Bayesian Policy learning during COVID-19.

Emmanuel Mamatzakis[a], Steven Ongenab and Mike Tsionas

July 2021

Abstract

The rapid spread of COVID-19 across the globe primed a variety of non-pharmaceutical interventions (NPIs). Given these NPIs, whether the SIR parameters followed a Bayesian learning, a random walk pattern or other type of learning with evolving epidemiological data over time has implications for policy learning literature. Using a sample of UK country specific data and also for 168 countries and 51,083 country-date observations (January 1, 2020 to January 9, 2021), we estimate a SIR model with time-varying β and γ parameters in three context of a dynamic panel vector autoregressive model. Although learning does not seem to be taking place, and despite the absence of evidence of governments' learning from the past, most policy measures are effective in reducing the values of the β and γ parameters. We also provide estimates of time-varying β and γ that can be used widely, and we develop novel testing procedures for testing for Bayesian learning.

[a]Department of Management, Birkbeck, University of London, Malet Street, Bloomsbury, London WC1E 7HX, UK. [b]Department of Banking and Finance, University of Zurich, Zurich 8032, Switzerland, Swiss Finance Institute, KU Leuven, Belgium and CEPR, UK. [c]Montpellier Business School, 2300 Avenue des Moulins, Montpellier, France & Lancaster University Management School, Lancaster LA1 4YX, UK.

1. Introduction

In response to the COVID-19 pandemic, a variety of non-pharmaceutical interventions (NPIs) were implemented and adapted over time. Whether policymakers across the world adapted their interventions based on feedback from epidemiological data is of primary interest to curb the pandemic and is also of importance to policy learning literature (Athey & Wager, 2020; Witting, 2017).

Though mimicry in NPI implementation across countries is valuable to lowering uncertainty (Jinjarak et al., 2020; Sebhatu et al., 2020), the subsequent adaptation of NPIs to emerging epidemiological data is important to managing time-varying β and γ parameters that are based on feedback from prior NPIs.

Though mimicry in NPI may be insurance to limiting judgment errors and lowering regret, the time-varying β and γ parameters are essential to improving timing and strengthening control. Time-varying β and γ parameters driven by NPIs have positive and negative consequences. Too often an adjustment could tax an already overstretched healthcare system and may affect the economic, social, and psychological well-being of citizens. Though a normative approach would suggest relying on scientific inquiry to adapt NPIs over time, policymakers may seek information to support their existing beliefs, define problems based on their beliefs, and learn from a limited set of experiences (Witting, 2017).

Conversely, influencing β and γ parameters could also improve planning and response to limit future waves. Though epidemiological models and literature on policy learning call for calibration of β and γ parameters through NPIs the extent of learning among policymakers through variegation in NPIs remains unexplored. The adaptive process of policymaking during COVID-19 is influenced by search and adaptation to limited information to improve understanding of the action-cause-effect associations

under noisy and rapidly evolving information. In the face of the unfamiliar and non-routine context of setting NPI, lowering judgmental errors and improving β and γ for the SIR model is essential.

By proposing a novel method, organized around Bayesian analysis of a time-varying parameter SIR model within a panel autoregression (VAR), we focus on learning-by-policymaking based on how policymakers managed time-varying β and γ parameters. This is an important question as it is critical for further work to understand whether any learning at all is taking place over time; whether policy instruments are significant in reducing the impact of COVID and, finally, in case there is learning whether it is optimal (Bayesian) or not. "Large" deviations from optimal learning would, of course, imply that conditional on the policy instruments, it was not possible to estimate accurately the fundamental parameters of the SIR model.

The proposed model aims to make the following research contributions. First, studies have focused on the efficacy of joint and individual NPIs (Bo et al., 2020), diffusion of NPIs across countries (Aravindakshan et al., 2020), and the political process of implementation of NPIs (Greer et al., 2020). Our model shifts the focus to learning from COVID-19 epidemiological data and changes to NPIs over time. The changes to β and γ parameters, contingent in policy-based learning, influences timing of NPI implementation and intensity. Second, we draw on the policy learning literature in economics and political science. During pandemic, policy learning is very critical, yet it is marred by uncertainty and incomplete information. By proposing and implementing Bayesian inference in a time-varying coefficient vector autoregressive model of SIR, learning based on leniency and stringency of NPIs is important, especially, given the World Health Organization (WHO) recommendation asking countries to learn from evolving country conditions.

We find that predictive Bayes factors in favor of Bayesian optimal learning and against the type of learning that can be calibrated from the data, dominate the second model which receives *some* support in the data, although the evidence is weak. So, we cannot establish *decisively* whether Bayesian learning takes place or not, although we do have *some* evidence against it.

It should be noted that we use international data in order to "gain strength" from the panel structure of the data as the monthly UK data have a very small number of observations.

2. Policy learning during COVID-19

Feedback and cues from the environment are drivers of policy learning (Witting, 2017). The policy learning environment is not only influenced by the normative needs to focus on scientific evidence, but also requires balancing of a variety of political, social, and economic factors that add complexity and volatility. Policy learning is bounded by influential elites, geographic and domain-specific forces that limit the efficacy of prescriptive learning models (Witting, 2017). With policymaking under COVID-19 occurring under variegated inputs from analysts, scientists, citizens, and interest groups. The epistemic diversity in inputs may limit the ability to validate (from different information and interest bases) and evaluate (due to evolving COVID-19 context) the action-effect-cause link.

At the same time, policy learning is ever more critical under COVID-19. Simply adopting and implementing NPIs through mimicry may not be sufficient over time. Calibrating such policies against emerging information is important to balance economic and social costs against health outcomes. Rooted in the notion of dual learning, policy learning (Sabatier, 1988) is based on reliance on heuristic and

analytical processing. Though analytical processing is guided by emerging epidemiological data, policy experience and the context add less meaningful filters through heuristic processing.

In general, optimal learning is Bayesian (Drugowitsch et al., 2019, Jaynes, 2003, Okasha, 2013; see also Tauber et al., 2017) as the Bayes update of beliefs given the prior and in the light of the data, summarizes the new information in the most effective and efficient way Therefore, it is a coherent approach to updating beliefs

In related research, Weible et al. (2010) find the learning potential is greatly reduced when individuals segregate into competing advocacy coalitions. In other words, they only maintain ties to like-minded others. Understanding the attributes of a learning situation is the second question that needs to be addressed to understand how individuals acquire, make sense of and disseminate information. Bayesian learning could be an important learning tool as past heuristics have limited benefits and analytical reasoning may not allow for a full balance of economic and social costs against health costs. Bayesian learning that allows for reliance on priors based on the confluence of analytical and heuristics actions occurring in the respective context. Because the tools of instrumental and social learning are seldom present in a pandemic situation, Bayesian learning relies on priors that are based on past outcomes and processes driven by a diverse set of inputs, interests, and actions based on non-trivial degrees of coordination, collaboration, and conflict. The priors are a reflection of convergent processes as policymakers try to make sense of the ambiguous situation, where the possibility of informed learning under time pressure is less feasible. Though instrumental learning is a norm in policy learning (May, 1992; Sabatier, 1988), we propose a model of policy learning.

2. The SIR model with time-varying parameters

The basic SIR model is

$$\frac{dS}{dt} = -\beta I(t)S(t), \tag{1}$$

$$\frac{dI}{dt} = I(t)[\beta S(t) - \gamma] \tag{2}$$

$$\frac{dR}{dt} = \gamma I(t), \tag{3}$$

where S, I, R denote the number of susceptible people, the number of infected, and the number of recovered persons. Here, β is the daily transmission rate, and γ is the daily transition rate from infected to recovered (which, so far, seems to be rather close to zero). In the first difference form, we have

$$S_{t+1} - S_t = -\beta I_t S_t, \tag{4}$$

$$I_{t+1} - I_t = I_t(\beta S_t - \gamma), \tag{5}$$

$$R_{t+1} - R_t = \gamma I_t. \tag{6}$$

It is well known that managing a SIR epidemic means modifying the constants β and γ.

To account for learning, we assume that the parameters β and γ are time-varying. However, we have data on several countries and the equations above cannot hold exactly so we introduce error terms for country $i \in \mathcal{I} = \{1, \dots, n\}$ and time $t \in \mathcal{T} = \{1, \dots, T\}$. We write $(i, t) \in \mathcal{I} \times \mathcal{T} \equiv \mathcal{J}$.

Therefore, we have the modified SIR model:

$$S_{i,t+1} - S_{i,t} = -\beta_{i,t} I_{i,t} S_{i,t} + v_{i,t,1}, \tag{7}$$

$$I_{i,t+1} - I_{i,t} = I_{i,t}(\beta_{i,t} S_{i,t} - \gamma_{i,t}) + v_{i,t,2}, \tag{8}$$

$$R_{i,t+1} - R_{i,t} = \gamma_{i,t} I_{i,t} \; \forall (i,t) + v_{i,t,3} \in \mathcal{J}. \tag{9}$$

Let the parameters be

$$\boldsymbol{\theta}_{i,t} = \begin{bmatrix} \beta_{i,t} \\ \gamma_{i,t} \end{bmatrix}. \tag{10}$$

For statistical inference we assume a panel vector autoregressive model for the parameters:

$$\begin{gathered} \boldsymbol{\theta}_{i,t} = \begin{bmatrix} a_{i,1} \\ a_{i,2} \end{bmatrix} + \begin{bmatrix} a_{i,11} & a_{i,12} \\ a_{i,21} & a_{i,22} \end{bmatrix} \boldsymbol{\theta}_{i,t-1} + \begin{bmatrix} \boldsymbol{x}'_{i,t-1}\boldsymbol{\alpha}_1 \\ \boldsymbol{x}_{i,t-1}\boldsymbol{\alpha}_2 \end{bmatrix} \begin{bmatrix} v_{i,t,4} \\ v_{i,t,5} \end{bmatrix} \\ \Rightarrow \end{gathered} \tag{11}$$

$$\theta_{i,t} = a_i + A_i \theta_{i,t-1} + \boldsymbol{X}_{i,t-1}\boldsymbol{\alpha} + \tilde{v}_{i,t}, \tag{12}$$

where $\boldsymbol{x}_{i,t}$ are a $k \times 1$ pre-determined regressors with coefficients, $\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2 \in \mathbb{R}^k$, $\boldsymbol{X}_{i,t-1} = \begin{bmatrix} \boldsymbol{x}'_{i,t-1} & \\ & \boldsymbol{x}'_{i,t-1} \end{bmatrix}$, $\boldsymbol{\alpha} = \begin{bmatrix} \boldsymbol{\alpha}_1 \\ \boldsymbol{\alpha}_2 \end{bmatrix}$ $\tilde{v}_{i,t} = \begin{bmatrix} v_{i,t,4} \\ v_{i,t,4} \end{bmatrix}$, and $v_{i,t,4}$ and $v_{i,t,5}$ are statistical error terms.[1]

The central question is whether there is learning in dealing with COVID. it is well known that parameters β and γ depend on social distancing, other government measures, as well as underlying fundamental characteristics in $x_{i,t}$. The first question we deal with is whether $\boldsymbol{\alpha}_1 = \boldsymbol{\alpha}_2 = \mathbf{0}$.

[1]It is possible to include $\boldsymbol{x}_{i,t}$ instead of $\boldsymbol{x}_{i,t-1}$ and, in fact, we test for it. When the interval of observation is short, this assumption can be easily defended as it takes actiom to implement announced policy measures.

The second and, perhaps, more important question is whether there is any Bayesian learning about $\boldsymbol{\theta}_{i,t}$ (that is, β and γ over countries and time) or a different type of learning –as we know Bayesian learning is the only coherent way of updating beliefs in the light of the data.

There are various posteriors that we can use in this context. First, define $\boldsymbol{\theta}_t = [\boldsymbol{\theta}_{i,t}\ \forall i \in \mathcal{I}]$. One can consider the posterior $p(\boldsymbol{\theta}_t|D_{t-1})$ where D_{t-1} is data up to period $t-1$. Another posterior can be $p(\boldsymbol{\theta}|D)$ where D denotes the entire data and $\boldsymbol{\theta} = [\boldsymbol{\theta'}_t\ \forall t \in \mathcal{T}]$. As $p(\boldsymbol{\theta}_t|D_{t-1})$ converges to $p(\boldsymbol{\theta}|D)$ this does not allow us to test for Bayesian learning.

2.1. Random walk behavior of the β and γ

Our first test for Bayesian learning is whether $\boldsymbol{\theta}_{i,t}$ follows a random walk with drift *conditional* on the $x_{i,t}$s, that is whether we have:

$$H\colon a_{i,12} = a_{i,21} = 0, \text{ for some or all } i \in \mathcal{I}. \tag{13}$$

In this case, we would have, from (11), we would have

$$\boldsymbol{\theta}_{i,t} = \begin{bmatrix} a_{i,1} \\ a_{i,2} \end{bmatrix} + \begin{bmatrix} a_{i,11} & 0 \\ 0 & a_{i,22} \end{bmatrix} \boldsymbol{\theta}_{i,t-1} + \begin{bmatrix} \boldsymbol{x'}_{i,t-1}\boldsymbol{\alpha}_1 \\ \boldsymbol{x}_{i,t-1}\boldsymbol{\alpha}_2 \end{bmatrix} \begin{bmatrix} v_{i,t,4} \\ v_{i,t,5} \end{bmatrix}. \tag{14}$$

Conditionally on the $\boldsymbol{x}_{i,t}$s, $\beta_{i,t}$ and $\gamma_{i,t}$ follow random walks with drifts $a_{i,1}$ and $a_{i,2}$:

$$\beta_{i,t} = a_{i,1} + a_{i,11}\beta_{i,t-1} + \boldsymbol{x'}_{i,t-1}\boldsymbol{\alpha}_1 + u_{i,t,3}, \tag{15}$$

$$\gamma_{i,t} = a_{i,1} + a_{i,22}\gamma_{i,t-1} + \boldsymbol{x'}_{i,t-1}\boldsymbol{\alpha}_2 + u_{i,t,4}. \tag{16}$$

If, indeed, (13) is correct for some or all $\mathrm{i} \in \mathcal{I}$, then some policy effects in ($\mathbf{x}_{\mathrm{i,t}}$) may be significant but conditional on them, no other actions are taken to correct the values of $\beta_{\mathrm{i,t}}$ and $\gamma_{\mathrm{i,t}}$. If (13) is rejected, then one might lean to believe that there are actions based on some type of learning that induce other sorts of policy actions to reduce the

values of $\beta_{i,t}$ and $\gamma_{i,t}$. How do we know this is Bayesian *learning*, however? The answer is that it comes through formal inference.

2.2. Comparing Bayesian learning with actual learning

Although Bayesian learning is known to be optimal, there might be other types of learning which we can estimate from calibrated time-varying parameters of the SIR model. In the absence of learning, we would expect the two parameters of the SIR model to follow random walks. Such other types of learning can be compared formally with optimal (Bayesian) learning. The comparison is performed formally through Bayes factors based on marginal likelihoods derived from Sequential Monte Carlo) also known as particle filtering techniques.

This would require other estimates of $\beta_{i,t}$ and $\gamma_{i,t}$ that can be calibrated from the data and, in turn, check whether these are "broadly" consistent with (14). Several works calibrate these parameters for the whole sample see, for example, Schaback (2020), and Cooper et al. (2020) set the parameters of the SIR model by visual inspection. Another approach is setting the model to estimate time-varying parameters as follows:

$$\hat{\beta}_{i,t} = R_{0,i,t}\hat{\gamma}_{i,t}, \tag{17}$$

where R_0 represents the famous "R-zero-index" (reproduction ratio, the average number of individuals infected by a single infected individual when everyone else is susceptible). Another estimate is

$$\hat{\beta}_{i,t} = R_{i,t}\hat{\gamma}_{i,t}, \tag{18}$$

where $\mathcal{R}_{i,t}$ is the adjusted reproduction number, defined as $R_{i,t} = R_{0,i,t}\frac{S_{i,t-1}}{N_i}$ (the average number of individuals infected by a single infected individual when a fraction

$\frac{S_{i,t-1}}{N_i}$ of individuals is susceptible.) and

$$\hat{\gamma}_{i,t} = \frac{R_{i,t+1} - R_{i,t}}{I_{i,t}}. \tag{19}$$

Perhaps it is more reasonable to set

$$\hat{\gamma}_i = T^{-1} \sum_{t=1}^{T} \left(\frac{R_{i,t+1} - R_{i,t}}{I_{i,t}} \right), \tag{20}$$

but this cannot be compared fully with our $\boldsymbol{\theta}_{i,t}$ unless we have a steady state which is a strong assumption. The estimates in (17) and (19) although noisier, provide at least a good benchmark of comparison with (14).

We assume that the error terms

$$\boldsymbol{v}_{i,t} \sim \mathcal{N}(0, \Sigma), \tag{21}$$

so, all errors are correlated. Our priors on the parameters are

$$\begin{aligned} p(a_i) &\propto \text{const.}, \\ p(A_i) &\propto \text{const.} \\ p(\Sigma) &\propto |\Sigma|^{-3/2}, \end{aligned} \tag{22}$$

see Zellner (1971, page 225 formula 8.9). For statistical inferences, we use Sequential Markov Carlo also known as Particle Filtering (see Technical Appendix A).

3. Data

We draw on three data sources. The NPI data is from the Oxford COVID-19 Government Response Tracker (OxCGRT) (Hale et al., 2020), and country-level controls are from the World Bank Development Indicators. The daily COVID-19 case data for the SIR model are from the Johns Hopkins University's Center for Civic Impact. OxCGRT collects publicly available information on 19 indicators of government responses related to containment and closure policies, economic policies,

and health system policies, which are combined into four indices ranging from 0 to 100. The indices include the number and strictness of government policies and do not indicate appropriateness or effectiveness response.

We control for GDP based constant 2010 US dollars, population density, median age, proportion of the population aged 65 and older, proportion of population age 70 and older, GDP per capita, cardiovascular death rate, diabetes prevalence, hospital beds per thousand people, life expectancy, and human development index. We also group the countries by regions due to a greater propensity to learn from regional countries: Western Europe, Eastern Europe, Southern Europe, Northern Europe, Asia & Pacific, and Americas. The sample descriptives are presented in Table S1.

Data on government interventions collected by (Hale et al., 2020) concern three main areas of interventions: a) containment and closure, b) health system, and c) economic stimulus. All the indicators are available on a daily and monthly basis.

The containment and closure interventions include eight sub-indicators: i) school closing, ii) workplace closing, iii) cancellation of public events, iv) restrictions on gatherings size, v) public transport closed, vi) stay at home requirements, vii) restrictions on internal movement, and viii) restrictions on international travel.

The second area of interventions include health system: i) public information campaigns, ii) testing policy, and iii) contact tracing. Since these policies help to cope with the pandemic quicker, they may be also discounted in stock prices.

The third area includes economic stimulus packages such as: income support, and debt or contract relief for households. These stimulus affect the economy through various channels. For instance, stimulus supports consumption and spending in times of distress; hence, they may significantly affect local equity markets.

Finally, besides the individual measures, we also consider the overall Stringency Index by Hale et al. (2020). The index aggregates the data pertaining is re-scaled to create a score between 0 and 100. This index provides a synthetic measure of the intensity of different non-medical government interventions during the pandemic.

Some of the policies in considered in this study can be implemented either as 1) targeted policies, limited to certain geographical region, category of business, or group of residents, or 2) general policies, applied to the entire country or population (for details, see Hale et al. 2020). We consider the scale' of these polices, and we introduce the additional general indicator to indicate whether the policy applies across the entire country or population.

All the changes in government policies are tracked daily and monthly. Therefore, when we perform the regressions based on weekly returns, we calculate the weekly averages for the considered period.

4. Results

In Figures 1 and 2 (see Supplement C) we present selected results about recursive posterior-mean-estimates of filtered β_t and γ_t. Figure 3 presents the plots of β_t and

γ_t by the regions. In Figures 4 and 5 we present recursive Bayes factors in favor of a random walk. Evidently, the odds in favor of random walk behavior in filtered β_t and γ_t are great and support the idea of a random walk. In Figures 6 and 7 we provide Bayes factors in favor of the estimates in (18) and (19), and against the Bayesian (learning) model. As these predictive Bayes factors are marginal, the Bayesian model receives some support in the light of the data, although the evidence is weak. So, there is probably no Bayesian learning on the part of the authorities.

Figure 1. Selected results

Note: Authors' estimations.

Figure 2. Selected countries A

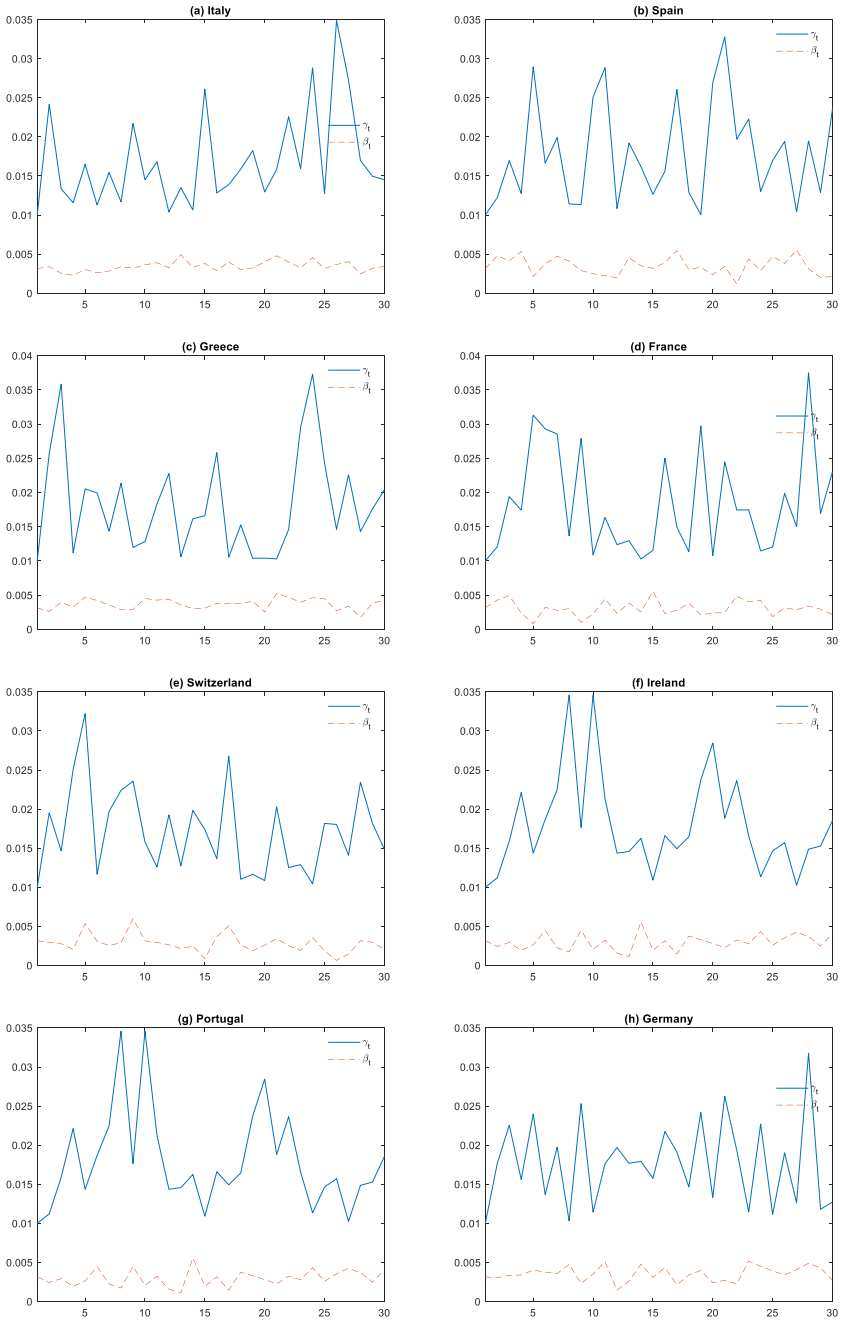

Note: Authors' estimations.

Figure 2. Relative beta β_t and γ_t effects by regions, end of sample period

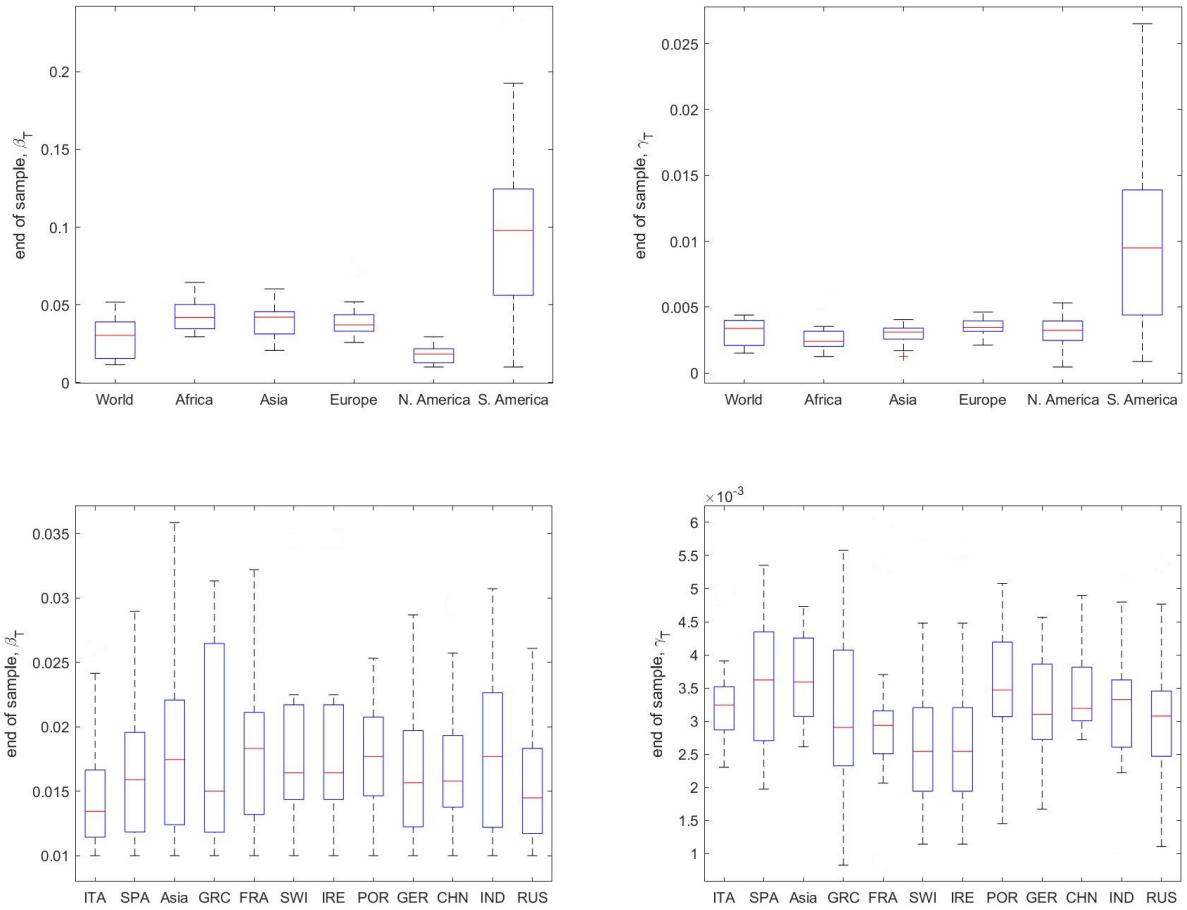

Figure 4. Selected countries B

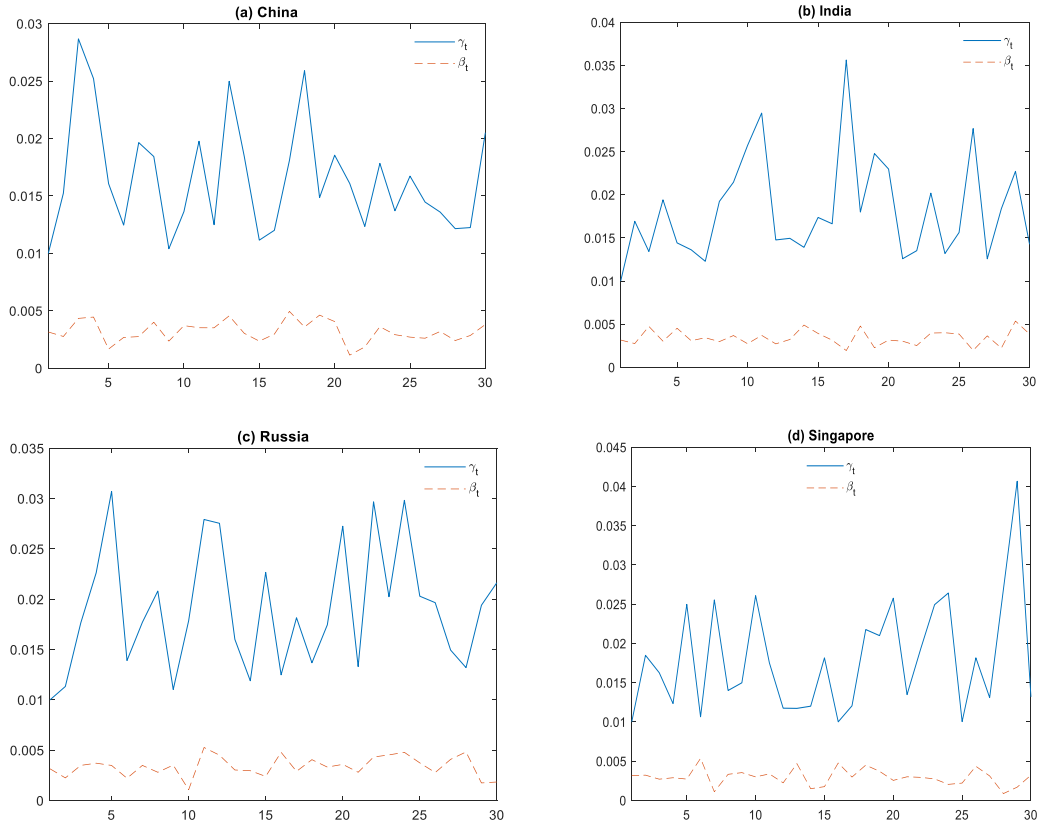

Note: Authors' estimations.

Figure 4. Recursive Bayes factors in favor of a random walk in both β_t and γ_t

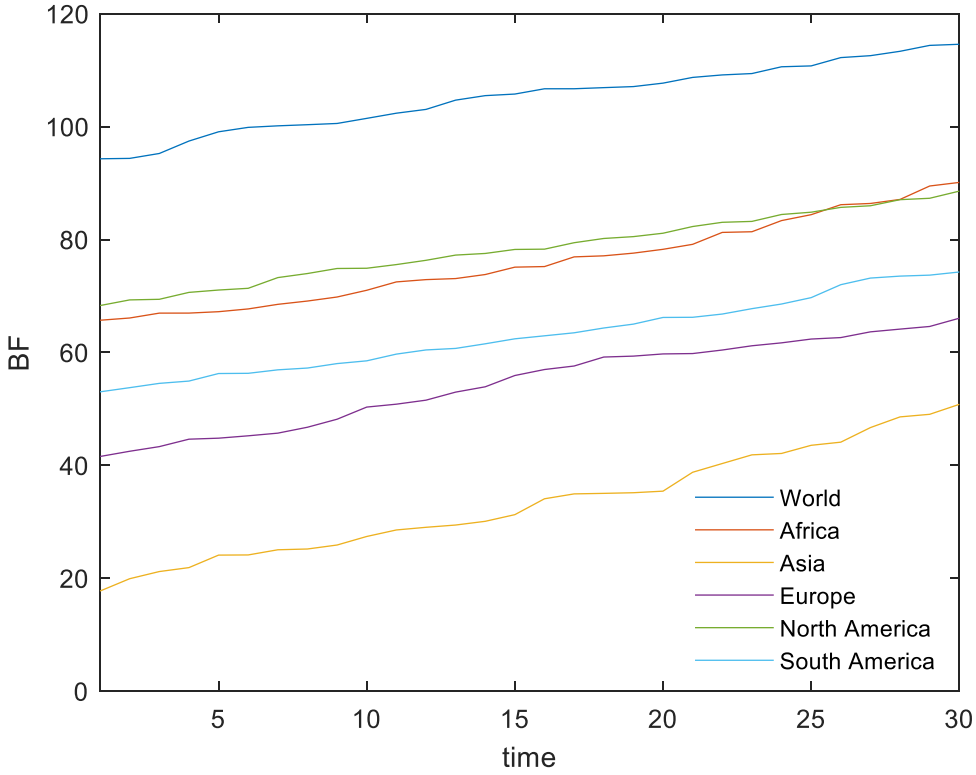

Note: Authors' estimations.

Figure 5. Recursive Bayes factors in favor of a random walk in both β_t and γ_t

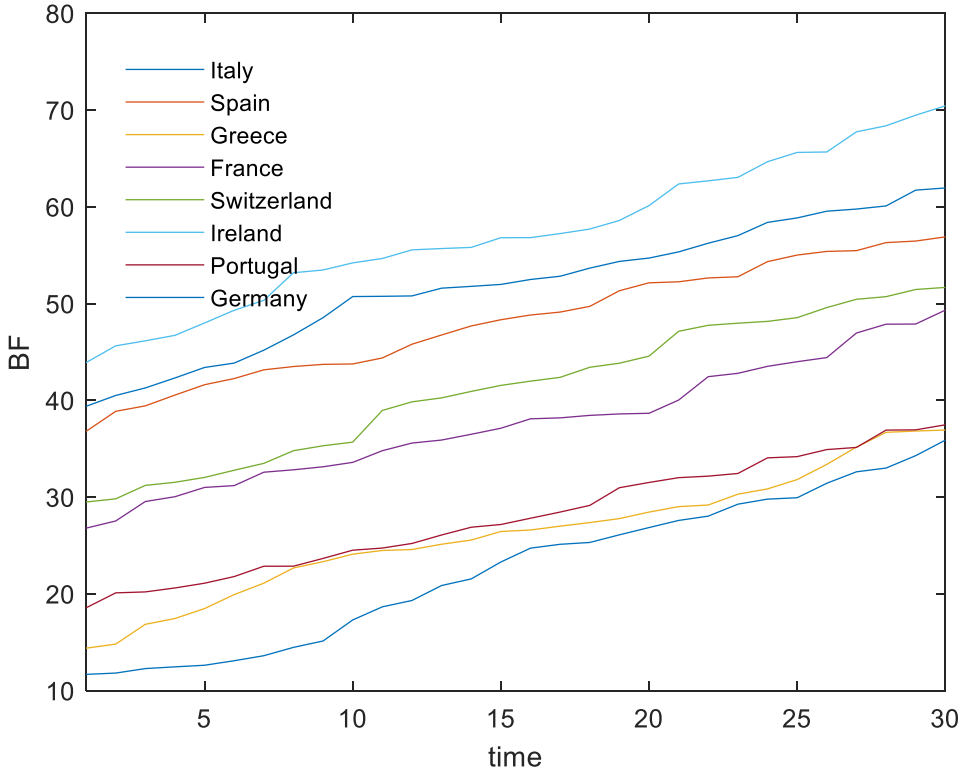

Note: Authors' estimations.

Figure 6. Recursive Bayes factors against Bayesian learning and in favor of calibrated time-varying values, A

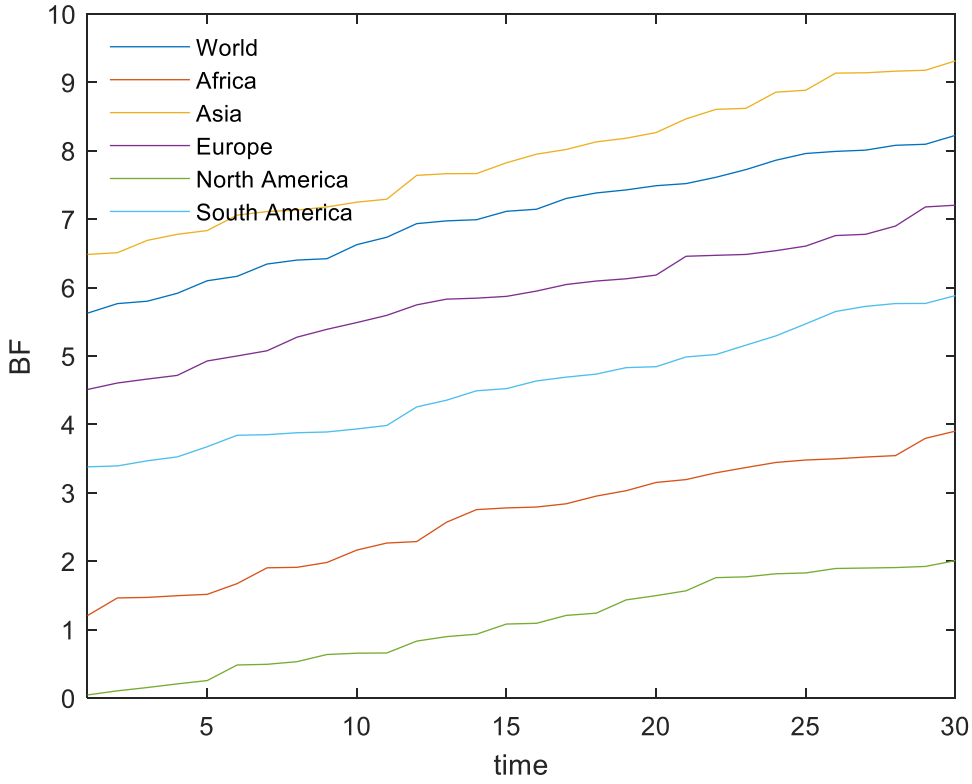

Note: Authors' estimations.

Figure 7. Recursive Bayes factors against Bayesian learning and in favor of calibrated time-varying values, B

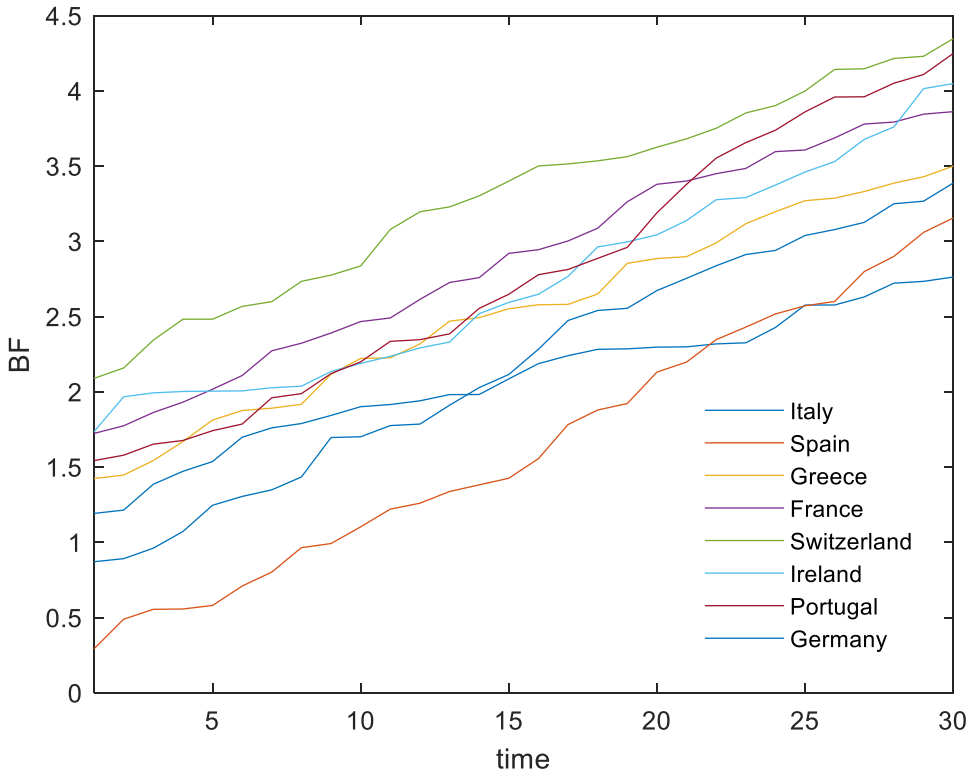

Note: Authors' estimations.

Figure 8. Recursive Bayes factors against Bayesian learning and in favor of calibrated time-varying values, C

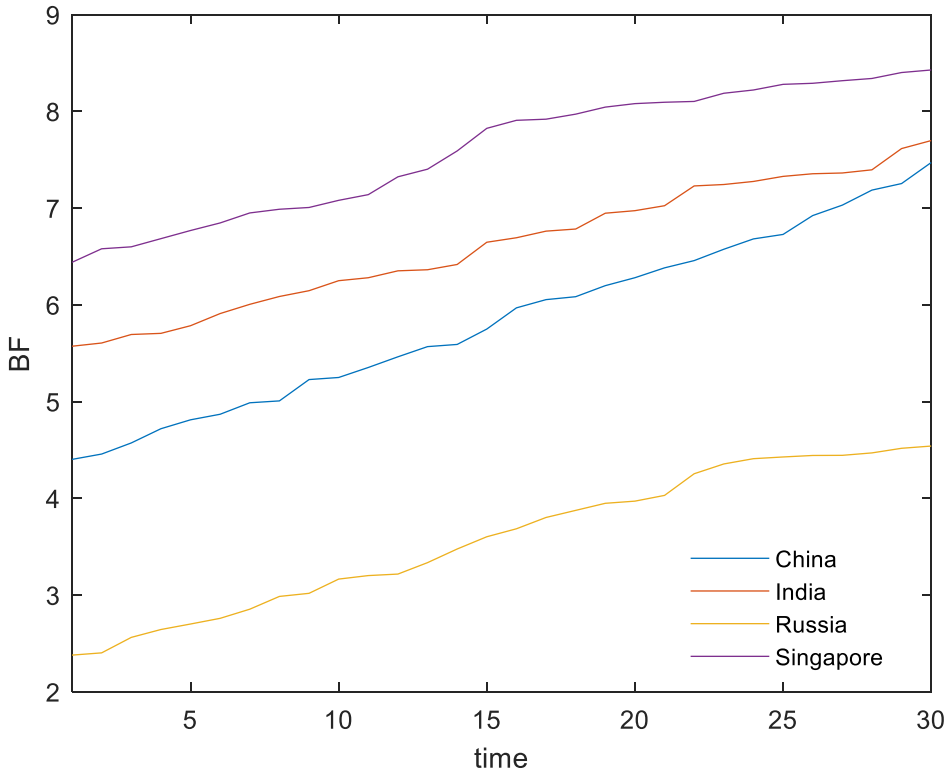

Note: Authors' estimations.

Figure 9. Recursive Bayes factors in favor of a model with time-varying and against constant but country-specific coefficients

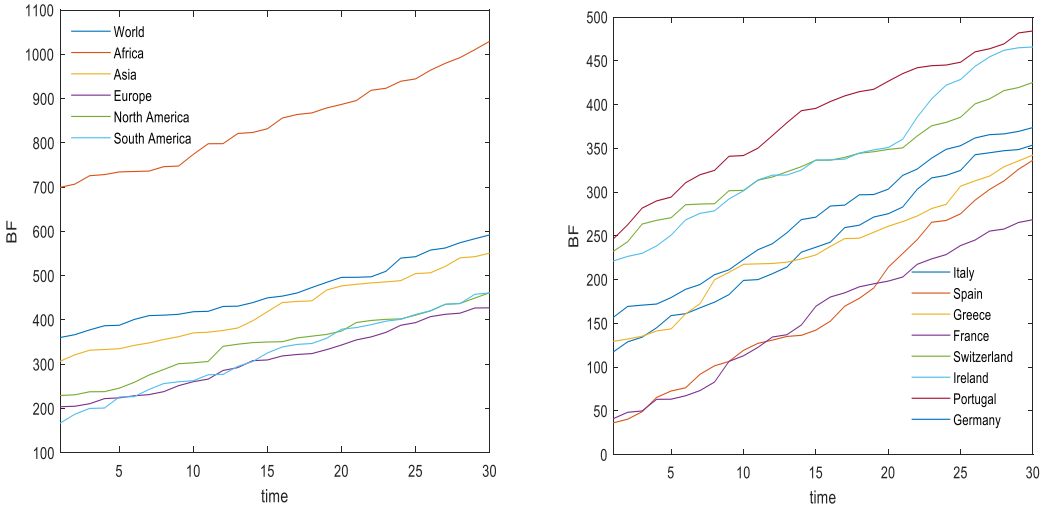

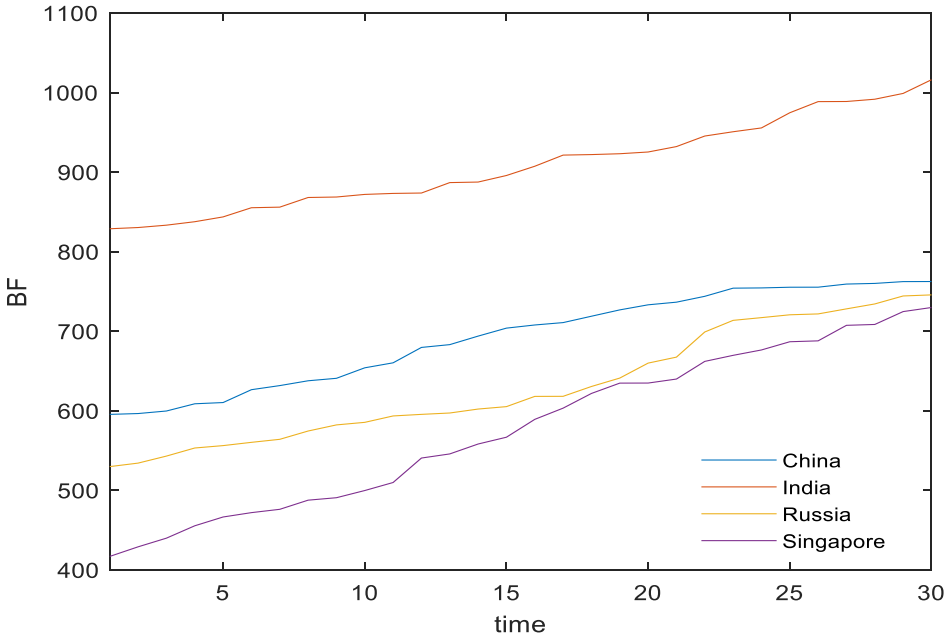

Note: Authors' estimations.

In Figure 9 we report Bayes factors in favor of restricted time-varying-parameter panel VAR and against certain more restricted models which are overwhelmingly rejected by the data including as well as panel VAR model without the policy covariates. A a random walk model without covariates is marginally rejected showing that a random walk hypothesis could be consistent with the data.

5. Repayment of household debt in the UK

Repayment of household debt in the UK (which is the country of interest), denotes D_t, (in logs) will be related to estimates of $\beta_{i,t}$ and $\gamma_{i,t}$, viz. the main epidemiological parameters using the following model:

$$D_{i,t} = a_0 + a_1\beta_{i,t-1} + a_2\gamma_{i,t-1} + \boldsymbol{x}'_{i,t-1}\boldsymbol{\delta}_1 + e_t, \quad (23)$$

where a_0, a_1, a_2 and $\boldsymbol{\delta_1}$ are unknown parameters, $\boldsymbol{x}_{i,t}$ has been introduced before the epidemiological parameters are lagged once to allow for the hypothesis that households use a one-month planning horizon and $e_{i,t}$ is an error term. As the number of monthly observations is small, we impose a tight prior on the parameters of (23), viz. the coefficients have normal $N(0,1)$ priors and the error variance follows the standard Jeffreys prior. The posterior means of a_1 and a_2 are respectively -0.0012 (0.002) and 0.0015 (0.000056) so, only the infection rate from infected to recovered seems to be statistically significant. All coefficients in $\boldsymbol{\delta}_1$ are statistically significant. In the Figure 10 below we present the plot of actual versus one-step-ahead (dash line) predictions of debt repayments. The one step ahead predictions closely follow the actual debt

repayments.

Figure 10. UK household repayment one step ahead predictions.

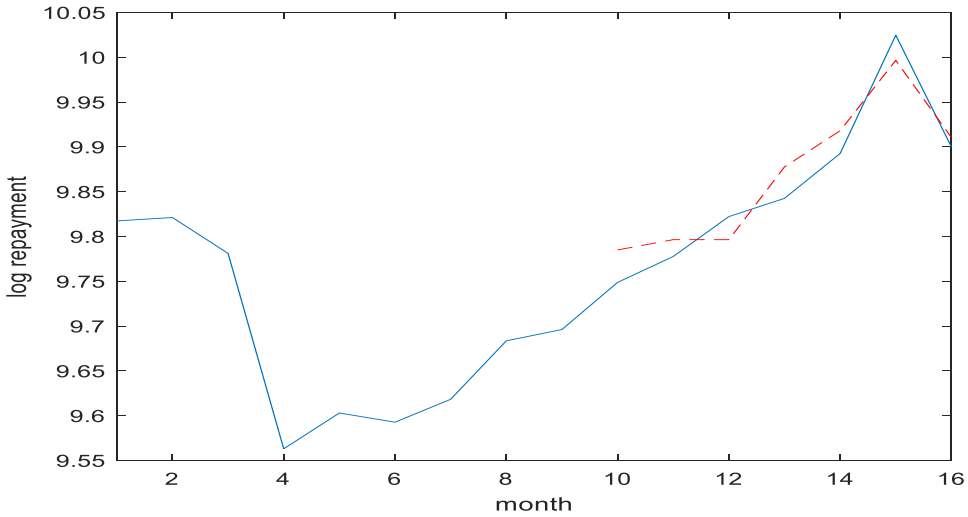

Note: Authors' estimations

In Figure 11 to 16 we report Impulse Response Functions (IRFs) of a VAR that shows the response of the main variable of our analysis household debt repayments to a plethora of Covid-19 related shocks. The IRFs concern 8 months ahead of one plus or minus standard deviation shock in the corresponding Covid-19 related shock. For example, Figure 11 shows that the response of household repayment to a shock in number of Covid-19 cases is positive over the first two months, though it is on declining trajectory, thereafter there is a roller coaster type of responses prior to convergence in three-month time. Similar patterns in the response of household debt repayments are observed to shocks of other variables in the remaining Figures, but for Figure 14, where the IRF shows that the response of household debt repayments to a shock in international movement restrictions is negative in the first two months. This implies that shocks in international movement restrictions would negatively affect household

debt repayments. So despite consistency in IRFs across all shocks there is also some variability that warrants further analysis.

Figure 11. IRF of response of household repayment to a shock in number of cases.

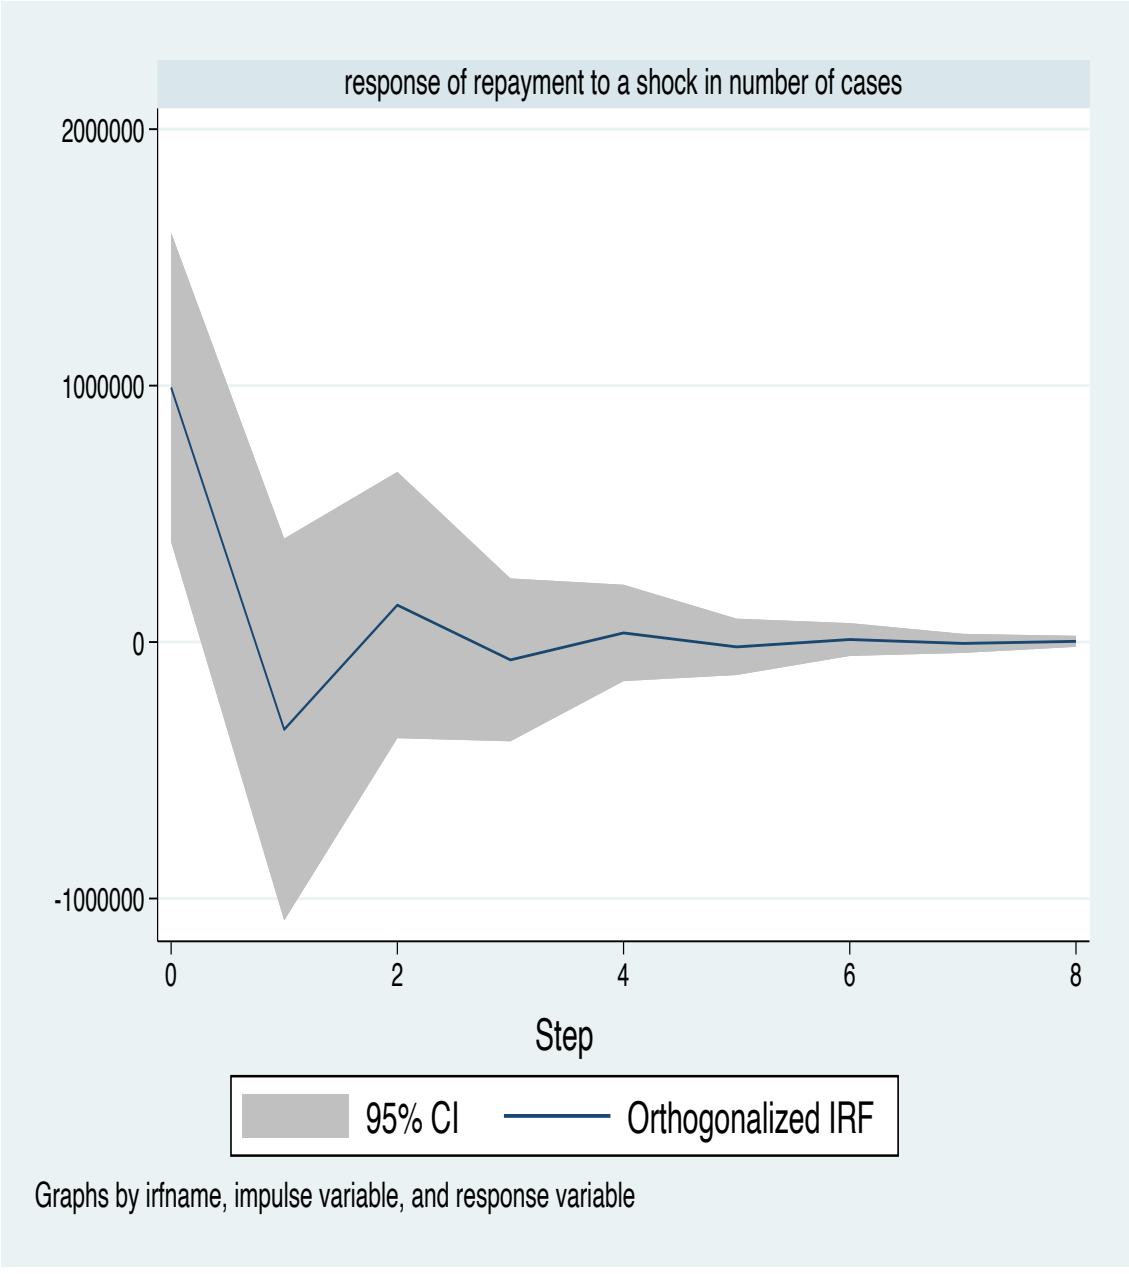

Note: Authors' estimations.

Figure 12. IRF of response of household repayment to a shock in number of hospitalisations.

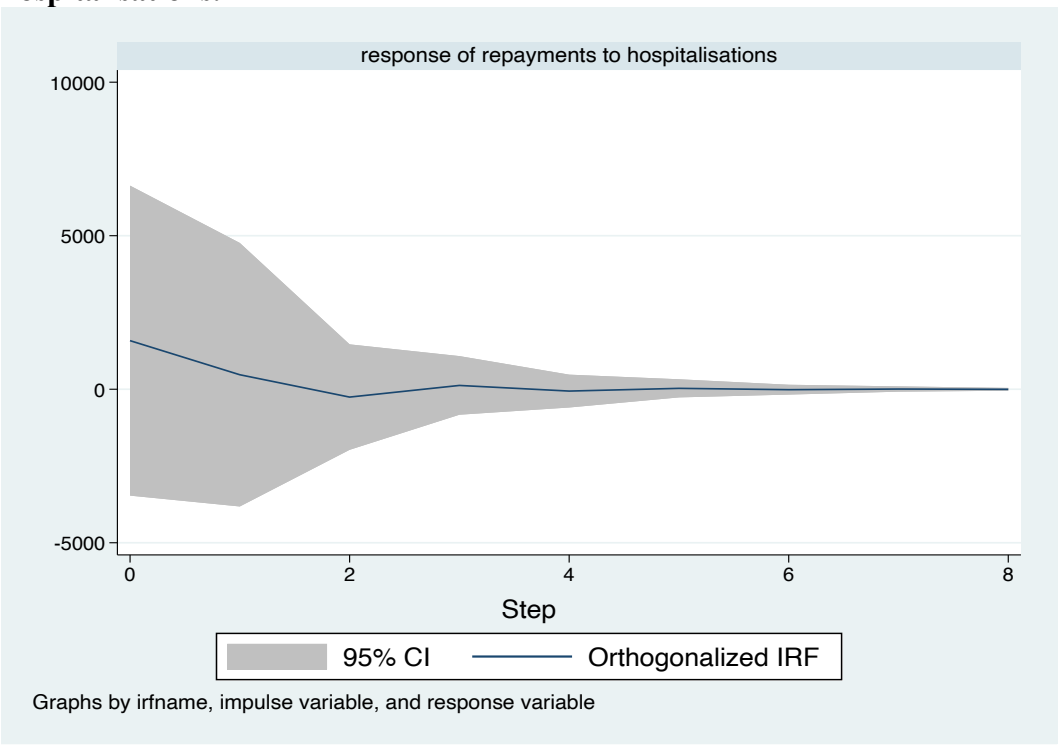

Note: Authors' estimations.

Figure 12. IRF of response of household repayment to a shock in number of deaths.

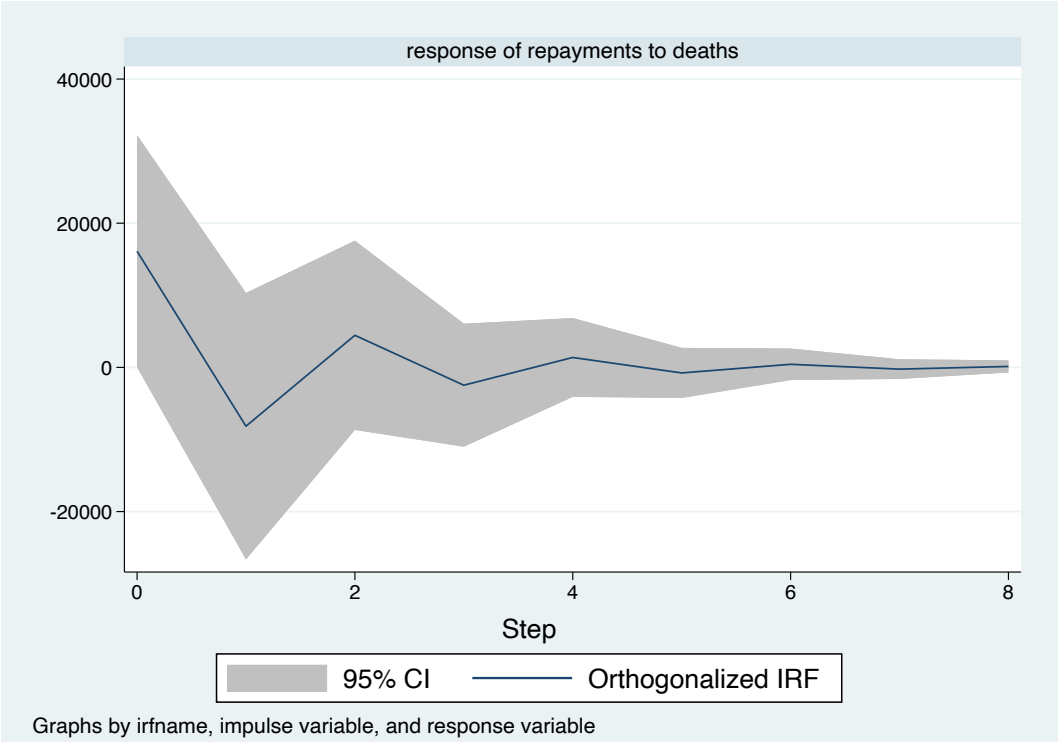

Note: Authors' estimations.

Figure 13. IRF of response of household repayment to a shock in internal movement restrictions.

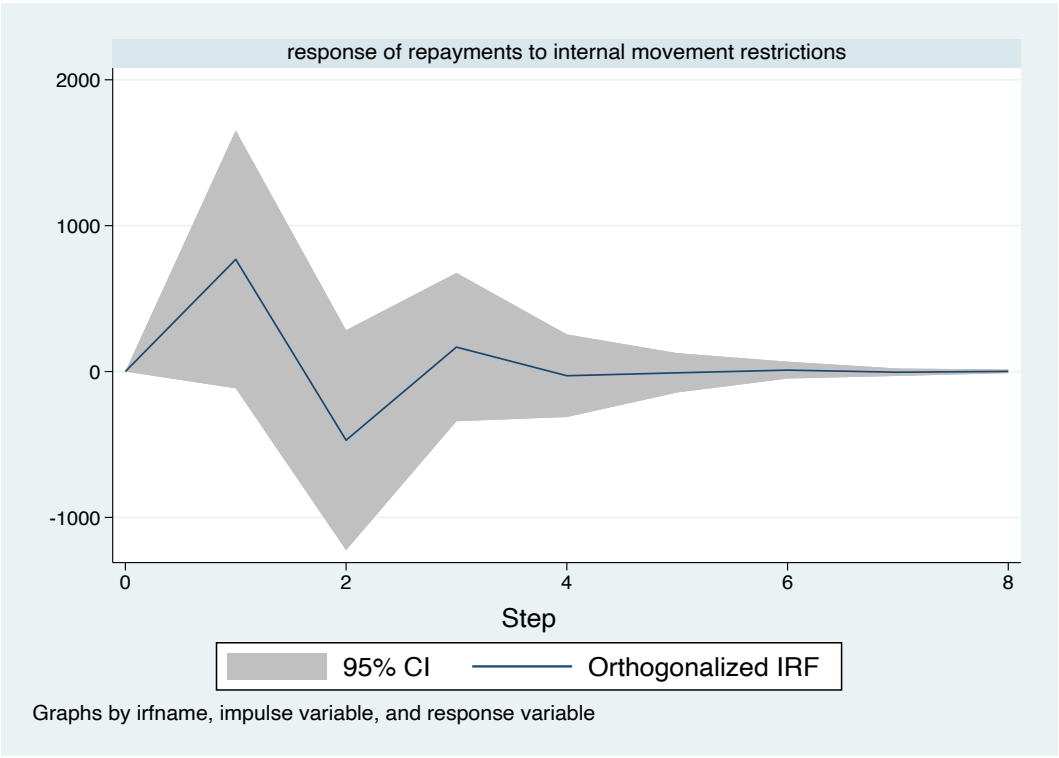

Note: Authors' estimations.

Figure 14. IRF of response of household repayment to a shock in international movement restrictions.

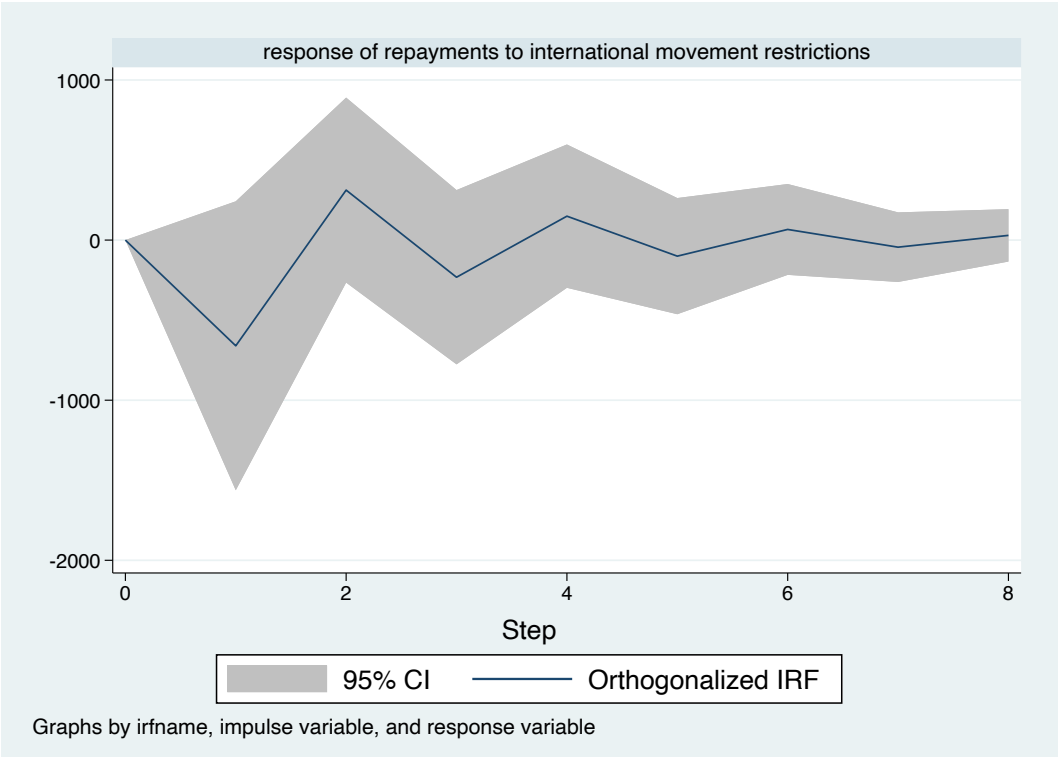

Note: Authors' estimations.

Figure 15. IRF of response of household repayment to a shock in stay home restrictions.

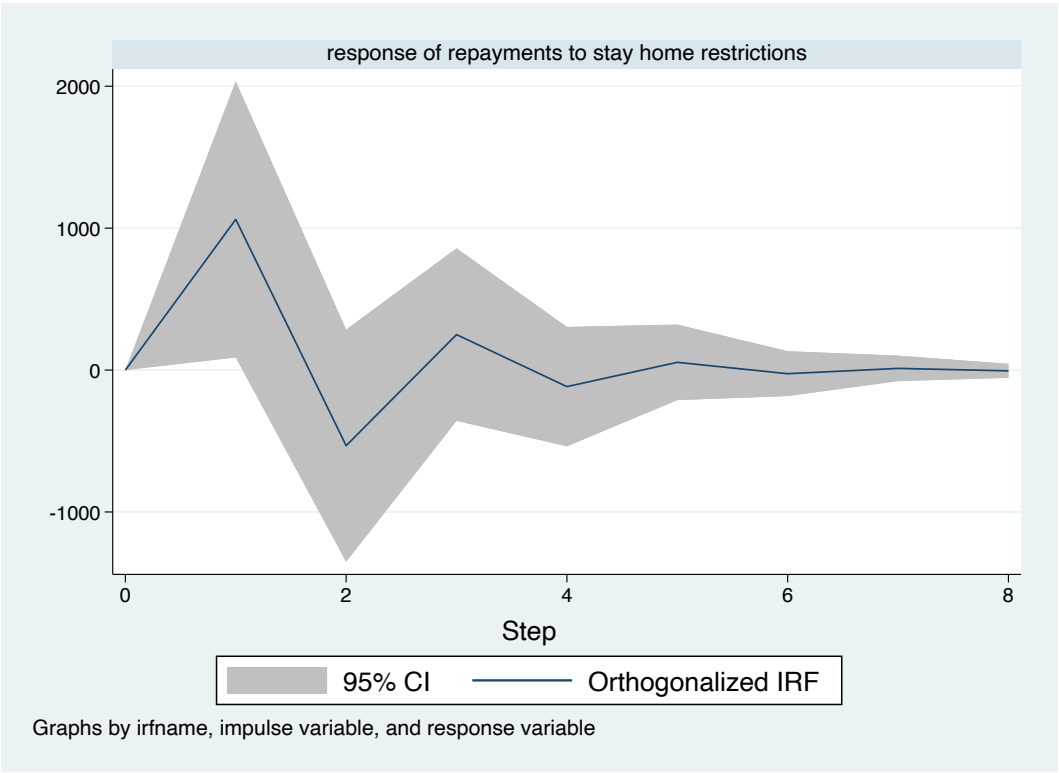

Note: Authors' estimations.

Figure 16. IRF of response of household repayment to a shock in workplace closing.

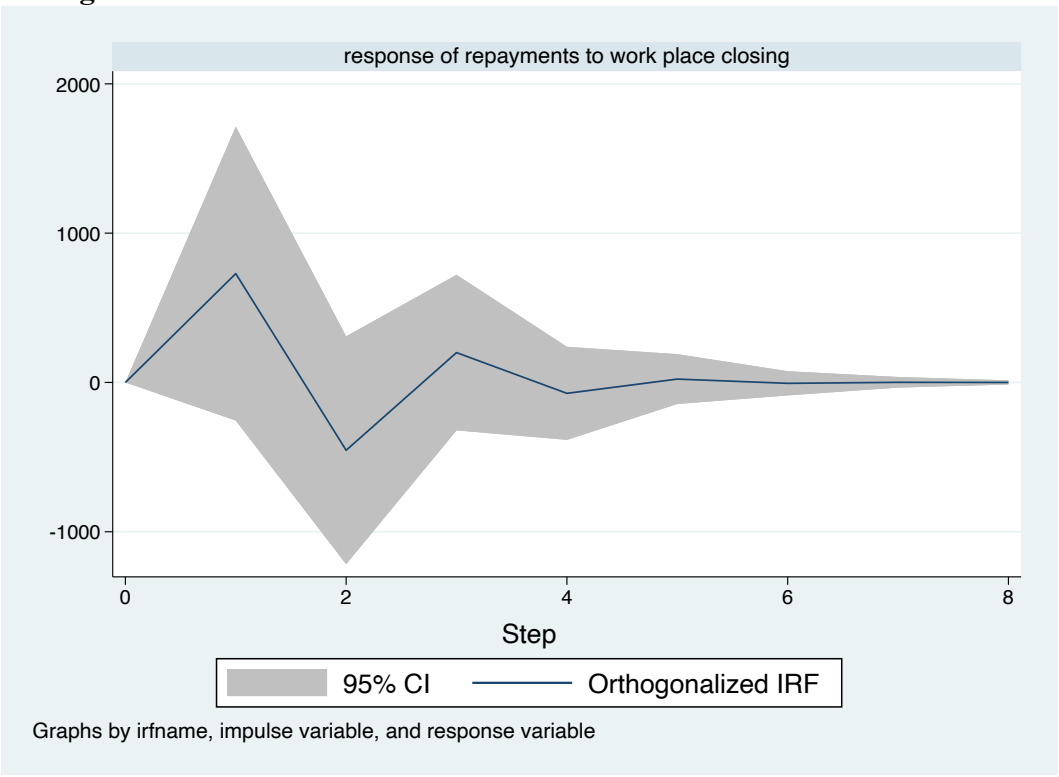

Note: Authors' estimations.

6. Concluding remarks

In this study, we have developed and implemented a time-varying parameter SIR model for COVID-19. Though heuristic and analytical learning are less feasible in a pandemic setting, aggregation of decisions over the COVID-19 emergence may drive Bayesian learning from previous priors. Our estimates of time-varying parameters can be of interest to a wider audience. We summarize our main results as follows. First, we find definite evidence that the proposed model with time-varying β_t and γ_t in the panel, VAR is better than a model with constant coefficients, conditional on the covariates (Figure 8), and with time-varying β_t and $\gamma_t$in the panel, VAR is *not* better than a random walk model conditional on the covariates (Figures 3—4). This provides some first evidence *against Bayesian learning*. Second, we find some, but in no way *definite*, evidence that the proposed with time-varying β, γ in panel VAR are better, in the light of the data, compared to a model with calibrated time-varying coefficients (Figures 5—7). This is weak evidence in favor of Bayesian learning, conditional on the covariates. The evidence is weak and therefore not decisive. Finally, from figures 1—3, β/γ less than 1 in most cases. Quantitative evidence on time-varying β, γ although no better (in a decisive way) than calibrated time-varying values implying that it is doubtful whether Bayesian (optimal) learning is taking place on the part of the authorities.

Our findings inform current discussions in policy learning during COVID-19. A more primary point of concern is the ability of policymakers to calibrate NPI responses to manage β/γ. However, we find that policymakers are unable to adapt their NPI response to flattening the curve. Though past research has highlighted that there is diffusion in policy adoption and calls for a focus on optimal adoption timing (Sears et al., 2020), our findings show that though adoption may have occurred sooner,

calibration is not present due to no support for Bayesian learning. Due to the inability to calibrate countries may have missed opportunities to fine-tune their NPI response. With changes between stringency and relaxation in NPIs, lack of Bayesian learning also implies mistiming in such policies. On a more secondary note, politicians taking credit for flattening the curve may be remiss on the fact that learning was minimal, if at all.

References

Andrieu, C., Doucet, A., & Holenstein, R. (2010). Particle Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 72*(3), 269-342.

Aravindakshan, A., Boehnke, J., Gholami, E., & Nayak, A. (2020). Preparing for a Future Covid-19 Wave: Insights and Limitations from a Data-Driven Evaluation of Non-Pharmaceutical Interventions in Germany. *Scientific reports, 10*(1), 1-14.

Athey, S., & Wager, S. (2020). Policy Learning with Observational Data. *Econometrica*.

Bo, Y., Guo, C., Lin, C., Zeng, Y., Li, H. B., Zhang, Y., . . . Kwok, K. O. (2020). Effectiveness of Non-Pharmaceutical Interventions on Covid-19 Transmission in 190 Countries from 23 January to 13 April 2020. *International Journal of Infectious Diseases, 102*, 247-253.

Chopin, N., & Singh, S. S. (2015). On Particle Gibbs Sampling. *Bernoulli, 21*(3), 1855-1883.

Cooper, I., Mondal, A., & Antonopoulos, C. G. (2020). A Sir Model Assumption for the Spread of Covid-19 in Different Communities. *Chaos, Solitons & Fractals, 139*, 110057.

Creal, D. (2012). A Survey of Sequential Monte Carlo Methods for Economics and Finance. *Econometric reviews, 31*(3), 245-296.

Creal, D. D., & Tsay, R. S. (2015). High Dimensional Dynamic Stochastic Copula Models. *Journal of Econometrics, 189*(2), 335-345.

DiCiccio, T. J., Kass, R. E., Raftery, A., & Wasserman, L. (1997). Computing Bayes Factors by Combining Simulation and Asymptotic Approximations. *Journal of the American Statistical Association, 92*(439), 903-915.

Drugowitsch, J, A.Mendonça, A. C., Z. F. Mainen, Z. F., and |A. Pouget (2019). Learning optimal decisions with confidence PNAS 116 (49), 24872-24880

Geweke, J. (1992). Evaluating the Accuracy of Sampling-Based Approaches to the Calculations of Posterior Moments. *Bayesian statistics, 4*, 641-649.

Godsill, S. J., Doucet, A., & West, M. (2004). Monte Carlo Smoothing for Nonlinear Time Series. *Journal of the American Statistical Association, 99*(465), 156-168.

Greer, S. L., King, E. J., da Fonseca, E. M., & Peralta-Santos, A. (2020). The Comparative Politics of Covid-19: The Need to Understand Government Responses. *Global public health, 15*(9), 1413-1416.

Hale, T., Webster, S., Petherick, A., Phillips, t., & Kira, B. (2020). *Oxford Covid-19 Government Response Tracker*.

Jinjarak, Y., Ahmed, R., Nair-Desai, S., Xin, W., & Aizenman, J. (2020). *Accounting for Global Covid-19 Diffusion Patterns, January-April 2020* (0898-2937).

Jaynes, E. T. (2003). Probability theory: the logic of science. Cambridge university press

Lewis, S. M., & Raftery, A. E. (1997). Estimating Bayes Factors Via Posterior Simulation with the Laplace—Metropolis Estimator. *Journal of the American Statistical Association, 92*(438), 648-655.

May, P. J. (1992). Policy Learning and Failure. *Journal of public policy, 12*(4), 331-354.

Okasha, S., 2013, The evolution of Bayesian updating, Philosophy of Science 80 (5), 745-757.

Sabatier, P. A. (1988). An Advocacy Coalition Framework of Policy Change and the Role of Policy-Oriented Learning Therein. *Policy sciences, 21*(2-3), 129-168.
Schaback, R. (2020). Modelling Recovered Cases and Death Probabilities for the Covid-19 Outbreak. *arXiv preprint arXiv:2003.12068*.
Sears, J., Villas-Boas, J. M., Villas-Boas, V., & Villas-Boas, S. B. (2020). Are We Stayinghome to Flatten the Curve? *Department of Agricultural and Resource Economics, first version April, 5*.
Sebhatu, A., Wennberg, K., Arora-Jonsson, S., & Lindberg, S. I. (2020). Explaining the Homogeneous Diffusion of Covid-19 Nonpharmaceutical Interventions across Heterogeneous Countries. *Proceedings of the National Academy of Sciences, 117*(35), 21201-21208.
Tauber, S., D. J. Navarro, A. Perfors, and M. Seyvers (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. Psychological Review 124 (4), 410-441.
Weible, C. M., Pattison, A., & Sabatier, P. A. (2010). Harnessing Expert-Based Information for Learning and the Sustainable Management of Complex Socio-Ecological Systems. *environmental science & policy, 13*(6), 522-534.
Whiteley, N., Andrieu, C., & Doucet, A. (2010). Efficient Bayesian Inference for Switching State-Space Models Using Discrete Particle Markov Chain Monte Carlo Methods. *arXiv preprint arXiv:1011.2437*.
Witting, A. (2017). Insights from 'Policy Learning'on How to Enhance the Use of Evidence by Policymakers. *Palgrave Communications, 3*(1), 1-9.
Zellner, A. (1971). *An Introduction to Bayesian Inference in Econometrics*. Retrieved from

Technical Appendix. MCMC and Particle filtering

We use a recent advance in sequential Monte Carlo methods known as the particle Gibbs (PG) sampler, see Andrieu et al. (2010). The algorithm allows us to draw paths of the state variables in large blocks. Particle filtering is a simulation-based algorithm that sequentially approximates continuous, marginal distributions using discrete distributions. This is performed by using a set of support points called ‘‘particles’’ and probability masses; see (D. Creal, 2012) for a review.

The PG sampler draws a single path of the latent or state variables from this discrete approximation. As the number of particles M goes to infinity, the PG sampler draws from the exact full conditional distribution. As mentioned in (Creal and Tsay, 2015): “The PG sampler is a standard Gibbs sampler but defined on an extended probability space that includes all the random variables that are generated by a particle filter. Implementation of the PG sampler is different than a standard particle filter due to the ‘‘conditional’’ resampling algorithm used in the last step. Specifically, for draws from the particle filter to be a valid Markov transition kernel on the extended probability space, Andrieu et al. (2010) note that there must be a positive probability of sampling the existing path of the state variables that were drawn at the previous iteration. The pre-existing path must survive the resampling steps of the particle filter. The conditional resampling step within the algorithm forces this path to be resampled at least once. We use the conditional multinomial resampling algorithm from Andrieu et al. (2010), although other resampling algorithms exist, see Chopin and Singh (2015)” (page 339).

We follow D. D. Creal and Tsay (2015). Suppose the posterior is $p(\theta, \Lambda_{1:T}|\boldsymbol{y}_{1:T})$ where $\Lambda_{1:T}$ denotes the latent variables whose prior can be described by $p(\Lambda_t|\Lambda_{t-1}, \theta)$. In the PG sampler we can draw the structural parameters

$\theta|\Lambda_{1:T}, \boldsymbol{y}_{1:T}$ as usual, from their posterior conditional distributions. This is important because, in this way, we can avoid mixture approximations or other Monte Carlo procedures that need considerable tuning and may not have good convergence properties. As such posterior conditional distributions, we omit the details and focus on drawing the latent variables.

Suppose we have $\Lambda_{1:T}^{(1)}$ from the previous iteration. The particle filtering procedure consists of two phases.

Phase I: Forward filtering (Andrieu et al., 2010).

- Draw a proposal $\Lambda_{i,t}^{(m)}$ from an importance density $q(\Lambda_{i,t}|\Lambda_{i,t-1}^{(m)}, \theta), m = 2, \dots, M$.
- Compute the importance weights:

$$w_{i,t}^{(m)} = \frac{p(y_{i,t}; \Lambda_{i,t}^{(m)}, \theta) p(\Lambda_{i,t}^{(m)}|\Lambda_{i,t-1}^{(m)}, \theta)}{q(\Lambda_{i,t}|\Lambda_{i,t-1}^{(m)}, \theta)}, m = 1, \dots, M. \quad \text{(A.1)}$$

- Normalize the weights: $\widetilde{w}_{i,t}^{(m)} = \frac{w_{it}^{(m)}}{\Sigma_{m'=1}^{M} w_{it}^{(m')}}, m = 1, \dots, M$.
- Resample the particles $\{\Lambda_{i,t}^{(m)}, m = 1, \dots, M\}$ with probabilities $\{\widetilde{w}_{i,t}^{(m)}, m = 1, \dots, M\}$.

In the original PG sampler, the particles are stored for $t = 1, \dots, T$ and a single trajectory is sampled using the probabilities from the last iteration. An improvement upon the original PG sampler was proposed by Whiteley et al. (2010), who suggested drawing the path of the latent variables from the particle approximation using the backwards sampling algorithm of Godsill et al. (2004). In the forwards pass, we store the normalized weights and particles and we draw a path of the latent variables as we

detail below (the draws are from a discrete distribution).

Phase II: Backward filtering (Chopin & Singh, 2015; Godsill et al., 2004).

- At time $t = T$ draw a particle $\Lambda_{i,T}^{*} = \Lambda_{i,T}^{(m)}$.
- Compute the backward weights: $w_{t|T}^{(m)} \propto \widetilde{w}_{t}^{(m)} p(\Lambda_{i,t+1}^{*} | \Lambda_{i,t}^{(m)}, \theta)$.
- Normalize the weights: $\widetilde{w}_{t|T}^{(m)} = \frac{w_{t|T}^{(m)}}{\sum_{m'=1}^{M} w_{t|T}^{(m')}}, m = 1, \dots, M$.
- Draw a particle $\Lambda_{i,t}^{*} = \Lambda_{i,t}^{(m)}$ with probability $\widetilde{w}_{t|T}^{(m)}$.

Therefore, $\Lambda_{i,1:T}^{*} = \{\Lambda_{i1}^{*}, \dots, \Lambda_{iT}^{*}\}$ is a draw from the full conditional distribution. The backwards step often results in dramatic improvements in computational efficiency. For example, Creal and Tsay (2015) find that $M = 100$ particles are sufficient There remains the problem of selecting an importance density $q(\Lambda_{i,t} | \Lambda_{i,t-1}, \theta)$. We use an importance density implicitly defined by $\Lambda_{i,t} = a_{i,t} + \sum_{p=1}^{P} b_{i,t} \Lambda_{i,t-1}^{p} + h_{i,t} \xi_{i,t}$ where $\xi_{i,t}$ follows a standard (zero location and unit scale) Student-t distribution with $\nu = 5$ degrees of freedom. That is, we use polynomials in $\Lambda_{i,t-1}$ of order P. We select the parameters $a_{i,t}, b_{i,t}$ and $h_{i,t}$ during the burn-in phase (using $P = 1$ and $P = 2$) so that the weights $\{\widetilde{w}_{i,t}^{(m)}, m = 1, \dots, M\}$ and $\{\widetilde{w}_{t|T}^{(m)}, m = 1, \dots, M\}$ are approximately not too far from a uniform distribution.

Chopin and Singh (2015) have analyzed the theoretical properties of the PG sampler and proved that the sampler is uniformly ergodic. They also prove that the PG sampler with backwards sampling strictly dominates the original PG sampler in terms of asymptotic efficiency.

Alternatively, when the dimension of the state vector is large, we can draw

$\Lambda_{i,1:T}$, conditional on all other paths $\Lambda_{-i,1:T}$ that are not path i. Therefore, we can draw from the full conditional distribution $p(\Lambda_{i,1:T}|\Lambda_{-i,1:T}, \boldsymbol{y}_{1:T}, \theta)$.

Implementation and recursive Bayes factors

Our implementation relies on 150,000 MCMC iteration with a burn-in length of 50,000 to mitigate possible start up effects, and we use 1,000 particles per MCMC iteration. The marginal likelihood is a direct by-product of the SMC algorithm so, recursive Bayes factors, which are ratios of marginal likelihoods, are easy to compute. The convergence of MCMC is tested successfully using the standard diagnostics of Geweke (1992).

To compute the Bayes factor in favor of (18) and (19), and against the Bayesian panel data time-varying parameters model, we plug in (7) and (8) the estimates from (18) and (19) into (7) – (9). We still estimate the covariance matrix $\boldsymbol{\Sigma}$ by Bayesian methods so that we can compute the marginal likelihood of this model easily using the Laplace approximation (DiCiccio et al., 1997; Lewis & Raftery, 1997). As the marginal likelihood of the Bayesian model is a by-product of SMC the two can be compared to obtain the Bayes factors. On the (DiCiccio et al., 1997) and Lewis and Raftery (1997)approximation, we proceed as follows: Given a likelihood function $L(\theta; Y)$ that depends on parameters $\theta \in \Theta \subseteq \mathbb{R}^d$ and data Y, a prior $p(\theta)$ and a posterior given by Bayes' theorem $p(\theta|Y) \propto L(\theta; Y)p(\theta)$ the marginal likelihood or evidence is a standard way for model selection and model comparison in a Bayesian framework. The marginal likelihood is $M(Y) = \int_\Theta L(\theta; Y)p(\theta)\mathrm{d}\theta$, viz. the integrating constant of the posterior: $p(\theta|Y) = \frac{L(\theta;Y)p(\theta)}{\int_\Theta L(\theta';Y)p(\theta')d\theta'}$. The marginal likelihood can be approximated using the identity (for all θ): $M(Y) = \frac{L(\theta;Y)p(\theta)}{p(\theta|Y)}$. DiCiccio et al. (1997) propose to

approximate the denominator with a normal distribution around the posterior mean, $\bar{\theta}$, yielding

$$M(Y) = \frac{L(\bar{\theta};Y)p(\bar{\theta})}{p(\bar{\theta}|Y)} = L(\bar{\theta};Y)p(\bar{\theta})(2\pi)^{d/2}|\bar{V}|^{1/2}, \quad \text{(A.2)}$$

where $\bar{V}$ is the posterior covariance matrix of θ. Both $\bar{\theta}$ and $\bar{V}$ can be estimated easily using MCMC output.

Supplementary Information

Supplementary information A: Variable descriptives
Supplementary information B: Effect of covariates
Supplementary information C: Additional model hypotheses testing

Supplementary information A: Variable descriptive

Table S1. Descriptive statistics (N = 41,706 country-date observations)

	mean	sd	min	max
Containment and closure policies				
School closing	2.0944	1.0303	0	3
workplace closing	1.5608	0.9575	0	3
Cancelled public events	1.5505	0.7236	0	2
Restrictions on gathering	2.7339	1.4338	0	4
Closed public transport	0.6736	0.7598	0	2
Stay at home requirements	1.1250	0.9331	0	3
Restrictions on internal movement	1.0404	0.9053	0	2
International travel controls	2.8152	1.1213	0	4
Economic policies				
Income support	0.9434	0.7694	0	2
Debt contract relief	1.1162	0.8224	0	2
Fiscal measures	188,000,000	9,940,000,000	-0.01	1,190,000,000,000
International support	20,700,000	4,090,000,000	0	834,000,000,000
Health system policies				
Public information campaigns	1.9041	0.3543	0	2
Testing Policy	1.7896	0.8165	0	3
Contact tracing	1.4810	0.6526	0	2
Emergency investment in health care	5008834.0000	350000000.0000	0	63000000000
Investment in vaccines	548055.8000	44500000.0000	0	7860000000
Facial coverings	2.1056	1.4305	0	4
Indices based on actions				
Stringency index	59.4632	22.6222	0	100
Government response index	54.4140	16.9675	0	89.17
Government response index for display	54.4140	16.9675	0	89.17
Containment health index	55.4277	17.4137	0	91.35
Containment health index for display	55.4277	17.4137	0	91.35
Economic support index	47.8253	31.1172	0	100
Economic support index for display	47.8253	31.1172	0	100
Controls				
GDP (constant 2010 US dollars)	451,000,000,000	1,300,000,000,000	1,170,000,000	11,500,000,000,000
Population density	211.0070	734.0610	1.9800	7915.7310
Median age	31.5465	8.8935	15.1000	48.2000
Age 65 and older	9.3107	6.3886	1.1440	27.0490
Age 70 and older	5.9627	4.4289	0.5260	18.4930
GDP per capita (constant 2010 US dollars)	20833.2500	20628.4500	661.2400	116935.6000
Cardiovascular death rate	253.6085	122.4997	79.3700	724.4170
Diabetes prevalence	7.7964	3.8890	0.9900	22.0200
Hospital beds per thousand	3.0014	2.4750	0.1000	13.0500
Life expectancy	73.8849	6.7793	53.2800	84.6300
Human development index	0.7323	0.1476	0.3540	0.9530

Supplementary information B: Effect of covariates

Table S2. Effect of covariates

Covariate	β_t	γ_t
Containment and closure policies		
School closing	-0.014	0.001
	(0.0034)	(0.0030)
workplace closing	-0.020	0.002
	(0.0012)	(0.0040)
Cancelled public events	-0.015	0.001
	(0.0040)	(0.0020)
Restrictions on gathering	-0.023	-0.002
	(0.0017)	(0.0010)
Closed public transport	-0.005	-0.001
	(0.0013)	(0.0003)
Stay at home requirements	-0.004	-0.001
	(0.0012)	(0.0001)
Restrictions on internal movement	-0.032	-0.005
	(0.0040)	(0.0020)
International travel controls	-0.036	0.000
	(0.0060)	(0.0001)
Economic policies		
Income support	-0.015	-0.001
	(0.0050)	(0.0007)
Debt contract relief	-0.003	-0.002
	(0.0020)	(0.0030)
Fiscal measures	-0.006	-0.005
	(0.0011)	(0.0016)
International support	-0.003	-0.001
	(0.0012)	(0.0006)
Health system policies		
Public information campaigns	-0.002	-0.001
	(0.0005)	(0.0004)
Testing Policy	-0.032	-0.005
	(0.0050)	(0.0010)
Contact tracing	-0.004	-0.001
	(0.0100)	(0.0004)
Emergency investment in health care	-0.035	-0.005
	(0.0040)	(0.0005)
Investment in vaccines	-0.004	-0.001
	(0.0010)	(0.0012)
Facial coverings	-0.032	-0.004
	(0.0120)	(0.0013)
Indices based on actions		
Stringency index	-0.005	-0.001
	(0.0013)	(0.0004)
Government response index	-0.017	-0.002
	(0.0200)	(0.0011)
Government response index for display	-0.003	-0.001
	(0.0013)	(0.0010)
Containment health index	-0.003	0.000
	(0.0010)	(0.0014)
Containment health index for display	-0.005	0.000
	(0.0013)	(0.0002)
Economic support index	-0.002	-0.001
	(0.0003)	(0.0007)
Economic support index for display	-0.004	-0.002
	(0.0025)	(0.0020)

Controls		
GDP (constant 2010 US dollars)	-0.003	-0.002
	(0.0028)	(0.0020)
Population density	0.004	0.000
	(0.0014)	(0.0002)
Median age	0.005	0.000
	(0.0100)	(0.0030)
Age 65 and older	0.003	0.000
	(0.0014)	(0.0002)
Age 70 and older	0.003	0.000
	(0.0005)	(0.0002)
GDP per capita (constant 2010 US dollars)	0.005	0.004
	(0.0040)	(0.0070)
Cardiovascular death rate	0.002	0.001
	(0.0010)	(0.0010)
Diabetes prevalence	0.004	0.001
	(0.001)	(0.0020)
Hospital beds per thousand	-0.005	-0.001
	(0.0012)	(0.0020)
Life expectancy	0.005	-0.003
	(0.0012)	(0.0030)
Human development index	-0.003	-0.001
	(0.0013)	(0.0020)

Supplementary information C: Additional model hypotheses testing

Hypothesis	Bayes factor
H: covariates jointly significant	$14.28\ 10^{13}$
H: time-invariant SIR with covariates	$2.59\ 10^{-4}$
H: time-invariant SIR without covariates	$3.52\ 10^{-7}$
H: Policy instruments lagged	$4.59\ 10^{-4}$
H: second-order panel VAR	$3.81\ 10^{-5}$
H: Omit cross-sectionally different parameters	$4.40\ 10^{-12}$
H: omit cross-sectional different parameter in panel VAR without covariates	$5.81\ 10^{-9}$
H: Random walk without covariates	11.212
H: Break (change of parameters in the middle of the sample)	$2.33\ 10^{-6}$

Note. Reported are Bayes factors in favor of the various hypotheses H. Bayes factors above 100, are considered as providing "decisive evidence" in favor of a hypothesis.