Dynamic time series modelling and forecasting of COVID-19 in Norway*

Gunnar Bårdsen[1,2] and Ragnar Nymoen[2]

[1]Norwegian University of Science and Technology
[2]University of Oslo

May 4, 2023

Abstract

A framework for forecasting new COViD-19 cases jointly with hospital admissions and hospital beds with COVID-19 cases is presented. This project, dubbed CovidMod, produced 21-days ahead forecasts each working day from March 2021 to April 2022, and forecast errors that were used to assess forecast accuracy. A comparison with the forecasts of the Norwegian Institute of Public Health (NIPH), with dates of origin in the same period, favours the CovidMod forecasts in terms of lower RMSFEs (Root Mean Squared Forecast Errors), both for new cases and for hospital beds. Another comparison, with the short term forecasts (7 day horizon) produced by a forecasting project at the University of Oxford, shows only little difference in terms of the RMSFEs of new cases. Next, we present a further development of the model which allows the effects of policy responses to a central model parameter to be forecasted by an estimated smooth-transition function. The forecasting performance of the resulting non-linear model is demonstrated, and it is suggested as a possible way forward in the development of relevant forecasting tools in general and for pandemics in particular.

Introduction

During the COVID-19 pandemic, there was a surge in both demand for and supply of forecasts of how the coronavirus was spreading in societies. Forecasts of new cases, and of new admissions to hospitals, were sometimes presented as premises for the many decisions that governments took to manage the crisis, in particular with regard to lockdowns and other non-pharmaceutical policy measures.

In Norway, the National Institute of Public Health (NIPH) published weekly "Situational awareness reports", with forecasts of the daily number of new infections, hospital beds and ventilator beds. The main sequence of forecasts had a 21-day horizon and were based on the institute's National Regional Model.[1] One key parameter was the reproduction number (R)

*Earlier versions of this work have been presented at the "25. Dynamic Econometrics Conference", London 13-14 April; at the "Network Meeting for Modellers" arranged by Norwegian Institute of Public Health, 3 November 2022; and at the start up workshop of the project "Model invariance and constancy in the face of large shocks to the Norwegian macroeconomic system", 22 March 2022. We would like to thank participants for comments and discussion. Norwegian Research Council project number 324472.

[1]See https://www.fhi.no/en/id/infectious-diseases/coronavirus/coronavirus-modelling-at-the-niph-fhi/.

which was calibrated so that the model was adapted to historical hospital admissions and test-data on a national level.[2]

The modelling group at NIPH revised reproduction numbers, in order to be able to capture new developments. Nevertheless there was never an expectation that the NIPH's 3-week forecast materialised unless the main drivers, eg., virus characteristics and infection control measures, remained as they were close to the data of origin of the forecast over the forecast horizon.

Changing parameters in the underlying processes is a common problem in forecasting. A given model might work very well within sample (the "training data set"), but yet produce poor forecasts if the structure of the underlying processes change in a time period after the preparation of the forecast. When modelling pandemic data, being the result of many changing factors, structural breaks occurring after the forecast preparation date are clearly a potential problem.

Ideal samples for estimation of the epidemiological R-number may be data from the early period of the pandemic, where the data were not so affected by lockdowns or other non-pharmaceutical policies and interventions, Korolev (2021). Yet, that estimated R-number would be of limited value in forecasting the pandemic after the intervention was started.

Even if all model-based forecasts are vulnerable to structural breaks and shocks that occur after the origin of the forecast, there might be strategies to derive forecasting models with a higher degree of robustness against breaks in relationships and parameters. Such approaches, which for the pandemic data may be complementary to epidemiological structural models, include trend models, stochastic trend models, autoregressive models, and robust forecasting methods, see Li and Linton (2021), Harvey and Kattuman (2021), Doornik et al. (2022),Mills (2022) amongst others.

Although an epidemiological model like the National Regional model is clearly a good model for understanding the deeper processes of the pandemic, it can still happen that the model's forecasts are less accurate than a time series oriented empirical model. The choice of model providing the most accurate forecast is an empirical question—depending upon the changing pandemic characteristics—and is likely to change. If the accuracy (quality) of the forecasts of the pandemic is at all of interest, it follows that the forecast errors of alternative models should be analysed and compared with improved forecast performance in mind.

This paper contains results from a project based on the premise that new cases (the measured incidence), hospital admissions, and the number of patients in hospital with COVID-19, can be modelled as a interrelated dynamic system — an approach not dissimilar to one often taken in macroeconometrics. More specifically, the model belongs to the class of autoregressive systems. Within this approach, the reproduction number R becomes not a parameter, but a time-varying variable within a model forecasting incidence and hospitalisations.

One important parameter in the three-weeks-ahead econometric forecasting model is the long-run (zero frequency) root of the characteristic equation of the process assigned to the development of new cases, Nymoen (2019, p. 322-323). This critical parameter can be estimated from the time series of new cases. The constancy of the critical parameter can be evaluated on a daily basis. The degree of invariance of the parameter is also testable and non-pharmaceutical policy responses as well as vaccinations and virus mutations are also incorporated.

After next presenting the conceptual framework, we explain how it is made operational in the form of a multiple-equation time series model, and how it (in the period between march 2021 and May 2022) was applied to produce daily 3-week ahead forecasts. One thing we note, is that the adaptation of the model structure to incorporate effects of policy measures, behavioural changes among the public etc, has been essential in order to produce forecasts with information value.

We also review comparisons of the accuracy of real time forecasts with NIPH-forecasts as

[2]As well as county level, but we only consider national forecasts in this paper.

well as with the Covid forecasts, as part a multi-country project in short-term forecasting, produced by a group of researchers at Oxford University.

While it is important to follow an adaptive modelling tactic when doing real time forecasting, it is important to take the opportunity to use the data as a training data set for a more strategic development of the forecasting model. Specifically, the paper introduces a new and extended version of the model, where there is a distinction between exogenous shocks/breaks, due to infections, and the effects of policy responses which are endogenous. The model produces results that are promising on a test data set, which indicates that it might be a useful complementary addition to the toolbox.

A model based on autoregressive processes

The model belongs to the class of autoregressive models. If we let Y_t denote new cases on day t, the starting point may be taken to be the dynamic model equation:

$$Y_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i Y_{t-i} + \gamma X_t + \epsilon_{,t} \tag{1}$$

where $\alpha_i, i = 0, 1., ...p$ and γ are parameters. X_t represents conditioning variables, which is interpreted broadly to include deterministic terms. Lags of X_t has been kept out the notation for simplicity. ϵ_t is an error term which we assume linearly independent of $Y_{t-1}, ..., Y_{t-1-p}$ and X_t.

Given the complexity of the true process that have generated $Y_t, t = 1, 2, ...$, the data generating process (DGP), compared to the simple linear form of the explanatory part of the model (1) it is plausible that the error-term process is non-stationary, with time dependent moments, although not integrated. Hence, right from the start it is clear that it will be near impossible to obtain a congruent model of the DGP. However, when the purpose is forecasting it is possible than models which are non-congruent can perform relatively well, in particular in comparison with models that are not similarly time-oriented. First and foremost, forecasting is a time-oriented activity, Granger and Newbold (1986).

Equation (1) can be re-written as:

$$Y_t = \alpha_0 + kY_{t-1} + \sum_{i=1}^{p-1} \alpha_i^{\dagger} \Delta Y_{t-i} + \gamma X_t + \epsilon_t, \tag{2}$$

where the coefficients $\alpha_i^{\dagger}$ $(i = 1, 2, ..., p-1)$ are combinations of the autoregressive coefficients and k is:

$$k = \sum_{i=1}^{p} \alpha_i$$

We note that $k = 1$ implies:

$$\Delta Y_t = \alpha_0 + \sum_{i=1}^{p-1} \alpha_i^{\dagger} \Delta Y_{t-i} + \gamma X_t + \epsilon_t, \tag{3}$$

hence there is a (low frequency) unit root in the homogeneous part of the equation generating Y_t, cf. Nymoen (2019, p. 322-323).

Moreover, since the relevant solutions of the stochastic difference are causal solutions, we note that $k > 1$ will result in explosive solutions, and $k < 1$ in stable solutions.

Hence, k is an important parameter. One could say that it is a critical parameter, since a value larger than one means that the process implies that the daily number of new cases will grow exponentially. As long as the coronavirus created serious illness among large swaths of the population, one consequence could be that the health care system became overflowed with patients.

The parameter k can be estimated from the data almost in real-time. Signs of changes in k can be investigated by the use of for example recursive or rolling regressions. It is possible to imagine k as a parameter that the health authorities would find useful to monitor, and as a parameter that they would want to push below 1 in order to check the explosive process and to secure that the health service system remained intact during the pandemic.

The estimates of k may change as a result of the evolutionary path taken by the virus, through the continuous development of new versions and through intermittent mutations. But as noted above, k may also change in response to non-pharmaceutical interventions (NPIs), and more gradually as a consequence of vaccination and of herd immunity.

If the autoregressive model is used to forecast new cases, feasible methods to identify and represent breaks in k (once they have occurred) need to be part of the forecasting practice and method. Of course, smaller and more gradual changes in k will be difficult to identify empirically. However, even getting the intermittent larger breaks in k into the forecasting model with little delay, may mean a difference between useful and useless forecasts during a pandemic.

As noted, a main distinction can be drawn between autonomous breaks in k due to new virus versions and mutations, and breaks that come as a result of NPIs. In our real-time forecasting, the practical method was to adjust the estimation of k based on the evidence that could be gleaned from the data as it came in, both autonomous and policy induced.

However, there are alternatives that could have improved on our informal methods. As one example of a way forward we show how effects of policy-induced structural breaks in k can be included in the forecasting model with the aid of non-linear (threshold) estimation.

Returning to equation (1), another way of re-writing it gives:

$$Y_t = \alpha_0 + \beta(\sum_{i=1}^{p-1} Y_{t-i}) + \sum\nolimits_{i=1}^{p-1} \alpha_i^{\ddagger} \Delta Y_{t-i} + \gamma X_t + \epsilon_t \tag{4}$$

where:

$$\beta = \frac{k}{p-1} \tag{5}$$

and $\alpha_i^{\ddagger}$ $(i = 1, 2, ..., p-1)$ are other combinations of the original autoregressive coefficients than $\alpha_i^{\dagger}$ $(i = 1, 2, ..., p-1)$ in (2).

The accumulated number of confirmed cases for day t is denoted S_{Yt}. It is given by the definition:

$$S_{Yt} = Y_t + S_{Yt-1}, \tag{6}$$

The change in accumulated cases over a period of q days, $\Delta_q S_{Y,t}$ is therefore the cumulative sum of new cases over q days:

$$\Delta_q S_{Y,t} = S_{Yt} - S_{Yt-q} = \sum_{i=0}^{q-1} Y_{t-i}, \tag{7}$$

and for $t-1$:

$$\begin{aligned} \Delta_q S_{Y,t-1} &= S_{Yt-1} - S_{Yt-1-q} \\ &= \sum_{i=0}^{q-1} Y_{t-1-i} = \sum_{i=1}^{q} Y_{t-i} \end{aligned} \tag{8}$$

If we set $q = p-1$,

$$\begin{aligned} \Delta_{(p-1)} S_{Y,t-1} &= S_{Yt-1} - S_{Yt-1-(p-1)} \\ &= S_{Yt-1} - S_{Yt-p} = \sum_{i=1}^{p-1} Y_{t-i} \end{aligned} \tag{9}$$

Equation (4) can therefore be expressed with the change $S_{Yt-1} - S_{Yt-p}$ on the right hand side:

$$Y_t = \alpha_0 + \beta(S_{Yt-1} - S_{Yt-p}) + \sum_{i=0}^{p-1} \alpha_i^{\ddagger} \Delta Y_{t-1-i} + \gamma X_t + \epsilon_t. \tag{10}$$

A working hypothesis could be that the amount of virus in the population, call it S_t^* (which is not observable) is increasing in the accumulated number of confirmed cases over a period of p days:

$$S_t^* = f(S_{Yt-1} - S_{Yt-p}),\ f' > 0, \tag{11}$$

which shows that "behind" the simple autoregressive model for new cases, it is the "stock" of infectious virus that generates the flow of new cases. Hence expression (10) and (6) highlights the positive feedback between between new cases and how much corona virus there is in the population.

Another reason why we have used (10) is more practical. As noted above, we modelled and forecasted the number of new cases reported daily in Norway. The time series was by the specimen collection date. There were typically 1-2 days delay between diagnosis and registration in the Norwegian Surveillance System for Communicable Diseases (MSIS). The number of cases was updated retroactively when new notifications arrive.[3]

The most important consequence for our forecasting was that it created a "ragged edge problem": when we prepared a forecast for $T+1$, $T+2$....,$T = 21$ the correct number of cases for important days in the initialisation period, in particular Y_T, Y_{T-1} and Y_{T-2} were not observable. We worked-around this problem by substituting the preliminary (and wrong) measurements by "nowcasted" values $\hat{Y}_T$, $\hat{Y}_{T-1}$ and $\hat{Y}_{T-2}$. For this we made us of the number of cases ordered by the date of registration, which were not subject to revision, together with lags of the numbers of cases by the specimen collection dates. For this reason, the summation in (4) improved the 21-day ahead forecasts compared to (2) which was quite vulnerable to inaccurate nowcasting of a particular single day, for example Y_T.

A flow chart view of the model

The model used for real time forecasting, includes four endogenous variables:

- Y_t, number of new confirmed cases with COVID-19 day t.[4]
- $S_{Y,t}$, accumulated number of confirmed cases COVID-19, day t.
- HA_t, number of new admissions to hospital with COVID-19, day t.[5]
- HB_t, number of persons in hospital with COVID-19 ("hospital beds") day t.[6]

Figure 1 shows a flow chart of the model, which was dubbed CovidMod. Early in the pandemic, new cases with COVID-19 were in practice imported, as in particular downhill skiers returned from South-Tyrol and other locations where there already were outbreaks. This early development is indicated by the line from the *COVID-19 abroad* node in the flow-chart. The "import channel" remained important during the the whole pandemic, but it was checked and in periods practically cut-off by forced and voluntary reductions in travel abroad and by the quarantine regime (indicated by the dotted line that goes from the node for *Mutations, behavioural changes and vaccines.*

[3] https://www.fhi.no/en/id/infectious-diseases/coronavirus/daily-reports/daily-reports-COVID19/
[4] https://statistikk.fhi.no/msis.
[5] https://www.fhi.no/sv/smittsomme-sykdommer/corona/dags--og-ukerapporter
[6] https://www.helsedirektoratet.no/statistikk/antall-innlagte-med-pavist-COVID-19-for-nedlasting

We have not explained new cases that were a direct result of international travel, and so changes in that ended up the error term in our model. However, once the COVID-19 had become a reality in Norway, our model could pick up the positive feed-back loop between new cases and the number of person with COVID-19 who could transmit the virus to other persons. This is indicated by the line with the double arrowheads between the $\Delta_q S_{Yt-1}$ node and Y_t node. This relationship, the pandemic feed-back loop, underwent numerous changes during the period. Non-pharmaceutical measures (lockdowns in particular), meant that fewer new cases resulted from the existing $\Delta_q S_{Yt-1}$ in the population, while the lifting of those measures worked in the other direction. As the vaccination rate grew, less social distancing was necessary to control the pandemic.

Mutations of the virus played a major role, and affected both the feed-back loop between $\Delta_q S_{Yt-1}$ and new cases. It also affected how new hospitalisations responded to more COVID-19 in the population, represented by the line between $\Delta_q S_{Yt-1}$ and HA_t. Eventually, the vaccination program became the single most important factor in breaking down that relationship, so that the society could open up again without any increase in new hospitalisations.

In CovidMod, the number of hospital beds (HB) is determined by the law-of-motion relationship:

$$HB_t = HA_t - \delta\, HB_{t-1} \tag{12}$$

where the "capital depreciation" parameter δ is estimated.

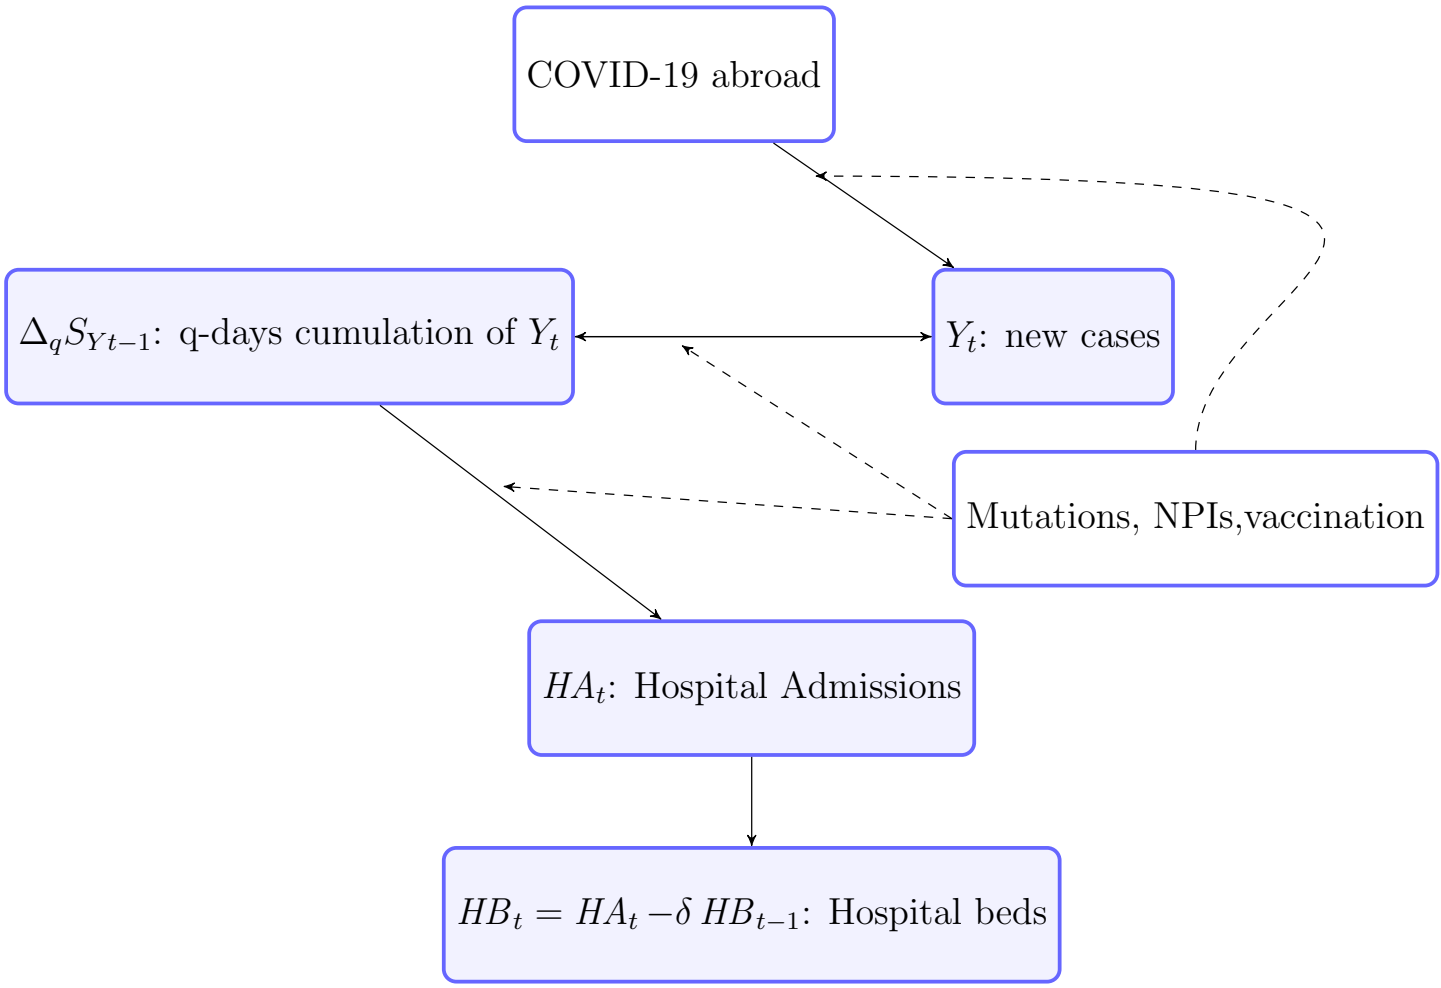

Figure 1: Variables and relationships of CovidMod. Relationships that may change (structural breaks) as a consequence of virus mutations, NPIs (inducing behavioural changes) and vaccination, are indicated by dotted lines.

As illustrated by the figure, and as pointed out by Mills (2022) there is a "causal structure" inherent in the relationship between new cases, hospital admissions and beds (and ultimately, deaths). We now explain how an multiple-equation model that captures the structure has been specified, estimated and used in practical daily forecasting.

Specification, estimation and simulation of the model

As noted above, the detailed model specification was developed through practical forecasting in real time. A pilot model version was developed in the winter of 2020, soon after the outbreak of the pandemic. In the summer the model became defunct, but when the forecasting project was re-started in March 2021, it was possible to build on the pilot version from 2021.

Then, from mid March 2021 and until early in May 2022, forecasts were produced five times weekly (ie., work-days). During that period, abrupt and more gradual changes in the dynamics of the pandemic were incorporated in the model by step-indicator variables (step dummies) and an by variables that measures the percentage of the population that have received vaccination (one of more doses).

It is our experience that keeping a keen eye on the model's forecast errors helped avoid unnecessary long response time, from structural changes in the DGP, to model revisions to accommodate consequential breaks. Conversely, it is hard to imagine that forecasts based on the idea of functional relationships, would not be damaged by failing to study past forecast errors.

The step-indicator variables can be divided in two categories:

- Abrupt exogenous events, like mutations (D).
- Non-pharmaceutical policy responses (G).

We next explain how variables of the two categories were included in the operative forecasting model. Since a new model was estimated for each of the forecasts, the project produced a whole sequel of models. As an example of the model specification in mid- November 2021, we consider the model estimated on the sample that begins with 21 February 2020 and ends with 22 November 2021.

The estimated equation for new cases Y_t, the empirical counterpart to equation (10), was:

$$\begin{aligned}
\widehat{Y}_t = & \underset{(0.001)}{0.072}\,(S_{Yt-1} - S_{Yt-14}) \\
& + \underset{(0.003)}{0.019}\,(S_{Yt-1} - S_{Yt-14})\,D_{\alpha t} \\
& + \underset{(0.006)}{0.014}\,(S_{Yt-1} - S_{Yt-14})\,D_{\delta t} \\
& - \underset{(0.002)}{0.013}\,(S_{Yt-1} - S_{Yt-14})\,G_{Mar21t} \\
& + \underset{(0.002)}{0.017}\,(S_{Yt-1} - S_{Yt-14})\,G_{Oct21t} \\
& - \underset{(0.011)}{0.051}\,(S_{Yt-1} - S_{Yt-14})\,VAC_t \\
& + \sum\nolimits_{i=0}^{13} \hat{\alpha}_i^{\ddagger} \Delta Y_{t-1-i} + \text{other short term factors} \\
& \text{Estimation period: } 21.2.2020 - 22.11.2021,\ 641 \text{ obs.}
\end{aligned} \tag{13}$$

The estimation method was OLS. Heteroscedasticity and autocorrelation consistent standard errors (Newey-West) are reported in parentheses below the estimated coefficients.

In this model equation, there is a reference estimate of β in the theoretical equation (10), namely 0.072 in the first row of (13).

In the second and third row, we find $D_{\alpha t}$ and $D_{\delta t}$, two step-dummies that are zero before the Alpha and Delta variants became dominant, and 1 after. They enter multiplicative with $(S_{Yt-1} - S_{Yt-14})$. The positive coefficients imply that the positive feed-back between new cases and the apremount of transmittable corona disease became stronger with theses two variants became dominant. When both dummies take the value 1, the implied estimate for β becomes 0.105 and the implied k-number becomes 1.36. Hence the process that generated new cases had become critical (explosive).

G_{Mar21t} represents the (renewed)-lockdown in March 2021, and G_{Oct21t} is a dummy for a partial opening-up of the society in October 2021. The interpretation is that these changes in non-pharmaceutical policies worked as expected: The lockdown weakened the positive feed-back loop. And the opening-up strengthened it.

The last of the multiplicative terms shows that a higher vaccination rate (VAC) had the expected dampening effect on the feed-back loop, and can be imagined as a force that puts downward pressure on β and the critical parameter k as VAC_t increases in a continuous manner towards 1. However the estimated coefficient is quite large in magnitude, and a value of 0.4 implies an effect that is comparable the estimated effect of the NPI measures.

The second estimated equation in the model is the conditional model of new admissions to hospital with COVID-19 (HA):

$$\begin{aligned}\widehat{HA}_t = & \underset{(0.081)}{0.276} HA_{t-1} + \underset{(0.0014)}{0.0035} Y_t \\ & + \underset{(0.00013)}{0.00039} \left(S_{Yt-3} - S_{Yt-19}\right) \\ & + \underset{(0.0001)}{0.0004} \left(S_{Yt-3} - S_{Yt-19}\right) D_{\alpha t} \\ & + \underset{(0.0001)}{0.0005} \left(S_{Yt-3} - S_{Yt-19}\right) G_{Oct21t} \\ & - \underset{(0.0002)}{0.0009} \left(S_{Yt-1} - S_{Yt-14}\right) VAC_{t-7} \\ & + \sum\nolimits_{i=0}^{2} \hat{\gamma}_i \Delta HA_{t-1-i} + \text{other short term factors} \\ & \text{Estimation period:14.7.2020} - \text{22.11.2021, 497 obs.}\end{aligned} \tag{14}$$

The main explanatory factor is the number of new cases in the form of the intermediate term cumulative number $S_{Yt-3} - S_{Yt-19}$. In addition there is an estimated impact effect of Y_t which at first can be difficult to rationalise, since there is normally a time lag between attracting the virus and being admitted to hospital. Although the statistical significance of Y_t was robust, there was a structural break at the end of April 2021. The estimated coefficient dropped, albeit not to zero. However, as there was no change in policies or in the virus at that point in time, we do no twant to classfiy that regime-shift as one or the other. Therefore, the reported coefficient of Y_t is the post-break estimate. Finally, as Y_t was forecasted by the use of (13), keeping it in equation (14) did not represent a cost in the operation of the forecasting model.

Interestingly, the estimated model equation contained effects of virus mutation (but only of the Alpha variant), lock-down and vaccination, with the expected signs of the coefficients.

The third empirical equation in the model estimated on 22 November 2021 was the simple law of motion for the number of hospital beds with COVID-19:

$$\widehat{HB}_t = HA_t + \underset{(0.0029)}{0.898}\, HB_{t-1} \tag{15}$$

$$\text{Estimation period:14.7.2020} - \text{22.11.2021, 497 obs.}$$

As noted, a realistic picture of the DGP of new cases Y and hospital admissions (HA), is that it is complex and changing. In comparison the model equations are simple. Hence, there are several factors that end up the the error-terms of the estimated equations. As mentioned, import of virus trough international travel has played a major role during the pandemic. This factor is not represented in CovidMod and is a source of simulation and forecast errors.

Figure 2 shows how the model estimated on 22 November 2021 simulated the historical development of the variables from 1 July 2021 to 20 November 2021.

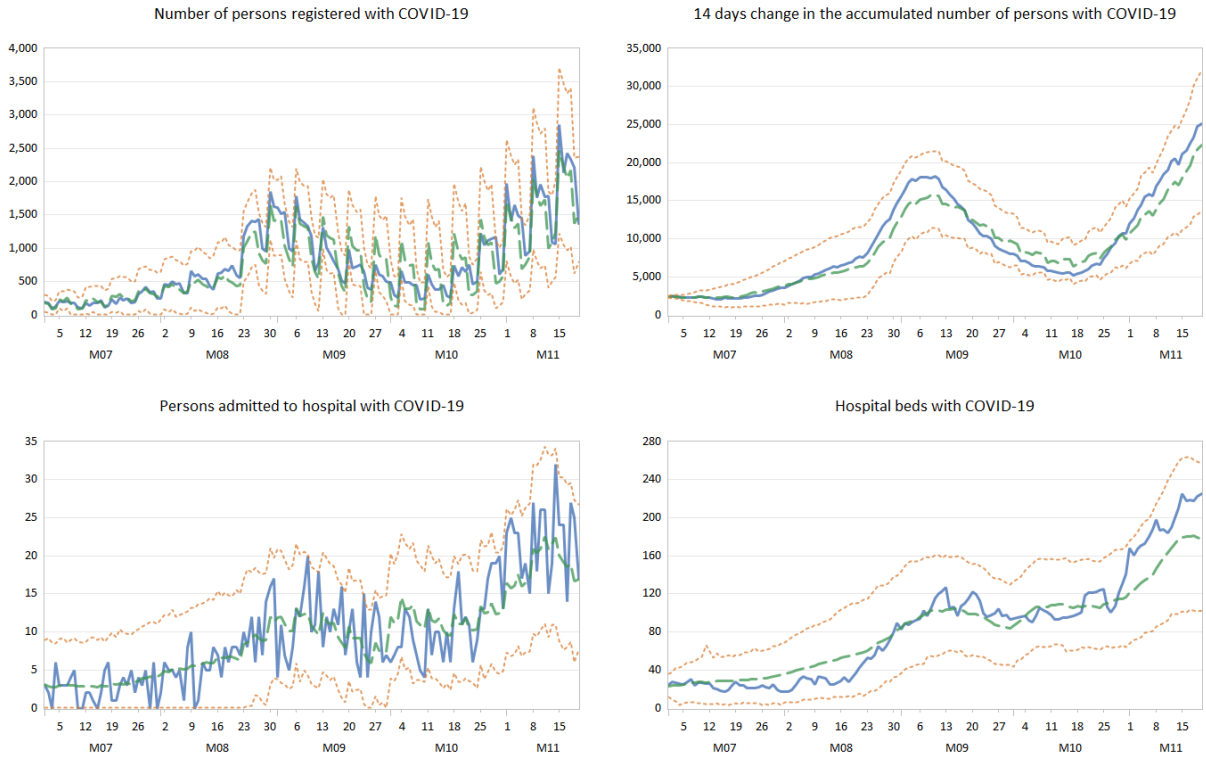

Figure 2: Example of dynamic simulation of CovidMod. Simulation started 1 July 2021. End of simulation was 20 November 2021. Line graphs represent actual values, dashed lines simulated values and dotted and lines represent upper and lower bounds of 90 % uncertainty intervals.

The model dynamic simulation shown in the figure demonstrates that the model was able to explain many properties of the actual data.

It may also be of interest to look at this simulation through the lens of the time varying reproduction number, the R-number. Following Harvey and Kattuman (2021), it can be calculated as the ratio between the sum of new cases over the last four days and the corresponding sum for the previous four days. More generally:

$$R_{t,r,p} = \frac{\sum_{j=0}^{p-1} Y_{t-j}}{\sum_{j=r}^{p+r-1} Y_{t-j}}, \tag{16}$$

where the sum in the denominator starts at lag r and the sums of numerator and denominator can overlap.

Simulated R-numbers can be obtained by substituting actual numbers of new cases by simulated variable values. For the simulation of number of new person registered with COVID-19 (ie., new cases) shown in Figure 2, the time varying R-number using $p = r = 4$ is plotted in Figure 3. The figure shows a simulated R-number close to one at the start of July 2021.

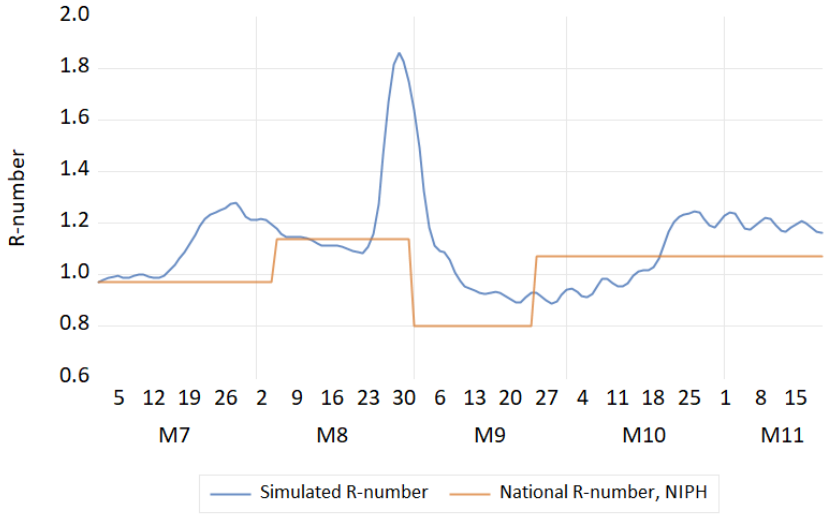

Figure 3: Dynamic simulation of a time-varying $R_{t,4,4}$, using the simulated number of new cases in Figure 2. Centred seven-day moving average. The national R-number for Norway produced by NIPH is also plotted, as a reference.

It then increased gradually during July, the main summer holiday month which meant that many travelled abroad, although less than in normal times also. Towards the end of August, there was a marked spike in the simulated R-number. During September 2021 the number was reduced towards one and even a little below one. There was a new increase in October, to a level above one, and this level lasted for the rest of the simulation period.

The national R-number produced by NIPH was visibly more constant over relatively long periods. It was estimated by NPIH to be 0.97 until 4 August, and to have been 1.14 during the rest of that month. The lowest estimated, 0.8, was for the period from 1 September to 24. September. From 25. September the national reproduction number was estimated to be 1.07.

It is not shown in Figure 3, but the next change the NPIH estimated R-number came 13. December 2021. It was then estimated to be 0.79.[7]

Real-time forecast performance

The model simulation in 2 was based on actual values of the vaccination indicator, and also the empirically identified breaks due to changes in the virus, and changes in the NPI regimes decided by the government, based on the advice of the experts working in the National Institute of Public Health and in the The Norwegian Directorate of Health.

Fundamentally, when the framework was used to forecast, in real time, new cases, admissions to hospital and hospital beds, it was not possible to condition on future values of any exogenous explanatory variables. The missing modelling of effects of mutations of the virus that occurred after the origin date of a forecast as well as effects of policy responses therefore contributed to the forecast errors of the daily real time forecasts.

The development of the vaccine indicator over a 21 day forecast horizon was also uncertain, and had to be forecasted. We did so by interpolation, between the actual value of the vaccine variable VAC close to the date of origin of a forecast and target level of the variable at a future point in time chosen in accordance with published plans for vaccination.

[7]Source: Table 1 (page 21) in *Situational awareness and forecasting for Norway*, dated 5. January 2022.

As mentioned, future effects of policy responses was not modelled in the version of the model that was used to produce daily forecasts. An approach to model the effects of policy responses over the forecast period is presented in the next section and seems to constitute an improvement of the model for forecasting.

Figure 4 shows forecasts and actuals for new cases in the period from 1 November 2021 to 10 April 2022. Actual numbers are shown in the red line graph. CovidMod forecasts are shown in the many black lines graphs. Each line graph represent a 21-day forecast, but since the forecasts were produced each workday, there is a considerable overlap. NIPH forecasts are shown in the green line graphs. The green line graphs are few in number. This is because publications of forecasts of COVID-19 "incidence" in the "Situational awareness and forecasting" reports became irregular i December, and in January 2022 they were discontinued. The last 3-week forecast of "incidence" was published in the report dated 5 January 2022.

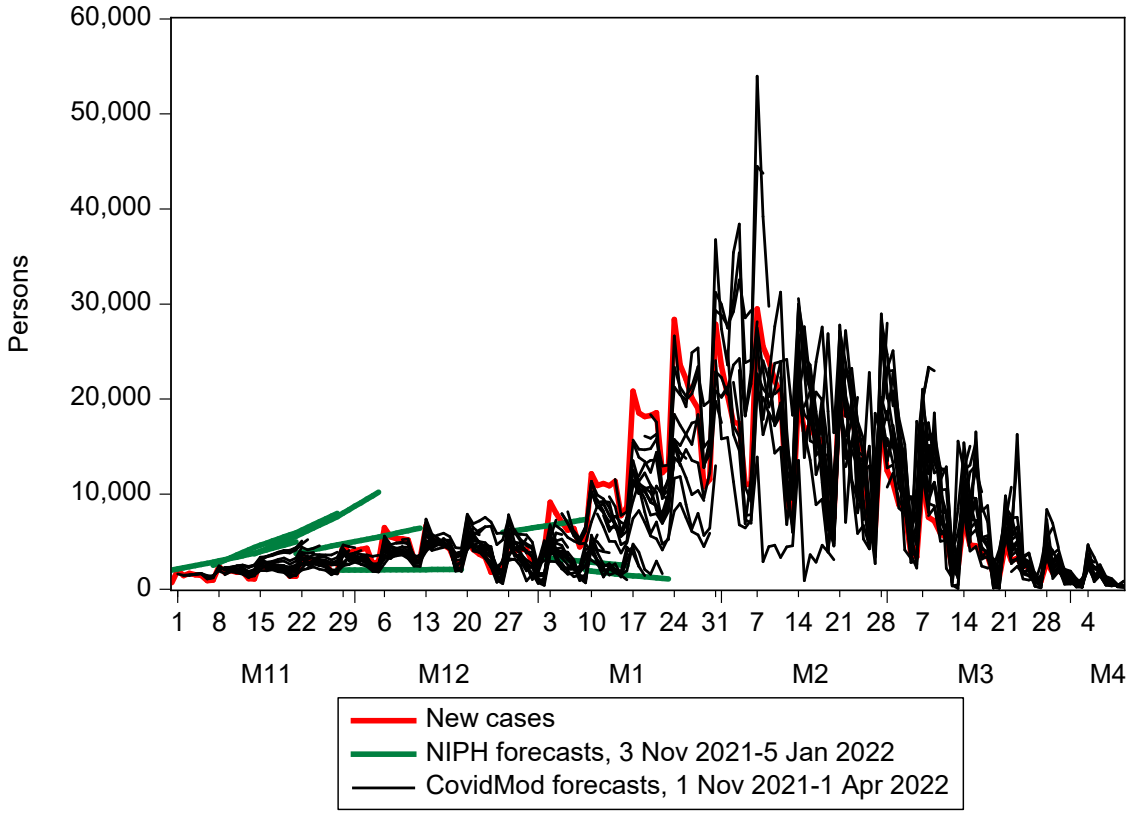

Figure 4: Forecasted number of new cases (daily data) and forecasts from CovidMod and NIPH (National regional model).

The CovidMod forecasts appear to have been relatively good during November and December, although the scale of the figure may make the forecasts to appear better than they were perceived in real-time. Early in January 2022, the actual number of daily new cases was significantly higher than the forecasts with date of origin two or three weeks before. This is generally recognised as the time when the Omicron variant of the virus became dominant in Norway.

The figure shows that not only did the forecasts with date of origin late in December and in early January have a higher average than earlier forecasts, but also that there were higher variations from one forecast to the next. This feature is explained by greater variation in the initial

conditions of the forecasts. The effects of Omicron and of the following non-pharmaceutical policy measures were also only gradually incorporated in the model, in particular in the light of the evidence of the forecast errors.

From 26 January 2022 people who had had a booster dose, or who had received 2 doses of coronavirus vaccine and had COVID-19 during the last 3 months, were no longer offered a confirmatory PCR test or rapid antigen test at a test centre after a positive self-test. Unquestionably, this affected the number of cases registered as the time series measured a lower proportion of those who were actually infected than before.

Nevertheless, the evolution of the number of new cases through February and March 2022 was quite well forecasted by CovidMod, at least judged by the plotted forecast graphs which are "on top of" the graph for actual number of new cases. Hence, important as it was, the structural break in the DGP of new cases that can be linked to the policy change of 26 January, was not in itself reason to step away from attempting to model the data and to forecast the series.

Figure 5 shows forecasts and actuals for hospital beds over the same period as in figure 4. 3-week forecasts from NIPH's National regional model were published throughout this period and are shown in the green line graphs in figure 5, while CovidMod forecasts are again shown in black line graphs. Note that the last observation of the time series for hospital beds (in red) is 22 March 2022 when the Norwegian Directorate of Health suddenly stopped the publication of those numbers.

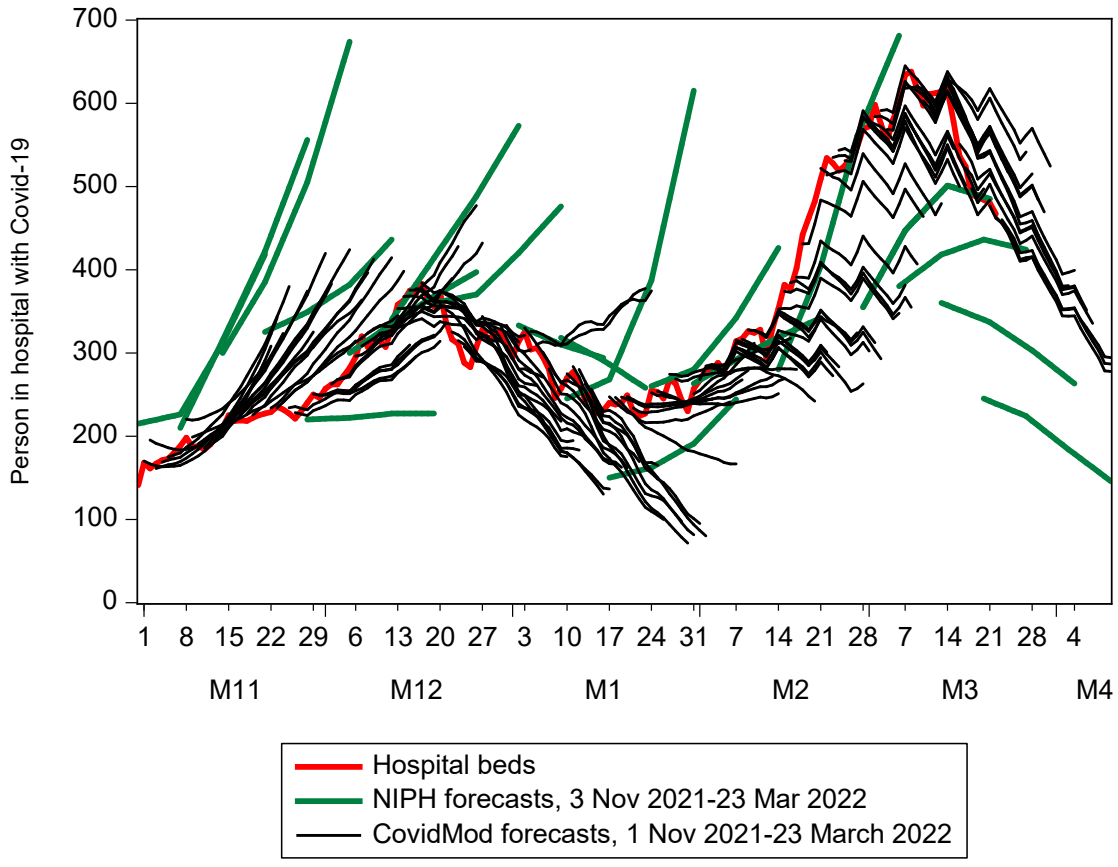

Figure 5: Forecasted number of hospital beds and forecasts from CovidMod and NIPH (National regional model).

The figure shows that both sets of forecasts have the "hedgehog" appearance which is familiar from economic forecasts. For example, the 14-21 day ahead CovidMod forecasts were systematically overshooting in November 2021, and undershooting in the development in mid-January and the first weeks of February. The same can be said about the NIPH forecasts. An additional remark that applies to those forecasts is that they sometimes got the level wrong "from the start", the latest four forecasts for example. It can be an indication that evaluation of initial conditions may have been an important part of the preparation of these forecasts.

The period of the pandemic covered by the previous graphs is interesting in itself, as it illustrates how the forecasts have adapted to a period of the pandemic characterised by a new and more infection version of the virus (Omicron) and a more open society but also a much higher degree of vaccination of the population.

However, it is also relevant to make formal comparison of forecast accuracy over a longer period of the pandemic. Norwegian Institute of Public Health (NPIH).

We next compare RMSFE (Root Mean Square Forecast Errors) from CovidMod and from the National and regional model of NPIH. The earliest forecasts have date of origin 19 March 2021 and the last originates from 1 December 2021. As mentioned above, this is when NIPH stopped the production of 21-day forecasts of "incidence" from their National and regional model.

For hospital beds the last forecasts are from 2 February 2022. Evaluation of later forecasts are not possible due to the stop in the publication of data for the number of persons in hospital with COVID-19. Hence, even though NPIH continued to publish forecasts of hospital beds, there are no forecast errors that can be evaluated.

All in all, this gave a sample of NIPH 33 forecasts errors for new cases and 44 for hospital beds. From the forecast record of CovidMod we have 178 forecast errors for new cases and 242 for hosital beds. As noted above, the CovidMod forecasts were prepared for each day in the working week, while NIPH forecasts were published weekly[8]

Table 1 shows the RMSFEs as functions of forecast horizon. For new cases, the RMSFEs for the NIPH's model are (more than) two times larger than for CovidMod, for all horizons. For hospital beds, the CovidMod forecasts again have the lowest RMSFEs. However, except for the shortest horizons, the NIPH model's RMSFEs are not significantly higher numerically speaking.[9]

Table 1: RMSFE for CovidMod forecasts with origin dates from 19 March 2021 to 1 December 2021 (new cases) and 2 February 2022 (hospital beds), and RMSFEs for NIPH forecasts with origins over the same period (n shows number of forecasts that the RMSFEs are based on).

	New cases		Hospital beds	
horizon	CovidMod ($n = 178$)	NIPH ($n = 33$)	CovidMod ($n = 242$)	NIPH ($n = 44$)
$h = 1$	163	897	7	50
$h = 7$	429	1398	31	60
$h = 14$	630	1836	62	80
$h = 21$	852	2401	96	108

Source: CovidMod, https://normetrics.no.
NIPH, Situational awareness and forecasting for Norway:
https://www.fhi.no/sv/smittsomme-sykdommer/corona/koronavirus-modellering/

[8]There are few examples of two weeks between forecasts.

[9]In an earlier assessment of hospital beds forecasts, Nymoen (2022) showed results which did not include data from 2022. Based on $n = 35$ forecast errors, CovidMod also then had the lowest RMSFs.

For new cases, we can make a RMSFE comparison with forecasts extracted from the multi-county COVID-19 forecasting project undertaken by Jennifer Castle, Jurgen Doornik, and David Hendry, at the University of Oxford.[10].

Those forecasts were for multiple countries and were produced by a automatised forecasting machinery developed by the research team, see Doornik et al. (2020); Castle et al. (2021b). The forecasted variables are confirmed cases and deaths. The forecast were short-term with a seven day horizon. The statistical forecasting device, dubbed Cardt, was highly adaptive and it rapidly adjusts to the latest information in the data, handling both stochastic trends and abrupt shifts, Castle et al. (2021a). It is based on decomposing the data into a trend, seasonal and irregular component, and then forecast the components separately before aggregating.[11], The research group in Oxford produced real time forecasts from 20 March 2020 to 16 June 2022. This gives a longer period of overlap with our forecasts than was the case for the forecasts made by NIPH. The Cardt forecasts were for the total number of confirmed cases and for deaths. However, because of the definition (6), we can obtain new cases from the Cardt forecasts as:

$$Y_t^{Cardt} = S_{Yt}^{Cardt} - S_{Y,t-1}^{Cardt} \tag{17}$$

where S_{Yt}^{Cardt} were the forecasts published on the internet. The difference between these forecasts and the NIPH data for new cases by date of registration gives the Cardt forcast errors for Norway. As noted above, CovidMod used data by date of specimen collection. That difference is however not large enough to undermine the comparability of the two sets of forecast errors.

Table 2: RMSFE for new cases, CovidMod and Norway forecasts from Castle, Doornik and Hendry multiple-countries and regions forecast. Forecast origins from 19 March to 1 April 2022 (CovidMod), and from 18 March 2021 to 6 April 2022 (Cardt).

Horizon	CovidMod ($n = 265$)	Cardt ($n = 83$)
$h = 1$	1658	5928
$h = 2$	1528	3485
$h = 3$	1385	2845
$h = 4$	1784	1890
$h = 5$	1829	2174
$h = 6$	2350	2662
$h = 7$	2661	2666

Source: Cardt: https://www.doornik.com/COVID-19/.
CovidMod: https://normetrics.no.

As shown in Table 2 the number of forecasts produced over the overlapping period was 265 CovidMod forecasts, and 83 Cardt forecasts. The forecast horizon is seven days in this table, determined by the short-term nature of the Cardt forecasts.

Compared to the new cases part of table 1 the scale of the RMSFEs in Table 2 is very much higher. This is because the new cases part of Table 1 had to be based on forecast errors from before the number of new cases started to re-scale (by a factor of 10 or more).

Table 2 shows that there are some differences between the RMSFEs of the two forecasting devices for really short horizons, one, two and three days ahead. However, there is in practice no difference between the RMSFEs for the seven days forecasts.

[10]https://www.doornik.com/COVID-19/

[11]Cardt is short for Calibrated Average of Rho, Delta and THIMA

As noted above, there is a break in the time series in the last week of January 2021, after it was decided that people who had had a booster dose (or who had received 2 doses of coronavirus vaccine and had COVID-19 during the last 3 months will) would no longer be offered a confirmatory PCR test. However, the break affect both the time series by the specimen collection date and the confirmed cases time series (forecasted by Cardt).

Extending the model with endogenous effects of policy interventions.

As noted above, future effects of policy responses was not modelled in the operative version of the model that was used to produce the daily forecasts. Instead, the practical method was to adjust the model ex-post, by estimating the effects of a policy change soon after it was included.

In a pandemic, interventions can take many forms. However, from the start of the coronavirus pandemic in the winter of 2020, and until the spring of 2020, policy interventions were consistently rationalised by the number of hospital beds occupied by patients with COVID-19 threatening to overwhelm the national health service system. This, it appears, represent a regularity in the data generating process that could made subject to econometric treatment.

Here, we present a way to include in the model the effects of policy responses. As noted, interventions can take many forms, and we only try to capture the effects of the policy actions that can be modelled by assuming that the policy variable is the number of hospital beds, and the deviation of that variable from a target level.

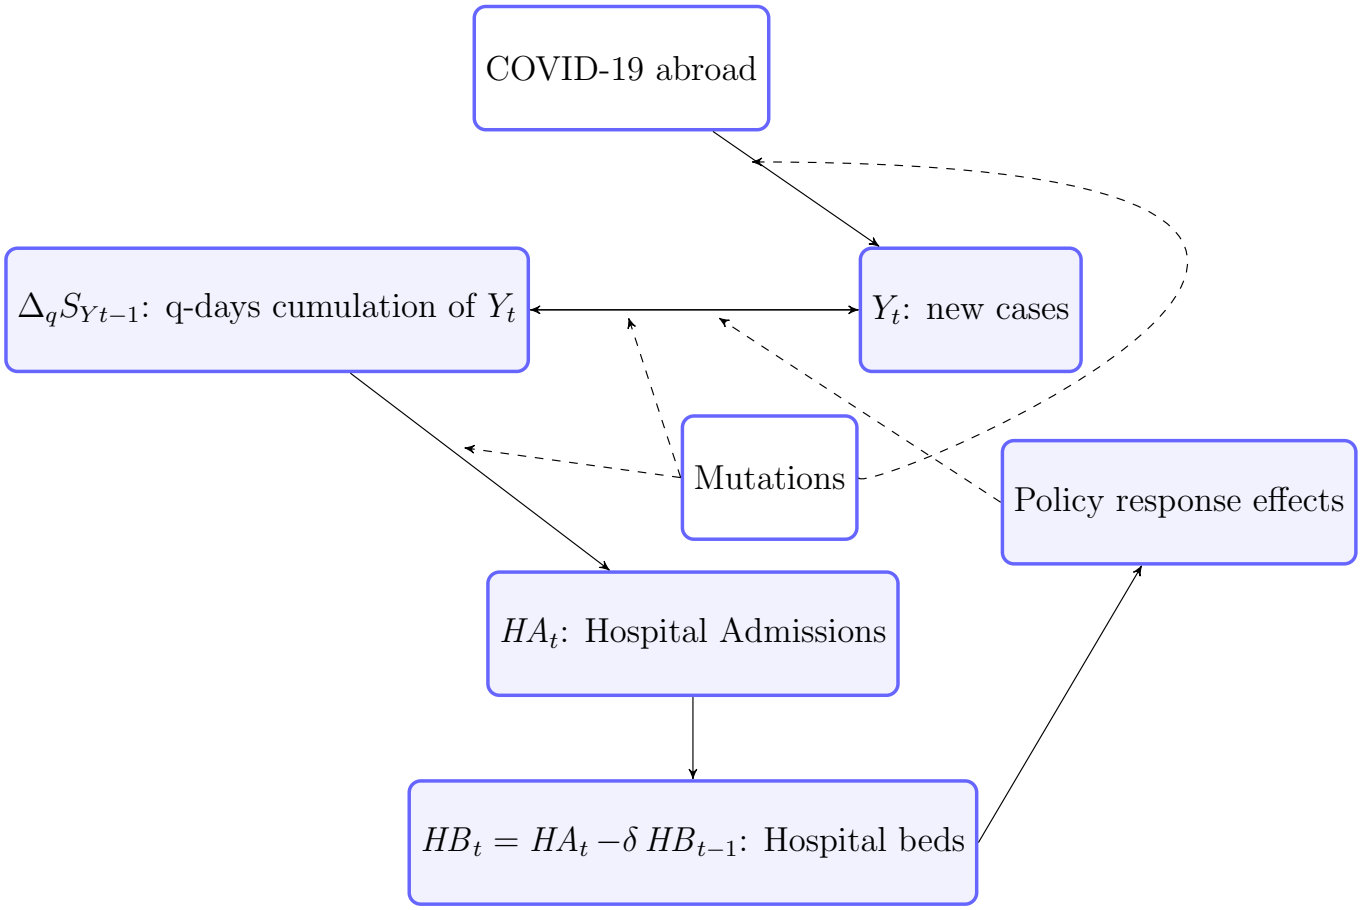

Figure 6: Variables and relationships of CovidMod-STR.

Figure 6 illustrates some of the differences between CovidMod, as illustrated by Figure 1

and the model with endogenous policy response effects. As just noted, this part of the model can be interpreted more broadly, as representing all policies that can be modelled as responses to the evolution of hospital beds with COVID-19 patients.

We now turn to some of the details about the implementation of the idea about endogenous policy response in the model, before we illustrate the forecasting potential of an fully specified and estimated version of the model.

Implementing endogenous policy responses

In modelling the effects of endogenous policy responses, we divide the breaks or the switching part of the model into exogenous regime changes D_t, like mutations, and effects of endogenous policy regime changes G_t, like lockdowns.

In order to illustrate the approach, consider the following stylized version of equation (10), where $X_t = (S_{Yt-1} - S_{Yt-p})$:

$$Y_t = \beta_0 + \beta_1 X_t + \beta_2 X_t D_t + \beta_3 X_t G_t + u_t \tag{18}$$

An exogenous event, or shock, makes the binary variable D_t change the effect of X_t on Y_t from β_1 to $\beta_1 + \beta_2$.

To model the endogenous policy effects G_t, we use the standard Smooth Transition Regression model (STR), see e.g. van Dijk et al. (2002) and Teräsvirta et al. (2010, Ch. 3.4). As is common, the transition function is the logistic specification

$$G\left(HB_{t-k}; \theta, HB^*\right) = \frac{1}{1 + \exp\left[-\theta\left(HB_{t-k} - HB^*\right)\right]}, \tag{19}$$

In our application, the transition function, that changes smoothly from 0 to 1 with an increasing policy target variable as argument. The policy variable is taken to be hospital beds HB_{t-k}, with a threshold value of HB^*.

As usual, equation (19) implies:
If $HB_{t-k} = HB^*$, then

$$G\left(HB_{t-k} - HB^*\right) = 0.5.$$

If $(HB_{t-k} - HB^*) \to -\infty$

$$G\left(HB_{t-k} - HB^*\right) \to 0,$$

and with $(HB_{t-k} - HB^*) \to \infty$

$$G\left(HB_{t-k} - HB^*\right) \to 1.$$

The steepness parameter of the transition function θ and the threshold value HB^* are both estimated.

With a very high θ, the effects of switching between zero or full policy effects can approximate a binary variable.

To continue the illustration: Imagine a contagious mutation represented by $\beta_2 D_t > 0$ in (18). The policy response might be a lockdown, bringing the infections down, with effect represented by $\beta_3 G_t < 0$. For illustration purposes, assume that $D_t = G_t = 1$. If then $\beta_3 = -\beta_2$, the effects of policy interventions mitigate the effects of the virus mutation.

To test out the suggested approach with endogenous policy effects, we estimated the model with data until 5 January, 2022, and evaluated the forecasting performance over the period 7 January, 2022 - 29 January, 2022.

The STR-model equation for new infections Y_t

The equation for Y_t is estimated as:

$$\begin{aligned}\hat{Y}_t &= \underset{(0.004)}{0.060}\,(S_{Y,t-1} - S_{Y,t-14}) \\ &+ \underset{(0.002)}{0.007}\,(S_{Y,t-1} - S_{Y,t-14})\,D_t \\ &- \underset{(0.004)}{0.007}\,(S_{Y,t-1} - S_{Y,t-14})\,G_t \\ &+ \sum\nolimits_{i=0}^{8} \hat{\alpha}_i^{\ddagger} \Delta Y_{t-1-i} + \text{ residual} \\ T &= 15.2.2020 - 5.1.2022,\ 691 \text{ obs.}\end{aligned} \tag{20}$$

where

$$\begin{aligned}D_t &= f(D_{It} + D_{\alpha t} + D_{\delta t} + D_{ot}) \\ G_t &= \frac{1}{1 + \exp\left[\underset{(0.032)}{-0.029}\left(HB_{t-5} - \underset{(48.575)}{295}\right)\right]}.\end{aligned}$$

HAC standard errors are reported in parentheses below the estimates.

The step dummies in $D_t = f(D_{It} + D_{\alpha t} + D_{\delta t} + D_{ot})$ represent exogenous breaks due to the occurrence of imported infections and the Alpha, Delta, and Omicron mutations, respectively:

$$\begin{aligned}D_{It} &= 1 \text{ from } 09.03.2020 \\ D_{\alpha t} &= 1 \text{ from } 02.03.2021 \\ D_{\delta t} &= 1 \text{ from } 16.07.2021 \\ D_{ot} &= 1 \text{ from } 28.12.2021\end{aligned}$$

The scale of incidence will be a function of the infection level in the population, which is however not observable. In equation (20) the change in the accumulated level of cases over a two-week period, $(S_{Yt-1} - S_{Yt-14})$ is used as a an indicator of the the infection level. For the policy effects function G_t, the threshold value $HB^{\star}$ is estimated to be 295 hospital beds.[12] Note that the effects of exogenous pandemic shocks D_t are mitigated by the endogenous policy effects G_t. The chosen specification indicates five days lag in the effects of policy.

Figure 7 shows the Hospitalisations HB_t, left scale, and the values of the smooth transition function G_t, right scale.

[12] In early 2020 the reported number of intensive care beds was 289 — see https://eurohealthobservatory.who.int/monitors/hsrm/all-updates/hsrm/norway/physical-infrastructure.

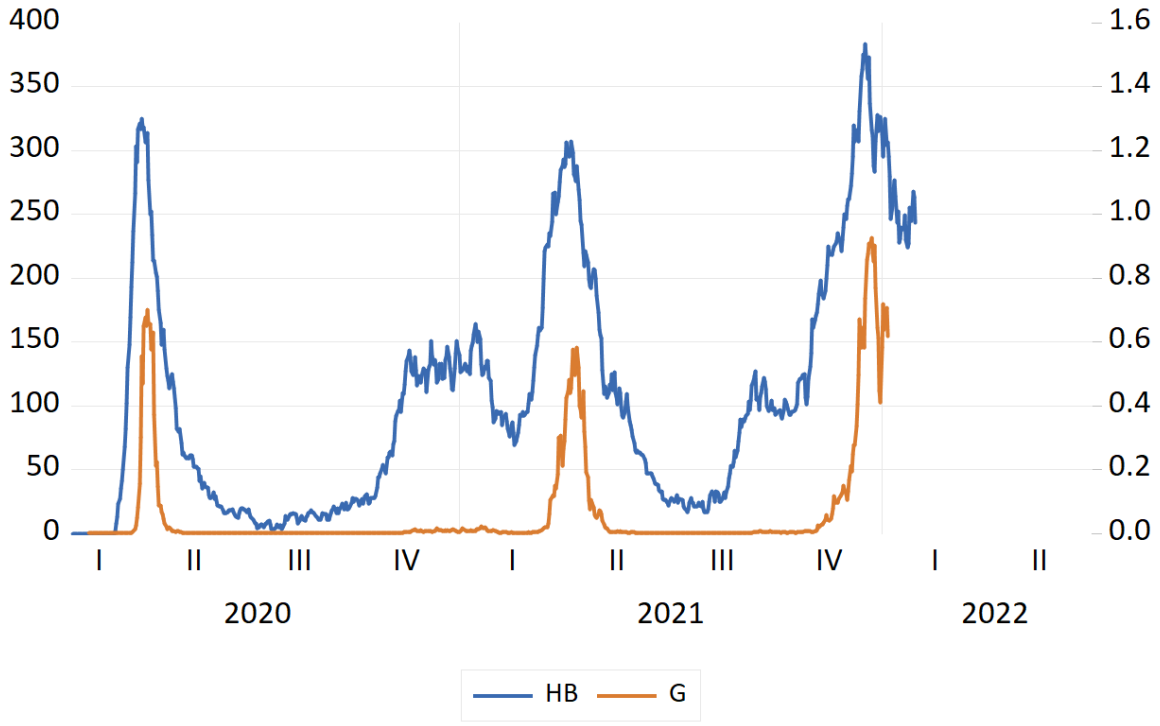

Figure 7: Hospitalisations HB_t (left scale) and smooth transition function G_t (right scale).

As shown in Figure 7, the policy effects, however, never completely counteract the exogenous shocks, since the G_t function never reaches the value 1,

Although not significant, the steepness parameter of the transition function θ is estimated to be close to 0.03, which indicates that the effects of the policy changes takes time to take effect.

To complete the model, the estimated versions of HA (hospital admissions) and HB (hospital beds) equations follows.

$$\begin{aligned}\widehat{HA}_t = &\underset{(0.0001)}{-0.0006}\,(S_{Y,t-3} - S_{Y,t-9})\,D_{o(t-6)} \\ &+ \underset{(0.001)}{0.012}Y_t \\ &- \underset{(0.001)}{0.007}Y_t \times D_{1,t} \\ &- \underset{(0.001)}{0.002}Y_t \times D_{2,t} \\ &+ \sum_{i=j}^{3} \alpha_j^{\ddagger\ddagger}\, HA_{t-j} + \text{residual}\end{aligned} \tag{21}$$

$$\widehat{HB}_t = HA_t + \underset{(0.004)}{0.901}\, HB_{t-1} \tag{22}$$

$$T = 14.7.2020 - 5.1.2022,\ 541 \text{ obs.}$$

$D_{1,t}$ = Step-dummy, 1 from 28 April 2021.

$D_{2,t}$ = Step-dummy, 1 from 15 November 2021.

Although qualitatively similar in terms of variables, the specification in (21 has a more complex structure of the effects of incidence, Y_t, than (14). A possible interpretation of the

dummies might be that more young people became infected, but with a lower frequency of hospital admittance.

Equation (22) is almost the same as (15).

Forecasting with the non-linear model, CovidMod-STR

In order to investigate the forecasting potential of the model with endogenous policy interventions, we simulated the model dynamically with 7 January 2022 as a the first forecast period and a (good) three weeks ahead (29 January, 2022 was the last solution period). The simulation did not condition on any exogenous variables in the simulation period, hence mimicking a real-time forecasting situation.

Figure 8 shows the three-week dynamic forecasts of the STR-version of CovidMod, with 7 January 2022 as the first day forecasted:

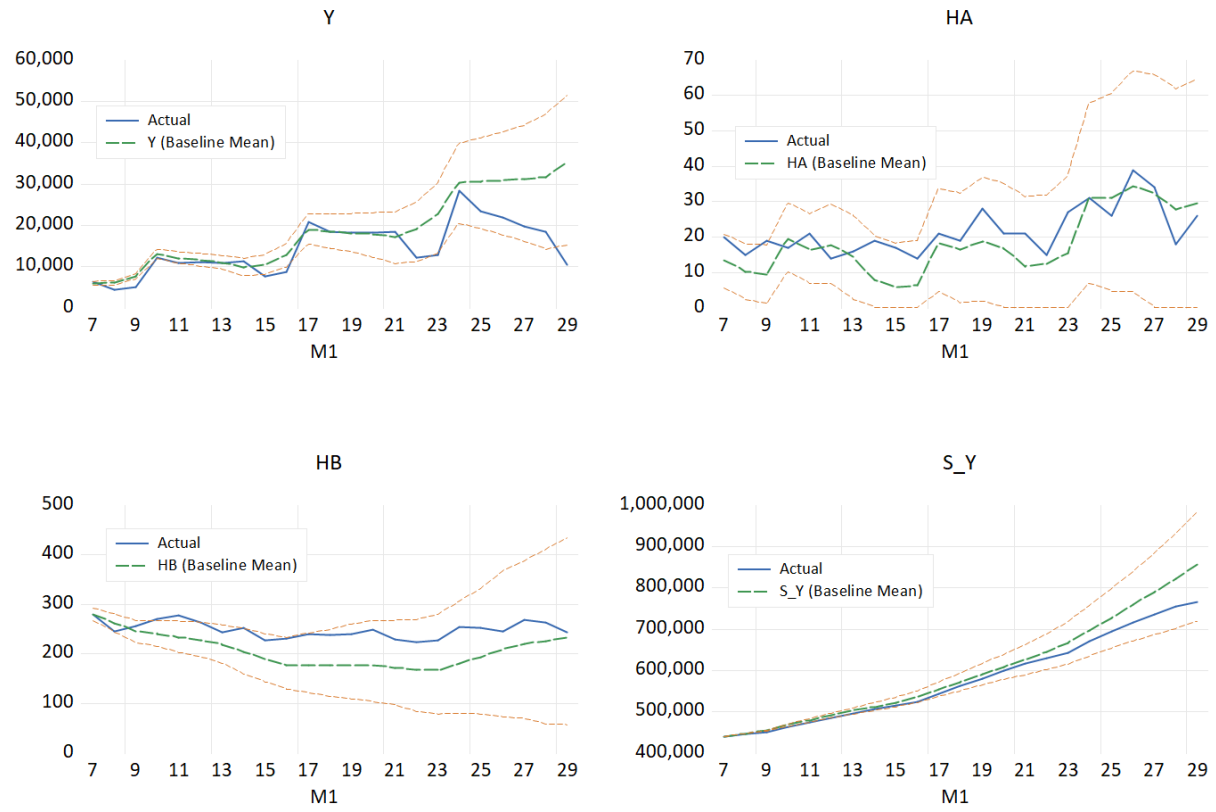

Figure 8: The CoviMod-STR forecasts and realizations. Line graphs represent actual values, dashed lines simulated values and dotted and lines represent upper and lower bounds of 90 % uncertainty intervals. The bound are simulated with parameter uncertainty.

As illustrated by the graphs in Figure 4, the pandemic had "changed gear" in the weeks before Christmas 2021, and the CovidMod forecast then struggled to adapt to the increased numbers of new cases and admissions to hospital. Figure 8 indicates that CovidMod-STR would have performed significantly better. Not only for new cases, but also for new hospital admissions (HA) and number of hospital beds. The accumulated number of new cases is also shown (in the fourth panel). Hence, taken together the estimation results for the STR-model for new cases, and the dynamic simulation, show that forecasting with a model that includes the effects of certain policies to contain the virus, is a feasible method.

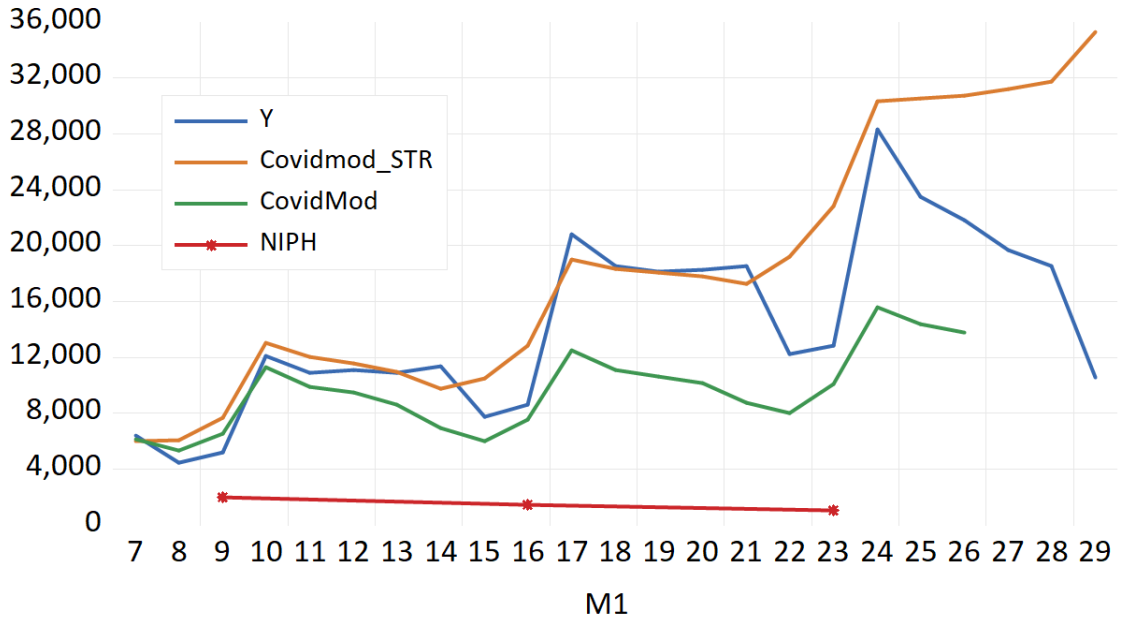

Figure 9: New cases (incidence). Source:"Situational awareness and forecasting for Norway". NIPH COVID-19 modelling team Week 1, 5 January 2022. Table 2. Linear interpolation between 7, 14 and 21 days ahead forecasts from forecast origin date.

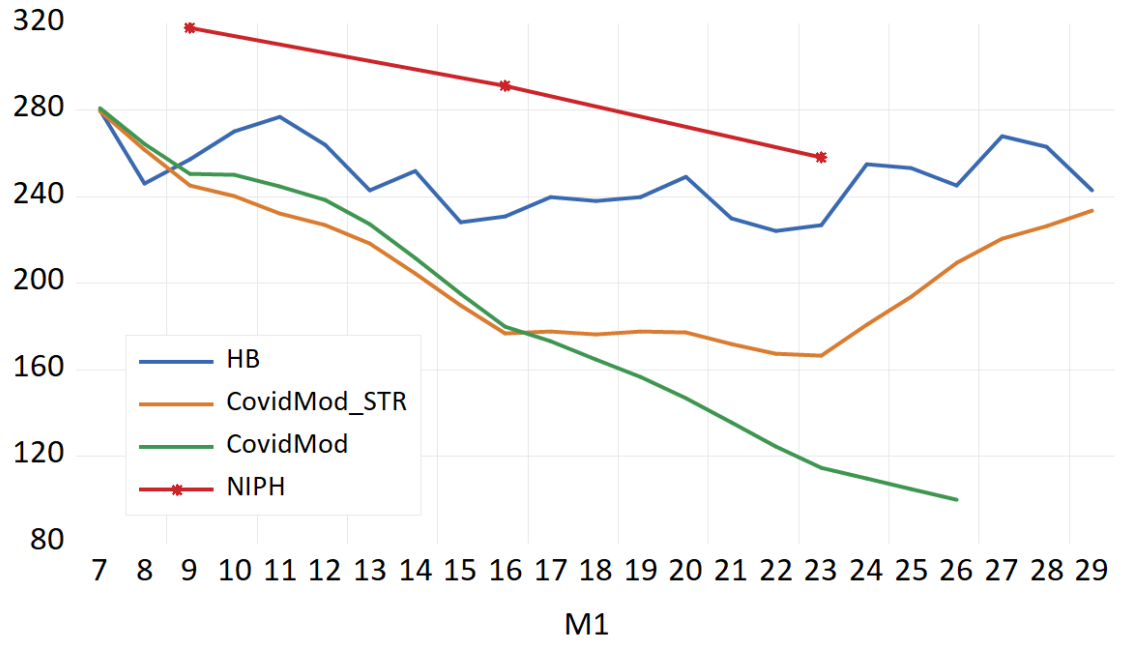

Figure 10: Hospitalisations. Source: "Situational awareness and forecasting for Norway" NIPH COVID-19 modelling team Week 1, 5 January 2022. Table 2 and Figure 5. Linear interpolation between 7, 14 and 21 days ahead forecasts from forecast origin date.

Figures 9 and 10 plot the CovidMod-STR (mean) forecasts for new cases and for hospital

beds, together with the NIPH and CovidMod forecasts from the first week of January 2022. The simulation results of the model with endogenous policy effects appear to be unbiased for new cases, for most of the period, while both CovidMod and NPIH under-forecasted the actual number of new cases in January.

When we look at the results for hospital beds, Figure 10, the picture is more complex. The two sets of CovidMod forecasts have biases (albeit insignificant when the uncertainty is taken into account), but we note that unlike the CovidMod forecasts, the forecasts of the STR version of the model recoup the level of the actual number of hospital beds at the end of the forecast period.

Table 3 shows RMSFEs for new cases and for hospital beds, for the non-linear version model (CovidMod-STR) and the other forecasting methods considered. The RMSFEs in Table 3 are for the one-off forecasts plotted in figures 9 and 10, ie., similar to having a sample of forecasts equal to one ($n = 1$) in Table 1. The table shows that compared to the (real-time) forecasts of

Table 3: RMSFEs for different forecasts with date of origin in the first week of January 2022

	New cases				Hospital beds		
horizon	CovidMod-STR	CovidMod	NIPH	Cardt	CovidMod-STR	CovidMod	NIPH
$h = 7$	1275	1327	7572	4047	37.7	19.8	53.2
$h = 14$	1743	4484	83221		45.2	50.2	52.9
$h = 21$	4586	7378	12557		49.5	106.7	46.3

Sources: See Tables 1 and 2.

NIPH, Situational awareness and forecasting for Norway:

CovidMod with the same data of origin, the RMSFEs for new cases increase less steeply with the forecasting horizon. For $h = 14$ and $h = 21$, the differences are quite large in favour of CovidMod-STR. As can be gleaned from the graph of the NPIH forecast in Figure 9, which was practically out of scale with the actual numbers of new cases in this period, the corresponding RMSFEs of were very large compared with those of CovidMod-STR in particular.

For completeness, the New cases part of the table also includes the RMSFE for the seven day Cardt-forecast with forecast origin in the first week of January 2022.

In the Hospital beds part of the table we see a somewhat different picture. CovidMod got the lowest RMSFE for the $h = 7$ horizon, while for $h = 14$ the RMSFE of CovidMod-STR was the lowest of the three methods. Interestingly, for $h = 21$ the RMSFE of the NIPH forecast was the lowest. This is as one could expect from the graphs in Figure 10, which show that although the NPIH forecast in question was way off the mark for the short-horizons, it nevertheless ended up quite close to the actual number of hospital beds three weeks after the production of the forecast. However, the RMSFE of CovidMod-STR was only marginally larger for the $h = 21$ horizon.

Discussion

From model based forecasting in economics, we are used to the fact that structural breaks in the economy is the main reason why forecasts become inaccurate and why downright forecast failures are not unusual.

Exogenous breaks that occur after the date of origin of a forecast are hard to predict. The best one often can do is to bring the breaks into the information set to be used for later forecasts, as soon as they have become detectable in the data. Which is why, in an operational context, it is important to carefully study the recent forecast errors.

Covid time series are also typical examples of changing data-generating processes, both because of mutations and policy responses. It is likely that forecasters who do not adjust their models, in response to evidence about such changes, will produce forecasts that are less accurate than they need to be.

In this paper we have presented a forecasting model that in an operational context allowed evidence of breaks to change the key parameters of the model. This then, we suggest, is the main explanation why, based on RMSFE comparisons, the real time forecast of CovidMod performed better than the NIPH forecasts, over the period investigated (from March and into 2022). It also indicates why the very adaptive Cardt method performs just as well as CovidMod in the forecasting av new cases over the 7-day horizons.

Another contribution that we make in the paper is to model effects of policy induced endogenous breaks as a function of a policy target variable by the use of a smooth transition function integrated in the system of equations. The example of forecasting with the non-linear model CovidMod-STR suggests that this approach seems to work in an operational context and might therefore be a useful tool in forecasting where policy reaction effects are relevant, like in pandemics.

References

Castle, J. L., M. P. Clements and D. F. Hendry (2021a). The value of robust statistical forecasts in the Covid-19 pandemic. *National Institute Economic Review*, *256*, 19–43. Https://doi.org/10.1017/nie.2021.9.

Castle, J. L., J. A. Doornik and D. F. Hendry (2021b). Forecasting principles from experience with forecasting competitions. *Forecasting*, *3*(2019-w01), 138–165.

Doornik, J. A., J. L. Castle and D. F. Hendry (2022). Short-term forecasting of the coronavirus pandemic. *International Journal of Forecasting*, *38*(2), 453–466.

Doornik, J. A., D. F. Hendry and J. L. Castle (2020). Statistical Short-term Forecasting of the COVID-19 Pandemic. *Journal of Clinical Immunology and Immunotherapy*, *6*(046).

Granger, C. W. J. and P. Newbold (1986). *Forecasting Economic Time Series*. Academic Press, San Diego.

Harvey, A. and P. Kattuman (2021). A farewell to R: time-series models for tracking and forecasting epidemics. *Journal of The Royal Society Interface*, *18*(182), 20210179.

Korolev, I. (2021). Identification and estimation of the SEIRD epidemic model for COVID-19. *Journal of Econometrics*, *220*(1), 63–85.

Li, S. and O. Linton (2021). When will the Covid-19 pandemic peak? *Journal of Econmetrics*, *220*(1), 130–157.

Mills, T. C. (2022). Modelling the link Between Covid-19 ccase, hospital admissions and deaths in England. *National Accounting Review*, *4*(1), 38–58. DOI: 10.3934/NAR.2022003.

Nymoen, R. (2019). *Dynamic Econometrics for Empirical Macroeconomic Modelling*. World Scientific, Boston.

Nymoen, R. (2022). Dynamisk modellering og framskrivning av nye smittede og innlagte med Covid-19 i Norge. *Samfunnsøkonomen*, *135*(1), 5–13.

Teräsvirta, T., D. Tjøstheim and C. W. J. Granger (2010). *Modelling Nonlinear Economic Time Series*. Oxford University Press.

van Dijk, D., T. Teräsvirta and P. H. Franses (2002). SMOOTH TRANSITION AUTOREGRESSIVE MODELS — A SURVEY OF RECENT DEVELOPMENTS. *Econometric Reviews*, *21*(1), 1–47.