SCISPACE
formerly Typeset

Open access • Posted Content • DOI:10.1101/2020.06.03.20120691

Modeling COVID-19 dynamics in Illinois under non-pharmaceutical interventions

— Source link

George N. Wong, Zachary J. Weiner, Alexei V. Tkachenko, Ahmed Elbanna ...+2 more authors

Institutions: University of Illinois at Urbana–Champaign, Center for Functional Nanomaterials

Published on: 17 Jun 2020 - medRxiv (Cold Spring Harbor Laboratory Press)

Topics: Markov chain Monte Carlo

Related papers:

- Modeling COVID-19 dynamics in Illinois under non-pharmaceutical interventions
- Bayesian model assessment for stochastic epidemic models
- Modelling the early phase of the Belgian COVID-19 epidemic using a stochastic compartmental model and studying its implied future trajectories.
- Uncertainty quantification in epidemiological models for the COVID-19 pandemic.
- Efficient Data Augmentation for Fitting Stochastic Epidemic Models to Prevalence Data

Share this paper:

View more about this paper here: https://typeset.io/papers/modeling-covid-19-dynamics-in-illinois-under-non-2c3ojmhb60

Featured in Physics

Modeling COVID-19 Dynamics in Illinois under Nonpharmaceutical Interventions

George N. Wong,[1,†] Zachary J. Weiner,[1,†] Alexei V. Tkachenko,[2] Ahmed Elbanna,[3] Sergei Maslov,[1,4,5,*] and Nigel Goldenfeld[1,5,*]

[1]*Department of Physics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[2]*Center for Functional Nanomaterials, Brookhaven National Laboratory, Upton, New York 11973, USA*
[3]*Department of Civil Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[4]*Department of Bioengineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[5]*Carl R. Woese Institute for Genomic Biology, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*

(Received 15 June 2020; revised 13 August 2020; accepted 22 October 2020; published 16 November 2020)

We present modeling of the COVID-19 epidemic in Illinois, USA, capturing the implementation of a stay-at-home order and scenarios for its eventual release. We use a non-Markovian age-of-infection model that is capable of handling long and variable time delays without changing its model topology. Bayesian estimation of model parameters is carried out using Markov chain Monte Carlo methods. This framework allows us to treat all available input information, including both the previously published parameters of the epidemic and available local data, in a uniform manner. To accurately model deaths as well as demand on the healthcare system, we calibrate our predictions to total and in-hospital deaths as well as hospital and ICU bed occupancy by COVID-19 patients. We apply this model not only to the state as a whole but also its subregions in order to account for the wide disparities in population size and density. Without prior information on nonpharmaceutical interventions, the model independently reproduces a mitigation trend closely matching mobility data reported by Google and Unacast. Forward predictions of the model provide robust estimates of the peak position and severity and also enable forecasting the regional-dependent results of releasing stay-at-home orders. The resulting highly constrained narrative of the epidemic is able to provide estimates of its unseen progression and inform scenarios for sustainable monitoring and control of the epidemic.

DOI: 10.1103/PhysRevX.10.041033

Subject Areas: Biological Physics, Computational Physics, Interdisciplinary Physics

I. INTRODUCTION

On January 24, 2020, the second known COVID-19 case to be diagnosed in the USA was reported in Chicago, Illinois. Community transmission of the disease was confirmed on March 8, 2020. During the subsequent ten days, the epidemic grew with a case doubling time of approximately 2.3 days, while testing capacity was essentially fixed. On March 21, 2020, a stay-at-home (SAH) order was issued for the entire state of Illinois and subsequently extended on March 31, 2020 and again on April 23, 2020. The order was lifted on May 30, 2020 [1]. Responsible relaxation of the mitigation of COVID-19 must be informed by realistic and well-calibrated epidemiological modeling of the outcomes of any scenarios under consideration—not just of the resulting (increased) death toll but also of the stress placed upon the healthcare system. The purpose of this report is to present such an analysis.

A variety of modeling approaches are used by hospitals, public health officials, and state governments. These range between phenomenological models that use a curve-fitting procedure to match data, such as the daily death rate, and mechanistic methods that model the trajectory of the epidemic as individuals transition through several disease and healthcare-bound stages [2–5]. Only mechanistic models are able to make justifiable predictions while accounting for changes in the epidemic environment, such as the imposition or relaxation of community mitigation efforts. Of these, compartmental models like the susceptible-infectious-recovered (SIR) models and susceptible-exposed-infectious-recovered (SEIR) extensions are widely

*Correspondence to: maslov@illinois.edu, nigel@illinois.edu.
†These authors contributed equally to this work.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

used. Compartmental models describe how fractions of a homogeneous, well-mixed population progress through different states of the disease, driven by interactions between infectious and susceptible individuals. In the simplest models, the dynamics is deterministic and the rates are constant in time, but many variants and extensions exist.

In order to be practically useful, models must be calibrated to observed data [4,6–8]. We calibrate the important dynamics of the model to several simultaneous streams of empirical data, including total and in-hospital deaths, as well as hospital and ICU bed occupancy by COVID-19 patients. To avoid biases resulting from non-uniform and nonconstant testing rates, which may be difficult to parametrize, we do not consider positive-case data. The resulting model is a description of the epidemic as it progresses through the hospital system in Illinois; as it is clear that a non-negligible number of COVID-19 deaths occur outside the hospital environment (e.g., in homes and nursing homes especially), we augment our model with an effective description of the net incidence of deaths due to COVID-19.

There are many limitations to the types of models that we and others use to describe COVID-19, and these have been explored extensively in the literature, especially with regard to spatial structure [9], superspreader events and individuals [10–13], and the structure of contact networks [14–16]. A geographical region as large as the state of Illinois is not well described as homogeneous, due to large variations in population density and prevalence of infection between the Chicago metropolitan area and the more rural regions in Central and Southern Illinois. Indeed, most of the known cases and deaths to date have occurred in the Chicago area (by roughly by an order of magnitude), so in practice, this region dominates our results for the state as a whole.

One cannot simply scale the results of models for the Chicago metropolitan area to the entire state, however, because the transmission characteristics depend on the frequency and duration of contacts, which likely vary significantly among geographical regions. For this reason, we simulate the dense urban areas and the three sparser, more rural areas separately. We note that a more refined treatment would account for transfers between these separate populations, as well as transfers into and out of the state as a whole; however, we do not currently model these processes. The number of cases in individual rural areas is sufficiently small, due to the early stay-at-home order, that a long phase of exponential growth in these regions was largely avoided. As a result, these populations are not well described by continuum, deterministic models. Nevertheless, by aggregating these populations, the numbers are large enough that deterministic, exponential growth, at least at the early stages, was visible.

In this work, we describe an estimate of the rise and fall of the epidemic within Illinois, taking into account the modulation of the transmission parameters due to social distancing. In the following sections, we first describe our extension of the SEIR model, which takes into account the long and variable delay times reported in the literature. After describing the procedure by which we calibrate the model to data, we argue for the robustness of short-term model predictions. Finally, we present and discuss our predictions of the epidemic trajectory through the summer of 2020, including the effects of the release of the stay-at-home order at the end of May.

II. MODEL DESCRIPTION

To facilitate a general treatment of delay times in the COVID-19 epidemic, we implement a non-Markovian model that derives from the classic Kermack-McKendrick age-of-infection model [17]. Age-of-infection models are similar to compartmental models, such as SIR or SEIR models, but allow transition delays between entering sequential states to be drawn from arbitrary probability density functions (see, e.g., Refs. [4,18]). We use non-Markovian models because their delays can be defined by an arbitrary number of timescales, in contrast to the single exponential rate parameter of compartment-based models. Non-Markovian models are thus, in principle, better able to reproduce the observed signal delays between different states, e.g., the flattening of the curve of hospital admissions compared to in-hospital deaths.

The use of deterministic rather than stochastic epidemic descriptions is generally justified when the modeled populations are large enough that relative daily changes are small and the number of individuals is large relative to one. This case means that a deterministic compartmental model has a self-consistency check because, once the epidemic size is of order unity, a stochastic model allows the epidemic to die out; in contrast, in a deterministic model, an epidemic will continue to evolve even with a unrealistic, fractional number of infectious individuals. We will see that for some regions of Illinois, our estimates are at the limit of validity of deterministic models.

We adapt the conventional age-of-infection model to include a number of delayed healthcare system observables of the epidemic, such as the number of patients in hospitals and intensive care units (ICUs), and, ultimately, the number of deaths.

A. Time-modulated age-of-infection model

The core of our model is the daily incidence (i.e., the number of newly infected individuals) in each demographic (age) group i, $j_i(t)$. This value determines the dynamics of susceptible individuals in that group $S_i(t)$ according to

$$\frac{\mathrm{d}S_i(t)}{\mathrm{d}t} = -j_i(t). \tag{1}$$

The incidence itself follows the renewal equation,

$$j_i(t) = R_t \frac{S_i(t)}{S(t)} \sum_m \xi_{im} \int_0^\infty \mathrm{d}\tau K_{\mathrm{gen}}(\tau) j_m(t-\tau). \quad (2)$$

Here, R_t is the time-dependent effective reproduction number, $S(t) = \sum_i S_i(t)$ is the total susceptible population, $K_{\mathrm{gen}}(\tau)$ is the probability density function (PDF) of generation intervals, and ξ_{im} is the contact matrix. For simplicity, we assume $\xi_{im} = 1$; i.e., all demographic groups infect each other at the same rate. We assume that N_i, the total number of individuals in the demographic group i, is approximately constant for the duration of the epidemic. In practice, this means our model neglects the effects of births, deaths due to causes unrelated to COVID-19, and mobility of the population, which is appropriate for the short timescales we model. We further assume that individuals are only infected once, i.e., that the duration of immunity to COVID-19 is longer than the timescale over which we simulate the epidemic. Our simulations begin on a day t_s (whose value we sample during parameter inference) at which point we impose that ten people are spontaneously infected, setting $j_i(t)$ in proportion to the age distribution of the population under study.

We parametrize the effective reproduction number R_t in terms of the basic reproduction number R_0, a seasonal forcing factor $F(t)$, a mitigation factor $M(t)$, and the susceptible population fraction $S(t)/N$ according to

$$R_t = R_0 F(t) M(t) \frac{S(t)}{N}. \quad (3)$$

The formal dependence of R_t on the total susceptible population $S(t)$ is corrected in Eq. (2) by a factor $S_i(t)/S(t)$, which accounts for possible variation of susceptibility between different demographic groups. In our model, the homogeneous factor $M(t)$ accounts for both sources of mitigation (including the effects of self-imposed isolation as well as government-mandated SAH orders, school closures, etc.) and changes to disease infectivity. We parametrize $M(t)$ as a piecewise cubic Hermite interpolating polynomial that smoothly interpolates from 1 at t_0 to $M(t_1)$ at t_1 and is otherwise constant. The mitigation factor $M(t_1)$, start time, and end time of the interpolation interval are fitted by our algorithm. We choose to parametrize $M(t)$ as a single transition (i.e., a single event) since this minimal model reduces the risk of overfitting to spurious trends. This choice is supported by the following observations: The adoption of social-distancing practices took place relatively rapidly over a one- to two-week period, and the stay-at-home order remained active after the initial transition, presumably suppressing the magnitude of changes to mitigation. We will observe that this choice is a sufficient approximation for the duration of the data we use for calibration. We evaluate these claims more explicitly in Sec. IV D.

To model seasonal effects, we follow Ref. [19], which estimates seasonal forcing from the observed seasonal variability of positive tests in three other endemic coronaviruses. We thus adopt the functional form

$$F(t) = 1 + A_{\mathrm{SF}} \cos\left(\frac{2\pi(t - t_{\mathrm{peak}})}{365}\right), \quad (4)$$

where A_{SF} denotes the strength of the forcing and t_{peak} sets the day of the year when seasonal forcing is strongest. From Ref. [19], we infer that seasonal forcing is strongest in the winter and set $t_{\mathrm{peak}} = \mathrm{January16}$.

Reference [19] finds evidence for $A_{\mathrm{SF}} = 0.2$; however, we account for uncertainty in this parameter by sampling over A_{SF} during parameter inference. Incorporating this uncertainty is critical: If mitigation were only able to reduce the effective reproduction number to roughly unity, then seasonal forcing could drive a second wave of the epidemic.

For timescales of only a few months, our parametrization of R_t includes a degeneracy: A change in the parameter A_{SF} may be compensated by adjusting the mitigation profile $M(t)$. The degeneracy is broken as data are collected over long timescales since $M(t)$ models relatively instantaneous changes in infectiousness and $F(t)$ produces an explicitly year-long, periodic modulation. Practically, this case implies that A_{SF} may account for both seasonal effects and concurrent slow variations in the mitigation factor.

B. Model topology

Because of limited and biased testing, neither the susceptible population $S_i(t)$ nor the daily incidence of new infection cases $j_i(t)$ is directly observable. Hence, we are forced to infer the dynamics of the epidemic using lagging and indirect indicators. These indicators include the total number of hospitalized (but not critical) patients, $H(t)$; the number of critically ill patients currently in ICU beds, $C(t)$; and the cumulative number of daily deaths in the hospital, $D(t)$.

Our model topology assumes that all hospital deaths occur in ICU rooms. In practice, this simplification would be invalid if either a significant number of individuals die immediately upon entering the hospital (i.e., if the true delay distribution between hospitalization and death is appreciably bimodal) or if the number of hospital decedents exceeds the number of individuals who are admitted to the ICU. In this sense, our inferred parameter values, e.g., the probability of a patient in critical care dying, should not be interpreted as having real-world meaning since the values accommodate approximations in order to fit all input data.

Furthermore, by separating the observables from the incidence dynamics, our model supposes that the hospitalization status of an individual does not affect their likelihood of infecting someone. This choice is appropriate if the number of hospitalized individuals is a small fraction of the total number of infected individuals or if the delay between infection and hospitalization is longer than the generation interval.

The dynamics of the system may be described by daily flux variables:

(i) $\sigma_i(t)$, the number of infected individuals who become symptomatic,
(ii) $h_i(t)$, the number of daily admissions to all hospitals,
(iii) $r_i(t)$, the daily number of patients discharged from all hospitals,
(iv) $c_i(t)$, the daily number of patients transferred from the main floor of a hospital to its ICU,
(v) $v_i(t)$, the daily number of patients transferred from the ICU to the main floor of a hospital, and
(vi) $d_i(t)$, the daily number of deaths in ICU rooms.

We do not directly model deaths that happen outside of hospitals but instead infer the ratio of these deaths to the hospital deaths during our fitting procedure, as described below. Figure 1 schematically depicts the topology of our model along with the names of all flux and cumulative variables.

The dynamics of any flux variable $y(t)$ defined above may be obtained from the variable $x(t)$ directly preceding it in the chain of events, as shown in Fig. 1:

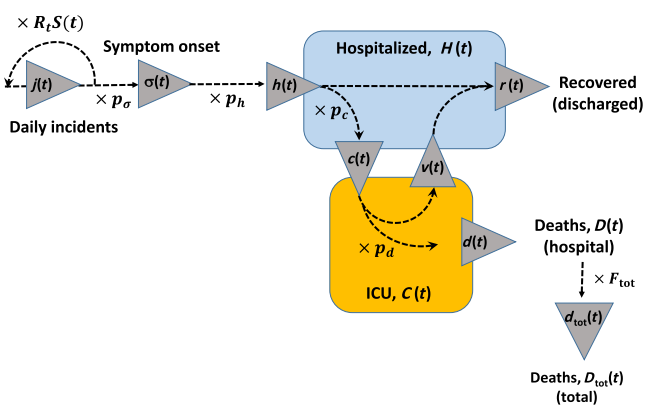

FIG. 1. Topology of our model, along with the names of all flux and state variables: the daily incidence, $j_i(t)$; the daily number of newly symptomatic individuals, $\sigma_i(t)$; the number of daily admissions to all hospitals, $h_i(t)$; the daily number of patients discharged from all hospitals, $r_i(t)$; the daily number of patients transferred from the main floor of a hospital to its ICU, $c_i(t)$; the daily number of patients transferred from the ICU to the main floor of a hospital, $v_i(t)$; the daily number of deaths in hospitals, $d_i(t)$; and the daily number of deaths in and out of hospitals, $d_{\mathrm{tot},i}(t)$. State variables are as follows: the total number of occupied hospital beds (main floor) $H_i(t)$ and the total number of occupied ICU beds $C_i(t)$. The other parameters of the model are the fractions of infected individuals who become symptomatic, $p_{\sigma,i}$; the fraction of symptomatic individuals who are hospitalized, $p_{h,i}$; the fraction of hospital patients who get to ICU, $p_{c,i}$; the fraction of ICU patients who will ultimately die, $p_{d,i}$; and the multiplier, F_{tot}, which converts between hospital deaths and all deaths in the state, including those outside of the hospital system. For the sake of legibility, we suppress age-group indices in the diagram.

$$y(t) = p_y \int_0^\infty \mathrm{d}\tau K_y(\tau) x(t-\tau). \tag{5}$$

Here, p_y is the proportion of individuals undergoing the transition $x \to y$ with time delays distributed according to a probability density function $K_y(t)$. Note that Eq. (2) has the same structure, except that both the input x and the output y variables are given by the daily incidence $j_i(t)$.

For the flux variables defined above, one obtains the following equations. The number of infected individuals who become symptomatic is

$$\sigma_i(t) = p_{\sigma,i} \int_0^\infty \mathrm{d}\tau K_\sigma(\tau) j_i(t-\tau), \tag{6}$$

where $p_{\sigma,i}$ is the (age-dependent) fraction of infected individuals who develop symptoms and $K_\sigma(\tau)$ is the PDF of the incubation period. The fraction $p_{h,i}$ of symptomatic individuals who are admitted to the hospital is

$$h_i(t) = p_{h,i} \int_0^\infty \mathrm{d}\tau K_h(\tau) \sigma_i(t-\tau). \tag{7}$$

The flux of hospital patients who become critically ill and are admitted to the ICU is given by

$$c_i(t) = p_{c,i} \int_0^\infty \mathrm{d}\tau K_c(\tau) h_i(t-\tau), \tag{8}$$

of which a fraction $p_{d,i}$ ultimately die according to

$$d_i(t) = p_{d,i} \int_0^\infty \mathrm{d}\tau K_d(\tau) c_i(t-\tau). \tag{9}$$

The fraction $1 - p_{d,i}$ of ICU patients who stabilize and return to the general ward of the hospital do so according to

$$v_i(t) = (1 - p_{d,i}) \int_0^\infty \mathrm{d}\tau K_v(\tau) c_i(t-\tau). \tag{10}$$

Both stabilized patients and hospital patients who recover without requiring critical care are discharged; thus, the influx of recovered individuals due to hospital discharges is given by

$$r_i(t) = \int_0^\infty \mathrm{d}\tau K_r(\tau) [(1 - p_{c,i}) h_i(t-\tau) + v_i(t-\tau)]. \tag{11}$$

Finally, to approximately account for patients who die outside of the hospital system, we introduce d_{tot}, a variable that tracks the total number of daily deaths both within and outside of the hospital. We connect total deaths to hospital deaths d according to a prefactor $F_{\mathrm{tot}} \geq 1$ and delay time τ_{tot} that reflects bureaucratic delays associated with issuing death certificates and publishing data on the Illinois Department of Public Health (IDPH) website. We observe

that these bureaucratic delays manifest themselves in strong day-of-the-week effects,

$$d_{\mathrm{tot},i}(t) = F_{\mathrm{tot}} \int_0^\infty \mathrm{d}\tau K_{\mathrm{tot}}(\tau) d_i(t-\tau). \tag{12}$$

In our simulations, we draw on clinical data to specify the age-dependent rates for hospitalization, ICU admission, and death. We report the details of this severity model in Table II. To account for differences between the literature values and what has occurred in Illinois, we introduce age-independent prefactors for the transition rates and fit them in our simulations. Although the relative severity values we use may not be accurate, in practice, this choice does not affect our model dynamics. Because the contact matrix ξ_{im} is constant across all interactions, and because the susceptibility is not a function of demographic group, the demographic-aggregated statistics are insensitive to the relative demographic ratios. Then, since we only fit data that has been summed over all age groups, we cannot observe any differences caused by inaccurate severity ratios.

The instantaneous occupation of hospital beds $H(t)$ and of ICU beds $C(t)$ may be obtained by integrating the incoming and outgoing fluxes as

$$H(t) = \sum_i \int_0^t \mathrm{d}t' [h_i(t') - c_i(t') + v_i(t') - r_i(t')], \tag{13}$$

$$C(t) = \sum_i \int_0^t \mathrm{d}t' [c_i(t') - v_i(t') - d_i(t')], \tag{14}$$

while the cumulative numbers of hospital and total deaths are

$$D(t) = \sum_i \int_0^t \mathrm{d}t' d_i(t'), \tag{15}$$

$$D_{\mathrm{tot}}(t) = \sum_i \int_0^t \mathrm{d}t' d_{\mathrm{tot},i}(t'). \tag{16}$$

III. PARAMETER INFERENCE

We calibrate our model to data by sampling over the high-dimensional model parameter space using a Markov chain Monte Carlo (MCMC) approach, as has been done by many others for epidemics, in general (see, e.g., Refs. [7,20]) and also for COVID-19 (see, e.g., Ref. [21] for applications in China, Ref. [22] for Mexico, and Ref. [23] for Italy, and also the related iterated filtering approach in Ref. [24]). While standard optimization techniques can also identify the best-fit model parameters, we use MCMC because it produces an estimate of the global posterior probability distribution. With the full distribution, we can motivate bounds on parameter uncertainty, explore correlations between parameters, and identify model idiosyncrasies. Access to the full distribution also provides a direct means to marginalize over some modeling uncertainties when forecasting the future trajectory of the epidemic. We generate a forecast by selecting an ensemble of sets of sample parameters from the posterior probability distribution and simulating the epidemic trajectory for each member of the ensemble.

A. Markov chain Monte Carlo methods

Given a set of model parameters $\vec{\Theta}$ and data $\vec{\mathcal{D}}$, the input to the MCMC sampler comprises a prior $\pi(\vec{\Theta})$ and a likelihood function $L(\vec{\mathcal{D}}|\vec{\Theta})$. The sampler computes the posterior probability $p(\vec{\Theta}|\vec{\mathcal{D}})$ for each point in parameter space according to Bayes' theorem,

$$p(\vec{\Theta}|\vec{\mathcal{D}}) = \frac{L(\vec{\mathcal{D}}|\vec{\Theta})\pi(\vec{\Theta})}{Z(\vec{\mathcal{D}})}. \tag{17}$$

The likelihood L represents the probability of observing the data $\mathcal{D}$ given a model with input parameters $\vec{\Theta}$, and the prior π encodes our expectation for the probability of a given set of model parameters $\vec{\Theta}$. Because we only compare points within the posterior distribution for an individual set of data, we neglect the constant model evidence by setting $Z(\vec{\mathcal{D}}) = 1$. In practice, the MCMC sampling recovers the log of the posterior probability distribution $H \equiv \ln p + \mathrm{const}$, which combines both types of inputs according to

$$H(\vec{\Theta}|\vec{\mathcal{D}}) = \ln[L(\vec{\mathcal{D}}|\vec{\Theta})] - \sum_\alpha \frac{(\Theta_\alpha - \Theta_\alpha^{(0)})^2}{2\Delta_\alpha^2}. \tag{18}$$

Here, the second term is the log of the prior over the model parameters; for each parameter, we either implement a Gaussian prior with mean expected value $\Theta_\alpha^{(0)}$ and variance Δ_α^2, or we use a flat prior, in which case the parameter does not appear in the sum.

This Bayesian framework enables a uniform treatment of all available input information, i.e., both observed time series data and the parameters of the model. We determine the prior means by averaging published clinical data weighted in proportion to the sample size of each study. To account for differences between reported estimates of parameters due to, e.g., possible variability of model parameters between different locations, the tolerance parameters Δ_α were estimated as *unweighted* root-mean-square deviations of the published data from their respective average values $\Theta_\alpha^{(0)}$ across the different studies. As a result, our procedure is flexible with respect to any local variability in model parameters. If high-quality local data

on parameter values are available for calibration, as might be the case for the duration of ICU stays or severity models, these data can be used directly with small values (or zero) for the respective tolerance parameters. By forcing parameters with known values to be more tightly constrained, unknown parameters will be automatically optimized with respect to all data types, and we can thus increase the fidelity of our model calibration result.

Many of our model outputs and data quantify daily incidences, e.g., the number of deaths per day. For these sorts of rate data, a Poisson likelihood estimator is appropriate. For a data point d at time t, the Poisson likelihood is given by

$$L(\lambda|t,d) = \left(\frac{e^{-\lambda(t)}\lambda(t)^d}{d!}\right)^{1/T}, \tag{19}$$

where the time-dependent rate $\lambda(t)$ is equal to the model output and T is the correlation time for the data. The likelihood over the full data set is the product over the likelihoods for each individual data point; thus, the total log-likelihood is given by

$$\ln\left[L(\vec{\mathcal{D}}|\vec{\Theta})\right] = \sum_i \frac{d_i \ln\lambda(t_i) - \lambda(t_i) - \ln(d_i!)}{T_i}. \tag{20}$$

In addition to being a well-motivated choice when comparing count data to rates, the Poisson likelihood is appealing because its effective uncertainty scales with the rate parameter. Thus, unlike with the L_2 norm, a single parameter specifies both the expectation value and the uncertainty of the measurement. In practice, we found that evaluating likelihoods using the L_2 norm did not significantly alter the qualitative features of our forecasts. In Appendix D, we provide a brief discussion of the Poisson (versus negative binomial) likelihood in the context of overdispersion.

We divide by T in Eq. (20) because we also calibrate against instantaneous hospital statistics, such as occupancy in the general ward and in the ICU. These data sources have natural correlation timescales: Occupancies correspond to smoothed averages since the majority of individuals who occupy a bed do so continuously over several days. We set T equal to the correlation timescales $\Theta^{(0)}$ from our priors. In particular, we assume a correlation of 6 days for occupied hospital beds and a correlation of 12.75 days for occupied ICU beds. We set $T = 1$ for the raw incidence data, i.e., for daily hospital deaths and for daily total deaths.

In Table I, we enumerate the model parameters we sample over and list the bounds on those parameters' values. We describe the shape of any prior distributions we impose. Parameters either describe the overall trajectory of the epidemic (i.e., the reproduction number and its time dependence) or the transition between individual states (i.e., the probability or time delay distribution associated with transitioning from one state to another). In our model, we use gamma distributions to describe delays and specify the mean and standard deviation for each distribution. Here, the mean τ and standard deviation σ of a gamma distribution are related to the standard shape and scale parameters by $k = \tau^2/\sigma^2$ and $\theta = \sigma^2/\tau$, respectively. We fix the generation interval mean and standard deviation to 4 and 3.25 days, respectively [25,26], while our incubation time distribution has a fixed mean of 5.5 days and a standard deviation of 2 days [27,28]. Parameters for all other delays are sampled. Notice that although our model accounts for asymptomatic individuals, we do not explicitly fit the fraction p_σ of infected individual who will become symptomatic since that quantity is degenerate with a combination of the infection-fatality rate (IFR) and the fraction of symptomatic individuals who eventually die [see Eq. (A1)].

Finally, we implement MCMC sampling using the PYTHON package EMCEE [34]. To improve sampling efficiency, we use ensemble move proposals based on the differential evolution [35], differential evolution snooker [36], and kernel density [37] proposal updates.

B. Posterior distributions and data fitting

We calibrate our model using data on hospital and ICU room occupancy by COVID-19 patients, the number of daily deaths of COVID-19 confirmed patients in hospitals, and the total number of daily deaths as publicly reported by the IDPH [38]. At the time of calibration, the hospitalization and ICU data were not publicly available, and they were provided to us by the IDPH. The MCMC sampling procedure produces a high-dimensional posterior probability. We use this posterior to identify the expectation values and uncertainties for each parameter with respect to the model and the data.

To summarize the posterior distribution in terms of epidemic trajectories, we take a representative sample of parameters $\vec{\Theta}_i$ according to their posterior probabilities. For each $\vec{\Theta}_i$, we simulate the full course of the epidemic, and we plot the resulting family of curves in aggregate. At every time point, we identify the median values of all measurable quantities (hospital and ICU rooms, and deaths) as well as quantiles corresponding to 68.4% and 95.6% confidence intervals. Because these quantiles are evaluated independently at each point in time, the visually recognizable curves do not correspond to actual epidemic trajectories.

IV. MODELING RESULTS

We use our age-of-infection model to describe the progression of COVID-19 in Illinois during 2020. We perform analyses for the state and for four distinct localities of the state. We also consider three separate scenarios of social distancing in our modeling.

TABLE I. Specification of sampler parameters and their priors. The mean and standard deviation of a delay-time PDF K_x are reported in days and denoted by τ_x and σ_x, respectively, and probabilities p_x denote the overall scaling of the age distribution $p_{x,i}$ as specified in Table II. In addition to the listed bounds on the various scaling factors p_x, we also enforce that the number of individuals drawn from a state does not exceed the total number presently in that state. We indicate strict bounds on parameter values and provide the mean $\Theta^{(0)}$ and standard deviation Δ for parameters' Gaussian priors where specified. Our priors on K_h, K_r, K_c, and K_d are informed by Refs. [28–32], and we use Ref. [33] to set our prior on lnIFR. The * indicates that we implement a Gaussian prior for lnIFR, with mean corresponding to IFR = 0.7% and variance of lnIFR set to 0.175.

Parameter	Description	Bounds	$\Theta^{(0)}$	Δ
R_0	Initial reproduction number	[1, 5]		
t_s	Date of first infection (days after 1/1/20)	[35, 65]		
t_0	Date mitigation begins to take effect (days after 1/1/20)	[60, 85]		
t_1	Date mitigation is first in full force (days after 1/1/20)	[70, 100]		
$M(t_1)$	Factor by which full mitigation decreases R from R_0	[0.05, 1]		
IFR	Infection fatality ratio (set by priors)	[0.25%,1.8%]	0.7%	0.175*
τ_h	Time delay distribution parameters for transition from	[0.5, 40]	6.5	2
σ_h	symptom onset to hospitalization	[0.5, 20]	4	2
τ_{disch}	Time delay distribution parameters for transition from	[0.5, 40]	6	2
σ_{disch}	hospitalization to hospital discharge	[0.5, 20]	4	2
p_c	Probability of ICU admission given hospitalization	[0.05, 4]		
τ_c	Time delay distribution parameters for transition from	[0.5, 10]	2	2
σ_c	hospitalization to ICU admission	[0.5, 10]	2	2
p_d	Probability of death given ICU admission	[0.05, 3]		
τ_d	Time delay distribution parameters for transition from	[4, 20]	12	3
σ_d	ICU admission to death	[1, 20]	8.5	3
τ_{rec}	Time delay distribution parameters for transition from	[4, 20]	12.75	3
σ_{rec}	ICU admission to recovery (and return to main floor)	[1, 20]	10	3
F_{tot}	Ratio of all deaths to hospital deaths	[1, 10]		
τ_{tot}	Time delay distribution parameters for delay between	[0, 10]	2	1
σ_{tot}	hospital deaths and all-death statistics	[0, 10]	2	1
A_{SF}	Strength of seasonal forcing term	[0, 0.2]		

A. Simulations for Illinois

In Fig. 2, we show the fits and predictions of our model for the entire state of Illinois, assuming that once implemented, the state-imposed and self-regulated social-distancing behavior of the population maintains the same level until the end of the simulation. We report the median and the 68.4% and 95.6% confidence intervals of several dynamical outputs of our model, obtained from an ensemble of forward simulations that sample the posterior distribution over model parameters. Figure 2 presents three separate calibrations: (a) using data up through April 20, 2020 and assuming a fixed seasonal forcing amplitude $A_{\text{SF}} = 0.2$ that reflected contemporary estimates, (b) again using data through April 20, 2020 but instead sampling over A_{SF} to allow for uncertainties in that value, and (c) using data through May 17, 2020 and sampling over A_{SF}.

Contrasting panels (a) and (b) demonstrates that fixing the amplitude of seasonal forcing to 20% results in a poor projection since the forecasts with the fixed amplitude disagree with the data by mid-May. Indeed, the uncertainty introduced by allowing A_{SF} to vary produces a more robust fit to the long plateau exhibited in all data sources. Next, comparing panels (b) and (c) shows that while short-term predictions are broadly consistent, the spread of model forecasts become narrower. In particular, including data up to May 17, as in panel (c), enables the model to identify that the plateau is beginning to bend, in contrast to the model shown in panel (b), in which a continued plateau is not precluded. We investigate the predictive power of our model and calibration procedure in the next subsection.

Hospital and ICU occupancy as well as deaths related to COVID-19 exhibit a long plateau spanning from the beginning of April through at least mid-May. This behavior is not just due to the fact that mitigation reduced R_t to almost exactly one but also because of the variance in when the lagging indicators respond to the rate of infection. Namely, while the various nonpharmaceutical interventions (NPIs) implemented in Illinois occurred on relatively short timescales (as can be seen in the sudden change of slope for the daily incidence of new infections in Fig. 2), the variance in the delays between when individuals become, e.g.,

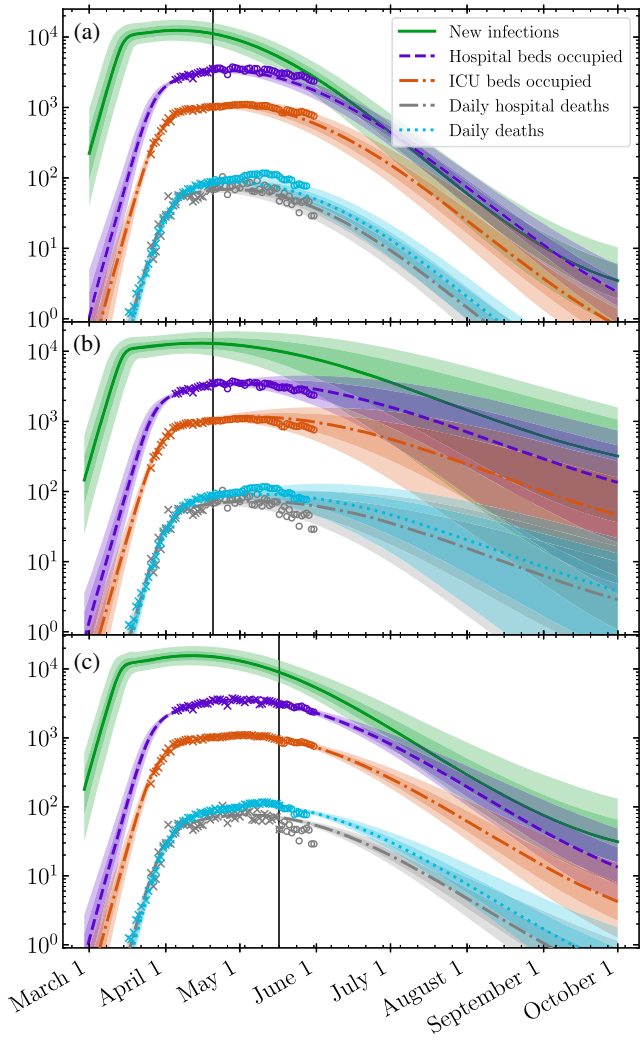

FIG. 2. Fits and predictions of our model of the entire state of Illinois under the continuation of the stay-at-home order and social-distancing measures, resulting from parameter inference, which (a) fixes $A_{\mathrm{SF}} = 0.2$ and calibrates to data through April 20, 2020, (b) samples A_{SF} and calibrates to data through April 20, 2020, and (c) samples A_{SF} and calibrates to data through May 17, 2020. In each case, the data used for calibration are denoted by crosses, while data from later dates (open circles) allow a comparison of model predictions with subsequent real-world observations. Each panel's vertical line marks the end of data used for calibration. Solid curves denote median model predictions at a given time, while the shaded regions denote the 68.4% (darker shading) and 95.6% (lighter shading) confidence intervals of a particular output. We depict the daily incidence of new infections (green, solid line), the number of occupied hospital beds by confirmed or tentative COVID-19 patients (purple, dashed line), the number of ICU beds occupied by confirmed COVID-19 patients (orange, dot-dashed line), the number of hospital deaths by confirmed COVID-19 patients (grey, dot-dashed line), and total daily deaths of COVID-19 patients (light blue, dotted line). We remove reporting artifacts by plotting daily-death data smoothed by a 7-day running average.

symptomatic and then hospitalized, introduces spread in the indicators' responses to changes in R_t. As an example, some portion of individuals infected before any mitigation takes place will continue to be admitted to the hospital well after mitigation occurs. Indeed, this variability compounds with subsequent transitions, such that daily deaths, being the final indicator, exhibit the most gradual change in incidence rate. Our model's generality to arbitrary delay distributions makes it particularly well suited to accurately capture this effect.

B. Robustness of predictions

In order to explore the predictive capabilities of the model, we present a series of benchmark simulations to compare the predictions of models calibrated with increasingly recent data. In Fig. 3, we show the fits to hospital beds occupied, ICU beds occupied, daily hospital deaths, and daily total deaths in the entire state of Illinois, calibrating with data up to April 1, 2020, April 8, 2020, April 20, 2020, and May 3, 2020. During the time period studied here, the Illinois stay-at-home order was issued, leading to an end to the exponentially growing phase of the epidemic and a flattening of the curve. Because of transition delays, the quantities to which we calibrate exhibit exponential growth as late as early April; for example, the number of daily deaths does not flatten until around April 10, 2020. Thus, the above-specified dates of calibration provide a test that measures the ability of our model to anticipate the bending of the curve.

As expected, simulations with earlier terminal calibration dates generate forecasts with larger, less restrictive confidence intervals: With data from only the earliest stages of the epidemic, neither continued exponential growth nor successful suppression of infections due to mitigation can be ruled out. Between April 1 and April 8, however, ICU occupation stopped growing exponentially, thereby providing the first indicator by which the model can infer the effect of mitigation. Indeed, even the April 8 forecast anticipates the subsequent flattening in the daily-deaths curve. Later forecasts (e.g., that of April 20) remain consistent with the April 8 model while also favoring a continuation of the plateau over a more rapid decline in use of hospital resources and deaths. Finally, the May 3 model is largely consistent with the April 20 one, but it begins to project a slight decline that agrees with the new data. This latest forecast should not be expected to lie strictly within the confidence interval of the previous one, as any future changes in mitigation cannot be anticipated.

We also point out that the April 1 forecast for hospital occupation spans several orders of magnitude on any given day [as seen in panel (a) of Fig. 3]. This result is due in part to the lack of hospitalization data that could provide direct constraints, but it is also an inevitable feature of a system that exhibits exponential growth dynamics. Furthermore,

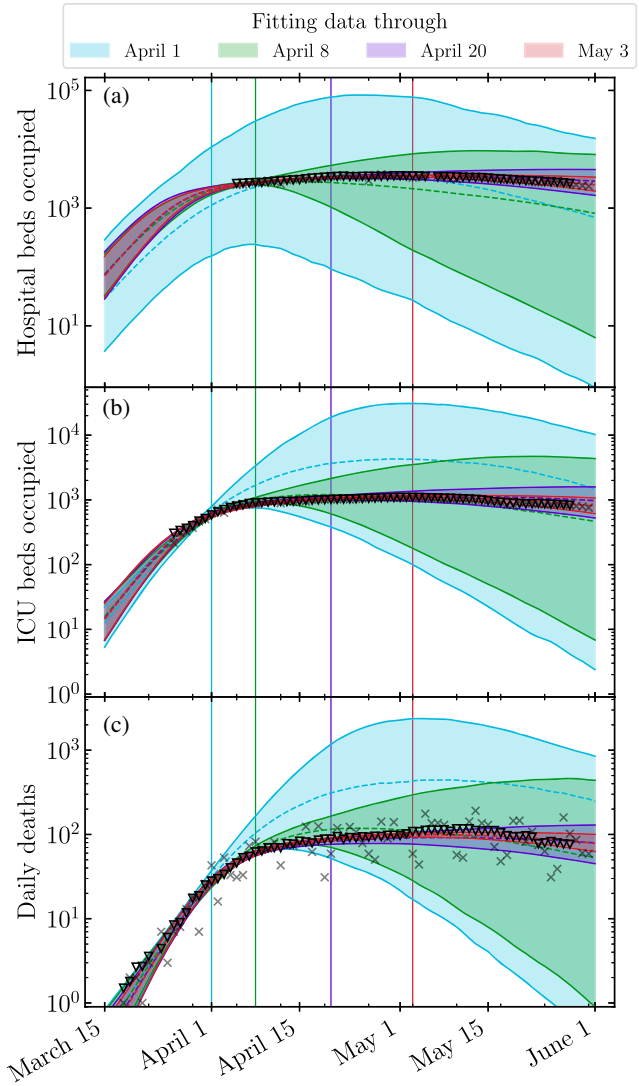

FIG. 3. Robustness of simulations for the COVID-19 epidemic in Illinois, evaluated by comparing 95.6% confidence regions of predictions for parameter inference calibrated to data through April 1 (light blue), April 8 (green), April 20 (purple), and May 3, 2020 (red). The panels depict the number of hospital beds occupied (a), the number of ICU beds occupied (b), and the total daily deaths (c). In each panel, triangles depict the data smoothed by a 7-day rolling average, and grey crosses are the actual data.

we hypothesize that the lack of hospitalization data before the beginning of the plateau is at least partially responsible for the April 1 forecast for ICU occupation and daily deaths being relatively unconstrained. Since hospitalizations serve as the earliest available indicator for the progression of the epidemic, we hypothesize that if the model were supplied with these data before April 1, the model would have been able to discern that the data were no longer consistent with an exponentially growing epidemic. This observation underscores the importance of rapid and reliable reporting of hospitalization data—or even better, robust and representative testing for positive cases—in the context of modeling epidemic dynamics.

The disappointingly short horizon of predictability for epidemic models when $R_t > 1$ shown here represents a fundamental limitation of forecasting, in much the same way that chaotic dynamics limits weather prediction, and this issue has been noticed in other independent studies [39–41]. However, the exponential sensitivity has a silver lining: Small changes to transmission can have large impacts on the overall trajectory and fatality of the disease.

In summary, the model curves fitted after April 8, 2020, i.e., with data from the plateau, follow the trends of the data well. We conclude that the model can be characterized as semiquantitative and that it is capable of capturing broad epidemic dynamics and fitting empirical data. In this sense, it can serve as a useful tool to make short-term predictions that may be useful for planning purposes.

C. Regional modeling

To account for differences in the epidemic in different parts of the state, we simulate the epidemic trajectory in the four distinct Restore Illinois regions separately. Each of these "super-regions" is composed of multiple emergency medical services (EMS) regions defined by the IDPH [42]. The Northeast super-region includes the city of Chicago and its suburbs, EMS regions 7 through 11. The North-Central super-region contains EMS regions 1 and 2, the Central super-region EMS includes regions 3 and 6, and finally, the Southern super-region comprises EMS regions 4 and 5. In the absence of detailed transportation data, we assume no population transfer between these super-regions, and so unlike, e.g., the model for the state of Georgia in Ref. [43] and Italy [23], our regional modeling is not a genuine metapopulation model for the state.

In Fig. 4, we show the fits and predictions of our model calibrated to data up to May 17, 2020 for each of these four regions. We report inferred model parameters in Table III. Our median estimates of the basic reproduction number R_0 at the start of the epidemic are consistently above 1, ranging from 2.4 ± 0.16 for the Northeastern region, including Chicago and its suburbs, to 1.7 ± 0.16 for the Central region, including the University of Illinois at Urbana-Champaign.

The per capita daily deaths and illnesses are at least tenfold higher in Chicago and its suburbs compared to the downstate areas of Illinois, which is likely due to increased contact density in the upstate region, reflected by a higher initial $R_0 \approx 2.3$ compared to approximately 1.8 in the downstate regions, coupled with the fact that mitigation began earlier relative to the start of the epidemic in some downstate regions. In regions where the virus entered the population later, the epidemic had a shorter phase of unmitigated exponential growth.

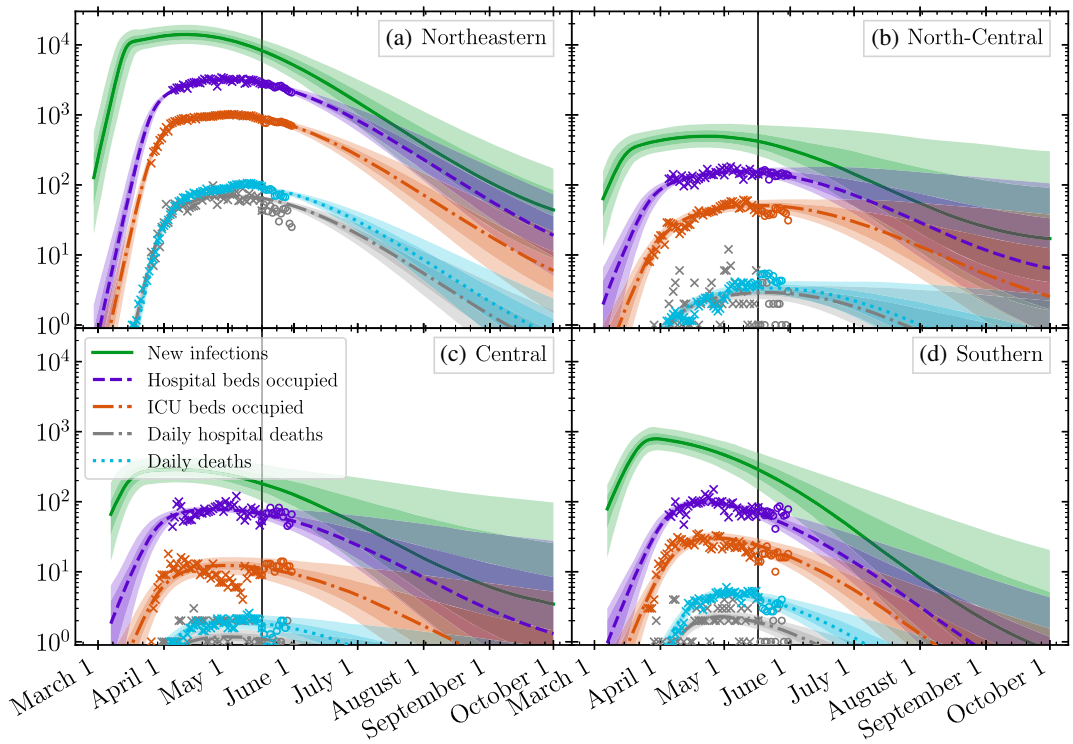

FIG. 4. Fits and predictions of our model under the baseline scenario with the stay-at-home order and social distancing maintained until October 1, 2020. The model is calibrated with data (crosses) through April 20, 2020 separately for each of the four super-regions of the state: (a) North-Central, (b) Northeast, (c) Central, and (d) Southern. The data from later dates (open circles) were not used in our fits, allowing a comparison of model predictions with real-world observations. Solid lines denote the median model prediction at a given time. The shaded regions denote the 68.4% (darker shading) and 95.6% (lighter shading) confidence intervals, obtained as quantiles of an ensemble of forward simulations, which sample the posterior distribution over model parameters. We depict the daily incidence of new infections (green, solid), the number of occupied hospital beds by confirmed or tentative COVID-19 patients (purple, dashed), the number of ICU beds occupied by confirmed COVID-19 patients (orange, dot-dashed), and total daily deaths of COVID-19 patients (light blue, dotted). We remove reporting artifacts by plotting daily-death data smoothed by a 7-day running average.

The values of R_t corresponding to the postmitigation basic reproduction number are very close to unity, reflecting the approximately constant number of hospital and ICU beds occupied by COVID-19 patients and daily deaths from COVID-19 in different super-regions in Illinois during much of April and May 2020.

D. Evaluation of parameter fits

In Fig. 5, we show a subset of the joint posterior probability distributions for model parameters relevant to the parametrization of the mitigation factor $M(t)$ as specified above, fitting to the data shown in Fig. 2. The correlations between some pairs of fitted parameters, e.g., between R_0 and the start date of the epidemic t_s, are reflected in the ellipsoidal shape of the posteriors, which is sensible: The later the epidemic begins, the larger the basic reproduction number must be in order to fit the data.

In Table III, we report the parameter values our model infers when fitting to data for different regions of the state and over different time ranges. We also report the effective reproduction numbers R_t as evaluated on May 1, 2020. On May 1, 2020, R_t appears to have barely dropped below unity, suggesting that mitigation efforts may have only marginally halted the exponential growth of the epidemic at that time.

E. Comparison with mobility data

While our model does not provide a microscopic description of social interactions and movement in the population, we may evaluate our fitted $M(t)$ relative to social mobility indices derived from cell phone data [44,45]. In the top panel of Fig. 6, we plot the time dependence of several mobility indices reported by Google [44] for the entire state of Illinois, measuring the change in visits to destination points categorized as retail and recreation, grocery and pharmacy, parks, transit stations, and workplaces. The Unacast mobility data [45] are shown in the bottom panel of Fig. 6, depicting an effective distancing metric and a measure of trips to so-called nonessential destinations.

Remarkably, although the model is supplied with no prior information on nonpharmaceutical interventions, the

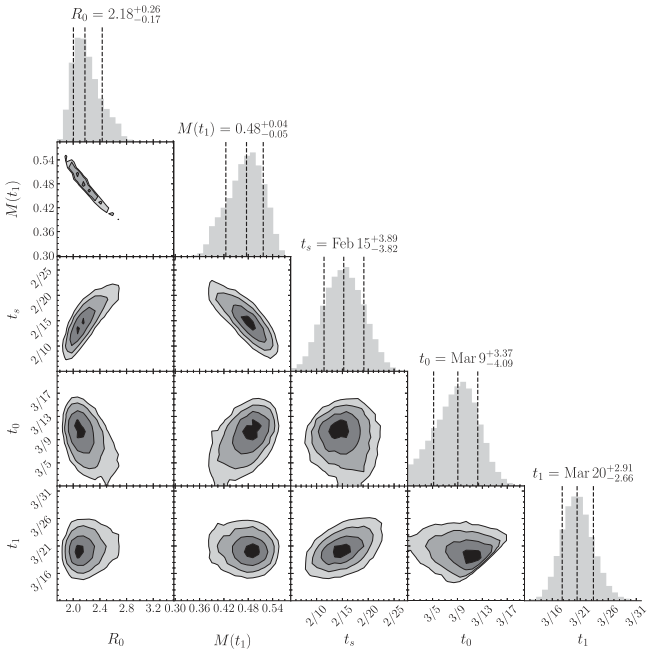

FIG. 5. Joint posterior distributions of pairs of the main parameters of our model fitted to the all-state data in Illinois up to May 17, 2020. The variables shown are the initial value of the reproduction number, R_0; the mitigation factor $M(t_1)$; the start date of the epidemic, t_s; the day the mitigation factor begins to deviate from 1, t_0; and the day the mitigation reaches its asymptotic value, t_1.

inferred dates and magnitude of mitigation agree with the start and end dates of the largest drop in both sets of mobility data. Comparing to the Google mobility data, our $M(t)$ curve exhibits an amplitude, start date, and end date consistent with indices corresponding to retail and recreation, transit stations, and workplaces. On the other hand, parks and grocery and pharmacy show a more modest reduction, which still matches the time frame of $M(t)$. In the Unacast data, both metrics also appear to match the time-dependent change in $M(t)$. Note that both data sets evince increased movement at later dates. Because here we parametrize mitigation as a single transition, our fits to $M(t)$ would not reproduce this recent increasing trend.

Our procedure is distinct compared to several previous efforts to incorporate mobility into models of COVID-19 dynamics, which either impose that changes in transmission coincide with known dates of nonpharmaceutical interventions or scale R_t according to reductions in mobility [46,47]. In contrast, we assume no causal relationship between R_t and mobility and simply perform a *post hoc* comparison between our model output and independent mobility data. The fact that our model recovers a trend in mitigation that is consistent with mobility measures speaks to its flexibility and calibration procedure.

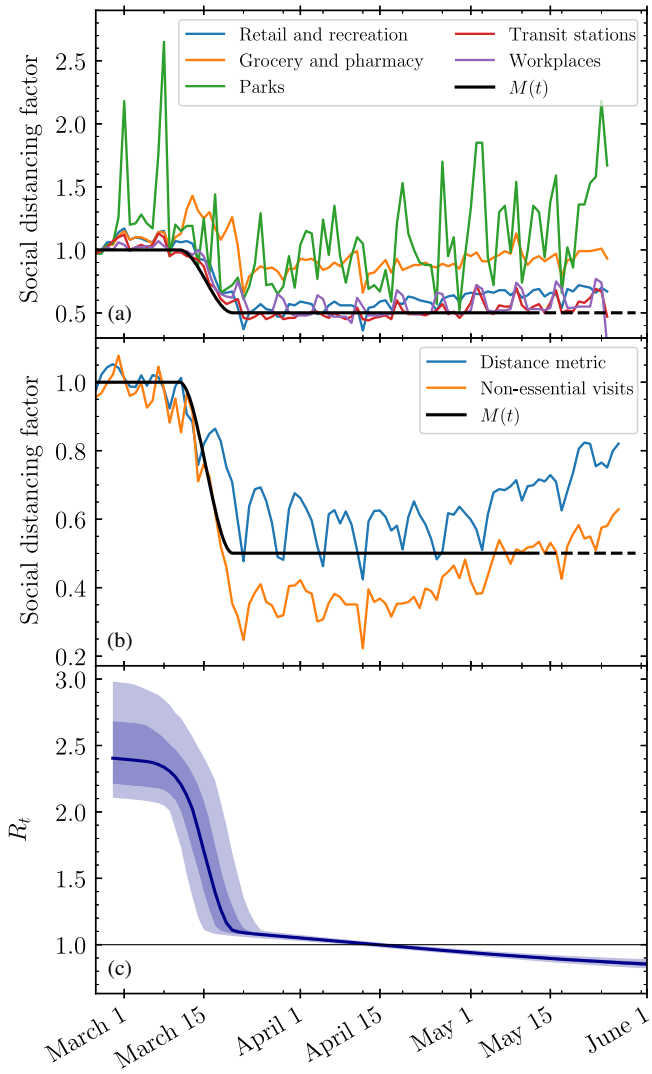

FIG. 6. The time-dependent mitigation factor $M(t)$ defined in Sec. II A as inferred from model fits compared to measures of mobility in Illinois provided by Google (a) and Unacast (b). In panel (c), we also plot the inferred evolution of R_t as defined in Eq. (3), which we draw from a calibration to data from the entire state of Illinois through May 17, 2020. Shaded regions denote the 68.4% (lighter shading) and 95.6% (darker shading) confidence intervals, while the solid curve denotes the median.

V. WHAT-IF SCENARIOS

We now consider possible future scenarios and alternative historical scenarios in which nonpharmaceutical interventions are lifted or were never implemented at all. The former enables a model-based assessment of the risk of, e.g., lifting stay-at-home orders on a certain date, while the latter demonstrates the impact that previously implemented measures have already had on outcomes. Our analysis focuses on two key measures for guiding and justifying policy decisions: the stress imposed on the healthcare system and the death toll.

A. No stay-at-home order

We first investigate the trajectory of the COVID-19 epidemic in Illinois in the absence of any form of social-distancing or mitigation measures, whether self-imposed or mandated by the government. We conclude that rate of hospital-bound deaths, ICU bed occupancy, and hospital bed occupancy would be higher than what actually took place by an order of magnitude or more, as shown in Fig. 7. The stay-at-home order and self-imposed social-distancing measures were clearly crucial to flattening the curve.

In Ref. [48], we made an early estimate of the ICU utilization in the city of Chicago under two scenarios: one in which the stay-at-home order was issued on March 20, 2020, and another in which the order was delayed by 20 days. Under the first scenario, the ICU utilization by COVID-19 patients never exceeded the number of ICU beds not occupied by other patients, while under the second scenario, it exceeded ICU capacity by nearly a factor of 10. This example highlights the cost of mistiming in NPIs [49]. In spite of inevitable uncertainties associated with these early estimates, that work correctly identified the timing of the peak in ICU demand to happen on or around April 22, 2020. However, the magnitude of this peak was underestimated in this study because the scenario assumed that the postmitigation value of $R_t = 0.9$ would be achieved by social distancing. In reality, the response of the population to the stay-at-home order in Chicago was somewhat weaker, resulting in a larger value of R_t and about a threefold higher peak ICU bed occupancy than we had predicted.

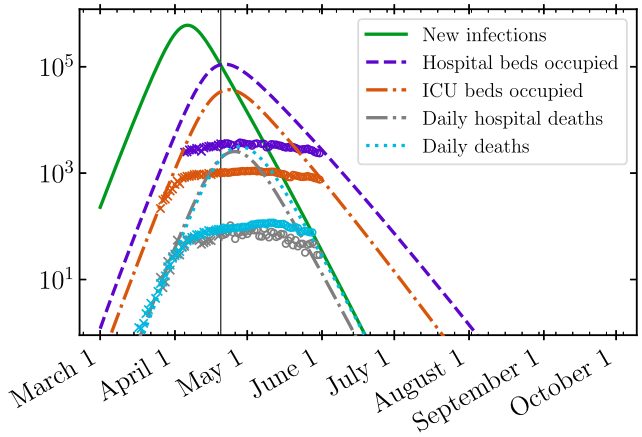

FIG. 7. Counterfactual simulation of the COVID-19 epidemic in the absence of government or population self-imposed social-distancing measures. The model is calibrated to data through April 20, 2020 (crosses); we then artificially set $M(t_1) = 1$, i.e., assume no mitigation ever takes place. The data from later dates (open circles) are not used in our fits, and they demonstrate the effect that real-world NPI strategies had on the epidemic. We depict the daily incidence of new infections (green, solid), the number of occupied hospital beds by confirmed or tentative COVID-19 patients (purple, dashed), the number of ICU beds occupied by confirmed COVID-19 patients (orange, dot-dashed), and total daily deaths of COVID-19 patients (light blue, dotted). We remove reporting artifacts by plotting daily-death data smoothed by a 7-day running average.

B. Partial removal of stay-at-home order

The state of Illinois lifted its original stay-at-home order on May 30, 2020 [1]. In Fig. 8, we consider two scenarios for lifting stay-at-home orders for the entire state of Illinois. In the first scenario, mitigation effects are completely removed, and $M(t) = 1$ for times t after June 1, 2020. We also consider the more conservative case that mitigation recovers by 30% to $M(t) = M(t_1) + 0.3(1 - M(t_1))$ for t after June 1, 2020. This second scenario assumes that a combination of self-regulation and remaining government-imposed mitigation measures—such as requiring masks, banning large gatherings, etc.,—results in only a partial reduction of the effective mitigation factor.

The first scenario exhibits a substantial second wave, with rapid and strong peaks in all quantities occurring successively through the month of July. In the second,

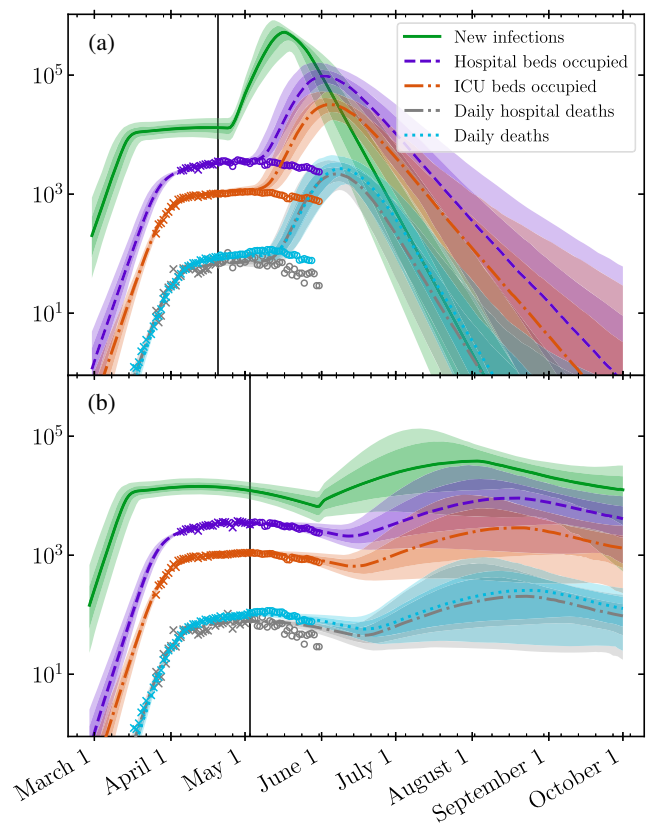

FIG. 8. Simulation of different scenarios for relaxation of the stay-at-home order for the entire state of Illinois. Each panel's vertical line marks the end of data used for calibration. We display two scenarios: (a) where the mitigation actor $M(t)$ returns to 1 on April 24, 2020, corresponding to a complete lack of social distancing, and (b) where on June 1, 2020, the effect of NPIs is reduced by 30%.

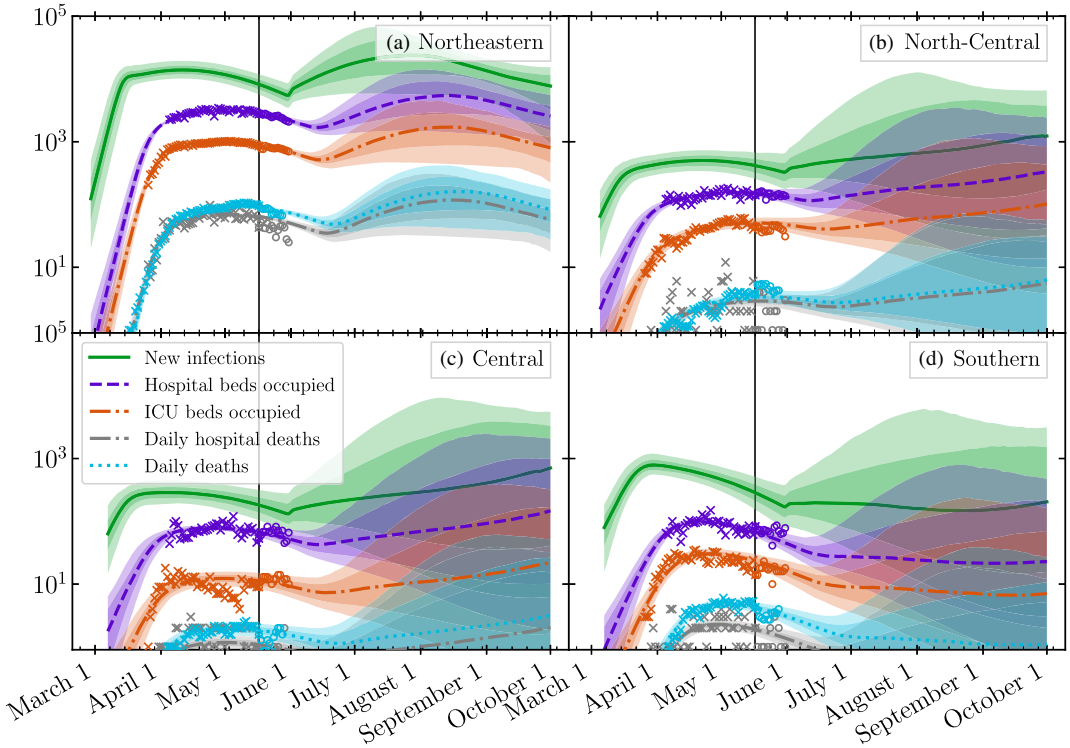

FIG. 9. Fits and predictions of our model under the a partial release of the stay-at-home order on June 1, 2020 to a mitigation factor $M(t) = M(t_1) + 0.3 \times (1 - M(t_1))$. The panels and plots denote the same regions and data sources as those in Fig. 4.

weaker (and perhaps more realistic) scenario, a second wave still occurs but does so later and with a reduced peak height. In this case, the peak demand for hospital and ICU beds and number of deaths are reduced.

An aggregated model of the entire state does not describe its heterogeneous population structure, which may be particularly important in forecasting beyond the lifting of the stay-at-home order. We now separately report the results of modeling the expiration of the stay-at-home order (with a 30% reduction in mitigation) in each of the four aforementioned super-regions, with the caveat that we are unable to take into account possible transfers of people between the regions. As in Fig. 9, we first calibrate models using data through May 17, 2020, imposing again that mitigation recovers by 30% to $M(t) = M(t_1) + 0.3(1 - M(t_1))$ for t after June 1, 2020.

In each super-region, a second wave is clearly visible, as all four populations currently still have a significant number of infectious people at large. Although the Northeast region, which includes Cook County and metropolitan Chicago, has the biggest second wave in absolute numbers, its relative impact on the epidemic trajectory is smaller merely because the epidemic has not declined to the same extent as the other three regions. Under the worst-case scenario, the more sparsely populated Southern, Central, and North-Central regions have relatively larger second waves because the epidemic as of yet is relatively less active in those regions compared to the Northeastern region, which includes Chicago and its suburbs.

We note that the magnitude of the second wave may be containable if rapid and efficient testing, contact tracing, and isolation mitigation strategies are employed [50,51]. Following standard protocols [52], we estimate that even if purely manual contact tracing is employed for all three stages of tracing (case identification, tracing, and follow-up), the number of contact tracers required is approximately 8.3 times the daily number of new cases identified. A considerable reduction in workforce and an increase in efficiency can be obtained by electronic measures. Nevertheless, the potential magnitude of the second wave in this scenario suggests that contact tracing will be challenging and require extraordinary resources to execute.

VI. DISCUSSION

Modeling plays an important role in the societal response to the COVID-19 pandemic, and a variety of approaches are used to inform public health policy, guide resource allocation, and plan nonpharmaceutical interventions [53]. The present study of the spread of COVID-19 in Illinois reveals both the strengths and limitations of modeling, and it provides potentially actionable insights into the future spread of the disease. We begin with some technical points and best practices that we have developed during our work.

Model calibration.—Our analysis highlights the importance of choosing appropriate data with which to calibrate models, and to perform calibration with precision. Because of the large number of parameters that inevitably enter epidemiological models, calibration requires parameter inference in a high-dimensional space with strong potential for improper fits resulting from failure to reach global optima. Although the MCMC methods we use are computationally intensive, they are relatively efficient in exploring high-dimensional, multimodal distributions, and they converge to well-behaved global posterior probability distributions. Bayesian inference enables the incorporation of previous studies (e.g., meta-analyses) to provide reasonable priors on parameters that are poorly constrained by the available data. As an example, although the data we calibrate to do not constrain the prevalence of the infection, we account for this uncertainty by informing our prior on the IFR from serological studies [33]. The IFR is an important variable in terms of disease outcomes, so model predictions must systematically account for the uncertainty in this variable; however, since we do not consider any data that can constrain IFR, the ranges we recover (both in terms of estimated IFR parameters and the number of new daily infections) are set entirely by the choice of our prior. In Appendix C, we provide a brief discussion of the consequences of deviations from the IFR prior by giving a model comparison that considers synthetic data.

In our analysis, we have taken data at their face values. A more thorough analysis could account for differences in the trustworthiness of data, e.g., programmatically dealing with uncertainties associated with classification of early hospital admissions as either COVID-positive or under investigation. Such an approach might also be able to calibrate against case data (i.e., the number of individuals who test positive for the virus), which is subject to strong ascertainment bias. In Appendix E, we show that ascertainment bias appears to decrease during the first months of the epidemic by comparing the model prediction for new infections to positive-case data (see Fig. 15).

Error estimation.—The MCMC procedure generates a complex, high-dimensional probability distribution that may be sampled to estimate future epidemic trajectories. While one might be tempted to simply generate the maximum likelihood trajectory for forward (in time) simulation, this can be misleading. The probability distribution of trajectories may have a maximum likelihood trajectory that is favored only slightly above other trajectories yet is noticeably different from the mean or median estimator. We thus sample trajectories forward in time, and at each time point, we evaluate the median and confidence intervals of outputs. Although the curves we depict do not correspond to a trajectory that would be realized in practice, together with confidence intervals, they provide a meaningful description of the range of forecast results.

Robustness.—In order to ascertain the predictive power of our model, we assessed to what extent it was able to make predictions beyond the range of data to which it had been calibrated. We first considered a rather severe test: How far in advance can the model predict the abrupt flattening of the epidemic curves that occurred after April 10, 2020? Unsurprisingly, the answer here is not particularly encouraging because the curves dramatically switched from exponential growth to a phase of much slower variation or plateau. We were only able to fit daily death and ICU occupancy data before the flattening since hospitalization data were not available until mid-April. As a result, none of the information feeding into the model dynamics was able to foreshadow the onset of the plateau.

Nevertheless, such self-consistency checks are important validations of the modeling process and should be attempted if there are early enough data available. Encouragingly, the range of our model estimates decreased in a hierarchical and consistent manner, with projections from earlier points in time bounding later projections.

Of course, it is not possible for our model to predict future changes to the strength of mitigation. If the behavioral changes in the host population occur slowly compared to the disease dynamics timescales, then the utility horizon of the model would be farther away than it would be otherwise. An example of rapid behavioral changes might be the onset or termination of mitigation measures, such as lockdown. Furthermore, our mitigation model only implements a single event and thus cannot account for future short-term changes. Real-world considerations, such as holidays and quarantine fatigue, would be inconsistent with our mitigation model approximation on long timescales. As such, we do not necessarily expect future predictions (which could account for future changes in mitigation) to lie within the confidence bounds of older ones.

Correlation with mobility data.—For policy and planning purposes, it is important to evaluate the effectiveness of various measures in mitigating the epidemic relative to, e.g., their economic costs. Although epidemiological modeling can, in principle, provide insight into the former issue, doing so would require a more fundamental description than the one presented here—for example, making use of agent-based models [12,13,16] or employing spatially structured and heterogeneous network descriptions [15]. Mean-field models, which assume a well-mixed population, do not capture the microscopic effects that could be correlated to specific mitigation strategies.

Nevertheless, mobility data present a potential means to guide and evaluate which policies and social interventions lead to the strongest reduction in disease transmission. While our treatment does not provide such a description or analysis and makes no claims about causality, the correlation between the effective measure of the impact of NPIs, $M(t)$ [see Eq. (3)], and mobility data presented in Fig. 6 is striking. Note that in contrast to Ref. [47], we do not use the

mobility data as an input to our calculations. Although our model makes no claims about the effect of changes in the reported mobility indices, this observed correlation encourages more detailed modeling to evaluate the impact of different social-distancing measures. Furthermore, our results support the utility of anonymized, aggregated mobility data as a potential low-latency measure of the impact of, e.g., government-mandated social-distancing measures on actual population movement, especially since such indicators may provide a predictor of the measures' influence on epidemic dynamics. A robust understanding of the mechanistic relationship between mobility measures and the spread of the infection could also guide efficient testing and contact-tracing strategies, for example, using risk-based surveillance methods [54,55].

Model generality.—Although we designed and tested our model using data for the state of Illinois, the model is general and can be fit to other geographical regions as well, as long as there is availability of hospital, ICU, and death data. Using the procedure described above, we calibrate model parameters to public data for hospital and ICU bed occupancy and daily deaths in the New York City data published by the New York City Department of Health and Mental Hygiene [56–58]. We present the results of the fit in Fig. 10. With no modification, the model generates a forecast that is remarkably consistent with subsequent data. These fits are subject to the stated limitations of our modeling. In particular, future changes in voluntary and state-mandated mitigation are, of course, not included, and our model, as described, cannot account for time-dependent changes to disease parameters (such as an increased treatment effectiveness, which could cause p_c and p_d to decrease with time). Furthermore, there are possible

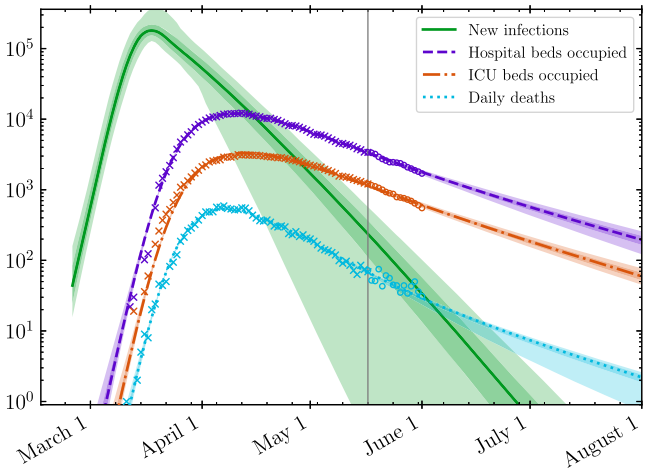

FIG. 10. Model fit and prediction for hospital and ICU bed occupancy [56,57] and deaths [58] in New York City using public data. Data from before the gray line on May 17, 2020 (crosses) are used for calibration. More recent data are plotted as circles. The model predictions shown here do not account for changes to mitigation after May 17, 2020.

indications of strong population heterogeneity and overdispersion in the unexpected smallness of τ_{disch} compared to the value for Illinois. A detailed discussion of these effects is beyond the scope of this paper. We report model parameter values in Table IV.

Spatial heterogeneity.—In reality, the entire state of Illinois is not a single well-mixed system, even if such approximations are frequently made [2,15,16]. Illinois has a densely populated region in the northeast, and more sparsely populated regions further west and to the south. Our modeling identified that the status of the epidemic (and the projections for scenarios in which mitigation is relaxed) differ among the four super-regions. In particular, in the Southern and Central regions, the epidemic is declining more rapidly because of a combination of social distancing and contact tracing using preexisting resources. Modeling these regions separately, and including population transfer between them, is essential to guide the implementation of region-dependent mitigation strategies and to provide input to policy makers to prevent the resurgence of the epidemic.

Relaxing mitigation.—Our simulations suggest that it is too soon to lift all social-distancing and mitigation restrictions, as the significant number of currently infectious individuals would make a second wave inevitable. The dynamics of a second wave and the approach to herd immunity may be crucially sensitive to effects beyond a mean-field description. For example, we do not account for superspreader events, which have played a major role in the spread of SARS and MERS [10,11,13] and are likely also relevant for COVID-19 [59]. Similarly, we do not account for heterogeneities in the population structure. Such features may accelerate the onset of herd immunity, which in turn decreases R_t. These differences have the potential to decrease the severity of the second wave and may enable more efficient containment. We defer a treatment of these effects to future work.

Additional limitations.— We end with a brief discussion of additional limitations. First, we note that our disease model is rigid in the sense that a hospitalized individual must pass through the ICU before dying, despite the fact that this may not be the case in the real world. Since our calibration procedure fits to all hospital deaths, it will be forced to adjust other parameters to compensate. Second, the importance of spatial structure, heterogeneities in population susceptibility, and social network structure is well appreciated in the literature [2,12,15]. We plan to investigate methods to model these effects and their impact on our results in future work. In addition, important sources of error in our existing model include the unknown impact of seasonal forcing, the discreteness of populations, and the effect of superspreader events and behavioral response.

In our initial modeling [48], we used seasonal forcing with $A_{\mathrm{SF}} = 0.2$ because similar effects had been documented for historical coronaviruses [19]. Because seasonal forcing varies slowly with time, it remains difficult to

verify this assumption with the observables we consider. Nevertheless, our parameter inference does not exhibit strong evidence for this level of seasonal forcing; the role of climate in the early stages of epidemics is a question of active investigation [60,61]. Furthermore, nonzero values for A_{SF} do not necessarily ascribe an equivalent seasonal modulation to the disease's infectivity.

In addition, we comment on the reported confidence intervals obtained from parameter inference for both stochastic and deterministic epidemic models [62] because these can depend on whether or not the model calibrations are performed for raw data or cumulative data. Generally speaking, the estimated intervals are unrealistically small when models are fit to cumulative data [62], with stochastic models giving slightly larger intervals. Our calculations are calibrated to raw data, not cumulative data, so we do not expect the uncertainties to be estimated inappropriately, in general. Regardless, any deterministic model of the epidemic trajectory will be inappropriate when few individuals are infected, in spite of apparent certainty in posteriors. Thus, in a regime where the deterministic model is not even appropriate as a first approximation, i.e., small populations, the estimate of uncertainties will not be a good indicator of the inapplicability of the deterministic model. Moreover, it is well known in ecology that the discreteness of populations—the fact that individuals are quantized and that birth-death processes are discrete—leads to important qualitative phenomena ranging from population cycles [63] to noise-induced pattern formation [64,65]. These effects are particularly important when numbers are small in the early and late stages of the epidemic because our mean-field model is incapable of representing the extinction state, i.e., when the number of infectious individuals drops below one. Once the epidemic is extinct in a particular region, it can only reemerge due to migration events, e.g., a super-spreader event like when university students return to campus. In this regime, a prudent regional health department uses contact tracing to contain outbreaks. Such mechanisms are not represented in our modeling.

The third major limitation is that modeling an epidemic is very different from modeling a physical system, even one as complex as a weather pattern. The transmission of an infectious disease involves a collaboration between the virus and the host population: The host population alters its behavior in response to its awareness of epidemic progress, leading to policy steps that may increase or decrease transmission, and self-regulation of social distancing by susceptible and vulnerable populations. If the host population behavior changes on a much slower timescale compared to the disease dynamics, the model predictions are likely to be more reliable than if the changes occur more rapidly. Thus, it is important to emphasize that predictions can easily be invalidated due to subsequent human actions that cannot be anticipated and will be impossible to model precisely.

VII. CONCLUSIONS

In this paper, we have presented a mathematical model and computational framework for recapitulating the COVID-19 epidemic. This model may be used to infer values for a set of effective parameters that drive trajectories representing the progression of the disease. We use our calibrated model to provide robust short-term forecasts of the epidemic trajectory in different regions of the state and explore the effects that steps to relax social-distancing measures may have, especially in the context of a second wave of the epidemic. The resulting highly constrained and quantitative narrative of the epidemic is a useful tool to inform scenarios for sustainable monitoring and control of the epidemic.

VIII. DATA AVAILABILITY

The model and calibration framework described above have been implemented in the open source PYTHON3 [66] package PYDEMIC. The source code for pydemic is freely available online at [67]. This work made use of NumPy [68], SciPy [69], PANDAS [70], EMCEE [34], corner.py [71], and MATPLOTLIB [72].

ACKNOWLEDGMENTS

We gratefully acknowledge discussions with David Ansell at Rush University Hospital, Mark Johnson at Carle Hospital, Katie Gostic and Sarah Cobey at University of Chicago, Jaline Gerardin at Northwestern University, Charles Gammie at the University of Illinois, and Josh Speagle at Harvard University. We also thank the two anonymous referees for many clarifying comments and discussions that greatly improved the text. The calculations we performed would have been impossible without the data kindly provided by the Illinois Department of Public Health through a Data Use Agreement with Civis Analytics. This work was supported by the University of Illinois System Office, the Office of the Vice-Chancellor for Research and Innovation, the Grainger College of Engineering, and the Department of Physics at the University of Illinois at Urbana-Champaign. Z. J. W. is supported in part by the U.S. Department of Energy (DOE) Computational Science Graduate Fellowship, provided under Grant No. DE-FG02-97ER25308. This work made use of the Illinois Campus Cluster, a computing resource that is operated by the Illinois Campus Cluster Program (ICCP) in conjunction with the National Center for Supercomputing Applications (NCSA), which is supported by funds from the University of Illinois at Urbana-Champaign. This research was partially completed at, and used resources of, the Center for Functional Nanomaterials, which is a U.S. DOE Office of Science Facility, at Brookhaven National Laboratory under Contract No. DE-SC0012704.

APPENDIX A: SEVERITY MODEL DESCRIPTION

We detail the parametrization of our age-dependent severity model below. As noted in Sec. II A, we track the number of individuals in different age groups separately, and thus we specify the fiducial transition probabilities for each age group independently. We further define an overall scaling prefactor p_x that governs the final transition probability $p_{x,i}$ as the product of the fiducial rate and the prefactor. Our fiducial transition rates $p_{x,i}/p_x$ are drawn from Refs. [29,73–75].

In Table II, we list the age-dependent values for the following probabilities: that infected individuals experience symptoms, $p_{\sigma,i}/p_\sigma$; that symptomatic individuals are hospitalized, $p_{h,i}/p_h$; that hospitalized patients enter the ICU, $p_{c,i}/p_c$; and that ICU patients expire, $p_{d,i}/p_d$. We also list the relative age distribution of the population in

TABLE II. Table of age-specific parameters, including the age distribution w_i and the various probabilities of transitioning between various states in our model. We present the unscaled distribution $p_{x,i}/p_x$, where p_x denotes the overall scaling that is sampled during parameter inference (or whose value is otherwise set as described in the text).

	[0, 10)	[10, 20)	[20, 30)	[30, 40)	[40, 50)	[50, 60)	[60, 70)	[70, 80)	[80, ∞)
w_i	0.120004	0.127891	0.139256	0.134948	0.121898	0.12725	0.116278	0.0727565	0.0397193
$p_{\sigma,i}/p_\sigma$	0.057	0.054	0.294	0.668	0.614	0.83	0.99	0.995	0.999
$p_{h,i}/p_h$	0.001	0.003	0.012	0.032	0.049	0.102	0.166	0.243	0.273
$p_{c,i}/p_c$	0.05	0.05	0.05	0.05	0.063	0.122	0.274	0.432	0.709
$p_{d,i}/p_d$	0.3	0.3	0.3	0.3	0.3	0.4	0.4	0.5	0.5

TABLE III. Table of inferred model parameters for disease severity and the time dependence of mitigation effects. Mitigation was modeled with a piecewise Hermite cubic interpolating polynomial as described in Sec. II A. Each column reports median parameter values from the MCMC algorithm for the different model fits, along with the largest of the upper and lower 1σ error bounds reported by the algorithm. In addition, we report the (derived) values and uncertainties of R_t [as defined in Eq. (3)] for each model, as evaluated on May 1, 2020. The column headers specify the modeled region of Illinois and the inclusive end date of the calibration data.

	Illinois (5/17)	Illinois (4/20)	Northeastern (5/17)	North-Central (5/17)	Central (5/17)	Southern (5/17)
R_0	2.2 ± 0.26	2.2 ± 0.25	2.3 ± 0.28	1.6 ± 0.38	1.6 ± 0.4	1.6 ± 0.33
t_s	Feb 15 ± 3.92	Feb14 ± 4.24	Feb 16 ± 3.73	Feb 16 ± 7.82	Feb 18 ± 8.25	Feb 17 ± 8.40
t_0	Mar 9 ± 4.13	Mar 10 ± 4.13	Mar 9 ± 4.04	Mar 8 ± 5.71	Mar 8 ± 6.08	Mar 11 ± 7.62
t_1	Mar 20 ± 2.94	Mar 21 ± 3.28	Mar 20 ± 3.01	Mar 24 ± 6.04	Mar 25 ± 6.58	Apr 4 ± 5.53
$M(t_1)$	0.48 ± 0.051	0.47 ± 0.048	0.45 ± 0.048	0.64 ± 0.12	0.61 ± 0.12	0.59 ± 0.099
R_t(May 1)	0.94 ± 0.01	0.97 ± 0.037	0.94 ± 0.011	0.99 ± 0.02	0.95 ± 0.023	0.9 ± 0.027
IFR	0.007 ± 0.0012	0.0071 ± 0.0012	0.007 ± 0.0012	0.007 ± 0.0012	0.007 ± 0.0012	0.007 ± 0.0012
τ_h	6.3 ± 2	6.2 ± 2	6.2 ± 2	6.1 ± 1.8	6.2 ± 1.8	6.9 ± 1.8
σ_h	2.3 ± 1.4	2.2 ± 1.6	2.4 ± 1.7	4.2 ± 1.9	4.2 ± 1.9	4.6 ± 1.7
τ_{disch}	7 ± 1.9	6.9 ± 1.9	7.1 ± 1.9	6.2 ± 1.7	6.6 ± 1.8	6.8 ± 1.9
σ_{disch}	6.1 ± 1.8	5.7 ± 2	6.4 ± 1.7	3.5 ± 1.7	4.6 ± 1.9	5 ± 1.8
p_c	0.61 ± 0.16	0.67 ± 0.18	0.62 ± 0.17	0.47 ± 0.14	0.32 ± 0.097	0.51 ± 0.14
τ_c	1.4 ± 0.97	1.7 ± 1.3	1.3 ± 0.93	3 ± 1.6	2.2 ± 1.5	2 ± 1.3
σ_c	2.2 ± 1.8	2 ± 1.8	2.3 ± 1.9	2.5 ± 1.6	2.5 ± 1.7	2.7 ± 1.7
p_d	1.4 ± 0.15	1.3 ± 0.16	1.4 ± 0.16	1.6 ± 0.3	1.9 ± 0.25	1.9 ± 0.26
τ_d	8.8 ± 0.96	7.7 ± 1	8.8 ± 0.98	13 ± 2.1	8.9 ± 1.6	11 ± 1.4
σ_d	4.6 ± 1.5	3.7 ± 1.5	4.7 ± 1.6	9.2 ± 2.4	7 ± 2.7	6.6 ± 2.2
τ_{rec}	11 ± 2.3	12 ± 2.5	11 ± 2.3	12 ± 2.8	12 ± 2.8	12 ± 2.8
σ_{rec}	6.8 ± 2.7	8.7 ± 2.9	6.9 ± 2.7	10 ± 2.8	10 ± 2.8	9.9 ± 2.9
F_{tot}	1.4 ± 0.033	1.3 ± 0.06	1.4 ± 0.034	1.2 ± 0.14	1.6 ± 0.3	2.1 ± 0.28
τ_{tot}	2.9 ± 0.46	2 ± 0.53	2.8 ± 0.47	2.2 ± 0.88	2.3 ± 0.91	2.8 ± 0.88
σ_{tot}	2.3 ± 0.88	1.9 ± 0.84	2.2 ± 0.86	2.1 ± 0.91	2 ± 0.92	1.9 ± 0.91
A_{SF}	0.15 ± 0.03	0.06 ± 0.071	0.11 ± 0.035	0.13 ± 0.065	0.12 ± 0.07	0.13 ± 0.073

TABLE IV. Same as Table III, but for fits for New York City.

	NYC (5/17)
R_0	3.4 ± 0.26
t_s	Feb 20 ± 1.69
t_0	Mar 3 ± 2.05
t_1	Mar 24 ± 2.38
$M(t_1)$	0.27 ± 0.039
R_t(May 1)	0.56 ± 0.042
IFR	0.0056 ± 0.00099
τ_h	10 ± 1.5
σ_h	6.2 ± 1
τ_{disch}	2.3 ± 1
σ_{disch}	9.4 ± 1.6
p_c	0.23 ± 0.11
τ_c	1.2 ± 0.31
σ_c	1.9 ± 0.76
p_d	2.1 ± 0.2
τ_d	8.3 ± 0.77
σ_d	17 ± 1
τ_{rec}	10 ± 3.5
σ_{rec}	12 ± 3.3
A_{SF}	0.13 ± 0.053

the United States, which we source from the UN [76]. As specified in Table I, we sample over the p_c and p_d scale factors.

Because the data to which we calibrate do not constrain the symptomatic population, we cannot observe p_h, so we fix it to one. We force our model to produce a given infection fatality ratio by setting p_σ to

$$p_\sigma = \frac{\mathrm{IFR}}{\sum_i w_i (p_{\sigma,i}/p_\sigma) p_{c,i} p_{h,i} p_{d,i} F_{\mathrm{tot}}}, \tag{A1}$$

where the value for $p_{\sigma,i}/p_\sigma$ on the right-hand side of the equation is taken first from Table II. This equation is valid if the initial infected population is distributed in proportion to w_i.

APPENDIX B: MODEL FIT POSTERIOR PROBABILITY DISTRIBUTION DETAILS

In Table III, we summarize the posterior probability distribution for all sample parameters enumerated in Table I. Figure 11 shows joint posterior distributions for all pairs over the complete set of parameters of our model fitted to all-state data up to May 17, 2020 as shown in the bottom panel of Fig. 2.

APPENDIX C: SYNTHETIC DATA TEST

In Fig. 12, we present the result of performing calibration on synthetic data to test our model's ability to recover a set of known parameters, and Fig. 13 shows a comparison of the posterior probability distribution estimates produced by the model compared to the truth values used to generate the synthetic data. We choose to generate synthetic data from our own model since the meaning of parameters may differ between different model implementations and network structures; thus, it may be difficult to compare inferred numerical parameters between different models. We introduce uncertainty into the synthetic data by Poisson sampling the time-dependent output model rates. We consider three different scenarios, each of varying difficulty.

In the first two "No Mitigation" scenarios, the peak is due to the onset of herd immunity. For calibration, only data from before the peak are used, so the timing of the peak can only be inferred via IFR from deaths. In the exponential growth phase of the epidemic, it is challenging to determine the timescales that govern removal from a state (e.g., being discharged from the hospital), so the quality of the forecast will be a strong function of the quality of the prior. We present both "Good Prior" and "Bad Prior" scenarios, corresponding to the cases where the synthetic data are generated from within the prior or otherwise. In the "Bad Prior" scenario, the data do not always lie within the forecast confidence interval. The likely strength of such disagreements is decreased in real-world scenarios in which the priors are well motivated. In the third scenario, we consider an epidemic that experiences mitigation and calibrate on data that come after the peak. In this scenario, the removal timescales are much better constrained by the calibration procedure since it has access to the downward trending part of the epidemic trajectory.

APPENDIX D: POISSON VERSUS NEGATIVE-BINOMIAL LIKELIHOOD

In Fig. 14, we evaluate the appropriateness of the Poisson estimator for the incidence data we use for calibration. While daily positive tests are clearly over-dispersed, the daily hospital deaths that are an input to our model are evidently well described by the Poisson distribution. For this reason, our use of a Poisson likelihood estimator is justified in our MCMC procedure.

APPENDIX E: CASE DATA ASCERTAINMENT BIAS

Because our model dynamics relies on an internal estimate for the number of infected individuals, we can compare the number of new daily positive tests (from data) to the model-predicted daily infections and estimate the testing ascertainment bias. In Fig. 15, we plot the ratio of daily positive tests to the model-predicted number of new infections. We note that dividing the number of positive tests by the number of new infections on the same day may be slightly misrepresentative since there is a delay between when an individual is infected and the time when they can first test positive; however, this delay is short compared to the timescale on which the dynamics changes, so we do not

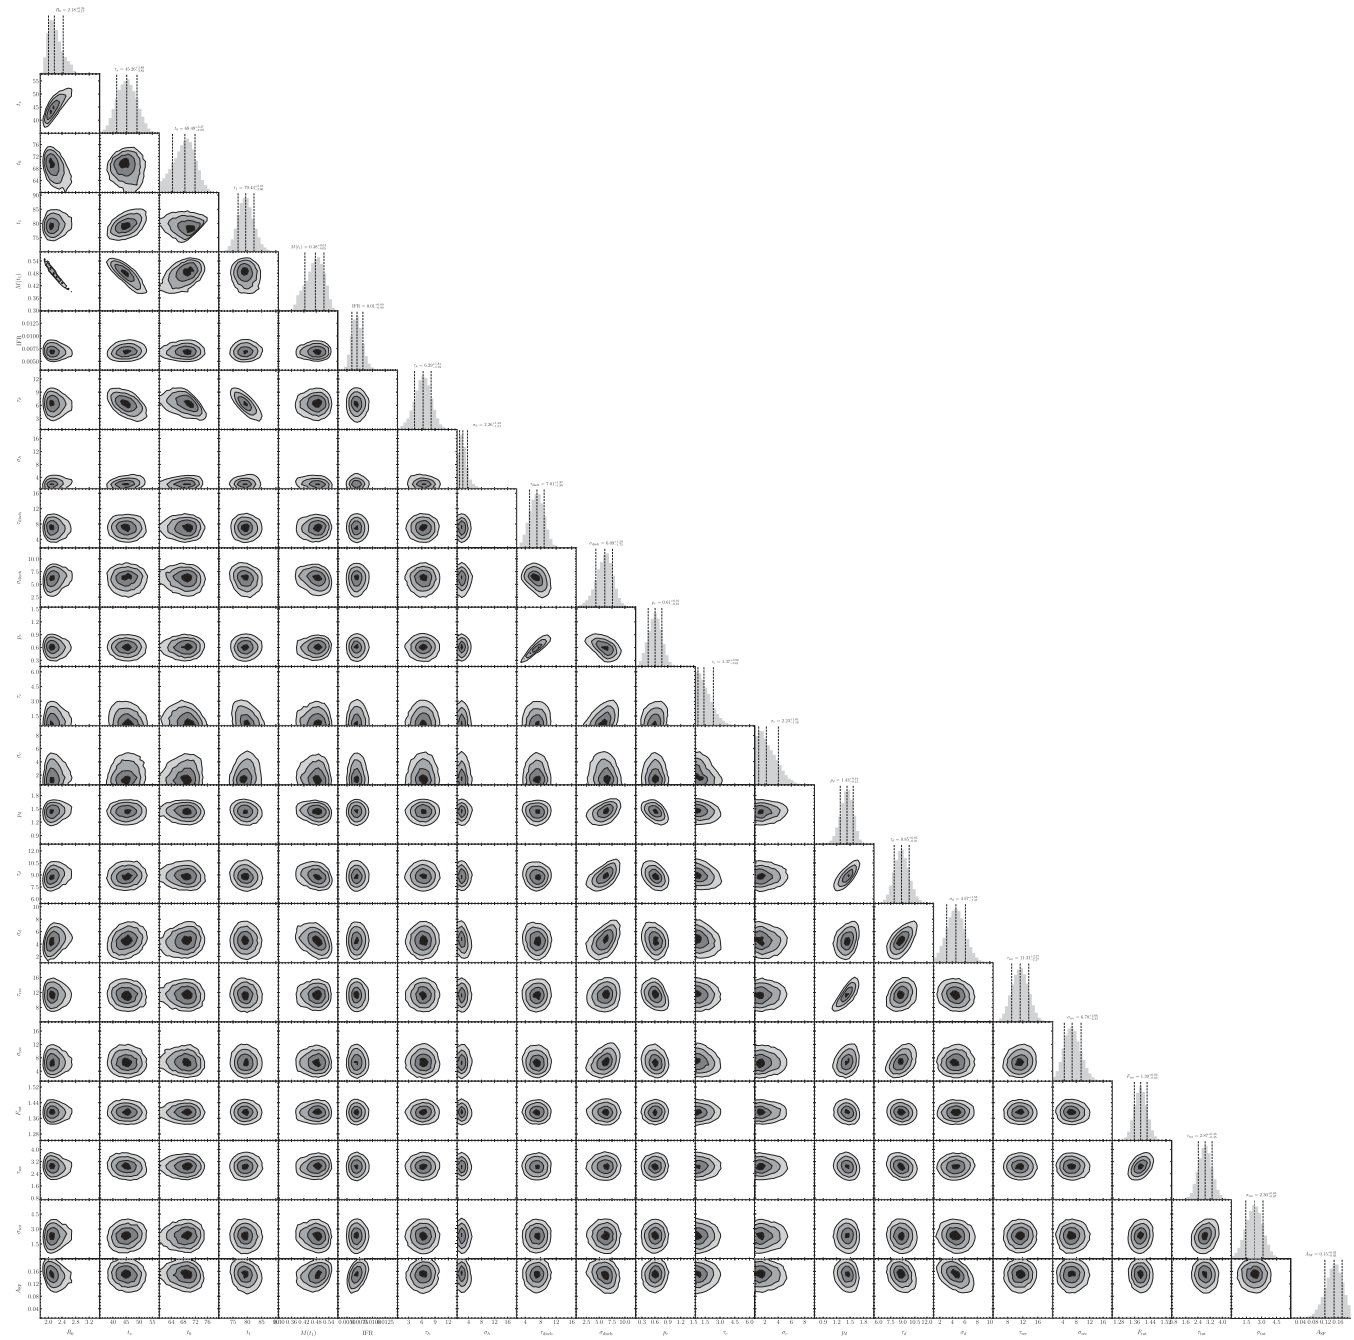

FIG. 11. Joint posterior distributions of pairs of the complete set of parameters of our model fitted to the all-state data up to May 17, 2020. The correlations between some pairs of fitted parameters such as, e.g., between R_0 and the start date of the epidemic are reflected in the ellipsoidal shape of the posteriors. This figure is an expanded version of Fig. 5.

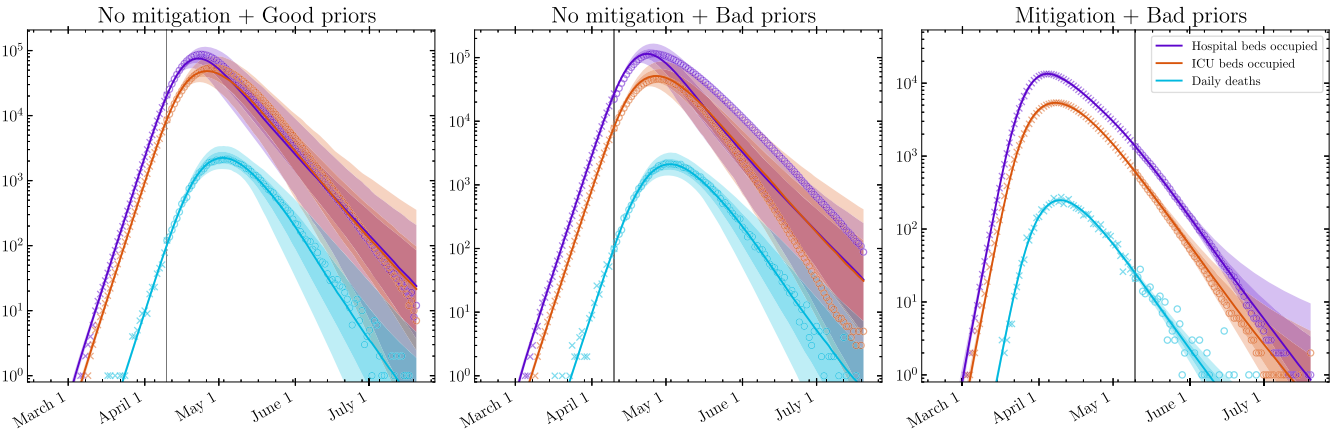

FIG. 12. Comparison of synthetic (generated) epidemic trajectories against calibration prediction for hospital beds occupied, ICU beds occupied, and daily deaths. Only data up to the vertical line (x markers) are used during the calibration procedure. The forecast for the downward part of the trajectories is heavily influenced by the removal timescales, which exhibit a low signal in the initial exponential phase, so the quality of the forecast for the "No Mitigation" scenarios is influenced by the quality of the prior. In the "Good (Bad) Prior" scenario, the removal timescales are chosen to be consistent (inconsistent) with the prior. Increasing the amount of data fed to the calibration procedure (as in the "Mitigation" scenario) allows the sampler to better infer the removal timescales.

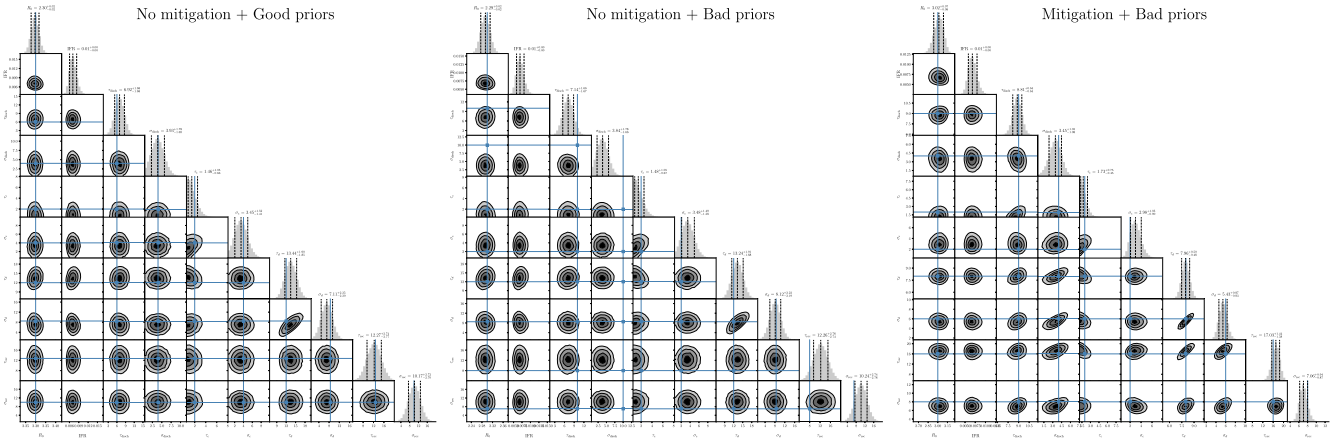

FIG. 13. Estimate of posterior probability distribution for synthetic data calibration tests shown in Fig. 12. Blue lines or points show the values of parameters used to generate the synthetic data set. In the "No Mitigation + Bad Priors" scenario, the truth values for some of the removal timescales were chosen to lie well beyond the prior since the calibration data did not include the downward part of the trajectory. By including data from after the peak, as is done for the "Mitigation + Bad Priors" scenario, the removal timescales can be better constrained.

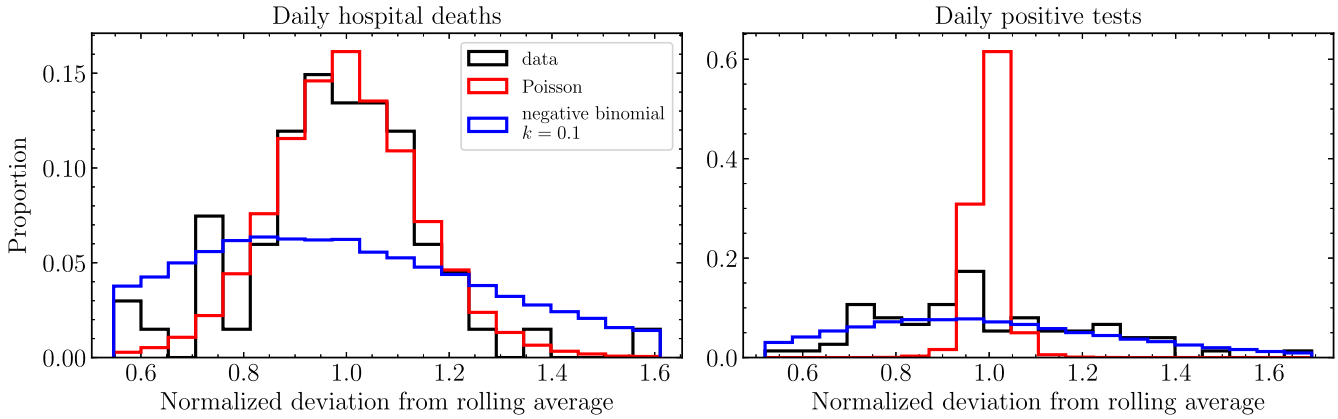

FIG. 14. Quantification of the dispersion of reported data in Illinois. For both the daily incidence of hospital deaths and of reported positive tests (right panel), we plot histograms of the raw data normalized by its 7-day rolling average. To compare this measure to theoretical distributions, we additionally plot the same measure of random samples of the Poisson distribution and negative binomial distribution (with dispersion coefficient $k = 0.1$), averaging over 1000 such samples. Note that we do not consider positive tests in our calibration procedure.

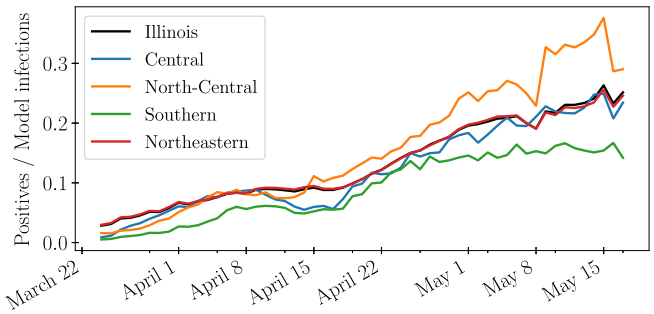

FIG. 15. Ratio of daily positive tests to estimated true number of daily infections in the different regions of Illinois. A steady-state ratio of one would imply that every new infection is caught by testing. We caution that the value of the ratio is influenced by the prior on IFR, which we do not calibrate and instead draw from the literature.

believe that this approximation severely affects the reported ratio. The steady increase in that ratio over time may correspond to a ramp-up in testing across the population during the early stages of the epidemic. Since the ratio relies on the model output for new daily infections, it is directly related to the prior we use for IFR, which determines the number of new daily infections from the number of daily deaths. Even after the initial ramp-up, we do not expect the ratio to remain constant since testing coverage may change and the demographics of infected-versus-tested individuals can evolve.

[1] All executive orders related to COVID-19 issued by the State of Illinois Governor J. B. Pritzker are listed here: https://www2.illinois.gov/government/executive-orders. The timeline of COVID-19 mitigation in Illinois is summarized at https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Illinois#Government_response.

[2] M. J. Keeling and P. Rohani, *Modeling Infectious Diseases in Humans and Animals* (Princeton University, Princeton, NJ, 2011).

[3] C. J. L. Murray, *Forecasting COVID-19 Impact on Hospital Bed-Days, ICU-Days, Ventilator-Days and Deaths by US State in the Next 4 Months*, medRxiv, https://doi.org/10.1101/2020.03.27.20043752.

[4] J. Ma, *Estimating Epidemic Exponential Growth Rate and Basic Reproduction Number*, Infect. Dis. Model. **5**, 129 (2020).

[5] J. T. Wu, K. Leung, and G. M. Leung, *Nowcasting and Forecasting the Potential Domestic and International Spread of the 2019-nCoV Outbreak Originating in Wuhan, China: A Modelling Study*, Lancet **395**, 689 (2020).

[6] C. Fraser, *Estimating Individual and Household Reproduction Numbers in an Emerging Epidemic*, PLoS One **2**, e758 (2007).

[7] G. Chowell, *Fitting Dynamic Models to Epidemic Outbreaks with Quantified Uncertainty: A Primer for Parameter Uncertainty, Identifiability, and Forecasts*, Infectious Disease Modelling **2**, 379 (2017).

[8] J. T. Wu, K. Leung, M. Bushman, N. Kishore, R. Niehus, P. M. de Salazar, B. J. Cowling, M. Lipsitch, and G. M. Leung, *Estimating Clinical Severity of COVID-19 from the Transmission Dynamics in Wuhan, China*, Nat. Med. **26**, 506 (2020).

[9] C. Viboud, O. N. Bjørnstad, D. L. Smith, L. Simonsen, M. A. Miller, and B. T. Grenfell, *Synchrony, Waves, and Spatial Hierarchies in the Spread of Influenza*, Science **312**, 447 (2006).

[10] J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, and W. M. Getz, *Superspreading and the Effect of Individual Variation on Disease Emergence*, Nature (London) **438**, 355 (2005).

[11] M. Small, C. Tse, and D. M. Walker, *Super-Spreaders and the Rate of Transmission of the SARS Virus*, Physica D (Amsterdam) **215**, 146 (2006).

[12] S. Bansal, B. T. Grenfell, and L. A. Meyers, *When Individual Behaviour Matters: Homogeneous and Network Models in Epidemiology*, J. R. Soc. Interface **4**, 879 (2007).

[13] Y. Kim, H. Ryu, and S. Lee, *Agent-Based Modeling for Super-Spreading Events: A Case Study of MERS-CoV Transmission Dynamics in the Republic of Korea*, Int. J. Environ. Res. Public Health **15**, 2369 (2018).

[14] Z. Dezső and A.-L. Barabási, *Halting Viruses in Scale-Free Networks*, Phys. Rev. E **65**, 055103 (2002).

[15] K. Rock, S. Brand, J. Moir, and M. J. Keeling, *Dynamics of Infectious Diseases*, Rep. Prog. Phys. **77**, 026602 (2014).

[16] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, *Epidemic Processes in Complex Networks*, Rev. Mod. Phys. **87**, 925 (2015).

[17] W. O. Kermack and A. G. McKendrick, *A Contribution to the Mathematical Theory of Epidemics*, Proc. R. Soc. A **115**, 700 (1927).

[18] J. Wallinga and M. Lipsitch, *How Generation Intervals Shape the Relationship between Growth Rates and Reproductive Numbers*, Proc. R. Soc. B **274**, 599 (2007).

[19] R. A. Neher, R. Dyrdak, V. Druelle, E. B. Hodcroft, and J. Albert, *Potential Impact of Seasonal Forcing on a SARS-CoV-2 Pandemic*, Swiss Medical Weekly **150**, 1112 (2020).

[20] A. Smirnova, L. deCamp, and G. Chowell, *Forecasting Epidemics through Nonparametric Estimation of Time-Dependent Transmission Rates Using the SEIR Model*, Bull. Math. Biol. **81**, 4343 (2019).

[21] C. C. Ku, T.-C. Ng, and H.-H. Lin, *Epidemiological Benchmarks of the COVID-19 Outbreak Control in China after Wuhan's Lockdown: A Modelling Study with an Empirical Approach*, SSRN Electron., https://doi.org/10.2139/ssrn.3544127 (2020).

[22] R. H. Mena, J. X. Velasco-Hernandez, N. B. Mantilla-Beniers, G. A. Carranco-Sapiéns, L. Benet, D. Boyer, and I. P. Castillo, *Using the Posterior Predictive Distribution to Analyse Epidemic Models: COVID-19 in Mexico City*, arXiv:2005.02294.

[23] M. Gatto, E. Bertuzzo, L. Mari, S. Miccoli, L. Carraro, R. Casagrandi, and A. Rinaldo, *Spread and Dynamics of the COVID-19 Epidemic in Italy: Effects of Emergency Containment Measures*, Proc. Natl. Acad. Sci. U.S.A. **117**, 10484 (2020).

[24] R. Li, S. Pei, B. Chen, Y. Song, T. Zhang, W. Yang, and J. Shaman, *Substantial Undocumented Infection Facilitates the Rapid Dissemination of Novel Coronavirus (SARS-CoV-2)*, Science **368**, 489 (2020).

[25] H. Nishiura, N. M. Linton, and A. R. Akhmetzhanov, *Serial Interval of Novel Coronavirus (COVID-19) Infections*, Int. J. Infectious Diseases **93**, 284 (2020).

[26] Z. Du, X. Xu, Y. Wu, L. Wang, B. J. Cowling, and L. A. Meyers, *Serial Interval of COVID-19 among Publicly Reported Confirmed Cases*, Emerg. Infect. Dis. **26**, 1341 (2020).

[27] S. A. Lauer, K. H. Grantz, Q. Bi, F. K. Jones, Q. Zheng, H. R. Meredith, A. S. Azman, N. G. Reich, and J. Lessler, *The Incubation Period of Coronavirus Disease 2019 (COVID-19) from Publicly Reported Confirmed Cases: Estimation and Application*, Ann. Intern. Med. **172**, 577 (2020).

[28] N. M. Linton, T. Kobayashi, Y. Yang, K. Hayashi, A. R. Akhmetzhanov, S.-m. Jung, B. Yuan, R. Kinoshita, and H. Nishiura, *Incubation Period and Other Epidemiological Characteristics of 2019 Novel Coronavirus Infections with Right Truncation: A Statistical Analysis of Publicly Available Case Data*, J. Clin. Med. **9**, 538 (2020).

[29] R. Verity *et al.*, *Estimates of the Severity of Coronavirus Disease 2019: A Model-Based Analysis*, Lancet Infect. Dis. (2020) [https://www.ncbi.nlm.nih.gov/pubmed/32240634].

[30] P. K. Bhatraju, B. J. Ghassemieh, M. Nichols, R. Kim, K. R. Jerome, A. K. Nalla, A. L. Greninger, S. Pipavath, M. M. Wurfel, L. Evans, P. A. Kritek, T. E. West, A. Luks, A. Gerbino, C. R. Dale, J. D. Goldman, S. O'Mahony, and C. Mikacenic, *Covid-19 in Critically Ill Patients in the Seattle Region—Case Series*, N. Engl. J. Med. (2020) [https://www.ncbi.nlm.nih.gov/pubmed/32227758].

[31] X. Yang, Y. Yu, J. Xu, H. Shu, J. Xia, H. Liu, Y. Wu, L. Zhang, Z. Yu, M. Fang, T. Yu, Y. Wang, S. Pan, X. Zou, S. Yuan, and Y. Shang, *Clinical Course and Outcomes of Critically Ill Patients with SARS-CoV-2 Pneumonia in Wuhan, China: A Single-Centered, Retrospective, Observational Study*, Lancet Resp. Med. **8**, 475 (2020).

[32] *ICNARC Report on COVID-19 in Critical Care*, https://www.icnarc.org/Our-Audit/Audits/Cmp/Reports (2020).

[33] G. Meyerowitz-Katz and L. Merone, *A Systematic Review and Meta-analysis of Published Research Data on COVID-19 Infection-Fatality Rates*, medRxiv, https://doi.org/10.1101/2020.05.03.20089854.

[34] D. Foreman-Mackey, W. Farr, M. Sinha, A. Archibald, D. Hogg, J. Sanders, J. Zuntz, P. Williams, A. Nelson, M. de Val-Borro, T. Erhardt, I. Pashchenko, and O. Pla, *EMCEE V3: A PYTHON Ensemble Sampling Toolkit for Affine-Invariant MCMC*, J. Open Source Software **4**, 1864 (2019).

[35] C. J. F. ter Braak and J. A. Vrugt, *Differential Evolution Markov Chain with Snooker Updater and Fewer Chains*, Stat. Comput. **18**, 435 (2008).

[36] C. J. F. ter Braak, *A Markov Chain Monte Carlo Version of the Genetic Algorithm Differential Evolution: Easy Bayesian Computing for Real Parameter Spaces*, Stat. Comput. **16**, 239 (2006).

[37] B. Farr, V. Kalogera, and E. Luijten, *A More Efficient Approach to Parallel-Tempered Markov-Chain Monte Carlo for the Highly Structured Posteriors of Gravitational-Wave Signals*, Phys. Rev. D **90**, 024014 (2014).

[38] Data were downloaded from https://www.dph.illinois.gov/covid19/covid19-statistics.

[39] I. Korolev, *Identification and Estimation of the SEIRD Epidemic Model for COVID-19*, https://ssrn.com/abstract=3569367.

[40] M. Castro, S. Ares, J. A. Cuesta, and S. Manrubia, *The Turning Point and End of an Expanding Epidemic Cannot Be Precisely Forecast*, Proc. Natl. Acad. Sci. U.S.A. **117**, 26190 (2020).

[41] W. C. Roda, M. B. Varughese, D. Han, and M. Y. Li, *Why Is It Difficult to Accurately Predict the COVID-19 Epidemic?*, Infectious Disease Modelling **5**, 271 (2020).

[42] See map at https://www.dph.illinois.gov/sites/default/files/resources/ems-regions-map.pdf.

[43] S. J. Beckett, M. Dominguez-Mirazo, S. Lee, C. Andris, and J. S. Weitz, *Spread of COVID-19 through Georgia, USA.*

Near-Term Projections and Impacts of Social Distancing via a Metapopulation Model, https://github.com/WeitzGroup/MAGEmodel_covid19_GA/blob/master/Report/GA_COVID19_assessment_21Apr2020.pdf.

[44] See https://www.blog.google/technology/health/covid-19-community-mobility-reports?hl=en.

[45] See https://www.unacast.com/post/rounding-out-the-social-distancing-scoreboard.

[46] S. Flaxman, S. Mishra, A. Gandy, H. Unwin, H. Coupland, T. Mellan *et al.*, *Report 13. Estimating the Number of Infections and the Impact of Non-pharmaceutical Interventions on COVID-19 in 11 European Countries* (Imperial College, London, 2020).

[47] H. J. T. Unwin *et al.*, *Report 23: State-Level Tracking of COVID-19 in the United States WHO Collaborating Centre for Infectious Disease Modelling MRC Centre for Global Infectious Disease Analytics*, 2020.

[48] S. Maslov and N. Goldenfeld, *Window of Opportunity for Mitigation to Prevent Overflow of ICU Capacity in Chicago by COVID-19*, https://www.medrxiv.org/content/early/2020/03/24/2020.03.20.20040048.

[49] D. H. Morris, F. W. Rossine, J. B. Plotkin, and S. A. Levin, *Optimal, Near-Optimal, and Robust Epidemic Control*, arXiv:2004.02209.

[50] J. Hellewell, S. Abbott, A. Gimma, N. I. Bosse, C. I. Jarvis, T. W. Russell, J. D. Munday, A. J. Kucharski, and R. M. Eggo, *Feasibility of Controlling COVID-19 Outbreaks by Isolation of Cases and Contacts*, The Lancet Global Health, 2020.

[51] L. Ferretti, C. Wymant, M. Kendall, L. Zhao, A. Nurtay, L. Abeler-Dorner, M. Parker, D. Bonsall, and C. Fraser, *Quantifying SARS-CoV-2 Transmission Suggests Epidemic Control with Digital Contact Tracing*, Science (2020) eabb6936 [https://www.ncbi.nlm.nih.gov/pubmed/32234805].

[52] *Fitzhugh Mullan Institute for Health Workforce Equity and George Washington University, Contact Tracing Workforce Estimator*, https://www.gwhwi.org/estimator-613404.html.

[53] I. Holmdahl and C. Buckee, *Wrong but Useful—What Covid-19 Epidemiologic Models Can and Cannot Tell Us*, New England J. Med., https://doi.org/10.1056/NEJMp2016822.

[54] K. D. Stärk, G. Regula, J. Hernandez, L. Knopf, K. Fuchs, R. S. Morris, and P. Davies, *Concepts for Risk-Based Surveillance in the Field of Veterinary Medicine and Veterinary Public Health: Review of Current Approaches*, BMC Health Services Res. **6**, 20 (2006).

[55] A. Foddai, A. Lindberg, J. Lubroth, and J. Ellis-Iversen, *Surveillance to Improve Evidence for Community Control Decisions During the COVID-19 Pandemic—Opening the Animal Epidemic Toolbox for Public Health*, One Health **9**, 100130 (2020).

[56] Data were downloaded from https://github.com/thecityny/covid-19-nyc-data.

[57] Data originally from https://www.thecity.nyc.

[58] Data were downloaded from https://github.com/nychealth/coronavirus-data.

[59] Y. Liu, R. M. Eggo, and A. J. Kucharski, *Secondary Attack Rate and Superspreading Events for SARS-CoV-2*, Lancet **395**, e47 (2020).

[60] K. S. Raines, S. Doniach, and G. Bhanot, *The Transmission of SARS-CoV-2 Is Likely Comodulated by Temperature and by Relative Humidity*, https://www.medrxiv.org/content/early/2020/05/26/2020.05.23.20111278.

[61] R. E. Baker, W. Yang, G. A. Vecchi, C. J. E. Metcalf, and B. T. Grenfell, *Susceptible Supply Limits the Role of Climate in the COVID-19 Pandemic*, https://www.medrxiv.org/content/early/2020/04/07/2020.04.03.20052787.

[62] A. A. King, M. D. de Cellès, F. M. Magpantay, and P. Rohani, *Avoidable Errors in the Modelling of Outbreaks of Emerging Pathogens, with Special Reference to Ebola*, Proc. R. Soc. B **282**, 0347 (2015).

[63] A. J. McKane and T. J. Newman, *Predator-Prey Cycles from Resonant Amplification of Demographic Stochasticity*, Phys. Rev. Lett. **94**, 218102 (2005).

[64] T. Butler and N. Goldenfeld, *Robust Ecological Pattern Formation Induced by Demographic Noise*, Phys. Rev. E **80**, 030902 (2009).

[65] D. Karig, K. M. Martini, T. Lu, N. A. DeLateur, N. Goldenfeld, and R. Weiss, *Stochastic Turing Patterns in a Synthetic Bacterial Population*, Proc. Natl. Acad. Sci. U.S.A. **115**, 6572 (2018).

[66] G. Van Rossum and F. L. Drake, *PYTHON3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).

[67] https://github.com/uiuc-covid19-modeling/pydemic.

[68] T. E. Oliphant, *A Guide to NumPy* (Trelgol Publishing, 2006), Vol. 1.

[69] P. Virtanen *et al.*, *SciPy 1.0: Fundamental Algorithms for Scientific Computing in PYTHON*, Nat. Methods **17**, 261 (2020).

[70] T. Pandas Development Team, pandas-dev/pandas: Pandas, https://doi.org/10.5281/zenodo.3509134.

[71] D. Foreman-Mackey, *corner.py: Scatterplot Matrices in PYTHON*, J. Open Source Software **24** (2016) [https://doi.org/10.5281/zenodo.45906].

[72] J. D. Hunter, *Matplotlib: A 2D Graphics Environment*, Comput. Sci. Eng. **9**, 90 (2007).

[73] Epidemiological Group of the New Coronavirus Pneumonia Emergency Response Mechanism of the China Center for Disease Control and Prevention, *Epidemiological Characteristic Analysis of New Coronavirus Pneumonia*, Chin. J. Epidemiol. **41**, 145 (2020).

[74] K. Mizumoto, K. Kagaya, A. Zarebski, and G. Chowell, *Estimating the Asymptomatic Proportion of Coronavirus Disease 2019 (COVID-19) Cases on Board the Diamond Princess Cruise Ship, Yokohama, Japan, 2020*, Eurosurveillance **25**, 10 (2020).

[75] H. Nishiura, T. Kobayashi, T. Miyama, A. Suzuki, S. Jung, K. Hayashi, R. Kinoshita, Y. Yang, B. Yuan, A. R. Akhmetzhanov, and N. M. Linton, *Estimation of the Asymptomatic Ratio of Novel Coronavirus Infections (COVID-19)*, https://www.medrxiv.org/content/early/2020/02/17/2020.02.03.20020248.

[76] Data were downloaded from https://data.un.org.