

A Online Appendix for The Causal Effects of Lockdown Policies on Health and Macroeconomic Outcomes

By Jonas E. Arias, Jesús Fernández-Villaverde, uan F. Rubio-Ramírez, and Minchul Shin

A.1 Implied Density over the Initial Share of Detected Cases

Figure A.1 plots the prior and posterior distributions of the initial share of detected cases, $\gamma_{n,0}$, implied by our model. While the prior assigns substantial probability mass to values as high as 0.6, the posterior distribution is concentrated in values below 0.4, indicating that, according to our model, the share of detected cases on March 15, 2020, was below 40% with high probability.

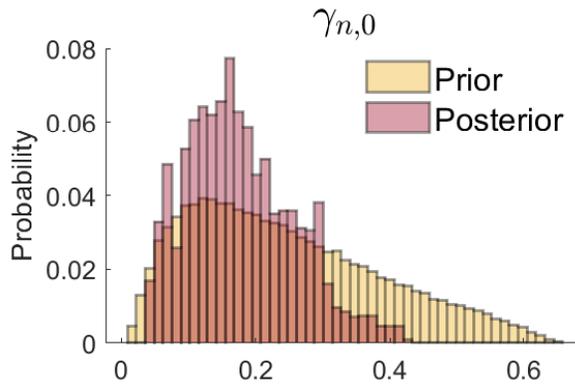


Figure A.1: Prior histograms (yellow) are based on 100,000 independent draws from the prior distribution presented in Table 2. Posterior histograms (red) are based on the MCMC chain with 90,000 posterior draws obtained after a burn-in period of 10,000 draws.

A.2 Prior versus Posterior

Figure A.2 shows the prior and posterior distributions for γ , θ_H , and θ_P . These parameters are inversely related to the average duration a person remains infectious, the average duration of stay in hospitals, and the average duration of stay at home while recovering from COVID-19. Figure A.2 reveals that the data are very informative about γ and θ_H and less informative about θ_P . Figure A.3 shows the prior and posterior distributions for σ_b , σ_h , σ_p , σ_g , and σ_n . These parameters govern the step size of the time-varying parameters of our model. Clearly, the data are informative about them. Finally, Figure A.4 shows the prior and posterior distributions for b_0 , d_{H_0} , d_{P_0} , g_{H_0} , S_0 , I_0 , n_0 , and μ . These parameters are the initial-value parameters and the share of false negatives parameter in the case of μ . Figure A.4 documents that data are informative about these 8 parameters as well. The prior for n_0 is truncated at 1 to rule out large

values of the permanent component of detected cases. Even so, the posterior indicates that there is not much probability mass near the truncation, suggesting that the upper bound for the prior could be relaxed without affecting our conclusions.

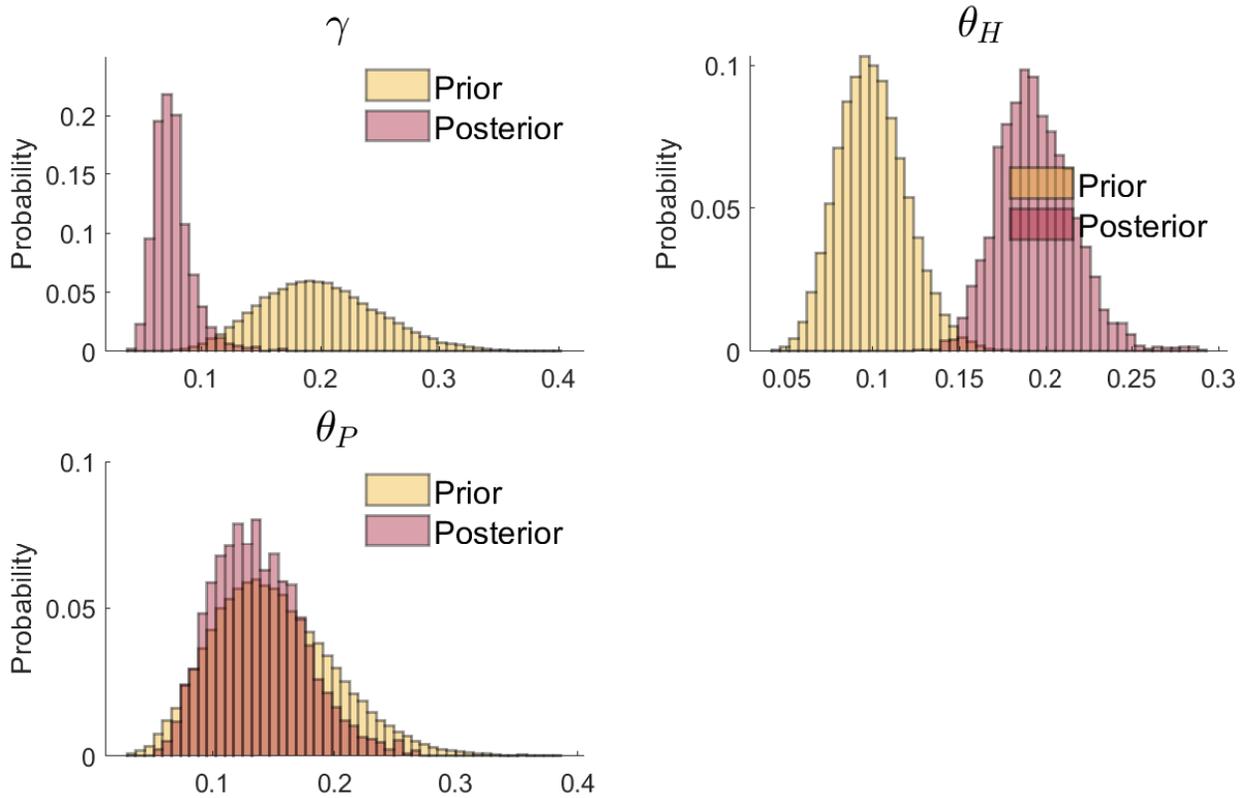


Figure A.2: Duration Parameters. Table 1 presents definitions of these parameters. Prior histograms (yellow) are based on 100,000 independent draws from the prior distribution presented in Table 2. Posterior histograms (red) are based on the MCMC chain with 90,000 posterior draws obtained after a burn-in period of 10,000 draws.

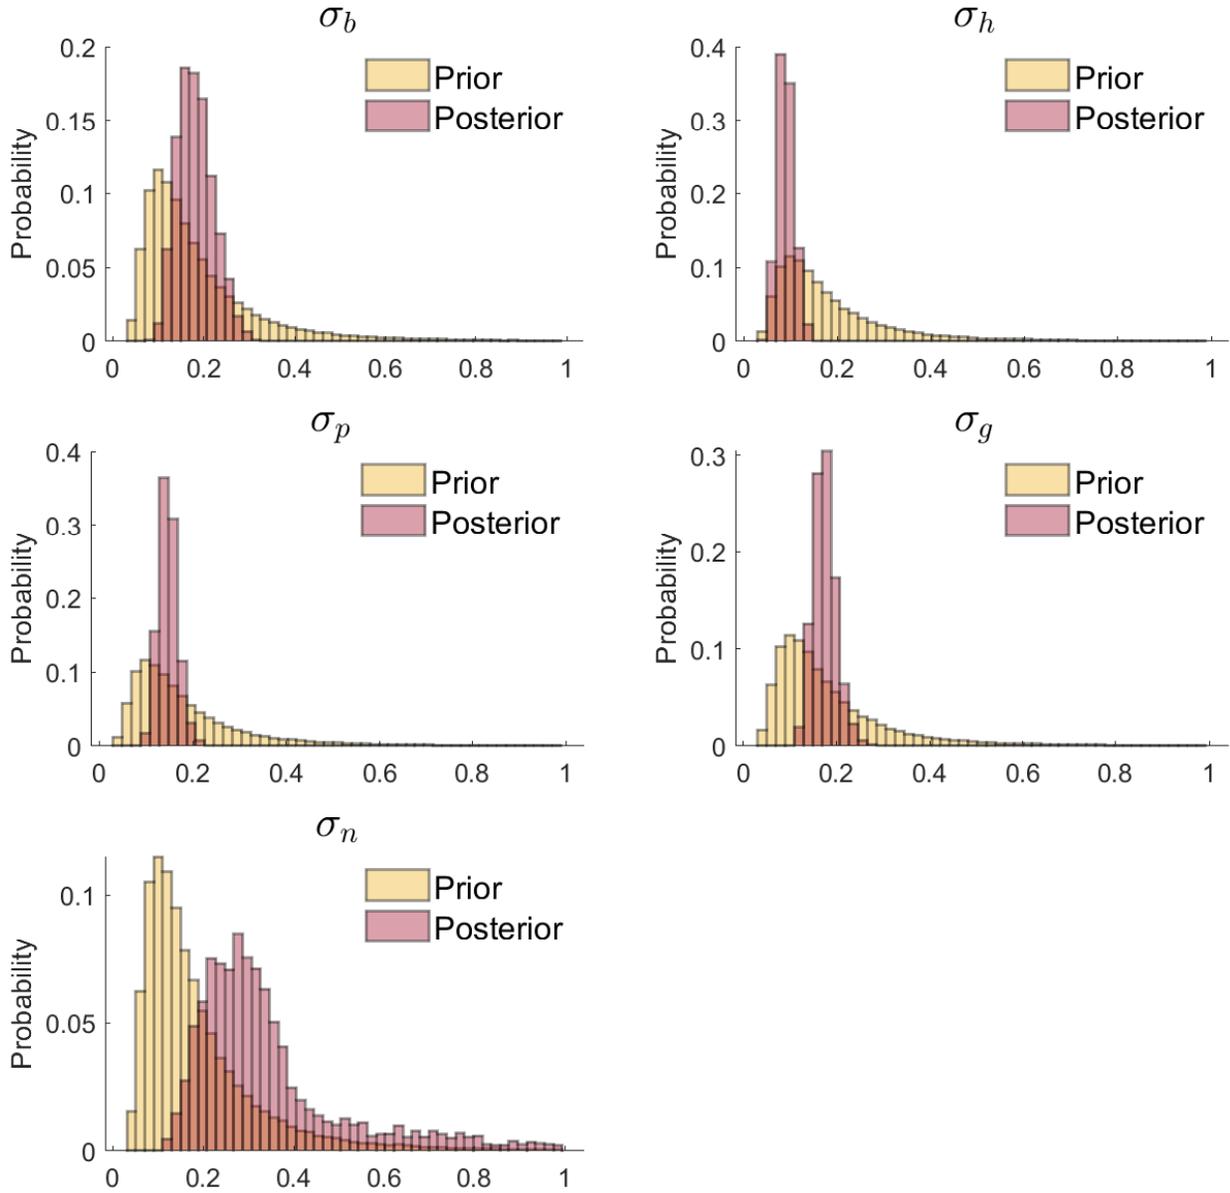


Figure A.3: Step-size Parameters. Table 1 presents definitions of these parameters. Prior histograms (yellow) are based on 100,000 independent draws from the prior distribution presented in Table 2. Posterior histograms (red) are based on the MCMC chain with 90,000 posterior draws obtained after a burn-in period of 10,000 draws.

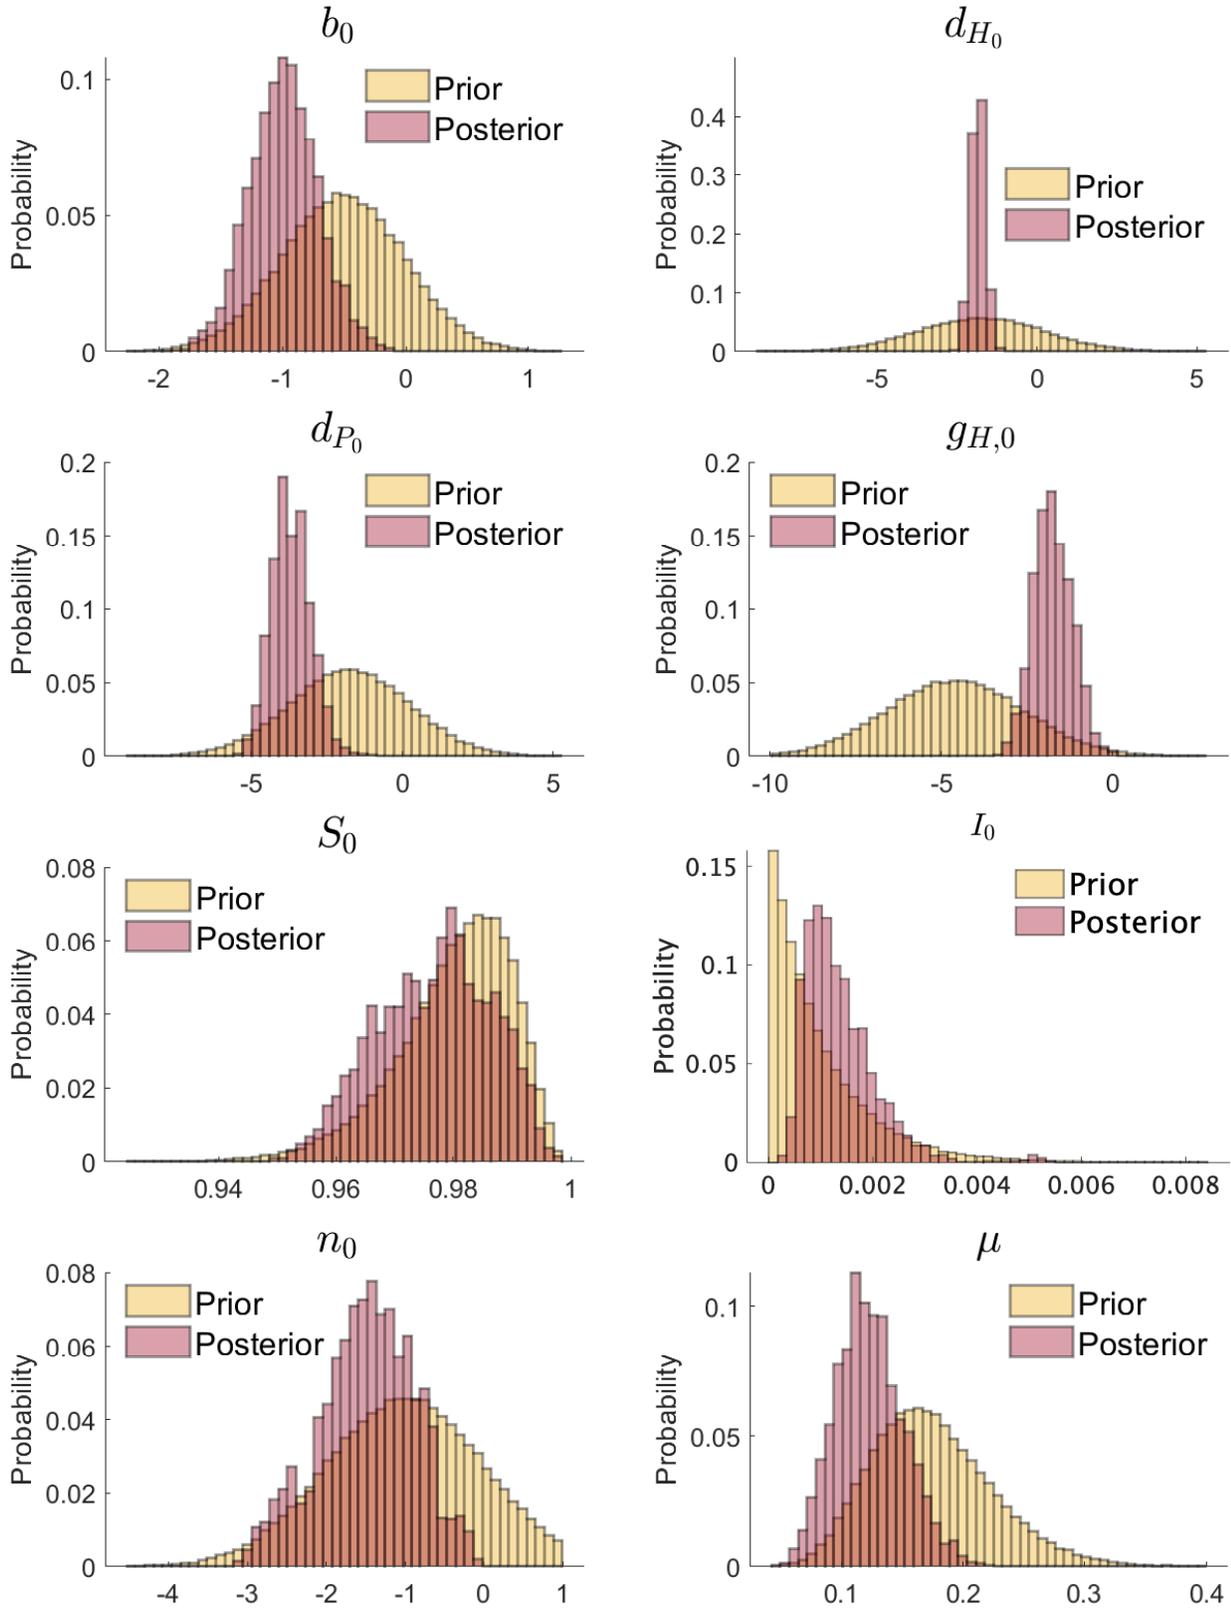


Figure A.4: Initial-value Parameters and Share of False Negatives Parameter. Table 1 presents definitions of those parameters. Prior histograms (yellow) are based on 100,000 independent draws from the prior distribution presented in Table 2. Posterior histograms (red) are based on the MCMC chain with 90,000 posterior draws obtained after a burn-in period of 10,000 draws.

A.3 Other States

Finally, we present the smoothed estimates of the time-varying death probabilities, the share of the population that is infectious, the share of the population that is recovering (in hospitals and at home), and the inflow of hospitalizations as a share of those that are no longer infectious.

Another state variable in our model is the time-varying death probabilities. Those probabilities can change for many reasons. We can enumerate a few. First, medical protocols vary. As health workers learn more about an infection, they can handle patients better, even in the absence of effective treatments. Second, hospitals experience different occupancy rates, with variations in the inflows and total capacity, as the supply of beds and ICU units responds to the crisis. Third, the demographics of patients can change, by varying either in terms of age or in terms of comorbidity levels.³¹

The left panel of Figure A.5 shows the in-hospital death probability (median smoothed, plus the 90% probability band), which went down from over 2.5% in March to less than 1% by early July. The sizeable third peak of COVID infections in the fall of 2020 increased that probability by only around 1.5%, suggesting a considerable degree of improvement in clinical outcomes.

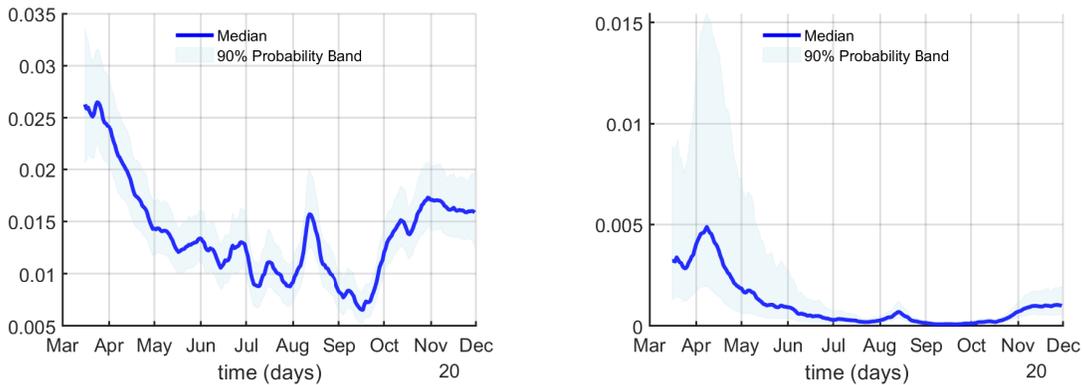


Figure A.5: Death probabilities, Belgium

The right panel of Figure A.5 shows the at-home death probability, which fell from around 0.5% in March to less than 0.1% by early July. Here the changing conditions at retirement communities, which were unprepared for the virulence of COVID-19 in the late winter of 2020, are probably at the core of the estimated variation in death probabilities.

Figure A.6 shows the time series of these smoothed variables throughout our sample. The share of infectious increases to 1.8% by the end of October: more than twice as large as the

³¹Imagine, for example, that individuals with a high probability of infection (e.g., due to their social networks) and high fatality rate (e.g., smokers) got infected in the first wave. As there are fewer of these individuals in the population when the second wave arrives, the measured death rates will mechanically fall.

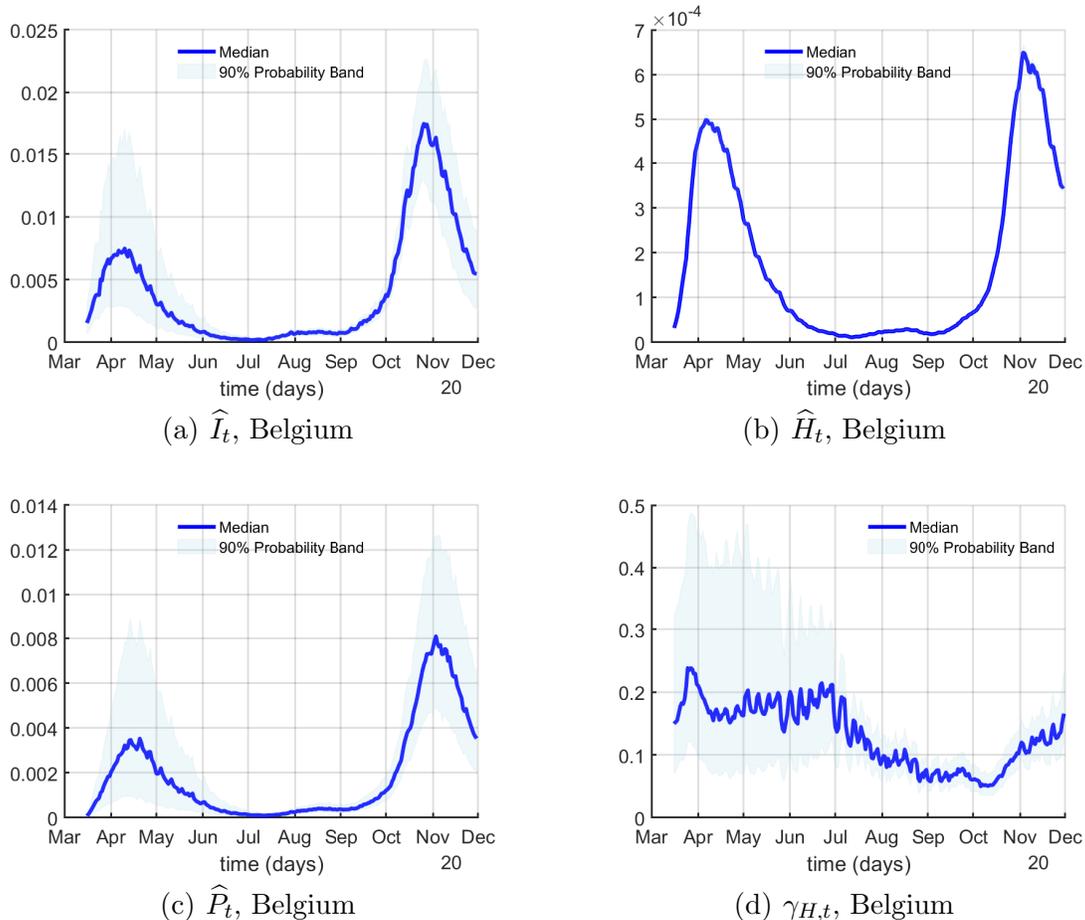


Figure A.6: Other state variables, Belgium

March 15 estimate of 0.7%. This sharp increase in the share of infectious may provide a rationale for the first nationwide lockdown imposed by the Belgian government on March 18, 2020. These measures appear to have had an effect as the number of infectious dropped at the beginning of April and continued to decline, reaching a trough at the end of June 2020. After that, the share of infectious began to increase but at a moderate pace up until the first week of September, when we observe a second exponential increase in the share of infectious, leading to a reintroduction of lockdown measures on October 16, 2020.

The shares of the population recovering in hospitals and at home broadly track the contour of the share of infectious. For example, the first peak of infections and hospitalizations occurs in early April. A similar pattern emerges for the share of the population recovering outside of hospitals. Notice that while the peak of the third wave of infections and the peak of those recovering at home are more than twice as high as the peak of the first wave, during the second wave, hospitalizations peak at a level only marginally higher than the peak of the first wave. This is consistent with the decline in the share of those recovering from COVID-19 in hospitals during the third wave relative to the first one.

A.4 Sentiment and Real GDP Growth Units

Table A.1 shows the OLS coefficients, the 95% confidence intervals for the coefficient estimates (in brackets), and the R^2 of a regression of four-quarter real GDP growth one quarter ahead on a constant and ENS_q over the sample 2000Q1-2019Q4, i.e., $RGDP_{q4q4_{q+1}} = \alpha + \gamma ENS_q + u_t$, where $q \in \{2000Q1, \dots, 2019Q3\}$.³²

Table A.1

Regressors	Coefficients
Constant	1.61 [1.36;1.87]
Economic News Sentiment	1.17 [0.84;1.51]
$R^2 = 0.39$	

Figure A.7 plots the quarterly average of the daily economic news sentiment index expressed in GDP units along with one-quarter-ahead four-quarter real GDP growth.

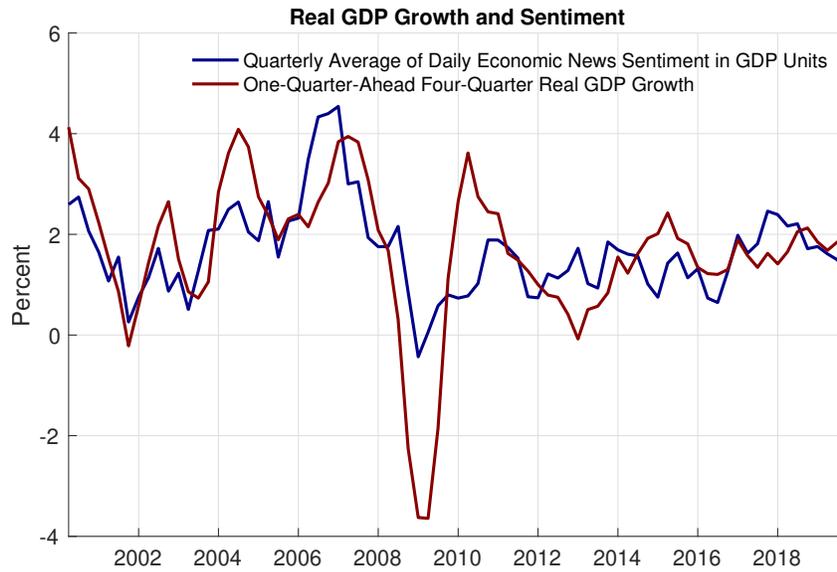


Figure A.7

³²The real GDP data for Belgium were retrieved from the FRED database, Federal Reserve Bank of St. Louis.

A.5 Robustness to State Uncertainty

Figure A.8 gauges the robustness of our IRFs to state uncertainty. First, we draw a sequence of the state variables from the posterior distribution of the smoothed state variables. Second, we estimate an SVAR as in Section 4.1 and compute the point-wise posterior median IRFs. We repeat these two steps 100 times. For each SVAR, the posterior median is based on about 1,000 independent draws of the structural parameters.

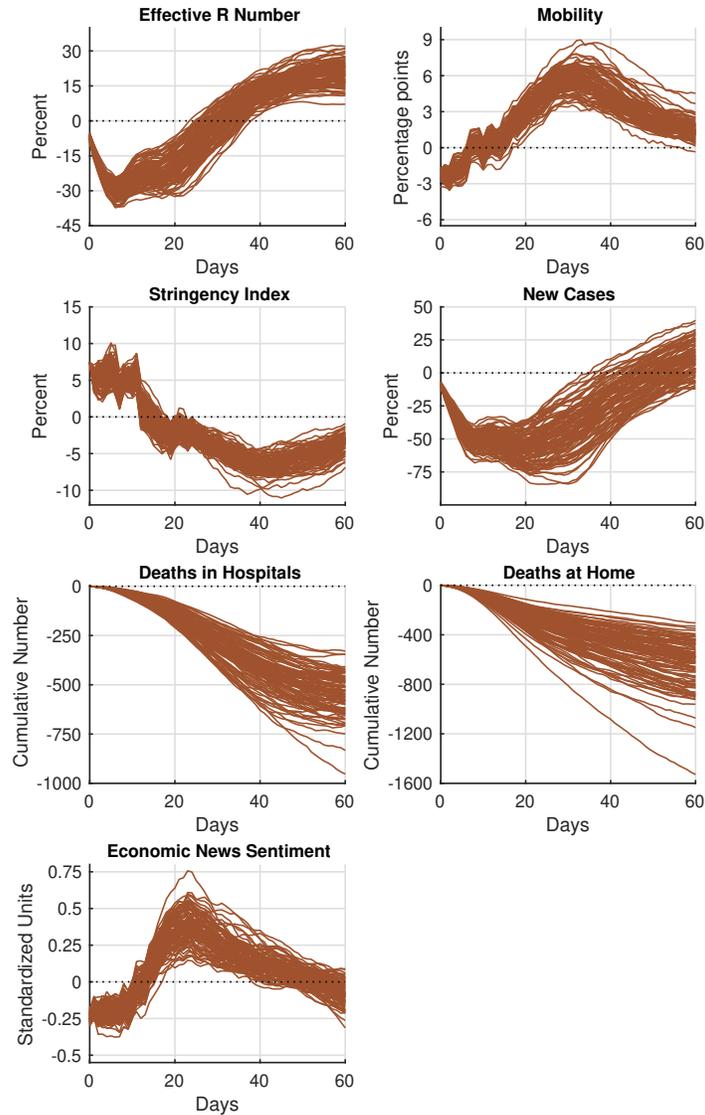


Figure A.8: IRFs to a stringency shock

A.6 Prior Robustness

Figure A.9 shows that the main conclusions from our SVAR analysis are robust to using the prior robust approach for SVARs proposed by [Giacomini and Kitagawa \(2018\)](#). The solid curves represent the point-wise posterior medians, the shaded areas represent the 68% equal-tailed point-wise probability bands, dotted curves represent the set of prior robust posterior means, and dashed-dotted curves depict the 68% robust credible regions. The figure is based on 1,000 independent draws of the reduced-form parameters and 100,000 orthogonal matrices draws for each reduced-form parameter.

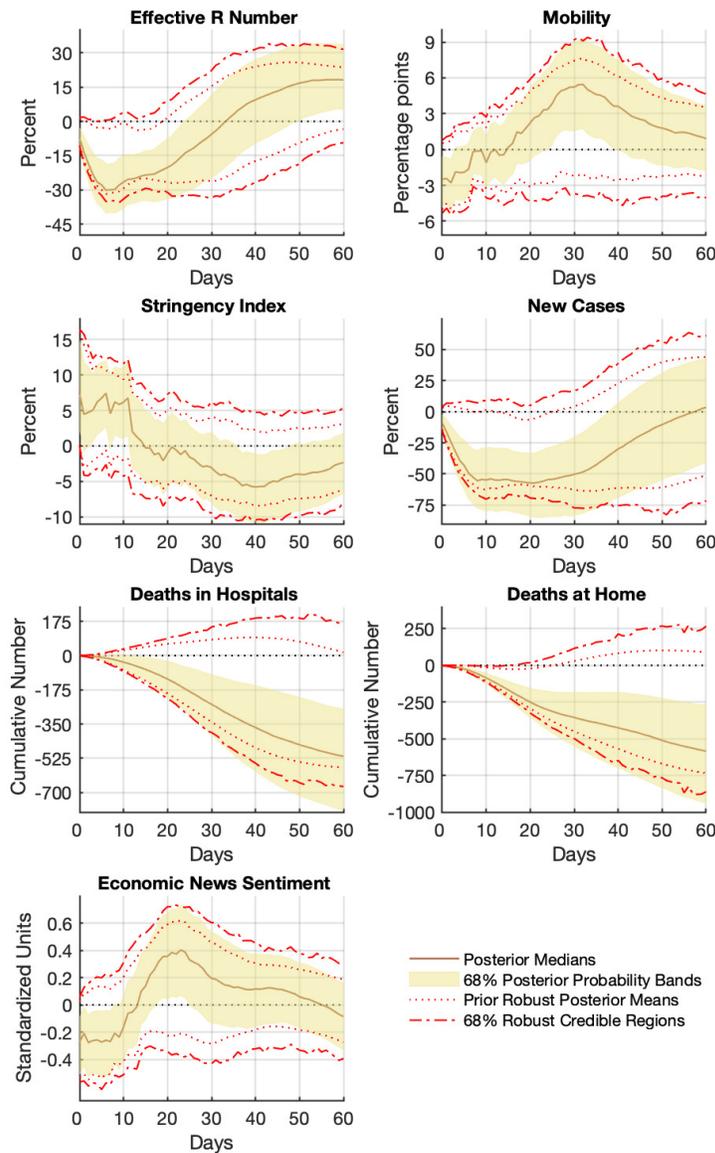


Figure A.9: IRFs to a stringency shock using a prior robust approach

A.7 Robustness to Higher Threshold

Figure A.10 shows that if the cutoff for high stringency is set higher, the main results do not change much. Notice that, in this case, the identified set for the IRF of the reproduction number in the high stringency regime is for the most part above the one in the low stringency regime. This is because, as the threshold for the stringency regime increases, the division between the early and the rest of the sample increases and hence the IRFs possibly reflect more contagious variants.

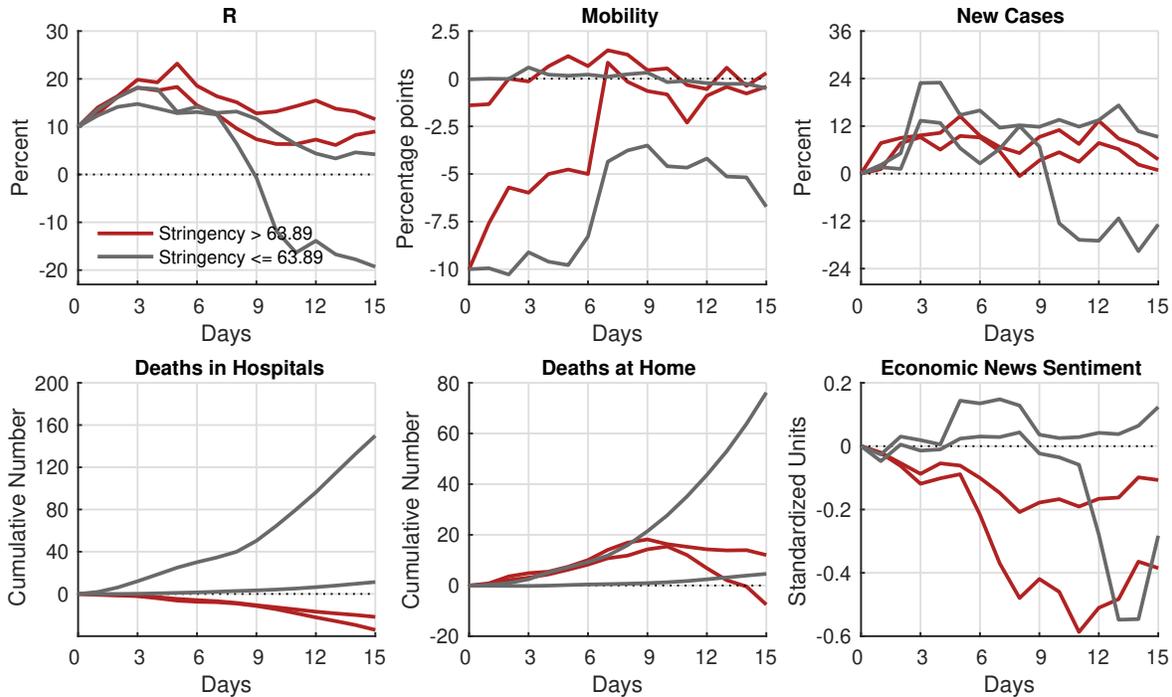


Figure A.10: IRFs to a reproduction shock