Adapting to the COVID-19 Pandemic Online Appendix

Michael Droste and James H. Stock

1 Model

This section describes the behavioral SIR model presented in our paper. There is a mass of N individuals (i.e. the US population, 328 million). Each individual belongs to one of five states at time t, denoted S (susceptible), E (exposed), I (infected), R (recovered), and D (deceased) respectively. The unit of time, indexed by t, is a day.

The transition equations for each compartment can be written as:

$$dS_t = -\beta_t \phi_t S_t \tfrac{I_t}{N} \tag{1}$$

$$dE_t = \beta_t \phi_t S_t \tfrac{I_t}{N_t} - \sigma E_t \tag{2}$$

$$dI_t = \sigma E_t - \gamma I_t - \delta_t I_t \tag{3}$$

$$dR_t = \gamma I_t \tag{4}$$

$$dD_t = \delta_t I_t \tag{5}$$

where β_t is an endogenous transmission factor, ϕ_t is an exogenous seasonal transmission factor, σ^{-1} is the mean delay (in days) between exposure and onset of infectiousness; γ^{-1} is the mean delay (in days) between onset and cessation of infectiousness, and δ_t is the (time-varying, daily) fatality rate for infected individuals.

The equation linking the transmission factor β to economic activity s_t is:

$$\beta_t = \exp(\beta_0 + \beta_1 s_t) \tag{6}$$

where the parameter β_0 represents transmissibility unrelated to economic activity s_t and the parameter β_1 represents transmissibility related to s_t.

The equation linking the economic activity s_t to yesterday's deaths DR_{t-1} is:

$$s_t = \kappa_0 + \kappa_1 DR_{t-1} \tag{7}$$

where the parameter κ_0 represents baseline economic activity to yesterday's deaths and the parameter κ_1 represents the relationship between yesterday's deaths DR_{t-1} and economic activity s_t.

The description of the model in this section assumes that the parameters β_0, β_1, κ_0, and κ_1 are fixed. This is an expositional choice that is meant to clarify that these are model parameters. Our estimation strategy allows these parameters to flexibly vary over time. Details on our calibration of the exogenous factors ϕ_t and δ_t and estimation of the model are given in section 3 of this appendix.

Given initial conditions and parameters, this model can be solved numerically in discrete time by repeatedly applying the transition equations. We solve a discrete-time analogue of this model with 14 steps per day.

2 Data

We estimate our model with daily time series on COVID-19 deaths and labor hours, which correspond to the time series D_t and s_t in our model.

The data on observed daily deaths in the United States is derived from the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University, available on a public GitHub repository.

The data on economic activity corresponds to (seasonally adjusted) aggregate weekly hours of production and nonsupervisory employees, which is produced by the US Bureau of Labor Statistics and available at monthly frequency on FRED. We interpolate this data to daily frequency using the daily private employment series made available by Chetty et al. (2020), also available in a public GitHub repository.

We explored several options for interpolation of the weekly hours series to daily frequency. An alternative interpolation using an index of time spent outside home (available through Google's Community Reports) yielded nearly identical results. Univariate methods of interpolating the hours series - for instance, cubic spline interpolation - also yield nearly identical results.

3 Estimation

The purpose of this section is to describe the estimation strategy used to produce time-varying estimates of the key parameters β_0, β_1, κ_0, and κ_1.

Calibration

We calibrate four objects that appear in our model: the seasonal transmission factor ϕ_t, the time-varying infection fatality rate δ_t, the infection onset latency parameter σ, and the recovery latency parameter γ.

The seasonal transmission factor ϕ_t is calibrated using information from Tzampougli and Loukidis (2020). These authors estimate that transmissibility of COVID-19 is subject to a 20% to 30% seasonal swing in transmission (peaking in the winter). We take the midpoint of this estimate and assume that seasonality is symmetric with a periodicity of exactly one year. The

seasonal transmission factor can therefore be expressed as:

$$\phi_t = 0.875 + 0.125\text{cos}(2\pi((t - t_0)/365))$$

where t_0 is a time index corresponding to January 1.

We calibrate the time-varying infection fatality rate using information from two sources. First, we calibrate an initial IFR for the population at 0.8%. Then $\delta_0 = \gamma \times \frac{\text{IFR}}{1-\text{IFR}}$. Following discussion reported in Ledford (2020), we assume that this 'baseline' IFR has fallen by 20% over the course of the pandemic.

The parameters γ and σ are calibrated using evidence from Kissler et al. (2020) as in Atkeson, Droste, Mina, and Stock (2020). γ is set to $1/5$ and σ is set to $1/4.87$, so that we assume the amount of time between exposure and infection and between infection and recovery both average about 5 days.

Rolling Estimation

We estimate these parameters in a flexible fashion by using system estimation with nonlinear least squares using a rolling 8-week estimation window.

The full sample period consists of data from March 15 to December 17, 2020. Our rolling estimation procedure considers eight week windows starting at the beginning of the sample, iterating the start and end dates forward by two weeks for each successive run. This procedure continues until we reach the end of the estimation sample.

For a given window, our estimates correspond to system estimation of equations (1)-(7) using nonlinear least squares, where daily deaths and economic activity are observed. There are six parameters: initial conditions for infected and exposed I_0 and E_0, plus the key parameters of interest β_0, β_1, κ_0, and κ_1. The initial stock of deceased individuals are set at historical values (since deaths are observed in our data), and the initial stock of recovered individuals is set to 0[1]. The remaining initial condition corresponding to the initial stock of susceptibles (S_0) is pinned down by a population adding-up constraint.

We report heteroskedasticity and autocorrelation-robust confidence intervals for these parameters by computing pointwise VAR-HAC standard errors for each window in the rolling estimation. HAC variances computed for each rolling window are very noisy because of the short window and the need for computing a 6-dimensional variance matrix. We therefore used the following approach for computing HAC standard errors: Start by fixing a starting date for the N-week window. At the estimated parameter values for that window, compute derivatives of the predicted values with respect to the parameters, along with residuals at the estimated values. Use these to construct the relevant matrices of derivatives, and derivatives-times-errors, for each day within the window. Save the values of that series at the center day, and then move the window forward by one day. Daily parameter estimates were computed by linear interpolation

[1]This is interpreted as low-frequency movement in the parameter β_0 in our estimates, which we view as inconsequential given the relatively small number of recovered individuals accumulating in each 8-week window.

of the weekly parameter estimates. A first-order vector autoregression is fit to the derivatives-times-errors data, and the VAR-HAC estimator is then computed. The resulting standard error is scaled for the window size then applied to each estimate. The results for a VAR of order 2 are slightly larger than the VAR(1) estimates reported here.

4 Sensitivity Analyses

We explored the sensitivity of our analysis to changes in the rolling estimation procedure; specifically, the window size. Our baseline results are estimated with 8-week rolling estimation windows. Figure 1 Panel A depicts observed weekly deaths in the United States and our model predictions under a baseline 8-week regression window. The black line corresponds to observed weekly deaths and the remaining lines indicate predictions from successive rolling regressions. Figure 1 Panel B utilizes a longer 12-week estimation window. As expected, model fit deteriorates significantly for larger window sizes. Intuitively, time variation in our parameters is necessary to fit the multiple peaks observed in COVID death data in the US, as described in the main text. We therefore focused on relatively short estimation windows.

Figures 2 and 3 plot our time-varying estimates of κ_1 and β_1, respectively, under 8-week and 6-week estimation windows. A 6-week estimation window produces estimates that are qualitatively very similar in both cases, though the confidence bands are somewhat wider. For completeness, we also report our time-varying estimates of κ_0 and β_0 in Figures 4 and 5, respectively, under the same 8-week and 6-week rolling estimation windows. There, too, the estimates are relatively robust to changes in window size.

Figure 1: Predicted vs. Observed Weekly Deaths, Rolling Estimation

Panel A. 8-week Estimation Windows (Baseline)

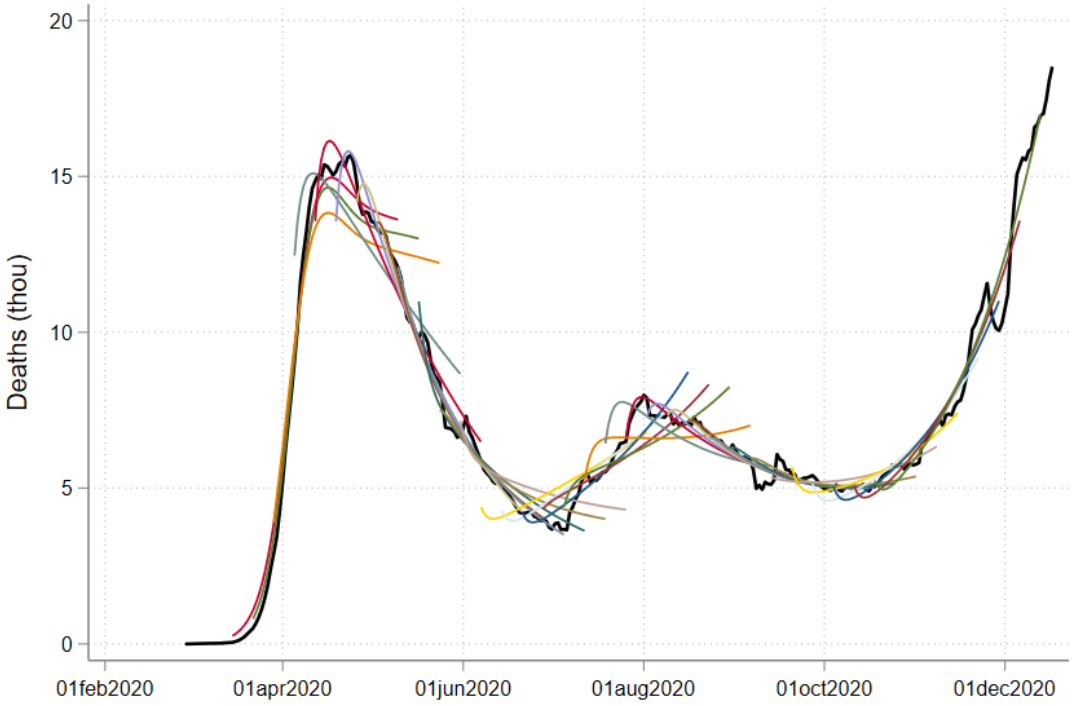

Panel B. 12-week Estimation Windows

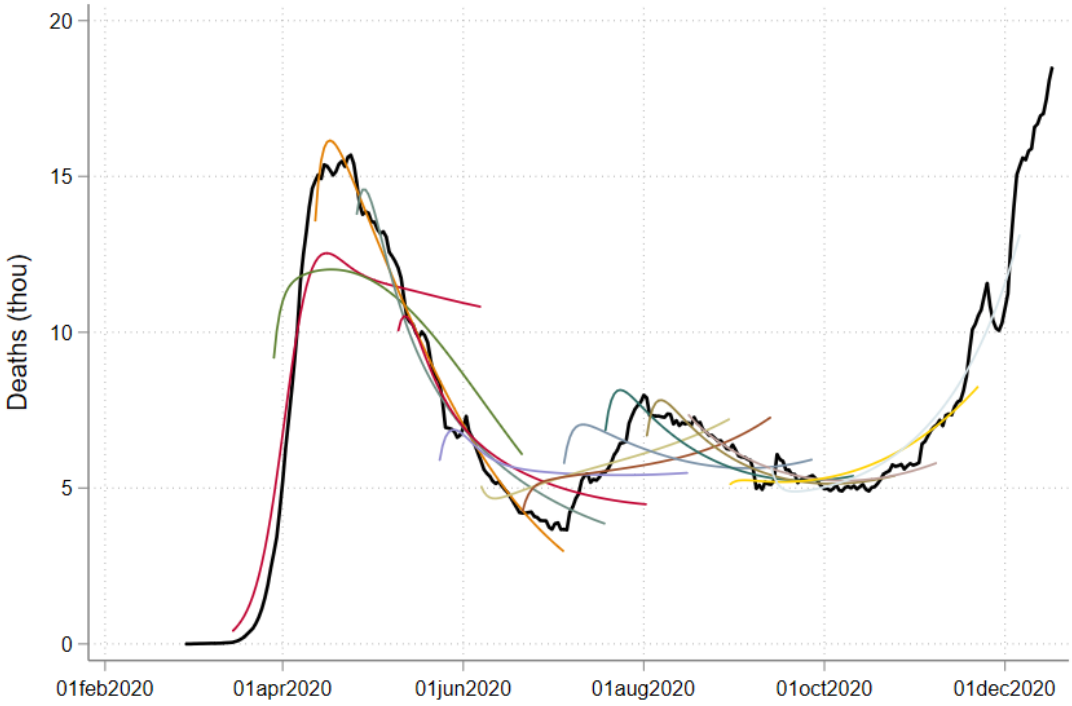

Notes: These figures plot observed weekly deaths (black line) vs. our rolling regression estimates described in section 3 of this appendix. Panel A plots our baseline predictions with 8-week (56-day) estimation windows. Panel B plots alternative predictions with 12-week (84-day) estimation windows.

Figure 2: Time-Varying Estimates of κ_1, Varying Rolling Estimation Window Size

Panel A. 8-week Estimation Windows (Baseline)

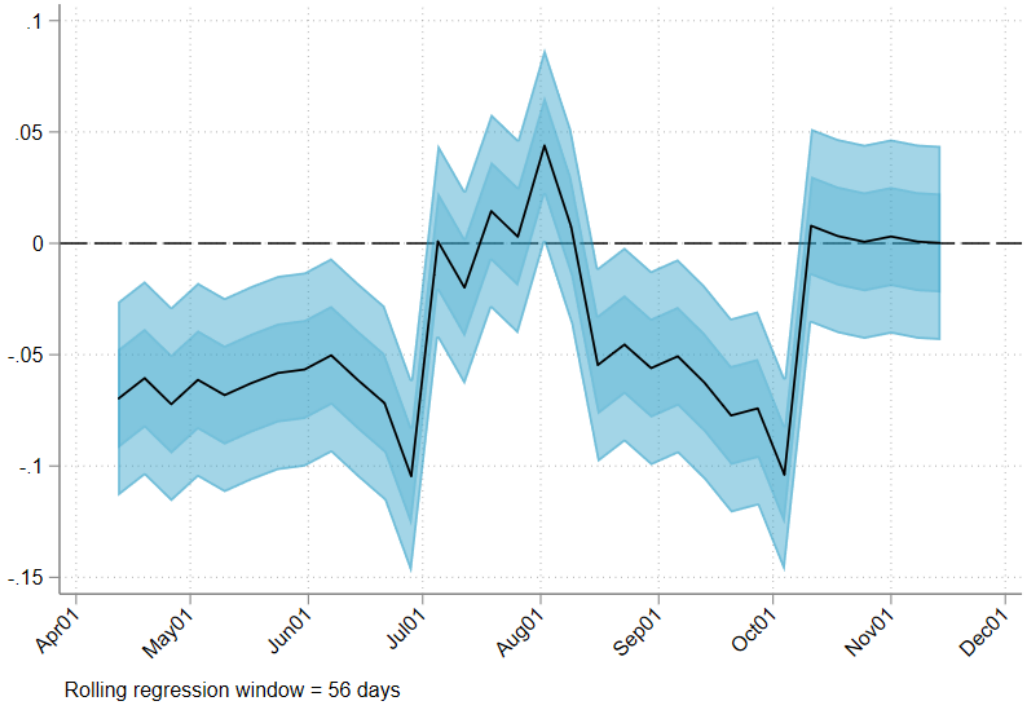

Panel B. 6-week Estimation Windows

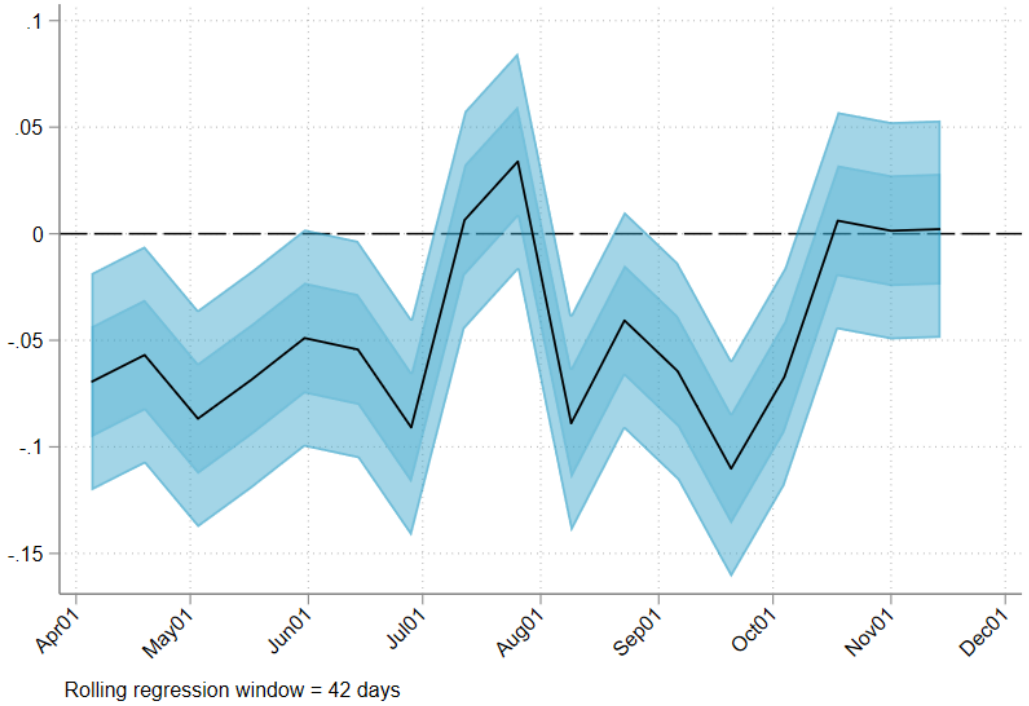

Notes: These figures our time-varying estimates of κ_1 under the baseline 8-week estimation window (Panel A) and an alternative 6-week window (panel B).

Figure 3: Time-Varying Estimates of β_1, Varying Rolling Estimation Window Size

Panel A. 8-week Estimation Windows (Baseline)

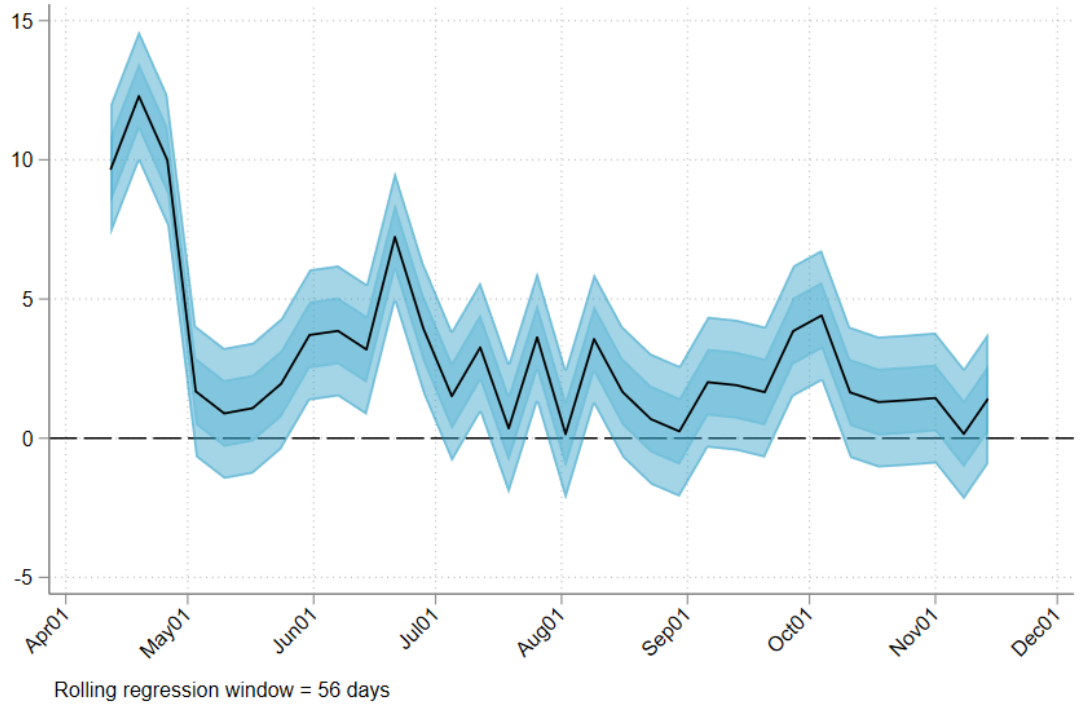

Panel B. 6-week Estimation Windows

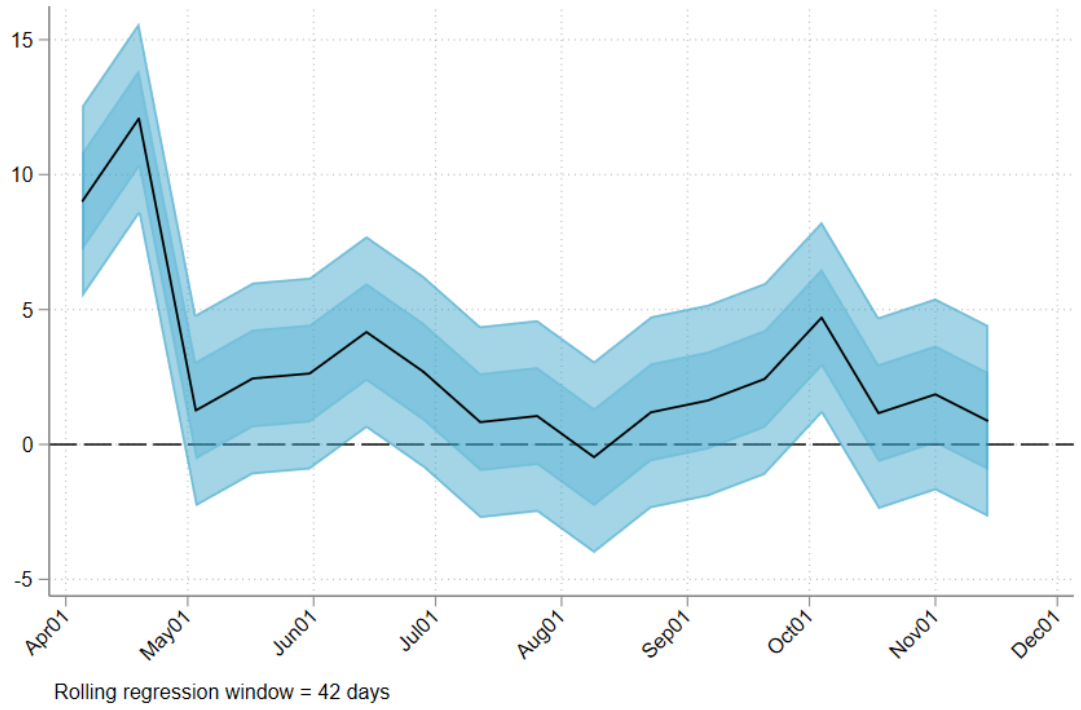

Notes: These figures our time-varying estimates of β_1 under the baseline 8-week estimation window (Panel A) and an alternative 6-week window (panel B).

Figure 4: Time-Varying Estimates of κ_0, Varying Rolling Estimation Window Size

Panel A. 8-week Estimation Windows (Baseline)

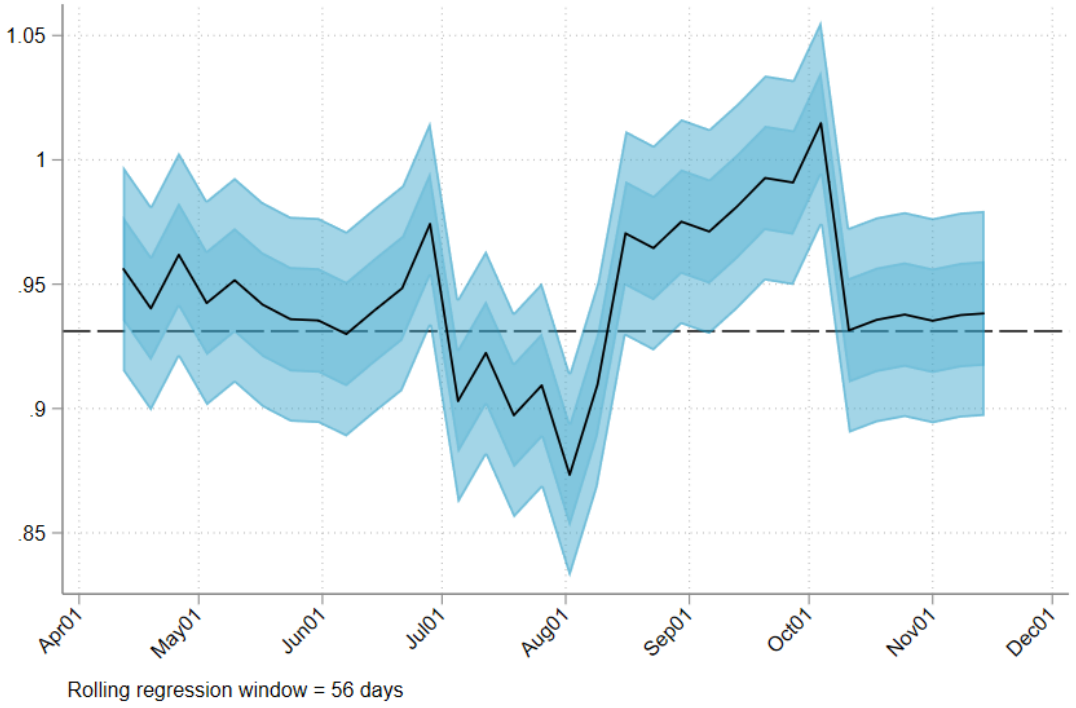

Panel B. 6-week Estimation Windows

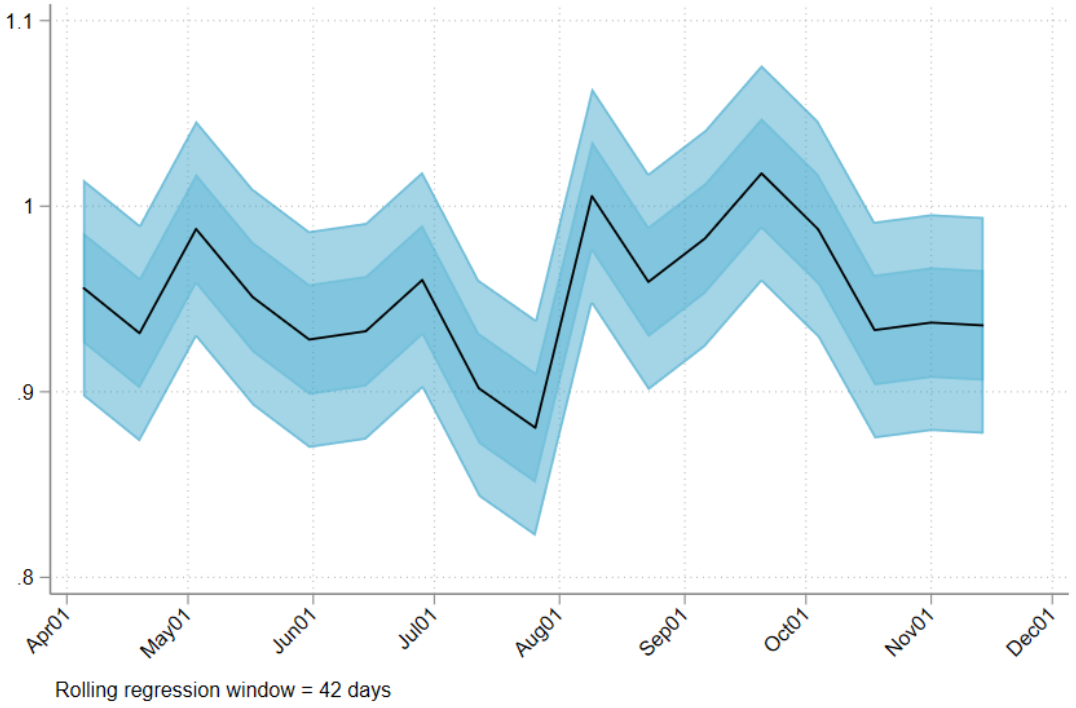

Notes: These figures our time-varying estimates of κ_0 under the baseline 8-week estimation window (Panel A) and an alternative 6-week window (panel B).

Figure 5: Time-Varying Estimates of β_0, Varying Rolling Estimation Window Size

Panel A. 8-week Estimation Windows (Baseline)

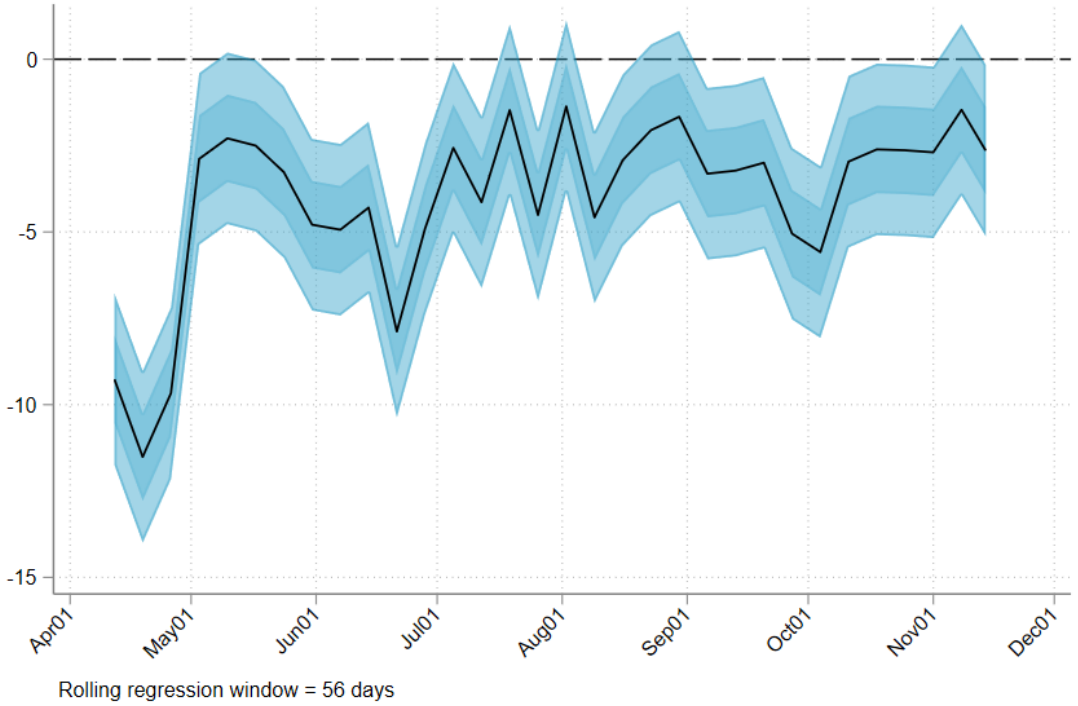

Panel B. 6-week Estimation Windows

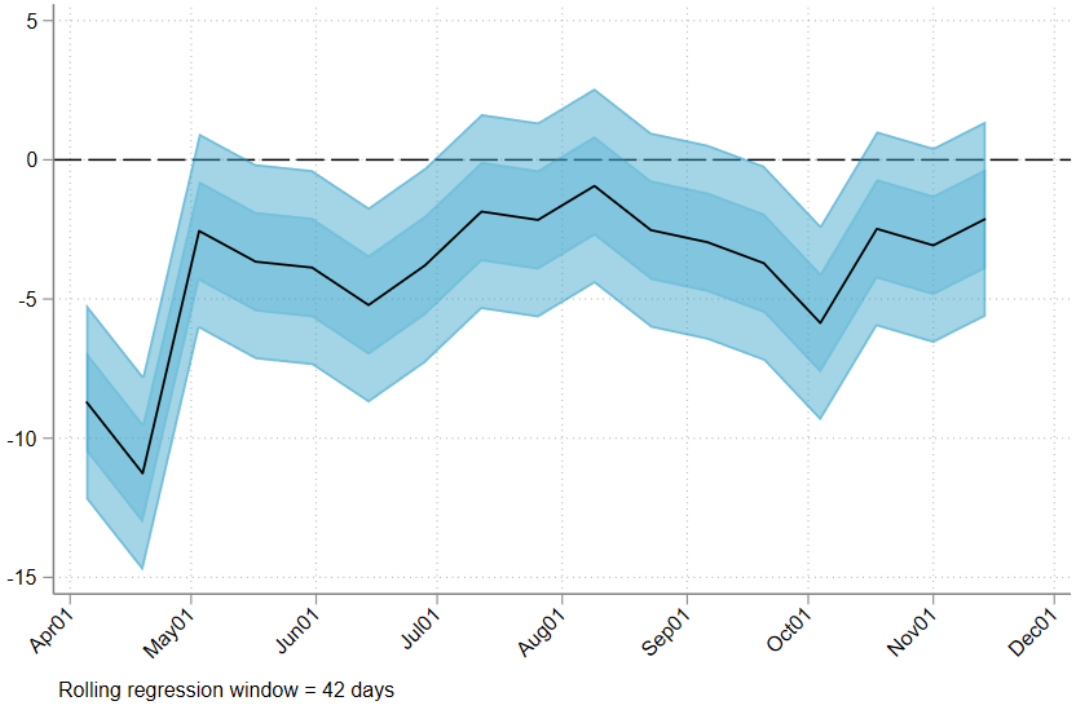

Notes: These figures our time-varying estimates of β_0 under the baseline 8-week estimation window (Panel A) and an alternative 6-week window (panel B).