

RESEARCH ARTICLE

Inferring the effective reproductive number from deterministic and semi-deterministic compartmental models using incidence and mobility data

Jair Andrade^{1*}, Jim Duggan²

1 Data Science Institute and School of Computer Science, National University of Ireland Galway, Ireland, **2** School of Computer Science, Ryan Institute and Data Science Institute, National University of Ireland Galway, Ireland

* jair.andrade@nuigalway.ie



OPEN ACCESS

Citation: Andrade J, Duggan J (2022) Inferring the effective reproductive number from deterministic and semi-deterministic compartmental models using incidence and mobility data. *PLoS Comput Biol* 18(6): e1010206. <https://doi.org/10.1371/journal.pcbi.1010206>

Editor: Benjamin Althouse, University of Washington, UNITED STATES

Received: October 6, 2021

Accepted: May 11, 2022

Published: June 27, 2022

Copyright: © 2022 Andrade, Duggan. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its [Supporting information](#) files. The code is available on GitHub at https://github.com/jandraor/time_varying_beta.

Funding: The project has received funding – through the School of Medicine, National University of Ireland Galway – from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 883285. The material presented and views expressed here are the responsibility of the author(s) only. The EU

Abstract

The effective reproduction number (\mathfrak{R}_t) is a theoretical indicator of the course of an infectious disease that allows policymakers to evaluate whether current or previous control efforts have been successful or whether additional interventions are necessary. This metric, however, cannot be directly observed and must be inferred from available data. One approach to obtaining such estimates is fitting compartmental models to incidence data. We can envision these dynamic models as the ensemble of structures that describe the disease’s natural history and individuals’ behavioural patterns. In the context of the response to the COVID-19 pandemic, the assumption of a constant transmission rate is rendered unrealistic, and it is critical to identify a mathematical formulation that accounts for changes in contact patterns. In this work, we leverage existing approaches to propose three complementary formulations that yield similar estimates for \mathfrak{R}_t based on data from Ireland’s first COVID-19 wave. We describe these Data Generating Processes (DGP) in terms of State-Space models. Two (DGP1 and DGP2) correspond to stochastic process models whose transmission rate is modelled as Brownian motion processes (Geometric and Cox-Ingersoll-Ross). These DGPs share a measurement model that accounts for incidence and transmission rates, where mobility data is assumed as a proxy of the transmission rate. We perform inference on these structures using Iterated Filtering and the Particle Filter. The final DGP (DGP3) is built from a pool of deterministic models that describe the transmission rate as information delays. We calibrate this pool of models to incidence reports using Hamiltonian Monte Carlo. By following this complementary approach, we assess the tradeoffs associated with each formulation and reflect on the benefits/risks of incorporating proxy data into the inference process. We anticipate this work will help evaluate the implications of choosing a particular formulation for the dynamics and observation of the time-varying transmission rate.

Commission takes no responsibility for any use made of the information set out. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: I have read the journal's policy and the authors of this manuscript have the following competing interests: 1. Prof. Jim Duggan is a member of the WHO Global Outbreak and Response Network (GOARN), through the involvement of the National University of Ireland Galway as a GOARN partner. (NGO) 2. Prof. Jim Duggan is a member of the Irish Epidemiological Modelling Advisory Group (IEMAG), and in this volunteering role, provides modelling advice to the Department of Health, Ireland. Also, the work presented is independent of (1) and has only benefited from an aggregated data set received from (2).

Author summary

Policymakers use the effective reproduction number (\mathcal{R}_t) to determine whether an epidemic is growing ($\mathcal{R}_t > 1$) or shrinking ($\mathcal{R}_t < 1$). One can estimate this quantity by simulating compartmental models fitted to data. These models can be seen as the ensemble of two structures: one that describes the course of a disease in an individual and another one that accounts for behavioural patterns. Nevertheless, these estimates are sensitive to the assumptions embedded in the model, such as the formulation of the time-varying transmission rate. In this paper, we couple an SEIR-type structure with three complementary formulations: 1) non-negative random-walks (Geometric Brownian Motion) 2) non-negative random-walks pulled toward a long-term goal (Cox-Ingersoll-Ross) 3) Gradual approximations towards a long-term goal (exponential smoothing). We refer to each coupling as a Data Generating Process (DGP). In essence, we simulate trajectories from these DGPs to identify plausible sets of transmission rates (based on incidence and mobility data) that explain Ireland's first COVID-19 wave. Here, we assume that mobility data is a proxy measurement for the transmission rate. These DGPs yield similar average estimates for \mathcal{R}_t , albeit with dissimilar degrees of uncertainty. Finally, we reflect on the tradeoffs of choosing each particular formulation.

Introduction

Since early 2020, SARS coronavirus 2 (SARS-CoV-2) has spread throughout the seven continents, causing a COVID-19 pandemic of catastrophic consequences, including the loss of millions of lives and jobs. In the early days of the pandemic, given the absence of vaccines and the lack of effective therapeutics, governments primarily relied on non-pharmaceutical interventions (NPIs) to reduce the transmission of SARS-CoV-2, thereby lowering the death toll. Although effective in preventing deaths [1], NPIs such as mobility restrictions and stay-at-home orders impose a burden on society with economic and psychological costs [2]. In addition to this, the effectiveness of these interventions wanes over time as compliance progressively diminishes. Following these considerations, policymakers strive to find an adequate balance between the interventions' severity and acceptable transmission levels. In this decision-making process, the effective reproduction number plays a crucial role. Briefly, the effective reproduction number, \mathcal{R}_t , is the time-varying average number of secondary cases caused by a primary case at a calendar time t [3, 4], and it is a theoretical indicator of the course of an infectious process [5]. Above the *epidemic threshold* ($\mathcal{R}_t > 1$), each infectious person leads to more than one secondary infectious person, and the disease is (re)emerging [6]; below that threshold, there is limited secondary transmission. In practice, policymakers can use \mathcal{R}_t in two ways. First, as a guide to assess in near real-time whether the interventions are succeeding ($\mathcal{R}_t < 1$) or whether it is required to increment the response's strength [4]. Second, in retrospective analyses to assess how policy decisions, population immunity, and other factors have impacted transmission at specific points in time [7].

Generally speaking, \mathcal{R}_t is the result of a combination of intrinsic (decline in susceptible individuals) and extrinsic (change in contact patterns due to the implementation of control measures) factors [4], for which there are no readily available measurements. One, therefore, must resort to statistical methods to obtain an approximation of this epidemic indicator. On one end of the spectrum, we find widely applicable and context-independent empirical methods such as the two-step Bayesian procedure proposed by Cori and colleagues [8, 9] and the likelihood-based estimation procedure proposed by Wallinga and Teunis [10]. At the other

end of the spectrum, we can infer \mathcal{R}_t from compartmental models calibrated to incidence data [11], which is the focus of this paper. In addition to serving as vehicles to obtain estimates, these mechanistic models are based on a scientific understanding of infectious disease dynamics [12], which one can interpret as a dynamic hypothesis of the underlying process that produces the observable behaviour patterns. This feature implies that fitting a compartmental model to data also tests a hypothesis that links structure to behaviour [13]. It thus follows that parameter estimates derived from this procedure have an interpretation in the real world. Notwithstanding these advantages, \mathcal{R}_t estimates from compartmental models are sensitive to data availability and assumptions in the model structure [7].

One such assumption is the transmission rate's dynamics. In the context of the COVID-19 pandemic, the assumption of a constant transmission rate is rendered unrealistic, apart from a few days in the initial phase of the outbreak [14, 15]. The rationale is that under the imminent surge of cases, governments implemented NPIs at early stages to reduce the number of contacts among the population. Modellers thus are required to describe formally the changes in the transmission rate over time. For instance, in measles studies [16–18], it is not unusual to assume *term-time forcing* structures [19], where the contact rate experiences sudden changes in time (e.g., because of school holidays). Other approaches have adopted *smoothly-varying functions* [19] to model the transmission rate in tuberculosis outbreaks [20]. In COVID-19 analyses, the transmission rate has been described as episodes of constant contact rates separated by change points where a transition occurs [14, 21]. These are likely once-off models, more appropriate for retrospective analyses, whose formulations are not designed to incorporate new data that account for policy changes (unless the structure is modified).

Nevertheless, ascertaining which deterministic formulation is the most adequate is far from straightforward. Its search involves several *trial-and-error* iterations and model comparisons until a satisfactory structure is found. If one aims for near real-time estimates, random-walk formulations offer a flexible device to *uncover* the underlying transmission rate dynamics [22]. This type of structure does not impose stringent constraints on the transmission's rate shape, facilitating the incorporation of new data without structural modifications. This approach has been applied to studying an influenza pandemic [22, 23] and Ebola outbreaks [24, 25]. Although random-walk models yield fits to incidence data, the match between observed and simulated data may be achieved at the expense of large uncertainty bounds. Moreover, under this framework, the inference of time-independent parameters requires burdensome computational efforts. More recently, the extensive research provoked by the COVID-19 pandemic prompted researchers to use non-traditional sources of data to infer the transmission rate. In particular, mobility data has been assumed as a proxy for the changes in the transmission rate [26]. In doing so, the dynamics exhibited by the transmission rate have an inherently plausible explanation (changes in human behaviour measured by mobile devices) so that models can more easily incorporate new incidence measurements. However, it should be mentioned that this approach entails a stringent assumption wherein one tacitly assumes a perfect correlation between changes in mobility data and effective contact patterns. Thus, discrepancies between actual and assumed transmission rates may result in unnecessary corrections to the estimates of other unknown parameters.

Consequently, this paper aims to draw upon the strengths of the approaches described above to formulate a complementary process for estimating \mathcal{R}_t from compartmental models. Specifically, we build three structures or Data Generating Processes (DGP) that accounts for Ireland's first COVID-19 wave. Two DGPs incorporate stochastic features in the transmission rate, whereas the other formulation is exclusively deterministic. These structures are complementary in the sense that the results obtained from one DGP inform the subsequent one. Below, we describe each DGP in detail, the inference process to obtain estimates for \mathcal{R}_t and

other unknown quantities (Fig 1), and finally, discuss the results. All the analysis is performed in R, mainly supported by the statistical packages *pomp* [27] and *Stan* [28]. The code is freely available at https://github.com/jandraor/time_varying_beta.

Results

Context

By the end of February 2020, more than sixty countries had detected at least one case of COVID-19 [29], including Ireland, where the first confirmed case was announced on *the 29th of February*. Twelve days after this event, the Irish Government ordered the closure of all schools, colleges, and childcare facilities, followed by a stricter *stay-at-home* mandate implemented on *the 27th of March*. These interventions resulted in low incidence and mortality rates, which allowed easing the restrictions from mid-May. In Fig 2A and 2B, respectively, we present the number of daily (y_d^1) and weekly cases (y_w^1) detected from the first report up to the point where the restrictions began to be lifted, a period that we refer to as the *first wave*.

In a nutshell, stay-at-home orders and similar measures aim to restrict the movements of a population so that the risk of exposure to a transmissible pathogen is reduced. Impractical several years ago, the advent of smartphones has permitted us to gauge patterns in population mobility in real-time. For instance, since *the 13th of January, 2020*, Apple has provided an index that quantifies the level of mobility by transportation type (driving, transit, and walking). Apple generates this data by counting the number of requests made to *Apple Maps* for directions. Fig 2C shows Ireland's daily driving mobility levels during the first wave (y_d^2), and Fig 2D, the value at the end of each week (y_w^2) from *the 29th of February 2020*. This dataset, along with the incidence reports (S1 Data), will serve as the basis to calibrate the proposed compartmental models below.

State-Space models

$$X_t \sim p_{X,t}^\theta(x_t|x_{t-1}) \quad (1)$$

$$Y_t \sim p_{Y,t}^\theta(y_t|x_t) \quad (2)$$

One can frame the inference process for compartmental structures following the terminology provided by state-space models (SSM) [30], also known as Partially observed Markov process models [31]. Through an SSM, one conceives a DGP as a generative probabilistic model that consists of two discrete-time Markovian mechanisms. The first mechanism (Eq 1) describes the evolution over time of the system's latent states (X), where X_t is drawn conditionally on the previous state of the latent process (X_{t-1}) according to the density $p_{X,t}^\theta(x_t|x_{t-1})$. Therefore, the DGP is a Markov chain [32], as the state of the latent variable at time t depends only on its previous state and the distribution from which it comes. In the literature, Eq 1 is often referred to as the *latent process model* [16] or the *system model* [33]. Intuitively, this formulation corresponds to the set of causal assumptions (*dynamic hypothesis*) that explains a phenomenon of interest in terms of states and transitions (rates). The process model may be defined in continuous or discrete time [31], but only its distribution at discrete times is considered ($X_t, X_{t+1}, X_{t+2}, \dots, X_{t+h}$), where $t \geq 1$ and h is an integer. For simplicity, we assume that X_0 is known.

In epidemiology, it is commonplace to represent the process model via compartmental structures in which individuals are categorised according to their infection status [34]. We refer to this categorisation as the *within-host profile*. Formally, one can employ a system of differential equations to build such compartmental models. The reader should recall that any

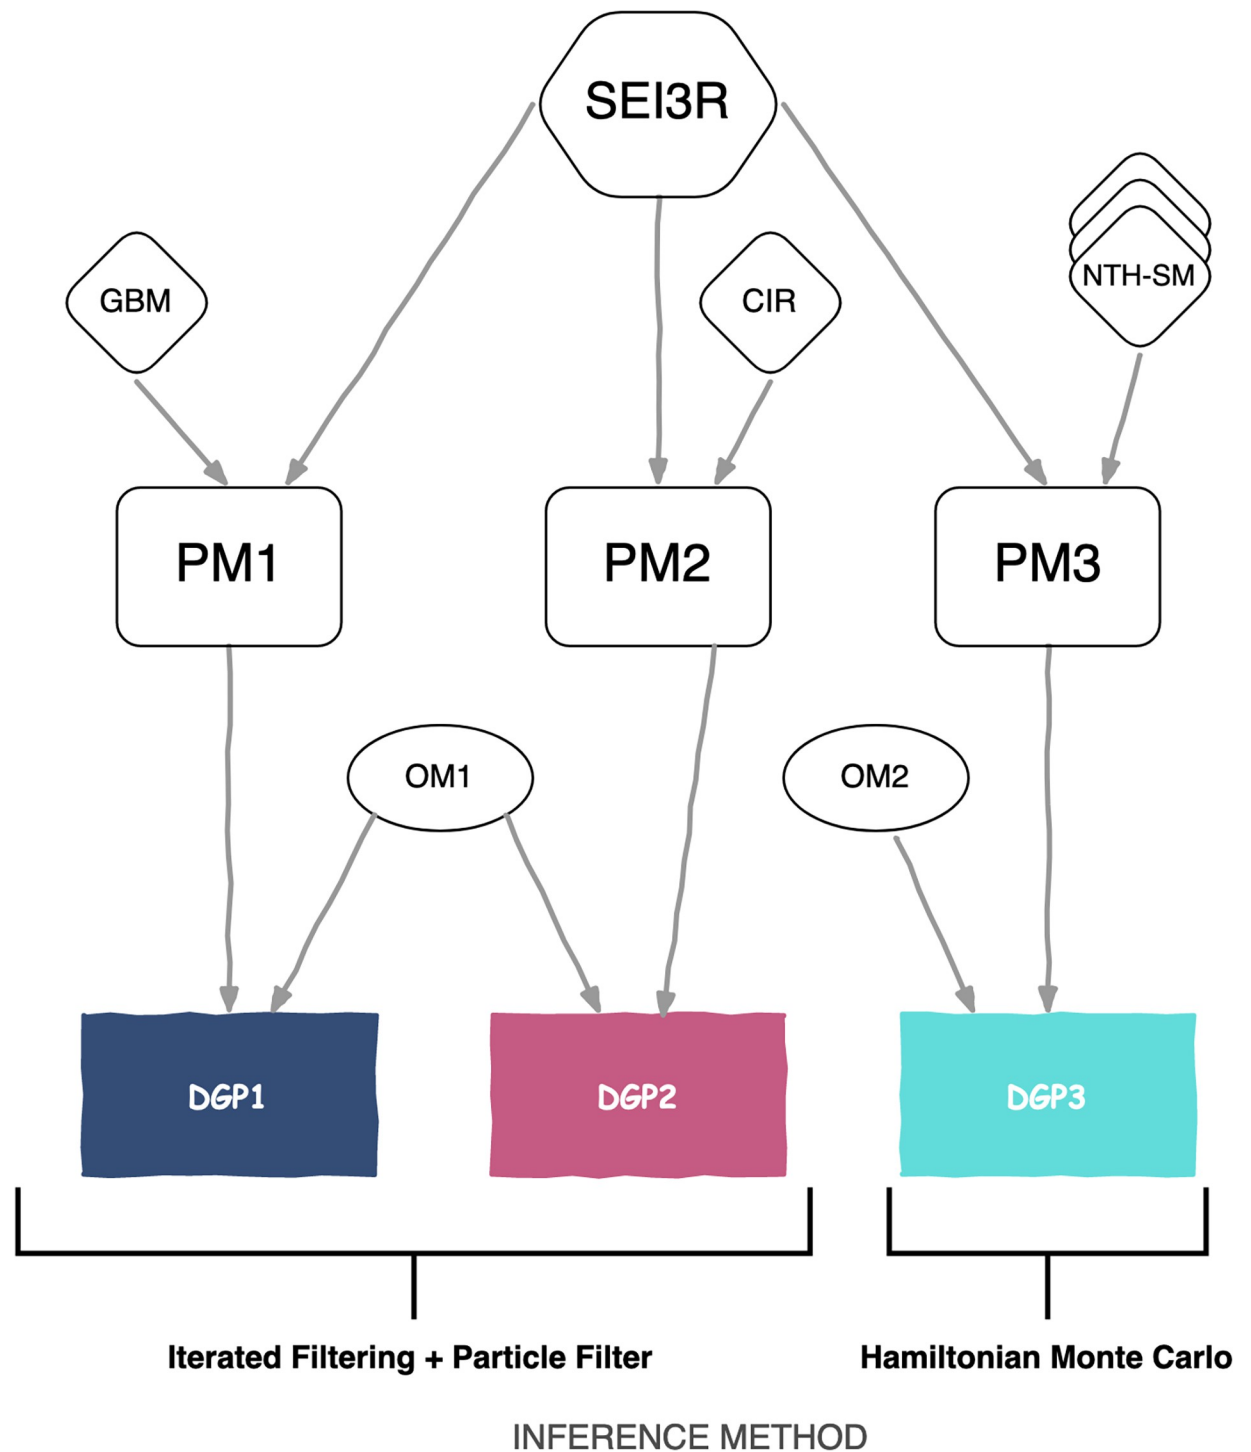


Fig 1. Schematic diagram of the data generating processes (DGPs) explored in this paper. This diagram aims to portray the DGPs as the ensemble of two components: a measurement or observational model (ellipse) and a process model (rounded rectangle). For instance, DGP1 is the amalgamation of the measurement model OM1 and the process model PM1. The process model is in turn the ensemble of two structures: a within-host profile (hexagon) and a time-dependent transmission rate (rhombus). Whereas all process models share a common within-host profile (SEI3R), they differ in the formulation of the transmission rate: Geometric Brownian Motion (GBM), Cox-Ingersoll-Ross (CIR), and n th-order exponential smoothing (NTH-SM). The inference method employed on each DGP depends upon the nature of the process model (Iterated Filtering + Particle Filter for stochastic structures and Hamiltonian Monte Carlo for deterministic ones).

<https://doi.org/10.1371/journal.pcbi.1010206.g001>

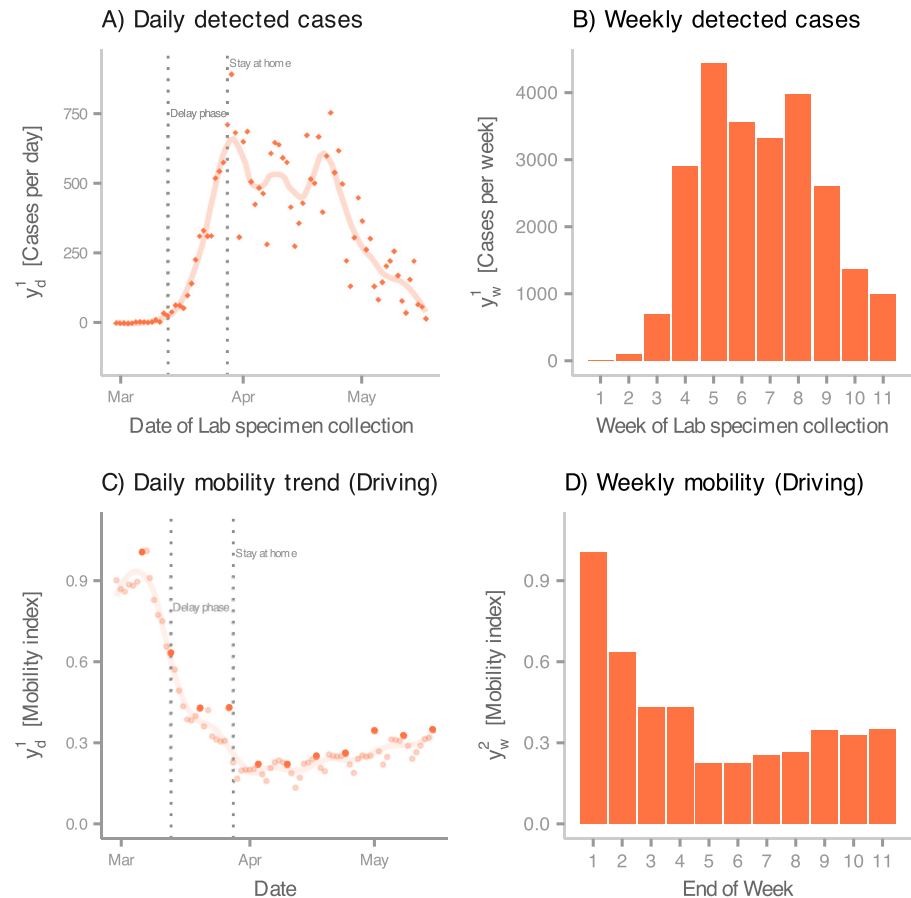


Fig 2. Incidence and mobility data. (A) Daily number (rhombus-shaped points) of COVID-19 cases detected during Ireland's first wave, from the 29th of February 2020 to the 17th of May 2020. The x-axis indicates the date in which the infected individuals were swabbed. The line represents the smoothed trend (via LOESS method) from the data. (B) Weekly number of COVID-19 cases detected in during Ireland's first wave. The x-axis indicates the number of weeks since the first case was detected. (C) Apple data for Ireland from the 29th of February 2020 to the 17th of May 2020. Points represent the normalised amount of daily requests for driving directions. These indexes are normalised to the value on the 28th of February 2020. We highlight points every 7 days. These highlighted points are used to calibrate DGP1 and DGP2. The line represents the smoothed trend (via LOESS method) from the data. (D) Normalised amount of daily requests for driving directions at the end of each week starting from the 29th of February 2020. These bars correspond to the highlighted points in C.

<https://doi.org/10.1371/journal.pcbi.1010206.g002>

system of ordinary differential equations $\frac{dx}{dt} = f(x)$ is Markovian. Here, we adopt the *SEI3R* profile [15, 35], an extension of the SEIR framework. Under this profile, we stratify individuals as susceptible (S_t), exposed (E_t), infectious, and recovered (R_t). We further disaggregate the infectious class by medical status, resulting in three compartments: *preclinical* (P_t), *clinical* (I_t), and *subclinical* (A_t) (see [Materials and methods](#) section for the complete description). The three DGPs presented in this paper share the *SEI3R* profile (Fig 1).

On the other hand, the exact state of the population at any given time is generally not observable and must be inferred from available data via statistical inference [36]. It is thus necessary to formally relate (Eq 2), at each discrete time ($t \geq 0$), latent states to *noisy* measurements via a *measurement* or *observational* model [33], where each Y_t is drawn conditionally on the most recent state of the latent variable, according to the density $p_{Y,t}^{\theta}(y_t|x_t)$. This work draws on incidence and mobility data to formulate such measurement models.

DGP1—Geometric Brownian Motion

$$\beta_t = \zeta Z_t \quad (3)$$

$$\frac{dZ}{dt} = \alpha Z_t dW \quad (4)$$

$$dW \sim Normal(0, \sqrt{dt}) \quad (5)$$

Thus far, we have not yet defined the time-varying effective contact rate or transmission rate (β_t). When defined, this component is integrated with the SEI3R profile to form a process model (Fig 1). For this and the other two DGPs, we formulate β_t as the product of two components (Eq 3). Here, ζ denotes the transmission rate's initial value. Namely, $\beta_0 = \zeta$. From this definition, it follows that Z_t represents the transmission rate's change over time relative to its initial value, where $Z_0 = 1$. In relation to Z_t dynamics, we initially opt for a flexible approach to build this first process model (PM1). Specifically, we define $\frac{dZ}{dt}$ in terms of Geometric Brownian Motion (GBM) with no *drift* (Eqs 4 and 5), an approach adopted in previous studies of influenza and Ebola [22–25]. This stochastic structure is a model for the change in a random process, dZ_t , in relation to the current value, Z_t , where the proportional change $\frac{dZ}{Z_t}$ follows Brownian motion [37]. That is, normal distributed random *jumps* (dW) moderated by a volatility parameter (α). We do not imply that the actual transmission rate follows a random walk. In fact, the expected value of Z_t is constant over time (Z_0); strictly speaking, a *martingale* [37]. In practice, however, we use this structure as a *scaffold* to obtain some idea of the non-linear structure of the process without committing to a particular form of non-linear model [38]. This procedure resembles the use of smoothing splines to estimate coefficients that are allowed to vary as smooth functions of other variables [39]. Although not a requirement for this work, smoothing splines also have a Bayesian interpretation under certain conditions [40]. In particular, we use the GBM structure to generate non-negative random walks from an initial value (Fig 3A displays a set of possible trajectories). The main benefit from random walks is that at each time t , we propose several possible paths that the transmission rate may take and then use the available data to determine their plausibility [22]. In doing so, we *unravel* the dynamics of the effective contact rate. Formally put, we approximate $p(x_t|y_{0:t})$, the *filtering distribution* [30] (see [Material and methods](#)).

As noted above, the measurement model is the link between the process model and the data whereby one quantifies (through likelihood densities) the relative consistency of each set of parameter values, or model configuration, with observations. This quantification allows us to perform inference on time-varying and time-independent parameters. Thus, any misspecification in the measurement formulation can lead to overly confident conclusions [12] or biased estimations. In light of its importance, we prevent the consequences of model misspecification by proposing and testing six candidates that account for the incidence data (y^1). Moreover, a subset of these candidates incorporate mobility data (y^2), assuming that this dataset is a proxy observation of the relative contact rate (Z_t).

Before defining each candidate, we clarify certain assumptions regarding the available datasets. On the one hand, for the incidence data (y^1), we posit that *actual* periodic (daily or weekly) symptomatic COVID-19 cases (C_t) stem from the transition $P_t \rightarrow I_t$. Our assumption implies that individuals seek the healthcare system for a diagnostic test as soon as they develop symptoms. Furthermore, under this formulation, it is implicit that *underreporting* is due to the non-identification of asymptomatic cases. On the other hand, for mobility data (y^2), we emphasise its proxy nature. Contrary to the incidence time series, it is not anticipated that

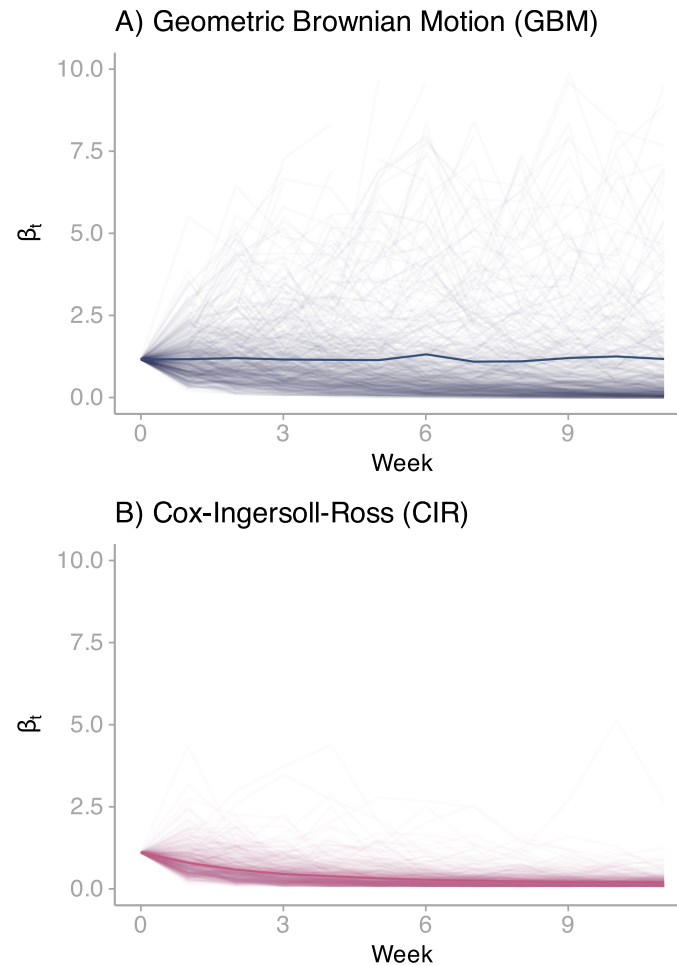


Fig 3. Brownian motion trajectories. (A) 200 simulations from a transmission rate described in terms of Geometric Brownian Motion. We generate these simulations from DGP1's Maximum Likelihood Estimate (MLE) using the Euler-Maruyama algorithm. The highlighted trend corresponds to the mean trajectory path. (B) 200 simulations from a transmission rate described in terms of the Cox-Ingersoll-Ross model. We generate these simulations from DGP2's Maximum Likelihood Estimate (MLE) using the Euler-Maruyama algorithm. The highlighted trend corresponds to the mean trajectory path.

<https://doi.org/10.1371/journal.pcbi.1010206.g003>

models yield faithful replications of Apple's mobility indexes. Should that be the case, we would have included this data (directly or parametrically) in the process model. However, we refrain from doing so as we deem there may be instances where the two elements are not strongly correlated. To illustrate this point, we consider the case in which the government relaxes social distancing mandates and individuals adopt a mask-wearing behaviour. Under these circumstances, the resulting increase in mobility and social contacts due to relaxed rules do not necessarily entail an equivalent effect on the effective contact rate given that individuals properly wear face coverings during their interactions. Hence, rigid structures in the process model may lead to unrealistic corrections in other parameters. As opposed to such inflexibility, we expect that the mobility data acts as a *nudge* on the transmission rate, guiding the latter towards the former only when plausible. In light of these considerations, for candidates 1 and 2, we formulate the observation of daily symptomatic COVID-19 cases (y_d^1) as independent Poisson and Negative Binomial counts, respectively. Then, we add an observational mechanism that relates Apple's daily driving data (y_d^2) to the transmission rate's relative level (Z_t),

Table 1. Measurement model candidates.

Id	Frequency	Incidence	Mobility	Converges	Fits incidence
1	Daily	Pois	No	No	N/A
2	Daily	Nbin	No	Yes	Yes
3	Daily	Pois	Yes	No	N/A
4	Daily	Nbin	Yes	Yes	No
5	Weekly	Pois	No	Yes	Yes
6	Weekly	Pois	Yes	Yes	Yes

<https://doi.org/10.1371/journal.pcbi.1010206.t001>

yielding candidates 3 and 4. Finally, even though King and colleagues [41] recommend that “models should be fit to raw, disaggregated data whenever possible and never to temporally accumulated data”, on candidates 1 and 3, we modify their periodicity from daily to weekly measurements, resulting in candidates 5 and 6. It should be noted that the use of weekly measurements has been performed previously in similar studies [22–25]. We refer the reader to [S1 Text](#) for the complete set of equations.

Having defined process and measurement structures, we proceed to the inference stage ([Table 1](#) summarises the results). Since non-linear SSM do not allow closed-form solutions [30] to calculate likelihood values, we must resort to simulation-based approaches such as *Sequential Monte Carlo*, also known as the *Particle Filter*. Naturally, these estimates must be robust so as to guide the inference process. By robustness, we refer to the quality that the Particle Filter returns similar likelihood values for various runs from a single model configuration. Furthermore, as with any Monte Carlo approach, it is expected that as the number of samples tends to infinity, the likelihood error (among various runs) converges to zero. To test this feature, we run the Particle Filter using the R package *pomp*, which implements the Sequence Importance Sampling algorithm [42]. In particular, through these runs, we evaluate likelihood estimates for each model candidate by varying the number of particles (samples), the integration step size, and the model configuration (see the complete analysis in [S1 Text](#)). The results indicate that measurement models that account for daily incidence observations as Poisson counts lead to unstable estimates. This finding suggests model misspecification in candidates 1 and 3, which are discarded from the pool.

To the remaining candidates, we estimate their latent states. Given its strength to infer time-varying random variables in the state space, the Particle Filter is also appropriate to numerically approximate (via samples) filtering distributions [33]. Nevertheless, drawing relevant samples requires plausible fixed-parameter values. Here, we assume that three parameters in PM1 are unknown: the effective contact rate at time 0 (ζ), the initial value of preclinical individuals (P_0), and the volatility parameter (α). Moreover, additional parameters may be required depending upon the specific measurement model. To infer such parameters, we employ the *Iterated Filtering* algorithm [31, 43]. This Maximum Likelihood estimation method has been designed to perform statistical inference on SSM and has been widely used to study infectious disease models [16, 17, 31, 41, 44]. Briefly, Maximum likelihood via Iterated Filtering (MIF) is a modified version of the Particle Filter, in which a sequence of filtering operations converges to the Maximum Likelihood Estimate (MLE). The key feature in this procedure is the set of stochastic perturbations applied to the unknown parameters in between the sequence of filtering operations, resulting in the selection of plausible parameter values in the light of the available data. Furthermore, the synergy between MIF and the Particle Filter permits us to calculate uncertainty bounds around the MLE. In particular, we use the *Profile Likelihood* method [45] and its refined version, the *Monte Carlo-adjusted profile* [46]. Ultimately, all of this information facilitates the construction of the parameters’ likelihood surface.

For each model, we leverage its likelihood surface to draw sets of point estimates from the neighbourhood surrounding the MLE [41]. These draws are subsequently plugged into the Particle Filter. In addition to likelihood estimates, *pomp* returns, for every run, a set of samples representing the filtering distribution at each time t . Then, we assign a weight to each run based on its relative likelihood. In doing so, we account for parameter uncertainty in the results. Finally, we summarise the results by computing weighted averages on the samples. This procedure allows us to calculate the uncertainty in the predicted latent states by the filtering distribution. The reader can find the complete set of results in S2–S5 Text.

The inference process on Candidate 2 (see S2 Text) reveals that this model yields a filtering distribution that fits the observed daily incidence. Interestingly, although Candidate 2's measurement model does not incorporate mobility data in its structure, the predicted relative contact rate captures the observed mobility indexes, albeit with a large degree of uncertainty. This finding supports the argument that such a dataset could be an adequate proxy for the relative contact rate. Then, one would logically expect that incorporating Apple's data into the measurement model (as we did for Candidate 4) would diminish the resulting uncertainty in the filtering distribution. However, the results (see S3 Text) show that the enhanced fit on the effective contact rate stems from unrealistic corrections to the predicted incidence, rendering Candidate 4 unreliable. On the other hand, we notice that Candidate 5's filtering distribution and parameter estimates convey similar insights to those of Candidate 2 (see S8 Text Section 1). Therefore, the change in periodicity does not result in severe loss of information. Yet more important, the crucial feature of the weekly formulation is that it allows integrating mobility data seamlessly into the measurement model (Candidate 6). This integration is accomplished without compromising the prediction on incidence counts and simultaneously reducing the uncertainty in the relative contact rate's fit. This behaviour differs from the unrealistic fit achieved by Candidate 4. We ascribe the resulting harmony between the two datasets to the stringency imposed by the Poisson distribution, which implicitly prioritises incidence counts over mobility indexes. In consequence, we select Candidate 6's measurement model (Eqs 6–8) as the structure (OM1) that completes DGP1's formulation (Fig 1).

$$\frac{dC}{dt} = \eta P_t - C_t \delta(t \bmod 7) \quad (6)$$

$$y_w^1 \sim \text{Pois}(C_t) \quad (7)$$

$$y_w^2 \sim \text{Normal}(Z_t, \tau) \quad (8)$$

Fig 4A presents a comparison between the predicted number of weekly symptomatic cases from DGP1 and observed incidence. Notice that this is a contrast between measurements (y_w^1) and a latent state (C_t). Although this approach is not generally applicable (comparing measurements to predicted latent states), in this case it is valid given that C_t corresponds to the mean of the measurement model (Eq 7). In Fig 4A, it can be seen that this model's filtering distribution captures the actual values in regions of high plausibility, thus yielding an accurate fit. This result helps us gain confidence in the model's structure as an adequate dynamic hypothesis to the studied phenomenon, considering that it can reproduce the observed behaviour [13]. Similarly, the estimated relative effective contact rate replicates to a large extent its assumed measurement values (Fig 4B). As expected, the filtering distribution does not capture all of the measurements (Weeks 9–11), given the proxy nature of the data. However, these results allow us to elucidate the trajectory of the effective contact rate, and in turn, the effective reproductive number (see Material and methods for the estimation of this quantity). It must be remarked that in the early

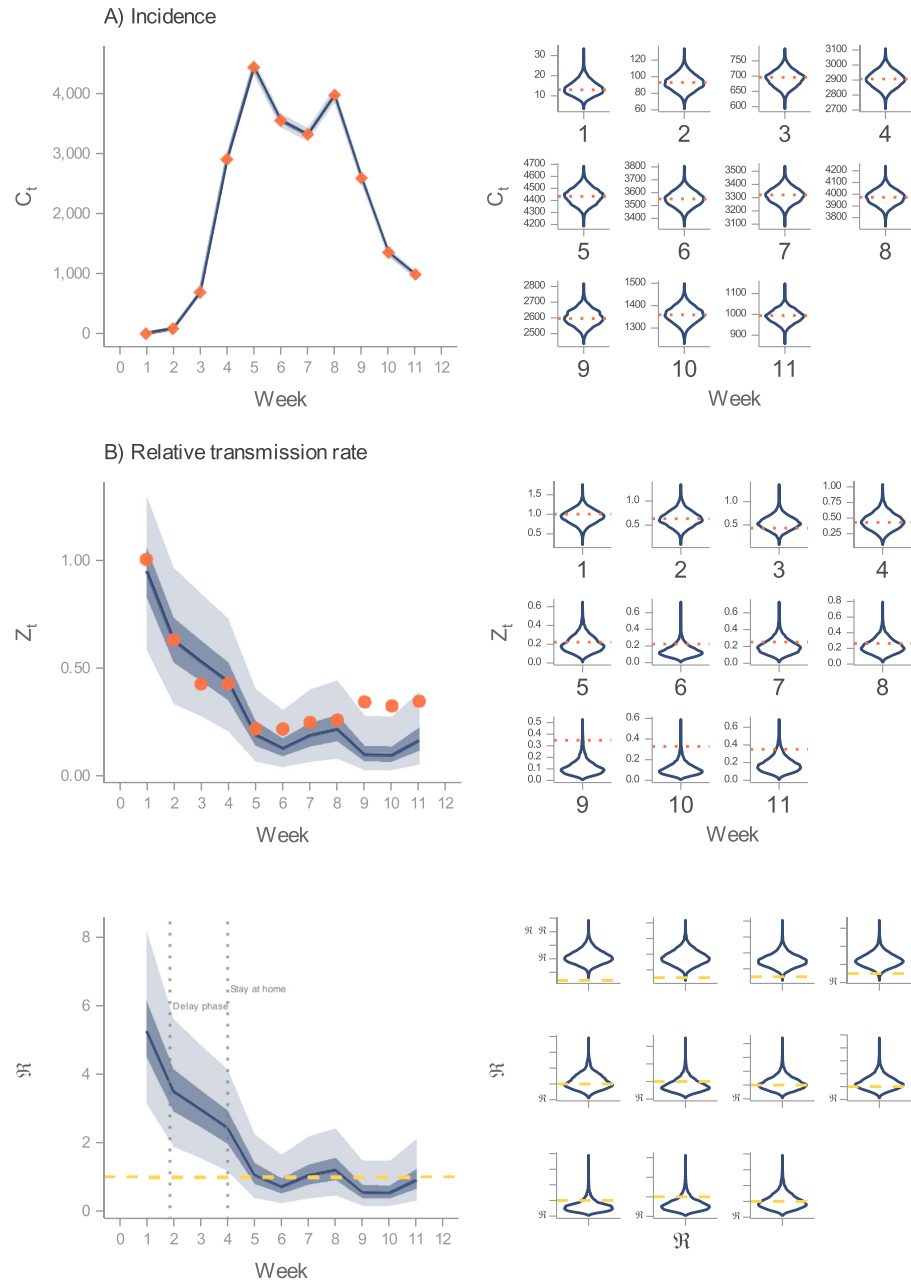


Fig 4. Inference on DGP1. In these three figures, the predicted values stem from DGP1's filtering distribution. Further, in the LHS, the solid line indicates the median, and the darker and lighter ribbons represent the 50% and 95% CI, respectively. (A) Comparison between the predicted incidence (solid line and ribbons in the LHS; violin plots in the RHS) and weekly detected cases from Ireland's first COVID-19 wave (rhombi in the LHS; horizontal dotted lines in the RHS). (B) Comparison between the predicted relative transmission rate (solid line and ribbons in the LHS; violin plots in the RHS) compared to Apple's mobility indexes in Ireland (points in the LHS; horizontal dotted lines in the RHS). (C) Predicted effective reproduction number (solid line and ribbons in the LHS; violin plots in the RHS). Horizontal dashed lines denote the epidemics threshold.

<https://doi.org/10.1371/journal.pcbi.1010206.g004>

stages of this outbreak, the dynamics of the transmission rate determined the level of \mathcal{R}_t . This characteristic occurs when the susceptible fraction is close to one, as was the case during the first wave [47]. In Fig 4C, we present the estimated \mathcal{R}_t , where it can be observed that the behaviour change (presumably caused by mobility restrictions and people's awareness) led to an \mathcal{R}_t close to or below the epidemics threshold, bringing about a lowering of the incidence rate.

DPG2—Cox-Ingersoll-Ross

$$\frac{dZ}{dt} = \nu(v - Z_t) + \sqrt{\alpha}Z_t dW \quad (9)$$

The dynamics of the transmission rate (Fig 4B) uncovered by DGP1 exhibit a compelling pattern. The transmission rate gradually decays for several weeks from its initial value until it levels off around a determined value. In other words, a pattern that resembles *goal-seeking behaviour* [48]. Based on this recognition, we formulate the relative transmission rate in terms of the Cox-Ingersoll-Ross (CIR) model [49]. This formulation (Eq 9) is a compromise between the rigidity of a deterministic structure and the flexibility offered by random walks. Under this structure, the randomly-moving quantity of interest (Z_t) is elastically pulled toward a central location or long-term goal, v . The strictly positive parameter ν determines the speed of adjustment. In practice, we can interpret the long-term goal as the minimum level of mobility that the restrictions can achieve and the adjustment parameter as the rate at which individuals adopt such mandates. Hence, inferring these parameters permit the characterisation of the implemented interventions, a piece of information that cannot be estimated from DGP1. The randomness in this process stems from the diffusion process (second term). That is, stochastic variations from the deterministic trend. More importantly, unlike those in the *Vasicek* and *Ornstein-Uhlenbeck* structures, this particular diffusion process precludes negative values [37], a *sine qua non* to describe transmission rates. Logically, we ensemble this structure with the SEI3R profile to build the process model (PM2). As with DGP1, we assess the convergence of likelihood estimates obtained from the amalgamation of PM2 and the previously defined six measurement model candidates (see S1 Text Section 3). The results reveal an identical pattern to that observed in DGP1. Therefore, it is warranted to integrate PM2 and OM1 (Fig 1) to form a DGP that we refer to as DGP2. In Fig 3B, we present simulated trajectories from this DGP, obtained from a single set of parameters (MLE).

The main objective for building DGP2 is to estimate its latent states conditional on the available data. To do so, we repeat the process applied to DGP1. Specifically, we first perform parameter inference and construct DGP2's likelihood surface using MIF and the Particle Filter. The next step consists of drawing samples from the MLE's neighbourhood to plug them into the Particle Filter. There is a slight alteration in this process, however. Previously, we selected parameter combinations that yielded likelihood values near the MLE to construct DGP1's neighbourhood. We then identified the bounds of these parameters to construct a four-dimensional hypercube. From this object, we obtained independent and uniformly distributed samples for each parameter. In light of DGP2's complex parameter space, we opt for a copula [50] instead of a hypercube. The copula is a multivariate cumulative distribution for which the marginal probability distribution of each variable is uniform, but there is dependence (correlation) among the random variables (unknown parameters). In doing so, we mitigate biases caused by point estimates that yield abnormal likelihood values. The reader can find the complete set of results in S6 Text.

Fig 5 displays the results obtained from the inference process carried out on DGP2. Qualitatively, the uncovered values match those obtained from DGP1. Namely, DGP2 produces an accurate fit of the incidence data (Fig 5A), and the inferred relative contact rate captures most

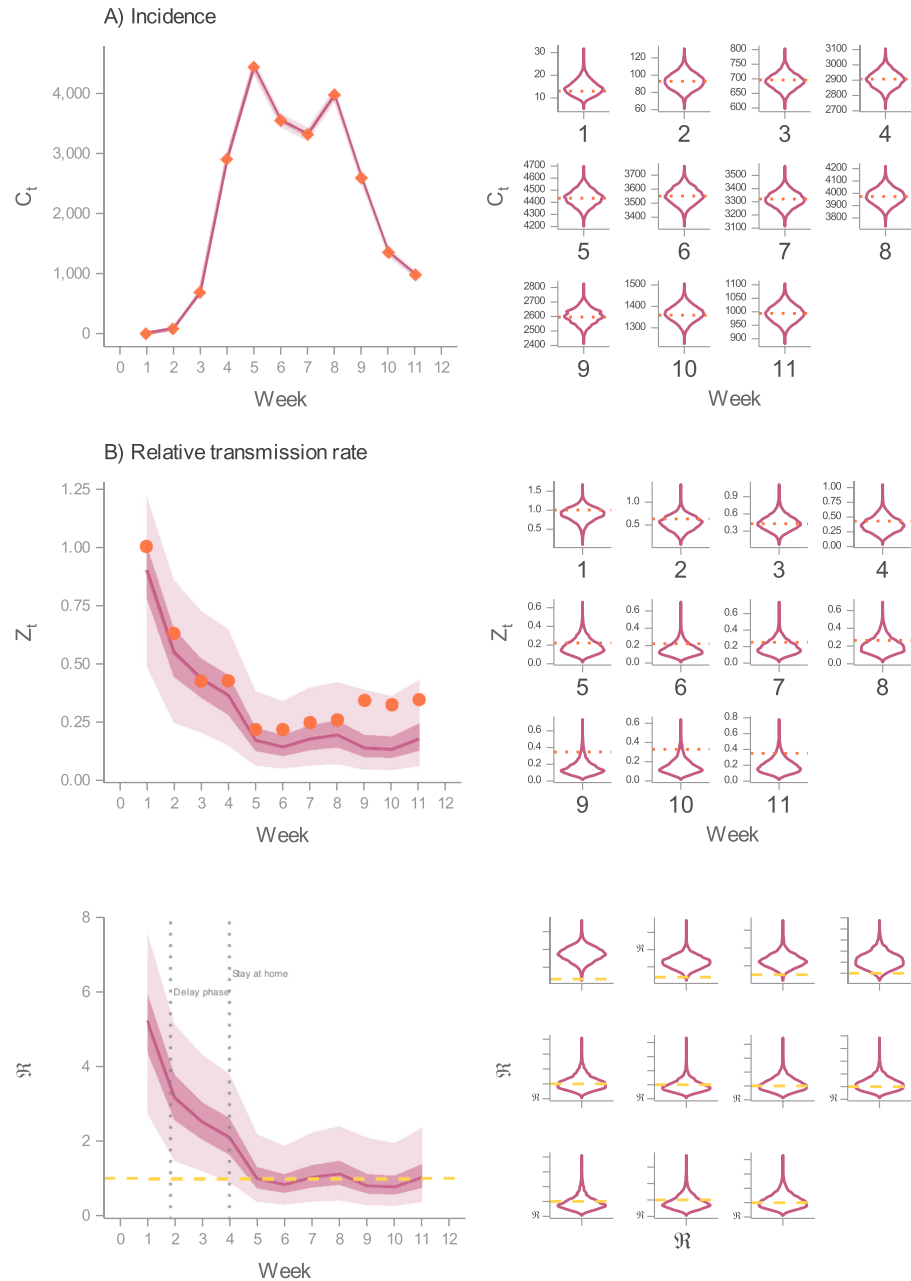


Fig 5. Inference on DGP2. In these three figures, the predicted values stem from DGP2's filtering distribution. Further, in the LHS, the solid line indicates the median, and the darker and lighter ribbons represent the 50% and 95% CI, respectively. (A) Comparison between the predicted incidence (solid line and ribbons in the LHS; violin plots in the RHS) and weekly detected cases from Ireland's first COVID-19 wave (rhombi in the LHS; horizontal dotted lines in the RHS). (B) Comparison between the predicted relative transmission rate (solid line and ribbons in the LHS; violin plots in the RHS) compared to Apple's mobility indexes in Ireland (points in the LHS; horizontal dotted lines in the RHS). (C) Predicted effective reproduction number (solid line and ribbons in the LHS; violin plots in the RHS). Horizontal dashed lines denote the epidemics threshold.

<https://doi.org/10.1371/journal.pcbi.1010206.g005>

of the mobility data (Fig 5B), resulting in a similar prediction of the effective reproduction number (Fig 5C). This outcome provides reassurance on the estimated transmission rate as an adequate account of the observed time series. That is, from two DGP that differ in the transmission rate's formulation, we estimate equivalent trajectories. DGP2, though, does not reduce significantly the uncertainty (see S6 Text Section 2.3.7) in the parameters (ν and v) that characterise the implemented NPIs.

DGP3—Adaptive expectations

$$\beta_t = \zeta Z_t^i \quad (10)$$

$$\frac{dZ^i}{dt} = \begin{cases} \frac{(v - Z_t^i)}{(v^{-1}/n)} & \text{for } i = n \\ \frac{(Z_t^{i+1} - Z_t^i)}{(v^{-1}/n)} & \text{for } i < n \end{cases} \quad (11)$$

The trajectories derived from the two previous DGPs (DGP1 and DGP2) suggest that it is reasonable to assume that the transmission rate's dynamics indeed follow a goal-seeking pattern (Eq 10). This conjecture is in agreement with the economic theory of *adaptive expectations*. First applied by Irving Fisher [51], this hypothesis posits that individuals gradually adjust their beliefs, and hence behaviour, in order to eliminate the discrepancy between the current state and a *desired* one [52]. In this case, such a discrepancy is the gap between individuals' behaviour at a given time t and the level of mobility that the restrictions (implicitly) aim to achieve. Mathematically, the n th-order information delay or *exponential smoothing* (Eq 11) provides a formal description of such an adjustment. This deterministic formulation describes the changes in current behaviour (Z_t^i) as the result of a series of intermediate exponential adjustments ($\frac{dZ_t^i}{dt}$), which one can interpret as the multiple stages intervening since the Government decrees mobility restrictions to the point where individuals alter their behaviour in accordance with the new rules. The delay order (n) represents the number of stages, where the most simple case ($n = 1$), the 1st-order information delay, is equivalent to the deterministic term in Eq 9. On the other hand, when $n \rightarrow \infty$, the dynamics follow a *term-time forcing* pattern.

To establish the exact number of stages, we evaluate the performance of nine candidate structures ($n = 1, \dots, 9$) in explaining the available data (incidence and mobility). From this evaluation, we ensemble the selected candidate with the SEI3R profile to generate the process model (PM3) of the third DPG (DPG3) presented in this paper (Fig 1). To complete DGP3's description, we formulate a measurement model (OM2) for the observed daily reported cases (y_a^1). As with DGP1 and DGP2, we assume these counts result from a Poisson distribution (Eqs 12 and 13). Moreover, OM2 does not include a structure relating mobility data to the relative transmission rate. We base this decision on the results shown in the previous sections. Since the mobility data is an imperfect predictor of the transmission rate, its inclusion in the inference process of a rigid deterministic structure may lead to *forced* model fits, resulting in undesired biases in parameter estimations. In relation to the inference process, since PM3 is deterministic, the inference of the filtering distribution becomes the estimation of DGP3's expected value. We approximate such expected value from a Bayesian perspective [53, 54] using Hamiltonian Monte Carlo [55] via Stan. The complete set of results can be found in S7 Text.

$$\frac{dC}{dt} = \eta P_t - C_t \delta(t \bmod 1) \quad (12)$$

$$y_1^d \sim \text{Pois}(C_t) \quad (13)$$

To illustrate the selection of DGP3's process model, we present the estimated expected values (fits) for each of the nine candidate structures (Fig 6). We depict expected values through simulated trajectories generated from one hundred draws from each model's posterior distribution. The results indicate that all of these structures yield similar fits to the incidence data. Using the mean absolute scaled error (MASE), a metric designed to measure the accuracy of time-series predictions [56], we notice diminishing marginal gains in accuracy as the order (of the number of stages) increases. These gains, though, are so tenuous that they do not provide clear guidance about which model to choose. To further complicate matters, the lower the delay order, the higher the likelihood value. Nevertheless, when we compare the expected relative transmission rate to mobility data, it can be seen that some structures approximate better the latter than others. If we accept the premise that mobility data is a proxy (supported by the results from DGP1 and DGP2), yet imperfect, measurement of the relative transmission rate, we can then lean towards the delay order that yields the lowest MASE ($n = 4$). From this structure's posterior distribution, we estimate, among others, the adjustment rate (v ; mean = 0.05, sd = 0.001), the minimum level of mobility (v ; mean = 0.11, sd = 0.005), and the effective reproduction number (discussed below). Notice that the particular form of the non-linear contact rate restricts the marginal distributions of v and v to such an extent that most of the probability mass concentrates on extremely narrow neighbourhoods. Despite this, those estimates resemble DPG2's MLE ($v = 0.05$, $v = 0.19$), which help us gain confidence in the overall process.

Acknowledging that the performance metrics above (MASE and likelihood values) do not lead to an unambiguous choice, we explore the implications of selecting an alternative measurement model. As it is widely known, the Poisson distribution is a discrete probability distribution in which the observation mean equals the variance [32]. Hence, using this distribution as a measurement model imposes a stringent assumption on the observational process of incidence counts. By contrast, the Negative Binomial distribution offers a more flexible framework to account for overdispersion in daily incidence. Moreover, the Negative Binomial converges to the Poisson distribution under a specific configuration. For this reason, we test the implications of this alternative formulation. See the complete set of results in S7 Text Section 4. Indeed, the posterior distribution suggests the presence of a small amount of overdispersion in the incidence data. However, such gain in realism is achieved at the expense of a *degenerate* posterior distribution. Succinctly, any of the model candidates coupled with the Negative binomial distribution yields a posterior distribution of two distinct modes, even from a single unknown parameter. This kind of behaviour is not unusual in Ordinary Differential Equation models. For instance, Gelman and colleagues [57] report a similar experience in the calibration of a simple mechanistic model of planetary motion.

In the set of bimodal distributions returned by Stan, we recognise two types of modes. One that corresponds to a region of unrealistic parameter values for which the HMC algorithm reveals pathological behaviour (divergences and low E-BFMI) [55] in the sampling procedure, rendering the inference from these samples unreliable. Conversely, the Markov chains that land in the other type of mode do not trigger any warnings from Stan. Furthermore, these *well-behaved* modes are located in regions similar to those found using the Poisson distribution. Following an exploratory analysis, we find that *well-behaved* modes and the set of posterior distributions obtained from the Poisson model provide similar (although not identical) information. Overall, the choice of the Poisson distribution and the delay order (4th) is the outcome of considering as a whole the information provided by the previous DGPs, and the two explored measurement models. This assessment, therefore, implies that we envision the




Fig 6. Inference on DGP3. Comparison between expected values and data. On the LHS, for each model, we show 100 overlapped simulations of the predicted incidence against daily case counts. On the RHS, for each model, we show 100 overlapped simulations of the predicted relative transmission rate against Apple’s mobility data. In this plot, we estimate the predicted values from the posterior distribution of each of the DGP3’s nine candidate process models.

<https://doi.org/10.1371/journal.pcbi.1010206.g006>

Poisson measurement model as an approximation that does not compromise the insights from the inference process. However, one cannot generalise this result to other applications. That is, taking the Poisson distribution as a default. On the contrary, it is imperative to test the assumptions embedded in any proposed measurement model and evaluate the trade-offs entailed by each alternative.

Discussion

Novel datasets that may assist modellers in gaining deeper insight into the dynamics of an infectious disease deserve a thorough examination. This task entails establishing adequate links between the data and a dynamical hypothesis. Far from trivial, one may derail the entire inference process by adopting a misspecified structure. For this reason, a robust approach involves the assessment of various levels of model complexity that account for the available data, which inevitably involves trade-offs [58]. This work highlighted the implications of committing to a particular model formulation. As seen above, DGP1 and DGP2 (DGPs with a stochastic process model) can only incorporate the mobility data if they are formulated in terms of weekly observations. Notwithstanding that this requirement reduces the number of data points available for the inference process, the loss of information is negligible. In contrast, a rigid structure such as DGP3 (whose process model is deterministic) restricts the use of mobility data only as a discriminant criterion.

With regards to the inference of fixed parameters, DGP1's *well-behaved* parameter space yields smooth quadratic profiles from which parameter uncertainty can be seamlessly calculated. Interestingly, when we amalgamate all the likelihood estimates, we obtain surfaces that resemble likelihood profiles. As a result, from three approaches (MCAP, profile, surface), we estimate similar confidence intervals. DGP2's parameter space is, on the other hand, of challenging exploration. In fact, the volume of high plausibility is so tightly concentrated that some regions in the MLE's neighbourhood yield vast negative log-likelihood values. To address this issue, we iterated over several hypercube sizes and densities until obtaining quadratic profiles, although not as smooth as those obtained from DGP1. Despite this hurdle, we obtain similar confidence intervals from the three quantification approaches. Regarding DGP3, given the Bayesian approach used to estimate its parameters, we refer to such uncertainty bounds as credible intervals. We obtain well-behaved quadratic posterior distributions for the nine candidate process models whose inference is backed by successful diagnostics unique to HMC. However, parameter estimates (posteriors) vary by the delay order, requiring a subjective assessment to determine which structure is more appropriate. Lastly, we consider the differences in computational burden between the inference methods (MIF + Particle Filter and HMC). Whereas performing parameter inference on DGP1 and DGP2 took roughly 14 and 20 hours, respectively; fitting DGP3's nine candidate models required 6 hours of computational time.

Likewise, the inference of the time-varying quantities deserves close inspection. DGP1 and DGP2 are more flexible than DGP3 in quantifying uncertainty. We illustrate this point with [Fig 7B and 7C](#). Here, we notice that DGP3 generates an estimate of the relative transmission rate and the effective reproduction number with narrow uncertainty intervals in comparison with those generated by the other DGPs. This apparent precision is the result of committing to a particular form of non-linear model, which imposes a stringent constraint in the shape of the transmission rate. By choosing the 4th-order information delay structure, we implicitly discard the possibility for the other formulations to be true, reducing the uncertainty in the estimations. However, we demonstrated that the nine delay orders account similarly for the incidence data, and to various degrees of accuracy, for the mobility data. Thus, we interpret the wide

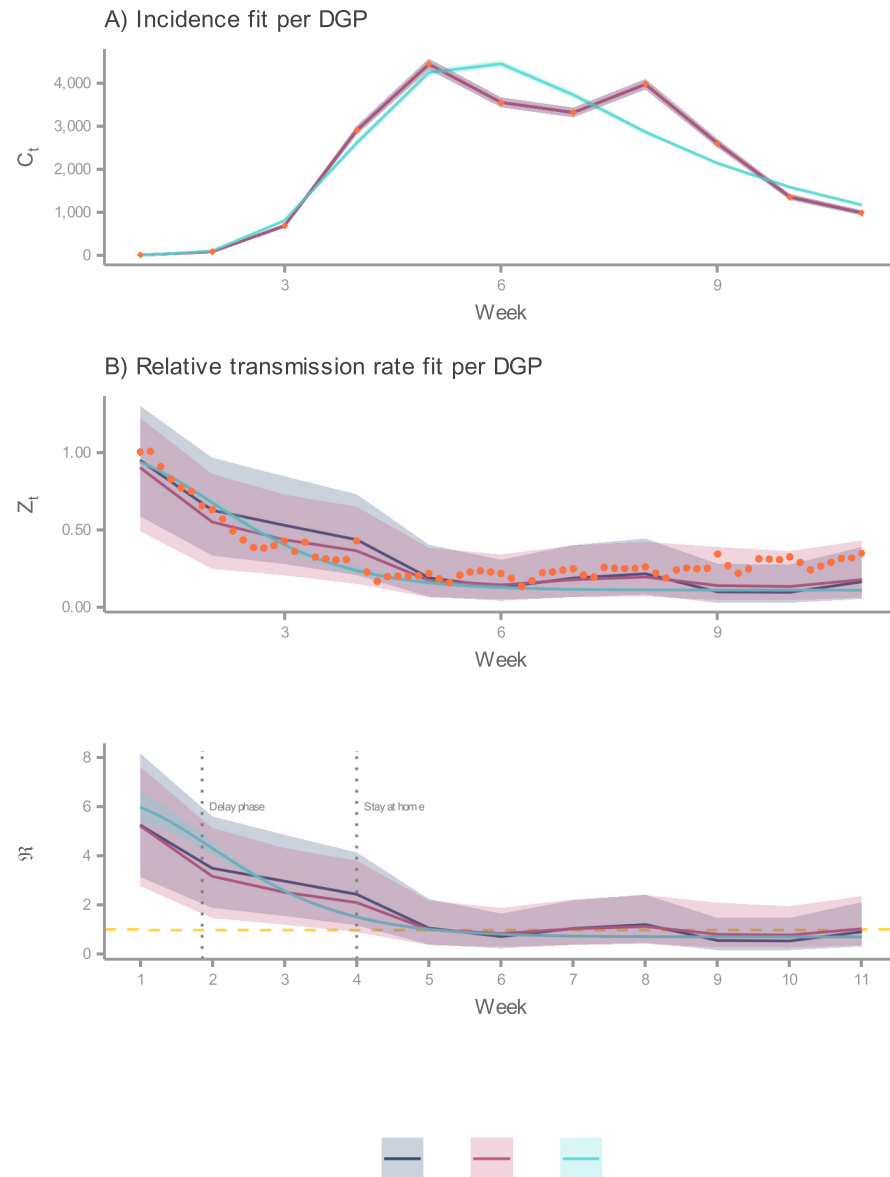


Fig 7. Comparison of predicted latent states. In this plot, predicted values stem either from the filtering distribution (DGP1 and DGP2) or the posterior predictive distribution (DGP3). Here, DGP3's process model corresponds to the structure that describes the transmission rate in terms of a 4th-order information delay. Further, solid lines indicates the median, and the ribbons represent the 95% CI. (A) Comparison between predicted incidences by DGP (solid lines and ribbons) and weekly detected COVID-19 cases (rhombi) in Ireland during the first wave. (B) Comparison between predicted relative transmission rates by DGP (solid lines and ribbons) and Apple's driving mobility indexes (points) in Ireland during the first wave. (C) Predicted effective reproductive numbers by DGP (solid lines and ribbons) during Ireland's first COVID-19 wave. The dashed horizontal line denotes the epidemics threshold.

<https://doi.org/10.1371/journal.pcbi.1010206.g007>

intervals generated by DGP1 and DGP2 as the uncertainty in the delay order plus the measurement error. This interpretation suggests that DGP3's plausible model candidates are subsumed under DGP1 and DGP2.

To conclude with this comparative analysis, we reflect on the role of DGPs presented in this paper. Owing to its flexible formulation, we can employ DGP1 for both retrospective and near real-time analysis (at least for the period where demographic processes do not significantly

impact the dynamics of the pandemic). In contrast, DGP2 and DGP3 formulations are context-dependent, restricted to retrospective analyses. Under this last role, we note that common patterns emerge from the three DGPs. Notwithstanding structural differences, all of them produce accurate fits to the incidence data (Fig 7A). Naturally, the stochastic process models replicate every feature in the data, whereas the deterministic one captures the underlying trend. Furthermore, the estimated medians for the relative transmission and the effective reproduction number (Fig 7B and 7C) tell similar stories. That is, individuals gradually decreased their movements following public health advice, which led to a decline in the transmission rate. This reduction pulled \mathcal{R}_t below the epidemics threshold, causing the incidence rate to subside. It should be noted that this mobility reduction levels off later in Ireland's first wave, suggesting a limit on the effectiveness of the implemented policies. We interpret this limit as the minimum mobility required for running essential services.

Finally, even though the primary interest of this work has been on estimating the effective reproduction number (\mathcal{R}_t), a by-product from this inference process is the approximation of the basic reproductive number (\mathcal{R}_0). This widely accepted metric [59] is defined as the average number of secondary infections produced when one infected individual is introduced into a totally susceptible population [3]. In the context of Ireland's COVID-19 epidemic, we derive similar \mathcal{R}_0 estimates from the three DGPs (DGP1: 95% CI[4.5–6.9], DGP2: 95% CI [4.4–6.8], DGP3: 95% CI[5.8–7.0]). These estimates are in close agreement with a previous modelling study on the COVID-19 pandemic in Ireland [35], albeit well above the initially reported $\mathcal{R}_0 = 2.2$ value from Wuhan [60]; a value that has been adopted as the reference point by the World Health Organization and other research groups [15, 61]. Other streams of research, however, argue that the initial estimate was low [62], and instead, advocate for higher values (4.5 [62]; 4.7–6.6 [63]). Moreover, the reader should recall that \mathcal{R}_0 is a context-dependent metric, and variations are expected due to population heterogeneity (e.g., age, spatial location, host genetics). In any case, we acknowledge the limitations that stem from the calibration of homogeneous population models, which require high \mathcal{R}_0 values to achieve accurate fits [17]. To address such limitations, future research should test the impact of disaggregating (by age or location) the structures presented in this paper. Another research avenue could explore the effect of replacing the deterministic rates in the within-host profile of these DGPs with stochastic ones that account for demographic and environmental effects.

Materials and methods

SEI3R profile

This within-host profile (Eqs 14–20) is formulated based on the work from Davies and colleagues [15]. Here, we assume that individuals are initially susceptible (S) and become exposed (E), at a rate λ , after effective contact with an infectious person (I, P, A). After a latent period (σ^{-1}), exposed individuals follow one of two paths. With probability ω , following a period (η^{-1}) of preclinical infectiousness (P), individuals develop full symptoms while transmitting the pathogen. This stage is known as the clinical infection state and lasts for γ^{-1} days. On the second path, with probability $1-\omega$, individuals enter a subclinical state (A) with none (asymptomatic) or mild symptoms (paucisymptomatic), who are not captured by the healthcare system. Individuals on this path recover after κ^{-1} days and are relatively (μ) less infectious than their counterparts on the clinical path. Finally, individuals from both paths eventually converge to the recovered state (R), in which they are no longer infectious and are immune to re-infection. In S2–S6 Text, we provide the values for fixed parameters and initial states and their respective

sources.

$$\frac{dS}{dt} = -S_t \lambda_t \tag{14}$$

$$\frac{dE}{dt} = S_t \lambda_t - \sigma E_t \tag{15}$$

$$\frac{dP}{dt} = \omega \sigma E_t - \eta P_t \tag{16}$$

$$\frac{dI}{dt} = \eta P_t - \gamma I_t \tag{17}$$

$$\frac{dA}{dt} = (1 - \omega) \sigma E_t - \kappa A_t \tag{18}$$

$$\frac{dR}{dt} = \kappa A_t + \gamma I_t \tag{19}$$

$$\lambda_t = \frac{\beta_t (P_t + I_t + \mu A_t)}{N_t} \tag{20}$$

Basic and effective reproductive number

To derive an analytical expression for the basic reproduction number (\mathfrak{R}_0) from the SEI3R profile, we employ the next generation matrix method [64]. That is, we rewrite the infected states' transitions (rates) in the form of two matrices. The first matrix \mathcal{F} corresponds to the rate of appearance of new infections in each compartment of infected individuals, and the second matrix \mathcal{V} corresponds to the rate of other transitions between compartments of infected individuals. From these matrices, we define the next generation matrix as $\mathcal{F}\mathcal{V}^{-1}$, whose largest eigenvalue (spectral radius) corresponds to \mathfrak{R}_0 [65]. We obtain the spectral radius's analytical solution (Eq 21) using the software system *Mathematica* (see Github repository). Following this expression, we can define (Eq 22) the effective reproductive number (\mathfrak{R}_t) as the product between \mathfrak{R}_0 and the susceptible fraction ($\frac{S_t}{N_t}$).

$$\mathfrak{R}_0 = \zeta Z_0 [\omega \gamma + \eta + (1 - \omega) \mu \kappa] \tag{21}$$

$$\mathfrak{R}_t = \mathfrak{R}_0 \frac{S_t}{N_t} \tag{22}$$

Filtering distribution

$$p(x_t | y_{1:t-1}) = \int p(x_t | x_{t-1}) p(x_{t-1} | y_{1:t-1}) dx_{t-1} \tag{23}$$

$$p(x_t | y_{1:t}) = \frac{p(y_t | x_t) p(x_t | y_{1:t-1})}{p(y_t | y_{1:t-1})} \tag{24}$$

The essence of the state-space approach is to estimate the state of a dynamical system using a

sequence of noisy measurements made on the system. We formulate this problem in terms of a recursive filter whose purpose is to construct the state's posterior probability density function (pdf) based on all available information, including the set of received measurements [33]. Formally, $p(x_t|y_{1:t})$. We refer to this pdf as the *filtering distribution*, whose inference process consists of two stages: *prediction* and *update*.

The prediction stage (Eq 23) draws on the *plug-and-play* property [17] to generate, from simulations of the process model $p(x_t|x_{t-1})$, a vector of predictions that describe the state at time t (x_t), which are conditional on the previously estimated state ($x_{t-1}|y_{t-1}$). Then, the update operation (Eq 24) uses the latest measurement to modify the prediction pdf ($p(x_t|y_{1:t-1})$). In practice, we assign weights to the prediction vector based on its plausibility, which is estimated from the measurement model $p(y_t|x_t)$. With these weights, we use the Sequence Importance Sampling algorithm [42] to produce samples that describe the filtering distribution. It is important to remark that this is a sequential process (hence the name Sequential Monte Carlo), executed every time a measurement is received. Moreover, in this simplified formulation, it is assumed that X_0 and θ (Eqs 1 and 2) are known. We refer the interested reader to [30, 33] for a complete treatment of this approach.

Supporting information

S1 Text. Likelihood convergence.

(PDF)

S2 Text. Inference on DGPI's Candidate 2.

(PDF)

S3 Text. Inference on DGPI's Candidate 4.

(PDF)

S4 Text. Inference on DGPI's Candidate 5.

(PDF)

S5 Text. Inference on DGPI's Candidate 6.

(PDF)

S6 Text. Inference on DGP2.

(PDF)

S7 Text. Inference on DGP3.

(PDF)

S8 Text. Output comparison.

(PDF)

S1 Data. Incidence data.

(XLSX)

Author Contributions

Conceptualization: Jair Andrade.

Data curation: Jair Andrade.

Formal analysis: Jair Andrade.

Funding acquisition: Jim Duggan.

Investigation: Jair Andrade.

Methodology: Jair Andrade.

Project administration: Jim Duggan.

Resources: Jim Duggan.

Software: Jair Andrade.

Validation: Jair Andrade.

Visualization: Jair Andrade.

Writing – original draft: Jair Andrade.

Writing – review & editing: Jim Duggan.

References

1. Flaxman S, Mishra S, Gandy A, Unwin HJT, Mellan TA, Coupland H, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*. 2020; 584(7820):257–261. <https://doi.org/10.1038/s41586-020-2405-7> PMID: 32512579
2. Douglas M, Katikireddi SV, Taulbut M, McKee M, McCartney G. Mitigating the wider health effects of covid-19 pandemic response. *BMJ*. 2020; 369. <https://doi.org/10.1136/bmj.m1557> PMID: 32341002
3. Anderson RM, May RM. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press; 1992.
4. Nishiura H, Chowell G. In: Chowell G, Hyman JM, Bettencourt LMA, Castillo-Chavez C, editors. *The Effective Reproduction Number as a Prelude to Statistical Estimation of Time-Dependent Epidemic Trends*. Dordrecht: Springer Netherlands; 2009. p. 103–121. Available from: https://doi.org/10.1007/978-90-481-2313-1_5.
5. Vynnycky E, White R. *An Introduction to Infectious Disease Modelling*. Oxford University Press; 2010.
6. Brett T, Ajelli M, Liu QH, Krauland MG, Grefenstette JJ, van Panhuis WG, et al. Detecting critical slowing down in high-dimensional epidemiological systems. *PLOS Computational Biology*. 2020; 16(3):1–19. <https://doi.org/10.1371/journal.pcbi.1007679> PMID: 32150536
7. Gostic KM, McGough L, Baskerville EB, Abbott S, Joshi K, Tedijanto C, et al. Practical considerations for measuring the effective reproductive number, Rt. *PLOS Computational Biology*. 2020; 16(12):1–21. <https://doi.org/10.1371/journal.pcbi.1008409> PMID: 33301457
8. Cori A, Ferguson NM, Fraser C, Cauchemez S. A New Framework and Software to Estimate Time-Varying Reproduction Numbers During Epidemics. *American Journal of Epidemiology*. 2013; 178(9):1505–1512. <https://doi.org/10.1093/aje/kwt133> PMID: 24043437
9. Thompson RN, Stockwin JE, van Gaalen RD, Polonsky JA, Kamvar ZN, Demarsh PA, et al. Improved inference of time-varying reproduction numbers during infectious disease outbreaks. *Epidemics*. 2019; 29:100356. <https://doi.org/10.1016/j.epidem.2019.100356> PMID: 31624039
10. Wallinga J, Teunis P. Different Epidemic Curves for Severe Acute Respiratory Syndrome Reveal Similar Impacts of Control Measures. *American Journal of Epidemiology*. 2004; 160(6):509–516. <https://doi.org/10.1093/aje/kwh255> PMID: 15353409
11. Andrade J, Duggan J. An evaluation of Hamiltonian Monte Carlo performance to calibrate age-structured compartmental SEIR models to incidence data. *Epidemics*. 2020; 33:100415. <https://doi.org/10.1016/j.epidem.2020.100415> PMID: 33212347
12. Bretó C. Modeling and Inference for Infectious Disease Dynamics: A Likelihood-Based Approach. *Statistical Science*. 2018; 33(1):57–69. <https://doi.org/10.1214/17-STS636> PMID: 29755198
13. Oliva R. Model calibration as a testing strategy for system dynamics models. *European Journal of Operational Research*. 2003; 151(3):552–568. [https://doi.org/10.1016/S0377-2217\(02\)00622-7](https://doi.org/10.1016/S0377-2217(02)00622-7)
14. Dehning J, Zierenberg J, Spitzner FP, Wibral M, Neto JP, Wilczek M, et al. Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions. *Science*. 2020; 369(6500). <https://doi.org/10.1126/science.abb9789> PMID: 32414780
15. Davies NG, Klepac P, Liu Y, Prem K, Jit M, Pearson CAB, et al. Age-dependent effects in the transmission and control of COVID-19 epidemics. *Nature Medicine*. 2020; 26(8):1205–1211. <https://doi.org/10.1038/s41591-020-0962-9> PMID: 32546824

16. Bretó C, He D, Ionides EL, King AA. Time series analysis via mechanistic models. *The Annals of Applied Statistics*. 2009; 3(1):319–348.
17. He D, Ionides EL, King AA. Plug-and-play inference for disease dynamics: measles in large and small populations as a case study. *Journal of The Royal Society Interface*. 2010; 7(43):271–283. <https://doi.org/10.1098/rsif.2009.0151> PMID: 19535416
18. Keeling MJ, Rohani P, Grenfell BT. Seasonally forced disease dynamics explored as switching between attractors. *Physica D: Nonlinear Phenomena*. 2001; 148(3):317–335. [https://doi.org/10.1016/S0167-2789\(00\)00187-1](https://doi.org/10.1016/S0167-2789(00)00187-1)
19. Liu X, Stechlinski P. Infectious disease models with time-varying parameters and general nonlinear incidence rate. *Applied Mathematical Modelling*. 2012; 36(5):1974–1994. <https://doi.org/10.1016/j.apm.2011.08.019>
20. Liu L, Zhao XQ, Zhou Y. A Tuberculosis Model with Seasonality. *Bulletin of Mathematical Biology*. 2010; 72(4):931–952. <https://doi.org/10.1007/s11538-009-9477-8> PMID: 20063125
21. Linka K, Peirlinck M, Kuhl E. The reproduction number of COVID-19 and its correlation with public health interventions. *Computational Mechanics*. 2020; 66(4):1035–1050. <https://doi.org/10.1007/s00466-020-01880-8> PMID: 32836597
22. Endo A, van Leeuwen E, Baguelin M. Introduction to particle Markov-chain Monte Carlo for disease dynamics modellers. *Epidemics*. 2019; 29:100363. <https://doi.org/10.1016/j.epidem.2019.100363> PMID: 31587877
23. Dureau J, Kalogeropoulos K, Baguelin M. Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics*. 2013; 14(3):541–555. <https://doi.org/10.1093/biostatistics/kxs052> PMID: 23292757
24. Funk S, Camacho A, Kucharski AJ, Eggo RM, Edmunds WJ. Real-time forecasting of infectious disease dynamics with a stochastic semi-mechanistic model. *Epidemics*. 2018; 22:56–61. <https://doi.org/10.1016/j.epidem.2016.11.003> PMID: 28038870
25. Camacho A, Kucharski A, Aki-Sawyer Y, White MA, Flasche S, Baguelin M, et al. Temporal Changes in Ebola Transmission in Sierra Leone and Implications for Control Requirements: a Real-time Modelling Study. *PLoS currents*. 2015; 7:ecurrents.outbreaks.406ae55e83ec0b5193e30856b9235ed2. <https://doi.org/10.1371/currents.outbreaks.406ae55e83ec0b5193e30856b9235ed2> PMID: 25737806
26. Davies NG, Abbott S, Barnard RC, Jarvis CI, Kucharski AJ, Munday JD, et al. Estimated transmissibility and impact of SARS-CoV-2 lineage B.1.1.7 in England. *Science*. 2021; 372 (6538). <https://doi.org/10.1126/science.abg3055> PMID: 33658326
27. King AA, Nguyen D, Ionides EL. Statistical Inference for Partially Observed Markov Processes via the R Package pomp. *Journal of Statistical Software, Articles*. 2016; 69(12):1–43.
28. Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, et al. Stan: A probabilistic programming language. *Journal of statistical software*. 2017; 76(1). <https://doi.org/10.18637/jss.v076.i01>
29. Ritchie H, Ortiz-Ospina E, Beltekian D, Mathieu E, Hasell J, Macdonald B, et al. Coronavirus Pandemic (COVID-19). *Our World in Data*. 2020.
30. Chopin N, Papaspiliopoulos O. An Introduction to Sequential Monte Carlo. *Springer Series in Statistics*. Springer International Publishing; 2020.
31. Ionides EL, Bretó C, King AA. Inference for nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*. 2006; 103(49):18438–18443. <https://doi.org/10.1073/pnas.0603181103> PMID: 17121996
32. Blitzstein JK, Hwang J. *Introduction to Probability, Second Edition*. Chapman & Hall/CRC Texts in Statistical Science. CRC Press; 2019.
33. Arulampalam MS, Maskell S, Gordon N, Clapp T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*. 2002; 50(2):174–188. <https://doi.org/10.1109/78.978374>
34. Keeling MJ, Rohani P. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press; 2011.
35. Gleeson JP, Brendan Murphy T, O'Brien JD, Friel N, Bargary N, O'Sullivan DJP. Calibrating COVID-19 susceptible-exposed-infected-removed models with time-varying effective contact rates. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2022; 380 (2214):20210120. <https://doi.org/10.1098/rsta.2021.0120>
36. Rasmussen DA, Ratmann O, Koelle K. Inference for Nonlinear Epidemiological Models Using Genealogies and Time Series. *PLOS Computational Biology*. 2011; 7(8):1–11. <https://doi.org/10.1371/journal.pcbi.1002136> PMID: 21901082
37. Wiersema UF. *Brownian Motion Calculus*. Wiley Finance. Wiley; 2008.

38. Priestley MB. STATE-DEPENDENT MODELS: A GENERAL APPROACH TO NON-LINEAR TIME SERIES ANALYSIS. *Journal of Time Series Analysis*. 1980; 1(1):47–71. <https://doi.org/10.1111/j.1467-9892.1980.tb00300.x>
39. Hastie T, Tibshirani R. Varying-Coefficient Models. *Journal of the Royal Statistical Society Series B (Methodological)*. 1993; 55(4):757–796. <https://doi.org/10.1111/j.2517-6161.1993.tb01939.x>
40. Kimeldorf GS, Wahba G. A Correspondence Between Bayesian Estimation on Stochastic Processes and Smoothing by Splines. *The Annals of Mathematical Statistics*. 1970; 41(2):495–502. <https://doi.org/10.1214/aoms/1177697089>
41. King AA, Domenech de Cellès M, Magpantay FMG, Rohani P. Avoidable errors in the modelling of outbreaks of emerging pathogens, with special reference to Ebola. *Proceedings of the Royal Society B: Biological Sciences*. 2015; 282(1806):20150347. <https://doi.org/10.1098/rspb.2015.0347> PMID: 25833863
42. Gordon NJ, Smith AFM, Salmond DJ. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*. 1993; 140:107–113(6). <https://doi.org/10.1049/ip-f-2.1993.0015>
43. Ionides EL, Bhadra A, Atchadé Y, King A. Iterated filtering. *The Annals of Statistics*. 2011; 39(3):1776–1802. <https://doi.org/10.1214/11-AOS886>
44. Wale N, Jones MJ, Sim DG, Read AF, King AA. The contribution of host cell-directed vs. parasite-directed immunity to the disease and dynamics of malaria infections. *Proceedings of the National Academy of Sciences*. 2019; 116(44):22386–22392. <https://doi.org/10.1073/pnas.1908147116>
45. Pawitan Y. In All Likelihood: Statistical Modelling and Inference Using Likelihood. In All Likelihood: Statistical Modelling and Inference Using Likelihood. OUP Oxford; 2013.
46. Ionides EL, Breto C, Park J, Smith RA, King AA. Monte Carlo profile confidence intervals for dynamic systems. *Journal of The Royal Society Interface*. 2017; 14(132):20170126. <https://doi.org/10.1098/rsif.2017.0126> PMID: 28679663
47. HPSC. Preliminary report of the results of the Study to Investigate COVID-19 Infection in People Living in Ireland (SCOPI): A national seroprevalence study, June–July 2020; 2020. Available from: <https://www.hpsc.ie/a-z/respiratory/coronavirus/novelcoronavirus/scopi/>.
48. Barlas Y, Yasarcan H. In: Quadrat-Ullah H, Spector JM, Davidsen PI, editors. A Comprehensive Model of Goal Dynamics in Organizations: Setting, Evaluation and Revision. Berlin, Heidelberg: Springer Berlin Heidelberg; 2008. p. 295–320. Available from: https://doi.org/10.1007/978-3-540-73665-3_15.
49. Cox JC, Ingersoll JE, Ross SA. A Theory of the Term Structure of Interest Rates. *Econometrica*. 1985; 53(2):385–407. <https://doi.org/10.2307/1911242>
50. Sklar A. Random Variables, Distribution Functions, and Copulas: A Personal Look Backward and Forward. Lecture Notes-Monograph Series. 1996; 28:1–14. <https://doi.org/10.1214/lnms/1215452606>
51. Fisher I. The Purchasing Power of Money: Its Determination And Relation to Credit Interest And Crises. Cosimo classics economics. Cosimo Classics; 2006.
52. Sterman J. Business Dynamics: Systems Thinking and Modeling for a Complex World. McGraw-Hill Higher Education. Irwin/McGraw-Hill; 2000.
53. Andrade J, Duggan J. A Bayesian approach to calibrate system dynamics models using Hamiltonian Monte Carlo. *System Dynamics Review*. 2021; 37(4):283–309. <https://doi.org/10.1002/sdr.1693>
54. Grinsztajn L, Semenova E, Margossian CC, Riou J. Bayesian workflow for disease transmission modeling in Stan. *Statistics in Medicine*. 2021; 40(27):6209–6234. <https://doi.org/10.1002/sim.9164> PMID: 34494686
55. Betancourt M. A Conceptual Introduction to Hamiltonian Monte Carlo; 2018.
56. Hyndman RJ, Koehler AB. Another look at measures of forecast accuracy. *International Journal of Forecasting*. 2006; 22(4):679–688. <https://doi.org/10.1016/j.ijforecast.2006.03.001>
57. Gelman A, Vehtari A, Simpson D, Margossian CC, Carpenter B, Yao Y, et al. Bayesian Workflow; 2020.
58. Funk S, King AA. Choices and trade-offs in inference with infectious disease models. *Epidemics*. 2020; 30:100383. <https://doi.org/10.1016/j.epidem.2019.100383>
59. Delamater P, Street E, Leslie T, Yang YT, Jacobsen K. Complexity of the Basic Reproduction Number (R_0). 2019; 25(1):1.
60. Li Q, Guan X, Wu P, Wang X, Zhou L, Tong Y, et al. Early Transmission Dynamics in Wuhan, China, of Novel Coronavirus-Infected Pneumonia. *New England Journal of Medicine*. 2020; 382(13):1199–1207. <https://doi.org/10.1056/NEJMoa2001316> PMID: 31995857

61. Petersen E, Koopmans M, Go U, Hamer DH, Petrosillo N, Castelli F, et al. Comparing SARS-CoV-2 with SARS-CoV and influenza pandemics. *The Lancet Infectious Diseases*. 2020; 20(9):e238–e244. [https://doi.org/10.1016/S1473-3099\(20\)30484-9](https://doi.org/10.1016/S1473-3099(20)30484-9) PMID: 32628905
62. Katul GG, Mrad A, Bonetti S, Manoli G, Parolari AJ. Global convergence of COVID-19 basic reproduction number and estimation from early-time SIR dynamics. *PLOS ONE*. 2020; 15(9):1–22. <https://doi.org/10.1371/journal.pone.0239800> PMID: 32970786
63. Sanche S, Lin YT, Xu C, Romero-Severson E, Hengartner N, Ke R. High Contagiousness and Rapid Spread of Severe Acute Respiratory Syndrome Coronavirus 2. 2020; 26(7):1470.
64. Diekmann O, Heesterbeek JAP, Metz JAJ. On the definition and the computation of the basic reproduction ratio R_0 in models for infectious diseases in heterogeneous populations. *Journal of Mathematical Biology*. 1990; 28(4):365–382. <https://doi.org/10.1007/BF00178324> PMID: 2117040
65. van den Driessche P. Reproduction numbers of infectious disease models. *Infectious Disease Modelling*. 2017; 2(3):288–303. <https://doi.org/10.1016/j.idm.2017.06.002> PMID: 29928743