applied the principle of partial identification to the infection rate and antibody tests, respectively; [Chernozhukov et al. \(2020\)](#) used the US state-level data to study determinants of social distancing behavior.

Across a wide spectrum of research, there is a rapidly emerging strand of literature based on a Susceptible–Infected–Recovered (SIR) model and its variants (e.g., [Hethcote, 2000](#), for a review of the SIR and related models). Many economists have embraced the SIR-type models as new tools to study the COVID-19 pandemic. [Avery et al. \(2020\)](#) provided a review of

[☆] We would like to thank the editor and an anonymous referee for helpful comments. This work is in part supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2018S1A5A2A01033487), the McMaster COVID-19 Research Fund (Stream 2), USA, the European Research Council (ERC-2014-CoG-646917-ROMIA) and the UK Economic and Social Research Council for research grant (ES/P008909/1) to the CeMMAP.

* Corresponding author.

E-mail addresses: sl3841@columbia.edu (S. Lee), yuan.liao@rutgers.edu (Y. Liao), myunghseo@snu.ac.kr (M.H. Seo), shiny11@mcmaster.ca (Y. Shin).

¹ The major outlets for economists are: arXiv working papers, NBER working papers, and CEPR's new working paper series called "Covid Economics: Vetted and Real-Time Papers" among others.

the SIR models for economists, calling for new research in economics. A variety of economic models and policy simulations have been built on the SIR-type models. See [Acemoglu et al. \(2020\)](#), [Alvarez et al. \(2020\)](#), [Atkeson \(2020\)](#), [Eichenbaum et al. \(2020\)](#), [Pindyck \(2020\)](#), [Stock \(2020\)](#), [Kim et al. \(2020\)](#), and [Toda \(2020\)](#) among many others.

One of central parameters in the SIR-type models is the contact rate, typically denoted by β .² It measures “the average number of adequate contacts (i.e., contacts sufficient for transmission) of a person per unit time” ([Hethcote, 2000](#)). The contact number β/γ is the product between β and the average infectious period, denoted by $1/\gamma$; the contact number is interpreted as “the average number of adequate contacts of a typical infective during the infectious period” ([Hethcote, 2000](#)).

The goal of this paper is to estimate the time-varying COVID-19 contact rate, say β_t . In canonical SIR models, β is a time-constant parameter. However, it may vary over time due to multiple factors. For example, as pointed by [Stock \(2020\)](#), self-isolation, social distancing and lockdown may reduce β . To estimate a SIR-type model, [Fernández-Villaverde and Jones \(2020\)](#) allowed a time-varying contact rate to reflect behavioral and policy-induced changes associated with social distancing. In particular, they estimated β_t using data on deaths at city, state and country levels. Their main focus was to simulate future outcomes for many cities, states and countries.

Researchers have also adopted nonlinear time-series models from the econometric toolbox. For example, [Li and Linton \(2020\)](#) analyzed the daily data on the number of new cases and the number of new deaths with a quadratic time trend model in logs. Their main purpose was to estimate the peak of the pandemic. [Liu et al. \(2020\)](#) studied the density forecasts of the daily number of active infections for a panel of countries/regions. They modeled the growth rate of active infections as autoregressive fluctuations around a piecewise linear trend with a single break. [Hartl et al. \(2020\)](#) used a linear trend model in logs with a trend break to fit German confirmed cases. [Harvey and Kattuman \(2020\)](#) used a Gompertz model with a time-varying trend to fit and forecast German and UK new cases and deaths.

In this paper, we aim to synthesize the time-varying contact rate with nonparametric time series modeling. Especially, we build a new nonparametric regression model for β_t that allows for a piecewise linear trend with multiple kinks at unknown dates. We analyze daily data from Johns Hopkins University Center for Systems Science and Engineering ([Dong et al., 2020](#), JHU CSSE) and suggest a particular transformation of data that can be regarded as a noisy measurement of time-varying β_t . Our measurement of β_t , which is constructed from daily data on confirmed, recovered and deceased cases, is different from that of [Fernández-Villaverde and Jones \(2020\)](#) who used only death data. We believe both measurements are complements to each other. However, the SIR model is at best a first-order approximation to the real world; a raw series of β_t would be too noisy to draw on inferences regarding the underlying contact rate. In fact, the raw series exhibits high degrees of skewness and time-varying volatility even after the log transformation.

To extract the time-varying signal from the noisy measurements, we consider nonparametric trend filters that produce possibly multiple kinks in β_t where the kinks are induced by government policies and changes in individual behavior. A natural candidate method that yields the kinks is ℓ_1 trend filtering (e.g., [Kim et al., 2009](#)). However, ℓ_1 trend filtering is akin to LASSO; hence, it may have a problem of producing too many kinks, just like LASSO selects too many covariates. In view of this concern, we propose a novel filtering method by adding a constraint on the maximum number of kinks to the popular [Hodrick and Prescott \(1997\)](#) (HP) filter. It turns out that this method produces a smaller number of the kink points than ℓ_1 trend filtering when both methods fit data equally well. In view of that, we call our new method the *sparse HP filter*. We find that the estimated kinks are well aligned with actual events in each country. To document and monitor outbreaks of COVID-19, we propose to use piecewise constant *contact growth rates* using the piecewise linear trend estimates from the sparse HP filter. They provide not only an informative summary of past outbreaks but also a useful surveillance measure.

The remainder of the paper is organized as follows. In Section 2, we describe a simple time series model of the time-varying contact rate. In Section 3, we introduce two classes of filtering methods. In Section 4, we have a first look at the US data, as a benchmark country. In Section 5, we present empirical results for five countries: Canada, China, South Korea, the UK and the US. In Section 6, we establish risk consistency of both the sparse HP and ℓ_1 trend filters. Section 7 concludes and appendices include additional materials. The replication R codes for the empirical results are available at <https://github.com/yshin12/sparseHP>. Finally, we add the caveat that the empirical analysis in the paper was carried out in mid-June using daily observations up to June 8th. As a result, some remarks and analysis might be out of sync with the COVID-19 pandemic in real time.

2. A time series model of the COVID-19 contact rate

In this section, we develop a time-series model of the contact rate. Our model specification is inspired by the classical SIR model which has been adopted by many economists in the current coronavirus pandemic.

We start with a discrete version of the SIR model, augmented with deaths, adopted from [Pindyck \(2020\)](#):

$$\begin{aligned}\Delta I_t &= \beta S_{t-1} I_{t-1} - \gamma I_{t-1}, \\ \Delta D_t &= \gamma_d I_{t-1}, \\ \Delta R_t &= \gamma_r I_{t-1}, \\ 1 &= S_t + I_t + D_t + R_t, \\ \gamma &= \gamma_r + \gamma_d,\end{aligned}\tag{2.1}$$

² It is also called the transmission rate by [Stock \(2020\)](#).

where the (initial) population size is normalized to be 1, S_t is the proportion of the population that is susceptible, I_t the fraction infected, D_t the proportion that has died, and R_t the fraction that has recovered. The parameter $\gamma = \gamma_r + \gamma_d$ governs the rate at which infectives transfer to the state of being deceased or recovered.

In the emerging economics literature on COVID-19, the contact rate β is viewed as the parameter that can be affected by changes in individual behavior and government policies through social distancing and lockdown. We follow this literature and let $\beta = \beta_t$ be time-varying.

Let C_t be the proportion of confirmed cases, that is $C_t = I_t + R_t + D_t$. In words, the confirmed cases consist of actively infected, recovered and deceased cases. Use the equations in (2.1) to obtain

$$\beta_t = Y_t := \frac{\Delta C_t}{I_{t-1} S_{t-1}}. \quad (2.2)$$

Assume that we have daily data on ΔC_t , ΔR_t and ΔD_t . From these, we can construct cumulative C_t , R_t and D_t . Then $S_t = 1 - C_t$ and $I_t = C_t - R_t - D_t$. This means that we can obtain time series of β_t from Y_t . We formally assume this in the following.

Assumption 1 (Data). For each t , we observe (C_t, R_t, D_t) .

By **Assumption 1**, we can construct $Y_t = \Delta C_t / (I_{t-1} S_{t-1})$. **Assumption 1** is a key assumption in the paper. We use daily data from JHU CSSE and they are subject to measurement errors, which could bias our estimates. In **Appendix A**, we show that the time series model given in this section is robust to some degree of under-reporting of confirmed cases. However, our estimates are likely to be biased if the underreporting is time-varying. For example, this could happen because testing capacity in many countries has expanded over the time period. Nonetheless, we believe that our measurement of Y_t primarily captures the genuine underlying trend of β_t . Moreover, because the SIR model in (2.1) is at best a first-order approximation, a raw series of Y_t would be too noisy to be used as the actual series of the underlying contact rate β_t . In other words, $\beta_t \neq Y_t$ in actual data and it would be natural to include an error term in Y_t . Because β_t has to be positive, we adopt a multiplicative error structure and make the following assumption.

Assumption 2 (Time-Varying Signal plus Regression Error). For each t , the unobserved random variable β_t satisfies

$$\log Y_t = \log \beta_t + u_t,$$

where the error term u_t has the following properties:

1. $\mathbb{E}[u_t | \mathcal{F}_{t-1}] = 0$, where \mathcal{F}_{t-1} is the natural filtration at time $t - 1$,
2. $\mathbb{E}[u_t^2 | \mathcal{F}_{t-1}] = \sigma_t^2 > 0$ for some time-varying conditional variance σ_t^2 .

Define

$$y_t := \log(\Delta C_t) - \log(I_{t-1}) - \log S_{t-1}. \quad (2.3)$$

Under **Assumption 2**, (2.2) can be rewritten as

$$y_t = \log \beta_t + u_t, \quad (2.4)$$

The time-varying parameter $\log \beta_t$ would not be identified without further restrictions. Because it is likely to be affected by government policies and cannot change too rapidly, we will assume that it follows a piecewise trend:

Assumption 3 (Piecewise Trend). The time-varying parameter $f_{0,t} := \log \beta_t$ follows a piecewise trend with at most κ kinks, where the set of kinks is defined by $\{t = 1, \dots, T : f_{0,t} - f_{0,t-1} \neq f_{0,t+1} - f_{0,t}\}$ and the locations of kinks are unknown.

The main goal of this paper is to estimate $\log \beta_t$ and its kinks under **Assumptions 1, 2 and 3**.

3. Filtering the COVID-19 contact rate

We consider two different classes of trend filtering methods to produce piecewise estimators of $f_{0,t} := \log \beta_t$. The first class is based on ℓ_1 trend filtering, which has become popular recently. See, e.g., **Kim et al. (2009)**, **Tibshirani (2014)**, and **Wang et al. (2016)** among others.

The starting point of the second class is the HP filter, which has been popular in macroeconomics and has been frequently used to separate trend from cycle. The standard convention in the literature is to set $\lambda = 1600$ for quarterly time series. For example, **Ravn and Uhlig (2002)** suggested a method for adjusting the HP filter for the frequency of observations; **de Jong and Sakarya (2016)** and **Cornea-Madeira (2017)** established some representation results; **Hamilton (2018)** provided criticism on the HP filter; **Phillips and Shi (2019)** advocated a boosted version of the HP filter via L_2 -boosting (**Bühlmann and Yu, 2003**) that can detect multiple structural breaks. We view that the kinks might be more suitable than the breaks for modeling β_t using daily data. It is unlikely that in a few days, the degree of contagion of COVID-19 would be diminished with an abrupt jump by social distancing and lockdown. The original HP filter cannot produce any kink just as ridge regression does not select any variable. We build the sparse HP filter by drawing on the recent literature that uses an ℓ_0 -constraint or penalty (see, e.g. **Bertsimas et al., 2016**; **Chen and Lee, 2018, 2020**; **Huang et al., 2018**).

3.1. ℓ_1 Trend filtering

In ℓ_1 trend filtering, the trend estimate f_t is a minimizer of

$$\sum_{t=1}^T (y_t - f_t)^2 + \lambda \sum_{t=2}^{T-1} |f_{t-1} - 2f_t + f_{t+1}|, \tag{3.1}$$

which is related to Hodrick and Prescott (1997) filtering; the latter is the minimizer of

$$\sum_{t=1}^T (y_t - f_t)^2 + \lambda \sum_{t=2}^{T-1} (f_{t-1} - 2f_t + f_{t+1})^2. \tag{3.2}$$

In this paper, the main interest is to find the kinks in the trend. For that purpose, ℓ_1 trend filtering is more suitable than the HP filtering. The main difficulty of using (3.1) is the choice of λ . This is especially challenging since the time series behavior of y_t is largely unknown.

The ℓ_1 trend filter is akin to LASSO. In view of an analogy to square-root LASSO (Belloni et al., 2011), it might be useful to consider a square-root variant of (3.1):

$$\left(\sum_{t=1}^T (y_t - f_t)^2 \right)^{1/2} + \lambda \sum_{t=2}^{T-1} |f_{t-1} - 2f_t + f_{t+1}|. \tag{3.3}$$

We will call (3.3) square-root ℓ_1 trend filtering. Both (3.1) and (3.3) can be solved via convex optimization software, e.g., CVXR (Fu et al., 2017).

3.2. Sparse hodrick–prescott trend filtering

As an alternative to ℓ_1 trend filtering, we may exploit Assumption 3 and consider an ℓ_0 -constrained version of trend filtering:

$$\begin{aligned} & \sum_{t=1}^T (y_t - f_t)^2 \\ & \text{subject to} \\ & \sum_{t=2}^{T-1} 1\{f_t - f_{t-1} \neq f_{t+1} - f_t\} \leq \kappa. \end{aligned} \tag{3.4}$$

The formulation in (3.4) is related to the method called best subset selection (see, e.g. Bertsimas et al., 2016; Chen and Lee, 2018). It requires only the input of κ . However, because of the nature of the ℓ_0 -(pseudo)norm, it would not work well if the signal-to-noise ratio (SNR) is low (Hastie et al., 2017; Mazumder et al., 2017). This is likely to be a concern for our measurement of the log contact rate.

To regularize the best subset selection procedure, it has been suggested in the literature that (3.4) can be combined with ℓ_1 or ℓ_2 penalization (Bertsimas and Van Parys, 2020; Mazumder et al., 2017). We adopt Bertsimas and Van Parys (2020) and propose an ℓ_0 -constrained version of the Hodrick–Prescott filter:

$$\begin{aligned} & \sum_{t=1}^T (y_t - f_t)^2 + \lambda \sum_{t=2}^{T-1} (f_{t-1} - 2f_t + f_{t+1})^2 \\ & \text{subject to} \\ & \sum_{t=2}^{T-1} 1\{f_t - f_{t-1} \neq f_{t+1} - f_t\} \leq \kappa. \end{aligned} \tag{3.5}$$

As in (3.4), the tuning parameter κ controls how many kinks are allowed for. Thus, we have a direct control of the resulting segments of different slopes. The ℓ_2 penalty term is useful to deal with the low SNR problem with the COVID 19 data. We will call (3.5) sparse HP trend filtering.

Problem (3.5) can be solved by mixed integer quadratic programming (MIQP). Rewrite the objective function in (3.5) as

$$\sum_{t=1}^T (y_t - f_t)^2 + \lambda \sum_{t=2}^{T-1} (f_{t-1} - 2f_t + f_{t+1})^2$$

subject to $z_t \in \{0, 1\}$, $t = 2, \dots, T - 1$, $\underline{f} \leq f_t \leq \bar{f}$, $\sum_{t=2}^{T-1} z_t \leq \kappa$, and

$$-Mz_t \leq f_{t-1} - 2f_t + f_{t+1} \leq Mz_t, \quad t = 2, \dots, T - 1.$$

This is called a big-M formulation that requires that

$$\max_t |f_{t-1} - 2f_t + f_{t+1}| \leq M.$$

We need to choose the auxiliary parameters \underline{f}, \bar{f} and M . We set $\underline{f} = \min y_t$ and $\bar{f} = \max y_t$. One simple practical method for choosing M is to set

$$M = \max_{t=2, \dots, T-1} |y_{t-1} - 2y_t + y_{t+1}|. \tag{3.6}$$

To implement the proposed method, it is simpler to write the MIQP problem above in matrix notation. Let \mathbf{y} denote the $(T \times 1)$ vector of y_t 's and $\mathbf{1}$ a vector of 1's whose dimension may vary. We solve

$$\min_{\mathbf{f}, \mathbf{z}} [(\mathbf{y} - \mathbf{f})^\top (\mathbf{y} - \mathbf{f}) + \lambda \mathbf{f}^\top \mathbf{D}^\top \mathbf{D} \mathbf{f}] \tag{3.7}$$

subject to $\mathbf{z} \in \{0, 1\}^{T-2}, \underline{f} \mathbf{1} \leq \mathbf{f} \leq \bar{f} \mathbf{1}, \mathbf{1}^\top \mathbf{z} \leq \kappa, -M\mathbf{z} \leq \mathbf{D}\mathbf{f} \leq M\mathbf{z}$, where \mathbf{D} is the $(T - 2) \times T$ second-order difference matrix such that

$$\mathbf{D} = \begin{bmatrix} 1 & -2 & 1 & & & & \\ & 1 & -2 & 1 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & 1 & -2 & 1 & \\ & & & & 1 & -2 & 1 \end{bmatrix}$$

with entries not shown above being zero. Let $\hat{\mathbf{f}}$ and $\hat{\mathbf{z}}$ denote the resulting maximizers. It is straightforward to see that $\hat{\mathbf{f}}$ also solves (3.5). Therefore, $\hat{\mathbf{f}}$ is the $(T \times 1)$ vector of trend estimates and $\hat{K} := \{t = 2, \dots, T - 1 : \hat{z}_t = 1\}$ is the index set of estimated kinks. The MIQP problem can be solved via modern mixed integer programming software, e.g., **Gurobi**. Because the sample size for y_t is typically less than 100, the computational speed of MIQP is fast enough to carry out cross-validation to select tuning parameters. We summarize the equivalence between the original and MIQP formulation in the following proposition.

Proposition 3.1. Define

$$\mathbb{F}(\kappa) := \{\mathbf{f} = (f_1, \dots, f_T) : \min_t f_t \geq \underline{f}, \max_t f_t \leq \bar{f}, \max_t |f_{t-1} - 2f_t + f_{t+1}| \leq M, \sum_{t=2}^{T-1} 1\{f_t - f_{t-1} \neq f_{t+1} - f_t\} \leq \kappa\}.$$

Let $\hat{\mathbf{f}}_{SHP} : \{\hat{f}_t : t = 1, \dots, T\}$ denote a solution to

$$\min_{\mathbf{f} \in \mathbb{F}(\kappa)} S_T(\mathbf{f}, \lambda) := (\mathbf{y} - \mathbf{f})^\top (\mathbf{y} - \mathbf{f}) + \lambda \mathbf{f}^\top \mathbf{D}^\top \mathbf{D} \mathbf{f}.$$

Let $\hat{\mathbf{f}}_{MIQP}$ denote a solution to (3.7). Then, both $\hat{\mathbf{f}}_{SHP}$ and $\hat{\mathbf{f}}_{MIQP}$ are equivalent in the sense that $\hat{\mathbf{f}}_{SHP} \in \mathbb{F}(\kappa), \hat{\mathbf{f}}_{MIQP} \in \mathbb{F}(\kappa)$, and $S_T(\hat{\mathbf{f}}_{SHP}, \lambda) = S_T(\hat{\mathbf{f}}_{MIQP}, \lambda)$.

3.3. Selection of tuning parameters

We first consider the sparse HP filter. There are two tuning parameters: λ and κ . It is likely that there will be an initial stage of coronavirus spread, followed by lockdown or social distancing. Even without any policy intervention, it will come down since many people will voluntarily select into self-isolation and there is a chance of herd immunity. Hence, the minimum κ is at least 1. If κ is too large, it is difficult to interpret the resulting kinks. In view of these, we set the possible values $\kappa \in \mathcal{K} = \{2, 3, 4\}$. For each pair of (κ, λ) , let $\hat{\mathbf{f}}_{-s}(\kappa, \lambda)$ denote the leave-one-out estimator of \mathbf{f}_s . That is, it is the sparse HP filter estimate by solving:

$$\begin{aligned} & \sum_{t=1, t \neq s}^T (y_t - f_t)^2 + \lambda \sum_{t=2}^{T-1} (f_{t-1} - 2f_t + f_{t+1})^2 \\ & \text{subject to} \\ & \sum_{t=2}^{T-1} 1\{f_t - f_{t-1} \neq f_{t+1} - f_t\} \leq \kappa. \end{aligned} \tag{3.8}$$

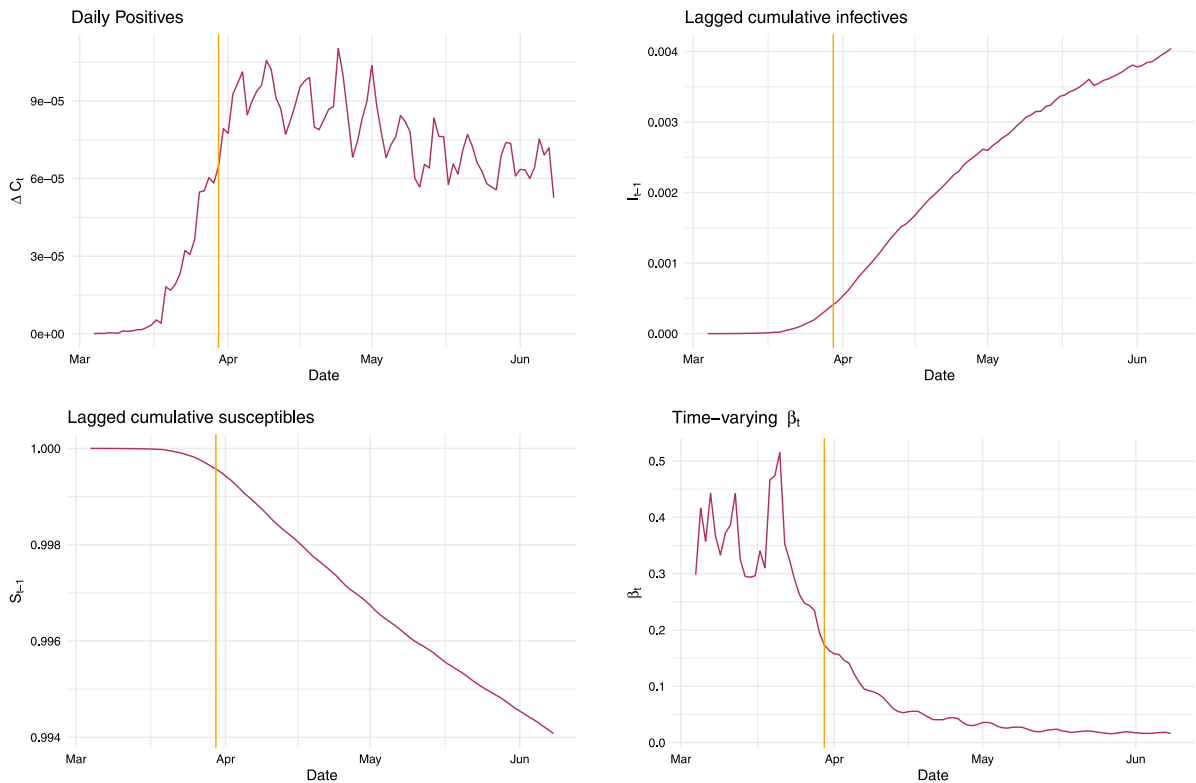


Fig. 1. US Data. Note: The orange vertical line denotes the lockdown date, March 30. The population size is normalized to be 1.

The only departure from (3.5) is that we replace the fidelity term $\sum_{t=1}^T (y_t - f_t)^2$ with $\sum_{t=1, t \neq s}^T (y_t - f_t)^2$. We choose the optimal (κ, λ) by

$$\min_{(\kappa, \lambda) \in \mathcal{K} \times \mathcal{L}} \sum_{t=1}^T \{y_t - \hat{f}_{-t}(\kappa, \lambda)\}^2, \tag{3.9}$$

where \mathcal{L} is the set for possible values of λ . We view λ as an auxiliary tuning parameter that mitigates the low SNR problem. Hence, we take \mathcal{L} to be in the range of relatively smaller values than the typical values used for the HP filter. In the numerical work, we let \mathcal{L} to a grid of equi-spaced points in the \log_2 -scale.

We now turn to the HP, ℓ_1 and square-root ℓ_1 trend filters. For each filter, we choose λ such that the fidelity term $\sum_{t=1}^T (y_t - f_t)^2$ is the same as that of the sparse HP filter. In this way, we can compare different methods holding the same level of fitting the data. Alternatively, we may choose λ by leave-one-out cross validation for each filtering method. However, in that case, it would be more difficult to make a comparison across different methods. Since our main focus is to find the kinks in the contact rate, we will fine-tune all the filters to have the same level of $\sum_{t=1}^T (y_t - f_t)^2$ based on the sparse HP filter’s cross validation result.

4. A first look at the time-varying contact rate

As a benchmark, we have a first look at the US data. The dataset is obtained via R package **coronavirus** (Krispin, 2020), which provides a daily summary of COVID-19 cases from Johns Hopkins University Center for Systems Science and Engineering (Dong et al., 2020). Following Liu et al. (2020), we set the first date of the analysis to begin when the number of cumulative cases reaches 100 (that is, March 4 for the US). To smooth data minimally, we take Y_t in (2.2) to be a three-day simple moving average: that is, $Y_t = (\hat{Y}_t + \hat{Y}_{t-1} + \hat{Y}_{t-2})/3$, where \hat{Y}_t is the daily observation of Y_t constructed from the dataset.³ Then, we take the log to obtain $y_t = \log Y_t$.

Fig. 1 has four panels. The top-left panel shows the fraction of daily positives, the top-right panel the fraction of lagged cumulative infectives, the bottom-left panel the fraction of lagged cumulative susceptibles, and the bottom-right $Y_t = \Delta C_t / (I_{t-1} S_{t-1})$. In the US, statewide stay-at-home orders started in California on March 20 and extended to 30 states

³ Liu et al. (2020) used one-sided three-day rolling averages; Fernández-Villaverde and Jones (2020) took 5-day centered moving averages.

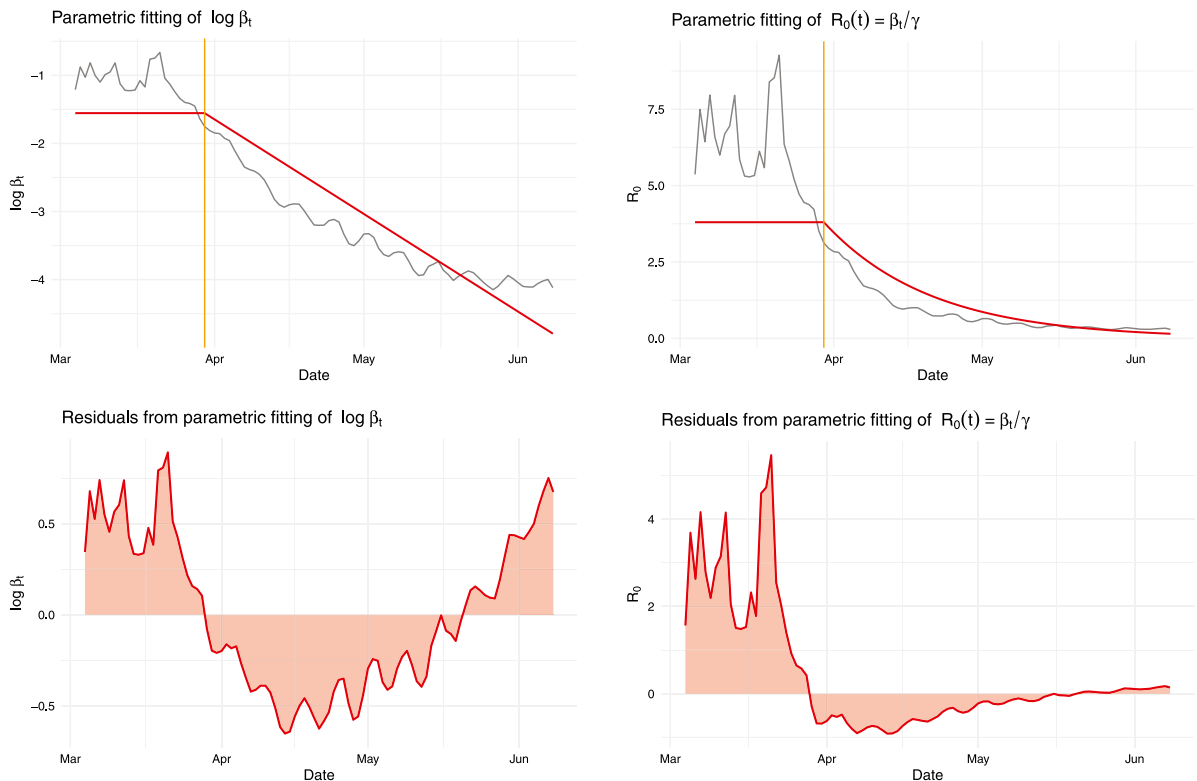


Fig. 2. $\log Y_t$ as a raw time series of $\log \beta_t$ and parametric fitting. Note: The orange vertical line denotes the lockdown date, March 30.

by March 30 (The New York Times, 2020d). The inserted vertical line in the figure corresponds to March 30, which we will call the “lockdown” date for simplicity, although there was no lockdown at the national level. As a noisy measurement of β_t , Y_t shows enormous skewness and fluctuations especially in the beginning of the study period. This indicates that the signal-to-noise ratio is high and is time-varying as well. This pattern of the data has motivated Assumption 2. Because S_{t-1} is virtually one throughout the analysis period (0.994 on June 8, which is the last date of the sample), $Y_t \approx \Delta C_t / I_{t-1}$, which is daily positives divided by the lagged infectives.

Fig. 2 shows the raw data along with parametric fitting. The top-left panel shows the logarithm of Y_t , which still exhibits some degree of skewness and time-varying variance. The fitted regression line is based on the following parametric regression model:

$$y_t = \alpha_0 + \alpha_1(t - t_0)1(t > t_0) + \varepsilon_t, \tag{4.1}$$

where t_0 is March 30. The simple idea behind (4.1) is that an initial, time-constant contact rate began to diminish over time after a majority of US states imposed stay-at-home orders.

In simple SIR models, the contact number β/γ is identical to the basic reproduction number denoted by R_0 , which is viewed as a key threshold quantity in the sense that “an infection can get started in a fully susceptible population if and only if $R_0 > 1$ for many deterministic epidemiology models” (Hethcote, 2000). Since β_t is time-varying in our framework, we may define a time-varying basic reproduction number by $R_0(t) := \beta_t/\gamma$.

The top-right panel shows the estimates of time-varying $R_0(t)$ ⁴:

$$\widehat{R}_0(t) := \exp[\widehat{\alpha}_0 + \widehat{\alpha}_1(t - t_0)1(t > t_0)]/\gamma, \tag{4.2}$$

where $\gamma = 1/18$ is taken from Acemoglu et al. (2020). This corresponds to 18 days of the average infectious period. The parametric estimates of $R_0(t)$ started above 4 and reached 0.15 at the end of the sample period.

The left-bottom panel shows the residual plot in terms of y_t and the right-bottom panel the residual plot in terms of $R_0(t)$. In both panels, the estimated residuals seem to be biased and show autocorrelation. Especially, the positive values of residuals at the end of the sample period are worrisome because the resulting prediction would be too optimistic.

⁴ The formula given in (4.2) is valid if errors are homoskedastic, which is unlikely to be true in actual data. However, we present (4.2) here because it is simpler. Our main analysis focuses on estimation of the kinks based on y_t , not on estimating $R_0(t)$. We use the latter mainly to appreciate the magnitude of the kinks.

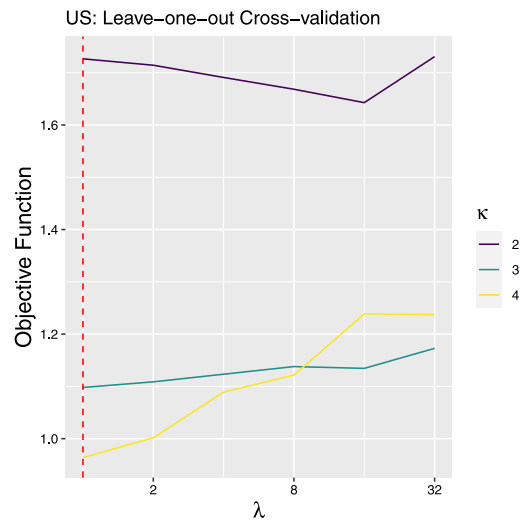


Fig. 3. Sparse HP Filtering: Leave-One-Out Cross Validation. Note: The red vertical line denotes the minimizer $(\hat{\kappa}, \hat{\lambda}) = (4, 1)$ of the cross-validation objective function. The x-axis is represented by the \log_2 scale.

5. Estimation results

In this section, we present estimation results for five countries: Canada, China, South Korea, the UK and the US. These countries are not meant to be a random sample of the world; they are selected based on our familiarity with them so that we can interpret the estimated kinks with narratives. We look at the US as a benchmark country and provide a detailed analysis in Section 5.1. A condensed version of the estimation results for other countries is provided in Section 5.2.

5.1. Benchmark: the US

Fig. 3 summarizes the results of leave-one-out cross validation (LOOCV) as described in Section 3.3. The range of tuning parameters was: $\kappa \in \{2, 3, 4\}$ and $\lambda = \{2^0, 2^1, \dots, 2^5\}$. We can see that the choice of κ seems to matter more than that of λ . Clearly, $\kappa = 2$ provides the worst result and $\kappa = 3$ and $\kappa = 4$ are relatively similar. The LOOCV criterion function was minimized at $(\hat{\kappa}, \hat{\lambda}) = (4, 1)$.

Based on the tuning parameter selection in Fig. 3, we show estimation results for the sparse HP filter in Fig. 4. The structure of Fig. 4 is similar to that of Fig. 2. The top-left panel shows estimates of the sparse HP filter along with the raw series of y_t and the parametric estimates shown in Fig. 2. The top-right panel displays counterparts in terms of $R_0(t)$. The bottom panels exhibit residual plots for the $\log \beta_t$ and $R_0(t)$ scales. The trend estimates from the sparse HP filter fit the data much better than the simple parametric estimates. The estimated kink dates are: March 16, March 20, April 14, and May 13. There are five periods based on them.

1. March 4–March 16: this period corresponds to the initial epidemic stage;
2. March 16–March 20: the contact rate peaked at the end of this period;
3. March 20–April 14: a sharp decrease of the contact rate is striking;
4. April 14–May 13: the contact rate decreased but less steeply;
5. May 13–June 8: it continued to go down but its slope got more flattened.

To provide narratives on these dates, President Trump declared a national emergency on March 13; The Centers for Disease Control and Prevention (CDC) recommended no gatherings of 50 or more people on March 15; New York City's public schools system announced that it would close on March 16; and California started stay-at-home orders on March 20 (The New York Times, 2020b,d). These events indicate that the second period was indeed the peak of the COVID-19 epidemic in the US. The impact of social distancing and stay-at-home orders across a majority of states is clearly visible in the third period. The fourth and fifth periods include state reopening: for example, stay-at-home order expired in Georgia and Texas on April 30; in Florida on May 4; in Massachusetts on May 18; in New York on May 28 (The New York Times, 2020c). In short, unlike the parametric model with a single kink, the nonparametric trend estimates detect multiple changes in the slopes and provide kink dates, which are well aligned with the actual events.

We now turn to different filtering methods. In Fig. 5, we show selection of λ for the HP, ℓ_1 and square-root ℓ_1 filters. As explained in Section 3.3, the penalization parameter λ is chosen to ensure that all different methods have the same level of fitting the data. Fig. 6 shows the estimation results for the HP filter. The HP trend estimates trace data pretty well




Fig. 4. Sparse HP Filtering. Note: The gray curve in each panel represents the data y_t . Sparse HP filtering solves (3.5) and the parametric fit uses the linear regression (4.1). The estimated kinks denoted by blue vertical lines are: March 16, March 20, April 14, and May 13. The orange vertical line denotes the lockdown date, March 30.




Fig. 5. Selection of λ . Note: The tuning parameter λ is chosen by minimizing the distance between two fidelities as described in Section 3.3. The selected tuning parameters for HP, ℓ_1 , and square-root ℓ_1 are 30, 0.9, and 0.5, respectively.

after late March, as clear in residual plots. However, there is no kink in the estimates due to the nature of the ℓ_2 penalty term in the HP filter. The tuning parameter was $\lambda = 30$, which is 30 times as large as the one used in the sparse HP filter. This is because for the HP filter, λ is the main tuning parameter; however, for the sparse HP filter, λ plays a minor role of regularizing the ℓ_0 constrained method.

In Fig. 7, we plot estimation results using ℓ_1 trend filtering. The results look similar to those in Fig. 6, but there are now 10 kink points: March 7, March 15, March 16, March 20, March 21, March 30, April 14, April 21, May 12, May 27.

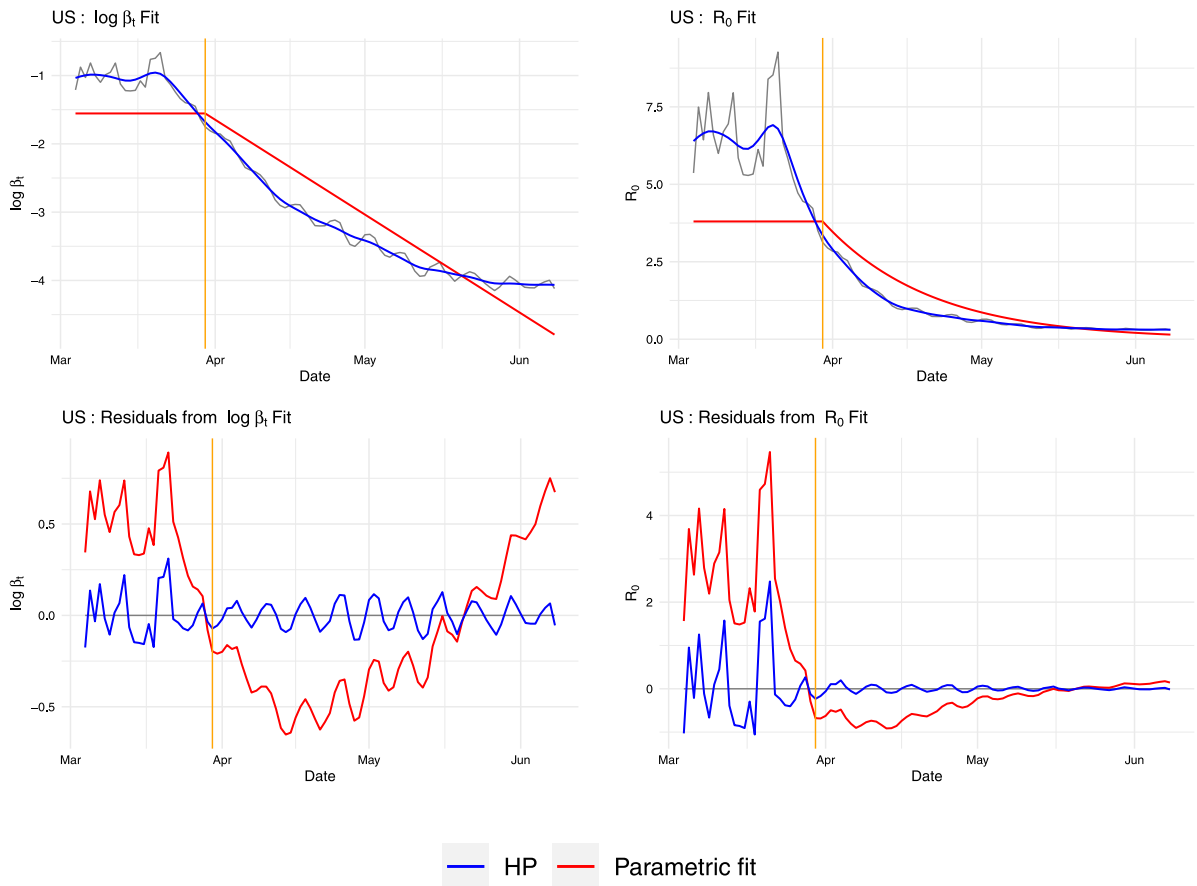


Fig. 6. HP Filtering. Note: HP filtering solves Eq. (3.2). The orange vertical line denotes the lockdown date, March 30.

They are dates t such that $|\Delta^2 \log \hat{\beta}_t| > \eta$, where Δ^2 is the double difference operator and $\eta = 10^{-6}$ is an effective zero.⁵ The tuning parameter $\lambda = 0.9$ was chosen by minimizing the distance between the fidelity of ℓ_1 and that of the Sparse HP. Recall that the sparse HP filter produces the kinks on March 16, March 20, April 14, and May 13. In other words, the ℓ_1 filter estimates 6 more kinks than the sparse HP filter when both fit the data equally well. It is unlikely that two adjacent dates (March 15–16 and March 20–21) correspond to two different regimes in the time-varying contact rate. This suggests that the ℓ_1 filter may over-estimate the number of kinks. Fig. 8 shows estimation results for the square-root ℓ_1 trend filters. The chosen $\lambda = 0.5$ was smaller than that of the ℓ_1 trend filter due to the change in the scale of the fidelity term; however, the trend estimates look very similar and the estimated kinks are identical between the ℓ_1 and square-root ℓ_1 trend filters. In Fig. 9, we plot the sparse HP filter estimates along with ℓ_1 filter estimates. Both methods have produced very similar trend estimates, but the number of kinks is substantially different: only 4 kinks for the sparse HP filter but 10 kinks for the ℓ_1 filter.

5.2. Other countries: Canada, China, South Korea and the UK

In this section, we provide condensed estimation results for other countries. We focus on the sparse HP and ℓ_1 filters whose tuning parameters are chosen as in the previous section. Appendices B and C contain the details of the selection of tuning parameters.

Fig. 10 shows the empirical results of Canada. The estimated kink dates are: March 18 and April 11. Based on them, we can classify observations into three periods:

1. March 6–March 18: This is an initial period of the epidemic in Canada. The contact rate peaked at the end of this period. Several lockdown measures started to be imposed.

⁵ The results are robust to the size of the effective zero and do not change even if we set $\eta = 10^{-3}$. Gurobi used for the Sparse HP filtering also imposes some effective zeros in various constraints. We use the default values of them. For example, the integer tolerance level and the general feasibility tolerance level are 10^{-5} and 10^{-6} , respectively.

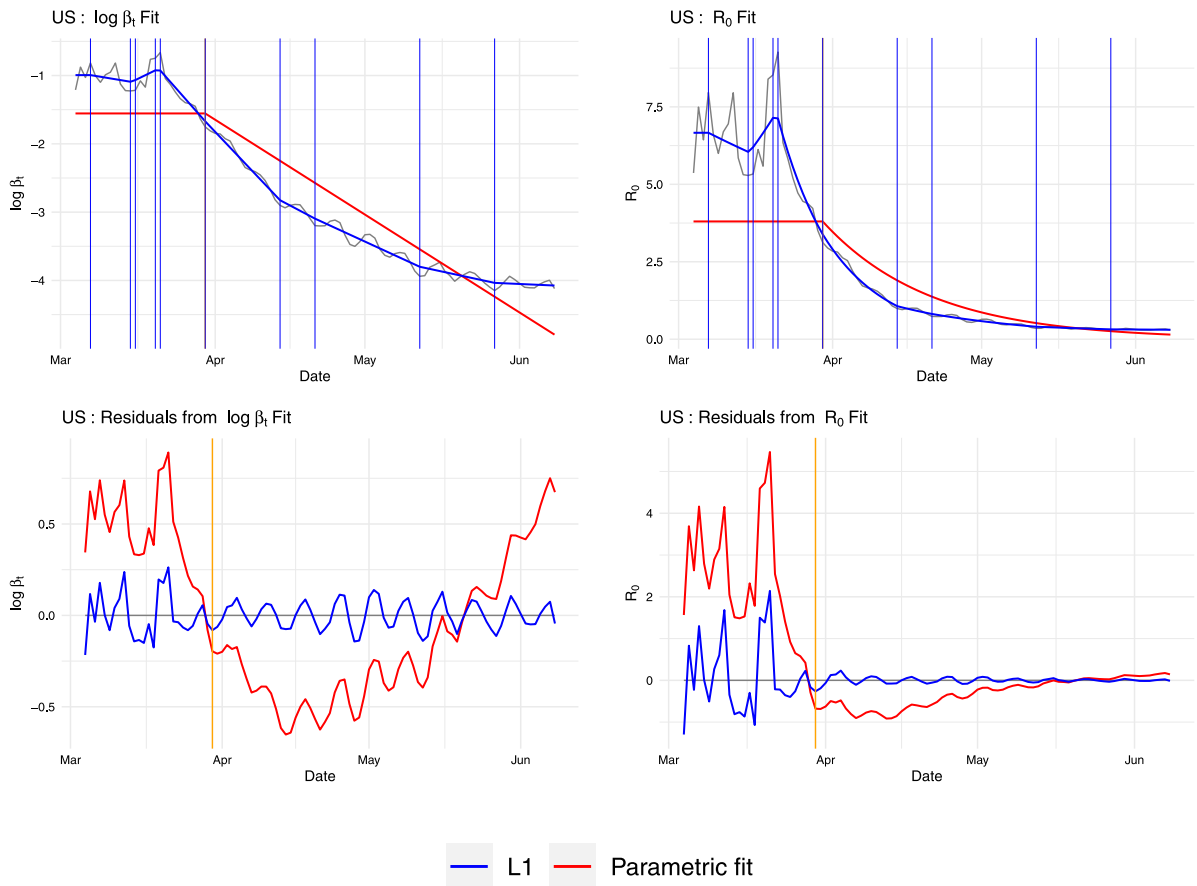


Fig. 7. ℓ_1 Filtering. Note: ℓ_1 filtering solves Eq. (3.1). The estimated kinks denoted by blue vertical lines are: March 7, March 15, March 16, March 20, March 21, March 30, April 14, April 21, May 12, and May 27. The orange vertical line denotes the lockdown date, March 30. The ℓ_1 -filtering kink dates are calculated by any t such that $|\Delta^2 \log \beta_t| > \eta$, where $\eta = 10^{-6}$ is an effective zero.

2. March 18–April 11: We observe a sharp decrease in the contact rate in this period. Additional measures were imposed.
3. April 11–June 8: The contact rate decreased but less steeply.

Quebec and Ontario are the two provinces hardest hit by COVID-19. In Quebec, daycares, public schools, and universities are closed on March 13 followed by non-essential businesses and public gathering places on March 15. Montreal declared state of emergency on March 27 (CTV News, 2020). Similarly, all public schools in Ontario were closed on March 12. The state of emergency was announced in Ontario on March 17 and it was ordered to close all non-essential businesses on March 23 (Global News, 2020). We set the lockdown date in Canada on March 13 as other provincial governments as well as the federal government started to recommend the social distancing measures strongly along with the cancellation of various events on the date (CBC News, 2020). These tight lockdown and social distancing measures seemed to contribute the sharp decline of the contact rate in the second period. Both governments started to announce the plans to lift the lockdown measures at the end of April, which corresponds to the third period. Lockdown fatigue would also cause the slower decrease of the contact rate. In sum, a series of social distancing measures have been effective to decrease the contact rate but with some lags. The sparse HP filtering separates these periods reasonably well. However, the ℓ_1 filtering overfits the model with 5 kinks.

Fig. 11 shows the results for China. Since the pandemic is almost over in China, we use the data censored on April 26th when the 3-day-average of newly confirmed cases is less than 10. The estimated kink dates are: January 28, March 14, March 24, and April 18. Based on them, we can classify observations into five periods:

1. January 23–January 28: This is an initial period of the epidemic in China. Since the official confirmation of the novel coronavirus on December 31, 2019, the confirmed cases had increased rapidly. President Xi presided and issued instructions on the epidemic control on January 20. The travel ban on Wuhan was imposed on January 23, 2020 in the period of the Lunar New Year holidays (The New York Times, 2020b). We set this date as the lockdown date.

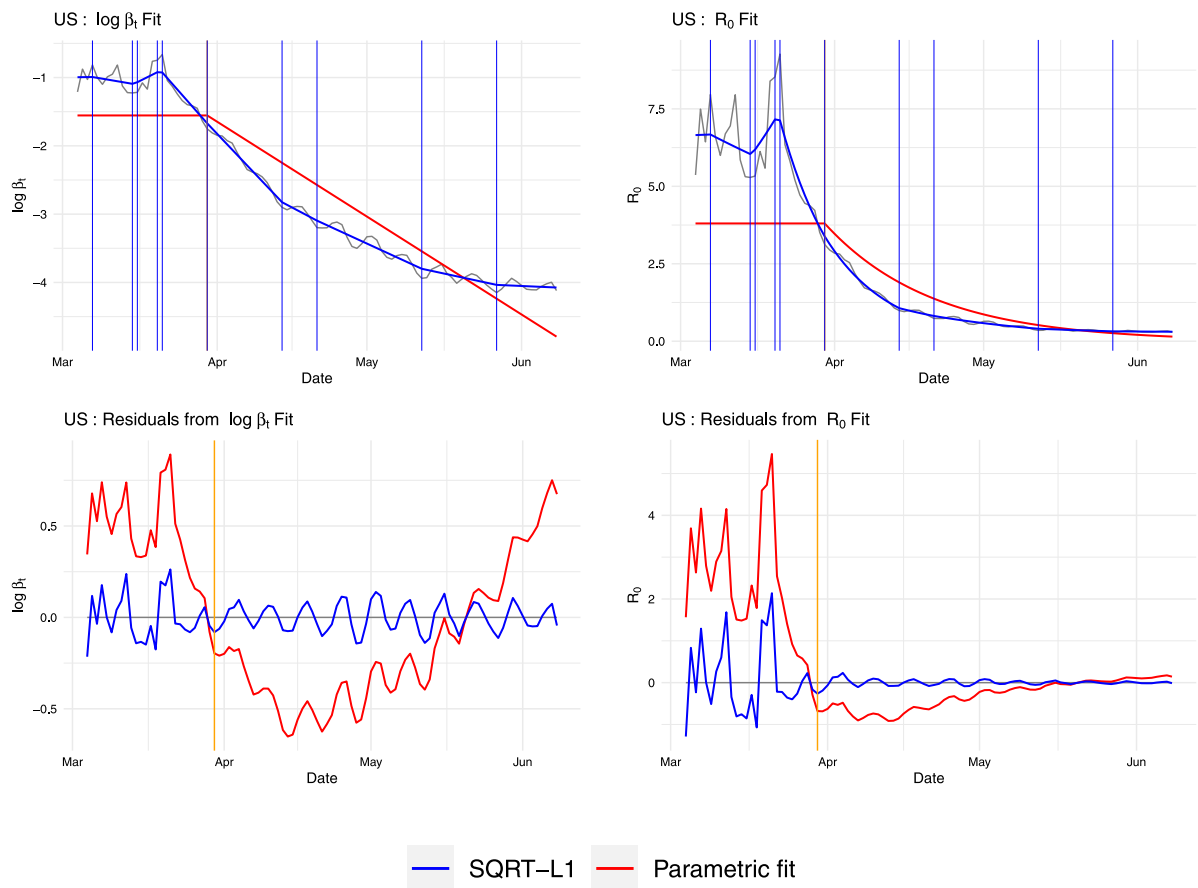


Fig. 8. Square-root ℓ_1 Filtering. Note: Square-root ℓ_1 filtering solves Eq. (3.3). The estimated kinks denoted by blue vertical lines are: March 7, March 15, March 16, March 20, March 21, March 30, April 14, April 21, May 12, and May 27. The orange vertical line denotes the lockdown date, March 30.

2. January 28–March 14: The contact rate shows a sharp decrease during this period. The Lunar New Year holiday was extended to February 2 across the country. China’s National Health Commission (NHC) imposed social distancing measures on January 26. By January 29, all 31 provinces in China upgraded the public health emergency response to the most serious level. By early February, nationwide strict social distancing policies were in place.
3. March 14–March 24: This period shows a V-turn of the contact rate in terms of the $\log \beta_t$ scale. It also shows an upward trend in the $R_0(t)$ scale but the level is lower than that in early February. The mass quarantine of Wuhan was partially lifted on March 19 (Bloomberg, 2020). Most provinces downgraded their public health emergency response level, where factories and stores started to reopen in this period.
4. March 24–April 18: The contact rate still increased but at a lower rate. It started to decrease again at the end of this period. We can see a slight increase in $R_0(t)$. The mass quarantine of Wuhan was lifted more and the travel to other provinces was allowed on April 8 (Bloomberg, 2020).
5. April 18–April 26: The contact rate went down quickly and was flattened at a low level. The last hospitalized Covid-19 patient in Wuhan was discharged on April 26 (Xinhua, 2020).

Fig. 12 shows the results for South Korea. For the same reason in China, we use the data censored on April 29. The estimated kink dates are: March 3, March 15, April 2, and April 21. Based on them, we can classify observations into five periods:

1. February 21–March 3: This period is the beginning of the coronavirus spread in South Korea. On February 21, Shincheonji Church of Jesus, a secretive church in South Korea was linked to a surge of infections in the country (The New York Times, 2020b). The sharp decline of $\log \beta_t$ could be due to the fact that the number of active infections is relatively small in this period and thus, $Y_t = \Delta C_t / (I_{t-1} S_{t-1})$ might not be properly measured.
2. March 3–March 15: A sharp decrease in $\log \beta_t$ in this period corresponds to Korean government’s swift reactions to the outbreak through active testing and contact tracing (The New York Times, 2020a; Aum et al., 2020; Kim et al.,

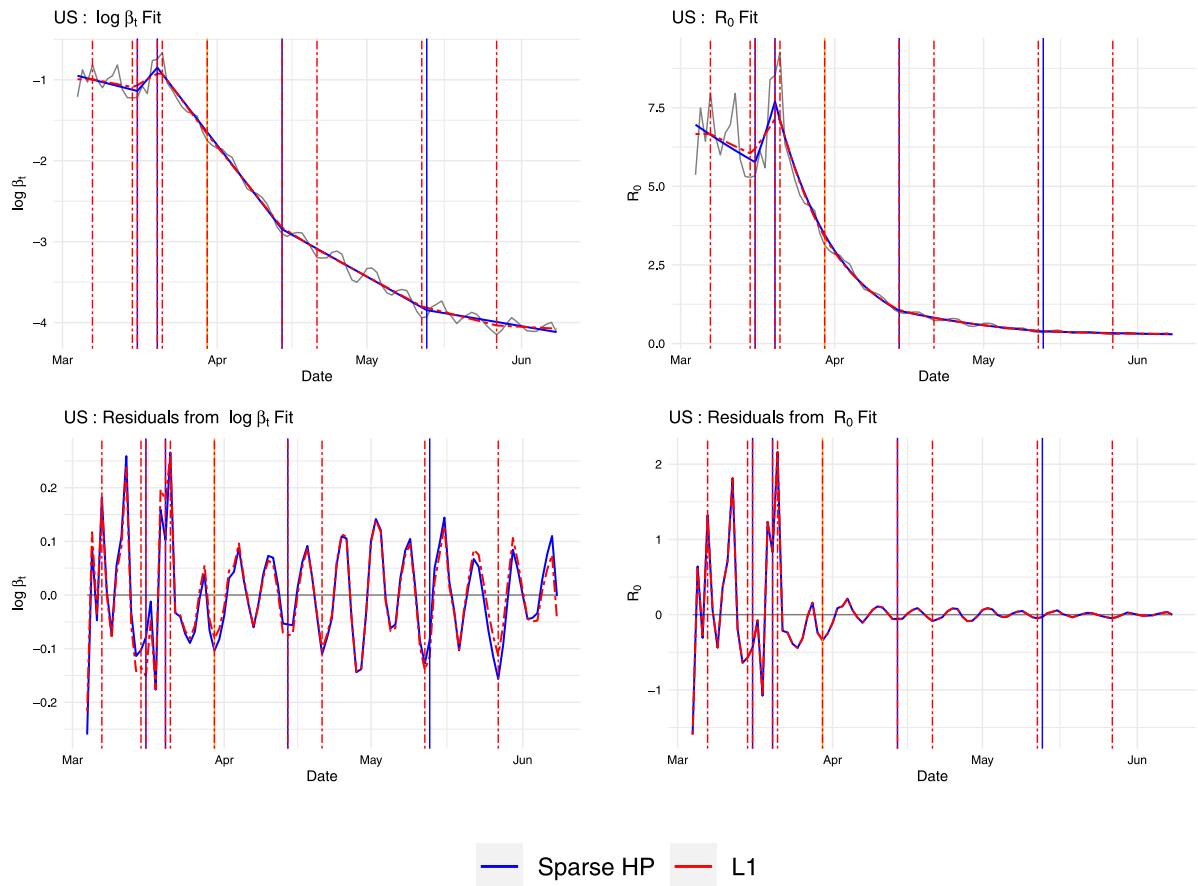


Fig. 9. Sparse HP and ℓ_1 Filtering for the US. Note: The Sparse HP kinks (blue) are: March 16, March 20, April 14, and May 13. The ℓ_1 kinks (red) are: March 7, March 15, March 16, March 20, March 21, March 30, April 14, April 21, May 12, and May 27. The orange vertical line denotes the lockdown date, March 30. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2020), highlighted by prompt containment of an outbreak started on March 8 at a call center in Seoul (Park et al., 2020).

3. March 15–April 2: This period shows a modest V-turn of the contact rate in terms of the $\log \beta_t$ scale but it is much less visible in the $R_0(t)$ scale.
4. April 2–April 21: This period displays a further reduction of the contact rate. A remarkable event was parliamentary elections on April 15 when 30 million people voted without triggering a new outbreak.
5. April 21–April 29: The contact rate was flattened at a low level.

Fig. 13 shows the empirical results of the UK. The estimated kink dates are: March 12 and March 14. Based on them, we can classify observations into three periods:

1. March 6–March 12: This is an initial period of the epidemic in the UK. The downward trend might be due to the fact that the cumulative number of confirmed cases is relatively small and therefore, its growth rate can be easily over-estimated.
2. March 12–March 14: This is still an early stage of the epidemic. The steep increase in the contact rate is again possibly due to the small number of the confirmed cases.
3. March 14–June 8: This period shows a steady and constant decrease in the contact rate. The lockdown measures began in the UK on March 23 (BBC News, 2020b). On May 10, the British prime minister Boris Johnson relaxed certain restrictions and announced the plan for reopening (BBC News, 2020a) but it keeps the downward trend.

Overall, the trend of the contact rate is quite similar to those of the US and Canada. The location of the kinks is around more in the initial periods but it shows the steady downward trend after the prime minister's lockdown announcement. This results in a smooth curve in the $R_0(t)$ scale. The trend estimates of the ℓ_1 filter are almost identical to those of the sparse HP filter; however, it indicates 10 kinks, which seem overly excessive.




Fig. 10. Sparse HP and ℓ_1 Filtering for Canada. Note: The Sparse HP kinks (blue) are: March 18 and April 11. The ℓ_1 kinks (red) are: March 17, March 18, March 24, April 11, and May 24. The orange vertical line denotes the lockdown date, March 13. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5.3. A measure of surveillance and policy implications

The sparse HP filter produces the kinks where the slope changes in the $\log \beta_t$ scale, thereby providing a good surveillance measure for monitoring the ongoing epidemic situation. The policy responses are based on various scenarios and the contact rate is one of the most important measures that determine different developments. As a summary statistic of the time-varying contact rate, we propose to consider the time-varying growth rate of the contact rate, which we call *contact growth rates*:

$$\xi(t) := \frac{\beta_t - \beta_{t-1}}{\beta_{t-1}} \times 100.$$

Recall that we have defined the time-varying basic reproduction number by $R_0(t) = \beta_t/\gamma$. Because γ is fixed over time, we have that

$$\xi(t) = \frac{R_0(t) - R_0(t-1)}{R_0(t-1)} \times 100.$$

Therefore, $\xi(t)$ can be interpreted as the *time-varying growth rate of the basic reproduction number*; it does not require the knowledge of γ and solely depends on β_t . Furthermore, by simple algebra,

$$\xi(t) = [\exp(\log \beta_t - \log \beta_{t-1}) - 1] \times 100, \tag{5.1}$$

which implies that $\xi(t)$ will be piecewise constant if $\log \beta_t$ is piecewise linear. This simple algebraic relationship shows that a change in the slope at the kink in the $\log \beta_t$ scale is translated to a break in the time-varying contact growth rates and therefore in growth rates of the time-varying basic reproduction number. When $\xi(t)$ is a large positive number, that will be a warning signal for the policymakers. On the contrary, if $\xi(t)$ is a big negative number, that may suggest that policy measures imposed before are effective to reduce the contagion.

Table 1 reports the time-varying contact growth rates in the five countries that we investigate, using the sparse HP trend estimates. For the US, the explosive growth rate of 7.5% in the second period is followed by the negative growth rates of -7.7%, -3.4%, and -1%, albeit at diminishing magnitudes. The trajectory of Canada is similar to that of the US. The growth rates of China fluctuated up and down: it started with a high positive 15% followed by -12%; a sharp V-turn

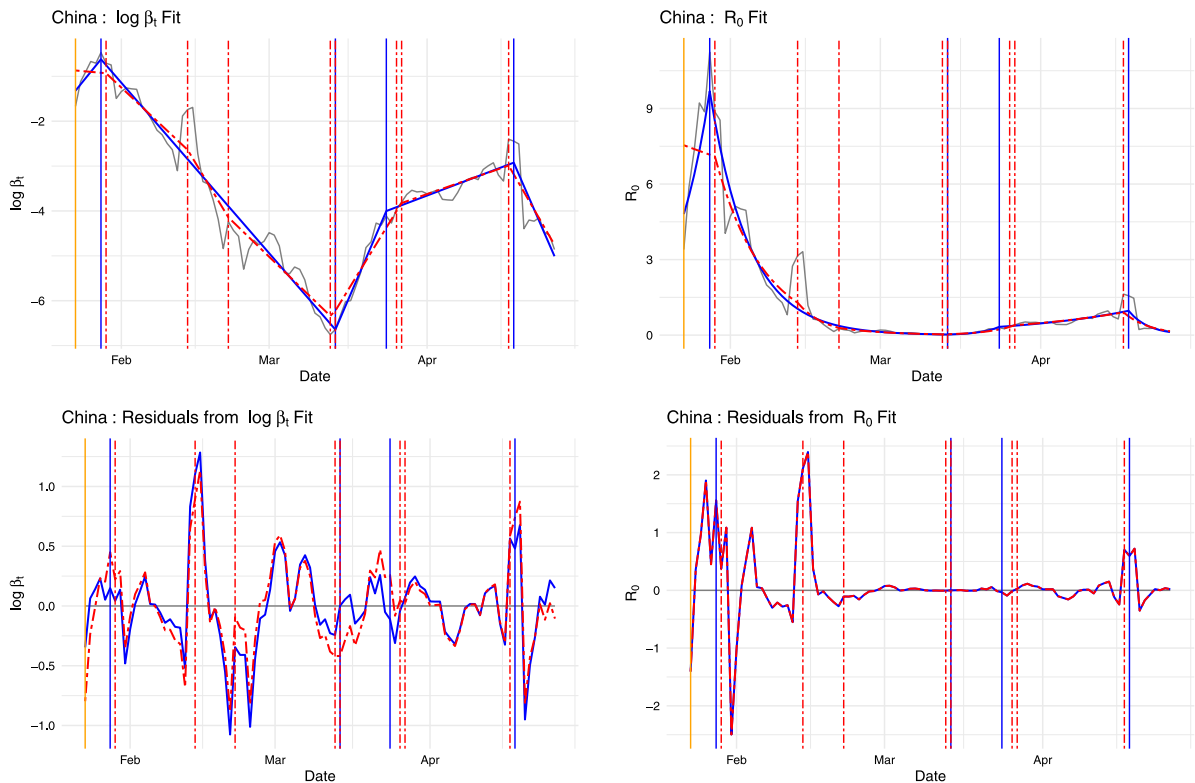


Fig. 11. Sparse HP and ℓ_1 Filtering for China. Note: The Sparse HP kinks (blue) are: January 28, March 14, March 24, and April 18. The ℓ_1 kinks (red) are: January 29, February 14, February 22, March 13, March 14, March 26, March 27, and April 17. The orange vertical line denotes the lockdown date, January 23. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Time-varying contact growth rates.

	US	Canada	China	South Korea	UK
Period 1	-1.55	7.08	15.04	-15.23	-10.96
Period 2	7.48	-5.02	-12.27	-20.34	31.10
Period 3	-7.67	-2.82	30.23	4.47	-4.70
Period 4	-3.39	NA	4.41	-7.88	NA
Period 5	-1.04	NA	-22.95	1.57	NA

Note: The growth rates, expressed as percentages, are obtained by (5.1) using the sparse HP trend estimates. The contact growth rates are also growth rates of $R_0(t)$. The kink dates separating distinct periods are different for each country and they are reported in Sections 5.1 and 5.2.

at the end of the second period (March 14) with the resulting explosive growth rate of 30%, followed by moderate 4% and impressive -23%. It might be the case that the up-and-down pattern observed in China is in part due to data quality issues since China was the first country to experience the pandemic. For South Korea, we can see the stunning drop of the growth rates culminating on March 15 (the end of the second period). A modest positive growth rate during period 3 is offset by a larger magnitude of negative growth rate in period 4. The UK has experienced steady - but not spectacular - negative growths over the sample period following a sharp fluctuation in mid-March. This hints the degrees of effectiveness of the UK lockdown policy. As early pandemic epicenters, China and South Korea experienced V-turns in the time-varying growth rates of basic reproduction number. Canada, the UK and the US may face similar trajectories as they reopen their countries. Our surveillance statistic can be a useful indicator to monitor a new outbreak of COVID-19. However, it will be mainly useful for a short-term projection of the contact growth rate because it is not designed to make long-term trend predictions.

6. Theory

In this section, we examine theoretical properties of the sparse HP and ℓ_1 filters in terms of risk consistency. Let $\|\cdot\|_0$ denote the usual ℓ_0 -(pseudo)norm, that is the number of nonzero elements, and let $\|\cdot\|_r$ and $\|\cdot\|_\infty$, respectively, denote the ℓ_r norm for $r = 1, 2$ and the sup-norm.




Fig. 12. Sparse HP and ℓ_1 Filtering for South Korea. Note: The Sparse HP kinks (blue) are: March 3, March 15, April 2, and April 21. The ℓ_1 kinks (red) are: March 3, March 12, March 15, March 16, April 2, April 3, and April 21. South Korea have not imposed any nation-wide lockdown measure. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6.1. Risk consistency of the sparse HP filter

Define

$$\mathcal{F} = \mathcal{F}(\kappa, M) := \{ \mathbf{f} : \|\mathbf{D}\mathbf{f}\|_0 \leq \kappa, \|\mathbf{D}\mathbf{f}\|_\infty \leq M \}, \tag{6.1}$$

where M is defined in (3.6). For each $\mathbf{f} \in \mathcal{F}$, define

$$S(\mathbf{f}) := \mathbb{E}_{\mathbf{y}} \left[\frac{1}{T} (\mathbf{y} - \mathbf{f})^\top (\mathbf{y} - \mathbf{f}) \right].$$

Let \mathbf{f}^* denote the ideal sparse filter in the sense that

$$\mathbf{f}^* \in \operatorname{argmin}_{\mathbf{f} \in \mathcal{F}} S(\mathbf{f}).$$

Let $\widehat{\mathbf{f}}$ denote the sparse HP filter defined in Section 3.2. Then,

$$R(\widehat{\mathbf{f}}, \mathbf{f}^*) := S(\widehat{\mathbf{f}}) - S(\mathbf{f}^*) \tag{6.2}$$

is always nonnegative. Following the literature on empirical risk minimization, we bound the excess risk R in (6.2) and establish conditions under which it converges to zero.

Recall that the sparse HP filter minimizes

$$Q_n(\mathbf{f}) := \frac{1}{T} (\mathbf{y} - \mathbf{f})^\top (\mathbf{y} - \mathbf{f}) + \frac{\lambda}{T} \mathbf{f}^\top \mathbf{D}^\top \mathbf{D} \mathbf{f}$$

subject to $\mathbf{f} \in \mathcal{F}$.

Let $S_n(\mathbf{f}) := T^{-1} (\mathbf{y} - \mathbf{f})^\top (\mathbf{y} - \mathbf{f})$. Write

$$\begin{aligned} R(\widehat{\mathbf{f}}, \mathbf{f}^*) &= S(\widehat{\mathbf{f}}) - Q_n(\mathbf{f}^*) + Q_n(\mathbf{f}^*) - S(\mathbf{f}^*) \\ &\leq S(\widehat{\mathbf{f}}) - Q_n(\widehat{\mathbf{f}}) + Q_n(\mathbf{f}^*) - S(\mathbf{f}^*) \end{aligned}$$

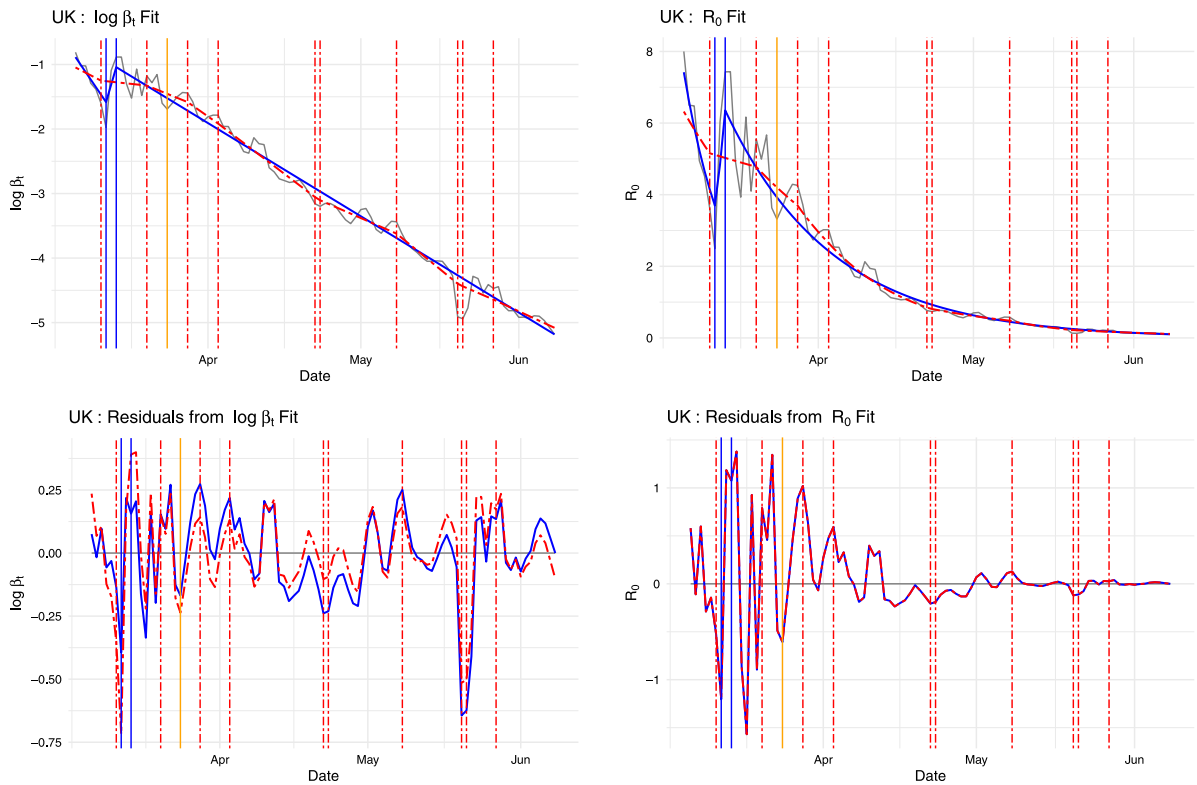


Fig. 13. Sparse HP and ℓ_1 Filtering for the UK. Note: The Sparse HP kinks (blue) are: March 12 and March 14. The ℓ_1 kinks (red) are: March 11, March 20, March 28, April 3, April 22, April 23, May 8, May 20, May 21, and May 27. The orange vertical line denotes the lockdown date, March 24. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$$\begin{aligned}
 &= S(\hat{\mathbf{f}}) - S_n(\hat{\mathbf{f}}) - \frac{\lambda}{T} \hat{\mathbf{f}}^\top \mathbf{D}^\top \mathbf{D} \hat{\mathbf{f}} + S_n(\mathbf{f}^*) + \frac{\lambda}{T} \mathbf{f}^{*\top} \mathbf{D}^\top \mathbf{D} \mathbf{f}^* - S(\mathbf{f}^*) \\
 &\leq 2 \sup_{\mathbf{f} \in \mathcal{F}} |S_n(\mathbf{f}) - S(\mathbf{f})| + 2 \frac{\lambda}{T} \sup_{\mathbf{f} \in \mathcal{F}} \mathbf{f}^\top \mathbf{D}^\top \mathbf{D} \mathbf{f}.
 \end{aligned}$$

Therefore, it suffices to bound two terms above. For the second term, we can use (3.6) and (6.1) to bound

$$2 \frac{\lambda}{T} \sup_{\mathbf{f} \in \mathcal{F}} \mathbf{f}^\top \mathbf{D}^\top \mathbf{D} \mathbf{f} \leq \frac{2\lambda M^2 \kappa}{T}.$$

We summarize discussions above in the following lemma.

Lemma 6.1. Let $\hat{\mathbf{f}}$ denote the sparse HP filter. Then,

$$R(\hat{\mathbf{f}}, \mathbf{f}^*) \leq 2 \sup_{\mathbf{f} \in \mathcal{F}} |S_n(\mathbf{f}) - S(\mathbf{f})| + \frac{2\lambda \kappa}{T} \max_{t=2, \dots, T-1} |y_{t-1} - 2y_t + y_{t+1}|^2.$$

To derive an asymptotic result, we introduce subscripts indexed by the sample size T , when necessary for clarification. Let \mathcal{G}_κ denote the set of every continuous and piecewise linear function whose slopes and the function itself is bounded by C_1 and C_2 , respectively, and the number of kinks is bounded by κ .

Assumption 4. Assume that \mathcal{F} in (6.1) satisfies

$$\mathcal{F}(\kappa, M) \subseteq \mathcal{F}_T := \{ \mathbf{f}_T = (f_{T,1}, \dots, f_{T,T}) : f_{T,t} = f(t/T), f \in \mathcal{G}_\kappa \}. \tag{6.3}$$

Moreover, $y_t = f_{T,t}^* + u_t$, where $\log \beta_t = f_{T,t}^*$, $\mathbf{f}_T^* \in \mathcal{F}_T$ and u_t satisfies $\sup_{t=1,2,\dots} \mathbb{E}|u_t|^p < \infty$ for some $p \geq 2$ and Assumption 2. Finally, $\lambda \kappa T^{-(1-1/p)} \rightarrow 0$ as $T \rightarrow \infty$.

Then, we have the following proposition.

Proposition 6.1. *Let Assumption 4 hold. Then, we have that as $T \rightarrow \infty$,*

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{T} \sum_{t=1}^T \{(y_t - f_t)^2 - \mathbb{E}(y_t - f_t)^2\} \right| \rightarrow_p 0 \tag{6.4}$$

and

$$\frac{\lambda \kappa}{T} \max_{t=2, \dots, T-1} |y_{t-1} - 2y_t + y_{t+1}|^2 \rightarrow_p 0. \tag{6.5}$$

Therefore, $R(\hat{f}, f^*) \rightarrow_p 0$.

Proposition 6.1 establishes the consistency in terms of the excess risk R . Assumption 4 provides sufficient conditions for (6.4) and (6.5). Condition (6.4) is a uniform law of large numbers for the class \mathcal{F} and condition (6.5) imposes a weak condition on λ . Proposition 6.1 follows immediately from Lemma 6.1 once (6.4) and (6.5) are established.

Proof of Proposition 6.1. Note that the summand in (6.4) can be rewritten $(y_t^2 - \mathbb{E}y_t^2) - 2f_t(y_t - \mathbb{E}y_t)$. Then, $y_t^2 - \mathbb{E}y_t^2 = u_t^2 - \sigma_u^2 + 2(f_{T,t}^* - f_t)u_t$ and $2f_t(y_t - \mathbb{E}y_t) = 2f_t u_t$. Furthermore, $T^{-1} \sum_{t=1}^T u_t^2 - \sigma_u^2 = o_p(1)$ due to the law of large numbers (LLN) for a martingale difference sequence (mds).

We now turn to $\sup_{f \in \mathcal{F}} (T^{-1} \sum_{t=1}^T f_t u_t)$. The marginal convergence is straightforward since $f_t u_t$ is an mds with bounded second moments due to the LLN for mds. Next, note that for a constant $\eta > 0$

$$\sup_{|f - f'|_{\infty} < \eta} \left| T^{-1} \sum_{t=1}^T (f_t - f'_t) u_t \right| \leq \eta \left(T^{-1} \sum_{t=1}^T |u_t| \right),$$

which implies the stochastic equicontinuity of the process indexed by $f \in \mathcal{F}_T$. Finally, recall Arzelà–Ascoli theorem, see e.g. Van Der Vaart and Wellner (1996), to conclude that \mathcal{G}_κ is totally bounded with respect to $|\cdot|_\infty$. Therefore,

$$\sup_{f \in \mathcal{F}_T} \left(T^{-1} \sum_{t=1}^T f_t u_t \right) = o_p(1)$$

by a generic uniform convergence theorem, e.g. Andrews (1992).

To show the condition (6.5), note that it is bounded by $16 \frac{\lambda \kappa}{T} \max_{t=1, \dots, T} y_t^2$, which is in turn $O_p(\lambda \kappa T^{-(1-1/p)})$ due to the moment condition on u_t .

6.2. Risk consistency of the ℓ_1 filter

The ℓ_1 trend filtering (3.1) can be expressed as

$$\tilde{f} := \arg \min_{f \in \mathbb{R}^T} \|y - f\|_2^2 + \lambda \|Df\|_1.$$

We now derive the deviation bound for $\|\tilde{f} - f^*\|_2$. First, the problem is equivalent to a regular LASSO problem as stated in Lemma 6.2.

Write $D = (D_3, D_2)$ where D_2 has two columns. Additionally, write

$$G_2 := \begin{pmatrix} D_3^{-1} \\ \mathbf{0} \end{pmatrix}, \quad g_1 := \begin{pmatrix} -D_3^{-1} D_2 \\ I_2 \end{pmatrix},$$

where $\mathbf{0}$ is $2 \times (T - 2)$, g_1 is $T \times 2$ and G_2 is $T \times (T - 2)$. Let $P_{g_1} = g_1(g_1^T g_1)^{-1} g_1^T$.

Lemma 6.2. *We have $\tilde{f} = y - \tilde{y} + \tilde{X}\hat{\theta}$, where $\tilde{y} := (I - P_{g_1})y$, $\tilde{X} := (I - P_{g_1})G_2$ and*

$$\hat{\theta} := \arg \min_{\theta} \|\tilde{y} - \tilde{X}\theta\|_2^2 + \lambda \|\theta\|_1.$$

Proof of Lemma 6.2. Let $D_1 = (\mathbf{0} : I_2)$ be a $2 \times T$ matrix, so that

$$\bar{D} := \begin{pmatrix} D \\ D_1 \end{pmatrix}$$

is upper triangular and invertible. Then, $G := \bar{D}^{-1} = (G_2, g_1)$. Then for a generic $f \in \mathbb{R}^T$, we can define

$$\alpha := \bar{D}f = \begin{pmatrix} Df \\ D_1 f \end{pmatrix} := \begin{pmatrix} \theta \\ a \end{pmatrix}.$$

So (θ, \mathbf{a}) also depend on \mathbf{f} and $\mathbf{G}\mathbf{a} = \mathbf{g}_1\mathbf{a} + \mathbf{G}_2\theta$. Then the problem is equivalent to: $\tilde{\mathbf{f}} = \mathbf{g}_1\hat{\mathbf{a}} + \mathbf{G}_2\hat{\theta}$, where

$$(\hat{\mathbf{a}}, \hat{\theta}) := \min_{\mathbf{a}, \theta} \|\mathbf{y} - (\mathbf{g}_1\mathbf{a} + \mathbf{G}_2\theta)\|_2^2 + \lambda\|\theta\|_1.$$

To solve the problem, we concentrate out \mathbf{a} : Given θ , the optimal \mathbf{a} is $(\mathbf{g}_1^\top \mathbf{g}_1)^{-1} \mathbf{g}_1^\top (\mathbf{y} - \mathbf{G}_2\theta)$ and the optimal $\mathbf{g}_1\mathbf{a}$ is $\mathbf{P}_{\mathbf{g}_1}(\mathbf{y} - \mathbf{G}_2\theta)$. Substituting, so the problem becomes a regular LASSO problem:

$$\min_{\theta} \|\tilde{\mathbf{y}} - \tilde{\mathbf{X}}\theta\|_2^2 + \lambda\|\theta\|_1.$$

Finally, $\tilde{\mathbf{f}} = \mathbf{P}_{\mathbf{g}_1}(\mathbf{y} - \mathbf{G}_2\hat{\theta}) + \mathbf{G}_2\hat{\theta} = \mathbf{y} - \tilde{\mathbf{y}} + \tilde{\mathbf{X}}\hat{\theta}$.

Next, let J denote the indices of t so that $f_{0,t-1} - f_{0,t} \neq f_{0,t} - f_{0,t+1}$ when $t \in J$; let J^c denote the indices of t so that $f_{0,t-1} - f_{0,t} = f_{0,t} - f_{0,t+1}$ when $t \in J$. Here, $\{f_{0,t} : t = 1, \dots, T\}$ denote the true elements of \mathbf{f} . For a generic vector $\theta \in \mathbb{R}^{T-2}$, let θ_J and θ_{J^c} respectively be its subvectors whose elements are in J and J^c . Now we define the restricted eigenvalue constant

$$\zeta := \inf_{\|\theta_{J^c}\|_1 \leq 9\|\theta_J\|_1} \frac{\|\frac{1}{\sqrt{T}}\tilde{\mathbf{X}}\theta\|_2^2}{\|\theta\|_2^2}.$$

Proposition 6.2. Let \mathbf{f}^* denote the true value of \mathbf{f} and $\mathbf{u} := \mathbf{y} - \mathbf{f}^*$. Suppose the event $2.5\|\mathbf{u}^\top \tilde{\mathbf{X}}\|_\infty < \lambda$ holds. Then on this event

$$R(\tilde{\mathbf{f}}, \mathbf{f}^*) \leq \frac{2}{T} \mathbf{u}^\top \mathbf{P}_{\mathbf{g}_1} \mathbf{u} + 2\|\frac{1}{T} \tilde{\mathbf{X}}^\top \tilde{\mathbf{X}}\|_\infty \left(\frac{18\lambda}{\zeta T} \|\mathbf{u}\|_0 \right)^2. \tag{6.6}$$

Proof of Proposition 6.2. Let $\theta^* = \mathbf{D}\mathbf{f}^*$. Consider the vector form of the model $\mathbf{y} = \mathbf{f}^* + \mathbf{u}$. Then $\tilde{\mathbf{y}} = \tilde{\mathbf{X}}\theta^* + \tilde{\mathbf{u}}$ where $\tilde{\mathbf{u}} = (\mathbf{I} - \mathbf{P}_{\mathbf{g}_1})\mathbf{u}$. By Lemma 6.2, $\tilde{\mathbf{f}} = \mathbf{y} - \tilde{\mathbf{y}} + \tilde{\mathbf{X}}\hat{\theta}$, where

$$\hat{\theta} := \arg \min_{\theta} \|\tilde{\mathbf{y}} - \tilde{\mathbf{X}}\theta\|_2^2 + \lambda\|\theta\|_1.$$

The standard argument for the LASSO deviation bound implies, on the event $2.5\|\mathbf{u}^\top \tilde{\mathbf{X}}\|_\infty < \lambda$,

$$\|\hat{\theta} - \theta^*\|_1 \leq \frac{18\lambda}{\zeta T} \|\mathbf{u}\|_0.$$

Finally, $\tilde{\mathbf{f}} - \mathbf{f}^* = \mathbf{P}_{\mathbf{g}_1} \mathbf{u} + \tilde{\mathbf{X}}(\hat{\theta} - \theta^*)$ implies

$$R(\tilde{\mathbf{f}}, \mathbf{f}^*) = \frac{1}{T} \|\tilde{\mathbf{f}} - \mathbf{f}^*\|_2^2 \leq \frac{2}{T} \mathbf{u}^\top \mathbf{P}_{\mathbf{g}_1} \mathbf{u} + 2\|\frac{1}{T} \tilde{\mathbf{X}}^\top \tilde{\mathbf{X}}\|_\infty \|\hat{\theta} - \theta^*\|_1^2.$$

To achieve risk consistency, λ has to be chosen to make the second term on the right-hand side of (6.6) asymptotically small and to ensure that the event $2.5\|\mathbf{u}^\top \tilde{\mathbf{X}}\|_\infty < \lambda$ holds with high probability. The first term on the right-hand side of (6.6) will converge to zero under mild conditions on \mathbf{u} . It is reassuring that the ℓ_1 trend filter fits COVID-19 data well in our empirical results.

6.3. Risk consistency of $\exp(\hat{f}_t)$ and $\exp(\tilde{f}_t)$

In this subsection, we obtain risk consistency of $\exp(\hat{f}_t)$ and $\exp(\tilde{f}_t)$. To do so, we first rewrite the excess risk in (6.2) as

$$R(\mathbf{f}, \mathbf{f}^*) = \frac{1}{T} \sum_{t=1}^T \mathbb{E}(f_t - f_t^*)^2,$$

for $\mathbf{f} = (f_1, \dots, f_T)^\top$ and $\mathbf{f}^* = (f_1^*, \dots, f_T^*)^\top$. We have proved in the previous sections that $R(\hat{\mathbf{f}}, \mathbf{f}^*) \rightarrow_p 0$ and $R(\tilde{\mathbf{f}}, \mathbf{f}^*) \rightarrow_p 0$. Then, under the assumption that there exists a constant $C < \infty$ such that $\max_t |f_t| + \max_t |f_t^*| < C$ for all $\{f_t\}$ on the parameter space, we get, uniformly for all \mathbf{f} on the parameter space,

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}(\exp(f_t) - \exp(f_t^*))^2 = \frac{1}{T} \sum_{t=1}^T \mathbb{E} \exp(2\tilde{f}_t) (f_t - f_t^*)^2 \leq \exp(2C) \frac{1}{T} \sum_{t=1}^T \mathbb{E}(f_t - f_t^*)^2,$$

where in the first equality we used the mean value theorem for some $\tilde{f}_t \in (f_t, f_t^*)$. Therefore,

$$R(\exp(\hat{\mathbf{f}}), \exp(\mathbf{f}^*)) \leq \exp(2C)R(\hat{\mathbf{f}}, \mathbf{f}^*) = o_p(1)$$

and the analogous result holds for $\exp(\tilde{\mathbf{f}})$.

7. Conclusions

We have developed a novel method to estimate the time-varying COVID-19 contact rate using data on actively infected, recovered and deceased cases. Our preferred method called the sparse HP filter has produced the kinks that are well aligned with actual events in each of five countries we have examined. We have also proposed contact growth rates to document and monitor outbreaks. Theoretically, we have outlined the basic properties of the sparse HP and ℓ_1 filters in terms of risk consistency. The next step might be to establish a theoretical result that may distinguish between the two methods by looking at kink selection consistency. It would also be important to develop a test for presence of kinks as well as an inference method on the location and magnitude of kinks and on contact growth rates. In the context of the nonparametric kernel regression of the trend function, [Delgado and Hidalgo \(2000\)](#) explored the distribution theory for the jump estimates but did not offer testing for the presence of a jump. Compared to the kernel smoothing approach, it is easier to determine the number of kinks using our approach, as we have demonstrated. Furthermore, the linear trend specification is more suitable for forecasting immediate future outcomes, at least until the next kink arises. The long-term prediction is more challenging and it is beyond the scope of this paper. Finally, it would be useful to develop a panel regression model for the contact rate at the level of city, state or country. These are interesting research topics for future research.

Appendix A. Under-reporting of positive cases

In Section 2, it is assumed that we observe (C_t, R_t, D_t) . In this appendix, we show that our time series model in Section 2 is robust to some degree of under-reporting of positive cases.

Assume that what we observe is only a fraction of changes in C_t . This assumption reflects the reality that a daily reported number of newly positive cases of COVID-19 is likely to be underreported. Suppose that we observe Δc_t in period t such that $\Delta c_t := \rho \Delta C_t$, where $0 < \rho < 1$ is unknown. Then,

$$c_t = \sum_{t=1}^T \Delta c_t = \rho \sum_{t=1}^T \Delta C_t = \rho C_t,$$

assuming that $c_0 = C_0 = 0$. In words, ρ is the constant ratio between reported and true cases. Formally, we make the following assumption.

Assumption 5 (Fraction Reporting). For each t , we observe (c_t, r_t, d_t) such that

$$c_t := \rho C_t, \quad r_t := \rho R_t \quad \text{and} \quad d_t := \rho D_t,$$

where $0 < \rho < 1$.

The two simplifying conditions in [Assumption 5](#) is that (i) ρ is identical among the three time series and (ii) ρ is constant over time. In reality, a fraction of reported deaths might be higher than that of reported cases; ρ might be time-varying especially in the beginning of the pandemic due to capacity constraints in testing. However, we believe that ρ is unlikely to vary over time as much as β_t changes over time; thus, we take a simple approach to minimize complexity. The common ρ can be thought of a broad measure of detecting COVID-19 in a community.

Define $i_t := c_t - r_t - d_t$ and $s_t := 1 - c_t$. Under [Assumption 5](#), the reported fraction infected at time t (i_t) is underestimated, but the reported fraction of the proportion that is susceptible at time t (s_t) is overestimated. Note that

$$g_t := \frac{\Delta c_t}{i_{t-1}} = \frac{\rho \Delta C_t}{\rho I_{t-1}} = \frac{\Delta C_t}{I_{t-1}}.$$

However,

$$s_{t-1} = 1 - \rho_{t-1} C_{t-1} \neq S_{t-1}.$$

In words, we have a measurement error problem on s_{t-1} but not on g_t . It follows from [\(2.2\)](#) that the observed g_t and s_{t-1} are related by

$$g_t = \beta_t s_{t-1} + v_t, \tag{A.1}$$

where

$$v_t = \beta_t (S_{t-1} - s_{t-1}) = \beta_t (\rho - 1) C_{t-1}. \tag{A.2}$$

The right-hand side of [\(A.2\)](#) is likely to exhibit an increasing trend since C_{t-1} is the cumulative fraction ever infected. To alleviate this problem, we now divide both sides of [\(A.1\)](#) by c_{t-1} , which is positive, to obtain

$$\frac{g_t}{c_{t-1}} = \beta_t \left[\frac{s_{t-1}}{c_{t-1}} + \frac{\rho - 1}{\rho} \right]. \tag{A.3}$$

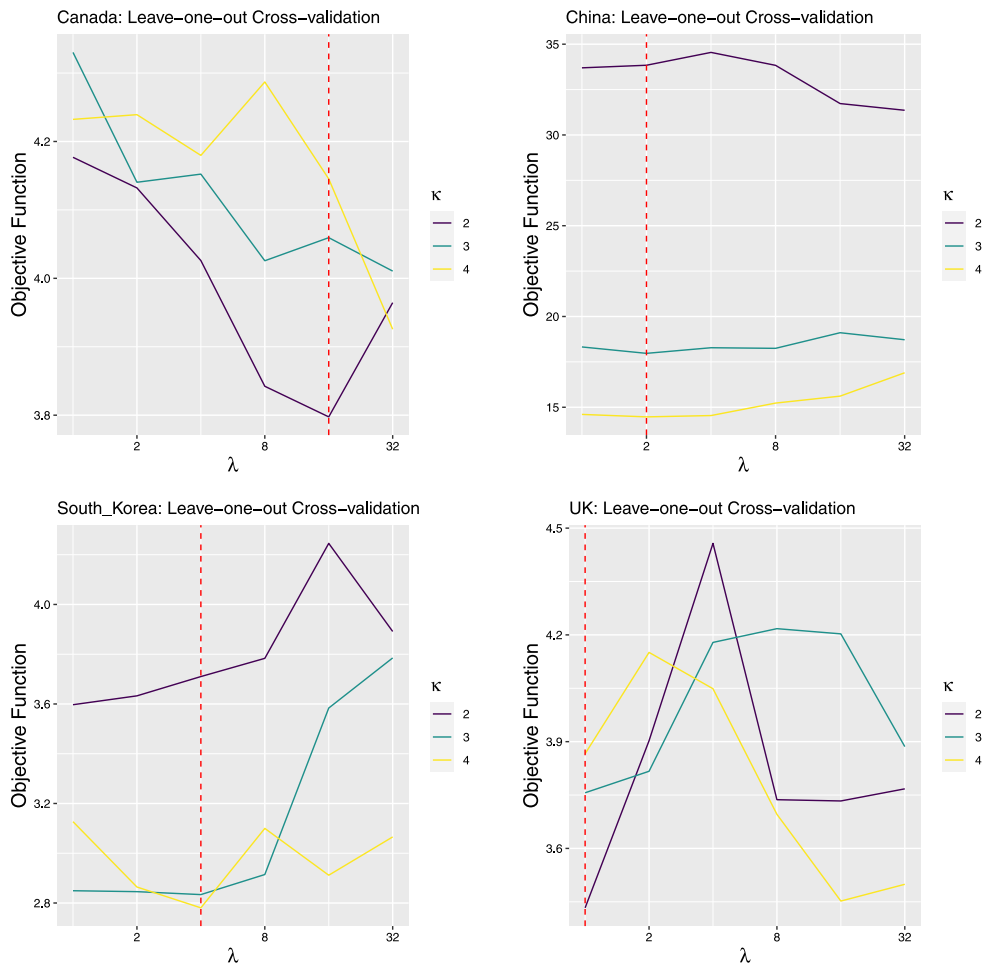


Fig. 14. Sparse HP Filtering: LOOCV for Other Countries. Note: The red dashed line denotes the minimizer of the cross-validation objective function: $(\hat{\kappa}, \hat{\lambda}) = (2, 16)$ for Canada; $(\hat{\kappa}, \hat{\lambda}) = (4, 2)$ for China; $(\hat{\kappa}, \hat{\lambda}) = (4, 4)$ for South Korea; and $(\hat{\kappa}, \hat{\lambda}) = (2, 1)$ for the UK. The analysis period is ended if the number of newly confirmed cases averaged over 3 days is smaller than 10: April 26 (China) and April 29 (South Korea). The grid points are: $\kappa \in \{2, 3, 4\}$ and $\lambda = \{2^0, 2^1, \dots, 2^5\}$. The x-axis is represented by the log₂ scale.

On one hand, if $\rho = 1$, (A.3) is identical to (2.2). On other hand, if $\rho \rightarrow 0$, the term inside the brackets on the right-hand side of (A.3) diverges to infinity.

In the intermediate case, it depends on the relative size between s_{t-1}/c_{t-1} and $(\rho - 1)/\rho$. We now use the UK data to argue that the latter is negligible to the former. According to the estimate by Office for National Statistics (2020), “an average of 0.25% of the community population had COVID-19 in England at any given time between 4 May and 17 May 2020 (95% confidence interval: 0.16% to 0.38%)”. In the UK data used for estimation, the changes in the number of cumulative positives between 4 May and 17 May 2020 is 0.08% of the UK population. Then, an estimate of $\rho = 0.08/0.25 = 0.32$, resulting in $(\rho - 1)/\rho = -2.12$. However, the sample maximum, median, minimum values of s_{t-1}/c_{t-1} are 572412, 804, and 264, respectively. Therefore, the correction term $(\rho - 1)/\rho$ is negligible and therefore, (A.3) reduces to

$$g_t \approx \beta_t s_{t-1}, \tag{A.4}$$

which is virtually the same as (2.2).

Appendix B. Sparse HP filtering: Leave-one-out cross-validation for Canada, China, South Korea and the UK

See Fig. 14.

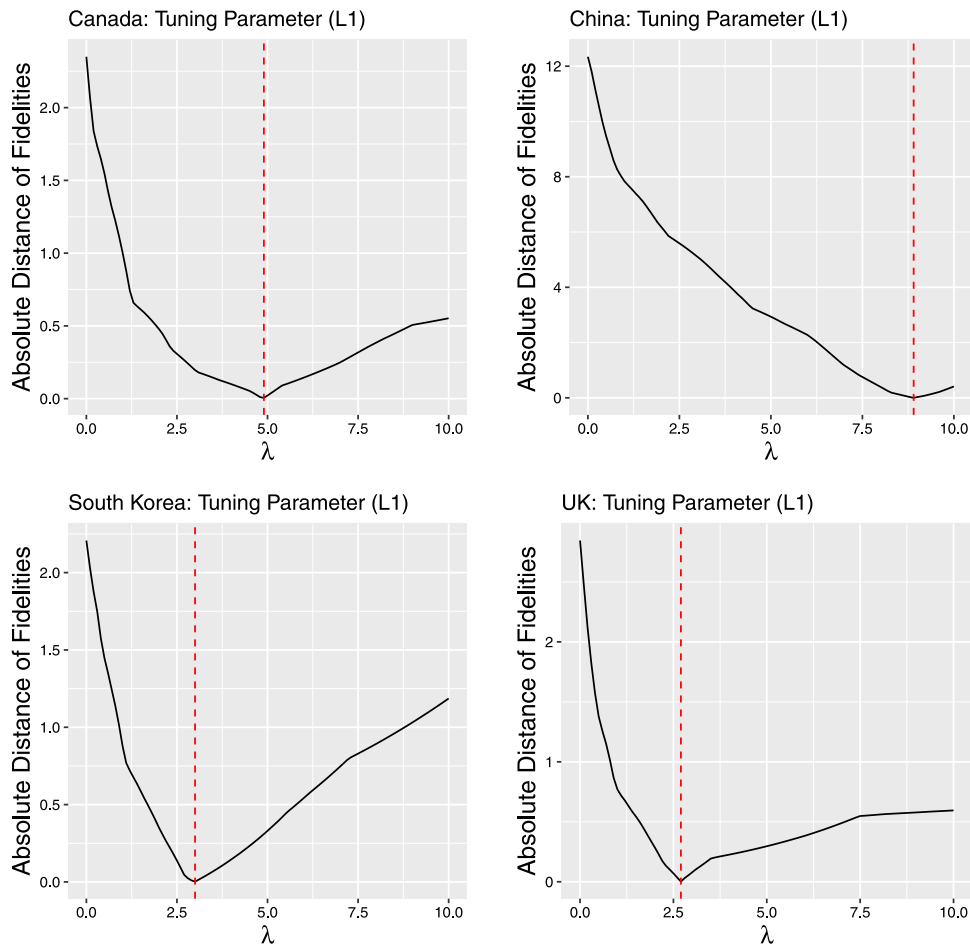


Fig. 15. ℓ_1 Trend Filtering: Selection of λ . Note: The red dashed line denotes the equalizer between the fidelity of the sparse HP filter and that of the ℓ_1 filter: $\hat{\lambda} = 4.9$ for Canada; $\hat{\lambda} = 8.9$ for China; $\hat{\lambda} = 3.0$ for South Korea; and $\hat{\lambda} = 2.7$ for the UK. The analysis period is ended if the number of newly confirmed cases averaged over 3 days is smaller than 10: April 26 (China) and April 29 (South Korea).

Appendix C. ℓ_1 Trend filtering: Selection of λ

See Fig. 15.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2020. Optimal Targeted Lockdowns in a Multi-Group SIR Model. Working Paper 27102, National Bureau of Economic Research.
- Alvarez, F.E., Argente, D., Lippi, F., 2020. A Simple Planning Problem for COVID-19 Lockdown. Working Paper 26981, National Bureau of Economic Research.
- Andrews, D.W., 1992. Generic uniform convergence. *Econometric Theory* 8 (2), 241–257.
- Atkeson, A., 2020. What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios. Working Paper 26867, National Bureau of Economic Research.
- Aum, S., Lee, S.Y.T., Shin, Y., 2020. COVID-19 Doesn't Need Lockdowns to Destroy Jobs: The Effect of Local Outbreaks in Korea. Working Paper 27264, National Bureau of Economic Research.
- Avery, C., Bossert, W., Clark, A., Ellison, G., Ellison, S.F., 2020. Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists. Working Paper 27007, National Bureau of Economic Research.
- BBC News, 2020a. Boris Johnson Speech: PM unveils 'conditional plan' to reopen society. <https://www.bbc.co.uk/news/uk-52609952>, Updated: 2020-05-10. (Accessed 14 June 2020).
- BBC News, 2020b. Coronavirus: Boris Johnson's address to the nation in full. <https://www.bbc.co.uk/news/uk-52011928>, Updated: 2020-03-23. (Accessed 14 June 2020).
- Belloni, A., Chernozhukov, V., Wang, L., 2011. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika* 98 (4), 791–806.
- Bertsimas, D., King, A., Mazumder, R., 2016. Best subset selection via a modern optimization lens. *Ann. Statist.* 44 (2), 813–852.
- Bertsimas, D., Van Parys, B., 2020. Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. *Ann. Statist.* 48 (1), 300–323.

- Bloomberg, 2020. China to lift lockdown over virus epicenter Wuhan on April 8. <https://www.bloomberg.com/news/articles/2020-03-24/china-to-lift-lockdown-over-virus-epicenter-wuhan-on-april-8>, Updated: 2020-03-24. (Accessed 15 June 2020).
- Bühlmann, P., Yu, B., 2003. Boosting with the L_2 loss. *J. Amer. Statist. Assoc.* 98 (462), 324–339.
- CBC News, 2020. Coronavirus: Here's what's happening in Canada and around the world on March 13. <https://www.cbc.ca/news/canada/coronavirus-updates-1.5496334>, Updated: 2020-03-13. (Accessed 15 June 2020).
- Chen, L.-Y., Lee, S., 2018. Best subset binary prediction. *J. Econometrics* 206 (1), 39–56.
- Chen, L.-Y., Lee, S., 2020. Binary classification with covariate selection through ℓ_0 -penalized empirical risk minimization. *Econom. J.* (forthcoming), <https://doi.org/10.1093/ectj/utaa017>.
- Chernozhukov, V., Kasaha, H., Schrimpf, P., 2020. Causal impact of masks, policies, behavior on early Covid-19 pandemic in the U.S. [arXiv:2005.14168](https://arxiv.org/abs/2005.14168).
- Cornea-Madeira, A., 2017. The explicit formula for the Hodrick-Prescott filter in a finite sample. *Rev. Econ. Stat.* 99 (2), 314–318.
- CTV News, 2020. COVID-19 in Quebec: A timeline of key dates and events. <https://montreal.ctvnews.ca/covid-19-in-quebec-a-timeline-of-key-dates-and-events-1.4892912>, Updated: 2020-05-23. (Accessed 14 June 2020).
- de Jong, R.M., Sakarya, N., 2016. The econometrics of the Hodrick-Prescott filter. *Rev. Econ. Stat.* 98 (2), 310–317.
- Delgado, M.A., Hidalgo, J., 2000. Nonparametric inference on structural breaks. *J. Econometrics* 96 (1), 113–144.
- Dong, E., Du, H., Gardner, L., 2020. An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infect. Dis.* 20 (5), 533–534.
- Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2020. The Macroeconomics of Epidemics. Working Paper 26882, National Bureau of Economic Research.
- Fernández-Villaverde, J., Jones, C.I., 2020. Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. <https://web.stanford.edu/~chadj/>, Updated: 2020-05-29 (Version 2.01). (Accessed 10 June 2020).
- Fu, A., Narasimhan, B., Boyd, S., 2017. CVXR: An R package for disciplined convex optimization. [arXiv:1711.07582](https://arxiv.org/abs/1711.07582).
- Global News, 2020. A timeline of the novel coronavirus in Ontario. <https://globalnews.ca/news/6859636/ontario-coronavirus-timeline/>, Updated: 2020-06-08. (Accessed 14 June 2020).
- Hamilton, J.D., 2018. Why you should never use the Hodrick-Prescott filter. *Rev. Econ. Stat.* 100 (5), 831–843.
- Hartl, T., Wälde, K., Weber, E., 2020. Measuring the impact of the German public shutdown on the spread of Covid-19. *Covid Econ. Vetted Real-Time Pap.* (1), 3 April 2020.
- Harvey, A., Kattuman, P., 2020. Time series models based on growth curves with applications to forecasting coronavirus. *Covid Econ. Vetted Real-Time Pap.* (24), 1 June 2020.
- Hastie, T., Tibshirani, R., Tibshirani, R.J., 2017. Extended comparisons of best subset selection, forward stepwise selection, and the Lasso. [arXiv:1707.08692](https://arxiv.org/abs/1707.08692).
- Hethcote, H.W., 2000. The mathematics of infectious diseases. *SIAM Rev.* 42 (4), 599–653.
- Hodrick, R.J., Prescott, E.C., 1997. Postwar U.S. business cycles: An empirical investigation. *J. Money Credit Bank.* 29 (1), 1–16.
- Huang, J., Jiao, Y., Liu, Y., Lu, X., 2018. A constructive approach to L_0 penalized regression. *J. Mach. Learn. Res.* 19 (10), 1–37.
- Kim, S.-J., Koh, K., Boyd, S., Gorinevsky, D., 2009. ℓ_1 trend filtering. *SIAM Rev.* 51 (2), 339–360.
- Kim, Y.-J., Seo, M.H., Yeom, H.-E., 2020. Estimating a breakpoint in the pattern of spread of COVID-19 in South Korea. *Int. J. Infect. Dis.* 97, 360–364.
- Krispin, R., 2020. Coronavirus: The 2019 Novel Coronavirus COVID-19 (2019-nCoV) Dataset, R package version 0.2.0. <https://github.com/RamiKrispin/coronavirus>.
- Li, S., Linton, O., 2020. When Will the Covid-19 Pandemic Peak?. Cambridge-INET Working Paper 2020/11, University of Cambridge, <https://www.inet.econ.cam.ac.uk/research-papers/wp-abstracts?wp=2011>.
- Liu, L., Moon, H.R., Schorfheide, F., 2020. Panel Forecasts of Country-Level Covid-19 Infections. Working Paper 27248, National Bureau of Economic Research.
- Ludvigson, S.C., Ma, S., Ng, S., 2020. Covid19 and the Macroeconomic Effects of Costly Disasters. Working Paper 26987, National Bureau of Economic Research.
- Manski, C.F., 2020. Bounding the Predictive Values of COVID-19 Antibody Tests. Working Paper 27226, National Bureau of Economic Research.
- Manski, C.F., Molinari, F., 2020. Estimating the COVID-19 infection rate: Anatomy of an inference problem. *Econom. J.* <http://dx.doi.org/10.1016/j.jeconom.2020.04.041>, forthcoming.
- Mazumder, R., Radchenko, P., Dedieu, A., 2017. Subset selection with shrinkage: Sparse linear modeling when the SNR is low. [arXiv:1708.03288](https://arxiv.org/abs/1708.03288).
- Office for National Statistics, 2020. Coronavirus (COVID-19) infection survey pilot: England. <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/bulletins/coronaviruscovid19infectionsurvey/pilot/england21may2020>, Release date: 2020-05-21. (Accessed 2 June 2020).
- Park, S., Kim, Y., Yi, S., Lee, S., Na, B., Kim, C., et al., 2020. Coronavirus disease outbreak in call center, South Korea. *Emerg. Infect. Diseases* 26 (8), An early release version is available at: <http://dx.doi.org/10.3201/eid2608.201274>.
- Phillips, P.C.B., Shi, Z., 2019. Boosting: Why you can use the HP filter. [arXiv:1905.00175](https://arxiv.org/abs/1905.00175).
- Pindyck, R.S., 2020. COVID-19 and the Welfare Effects of Reducing Contagion. Working Paper 27121, National Bureau of Economic Research.
- Ravn, M.O., Uhlig, H., 2002. On adjusting the Hodrick-Prescott filter for the frequency of observations. *Rev. Econ. Stat.* 84 (2), 371–376.
- Stock, J.H., 2020. Data Gaps and the Policy Response to the Novel Coronavirus. Working Paper 26902, National Bureau of Economic Research.
- The New York Times, 2020a. How South Korea flattened the curve. <https://www.nytimes.com/2020/03/23/world/asia/coronavirus-south-korea-flatten-curve.html>, Updated: 2020-04-10. (Accessed 15 June 2020).
- The New York Times, 2020b. How the coronavirus pandemic unfolded: a timeline. <https://www.nytimes.com/article/coronavirus-timeline.html>, Updated: 2020-06-09. (Accessed 13 June 2020).
- The New York Times, 2020c. See how all 50 states are reopening. <https://www.nytimes.com/interactive/2020/us/states-reopen-map-coronavirus.html>, Updated: 2020-06-12. (Accessed 13 June 2020).
- The New York Times, 2020d. See which states and cities have told residents to stay at home. <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html>, Updated: 2020-04-20. (Accessed 6 June 2020).
- Tibshirani, R.J., 2014. Adaptive piecewise polynomial estimation via trend filtering. *Ann. Statist.* 42 (1), 285–323.
- Toda, A.A., 2020. Susceptible-infected-recovered (SIR) dynamics of covid-19 and economic impact. *Covid Econ. Vetted Real-Time Pap.* (1), 3 April 2020.
- Van Der Vaart, A.W., Wellner, J.A., 1996. Weak convergence. In: *Weak Convergence and Empirical Processes*. Springer, New York, NY, pp. 16–28.
- Wang, Y.-X., Sharpnack, J., Smola, A.J., Tibshirani, R.J., 2016. Trend filtering on graphs. *J. Mach. Learn. Res.* 17 (1), 3651–3691.
- Xinhua, 2020. Fighting Covid-19, China in action. http://www.xinhuanet.com/english/2020-06/07/c_139120424.htm, Updated: 2020-03-24. (Accessed 15 June 2020).