

RESEARCH ARTICLE

Optimal algorithms for controlling infectious diseases in real time using noisy infection data

Sandor Beregi , Kris V. Parag *

Department of Infectious Disease Epidemiology, School of Public Health, Imperial College London, London, United Kingdom

* k.parag@imperial.ac.uk



Abstract

Deciding when to enforce or relax non-pharmaceutical interventions (NPIs) based on real-time outbreak surveillance data is a central challenge in infectious disease epidemiology. Reporting delays and infection under-ascertainment, which characterise practical surveillance data, can misinform decision-making, prompting mistimed NPIs that fail to control spread or permitting deleterious epidemic peaks that overload healthcare capacities. To mitigate these risks, recent studies propose more data-insensitive strategies that trigger NPIs at predetermined times or infection thresholds. However, these strategies often increase NPI durations, amplifying their substantial costs to livelihood and life-quality. We develop a novel model-predictive control algorithm that optimises NPI decisions. We jointly minimise the cumulative risks and costs of interventions of different stringency over stochastic epidemic projections. Our algorithm is among the earliest to realistically incorporate uncertainties underlying both the generation and surveillance of infections. We find, except under extremely delayed reporting, that our projective approach outperforms data-insensitive strategies and show that earlier decisions strikingly improve real-time control with reduced NPI costs. Moreover, we expose how surveillance quality, disease growth and NPI frequency intrinsically limit our ability to flatten epidemic peaks or dampen endemic oscillations and reveal why this potentially makes Ebola virus more controllable than SARS-CoV-2. Our algorithm provides a general framework for guiding optimal NPI decisions ahead-of-time and identifying the key factors limiting practical epidemic control.

OPEN ACCESS

Citation: Beregi S, Parag KV (2025) Optimal algorithms for controlling infectious diseases in real time using noisy infection data. PLoS Comput Biol 21(9): e1013426. <https://doi.org/10.1371/journal.pcbi.1013426>

Editor: Tom Britton, Stockholms Universitet, SWEDEN

Received: August 8, 2024

Accepted: August 11, 2025

Published: September 3, 2025

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcbi.1013426>

Copyright: © 2025 Beregi, Parag. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: The code generating the results presented here is available at <https://github.com/sandorberegi/Epidemic-control-with-noisy-real-time-data>.

Author summary

In our work, we tackle the challenge of determining the best time to enforce or relax non-pharmaceutical interventions (NPIs), such as mandatory mask wearing, social distancing or quarantine, to manage the spread of infectious diseases. Making an optimal

Funding: SB and KVP acknowledge funding from the MRC Centre for Global Infectious Disease Analysis (reference MR/X020258/1), funded by the UK Medical Research Council (MRC). This UK funded award is carried out in the frame of the Global Health EDCTP3 Joint Undertaking. The funders had no role in study design, data collection and analysis, decision to publish, or manuscript preparation.

Competing interests: The authors have declared that no competing interests exist.

decision on NPIs requires balancing the risks and the burden of prevalent infections on the healthcare systems against the costs of restrictive measures to livelihood and life-quality. Real-world data used to inform these decisions can often be unreliable due to delays in reporting and missed cases. This can lead to NPIs being implemented too late or too soon, and as such, failing to contain the outbreak or unnecessarily disrupting daily life. We introduced a novel algorithm that projects future scenarios based on current data to optimise NPI decisions across interventions with different overall stringency and costs. Our results show that our method can effectively reduce the duration and cost of NPIs while better controlling the spread of infections than more traditional approaches of having fixed thresholds or NPI schedules. Our approach optimises these decisions even when data is uncertain and is a versatile tool that can adapt to changes in the epidemic dynamics, such as the appearance of new variants. Moreover, we highlight how the quality of surveillance, the growth rate of the disease, and the frequency of NPIs play crucial roles in managing outbreaks and why this potentially makes Ebola virus more controllable than SARS-CoV-2.

Introduction

When and how should we intervene in order to most effectively manage an emerging infectious disease? This is a question that is at the core of public health policy-making and has been the subject of ongoing debate [1–3]. This decision problem is especially crucial during the early stages of an outbreak when there is no or limited immunity in the population and vaccines or other pharmaceutical remedies are unavailable. In this situation, the main control measures are non-pharmaceutical interventions (NPIs), such as mandatory social distancing, mask-wearing, lockdowns and travel restrictions [4–7].

Outbreak management policies need to balance the risks from mistimed and ineffective intervention decisions with the likely costs of those decisions. An NPI that is applied too slowly or removed too quickly risks large epidemic peaks or rebounds that overburden healthcare systems [8]. However, more conservative approaches may prompt long periods of restrictions that incur costs due to closed economic sectors and borders as well as limited mobility [9].

Optimising the counteracting costs and risks of NPIs is a challenging and enduring problem. This problem is further exacerbated by the practical constraints of real-time surveillance. The data available for an unfolding epidemic are subject to multiple sources of noise and uncertainty that fundamentally limit our ability to infer the state of the epidemic [1,10–12]. Solutions therefore require evidence-based research into the benefits, risks and societal costs of different NPIs and public health policies [13–17] as well as rigorous algorithms that can integrate outcomes of that research with uncertain knowledge of the epidemic state to guide decision-making [18–21].

Here, we focus on the latter issue and investigate how optimal, data-driven policies can be derived from real-time surveillance data. Our study analyses the overall effect and costs of tiered restrictions rather than modelling the mechanisms underpinning individual preventive measures. Our framework can include multiple individual interventions or intervention packages with known efficacies. Studies such as [5,22,23] retrospectively modelled the effect of numerous interventions via their impact on the reproduction number. Our model combines those effects with costs and other benefits (e.g., peak size or endemic load) to make decisions prospectively about which interventions from among a suite of possible options should be initiated or removed at any given policy review time. We leverage ideas from control theory

and reinforcement learning and expose exactly how uncertainties in practical surveillance intrinsically limit the optimal policies. This approach, which uses feedback control, dynamically updates NPI choices by feeding back data on the incidence of new infections that should reflect the most recent epidemic state.

However, the reporting of the incidence of infections is subject to delay and under-ascertainment that is often inherent to real surveillance systems. Under-ascertainment of infections can result from asymptomatic and mild infections, which are rarely observed, or from limitations on testing capacities [24]. Consequently, we only receive reports on a random fraction of all new infections. Delays can emerge from the lag between infection and symptom onset or confirmation as well as latencies in testing and processing test results. The consequence of this is that the reported time series of cases (or a related proxy for infections) are stochastically behind the actual incidence [1,25,26]. We also highlight that a clear distinction between our approach and much of the existing literature is our use of the incidence of new cases or infections instead of prevalence to inform control actions. Establishing prevalence is difficult and often requires additional efforts and testing programmes, such as the REACT scheme in the UK [27]. By framing our optimisation in terms of incidence we focus on more generalisable decision-making that only requires routinely available data and better align with existing real-time response frameworks.

These sources of noise and uncertainty sparked an ongoing debate on what is the best approach for controlling epidemics in real time. At least three challenges have influenced this debate. First, although feedback control is widely used to solve real-time problems in electrical and mechanical engineering [28–31] these strategies can become destabilised by noise and delays [1,32–36]. Second, the timing of public health interventions is critical to their efficacy and hence their associated risks and costs [37,38]. Last, integrating costs, risks and noise within a framework is difficult and often intractable for deriving insights.

In view of these challenges, some studies have proposed feedback-independent methods that are insensitive to noise and uncertainty in real-time epidemic surveillance. One such approach is to implement a pre-set sequence of cyclic switching between lockdowns and periods with no restrictions [39]. This was proposed as a strategy to exit full-lockdown more reliably [40]. Other works have focussed on optimising the timing of specific interventions considering a ‘one-shot’ control with the start and ending time to be optimised [38,41,42], highlighting the importance of timing for efficiency and the detrimental effect of delays in intervention. This makes a case for using evidence-based policies that dynamically optimise interventions to available-data. Some recent studies support this optimal timing approach informed by real-time data over feedback-independent methods even for uncertainty in data [43,44]. However, those works do not consider the intrinsic stochasticity of the epidemic and the complex and non-intuitive tradeoffs required to balance the desired public health outcomes against the costs of interventions.

The last challenge, which stems from the complexity of the decision-making process given the uncertainties in data, transmission details and the likely effect of actions, has meant that the majority of approaches in the field on cost-optimal control only consider deterministic models or limited modelling of noise. As a result, there is scope for fully stochastic but rigorous decision-making and modelling frameworks that can guide interventions by providing insight into how cost-optimal choices and various uncertainties interact.

We consider this real-time control problem in a probabilistic setting, where the epidemic is modelled by a renewal branching process [45]. This is more realistic than the deterministic approach, generalisable to multiple diseases and reflects on the intrinsic variability of infections between individuals. We parameterise our renewal models to describe the dynamics

of epidemics of COVID-19 and Ebola virus disease. We propose a model predictive optimal control strategy that minimises the costs of NPIs jointly with those generated from the infections projected to occur under the renewal process given our NPI choices. Our control approach is based on real-time incidence data which is delayed and under-ascertained, and incorporates the stochasticity of the epidemic generation process. We also include other key factors that limit real-time control, such as constraints on how frequently NPI policies can be changed and restrict control actions that can be applied to a finite set of NPI options rather than continuous levels of effort.

We assess what limitations data quality imposes on the viability of real-time feedback control for epidemic management and how this is influenced by disease growth dynamics. We compare the performance of our proposed optimal control algorithm with two benchmark control strategies that apply decisions based on chosen thresholds or times. We demonstrate that our algorithm not only outperforms these approaches but can adapt to unexpected changes such as the emergence of new variants or reduced effectiveness of NPIs due to behavioural changes.

Methods

Epidemic governing equations

We model the spread of the disease in a population as a generalisation of the standard renewal branching process [45]. This model is used both to make projections that inform optimal control and to simulate ‘ground truth’ epidemic trajectories. The renewal branching process is a stochastic model describing how the incidence of new infections on day t , I_t , depends on past infections at times $s \leq t$ and the characteristics of the disease. This is captured by the Poisson distribution

$$I_t \sim \text{Pois}\left(\sum_{s=0}^t I_s R_s w_{t-s}\right), \quad (1)$$

where R_t is the effective reproduction number on day t , with the set of weights, w_{t-s} for all s , obtained from the generation time distribution [46] of the disease. We assume that the generation time distribution is known or estimated from other paired transmission data. The weight w_{t-s} is the probability that a secondary infection occurs $t-s$ days after its primary infection. As is standard practice, we model the stochasticity of the generation time with a Gamma distribution

$$t_{\text{gen}} \sim \text{Gamma}(\alpha_{\text{gen}}, \beta_{\text{gen}}). \quad (2)$$

The shape and scale factors α_{gen} and β_{gen} parametrise the probability density function $p_{\text{gen}}(t)$. The weights w_{t-s} used in Eq (1) are then calculated as $w_{t-s} = p_{\text{gen}}(t-s) / \sum_{k=0}^t p_{\text{gen}}(k)$. We consider two generation time distributions, with respective parameters provided in Table 1, which are commonly used to describe epidemics of Ebola virus disease [47] and COVID-19 [48].

The above formulation generalises the standard renewal model, which describes incidence $I_t \sim \text{Pois}(R_t \sum_{s=0}^t I_s w_{t-s})$ with the sum $\sum_{s=0}^t I_s w_{t-s}$ referred to as total infectiousness of all infectious individuals [49]. However, this classical formula implies that any intervention applied to curb the spread of the disease results in an immediate change in the reproduction number. We apply the generalised formula of Eq (1) to model scenarios where the infectiousness of a population and the impact of interventions depend on both the history of infections and reproduction numbers. The latter dependence has a smoothing effect that better captures

the finite time effects of realistic interventions on transmissibility. A similar generalisation was introduced in [50].

The effective reproduction number is derived from the basic reproduction number R_0 , which describes how many people an infected individual is expected to infect in a fully susceptible population. When an NPI is introduced the basic reproduction number R_0 is multiplicatively changed by a factor c_t , yielding

$$R_{0t} = R_0 c_t. \quad (3)$$

We consider an action space with three possible NPIs: no intervention ($c_t = 1$), limited social distancing ($c_t = 0.5$) and full lockdown ($c_t = 0.2$). While this is a simplified classification of NPI types, our control framework allows for more possible actions to be easily modelled (e.g., we can introduce an arbitrary number of c_t options).

Although in reality the factor c_t is unknown and needs to be estimated, we assume throughout this study, that the effect of any NPI on the reproduction number is known and without uncertainty. Our framework does allow for the inclusion of uncertainty on these effects and we present analyses under stochastically varying c_t within our results. However, our goal is not to describe precisely how interventions attenuate transmissibility, but instead to derive insights into how surveillance quality influences the optimal timing and application of interventions generally. Consequently, our choices of c_t are sensible (i.e., values are within realistic ranges from the literature) but not specific. The actual effect of an intervention can vary with location, context and socio-demographic structure. Generally, we do not include the uncertainty on c_t to isolate the influence of the surveillance noise. However, if specific estimates of c_t or its uncertainty are available (e.g., from [5,23]) or derived from auxiliary data, these can be seamlessly integrated within our framework for more precise results.

Optimal model predictive control (MPC) of epidemics

Our study focuses on epidemic control at the early stages of an epidemic with no or limited immunity in the population and without any available pharmaceutical remedies. Consequently, the main control measures are NPIs, such as social distancing, mask-wearing, stay-at-home orders, business closures and travel restrictions [22]. These measures limit disease transmission with the aim of reducing the likely numbers of severe infections below health-care capacities and minimising expected morbidity and mortality. However, these benefits must be balanced against the costs induced by those NPIs, which may include economic downturns and loss of livelihood.

An ideal control policy would keep the incidence of new infections at a manageable (target) level, optimally balancing the costs of treating infections and implementing interventions. However, this is non-trivial both because disease dynamics can change in real time and our ability to track those changes is strongly limited by the quality of available surveillance data. To achieve this, we propose a *model predictive control (MPC)* framework for optimising epidemic interventions based on real-time incidence data. MPC utilises a mathematical model to project the dynamical behaviour of the controlled system [51].

The MPC algorithm we propose aims to curb disease spread, jointly minimising the risks and costs arising from infections and applied interventions. We outline our control framework in panel (a) of Fig 1, which consists of the following elements: a plant (the controlled system) with observable states, state-transition probabilities, an agent with an action space defining possible control actions, and a reward function.

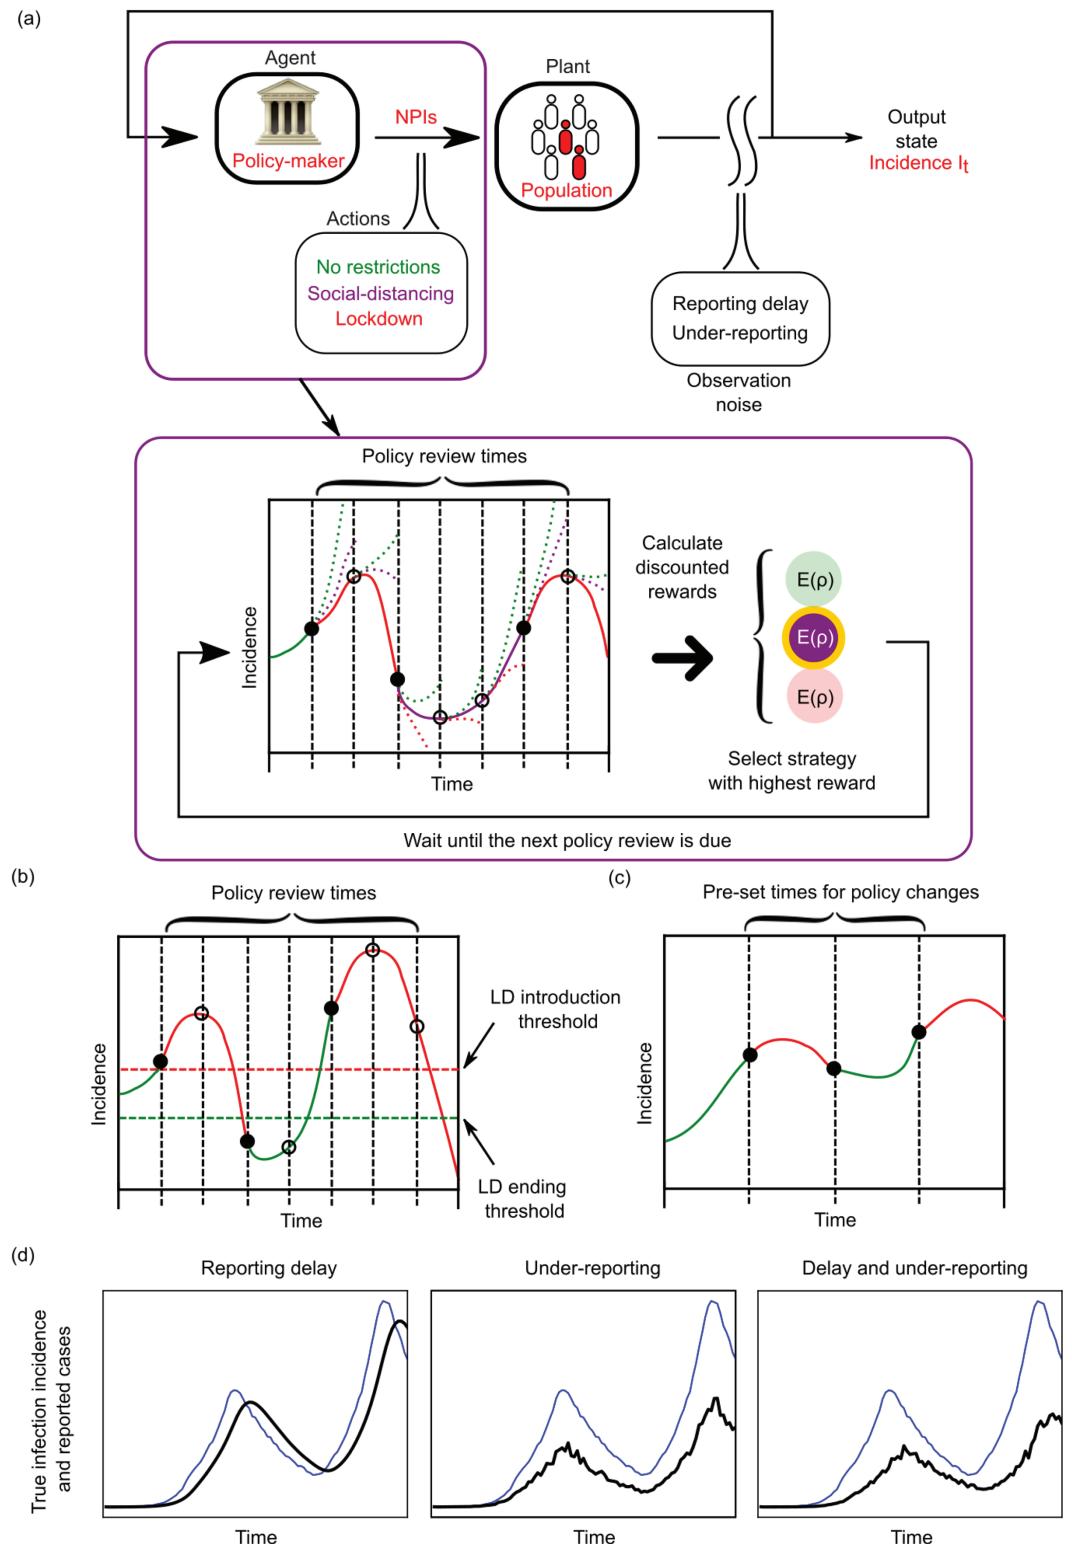


Fig 1. Panel (a): Schematic diagram of model predictive control for optimising epidemic interventions. The top chart introduces the elements of the feedback loop where the actions of the agent are chosen according to the incidence of new infections (or a proxy such as new cases) which is the monitored output state. The highlighted panel explains the model predictive method for selecting the optimal NPI from the action space which is based on using short-horizon projections and the expected reward, $E(\rho)$, under those projections for various strategies. The strategy with the highest reward (minimum cost) is implemented. The reward or cost here usually depends on how far the epidemic state is from our desired objectives, which may aim at jointly reducing both severe epidemic outcomes and intervention intensity. Panels (b) and (c) show

event-triggered and time-triggered alternative strategies, with NPIs implemented or relaxed based on incidence thresholds (the event trigger) or according to a pre-defined schedule (the time trigger). Panel (d) illustrates how realistic surveillance imperfections such as reporting delay and under-reporting distort the true incidence of infections into the incidence of cases, which we practically must use to inform decision-making. Thick black curves represent the reported cases, while thin blue curves indicate true infection incidence. Reporting delay manifests (approximately) as a time-lag with respect to the true incidence curve, whereas under-reporting results in a stochastic downscale of the incidence curve along the vertical axis.

<https://doi.org/10.1371/journal.pcbi.1013426.g001>

Our algorithm is analogous to a Markov Decision Process [52] in which the plant is the population where the disease is spreading, while the output state monitored is the incidence (number of daily new infections) I_t . The state transition probabilities, i.e. the probability of transitioning to any I_{t+1} from any given I_t are not explicitly defined, but are implicitly determined by the Poisson distribution of the renewal model (see Eq (1)). The control framework we use here largely overlaps with Markov decision processes (MDPs) [53]. However, the renewal model utilises both the immediate and past incidence. This is not exactly Markov but may be reconfigured into an MDP if higher dimensional state spaces are used [54,55].

The agent in the context of an epidemic is the public health policy-maker i.e., the individual or group responsible for proposing or removing NPIs, while the action-space comprises the possible NPI choices. We consider 3 levels of interventions that we class as *no intervention*, *social distancing* and *full lockdown*. This broadly models stepped interventions which were common across the COVID-19 pandemic. These include the three tier system that England used to enforce localised NPIs in 2020, the 4-level alert system applied by New Zealand and related policies taken by Italy, France, Canada and others [56,57]. Our framework computes decisions based on the projected reward over a fixed time-horizon which incorporates the costs of possible actions in our decision space and their risks in terms of expected infections.

In the reward function, we account for costs arising from the economic impact of NPIs and the risks associated with high numbers of incident infections. We consider a target incidence level that defines some manageable infection level and define the absolute error $I_{\text{err}} = |I_t - I_{\text{target}}|$. This target may relate to healthcare capacities e.g., setting a level of incidence such that the expected hospitalisations resulting from that incidence do not overwhelm healthcare resources. Although, studies rarely consider a target incidence level, our aim is to understand and characterise the intervention tradeoffs (e.g., timing choices) that can jointly limit expected infections and the costs of those interventions.

Setting $I_{\text{target}} = 0$ is analogous to an elimination target, which models the broad aims of pandemic policies employed by New Zealand [58] and China [59], for example. An $I_{\text{target}} > 0$ recognises that elimination is difficult, particularly in the face of infection reintroductions and so refocuses on stabilising healthcare burdens to sustainable levels that balance the supply and demand of health resources. Additionally, as we want to minimise the risk of large infection peaks and overshoots, our reward function also includes a penalty term ϕ_{over} that activates when $I_t > 1.5I_{\text{target}}$ but is zero otherwise.

While regulating disease spread within the limits of healthcare capacities is of paramount importance, interventions that restrict mobility or close businesses and trade generate substantial economic and other costs. We model this with a term ϕ_t attached to every element of the action-space. There is no cost under no restrictions and the cost of full lockdown is assumed to be 15-times larger than that of limited social distancing. While some studies into COVID-19 NPIs suggest stringent interventions are 5–6 times more costly than more limited measures [9], our factor was chosen to more markedly distinguish between our two NPI tiers

so that general qualitative insights could be better derived. Including all the above components, the reward function on day t is calculated as the negative quantity

$$r_t = -\delta|I_t - I_{\text{target}}| - \phi_t - \phi_{\text{over}}. \quad (4)$$

The agent's task is to choose the action which maximises the expected reward. We note that there are alternative reward (or cost) functions that may also be capable of achieving the same target (e.g. using different norms of the error from the target instead of the absolute error in the reward function). Although the assessment of different cost functions is out of the scope of this study, Eq (4) is representative of the various elements that may shape realistic decision-making, with its components drawn from or inspired by [9,37,44,60]. We also consider practical limitations to decision-making. While we use daily incidence data to inform our epidemic model, we allow policy review to only occur every 7 or 14 days, i.e., the agent can only change control actions every week or fortnight. The time between policy updates is t_{rev} and reflects practical intervention constraints, e.g., both in terms of logistics and ensuring compliance, policymakers may not want to switch NPIs any faster than weekly. We also impose a practical constraint on reward optimisation by considering only finite time horizons for assessing the costs of any action. We denote this projection horizon t_{proj} . This models the fact that only short-term forecasts are known to be reliable for epidemic decision-making [61].

The agent calculates the expected reward for each action by simulating the epidemic with all possible control states until t_{proj} and taking the total temporally discounted reward

$$\rho = \sum_{s=0}^{t_{\text{proj}}} r_{t+s} \gamma^s, \quad (5)$$

where $\gamma < 1$ is the temporal discount factor. A higher γ means that the agent is more concerned about long-term rewards, whereas a smaller γ means that shorter term benefits are emphasised. Since the epidemic dynamics are stochastic, the total discounted reward ρ is probabilistic. We therefore compute the expected total reward $E(\rho)$ over an ensemble of simulations. This also allows us to factor in the intrinsic variability of the epidemic generating process (e.g., from the random times between infections). This joint target-cost optimisation process is iteratively done in real-time via the feedback loop in Fig 1 and makes use of short-term projections with a receding horizon [62].

If the projection horizon is longer than the policy review period, i.e., $t_{\text{proj}} > t_{\text{rev}}$, then, we can also propose sequences of actions over the projection horizon to be taken by the agent. We may then compute the expected reward for each action sequence but only implement the first action of the sequence with the best projected reward whilst subsequent actions are reconsidered at following policy revisions. However, for the scenarios we consider, we found that this approach increases computational complexity but does not improve performance. Consequently, for these longer projections, we maintain our original approach of only evaluating single possible actions and their consequences across the horizon. We tuned the projection horizon t_{proj} and the control gain δ by Bayesian optimisation to achieve the best performance without surveillance noise. We fitted the Gaussian process surrogate model using the mean total reward from 100 simulations for time-horizons of 21 weeks for each parameter combination we surveyed and used the surrogate model to find the optimal projection horizon and control gain values. We collect all the key epidemic and control algorithm parameters in Table 1.

Table 1. Parameters of the epidemic model and the control algorithms for COVID-19 and Ebola virus.

Parameter name		COVID-19	Ebola virus	Unit
Basic reproduction number	R_0	3.5	2.5	1
Mean generation time	$\alpha_\tau \beta_\tau$	6.5	15.0	days
Generation time variance	$\alpha_\tau \beta_\tau^2$	13.65	31.5	days ²
Infection incidence error gain	δ	0.00026	0.00065	$\frac{1}{\text{cases/day}}$
Projection window ($t_{\text{rev}} = 7$ days)	t_{proj}	12	12	days
Projection window ($t_{\text{rev}} = 14$ days)	t_{proj}	14	14	days
Cost of no intervention	ϕ_t	0	0	1
Cost of social distancing	ϕ_t	0.01	0.01	1
Cost of full lockdown	ϕ_t	0.15	0.15	1
Cost of large overshoots	ϕ_{over}	5.0	5.0	1
Cost discount factor	γ	0.95	0.95	1
Threshold for imposing full lockdown	C_{LD}	2500	3000	cases/day
Threshold for relaxing full lockdown	C_{relax}	1500	3500	cases/day
Lockdown length in cycle	t_{LD}	45	36	days
No intervention length in cycle	t_{relax}	9	21	days

<https://doi.org/10.1371/journal.pcbi.1013426.t001>

Alternative control strategies

We compare the performance of the above proposed optimal control algorithm with two simpler control strategies: an *event-triggered* feedback control and a cyclic *time-triggered* control strategy. These strategies represent two fundamental approaches for controlling epidemics in real time, which have either been applied or proposed earlier. Similar event and time-triggered strategies also have a wider application across engineering and biology [63–65]. Note that for all strategies that we consider, we allow tuning so that the strategies can stabilise the observed incidence.

The event-triggered control applies or relaxes lockdowns whenever reported incidence crosses a predefined threshold, which is often heuristically set in practice. This crossing constitutes an event. We illustrate this approach in panel (b) of Fig 1 for an example incidence time series. Event-triggered control approaches have previously been used to enact interventions, e.g. for influenza [66] and were considered for triggering NPIs to suppress COVID-19 in the UK [4]. Although this strategy applies limited feedback based on the most recent incidence, it is unable to leverage the information in the full past time series or assess the likely future outcomes of its decisions.

In contrast, the cyclic time-triggered control strategy (see panel (c) of Fig 1) implements a pre-defined sequence of actions, which is not based on any direct feedback from the epidemic dynamics. The periods with full lockdown or no restrictions can be arbitrarily long, e.g. a 20/10 cyclic control strategy repeats 20 days of full lockdown followed by 10 days of no restrictions. This strategy was proposed as an effective means of COVID-19 control when surveillance data are poor quality and hence unreliable for informing decisions [40].

Surveillance noise and uncertainty in incidence data

Ideally, the agent would make decisions about possible NPIs based on the infection incidence in the population. Unfortunately, infection data are rarely available and a proxy such as the incidence of confirmed cases or deaths is commonly used. We focus here on the daily incidence of cases C_t but note that other proxies have analogous descriptions [50,67]. These proxies are commonly subject to practical surveillance imperfections, which we define via a

stochastic reporting delay τ and a stochastic reporting rate ν . In our model of the surveillance imperfections, the true infection incidence data I_t is first distorted by delay, then we consider under-ascertainment of the delayed cases (see Fig 2).

Reporting delay describes the lag between an infection and its proxy. For cases this includes latencies such as the time taken between infection and presenting symptoms or confirmation via testing. In our framework, we model the reporting delay for a single case using a Gamma distribution

$$\tau \sim \text{Gamma}(\alpha_\tau, \beta_\tau), \quad (6)$$

with shape and scale factors α_τ and β_τ , respectively. The mean reporting delay is then $\alpha_\tau \beta_\tau$ while the variance is $\alpha_\tau \beta_\tau^2$. To control the mean delay τ_{mean} and dispersion α directly we re-parametrise the distribution by the choice $\alpha_\tau = \alpha$ and $\beta_\tau = \tau_{\text{mean}}/\alpha$. We plot the probability density functions used to model reporting delay are visualised in the top row of Fig S3 in the Supplement. The cases \tilde{C}_t reported with delay on day t then result from a weighted sum of past incidence at day s and the probability that it takes $t-s$ time units before those infections are reported or confirmed as cases. This follows as

$$\tilde{C}_t \sim p^\tau(s) = \text{Pois}\left(\sum_{s=0}^t I_s w_{t-s}^\tau\right), \quad (7)$$

where the weight factors w_{t-s}^τ are derived from the Gamma distribution of the reporting delay, $w_{t-s}^\tau = p^\tau(t-s)/\sum_{k=0}^t p^\tau(k)$.

The reporting rate models incomplete or under-reporting, which captures the fact that proxies commonly represent only a fraction of infections. For example, asymptomatic and less severe infections are unlikely to appear as cases. This means that only a fraction ν of the delayed infection incidence \tilde{C}_t is reported

$$C_t = \nu_t \tilde{C}_t, \quad (8)$$

with $\nu_t \in [0, 1]$. In our model, we assume that the number of reported cases follows a Beta-binomial distribution

$$C_t \sim \text{BetaBin}(\tilde{C}_t, \alpha_\nu, \beta_\nu). \quad (9)$$

Consequently, the expected number of reported cases is $\tilde{C}_t \alpha_\nu / (\alpha_\nu + \beta_\nu)$ while the variance is $\tilde{C}_t \alpha_\nu \beta_\nu (\alpha_\nu + \beta_\nu + \tilde{C}_t) / ((\alpha_\nu + \beta_\nu)^2 (\alpha_\nu + \beta_\nu + 1))$. We refer to the ratio of the expected reported cases and the true infection incidence as the mean reporting ratio $\nu := E[C_t / \tilde{C}_t]$

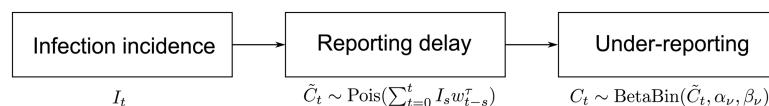


Fig 2. Models of realistic epidemic surveillance. The true infection incidence data I_t is first distorted by a probabilistic delay modelled by a convolution with w_{t-s}^τ , which are probabilities from a Gamma distribution. Under-ascertainment then occurs by downsampling these delayed cases \tilde{C}_t using a Beta-binomial distribution. This yields the reported daily cases C_t , which is frequently used as a proxy for the unobservable I_t . In some simulations, we turn either reporting delay or under-reporting off. If there is no reporting delay, $\tilde{C}_t = I_t$ and similarly, if there is no under-reporting $C_t = \tilde{C}_t$.

<https://doi.org/10.1371/journal.pcbi.1013426.g002>

which is constant in time. To directly control the mean reporting ratio ν_{mean} and the dispersion a we choose $\alpha_\nu = a$ and $\beta_\nu = a(1 - \nu_{\text{mean}})/\nu_{\text{mean}}$. Here the distribution of the reporting rate depends on the true number of cases. In order to visualise the reporting rate independently from the number of infections, we show the probability density functions of the equivalent Beta distributions of the infection reporting rate in the bottom row of [Fig S3](#) in the Supplement.

In some cases, we investigate the isolated effect of reporting delay or under-reporting, which means that for these simulations the other source of surveillance imperfections is turned off. If there is no reporting delay, $\tilde{C}_t = I_t$ and similarly, if there is no under-reporting $C_t = \tilde{C}_t$. These stochastic delay [68] and under-reporting [69] models have been widely used to describe surveillance noise in the literature, as well as serve as the starting point for deconvolution and nowcasting methods that attempt to correct for these noise sources [50,70–72].

Estimation of the reproduction number

When projecting likely infections (or proxies) over a horizon in our algorithm, we assume knowledge of the effect of NPIs on the reproduction number. This is captured by the coefficients c_t . However, we do not assume knowledge of the true basic reproduction number and so must estimate this quantity from past data.

We start by inferring the time-varying effective reproduction number by applying the formula [49]

$$R_t \approx \frac{\sum_{s=1}^{t_{\text{est}}} I_{t-s}}{\sum_{s=1}^{t_{\text{est}}} \Lambda_{t-s}}, \quad (10)$$

where Λ is the total infectiousness and is calculated as

$$\Lambda_s = \sum_{s=0}^t I_{t-s} w_s, \quad (11)$$

with weights w_s derived from the generation time distribution. Then we recover the basic reproduction number R_0 by factoring in the history of applied control measures

$$c_{\text{est}} = \frac{\sum_{t=0}^{t_{\text{est}}} c_{t-s} w_s}{\sum_{s=0}^{t_{\text{est}}} w_s} \quad (12)$$

and hence $R_{0\text{est}} = R_t/c_{\text{est}}$.

The quality of our estimates in [Eqs \(10\), \(11\)](#) and [\(12\)](#) depends on the length of the estimation window t_{est} . Short windows are more sensitive to stochastic fluctuations in incidence, while long windows over-smooth estimates and delay projections [49,73]. We apply $t_{\text{est}} = 5$ days, which appears to be a good compromise between the two extremes.

Results

Optimal MPC performance

We demonstrate the performance of our MPC algorithm first on perfectly observed incidence data and then explore the influence of practical surveillance limitations. Our aim here is to explore the best and worst case limits that surveillance induces on practical, feedback-based control. [Fig 3](#) presents four scenarios using MPC to mitigate an infectious disease with generation time (mean: 6.5 days [46,48]) and basic reproduction number ($R_0 = 3.5$ [74]) chosen

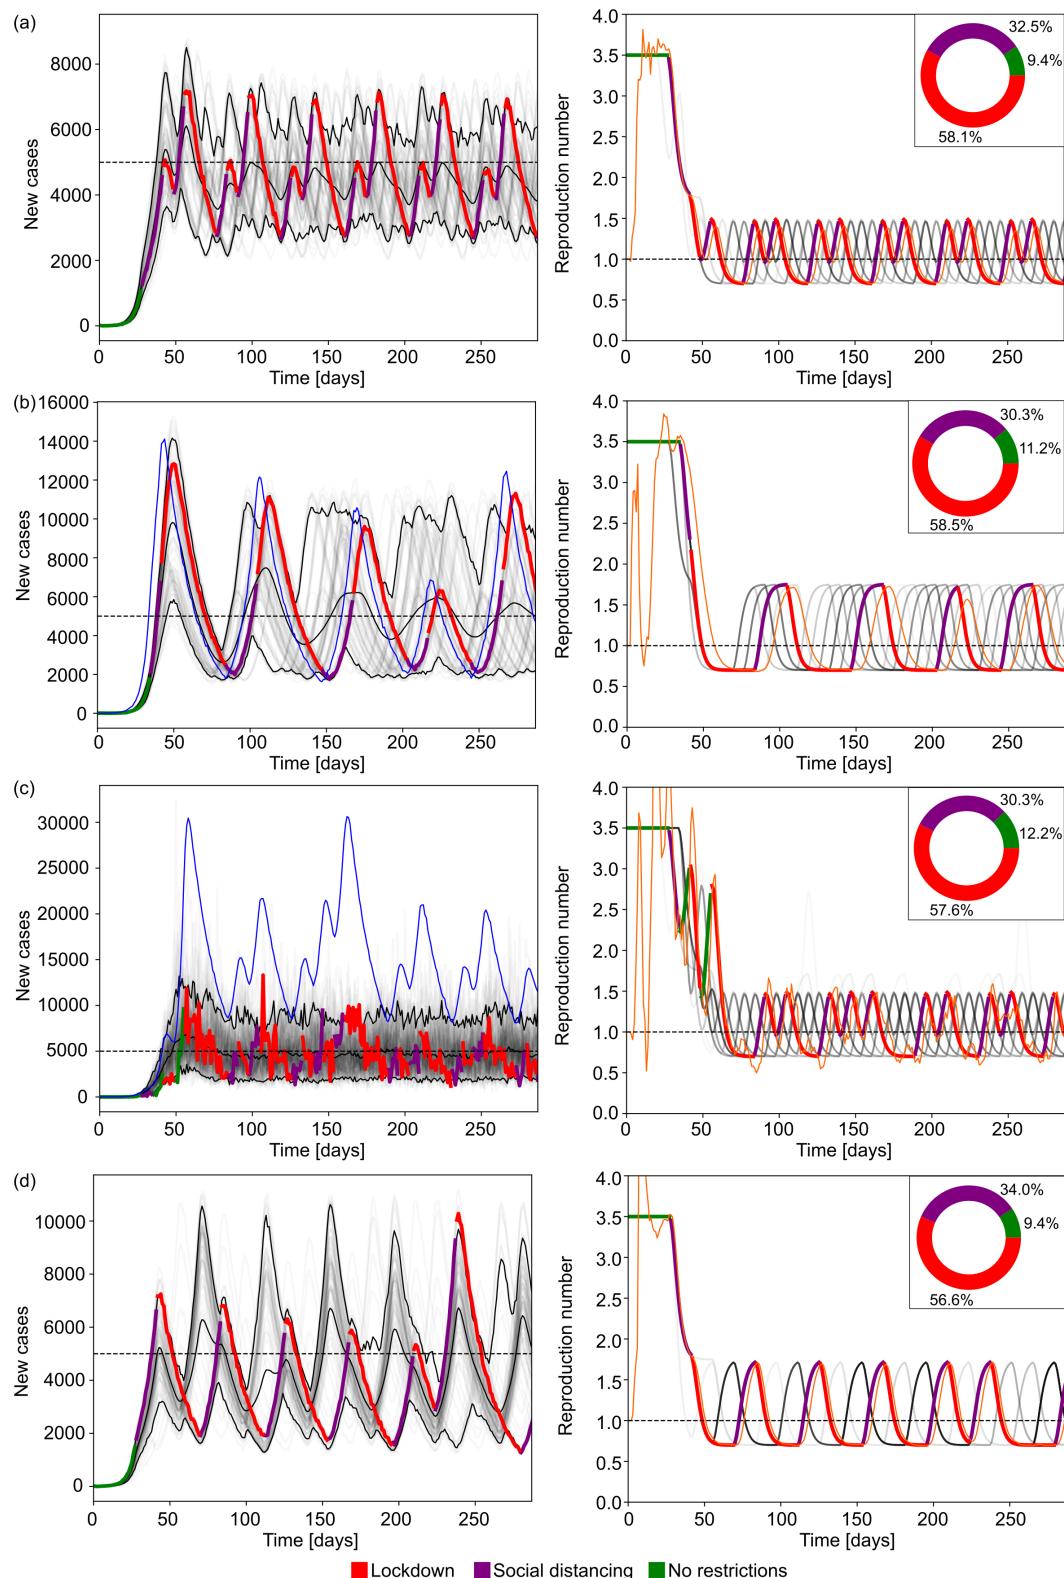


Fig 3. Simulation results with optimal epidemic control. The left panels show reported cases from ensembles of 100 simulations using generation times and reproduction numbers estimated for COVID-19. The faded curves show different individual realisations of the epidemic with the three black curves marking the 5% and 95% percentiles of the ensemble and the mean reported cases. The horizontal dashed line shows the incidence target. The highlighted thick curve of reported cases is coloured based on the NPI implemented on a given day. In rows (b) and (c), the blue thin curve indicates the true incidence corresponding to that highlighted simulation. In rows (a) and (d) where we do not simulate surveillance noise,

infection incidence are identical to the reported cases. The right column shows similar diagrams for the effective reproduction number. In that column the faded grey curves and the highlighted curve represents the true effective reproduction number whereas the thin orange curve indicates the estimated value from the reported cases in the highlighted realisation of the epidemic. The inset pie charts in the right column indicate the ratio of days spent under a given NPI across the full simulation ensemble. Row (a) represents the baseline case without delay or under-reporting with a policy review period of $t_{rev} = 7$ days. Compared to the baseline case, panel (b) shows simulations with reporting delay (mean 7 days, shape parameter $\alpha = 5.0$), panel (c) simulations with under-reporting (with mean reporting rate 0.3, dispersion $a = 8.0$). Panel (d) has no observation noise but provides simulations under an increased policy review period of $t_{rev} = 14$ days.

<https://doi.org/10.1371/journal.pcbi.1013426.g003>

to match those previously estimated for COVID-19. Row (a) shows the ideal case where the agent has access to the true incidence of infections. Here, the MPC strategy is able to keep incidence near the target incidence level (5000 new infections/day). Note that we cannot precisely track the target even in this ideal data setting due to intrinsic stochasticity in the epidemic, a finite action space of possible interventions and a policy review period that is at least a week. The resulting fluctuations about the target are a measure of the fundamental control performance under these settings.

In rows (b) and (c) we demonstrate how reporting delay or under-reporting in isolation affects control performance (cf. panel (d) in Fig 1). As indicated by the panels in the right column, imperfect monitoring of infection incidence also alters the estimated reproduction numbers leading to discrepancies between the observed and true epidemic states. Row (b) shows the effect of reporting delay with a mean of 7 days. In this case, the MPC strategy is still able to keep incidence at a manageable level, however, the delay in observing cases results in late interventions that spur higher peaks in incidence and larger fluctuations once the epidemic is under control. The thick highlighted curve shows the cases that inform the agent or decision-maker while the thin highlighted curve are the true (unknown) infections.

Row (c) shows the effect of under-reporting with a mean reporting rate of 0.3, i.e., only 30% of infections are reported as cases. Our MPC algorithm is able to achieve the target level but the true incidence fluctuates at a higher level due to the under-reporting noise process which, on average scales down infections to reported cases according to the mean reporting rate. The stochasticity of the under-reporting also occasionally misleads the MPC strategy to believe that the epidemic is under control and the incidence is below the target level or the effective reproduction number is smaller than its true value, causing higher epidemic peaks than in the ideal surveillance scenario. Interestingly, the average proportion days under each NPI across the simulation time horizon, i.e. the proportion of days spent with each NPI in place across all simulated scenarios, is similar across scenarios (a-c), indicating that the deviations in the performance of the optimal MPC strategy are due to mistimed interventions resulting from the noise in surveillance, causing either overly conservative or relaxed policies.

Row (d) in Fig 3 shows the effect of increasing the policy review period from 7 to 14 days with no observation noise. In this case, the MPC strategy is still able to keep incidence at a manageable level, however, the fluctuations in daily incidence, and consequently the peak and the bounding envelope of the later stabilised epidemic are larger as the agent has a reduced ability to intervene and update control actions. The overall effect of increasing the policy review period is similar to having a reporting delay as ultimately, both lead to delayed responses.

The limits of control due to delayed reporting

We assess how the delay in reporting limits the performance of each control strategy. We consider scenarios with different but stationary reporting delay distributions. The delay for a single infection follows a Gamma distribution $\tau \sim \text{Gamma}(\alpha_\tau, \beta_\tau)$, with shape and scale factors α_τ and β_τ , respectively. The mean reporting delay is then $\alpha_\tau \beta_\tau$ while the variance is $\alpha_\tau \beta_\tau^2$. This is a common model of reporting delays and has been used to describe surveillance of COVID-19 and Ebola virus disease among others [46,75–77]. To control the mean delay τ_{mean} and dispersion α directly we re-parametrise the distribution by the choice $\alpha_\tau = \alpha$ and $\beta_\tau = \tau_{\text{mean}}/\alpha$. This means that for a given mean reporting delay, the variance is inversely proportional to the dispersion parameter α , i.e., larger values of α correspond to more deterministic reporting delays.

We consider 6 mean reporting delays ranging from 3.5 to 21 days, each with 4 different variances. This allows us to characterise how the mean and the dispersion of the reporting delay limit optimised and heuristic control strategies. For each scenario, we run an ensemble of 1000 simulations, each with a different random seed. Fig 4 plots key results for simulations under COVID-19 disease parameters. Each column of panels depicts the performance of a different control strategy: MPC, event-triggered, and time-triggered cyclic control. The first row plots the peak incidence, the second row shows the size of the steady-state solution envelope (the range of oscillations after the epidemic is under control) and the third row charts the average costs of NPIs. When calculating the steady-state envelope, we look at the maximum and minimum of daily cases after the incidence of cases falls below the target and the effective reproduction number is below 1. It is not possible to determine a steady-state envelope for every epidemic realisation of every scenario as under some parameter combinations the control algorithm may fail to stabilise the outbreak. If this happens we use the peak value as the size of the solution envelope. We find that the distribution of peaks and steady-state envelopes are multi-modal for certain parameter combinations. This results from applying a controller that has a discrete action space and policy review time longer than the algorithm time step (which is daily).

For the comparison in Fig 4, we tuned the parameters of the different control algorithms such that they involve similar intervention costs. For the time-triggered approach, we selected a rolling policy of 45 days of full lockdown followed by 9 days of no intervention (i.e. a 45/9 cyclic policy), which is enacted 39 days after the simulated epidemic started. This is to ensure that the long term average effective reproduction number is approximately 1, so we observe the epidemic reaching a near steady-state. Arbitrary choices of policy lengths may fail to stabilise the epidemic, resulting in continued (but slower) growth or rapid suppression that is costly.

Fig 4 indicates that reporting delays have a detrimental effect to feedback-control strategies i.e., the MPC and event-triggered control approaches, as mistimed action leads to drastically higher incidence peaks and wider steady-state envelopes. In some realisations, the desired target is exceeded by a factor as large as 100. Interestingly and perhaps counter to intuition, the detrimental effect of reporting delay is less if the variance of the reporting delay is high. Low variance means that the reporting delay is almost deterministic (see results with higher values of the dispersion parameter α), while high variance implies that we have a larger portion of cases where the delay is small (and also more cases subject to very large delays, for the same overall mean lag). As a result, high variance delay distributions allow some access to more recent infection information, which can aid decision making. Row (c) of Fig 4 shows the average costs spent on interventions for each strategy under different delay settings. As

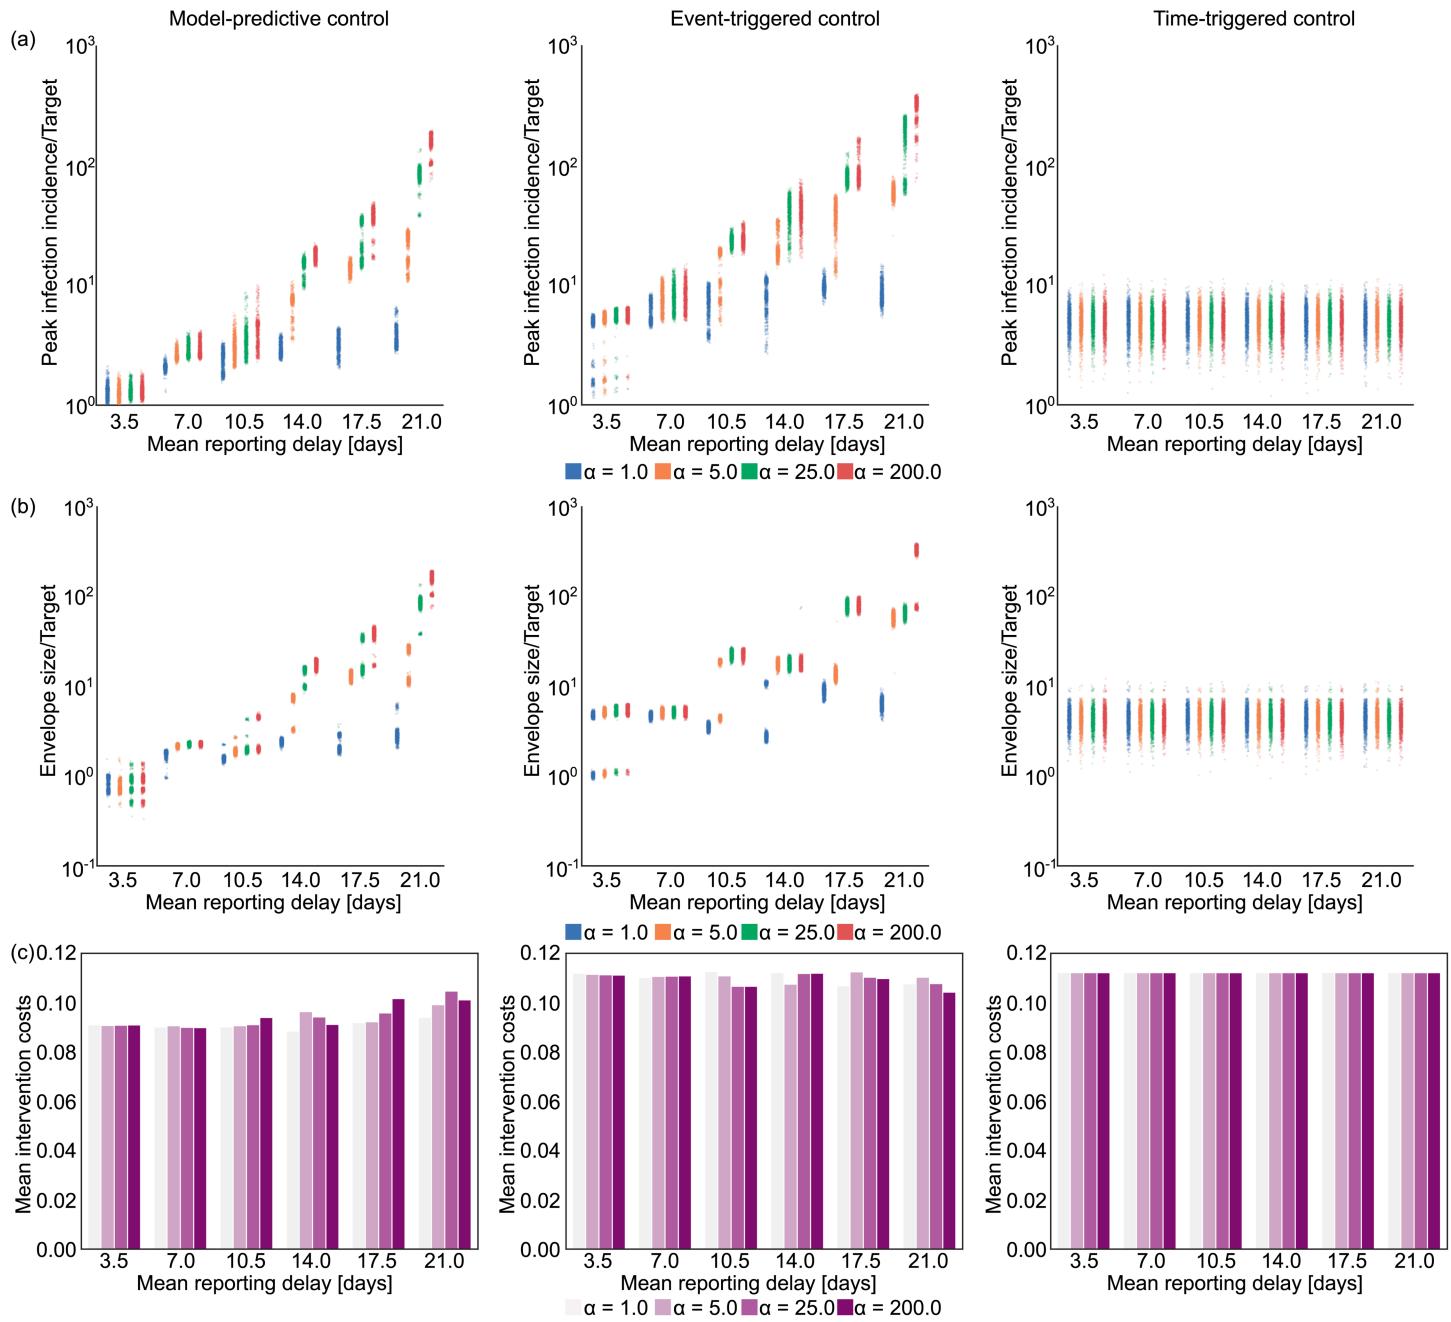


Fig 4. The impact of increasing reporting delay on optimal control. The left panels show results for our MPC algorithm, while the middle and right panels respectively present equivalent outputs under event-triggered (lockdown and relaxation thresholds of $C_{LD} = 2500$ new cases/day, and $C_{relax} = 1500$ new cases/day) and time-triggered (cycle of 45 days lockdown, 9 days of no restrictions, starting on day 38) strategies. Row (a) shows scatterplots for peak incidence and row (b) illustrates the steady-state envelope size relative to the target incidence across different time delay distributions displayed in the top row of Fig S3. The horizontal axis represents different mean reporting delays with colours depicting different dispersion levels. Larger values of the parameter α indicate more deterministic delays. Row (c) shows the mean intervention costs for each ensemble of 1000 epidemics, simulated under estimated parameters from COVID-19. An equivalent analysis for simulated Ebola virus disease epidemics is presented in Fig S1 of the Supplement.

<https://doi.org/10.1371/journal.pcbi.1013426.g004>

expected, we find the optimal control strategy is notably more economical than either the event-triggered or the time-triggered method.

Time-triggered cyclic control is insensitive to the reporting delay as it follows a pre-defined sequence of actions irrespective of the observed incidence trajectory. As a result, while the costs of the interventions are higher than those of MPC, it is possible to devise predefined intervention strategies that outperform optimal feedback-control under large case reporting delays with means of 10.5–14.0 days or more, depending on the delay dispersion. This exposes the limits that surveillance quality can impose on our ability to reactively control an epidemic. Importantly, as long as the dispersion of our reporting delays is sufficiently large ($\alpha \approx 1$) then optimal MPC strategies offer a cost-effective intervention approach for any mean delay.

In the Supplement, we also present the results of the same analysis for Ebola virus (See [Fig S1](#)). We find that the performance deterioration caused by reporting delays follows similar trends to what we observed for COVID-19 in [Fig 4](#). However, a key difference is that Ebola virus has a longer generation time, which results in slower dynamics, that reduce some of the negative impact of reporting delays. For comparison, with the MPC approach and a near deterministic ($\alpha = 200$) delay distribution with a mean of 21 days, we obtain peak infection incidence 100–1000 times larger than the target for COVID-19. In contrast, the same delay only leads to peaks of about 4–5 times the target for Ebola virus disease.

The limits of control due to under-reporting

Having isolated the influence of reporting delays, we now examine the impact of under-reporting. This is another common and important surveillance imperfection in which not all newly infected individuals are reported as cases leading to under-ascertainment of the true numbers of infections and hence the size of the epidemic. We model the number of daily reported cases with a Beta-binomial distribution $C_t \sim \text{BetaBin}(I_t, \alpha_\nu, \beta_\nu)$, with shape parameters of α_ν and β_ν . The expected number of reported cases is then $I_t \alpha_\nu / (\alpha_\nu + \beta_\nu)$ while the variance is $I_t \alpha_\nu \beta_\nu (\alpha_\nu + \beta_\nu + I_t) / ((\alpha_\nu + \beta_\nu)^2 (\alpha_\nu + \beta_\nu + 1))$. We refer to the ratio of the expected reported cases and the true infection incidence as the mean reporting ratio $\nu := E[C_t / I_t]$ which is constant in time. To directly control the mean reporting ratio ν_{mean} and the dispersion, we choose $\alpha_\nu = a$ and $\beta_\nu = a(1 - \nu_{\text{mean}}) / \nu_{\text{mean}}$. This means that larger values of the dispersion parameter a result in a smaller dispersion in reporting rate.

We consider 6 different mean reporting ratios from 0.1 to 0.85 each with 4 different variances. A smaller reporting fraction means larger under-reporting. [Fig 5](#) plots the performance of the MPC, event-triggered and cyclic strategies in successive columns. For these diagrams, the target incidence is scaled up according to the mean reporting ratio to have comparable results.

We find that as the variance in under-reporting increases the performance of optimal feedback strategies deteriorates. This follows because the under-reported case curve is effectively a stochastic downscaling of the true infection incidence curve. As a result, larger stochasticity in fluctuations for a given mean reporting rate more substantially distorts the observed incidence and misleads control. Low variance means that the under-reporting is more deterministic, so that reported case incidence better resembles the shape of the true infection incidence curve.

Both our MPC and the event-triggered strategy show similar patterns in how the peak and envelope of the controlled epidemics vary with the under-reporting statistics. However, the MPC algorithm achieves better performance with smaller intervention costs. This improvement derives from the MPC approach leveraging all the available historical information in the incidence curves. Time-triggered cyclic control is not affected by under-reporting because it

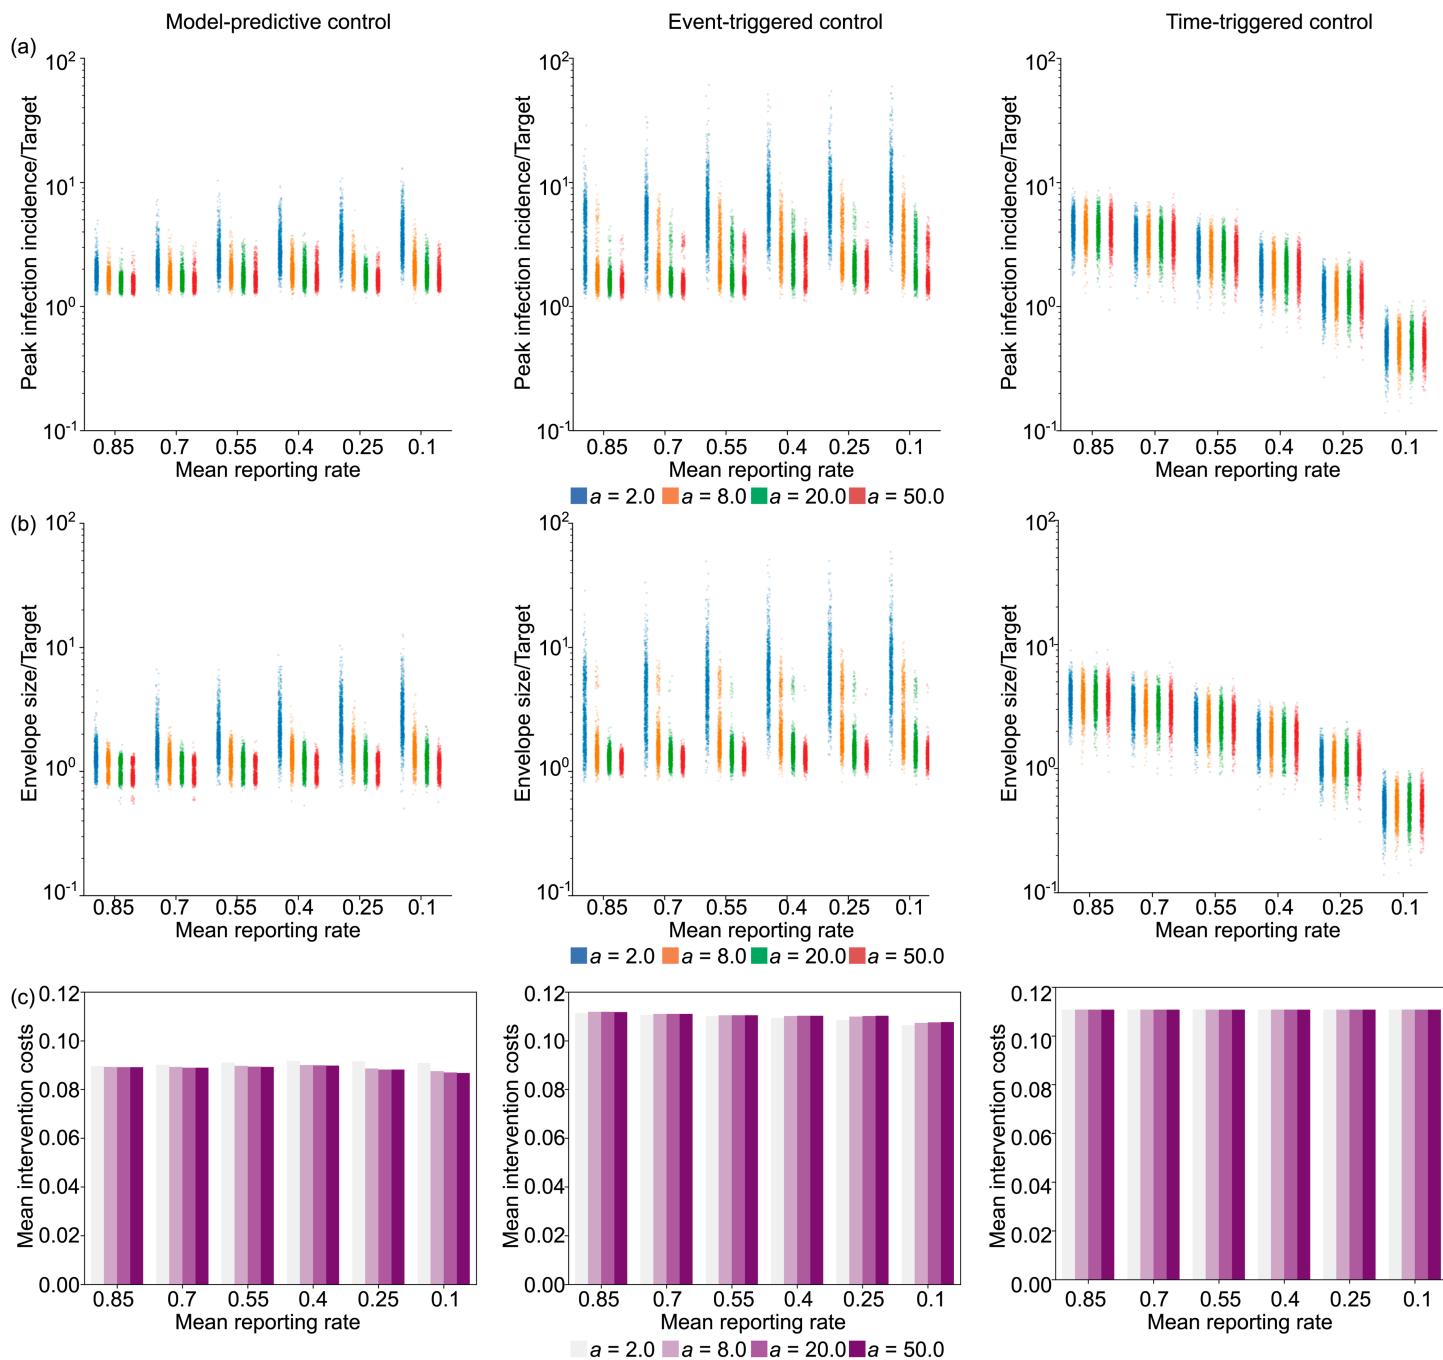


Fig 5. The impact of increasing under-reporting on optimal control. The left panels show results for our MPC algorithm, while the middle and right panels respectively present equivalent outputs under event-triggered (lockdown and relaxation thresholds of $C_{LD} = 2500$ new cases/day, and $C_{relax} = 1500$ new cases/day) and time-triggered (cycle of 45 days lockdown, 9 days of no restrictions, starting on day 38) strategies. Row (a) shows scatterplots for peak incidence, row (b) illustrates the steady-state envelope size relative to the target incidence for different reporting rate distributions displayed in the bottom row of Fig S3. The horizontal axis represents different mean reporting ratios with colours depicting different dispersion levels. Larger values of the parameter a belong to more deterministic case reporting distributions (i.e., constant reporting). Row (c) shows the mean intervention costs for each ensemble of 1000 epidemics, simulated under estimated parameters from COVID-19. An equivalent analysis for simulated Ebola virus disease epidemics is presented in Fig S2 of the Supplement.

<https://doi.org/10.1371/journal.pcbi.1013426.g005>

is agnostic to the incidence data and how it is reported. The apparent variation in the performance of the cyclic control strategy is due to the scaling of the target according to the under-reporting rate. While under moderate noise and uncertainty, our MPC approach is more effective and cost efficient than either of the heuristic methods, cyclic control can outperform MPC under extreme settings where the reporting rate is very low and the variance is high.

Comparing the effect of under-reporting on control of COVID-19 and Ebola virus (see Fig S2 in the Supplement) we find similar trends for both pathogens. Even though the drop in performance that we observe is less pronounced than that due to reporting delay, the slower dynamics of Ebola virus still makes it more controllable than COVID-19 when uncertainty in case under-reporting is present. This is evidenced by the solution peaks and steady-state envelopes being distributed in a smaller interval across the simulation ensembles for Ebola virus disease than for COVID-19.

Integrating noise and intervention frequency

Having explored the performance limits induced by reporting delays and under-reporting in isolation, we now consider their combined impact. We analyse ensembles of simulations under realistic noise distributions for COVID-19 and Ebola virus disease. The uncertainty in case reporting data can vary markedly depending on the context and national or regional differences in how surveillance is conducted. In [1], reporting delays of 9–12 days were estimated for COVID-19 in Italy, whereas [78] inferred case-reporting rates between 7–38 % across France. Based on these, we consider a mean reporting delay of 10.5 days with a dispersion of ($\alpha = 5.0$) and a mean reporting rate of 0.3 with a dispersion of $a = 8.0$.

Rows (a) and (b) in Fig 6 show results for simulated COVID-19 epidemics under realistic noise distributions and subject to policy review periods of 7 and 14 days, respectively. We find that the MPC strategy stabilises the epidemic around the target, but fluctuations are considerably larger than in the baseline case (see Fig 3a) due to the delay and under-reporting. This results in an overshoot of about 5-times the target incidence under a 7-day policy review period and around 10-times for a 14-day policy review period. Comparing these results with the simulations in panels (c) and (d) of Fig 6 for epidemics simulated under Ebola virus parameters, we observe that, the MPC strategy is more effective in controlling these epidemics and the effect of noise is less detrimental to controllability.

For the Ebola virus epidemics (panel (c)) we find that running the MPC algorithm with a 7-day policy review period causes a 3-times overshoot of the target incidence of cases, which remarkably, only slightly deteriorates to about 4-times when a 14 policy review period is used (see panel (d)). This occurs because the longer generation time of Ebola virus results in a slower epidemic that allows the MPC algorithm more time to effectively adapt to changes in the epidemic dynamics. Consequently, we must act more swiftly when responding to diseases with shorter generation times as their faster growth can quickly destabilise data-informed policies. Note that in scenarios where the same epidemic parameters were used ((a) and (b) for COVID-19, (c) and (d) for Ebola virus disease) we observe that our MPC algorithm enacted and sustained NPIs for similar time ratios (with COVID requiring more restrictions than Ebola). This confirms that differences in performance are due to timing instead of overly conservative or relaxed strategies.

Sensitivity to uncertainties and changing epidemic dynamics

In the above subsections, we characterised the performance boundaries that practical surveillance place on our optimised MPC approach. Here we examine the robustness of our algorithms to uncertainties and unanticipated changes in disease transmission and intervention

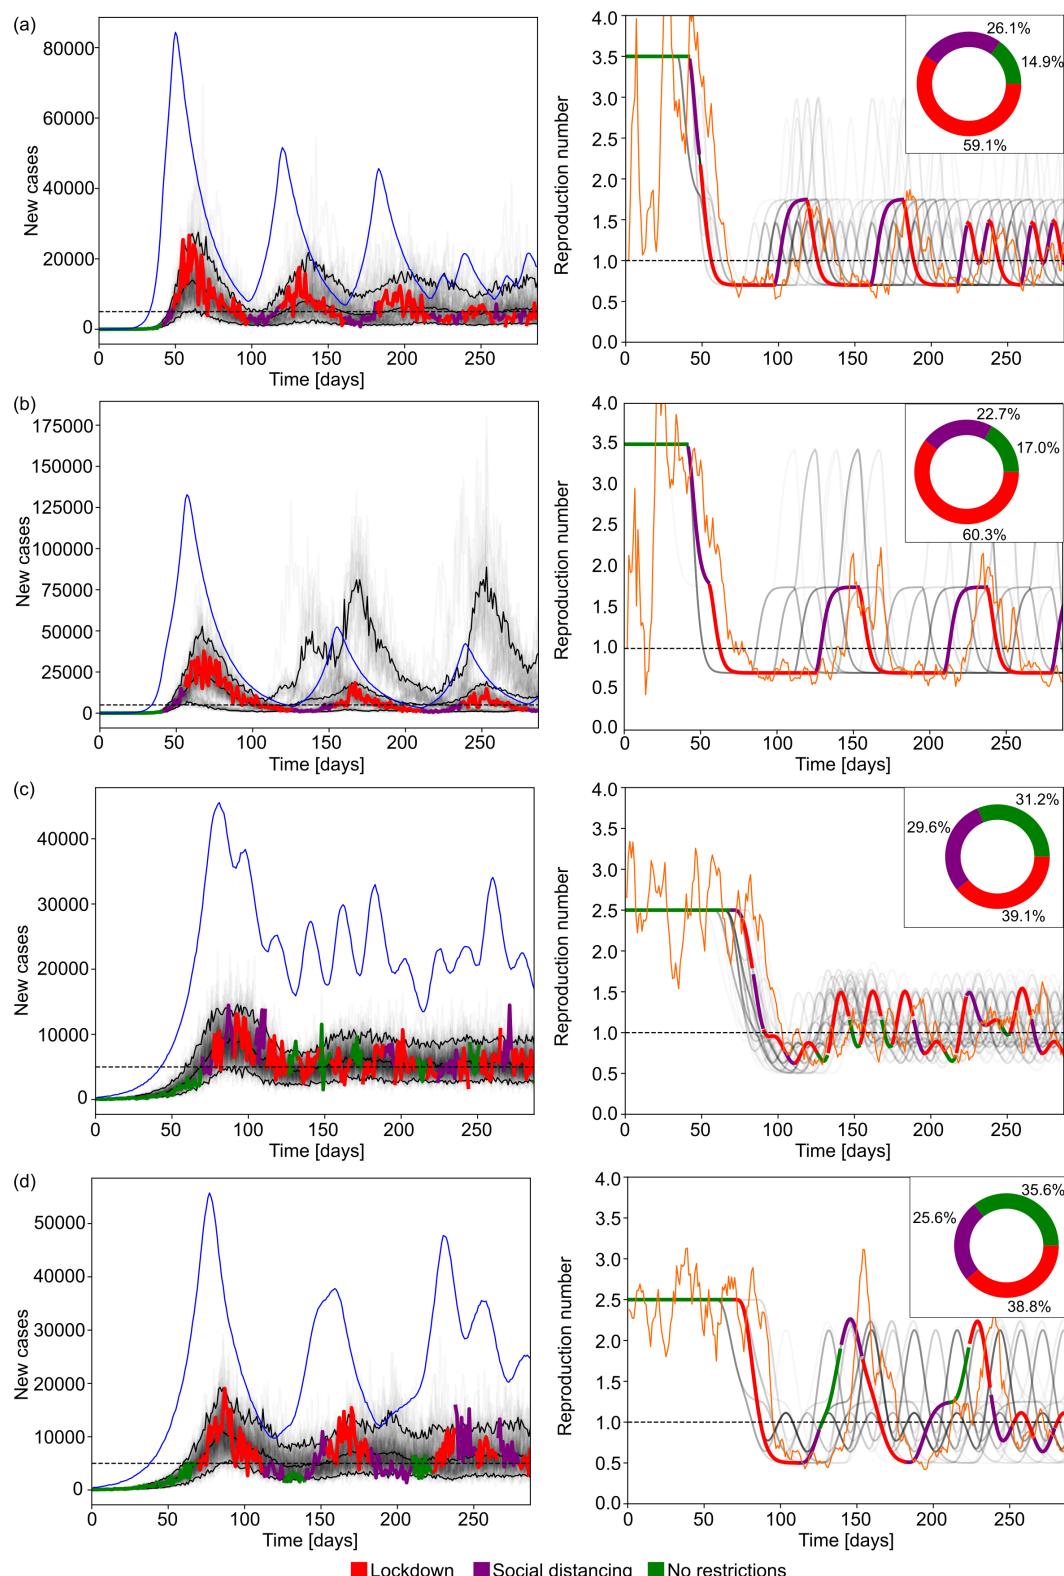


Fig 6. Optimal control and policy review for realistic simulated epidemics. Rows (a) and (b) represent epidemics simulated under a COVID-like generation time and basic reproduction number with 7 and 14 days policy review periods, respectively. Rows (c) and (d) represent epidemics simulated under Ebola-like parameters with 7 and 14 days policy review periods, respectively. Because Ebola virus disease has a longer generation time, we discard a burn in period of 14 weeks to allow the epidemic to grow towards the initial target. The left column show reported cases from ensembles of 100 simulations with mean delay 10.5 days, delay dispersion $\alpha = 5.0$, mean reporting rate 0.25, $a = 8.0$. These settings reflect

estimates of realistic surveillance noise from the literature. The faded curves show different individual realisations of the epidemic with the 3 black curves marking the 5% and 95% percentiles of the ensemble as well as the mean reported cases. The horizontal dashed line shows the incidence target. The highlighted thick curve of reported cases is coloured based on the NPI implemented on a given day. The blue thin curve indicates the true incidence corresponding to that highlighted simulation. The right column shows similar diagrams for the effective reproduction number. The faded grey curves and the highlighted curve represents the effective reproduction number whereas the thin orange curve indicates the estimated value from the reported cases for the highlighted realisation of the epidemic. The inset pie charts in the right column indicate the ratio of days spent under a given NPI across the full simulation ensemble.

<https://doi.org/10.1371/journal.pcbi.1013426.g006>

effect size. For example, if a new strain of the pathogen emerges (e.g., new variants appeared several times during the COVID-19 pandemic [79]) this can change the basic reproduction number R_0 of the epidemic and hence its controllability. The MPC algorithm has the flexibility to handle these changes in disease parameters as it infers the reproduction number from data and uses updated projections against our desired targets to dynamically adjust control policies. We present simulations in Fig 7, panel (a) in which the basic reproduction number for COVID-19 increased from 3.5 to 4.5 to model the emergence of a new, more transmissible variant. In this case, as the reproduction number is re-estimated from the most recent case counts at every policy review, the MPC algorithm retains control over the outbreak and keeps infections around the desired target.

Our MPC algorithm is also capable of accommodating uncertainty in the efficacy of NPIs i.e., unexpected variations in the overall reduction on transmission that an NPI achieves. This is more realistic than assuming a fixed reduction in transmission because difficult to model factors such as population behaviours often influence the actual efficacy of any intervention. Since our study focuses on the limiting effect of surveillance noise on feedback-control strategies, we did not include uncertainty in NPI efficacy in most of our simulations. This was necessary to isolate and expose the effect of surveillance noise on control performance. In Fig 7 panel (b), we include uncertainty around the expected effectiveness of NPIs. Here, instead of assuming fixed reduction factors, we sample from a Beta distribution where the mean is aligned with the corresponding reductions in R for each NPI except for 'No restrictions' where we keep the reproduction number at the baseline, R_0 . Similar NPI effect distributions can be derived from studies like [5,22] or [23]. Our results indicate that whilst uncertainty in NPI efficacy may occasionally lead to suboptimal action choices that cause larger peaks in infection incidence, overall the MPC algorithm still effectively controlled the epidemic by actively reacting to the discrepancies between the actual new cases and the target. This demonstrates both the robustness of our approach and the benefits of feedback control.

Discussion

Here we proposed a model predictive control strategy for optimising epidemic interventions that uses incidence data in real time. Our approach is one of the first to design feedback and cost-minimal strategies that integrate both the intrinsic stochasticity of the transmission process and the practical noise that is ubiquitous to real surveillance. Our results indicate that, within the limitations of the data quality, model predictive optimal control is a viable strategy for cost-effectively guiding intervention decisions in real time. Comparing it to earlier reference approaches, the MPC strategy appreciably outperforms both event-triggered feedback control and time-triggered control.

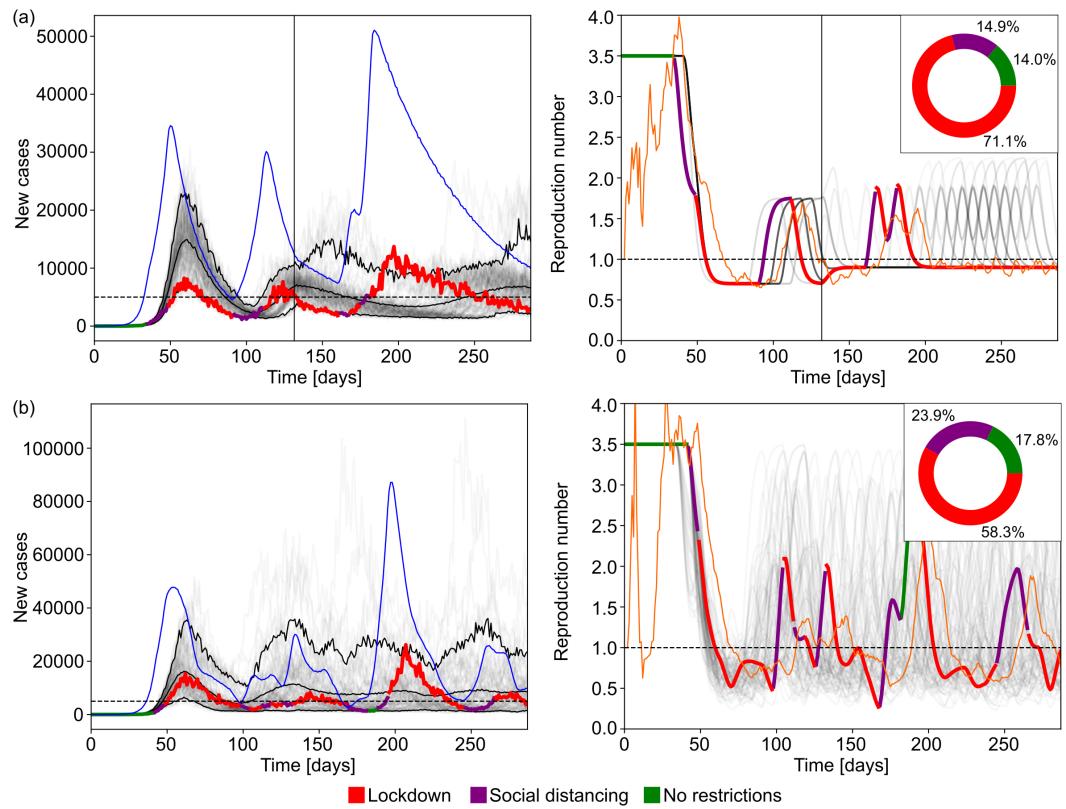


Fig 7. Optimal control for realistic simulated epidemics with uncertainty in the estimated reproduction number.

Row (a) shows epidemics simulated under COVID-like generation time where the basic reproduction number is changed from 3.5 to 4.5 on day 130 (black vertical line). Row (b) represents epidemics simulated under COVID-like generation time and a fixed basic reproduction number with uncertainty in the effect of NPIS on reduction in disease transmissions. Instead of fixed reduction factors, the factor c_t altering the basic reproduction number as $R_{0t} = R_0 c_t$ is sampled from a Beta distribution $\text{Beta}(50, 50 \times (1 - c)/c)$, where c is the mean reduction in reproduction number caused by the NPI used. The mean and variance of this Beta distribution are c and $c^2(c - 1)/(50 + c)$, respectively. The exception is 'No restrictions', which has $c_t = 1$. The left column shows reported cases from ensembles of 100 simulations with mean delay of 10.5 days, delay dispersion $\alpha = 5.0$, mean reporting rate 0.25, $a = 8.0$. The faded curves show different individual realisations of the epidemic with the three black curves marking the 5% and 95% percentiles of the ensemble as well as the mean reported cases. The horizontal dashed line shows the incidence target. The highlighted thick curve of reported cases is coloured based on the NPI implemented on a given day. The blue thin curve indicates the true incidence corresponding to that highlighted simulation. The right column shows similar diagrams for the effective reproduction number. The faded grey curves and the highlighted curve represents the effective reproduction number whereas the thin orange curve indicates the estimated value from the reported cases in the highlighted realisation of the epidemic. The inset pie charts in the right column indicate the ratio of days spent under a given NPI across the full simulation ensemble.

<https://doi.org/10.1371/journal.pcbi.1013426.g007>

Noise in surveillance data has a detrimental effect on MPC and the event-triggered approaches, impacting the quality of the real-time signals that are fed back to inform intervention choices. However, because the MPC approach considers the complete epidemic state (event-triggered control only uses recent states) and optimises decisions across stochastic projections of epidemic dynamics, it is able to better extract and leverage the information within the incidence data. This allows it to simultaneously achieve better or equivalent noise robustness while utilising a smaller intervention budget than the event triggered approach, which is limited in performance by its relatively inflexible feedback approach i.e., it imposes or relaxes NPIS on fixed thresholds.

If noise levels are extreme, for example when delays are large and relatively deterministic, time-triggered strategies, which schedule NPIs without directly considering real time data, can be more effective in limiting peak incidence as this method acts at the same preset time under all circumstances. This marks the limits of data quality for feedback-control strategies and indicates that in these scenarios the available data are fundamentally too unreliable for guiding decision-making. However, the time-triggered strategy comes with notable drawbacks because its design is heavily reliant on having accurate prior knowledge of epidemic parameters and NPI efficacies. This makes it extremely non-robust to changes and uncertainties, e.g., if transmissibility is different from what is expected, then the time-triggered approach may fail to stabilise the epidemic. Whilst MPC relies on several assumptions about transmission dynamics it does not share this lack of robustness. We therefore find strong evidence that the additional complexity, relative to reference strategies, such as event and time-triggered approaches, required to compute and perform MPC strategies brings substantial advantages. Moreover, because this optimisation is sequential it adapts well to unexpected changes or uncertainties, offering important robustness to the many unknowns during an unfolding epidemic.

For example, if a new variant emerges that changes the reproduction number, the MPC algorithm is able to adapt and bring the epidemic under control as long as there are actions (i.e., possible interventions) in the action space capable of forcing the effective reproduction number below 1. In the scenario where our strongest NPI is unable to achieve this reduction, our control algorithm will not be able to stabilise the epidemic. However, this would indicate that these measures are insufficient to control the epidemic, providing useful and timely evidence for enacting more stringent measures.

Another example of varying transmission parameters occurs when immunity is acquired by infection. This reduces the susceptible population and decreases the effective reproduction number. In our modelling framework, this can be easily included by setting $R_t = R_0 S/N$, where S and N are the susceptible and the total population sizes, respectively. Assuming perfect immunity from infection the susceptible population size is calculated as $S_t = N - \sum_{s=0}^t I_s$, which is a common method for including susceptibles in renewal models [23,80]. This would not affect the MPC algorithm's operation, as its only link to susceptible depletion is implicit, through changes in the estimated reproduction number R_t , which naturally declines as immunity builds through infection. As the projections used to determine optimal control actions are short-term, and the reproduction-number estimate is updated at every decision point, susceptible depletion has minimal impact over these time frames and so does not affect the optimisation over policy choices. Note that we excluded this from our simulations in order to isolate the impact of the NPIs and to allow fair comparison among epidemics that have dynamics on differing timescales.

We also demonstrated that even though it results in larger oscillations, overall, the MPC algorithm retains control over the epidemic even when the efficacies of NPIs are unreliable known or have substantial uncertainty. This result corroborates findings in [44] and highlights why having an adaptive strategy that considers data and action in feedback is beneficial. Furthermore, the inherent uncertainties in policy efficacy make a strong case for data-driven control and forecasting frameworks allowing for the efficient testing of various scenarios to support informed decision-making.

Our results also emphasise that while noise can degrade even optimal MPC strategies, these rigorously optimised control approaches are valuable in almost all settings. Note that we did not correct for these sources of noise when assessing their detrimental effects on epidemic controllability. In real-time epidemic analysis, reporting delays and underreporting

are typically recognised and, where possible, auxiliary data are often used to inform corrections. Our aim here, however, was not to simulate exactly what analysts should or would do in real-time, but to understand the intrinsic performance bounds that imperfections in routinely collected data impose. While several studies have focussed on estimating and compensating for under-reporting [24] and reporting delays [81–83], the additional knowledge about the reporting process or the auxiliary data required [84,85] may often be unavailable, expensive to collect or only become available later in epidemics. Consequently, we focussed on characterising performance when little else is known about the epidemic than its time series of cases.

While surveillance data corrections can improve estimates, they can also introduce additional bias if the underlying assumptions are incorrect [86,87]. The feasibility of adjusting for imperfect reporting therefore varies greatly by location and by disease. Only a few pathogens, most notably COVID-19, benefited from the dense, near-real-time monitoring needed to support rigorous corrections. Even then, some forms of under-reporting, such as infections that never manifest symptoms, remained challenging to quantify [88]. Moreover, there are also inherent practical delays involved with the public announcement and implementation of NPIs from the point of decision [89], which would result in delayed action even with perfect knowledge of the epidemic states. For example, having a policy review period that is several weeks for a pathogen with dynamics that vary on daily timescales can mean that interventions are inevitably late or suboptimal. This effectively causes an additional lag in the feedback loop and may itself prevent controllability. Our study helps understanding a range of scenarios between the best and worst cases and identifies the kinds of noise that most impair epidemic control.

We also found that the speed of disease spread is an important factor that influences both the limiting impact of surveillance imperfections and sets the required timescales for a suite of interventions to be successful. It is easier to control a slower spreading pathogen like Ebola virus disease (mean generation time ~ 15 days), as compared to COVID-19 (mean generation time ~ 6 days). Generally, for a slower spreading disease there is more tolerance for having larger reporting delays and less frequent policy reviews. Accordingly, there is also less sensitivity to the timing of interventions. The negative effect of having a lower policy review frequency (or longer time between policy updates) implies that it is ideal to review intervention decisions as often as possible (hence allowing for a more continuous feedback loop). However, because NPIs are intrusive and costly, doing so would probably drive changes in public behaviour that have a rebound effect on the effectiveness of the NPIs [90,91]. For example, the level of adherence to closure or social distancing policies may wane due to fatigue or perceived risk [92,93]. This illustrates the complex multi-loop feedback circuits that can exist and will form the subject of our future study.

Although our MPC approach is adaptive and efficient at controlling unfolding epidemics, it currently leaves several factors unmodelled. Specifically, we do not consider how dynamic changes in behaviour, mobility or other individual-level variations in response to policy and epidemic data alter transmissibility. We also assume that the population is well-mixed which means we neglect heterogeneity in transmission (e.g., superspreading) as well as spatial and sociodemographic differences that can modulate spread. However, our MPC framework can accommodate some of these sources of heterogeneity (e.g., we can easily include superspreading via more dispersed renewal models). Our future work will focus on better understanding how heterogeneity may affect intervention choices.

While the realism of our study is dependent on having accurate knowledge of relative costs (both economic and due to health outcomes), our framework is flexible and easily incorporates estimates of these quantities as well as finer resolution intervention options (e.g., we can

directly expand our action space to include NPIs with intermediate stringency such as mandatory wearing of masks, restrictions to larger meetings, or limiting in-person attendance to work or education). We also do not assess the impact of model errors beyond those stemming from imperfect surveillance and always assume that our model is parameterised using the best available knowledge for the generation time distribution and the basic reproduction number. Our goal was simply to construct a general framework that can qualitatively explore and evaluate how optimal and suboptimal but known control strategies depend on realistic surveillance limitations.

Our results unequivocally demonstrate that timing is a crucial factor in intervention efficiency. The same interventions or interventions under the same overall budget applied differently across time can yield markedly different disease control outcomes due to this sensitivity to timing. Even when optimal algorithms such as our MPC approach are applied, mistimed action (due to delays in data or irregular and slow NPI reviews), can be detrimental and cause substantial reductions in policy effectiveness, large epidemic peaks and high endemic loads. Ascertainment of infections is also important but even suboptimal strategies are more robust to this type of noise than delays. Consequently, improving the speed of epidemic detection and response systems should be a priority for disease surveillance and policy.

Supporting information

S1 Fig. The impact of increasing reporting delay on optimal control. The left panels show results for our MPC algorithm, while the middle and right panels respectively present equivalent outputs under event-triggered (lockdown and relaxation thresholds of $C_{LD} = 3000$ new cases/day, and $C_{relax} = 3500$ new cases/day) and time-triggered (cycle of 36 days lockdown, 21 days of no restrictions, starting on day 143) strategies. Row (a) shows scatterplots for peak incidence and row (b) illustrates the steady-state envelope size relative to the target incidence across different time delay distributions displayed in the top row of Fig S3. The horizontal axis represents different mean reporting delays with colours depicting different dispersion levels. Larger values of the parameter α indicate more deterministic delays. Row (c) shows the mean intervention costs for each ensemble of 1000 epidemics, simulated under estimated parameters from Ebola virus disease.

(EPS)

S2 Fig. The impact of increasing under-reporting on optimal control. The left panels show results for our MPC algorithm, while the middle and right panels respectively present equivalent outputs under event-triggered (lockdown and relaxation thresholds of $C_{LD} = 3000$ new cases/day, and $C_{relax} = 3500$ new cases/day) and time-triggered (cycle of 36 days lockdown, 21 days of no restrictions, starting on day 143) strategies. Row (a) shows scatterplots for peak incidence, row (b) illustrates the steady-state envelope size relative to the target incidence for different reporting rate distributions displayed in the bottom row of Fig S3. The horizontal axis represents different mean reporting ratios with colours depicting different dispersion levels. Larger values of the parameter a belong to more deterministic case reporting distributions (i.e., constant reporting). Row (c) shows the mean intervention costs for each ensemble of 1000 epidemics, simulated under estimated parameters from Ebola virus disease.

(EPS)

S3 Fig. Probability density functions (PDFs) of the simulated surveillance noise. Top row: case reporting delay distributions modelled using Gamma distributions. The shape parameter

α controls dispersion (larger α results in lower variance). Numbers above the curves indicate the mean of each distribution. Bottom row: infection reporting rate distributions modelled using Beta distributions. The parameter a controls dispersion (larger a results in lower variance). Numbers above the curves indicate the mean reporting rate.

(EPS)

Author contributions

Conceptualization: Kris V. Parag.

Formal analysis: Sandor Beregi.

Funding acquisition: Kris V. Parag.

Investigation: Sandor Beregi, Kris V. Parag.

Methodology: Sandor Beregi, Kris V. Parag.

Software: Sandor Beregi.

Supervision: Kris V. Parag.

Validation: Sandor Beregi, Kris V. Parag.

Visualization: Sandor Beregi.

Writing – original draft: Sandor Beregi.

Writing – review & editing: Sandor Beregi, Kris V. Parag.

References

1. Casella F. Can the COVID-19 epidemic be controlled on the basis of daily test reports?. IEEE Control Syst Lett. 2021;5(3):1079–84. <https://doi.org/10.1109/lcsys.2020.3009912>
2. Mendez-Brito A, El Bcheraoui C, Pozo-Martin F. Systematic review of empirical studies comparing the effectiveness of non-pharmaceutical interventions against COVID-19. J Infect. 2021;83(3):281–93. <https://doi.org/10.1016/j.jinf.2021.06.018> PMID: 34161818
3. Bendavid E, Patel CJ. Epidemic outcomes following government responses to COVID-19: insights from nearly 100,000 models. Sci Adv. 2024;10(23):eadn0671. <https://doi.org/10.1126/sciadv.adn0671> PMID: 38838157
4. Ferguson N, Laydon D, Nedjati Gilani G, Imai N, Ainslie K, Baguelin M, et al. Report 9: impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. 2020. <https://doi.org/10.25561/77482>
5. Brauner JM, Mindermann S, Sharma M, Johnston D, Salvatier J, Gavenčiak T, et al. Inferring the effectiveness of government interventions against COVID-19. Science. 2021;371(6531):eabd9338. <https://doi.org/10.1126/science.abd9338> PMID: 33323424
6. Thu TPB, Ngoc PNH, Hai NM, Tuan LA. Effect of the social distancing measures on the spread of COVID-19 in 10 highly infected countries. Sci Total Environ. 2020;742:140430. <https://doi.org/10.1016/j.scitotenv.2020.140430> PMID: 32623158
7. Ensorink M, Kupferschmidt K. With COVID-19, modeling takes on life and death importance. Science. 2020;367(6485):1414–5. <https://doi.org/10.1126/science.367.6485.1414-b> PMID: 32217707
8. Huberts N, Thijssen J. Optimal timing of interventions during an epidemic. SSRN Journal. 2020. <https://doi.org/10.2139/ssrn.3607048>
9. Haw DJ, Forchini G, Doohan P, Christen P, Pianella M, Johnson R, et al. Optimizing social and economic activity while containing SARS-CoV-2 transmission using DAEDALUS. Nat Comput Sci. 2022;2(4):223–33. <https://doi.org/10.1038/s43588-022-00233-0> PMID: 38177553
10. Haw DJ, Morgenstern C, Forchini G, Johnson R, Doohan P, Smith PC, et al. Data needs for integrated economic-epidemiological models of pandemic mitigation policies. Epidemics. 2022;41:100644. <https://doi.org/10.1016/j.epidem.2022.100644> PMID: 36375311

11. Gallo L, Frasca M, Latora V, Russo G. Lack of practical identifiability may hamper reliable predictions in COVID-19 epidemic models. *Sci Adv.* 2022;8(3):eabg5234. <https://doi.org/10.1126/sciadv.abg5234> PMID: 35044820
12. Jayatilleke K. Challenges in implementing surveillance tools of High-Income Countries (HICs) in Low Middle Income Countries (LMICs). *Curr Treat Options Infect Dis.* 2020;12(3):191–201. <https://doi.org/10.1007/s40506-020-00229-2> PMID: 32874140
13. Davies NG, Kucharski AJ, Eggo RM, Gimma A, Edmunds WJ, Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group. Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study. *Lancet Public Health.* 2020;5(7):e375–85. [https://doi.org/10.1016/S2468-2667\(20\)30133-X](https://doi.org/10.1016/S2468-2667(20)30133-X) PMID: 32502389
14. Miles DK, Stedman M, Heald AH. “Stay at home, protect the national health service, save lives”: a cost benefit analysis of the lockdown in the United Kingdom. *Int J Clin Pract.* 2021;75(3):e13674. <https://doi.org/10.1111/ijcp.13674> PMID: 32790942
15. Sandmann FG, Davies NG, Vassall A, Edmunds WJ, Jit M, Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group. The potential health and economic value of SARS-CoV-2 vaccination alongside physical distancing in the UK: a transmission model-based future scenario analysis and economic evaluation. *Lancet Infect Dis.* 2021;21(7):962–74. [https://doi.org/10.1016/S1473-3099\(21\)00079-7](https://doi.org/10.1016/S1473-3099(21)00079-7) PMID: 33743846
16. Lison A, Banholzer N, Sharma M, Mindermann S, Unwin HJT, Mishra S, et al. Effectiveness assessment of non-pharmaceutical interventions: lessons learned from the COVID-19 pandemic. *Lancet Public Health.* 2023;8(4):e311–7. [https://doi.org/10.1016/S2468-2667\(23\)00046-4](https://doi.org/10.1016/S2468-2667(23)00046-4) PMID: 36965985
17. Hart WS, Buckingham JM, Keita M, Ahuka-Mundeke S, Maini PK, Polonsky JA, et al. Optimizing the timing of an end-of-outbreak declaration: Ebola virus disease in the Democratic Republic of the Congo. *Sci Adv.* 2024;10(27):eado7576. <https://doi.org/10.1126/sciadv.ado7576> PMID: 38959306
18. Ash T, Bento AM, Kaffine D, Rao A, Bento AI. Disease-economy trade-offs under alternative epidemic control strategies. *Nat Commun.* 2022;13(1):3319. <https://doi.org/10.1038/s41467-022-30642-8> PMID: 35680843
19. Gubar E, Policardo L, Sánchez Carrera EJ, Taynitskiy V. On optimal lockdown policies while facing socioeconomic costs. *Ann Oper Res.* 2023;337(3):959–92. <https://doi.org/10.1007/s10479-023-05454-8>
20. Shearer FM, Moss R, McVernon J, Ross JV, McCaw JM. Infectious disease pandemic planning and response: Incorporating decision analysis. *PLoS Med.* 2020;17(1):e1003018. <https://doi.org/10.1371/journal.pmed.1003018> PMID: 31917786
21. Shea K, Borchering RK, Probert WJM, Howerton E, Bogich TL, Li S-L, et al. Multiple models for outbreak decision support in the face of uncertainty. *Proc Natl Acad Sci U S A.* 2023;120(18):e2207537120. <https://doi.org/10.1073/pnas.2207537120> PMID: 37098064
22. Haug N, Geyrhofer L, Londei A, Dervic E, Desvars-Larrive A, Loreto V, et al. Ranking the effectiveness of worldwide COVID-19 government interventions. *Nat Hum Behav.* 2020;4(12):1303–12. <https://doi.org/10.1038/s41562-020-01009-0> PMID: 33199859
23. 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):257–61. <https://doi.org/10.1038/s41586-020-2405-7> PMID: 32512579
24. Meadows AJ, Oppenheim B, Guerrero J, Ash B, Badker R, Lam CK, et al. Infectious disease underreporting is predicted by country-level preparedness, politics, and pathogen severity. *Health Secur.* 2022;20(4):331–8. <https://doi.org/10.1089/hs.2021.0197> PMID: 35925788
25. Parag KV. How to measure the controllability of an infectious disease? *Phys Rev X.* 2024;14(3):031041. <https://doi.org/10.1103/physrevx.14.031041>
26. Peak CM, Childs LM, Grad YH, Buckee CO. Comparing nonpharmaceutical interventions for containing emerging epidemics. *Proc Natl Acad Sci U S A.* 2017;114(15):4023–8. <https://doi.org/10.1073/pnas.1616438114> PMID: 28351976
27. Elliott P, Whitaker M, Tang D, Eales O, Steyn N, Bodinier B, et al. Design and implementation of a national SARS-CoV-2 monitoring program in England: REACT-1 Study. *Am J Public Health.* 2023;113(5):545–54. <https://doi.org/10.2105/AJPH.2023.307230> PMID: 36893367
28. Wang L. PID control system design and automatic tuning using MATLAB/Simulink: design and implementation using MATLAB/Simulink. Wiley-Interscience.
29. Rabath CA, Léchevin N. Discrete-time control system design with applications. Springer.
30. Stépán G, Haller G. Quasiperiodic oscillations in robot dynamics. *Nonlinear Dyn.* 1995;8(4):513–28. <https://doi.org/10.1007/bf00045711>

31. Orosz G, Wilson RE, Stépán G. Traffic jams: dynamics and control. *Philos Trans A Math Phys Eng Sci.* 2010;368(1928):4455–79. <https://doi.org/10.1098/rsta.2010.0205> PMID: 20819817
32. Kyrychko YN, Blyuss KB, Hövel P, Schöll E. Asymptotic properties of the spectrum of neutral delay differential equations. *Dynamical Systems.* 2009;24(3):361–72. <https://doi.org/10.1080/14689360902893285>
33. Sykora HT, Sadeghpour M, Ge JI, Bachrathy D, Orosz G. On the moment dynamics of stochastically delayed linear control systems. *Intl J Robust & Nonlinear.* 2020;30(18):8074–97. <https://doi.org/10.1002/rnc.5218>
34. Young L-S, Ruschel S, Yanchuk S, Pereira T. Consequences of delays and imperfect implementation of isolation in epidemic control. *Sci Rep.* 2019;9(1):3505. <https://doi.org/10.1038/s41598-019-39714-0> PMID: 30837533
35. Albi G, Pareschi L, Zanella M. Control with uncertain data of socially structured compartmental epidemic models. *J Math Biol.* 2021;82(7):63. <https://doi.org/10.1007/s00285-021-01617-y> PMID: 34023964
36. Mejía C, Salazar E, Camacho O. A comparative experimental evaluation of various Smith predictor approaches for a thermal process with large dead time. *Alexandria Engineering Journal.* 2022;61(12):9377–94. <https://doi.org/10.1016/j.aej.2022.03.047>
37. Britton T, Leskelä L. Optimal intervention strategies for minimizing total incidence during an epidemic. *SIAM J Appl Math.* 2023;83(2):354–73. <https://doi.org/10.1137/22m1504433>
38. Morris DH, Rossine FW, Plotkin JB, Levin SA. Optimal, near-optimal, and robust epidemic control. *Commun Phys.* 2021;4(1). <https://doi.org/10.1038/s42405-021-00570-y>
39. Meidan D, Schulmann N, Cohen R, Haber S, Yaniv E, Sarid R, et al. Alternating quarantine for sustainable epidemic mitigation. *Nat Commun.* 2021;12(1):220. <https://doi.org/10.1038/s41467-020-20324-8> PMID: 33431866
40. Bin M, Cheung PYK, Crisostomi E, Ferraro P, Lhachemi H, Murray-Smith R, et al. Post-lockdown abatement of COVID-19 by fast periodic switching. *PLoS Comput Biol.* 2021;17(1):e1008604. <https://doi.org/10.1371/journal.pcbi.1008604> PMID: 33476332
41. Di Lauro F, Kiss IZ, Miller JC. Optimal timing of one-shot interventions for epidemic control. *PLoS Comput Biol.* 2021;17(3):e1008763. <https://doi.org/10.1371/journal.pcbi.1008763> PMID: 33735171
42. Tran Kiem C, Crépey P, Bosetti P, Levy Bruhl D, Yazdanpanah Y, Salje H, et al. Lockdown as a last resort option in case of COVID-19 epidemic rebound: a modelling study. *Euro Surveill.* 2021;26(22):2001536. <https://doi.org/10.2807/1560-7917.ES.2021.26.22.2001536> PMID: 34085634
43. Di Lauro F, Kiss IZ, Rus D, Della Santina C. Covid-19 and flattening the curve: a feedback control perspective. *IEEE Control Syst Lett.* 2020;5(4):1435–40. <https://doi.org/10.1109/LCSYS.2020.3039322> PMID: 37974563
44. van Heusden K, Stewart GE, Otto SP, Dumont GA. Effective pandemic policy design through feedback does not need accurate predictions. *PLOS Glob Public Health.* 2023;3(2):e0000955. <https://doi.org/10.1371/journal.pgph.0000955> PMID: 36962799
45. Grassly NC, Fraser C. Mathematical models of infectious disease transmission. *Nat Rev Microbiol.* 2008;6(6):477–87. <https://doi.org/10.1038/nrmicro1845> PMID: 18533288
46. Parag KV, Donnelly CA. Fundamental limits on inferring epidemic resurgence in real time using effective reproduction numbers. *PLoS Comput Biol.* 2022;18(4):e1010004. <https://doi.org/10.1371/journal.pcbi.1010004> PMID: 35404936
47. Van Kerkhove MD, Bento AI, Mills HL, Ferguson NM, Donnelly CA. A review of epidemiological parameters from Ebola outbreaks to inform early public health decision-making. *Sci Data.* 2015;2:150019. <https://doi.org/10.1038/sdata.2015.19> PMID: 26029377
48. Chen D, Lau Y-C, Xu X-K, Wang L, Du Z, Tsang TK, et al. Inferring time-varying generation time, serial interval, and incubation period distributions for COVID-19. *Nat Commun.* 2022;13(1):7727. <https://doi.org/10.1038/s41467-022-35496-8> PMID: 36513688
49. Cori A, Ferguson NM, Fraser C, Cauchemez S. A new framework and software to estimate time-varying reproduction numbers during epidemics. *Am J Epidemiol.* 2013;178(9):1505–12. <https://doi.org/10.1093/aje/kwt133> PMID: 24043437
50. Azmon A, Faes C, Hens N. On the estimation of the reproduction number based on misreported epidemic data. *Stat Med.* 2014;33(7):1176–92. <https://doi.org/10.1002/sim.6015> PMID: 24122943
51. Schwenzler M, Ay M, Bergs T, Abel D. Review on model predictive control: an engineering perspective. *Int J Adv Manuf Technol.* 2021;117(5–6):1327–49. <https://doi.org/10.1007/s00170-021-07682-3>
52. Uther W. *Markov Decision Processes.* Springer US. 2011. p. 642–6.

53. Toussaint M, Storkey A. Probabilistic inference for solving discrete and continuous state Markov decision processes. In: Proceedings of the 23rd International Conference on Machine Learning - ICML '06. ACM Press; 2006.
54. Baker CTH. Retarded differential equations. *Journal of Computational and Applied Mathematics*. 2000;125(1–2):309–35. [https://doi.org/10.1016/s0377-0427\(00\)00476-3](https://doi.org/10.1016/s0377-0427(00)00476-3)
55. Beregi S, Takacs D, Stepan G. Bifurcation analysis of wheel shimmy with non-smooth effects and time delay in the tyre–ground contact. *Nonlinear Dyn*. 2019;98(1):841–58. <https://doi.org/10.1007/s11071-019-05123-1>
56. Smith LE, Potts HWW, Amlot R, Fear NT, Michie S, Rubin GJ. Tiered restrictions for COVID-19 in England: knowledge, motivation and self-reported behaviour. *Public Health*. 2022;204:33–9. <https://doi.org/10.1016/j.puhe.2021.12.016> PMID: 35144152
57. Campbell M, Marek L, Wiki J, Hobbs M, Sabel CE, McCarthy J, et al. National movement patterns during the COVID-19 pandemic in New Zealand: the unexplored role of neighbourhood deprivation. *J Epidemiol Community Health*. 2021;75(9):903–5. <https://doi.org/10.1136/jech-2020-216108> PMID: 33727245
58. Kung S, Hills T, Kearns N, Beasley R. New Zealand's COVID-19 elimination strategy and mortality patterns. *Lancet*. 2023;402(10407):1037–8. [https://doi.org/10.1016/S0140-6736\(23\)01368-5](https://doi.org/10.1016/S0140-6736(23)01368-5) PMID: 37634521
59. Tang J-L, Abbasi K. What can the world learn from China's response to covid-19?. *BMJ*. 2021;375:n2806. <https://doi.org/10.1136/bmj.n2806> PMID: 34853017
60. Dutta R, Gomes SN, Kalise D, Pacchiardi L. Using mobility data in the design of optimal lockdown strategies for the COVID-19 pandemic. *PLoS Comput Biol*. 2021;17(8):e1009236. <https://doi.org/10.1371/journal.pcbi.1009236> PMID: 34383756
61. Prasse B, Achterberg MA, Van Mieghem P. Accuracy of predicting epidemic outbreaks. *Phys Rev E*. 2022;105(1–1):014302. <https://doi.org/10.1103/PhysRevE.105.014302> PMID: 35193247
62. Drgoňa J, Arroyo J, Cupeiro Figueroa I, Blum D, Arendt K, Kim D, et al. All you need to know about model predictive control for buildings. *Annual Reviews in Control*. 2020;50:190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>
63. Astrom KJ, Bernhardsson BM. Comparison of Riemann and Lebesgue sampling for first order stochastic systems. In: Proceedings of the 41st IEEE Conference on Decision and Control. IEEE; 2002. p. 2011–6. <https://doi.org/10.1109/cdc.2002.1184824>
64. Rabi M, Moustakides GV, Baras JS. Adaptive sampling for linear state estimation. *SIAM J Control Optim*. 2012;50(2):672–702. <https://doi.org/10.1137/090757125>
65. Parag KV. On signalling and estimation limits for molecular birth-processes. *J Theor Biol*. 2019;480:262–73. <https://doi.org/10.1016/j.jtbi.2019.07.007> PMID: 31299332
66. Reich NG, Cummings DAT, Lauer SA, Zorn M, Robinson C, Nyquist A-C, et al. Triggering interventions for influenza: the ALERT algorithm. *Clin Infect Dis*. 2015;60(4):499–504. <https://doi.org/10.1093/cid/ciu749> PMID: 25414260
67. Parag KV, Donnelly CA, Zarebski AE. Quantifying the information in noisy epidemic curves. *Nat Comput Sci*. 2022;2(9):584–94. <https://doi.org/10.1038/s43588-022-00313-1> PMID: 38177483
68. Gostic KM, McGough L, Baskerville EB, Abbott S, Joshi K, Tedijanto C, et al. Correction: practical considerations for measuring the effective reproductive number, Rt. *PLoS Comput Biol*. 2021;17(12):e1009679. <https://doi.org/10.1371/journal.pcbi.1009679> PMID: 34879070
69. Gamado KM, Streftaris G, Zachary S. Modelling under-reporting in epidemics. *J Math Biol*. 2014;69(3):737–65. <https://doi.org/10.1007/s00285-013-0717-z> PMID: 23942791
70. Bastos LS, Economou T, Gomes MFC, Villela DAM, Coelho FC, Cruz OG, et al. A modelling approach for correcting reporting delays in disease surveillance data. *Stat Med*. 2019;38(22):4363–77. <https://doi.org/10.1002/sim.8303> PMID: 31292995
71. McGough SF, Johansson MA, Lipsitch M, Menzies NA. Nowcasting by Bayesian smoothing: a flexible, generalizable model for real-time epidemic tracking. *PLoS Comput Biol*. 2020;16(4):e1007735. <https://doi.org/10.1371/journal.pcbi.1007735> PMID: 32251464
72. Goldstein E, Dushoff J, Ma J, Plotkin JB, Earn DJD, Lipsitch M. Reconstructing influenza incidence by deconvolution of daily mortality time series. *Proc Natl Acad Sci U S A*. 2009;106(51):21825–9. <https://doi.org/10.1073/pnas.0902958106> PMID: 20080801
73. Parag KV, Donnelly CA. Using information theory to optimise epidemic models for real-time prediction and estimation. *PLoS Comput Biol*. 2020;16(7):e1007990. <https://doi.org/10.1371/journal.pcbi.1007990> PMID: 32609732
74. Park M, Cook AR, Lim JT, Sun Y, Dickens BL. A systematic review of COVID-19 epidemiology based on current evidence. *J Clin Med*. 2020;9(4):967. <https://doi.org/10.3390/jcm9040967> PMID: 32244365

75. Shim E, Choi W, Song Y. Clinical time delay distributions of COVID-19 in 2020-2022 in the Republic of Korea: inferences from a nationwide database analysis. *J Clin Med.* 2022;11(12):3269. <https://doi.org/10.3390/jcm11123269> PMID: 35743340
76. Tariq A, Lee Y, Roosa K, Blumberg S, Yan P, Ma S, et al. Real-time monitoring the transmission potential of COVID-19 in Singapore, March 2020. *BMC Med.* 2020;18(1):166. <https://doi.org/10.1186/s12916-020-01615-9> PMID: 32493466
77. Akhmetzhanov AR, Lee H, Jung S-M, Kayano T, Yuan B, Nishiura H. Analyzing and forecasting the Ebola incidence in North Kivu, the Democratic Republic of the Congo from 2018-19 in real time. *Epidemics.* 2019;27:123–31. <https://doi.org/10.1016/j.epidem.2019.05.002> PMID: 31080016
78. Pullano G, Di Domenico L, Sabbatini CE, Valdano E, Turbelin C, Debin M, et al. Underdetection of cases of COVID-19 in France threatens epidemic control. *Nature.* 2021;590(7844):134–9. <https://doi.org/10.1038/s41586-020-03095-6> PMID: 33348340
79. Carabelli AM, Peacock TP, Thorne LG, Harvey WT, Hughes J, COVID-19 Genomics UK Consortium, et al. SARS-CoV-2 variant biology: immune escape, transmission and fitness. *Nat Rev Microbiol.* 2023;21(3):162–77. <https://doi.org/10.1038/s41579-022-00841-7> PMID: 36653446
80. Unwin HJT, Mishra S, Bradley VC, Gandy A, Mellan TA, Coupland H, et al. State-level tracking of COVID-19 in the United States. *Nat Commun.* 2020;11(1):6189. <https://doi.org/10.1038/s41467-020-19652-6> PMID: 33273462
81. Contreras S, Biron-Lattes JP, Villavicencio HA, Medina-Ortiz D, Llanovarced-Kawles N, Olivera-Nappa Á. Statistically-based methodology for revealing real contagion trends and correcting delay-induced errors in the assessment of COVID-19 pandemic. *Chaos Solitons Fractals.* 2020;139:110087. <https://doi.org/10.1016/j.chaos.2020.110087> PMID: 32834623
82. Albani VVL, Albani RAS, Massad E, Zubelli JP. Nowcasting and forecasting COVID-19 waves: the recursive and stochastic nature of transmission. *R Soc Open Sci.* 2022;9(8):220489. <https://doi.org/10.1098/rsos.220489> PMID: 36016918
83. Miller AC, Hannah LA, Futoma J, Foti NJ, Fox EB, D'Amour A, et al. Statistical Deconvolution for Inference of Infection Time Series. *Epidemiology.* 2022;33(4):470–9. <https://doi.org/10.1097/EDE.0000000000001495> PMID: 35545230
84. Beesley LJ, Osthuis D, Del Valle SY. Addressing delayed case reporting in infectious disease forecast modeling. *PLoS Comput Biol.* 2022;18(6):e1010115. <https://doi.org/10.1371/journal.pcbi.1010115> PMID: 35658007
85. Cauchemez S, Bosetti P, Cowling BJ. Managing sources of error during pandemics. *Science.* 2023;379(6631):437–9. <https://doi.org/10.1126/science.add3173> PMID: 36730404
86. Lison A, Abbott S, Huisman J, Stadler T. Generative Bayesian modeling to nowcast the effective reproduction number from line list data with missing symptom onset dates. *PLoS Comput Biol.* 2024;20(4):e1012021. <https://doi.org/10.1371/journal.pcbi.1012021> PMID: 38626217
87. Charniga K, Park SW, Akhmetzhanov AR, Cori A, Dushoff J, Funk S, et al. Best practices for estimating and reporting epidemiological delay distributions of infectious diseases. *PLoS Comput Biol.* 2024;20(10):e1012520. <https://doi.org/10.1371/journal.pcbi.1012520> PMID: 39466727
88. Shearer FM, Lipsitch M. The importance of playing the long game when it comes to pandemic surveillance. *Proc Natl Acad Sci U S A.* 2025;122(15):e2500328122. <https://doi.org/10.1073/pnas.2500328122> PMID: 40203044
89. Keeling MJ, Dyson L, Tildesley MJ, Hill EM, Moore S. Comparison of the 2021 COVID-19 roadmap projections against public health data in England. *Nat Commun.* 2022;13(1):4924. <https://doi.org/10.1038/s41467-022-31991-0> PMID: 35995764
90. Wang Z, Andrews MA, Wu Z-X, Wang L, Bauch CT. Coupled disease-behavior dynamics on complex networks: A review. *Phys Life Rev.* 2015;15:1–29. <https://doi.org/10.1016/j.plrev.2015.07.006> PMID: 26211717
91. Phillips B, Bauch CT. Early warning indicators of epidemics on a coupled behaviour-disease model with vaccine hesitance and incomplete data. *JDG.* 2023;10(1):49–86. <https://doi.org/10.3934/jdg.2022024>
92. Franzen A, Wöhner F. Fatigue during the COVID-19 pandemic: evidence of social distancing adherence from a panel study of young adults in Switzerland. *PLoS One.* 2021;16(12):e0261276. <https://doi.org/10.1371/journal.pone.0261276> PMID: 34890414
93. Petherick A, Goldszmidt R, Andrade EB, Furst R, Hale T, Pott A, et al. A worldwide assessment of changes in adherence to COVID-19 protective behaviours and hypothesized pandemic fatigue. *Nat Hum Behav.* 2021;5(9):1145–60. <https://doi.org/10.1038/s41562-021-01181-x> PMID: 34345009