NPI Models Explained and Complained

***Abstract*—Numerous modelling efforts have attempted to characterize the effects of different non-pharmaceutical interventions (NPIs) on the Covid-19 spread. Arguably the most famous is one published in Nature by an Imperial College group. A slight variation of it was later published in Science by a group of Oxford researchers. Both publications are based on hierarchical Bayesian modelling that aims to explain observed data by information on enacted NPIs. Due to the Bayesian approach, the models become quite complex and opaque, with many priors that have been assigned more or less ad hoc, and there are even priors on the prior parameters. We show how these models can be recast into the classic linear regression framework. This enables us to transparently analyze basic concepts such as persistency of excitation, identifiability, and model sensitivity.**

Fredrik Gustafsson
Department of Electrical Engineering
Linkoping University
Linköping, Sweden
fredrik.gustafsson@liu.se

Kristian Soltesz
Department of Automatic Control
Lund University
Lund, Sweden
kristian@control.lth.se

THE SIR MODEL REVISITED

We will refer to the two studied non-pharmaceutical intervention (NPI) models as the *Nature* [1] and *Science* [2] model, respectively. In the presentation we focus on the former, although our methodology remains applicable to either and we will present results obtained using both models.

Within the models, NPIs are typed as school closure, crowd size limit, lockdown, etc. The purpose of the model is then to explain the epidemic trajectory based on enactment of the NPIs.

Before delving into the details of the models, let us briefly revisit the classic Susceptible, Infected, and Recovered (SIR) compartment model [3] that lies at the core of many more advanced epidemiological models, including the ones considered here. It is a lumped-parameter model that can be applied on a societal or subsocietal level and describes how a considered population is partitioned into susceptible S, infectious I, and removed (recovered and immune $\cup$ deceased) R fractions.[1] The population is normalized according to (1a), and the dynamics are:

$$1 = S + I + R, \tag{1a}$$

$$\frac{dS}{dt} = -\beta SI, \tag{1b}$$

$$\frac{dI}{dt} = \beta SI - \gamma I, \tag{1c}$$

$$\frac{dR}{dt} = \gamma I. \tag{1d}$$

The equations governing the epidemic trajectory are determined by an infection parameter β and a recovery parameter γ.

The famous basic reproduction number $R_0 = \beta/\gamma$ defines how many secondary infections are expected from one primary infection[2], when $S \gg I$. The adjective *basic* is with respect to some considered action, such as an NPI or set of NPIs. In contrast to the growth *rate*[3] $r_0 = \beta - \gamma$, R_0 is unit-less and decoupled from time (making it a less obvious choice for measuring time-dependent growth in the first place).

Arguably the simplest way to model NPI effectiveness is to investigate how enacting an NPI affects the spread parameter β, which coincides with how it affects R_0 (or r_0) if γ is constant. Since β cannot be directly measured (either), an observation model involving a measurable signal related to it is needed. A very simple such observation model would be to assume a constant infection-to-fatality ratio (IFR), and that all deaths occur a fixed time τ_{ID} following infection.

Figure 1 shows the reported daily deaths in Sweden [4] during the first wave of the pandemic in green. The red curve is fitted using nonlinear least squares (NLS), under the assumption that β was constant throughout the first wave. The blue curve is the fit that minimizes the quadratic NLS loss using one NPI (change in β), with date chosen to optimize curve fit.

Although the results look convincing, the model is useless in all aspects other than for curve fitting.

One artefact is that the parameters generating the red curve explain the decline of the first wave as a consequence of herd immunity, with 99.9% of the population having been infected by October 2020, while the blue curve mainly explains the decline in deaths through an effective NPI enacted on April 7,

[1] The basic SIR model does not distinguish between infected and infectious, but such additional state partitioning is straightforward, as are the partitionings aimed at tracking different subpopulations.

[2] A simple derivation is based on $R_0 = -dS/dR = \beta S/\gamma \approx \beta/\gamma$ when $S \approx 1$, so every removed person gives rise to R_0 new infections.

[3] When $S \approx 1$ the solution to (1c) is $I(t) \approx e^{(\beta-\gamma)(t-t_0)} I(t_0)$.

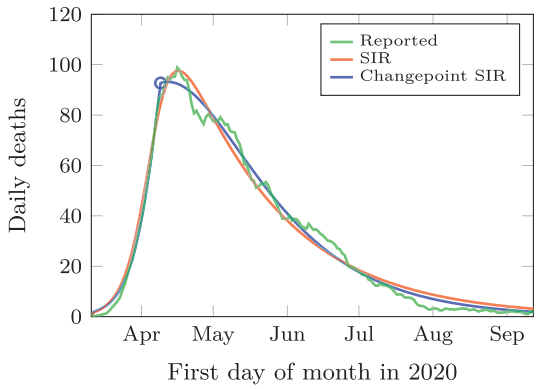

Figure 1
NLS fit of the SIR model (1) to officially reported daily deaths in Sweden during the first pandemic wave. Deaths are shown in green, the SIR model fit in red, and the best NLS fit of an SIR model, where β is allowed to change twice, is shown in blue. The marker indicates the instance of the parameter change.

2020, with 50% of the population having undergone infection by October 2020.

An important conclusion from this example is that the model fit cannot alone be used to validate a model—not even for the very simple model considered above. The lack of *informative* data for model validation is a fundamental problem, in particular during the early stage of an (infectious disease) epidemic.

MODELLING FRAMEWORK

The basic dynamics of how NPIs are modeled to affect reported death (or case) data in the *Nature* and *Science* models are illustrated in Figure 2. Both models share the same basic equations as will be outlined in this section. The differences are in the details of NPI types, time span, and the strategies to estimate parameters to fit the data.

SOCIETAL MODEL

NPI k enacted on day t in country c is modeled to induce an undelayed step change of magnitude ($\alpha_{k,c}$) in the reproduction number within the country:

$$R_c(t) = R_{0,c} \cdot \exp\left(-\sum_{k=1}^{N_\sigma} \alpha_{k,c} \sigma_{k,c}(t) \right), \tag{2}$$

where σ is the binary indicator function, and N_σ is the number of NPI types. The country-specific basic reproduction number $R_{0,c}$ is treated as an unknown parameter to be identified from data together with the effectiveness parameters $\alpha_{k,c}$.

If $\alpha_{k,c} = \alpha_k$ is the same for all countries $c = 1, \ldots, N_c$, the model for each NPI $k = 1, \ldots, N$... is referred to as *fully pooled*. If the effectiveness parameters are allowed full international flexibility, the model is said to be *unpooled*. Models between these extremes are referred to as *partially pooled*. The main differences between the *Nature* and *Science* models reside with the definition of the NPI types, and the pooling assumptions on individual NPIs.

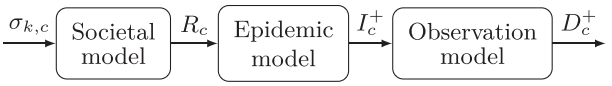

Figure 2
Block diagram of the considered NPI models showing their three principal components and intermediate signals: Enactment indicator of NPI k, σ_k, reproduction number R, daily infections I^+, reported daily deaths D^+. Subscript c denotes country.

EPIDEMIC MODEL

Two time-distributions play a central role in the NPI effectiveness modelling. They model the duration τ between

- primary and secondary infection following the serial (or generation) interval distribution $p_{II}(\tau)$;
- infection and death following the distribution $p_{ID}(\tau)$.

These distributions are assumed to be time-invariant, and prior assumptions on their shape within the *Nature* model are reviewed further below in the section "Priors".

The number of individuals $I_c^+(t)$ that become infectious on day t in country c can now be expressed as

$$I_c^+(t) = \underbrace{R_{0,c} \frac{N_c - \sum_{\tau=0}^{t-1} I_c^+(\tau)}{N_c}}_{R_{e,c}(t)} \sum_{\tau=0}^{t-1} I_c^+(\tau) p_{II}(t-\tau). \tag{3}$$

Here, the effective reproduction number $R_{e,c}$ accounts for both NPI and herd immunity effects. Equation (3) constitutes a nonlinear auto-regressive model since previous values of I_c^+ are combined in a nonlinear fashion to determine $I_c^+(t)$.

To get some intuition for (3), consider the special case where $p_{II}(\tau) = \delta(\tau - k)$ is a Dirac distribution. That is, each infectious individual spreads the disease to (on average) $R_c(t) I_c^+(t) / N_c$ susceptible individuals, k days after becoming infected. If R_c is constant, the spread can be locally approximated with

$$I_c^+(t) \approx R_c^{(t-t_0)/k} I_c^+(t_0) / N_c, \tag{4}$$

where the exponential growth rate $r_c = \log(R_c)/k$ clearly shows the strong influence of the delay k.

OBSERVATION MODEL

Finally, we have an observation model for reported daily deaths $D_c^+(t)$ in country c on day t, based on an assumed distribution $p_{ID}(\tau)$ of the time between infection and death:

$$D_c^+(t) = \text{IFR} \sum_{\tau=0}^{t-1} I_c^+(\tau) p_{ID}(t-\tau), \tag{5}$$

If case data is incorporated into the model, the infection-to-case ratio plays an analogous role to the IFR. More generally, each of the distributions of the epidemic model in the section "Epidemic Model" is associated with a normalization factor of this kind, since not all primary infections result in secondary infections, not all infected individuals die, etc.

MODEL ESTIMATION

The model is defined through (2)–(5). How can it be used to estimate the parameters $R_{0,c}$ and $\alpha_{k,c}$ for $c = 1, \ldots, N_c$ and $k = 1, \ldots, N_\alpha$ from D_c^+ time series data?

BAYESIAN APPROACH

Both the *Nature* and *Science* estimation methods are so-called hierarchical Bayesian models that are fitted using massive Monte Carlo sampling [5]. Our contribution is to recast them into the linear regression framework rather than to analyze the estimation method they originally rely on. However, we believe it is adequate to summarize the main design parameters of the Bayesian approach.

PRIORS

Prior assumptions on distributions are needed, and in the final published version, the *Nature* model assumes the following priors[4] for the time distributions of the epidemic model reviewed in the section "Epidemic Model":

$$p_{II}(\tau) \sim \Gamma(6.5, 0.62), \tag{6a}$$

$$p_{ID}(\tau) \sim \Gamma(5.1, 0.86) + \Gamma(17.8, 0.45). \tag{6b}$$

Here, $\Gamma(a, b)$ denotes the Gamma distribution with mean a, coefficient of variation b, and standard deviation ab.

The effectiveness prior for NPI type k in the *Nature* model is

$$\alpha_k + \frac{\log(1.05)}{6} \sim \Gamma(1/6, 1), k = 1, \ldots, 6, \tag{7}$$

with the motivation that

$$\sum_{k=1}^{6} \alpha_k \sim U(0, 1.05). \tag{8}$$

That is, there is a possibility that the interventions will increase the reproduction number by a factor 1.05, but most of the prior is assigned to a significant decrease.

Finally, the prior for the basic reproduction numbers were chosen $R_{0,c} \sim \mathcal{N}\left(2.4, |\kappa|\right)$, $c = 1, \ldots, N_c$ where $\kappa \sim \mathcal{N}(0, 0.5)$.

REGULARIZING EFFECT OF THE GAMMA PRIOR

As was illustrated by Nic Lewis in his blog [6], the Gamma prior has a regularization effect. Suppose that a pooled model ($\alpha_{k,c} = \alpha_k$) is used and that the data is consistent with a posterior where $\alpha_1 + \ldots + \alpha_{N_\alpha} = 1.75$. This corresponds to a factor $1 - e^{1.75} = 0.83$ decrease of the reproduction number should all NPIs be simultaneously enacted, and it happens to coincide with the median between Markov chain Monte Carlo (MCMC) samples reported in [1]. Then, the prior for the sum becomes orders of magnitude larger if one α_k dominates (e.g., $\alpha_1 = 1.70$), compared to when they all are of similar size (e.g., $\alpha_k = 1.75/N_\alpha$, $k = 1, \ldots, N_\alpha$). This makes the Gamma prior strongly biased towards one (or a few) NPIs explaining the data, and in the case of [1] the lockdown NPI was singled out as the by far dominating explanation of the decrease in viral reproduction. Figure 3 illustrates the effect for the visualizable case of $N_\alpha = 2$.

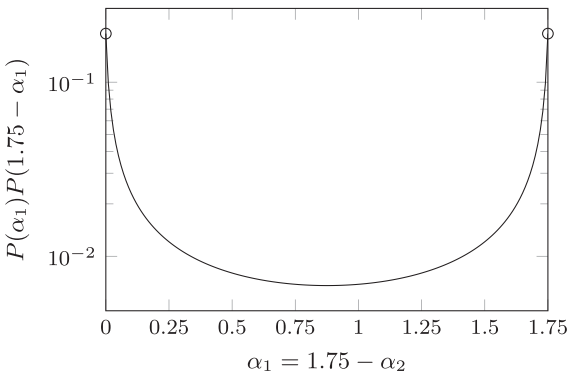

Figure 3
Probability for $\alpha_1 + \alpha_2$ for different splits between priors α_1 and $\alpha_2 = 1.75 - \alpha_1$, individually distributed according to (7).

MCMC SAMPLING PRINCIPLE

To compute the (posterior distributions of the) parameters, the *Nature* and *Science* models rely on sampling from the parameter priors, evolving (3) and (2), and then evaluating the likelihood for each sample in a Bayesian MCMC framework, as very sketchily summarized by the following basic steps:

1. Draw a candidate for the parameters from the prior distributions. This is the most important step and there are many different sampling strategies that could be considered.
2. Simulate the model with the parameter candidate.
3. Compute the likelihood for the observed mortality data.
4. Generate a random number $u \sim U(0, 1)$.
5. If the log-likelihood ratio has increased more than u, the parameter candidate is accepted, otherwise it is rejected.
6. Continue until a predefined number of parameter candidates have been accepted, excluding the burn-in phase before the MCMC has converged to stationarity.

LINEAR REGRESSION NPI MODEL

SOCIETAL MODEL AS A LINEAR REGRESSION

We note that (2) can be cast as a linear regression in the log domain:

$$\log R_c(t) = \log R_{0,c} - \sum_{k=1}^{N_\alpha} \alpha_{k,c}\ \sigma_{k,c}(t). \tag{9}$$

Taking the logarithm of (3) and using it to eliminate $R_c(t)$ yields

[4] The reason that p_{ID} is the sum of two Gamma-distributed variables is that the *Nature* model breaks it down into the infection-to-symptom-onset distribution p_{IO} (first term), that is convoluted with the symptom-onset-to-death distribution p_{OD} (second term).

$$\log I_c^+(t) = \log R_{0,c} - \sum_{k=1}^{N_\alpha} \alpha_{k,c}\sigma_{k,c}(t) + \log\left(1 - \frac{\sum_{\tau=0}^{t-1} I_c^+(\tau)}{N_c}\right) + \log\left(\sum_{\tau=0}^{t-1} I_c^+(\tau) p_{II}(t-\tau)\right). \quad (10)$$

This fits the classical linear regression framework

$$z_c(t) = h_c(t)\theta + w(t). \quad (11a)$$

The left-hand-side is the auto-regressive process

$$z_c(t) = \log I_c^+(t) - \log\left(\sum_{\tau=0}^{t-1} I_c^+(\tau) p_{II}(t-\tau)\right) \quad (11b)$$

$$+\log\left(1 - \frac{\sum_{\tau=0}^{t-1} I_c^+(\tau)}{N_c}\right). \quad (11c)$$

The right-hand side of (11a) comprises of the regression matrix and parameter vector,

$$h_c(t) = \left[\sigma_{1,c}(t)\ldots\ \sigma_{N_\alpha,c}(t)\quad e_c^\top\right], \quad (11d)$$

$$\theta = \left[\alpha_1\ldots\alpha_{N_\alpha}\quad \log(R_{0,1})\ldots\log(R_{0,N_c})\right]^\top, \quad (11e)$$

where e_c is the cth unit vector. Equation (11e) corresponds to the fully pooled model, and partially pooled or unpooled formulations only differ structurally in that they have a larger number of parameters.

Vectorizing the data for the N_c countries, we obtain the more compact form

$$Z(t) = H(t)\theta + W(t), \quad (12)$$

where $Z(t) = \left[z_1(t)\ldots z_{N_c}(t)\right]^\top$ is a function of mortality data and $H(t) = \left[h_1^\top(t)\ldots h_{N_c}^\top(t)\right]^\top$ depends on the NPIs only. The model error $W(t) = \left[w_1(t)\ldots w_{N_c}(t)\right]^\top$ is the realization of an observation noise process, and its variance λ can be interpreted as the best model fit for this particular model structure.

LEAST SQUARES SOLUTION

Assuming that W in (12) consists of independent and identically distributed samples from a Gaussian process, the maximum-likelihood (ML) estimate $\hat{\theta}$ of θ is obtained by minimizing the quadratic ordinary least-squares (OLS) loss, with closed-form solution

$$J = \sum_t H(t)^\top H(t) \quad (13a)$$

$$\hat{\theta} = J^{-1}\sum_t H(t)^\top Z(t), \quad (13b)$$

$$\mathrm{Cov}\left(\hat{\theta}\right) = \lambda J^{-1}. \quad (13c)$$

Here, $\lambda = \mathrm{Var}\left(z_{t,c}\right)$ denotes the variance of the transformed data, assuming it to be the same for all times and countries, and J is the Fisher information matrix (FIM). Structural identifiability is determined by the rank of the FIM, and persistency of excitation is measured by its condition number.

In the *Nature* and *Science* models, observation noise was not introduced as in (12), but instead implicitly generated through the random variables assigned with priors according to the section "Priors". To understand how the observation noise (or LS residual) W relates to the stochastics of the original model, assume we have access to the posteriors that maximize the likelihood within the Bayesian formulation. Then (in theory) I_c^+ can be computed through deconvolution of (5) and used to construct the X and Z of (12). Having access also to the posteriors of θ, $W(t) = Z(t) - H(t)\,\theta$ can be evaluated. It is thus important to note that we cannot (at least directly) use the linear regression formulation for identifying θ. However, and importantly, we can use it to analyze sensitivity of the identified system to small perturbations, that could arise from uncertainty or error in NPI actuation date (σ) or death (D^+) reporting.

IDENTIFIABILITY OF THE NATURE MODEL

The *Nature* model defined the following $N_\alpha = 5$ NPI types[5]:

1. Social distancing encouraged;
2. Self isolation;
3. School closure;
4. Public events (banned);
5. Lockdown.

Figure 4 (top) shows the dates that the different NPIs were enacted within the published version [1] of the *Nature* model. Countries are labeled by their ISO 3166-1 alpha-2 codes. Note in particular that all NPIs were enacted within a short time window. The corresponding reported daily death time series D^+ are also shown in the same figure.

The inverse FIM for the model of the section "Linear Regression NPI Model" has the block structure

$$J^{-1} = \begin{bmatrix} J_{\alpha\alpha}^{-1} & J_{\alpha R_0}^{-1} \\ [2pt] J_{R_0\alpha}^{-1} & J_{R_0R_0}^{-1} \end{bmatrix}. \quad (14)$$

Let us here focus on the block $J_{\alpha\alpha}^{-1} \propto \mathrm{Cov}(\alpha)$ that defines the covariance of the NPI effectiveness parameters.

For the data used in [1] and shown in Figure 4, it evaluates to

$$J_{\alpha\alpha}^{-1} = \lambda \begin{bmatrix} 0.044 & -0.019 & -0.015 & -0.007 & -0.005 \\ -0.019 & 0.034 & -0.003 & -0.004 & -0.002 \\ -0.015 & -0.003 & 0.045 & -0.009 & -0.013 \\ -0.007 & -0.004 & -0.009 & 0.027 & -0.006 \\ -0.005 & -0.002 & -0.013 & -0.006 & 0.025 \end{bmatrix}. \quad (15)$$

[5] The careful reader might have noticed that (7)–(8) correspond to $N_\alpha = 6$. This is related to how (a particular version of) the *Nature* model pools data, with a "bonus" NPI for the last intervention introduced in each country, and further explained in [7].

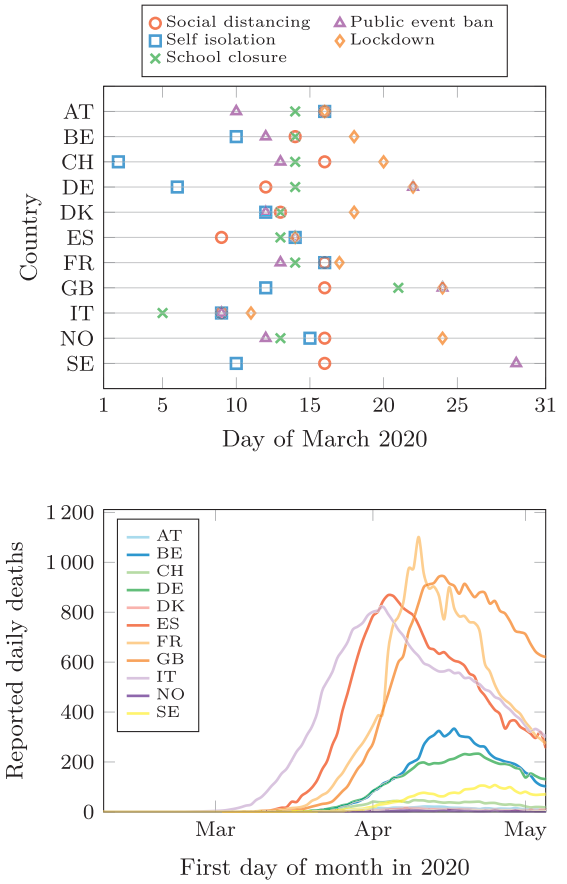

Figure 4
NPI enactment dates (top) and reported daily deaths (bottom) used in the published version [1] of the *Nature* model.

The condition number of this matrix is 36, but more severe for the model is that the condition number for the full J matrix is 600. This ill-conditioning[6] can be hard to detect directly from (15), where according to (13c), $\log \alpha_k$ can be estimated with a variance less then 0.04 times that of the observation noise λ (e.g., model fit).

The SVD of the inverse FIM (covariance matrix) block defined by $J_{\alpha\alpha}^{-1} = U\Sigma U^\top$ is given by

$$\Sigma = \lambda \operatorname{diag}\begin{bmatrix} 0.0018 & 0.0649 & 0.0498 & 0.0263 & 0.0331 \end{bmatrix}, \tag{16a}$$

$$U = \begin{bmatrix} -0.046 & 0.430 & 0.552 & -0.703 & -0.118 \\ 0.390 & 0.402 & 0.465 & 0.635 & -0.259 \\ 0.467 & -0.021 & 0.231 & -0.097 & 0.848 \\ 0.482 & -0.108 & 0.351 & -0.663 & -0.440 \\ 0.461 & -0.722 & -0.337 & 0.365 & -0.138 \end{bmatrix}. \tag{16b}$$

The first column of U, that corresponds to the smallest singular value, points less than 5° from the direction $\sum_k \alpha_k$, indicating that the summed effect of all NPI types constitutes the linear combination that can be identified with highest certainty. This should not be surprising, since all NPIs were enacted within a short time window in all but one country, as seen in Figure 4. Similarly, the last column of U reveals the linear combination of NPIs that is associated with the highest uncertainty (variance).

[6] The Stan code [8] used in the *Nature* model [1] throws a large number of warnings for numerical ill-conditioning; that might be a consequence of this.

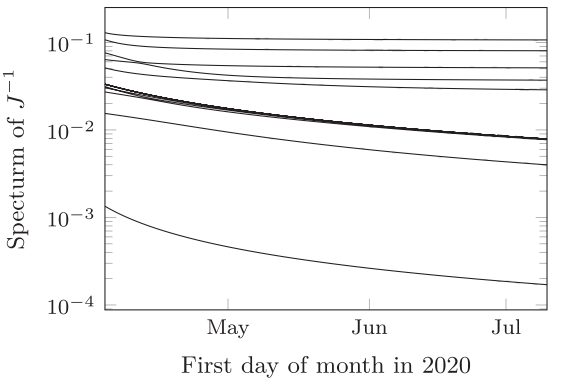

Figure 5
Evolution of the 16 (N_c = 11 R_0 values, N_α = 5 NPIs) singular values of the full inverse FIM J^{-1} in (13a) for the *Nature* model.

Returning to the unfactored covariance matrix (15) we note that the variance of the normalized sum

$$\operatorname{Var}\left(\sum_{k=1}^{N_\alpha} \frac{\hat{\alpha}_k}{\sqrt{5}}\right) = \frac{1}{5}\operatorname{Var}\left(\mathbf{1}^\top \alpha\right) = \frac{1}{5}\mathbf{1}^\top J_{\alpha\alpha}^{-1}\mathbf{1}\lambda = 0.002\lambda, \tag{17}$$

is an order of magnitude smaller than the smallest diagonal element of the unfactored matrix (15). This further indicates that the summed effect is much easier to identify than the effect of individual NPIs.

IDENTIFIABILITY OVER TIME

It can be argued that the illustrated identifiability issues result from a lack of data early in the pandemic, and that better estimates could have been obtained if the models were executed at a later time when more data was available. Let us therefore investigate how identifiability of the *Nature* and *Science* models—and a third related model that we are yet to introduce—has evolved throughout (the first pandemic year) 2020. We use exactly the same linear regression framework for all three models; only the NPI definitions and time frames differ.

THE NATURE MODEL

Figure 5 shows the day-by-day evolution of the singular values of the covariance matrix (15), as more data from the originally used data source [8] became available.[7]

Uncertainty in the principal directions of J^{-1} decreases during the spring of 2020, but then levels out, indicating that iden-

[7] We have blinded out data to emulate past dates. This is very similar, but not exactly identical, to using data causally available on those past dates, since most sources of Covid-19-related time series apply retrospective adjustments.

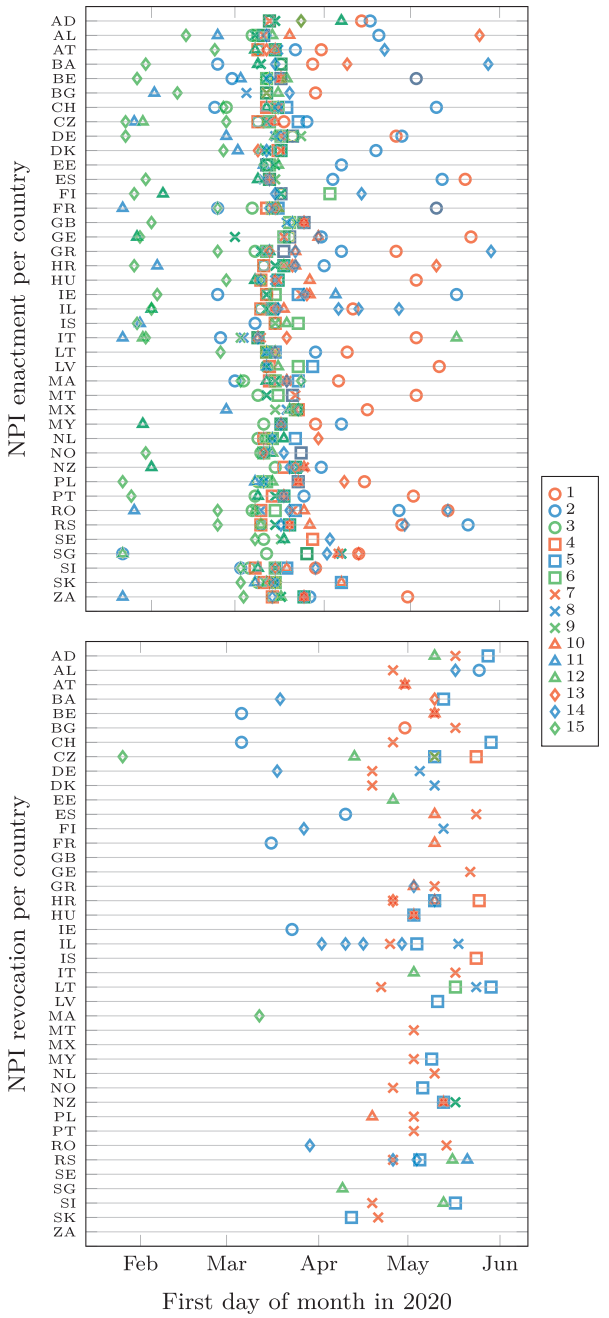

Figure 6
Dates when individual NPIs were enacted (top) and revoked (bottom) within the *Science* model. The NPI enumeration is consistent with [2].

tifiability issues, as pointed out in e.g., [7] and [6], have not improved markedly due to data available since the acceptance date of [1].

THE SCIENCE MODEL

A rule of thumb in system identification, when the regressors can be designed, is that they should resemble white noise. Such a design can be excluded for NPIs in pandemics, but the *Science* model features an input that is more persistently exciting

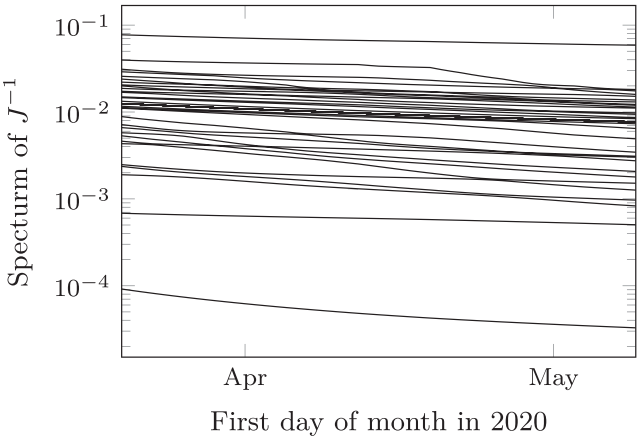

Figure 7
Evolution of the 56 (N_c = 41 R_0 parameters, N_α = 15 NPIs) singular values of the full inverse FIM J^{-1} in (13a) for the *Science* model.

than that of the *Nature* model, since the former was secured at a later point in time and has a regressor that keeps track not only of enactment, but also revoking of individual NPIs. The corresponding dates are shown in Figure 6, where it can also be seen that the number of NPI types is larger, compared to the *Nature* model.

The evolution of the singular values of the corresponding J^{-1} is shown in Figure 7. It shows that variance in the least certain principal direction has decreased roughly a factor of three.

THE INDEX MODEL

A third model [9], here referred to as the "Index" model, differs from the *Nature* and *Science* models in that its individual NPI effectiveness parameters are set by the modellers rather than estimated from data. The sum of these parameters over enacted NPIs at any given time forms the scalar "stringency index", taking values between 0 and 100.

Data to compute the index is taken from an impressive database, further described in [9], that keeps track of over 100 NPI (sub)types. The index is shown in Figure 8 (top). To reduce complexity of the graphics, we picked out the same 11 countries as used in the *Nature* model.

We apply our linear regression framework for the *Science* model (with the selected 11 countries) using this index as the sole NPI. This model thus has one α-parameter (the index) and 11 country-specific R_0 parameters. The singular values of the inverse FIM are shown in Figure 8 (bottom). Only three distinguishable lines are seen, since 10 of the 12 singular values are identical for this model structure. The careful reader may already have observed that the number of lines in Figures 5 and 7 are fewer than $N_c + N_\alpha$ for similar reasons.

The largest singular value remains steady at a similar level to where the corresponding value of the *Science* model levels out, and roughly 10 times lower than the corresponding level for the *Nature* model. This corresponds to a one-dimensional subspace, within which parameter values can move without affecting the model output much.

DISCUSSION

WHAT HAVE WE LEARNT?

The modelling of NPIs as instantaneously changing the reproduction number received much attention early in the Covid-19 pandemic, starting with the Imperial College COVID-19 Response Team (ICCRT) report [10] and subsequent publication [1], that categorized interventions into five NPI types. Identifiability issues due to high sensitivity of the model with the data available at the time (spring of 2020) were apparent and was addressed early in a technical blog post by Nic Lewis [6], and in a response by us [7] published alongside the original work [1]. The focus of that response was on how the partial pooling of national data within the model had incrementally changed in its official code base (8) between publication of the ICCRT report [10] and the *Nature* paper [1], and how these changes resulted in masking an apparent identifiability issue.

Here we have intentionally stayed away from the intricacies of how different models—or versions of the same model—have chosen to pool national data or define the NPI types to include, alongside the criteria associated with enactment of these NPI types within the models. Instead, we have taken one step back and regarded the model structure, and particularly that it is very closely related to a linear regression. Applying standard information theoretic analysis, we have then illustrated that the *information* in the data available in the spring of 2020—and presumably early during possible future epidemics of novel pathogens—is insufficient to uniquely distinguish the effects of (linear combinations of) NPIs.

While both the *Nature* and *Science* models attempt to identify the effects of NPIs on the reproduction number, the Index model instead provides an index based on a carefully curated NPI dataset. Within the herein considered framework, this index corresponds to preassigning the effectiveness parameters rather than identifying them from data. This obviously resolves the identifiability issue, but results in a signal (the index) that correlated very poorly with observed case or death data. This bias is visually apparent when comparing Figure 8 (top) to Figure 4 (bottom).

HAS IDENTIFIABILITY IMPROVED OVER TIME?

A natural question to ask is if more careful definition of the NPIs and choice of national data pooling could alleviate the aforementioned problem, and whether poor identifiability of NPI effects was merely a consequence of the models being applied too early, when only insufficient data was available. In relation to these questions, it is interesting to note that a model [1] similar to [2] was later published in *Science*. Without additional background, its diametrically differing conclusions regarding the effectiveness of lockdowns in particular (highly effective versus at best mediocre) could be interpreted as the sequel having addressed the above issues. To this end, we have plotted how the singular values of the effectiveness parameter covariance has evolved within the linear regression interpretation of the two models, with more data becoming available. These plots tell us how confidently identifiable the corresponding principal parameter space directions are with respect to the observation noise. Particularly, the condition number, being the ratio between the largest and smallest of them, tells us the identifiability ratio between the least and most certain principal direction. For both models we see that the information remains poorly conditioned, albeit much improved as compared to the spring of 2020. Looking at the corresponding principal directions, we could also see that the most certain principal direction corresponds to the sum of all NPIs within the *Nature* model.

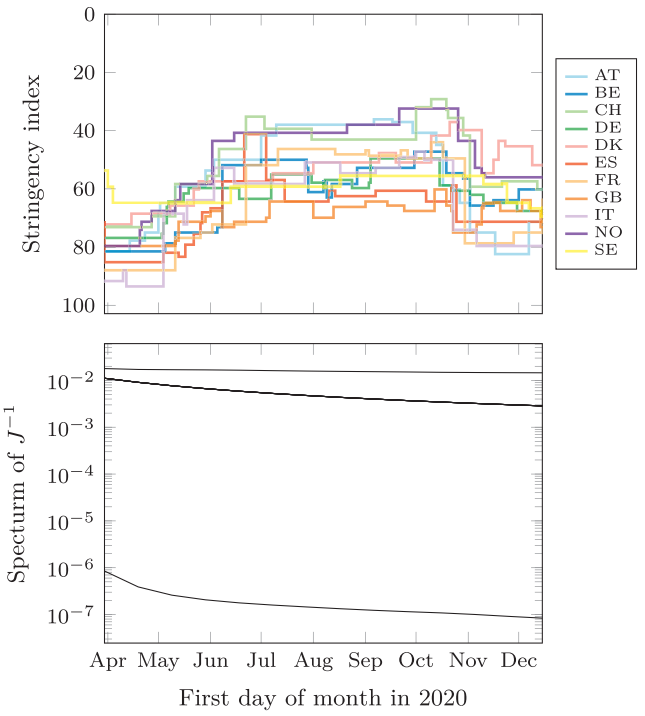

Figure 8

Stringency index according to the Index model for the same 11 countries that were included in the *Nature* model (top). Evolution of the 12 (N_c = 11 R_0 values, N_α = 1 NPI) singular values of the full inverse FIM J^{-1} in (13a) for the *Science* model (bottom). Note that 10 of the 12 singular values are identical for this special structure.

CONCERNS BEYOND THE DATA

So far our analysis has pointed out severe issues with the models while tacitly assuming them to be structurally sound. This is a very generous assumption.

For starters, all three models come with a linearity assumption, in that the combined effect of two NPIs equals the sum of the effects they would inflict if enacted individually. For example, the Index model defines the NPIs "close public transport" and "stay at home requirements". An educated guess is that the impact on virus spread of the former is decreased should the latter be enacted.

To continue, there is a time-invariance assumption leaving no room for saturation effects or improvements over time within the healthcare sector, mutations resulting in more or less transmissible or harmful pathogen variants, etc.

However, the most pushing point is that of causality. Since neither of the models leave room for any extrinsic variable but the NPIs to explain virus spread, they are bound to explain the decline of the first pandemic wave with the NPIs (and provided that the data is uninformative are most certain about the combined NPI effect). To what extent have changes in behavior, other than those enforced by legislation, affected the pandemic trajectory? What has the role of the change in season been on viral transmission? These are hard questions, and ones that remain unanswerable within the considered models.

CONCLUSION

In the end, it all comes down to fundamental properties of the true dynamics, the model, and the data:

- Good curve fitting does not validate a model. Cross validation on fresh data is the preferred procedure, and when not possible (for example early in an epidemic), identifiability and sensitivity analyses need to be carefully conducted.
- Good curve fitting does not imply identifiability, only that the model is complex (i.e., flexible) enough to describe observed data.
- Identifiability does not imply good curve fitting, only that the model is not too complex compared to what is measured.
- Uninformative data, resulting from a lack in persistence of excitation, implies that no conclusions whatsoever can be drawn on the model.
- Modelling assumptions directly impact the usefulness of the resulting model. This holds true also for models that can be reliably identifiable from data.

We have demonstrated that several NPI models, including ones published in *Nature* and *Science*, fall short in all these aspects.

REFERENCES

1. Flaxman, S., Mishra, S., Gandy, Juliette, A. H., Unwin, T., Mellan, T. A., et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, Vol. 584, (2020), 257–261.
2. Brauner, J., Mindermann, S., Sharma, M., Johnston, D., Salvatier, J., Gavenčiak, T., et al. Inferring the effectiveness of government interventions against COVID-19. *Science*, Vol. 371, 6531 (2021).
3. Kermack, W., and McKendrick, A. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 115, 772 (1927), 700–721.
4. Public Health Agency of Sweden. Folkhalsonyndigheten_covid19 [dataset], https://www.arcgis.com/sharing/rest/content/items/b5e7488e117749c19881cce45db13f7e/data, last access May 5, 2021.
5. Meyn, S., and Tweedie, R. *Markov Chains and Stochastic Stability*. London: Springer-Verlag, 1993.
6. Lewis, N. Did lockdowns really save 3 million COVID-19 deaths, as Flaxman et al. claim?, https://judithcurry.com/2020/06/21/did-lockdowns-really-save-3-million-covid-19-deaths-as-flaxman-et-al-claim, last access May 4, 2021.
7. Soltesz, K., Gustafsson, F., Timpka, T., Jaldén, J., Jidling, C., Heimerson, A., et al. The effect of interventions on COVID-19. *Nature*, Vol. 588, E26-E28 (2020).
8. Imperial College COVID-19 Response Team (ICCRT). Covid19model, https://github.com/ImperialCollegeLondon/covid19model, last access Jun. 1, 2020.
9. Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., et al. A global panel database of pandemic policies (Oxford COVID-19 government response tracker). *Nature Human Behaviour*, Vol. 5, (2021), 529–538.
10. Flaxman, S., Mishra, S., Gandy, A., Unwin, H., Coupland, H., Mellan, T., et al. Report 13 UPDATE: Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries—Technical description update. Imperial College London, Tech. Rep., 2020. Available: http://spiral.imperial.ac.uk/handle/10044/1/77731, last access Apr. 24, 2020.

Fredrik Gustafsson is professor in Sensor Informatics at Department of Electrical Engineering, Linköping University, since 2005. He received the M.Sc. degree in electrical engineering 1988 and the Ph.D. degree in Automatic Control, 1992, both from Linköping University. His research interests are in stochastic signal processing and related areas. He is a co-founder of the companies NIRA Dynamics (automotive safety systems), Softube (audio effects) and Senion (indoor positioning systems).

Kristian Soltesz is associate professor with the Department of Automatic Control, Lund University, since 2019. He received the M.Sc. and Ph.D. degrees in Automatic Control in 2008 and 2013, both from Lund University. His research interests span various topics within medical control systems. Kristian is a member of the IEEE Committee on Healthcare and Medical systems, has been engaged in several COVID-19-related research topics during the past year and presented a keynote talk on the topic at the FUSION 2020 conference. One year later, he says: "*[…] It is inspiring to see the level of societal engagement […] within many scientific communities, and notably amongst epidemiologists applying statistical signal processing and dynamic modelling methods. […] [M]odelling early in a pandemic could also serve to guide resource allocation, should a similar situation arise in the future.*" —Excerpt from an interview by M. Mihaylova, May 2021.