

Approximate Bayesian Computation for an Explicit-Duration Hidden Markov Model of COVID-19 Hospital Trajectories

Gian Marco Visani¹

GIAN_MARCO.VISANI@TUFTS.EDU

Alexandra Hope Lee¹

ALEXANDRA.LEE@TUFTS.EDU

Cuong Nguyen¹, M.S.

CUONG.NGUYEN@TUFTS.EDU

David M. Kent², M.D., C.M., M.Sc.

DKENT1@TUFTSMEDICALCENTER.ORG

John B. Wong³, M.D.

JWONG@TUFTSMEDICALCENTER.ORG

Joshua T. Cohen⁴, Ph.D.

JCOHEN@TUFTSMEDICALCENTER.ORG

Michael C. Hughes¹, Ph.D.

MICHAEL.HUGHES@TUFTS.EDU

¹ *Dept. of Computer Science, Tufts University, Medford, MA, USA*

² *Predictive Analytics and Comparative Effectiveness Center, Tufts Medical Center, Boston, MA*

³ *Division of Clinical Decision Making, Tufts Medical Center, Boston, MA*

⁴ *Center for the Evaluation of Value and Risk in Health, Tufts Medical Center, Boston, MA*

Abstract

We address the problem of modeling constrained hospital resources in the midst of the COVID-19 pandemic in order to inform decision-makers of future demand and assess the societal value of possible interventions. For broad applicability, we focus on the common yet challenging scenario where patient-level data for a region of interest are *not* available. Instead, given daily admissions counts, we model aggregated counts of observed resource use, such as the number of patients in the general ward, in the intensive care unit, or on a ventilator. In order to explain how individual patient trajectories produce these counts, we propose an aggregate count explicit-duration hidden Markov model, nicknamed the ACED-HMM, with an interpretable, compact parameterization. We develop an Approximate Bayesian Computation approach¹ that draws samples from the posterior distribution over the model’s transition and duration parameters given aggregate counts from a specific location, thus adapting the model to a region or individual hospital site of interest. Samples from this posterior can then be used to produce future forecasts of any counts of interest. Using data from the United States and the United Kingdom, we show our mechanistic approach provides competitive probabilistic forecasts for the future even as the dynamics of the pandemic shift. Furthermore, we show how our model provides insight about recovery probabilities or length of stay distributions, and we suggest its potential to answer challenging what-if questions about the societal value of possible interventions.

1. Introduction

The ongoing COVID-19 pandemic has imposed immense costs to human health, quality of life, and the economy. There remains a pressing need for forecasting models that can reliably predict demand for scarce hospital resources, such as general-ward beds, ICU beds, and ventilators. Many previous COVID-19 forecasts have failed (Ioannidis et al., 2020). To be successful, it is especially important that models are likely to extrapolate to the

1. Code URL: <https://www.github.com/tufts-ml/aced-hmm-hospitalized-patient-trajectory-model>

future rather than simply repeat the past, as the pandemic evolves through new waves, new treatments, new disease variants, and the introduction of vaccines. This need favors *mechanistic models* (Baker et al., 2018) that try to capture how patients actually move through the hospital, rather than “black-box” predictive models.

In addition to helping hospital administrators or regional authorities plan ahead, mechanistic forecasting models can also help assess the *societal value* of possible interventions. Throughout this pandemic, many interventions have been considered to improve outcomes for infected and hospitalized individuals, such as screening strategies, pharmaceuticals (Beigel et al., 2020; Chen et al., 2021), physical treatments (Koeckerling et al., 2020), and more (Zhang et al., 2020; Ahamad et al., 2021). Estimating the societal value of these interventions remains challenging, as they can impact many different stages of care.

Our motivating problem is to assess which interventions might avoid short-term future demand exceeding hospital capacity, as aggressive shut-downs harm other aspects of health, reduce quality of life, and restrict economic activity (Neumann et al., 2020). Policy makers, pharmaceutical companies, and funding agencies can use mechanistic forecasting models to identify interventions to prioritize; such models can also inform health technology assessment and hence questions of reimbursement. The mechanistic aspect of the model is crucial: a model can shed light on an intervention’s value only if it supports “*what-if*” queries about how an intervention would change patient trajectories. To make accurate decisions, it is also crucial that the model can be *adapted* to the hospital dynamics in the region of interest.

In this work, we develop and evaluate a *mechanistic* probabilistic model of a COVID-19 patient’s trajectory through stages of care in the hospital. This model can be used to make data-driven forecasts of the future daily census counts at different stages of care. That is, we can model the total count of patients in general ward, in the intensive care unit (ICU), or in the ICU on a ventilator. Importantly, the parameters driving this forecast are *interpretable* in the sense that they directly correspond to underlying mechanisms, such as transitions between stages of care and dwell-time distributions. Further, we show how our model supports asking what-if queries, such as “what if the average stay on the ventilator decreased by 2 days?”, and how it quantitatively answers such questions via downstream forecasts of daily demand for hospital beds or ventilators. These forecasts could be used at a regional or hospital level to assess if utilization will exceed crucial thresholds.

Desiderata. In addition to seeking a *mechanistic* model with interpretable parameters, another key design goal is *portability* to health systems across the globe. We wish to use the model to accurately assess the value of an intervention in a given specific hospital, region, or country. Healthcare system heterogeneity means that in different regions of interest, different data may be available due to logistical obstacles and legal concerns. To achieve portability, we assume *no* patient-level data are available (not even individualized length of stay information). Instead, our models rely only on aggregated daily count data.

Another goal is to explicitly capture our uncertainty given limited data by learning a *posterior distribution* over parameters rather than a point estimate. We take a Bayesian approach by expressing our prior beliefs, updating those beliefs given some observed data, and then making predictions and decisions by averaging over the remaining uncertainty (Gelman et al., 2013). Usually, fitting a posterior requires either optimization methods based on variational inference (VI) (Wainwright and Jordan, 2008), or sampling methods based

on Markov chain Monte Carlo (MCMC) (Andrieu et al., 2003). However, for our model and many like it, VI and MCMC have high derivation and implementation costs even for statistical experts. These costs make VI and MCMC difficult for those who need our models the most: data scientists inside healthcare organizations. Key discrete variables in our model (difficult to marginalize away) preclude the easy application of recent “automatic” VI or MCMC (Hoffman and Gelman, 2014; Kucukelbir et al., 2017). Thus, to achieve accessible and extensible posterior estimation, we use simulation methods via Approximate Bayesian Computation (ABC) (Marjoram et al., 2003; Marin et al., 2012; Cranmer et al., 2020). ABC algorithms are easy to implement given a simulation that samples from the model. Furthermore, ABC allows us to design a custom distance function to assess fit, which we tailor to our application. This would be difficult with an explicit likelihood in VI and MCMC.

Contributions. This study makes three key contributions.

- First, we develop a new model capable of forecasting the daily counts of patients at various stages of care (general ward, intensive care unit (ICU), on ventilator in the ICU, death) when given the daily admissions counts. The model parameters are fully specified by fewer than 20 numbers, each one interpretable by practicing clinicians because they correspond directly to how patients move through typical hospitals.
- Second, we develop routines for fitting posterior distributions over parameters given only aggregated daily count data from any region of interest. Building on Approximate Bayesian Computation (ABC) extensions of Metropolis-Hastings MCMC (Marjoram et al., 2003), we show how distance metrics and annealing schedules for acceptance thresholds can be carefully designed to generalize well to new target regions. We further show how to scale up to *tens of thousands* of hospital stays. Because our ABC approach is simpler to implement than other possibilities, other researchers can adopt our work or even extend our model without advanced expertise in Bayesian learning.
- Third, we *validate* our model and posterior estimation procedures on real hospital census counts from both the United States and the United Kingdom. We show that the model can provide competitive forecasts of daily counts for both a single hospital and for an entire geographic region, even when the “training” period shows rising numbers but the “testing” period exhibits falling counts due to underlying shifts in policy and population behavior. Beyond accuracy, we show that our approach can deliver *insight* into underlying mechanistic parameters and can be interrogated with “*what-if*” *scenarios* to understand the likely quantitative impact of possible interventions.

Generalizable Insights about Machine Learning in the Context of Healthcare

Many healthcare modeling problems are beset by two challenges. First, it is difficult to acquire large volumes of patient-level data (due to privacy concerns or logistical obstacles). However, relevant *aggregated* data with no patient-specific information are often publicly available. Second, it is usually easy to specify a mechanistic probabilistic model that can *simulate* individual patient trajectories and produce aggregated counts, but difficult to *learn* the posterior distribution over model parameters. Our work can be seen as a case study of lessons learned in how to apply an under-utilized methodology – Approximate Bayesian Computation (ABC) – to address both concerns in a healthcare context.

2. Related Work

Projecting the spread of COVID-19 and its impact on hospital utilization has been widely pursued. We briefly survey several threads of related work: regional modeling, local modeling (at either the hospital-level or patient-level), combinations of regional and local models, and efforts to assess societal value.

Regional modeling. Many studies have forecasted hospital utilization within a large geographic region (Murray et al., 2020; Jewell et al., 2020; Reiner et al., 2021). Some regional models use Susceptible-Infected-Recovered (SIR) models to predict regional case rates (Zou et al., 2020; Li et al., 2020), which can then inform bed usage and ventilator usage. Others (Murray et al., 2020) focus on predicting mortality rates given observed deaths, then use an internal simulation to estimate hospital utilization given fatalities.

A common limitation of regional models is their simplified and inflexible characterizations of the care pathways within hospitals. For example, CHIME² (Weissman et al., 2020) allows users to specify the proportions of patients requiring ICU care and mechanical ventilation, as well as the average length of stay in the hospital and in the ICU. But the model does not accommodate *distributions* for the “dwell times” at each level of care; nor does it allow for different durations among patients who recover and who deteriorate. The Imperial College model (Flaxman et al., 2020) connects deaths to key epidemiological parameters, but does not capture hospital resources. The published forecasts of the popular “IHME model” (Reiner et al., 2021), maintained by the Institute for Health Metrics and Evaluation (IHME) at the University of Washington, rely on rigid assumptions about how patients progress through the hospital (IHME COVID-19 Forecasting Team, 2020, Suppl. Sec. 8). For example, length-of-stay is not probabilistic but fixed for each class of patients (details in App. F). Users of these forecasts cannot easily modify these internal parameters or *learn* them for a location of interest; instead, the IHME itself must release updates. While the IHME model has at times performed well at its target task of accurate forecasting of mortality (Friedman et al., 2020), arguably it has been outperformed by other efforts (Gu, 2020).

Hospital-level modeling. At the hospital level, Lee et al. (2021) develop a non-mechanistic prediction model for univariate counts of COVID-19 patients at a specific hospital site based on auto-regressive statistical models. While this work is portable, we expect our mechanistic model to generalize better (as confirmed in later experiments).

Epstein and Dexter (2020) developed a mechanistic model for making hospital-specific short-term predictions of daily census counts. Given access to length-of-stay (LOS) information at that hospital from previous patients, they can forecast the length-of-stay of current patients at various stages of severity. The LOS data requirement limits portability compared to our approach.

Patient-level modeling. At the patient level, the multi-state survival model by Roimi et al. (2021) predicts how a patient moves day-by-day between clinical states (critical, severe, or moderate) using covariates such as age and sex. They use this model to forecast hospital demand in Israel. Our proposed model is simpler and does not depend on any patient-level data. Unlike our approach, it would be difficult to transfer Roimi et al. (2021)’s model across regions because acquiring the necessary data is challenging.

2. <https://penn-chime.phl.io>

Multi-level modeling. Qian et al. (2020) present a model for COVID-19 at the national, regional, hospital and individual levels, targeted at the United Kingdom. Their COVID-19 Capacity Planning and Analysis System (CPAS) makes hospital stay forecasts for current patients via patient-level covariates (e.g. age, sex, comorbidities). They simulate the hospital stays of future patients via a regional-level trend forecaster to predict admissions. While admirable, this kind of model is only possible in a nationalized health system with access to patient-level data. For many regions of interest, such data would not be available.

Efforts to assess societal value. We have identified three health technology assessments of COVID-19 interventions that include hospital utilization projections. Congly et al. (2020) evaluated the administration of either remdesivir or dexamethasone to various groups of patients hospitalized to treat their COVID 19 infections. Assumptions for the model came from the literature. Gandjour (2021) assessed the addition of ICU capacity in German hospitals, using assumptions derived from information released by the German Ministry of Health. Finally, Jo et al. (2020) examined use of remdesivir and dexamethasone in South Africa. The assessment relied on clinical trial results to project therapeutic impact on length of stay and clinical outcomes. But the analysis did not explicitly describe its modeling of patient progression in the hospital setting. We emphasize that none of these analyses performed any *learning or calibration* of models to ensure consistency with observed data from the region of interest. Thus, our study advances the state-of-the-art for assessing societal value by applying Approximate Bayesian Computation methods to ensure that models accurately explain data from the target site.

3. Model

We now describe our proposed probabilistic model for an individual COVID patient’s trajectory through the hospital. We call our model the Aggregate Counts Explicit Duration Hidden Markov Model, or “ACED-HMM” for short. While we specialize to modeling hospitalized patients with confirmed cases of COVID-19 here, the model could conceptually be used for any disease with appropriate changes to the stages of care considered as well as modified assumptions about transitions and durations.

Our proposed model represents the patient’s current health state and stage of care in the hospital over a period of T days, indexed by $t \in \{1, 2, \dots, T\}$. There are 5 possible *stages*, each indicated by a unique symbol. We assume each patient occupies exactly one stage on each day. Fig. 1 diagrams the possible pathways through these stages and associated transition probabilities. There are 3 stages in the *intermediate* set $\mathcal{S} = \{\mathbf{G}, \mathbf{I}, \mathbf{V}\}$: stage \mathbf{G} indicates the patient is the general ward, stage \mathbf{I} indicates the patient is in the intensive care unit (but not yet on mechanical ventilation), and stage \mathbf{V} indicates the patient is actively receiving mechanical ventilation in intensive care. There are additionally 2 possible *absorbing or terminal* stages: “terminal death” \mathbf{T} or “fully recovered” \mathbf{R} .

For each patient, indexed by n , our model generates a sequence of care segments indexed by ℓ , with total length L_n (not all patients will experience the same number of segments). Each segment is defined by a *stage indicator* $s_{n\ell} \in \{\mathbf{G}, \mathbf{I}, \mathbf{V}, \mathbf{R}, \mathbf{T}\}$, a *health indicator* $h_{n\ell} \in \{0, 1\}$ corresponding to a “declining” or “recovering” trajectory within that segment, and a *duration* integer $\Delta_{n\ell} \in \{1, 2, \dots, D\}$ representing the number of *days* spent in that care segment. Our model generates patient n ’s sequence of stage indicators $s_{n,1:L_n}$,

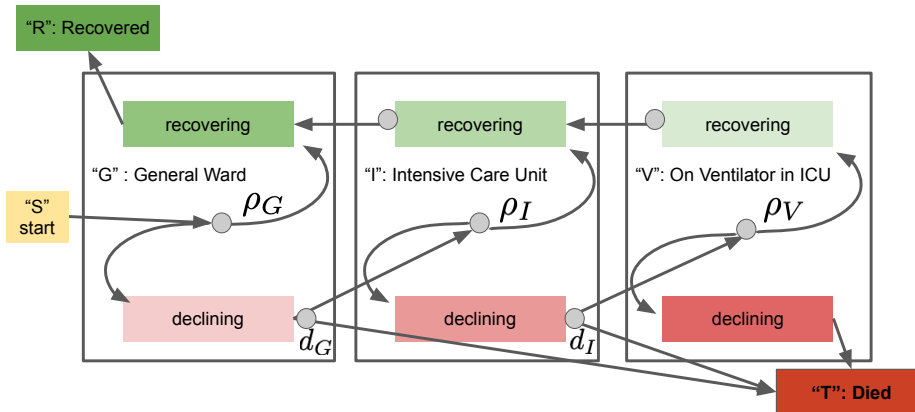


Figure 1: Diagram of the proposed Markovian model for a COVID-19 patient’s trajectory through the hospital. Rectangles indicate possible states, defined by the patient’s current stage of care s (general ward \mathbf{G} , off ventilator in ICU \mathbf{I} , on ventilator in ICU \mathbf{V}) as well as health trajectory h (recovering or declining). Small circles indicate stochastic binary choices, with one path marked by transition probability. Recovery transition parameters ρ_k give the chance of switching to a recovery path, while death probabilities d_k mark the chance of dying early.

health indicators $h_{n,1:L_n}$, and durations $\Delta_{n,1:L_n}$ via the generative process defined below. For simplicity, we drop the patient index n when it is clear we are talking about one patient.

Our model assumes that for each patient, we are given an integer $a_n \in \{1, 2, \dots, T\}$ indicating the day of *admission* to the hospital. Throughout this paper, we assume all patients begin in the general ward on their first day. However, if desired it is easy to alternatively allow some incoming patients to directly begin in the ICU (\mathbf{I}) or on the ventilator (\mathbf{V}), which might be needed to model transfers to a large facility with specialized care.

3.1. The ACED-HMM Generative Model

Initialization. We begin in the general ward, $s_1 \leftarrow \mathbf{G}$, and sample the health indicator $h_1 \sim \text{Bern}(\rho_{\mathbf{G}})$, where parameter $\rho_{\mathbf{G}}$ gives the probability of recovery in the general ward.

Generating stage indicators. We generate the *stage indicator* $s_\ell \in \{\mathbf{G}, \mathbf{I}, \mathbf{V}, \mathbf{R}, \mathbf{T}\}$ for each subsequent segment $\ell \in 2, 3, \dots$ given the previous health indicator $h_{\ell-1}$ via:

$$s_\ell \leftarrow \begin{cases} \mathbf{T} & \text{w. prob. } d_{s_{\ell-1}} \text{ if } h_{\ell-1} = 0 \\ \text{next_stage}(s_{\ell-1}, h_{\ell-1}) & \text{otherwise} \end{cases} \quad (1)$$

where parameter d_s defines the chance of early death for each intermediate stage $s \in \mathcal{S}$. The deterministic `next_stage` procedure always advances in order $[\mathbf{G}, \mathbf{I}, \mathbf{V}, \mathbf{T}]$ if health is declining ($h_{\ell-1} = 0$), and advances in order $[\mathbf{V}, \mathbf{I}, \mathbf{G}, \mathbf{R}]$ if health is recovering ($h_{\ell-1} = 1$). See Fig. 1 for visual illustration of these stage pathways. Once we reach either absorbing stage (\mathbf{R} or \mathbf{T}), we halt all further generation (this event defines the sequence length L_n).

Generating health indicators. We generate the health indicator $h_\ell \in \{0, 1\}$ for each subsequent segment $\ell \in 2, 3, \dots$ given the previous health $h_{\ell-1}$ and the new stage s_ℓ via:

$$h_\ell \leftarrow \begin{cases} 1 & \text{if } h_{\ell-1} = 1 \\ u & \text{if } h_{\ell-1} = 0, \quad u \sim \text{Bern}(\rho_{s_\ell}), \end{cases} \quad \ell \in \{2, 3, \dots\}. \quad (2)$$

After each declining segment ($h = 0$), we switch to recovery with a stage-specific recovery probability $\rho_s \in [0, 1]$ for each intermediate stage $s \in \mathcal{S}$, but otherwise keep declining. Once on a path to recovery ($h = 1$), we assume the patient always continues to recover.

Assuming recovery continues unabated is a simplifying assumption. The chance of readmission to the ICU or reintubation on the ventilator is non-zero and has evolved with care practices substantially over the course of the pandemic. Among all hospitalized patients with COVID-19, only a fraction progress to the ICU and later recover to the general ward. Among these initially recovered patients, our clinical experts report that for the focus period of our later case studies (Nov. 2020 - Feb. 2021) only a small fraction – roughly 10% – would be later sent back to the ICU. Because the overall fraction of patients on pathway to recovery who then decline is small, we suggest that our model, while not capturing all possible pathways, still captures most while benefiting from simplicity in estimation and implementation. Allowing recovering patients to later decline is a future research opportunity and would require additional parameters that are currently unavailable. Our model and ABC learning could handle this extension: we would simply add a probability of reversal to our Markovian model simulation with an appropriate prior, but would also then need to understand if progression or recovery then differs.

Generating durations. We generate the duration Δ_ℓ (in days) for segment ℓ as:

$$\Delta_\ell | s_\ell, h_\ell \sim \text{Cat}(\pi_1^{s_\ell, h_\ell}, \dots, \pi_D^{s_\ell, h_\ell}). \quad (3)$$

where probability vector $\pi^{s,h}$ defines a categorical probability mass function over the first D integers $\{1, 2, \dots, D\}$ representing possible durations spent at intermediate stage $s \in \mathcal{S}$ while in health state $h \in \{0, 1\}$. Notably, this model flexibly allows recovering ($h = 1$) and declining ($h = 0$) patients to have different duration distributions at each stage.

For tractability we define a maximum number of days D allowed in any single segment. Since there are at most 5 possible segments before a patient reaches a terminal state (see Fig. 1), the maximum hospital stay is $5D$ days. In all experiments, we set $D = 22$. Designating this value as the maximum duration of any segment of care reflects COVID-19 length-of-stay statistics as of September 2020 from the U.S. CDC (CDC, 2020, Table 2), which only reports 25th, 50th, and 75th percentiles. For patients admitted to the ICU, the 75th percentile estimated for *total hospital length of stay* ranges from 20 days to 25 days across age groups (corresponding to at least 2 segments of our model: general-ward plus ICU). For patients not admitted to the ICU, stays are typically much shorter. By setting $D = 22$ in our analyses, the total length of stay could reach up to 66 days for patients admitted to the ICU at any time during hospitalization, and up to 110 days for general ward patients who eventually require mechanical ventilation. Truncation to at most 22 days per segment should thus affect few real patients.

At the suggestion of reviewers, we investigated higher truncation levels (full details in App. B.9). While results were almost all indistinguishable, we did find that using a larger duration ($D = 44$) produced slightly better forecasts for ventilator counts. Because the computational cost of a larger truncation limit is negligible (except for sampling each duration, runtime and storage scales with the number of segments not D), we recommend that future analyses use $D = 44$ if possible. Our public ACED-HMM implementation allows the user to set arbitrary upper bounds for each duration distribution via a plain-text configuration file.

Tractable parameterization of durations. To define each duration distribution, we could learn a separate D -dimensional probability vector $\pi^{s,h}$ for each state (indexed by stage s and health h). We instead explore a simpler formulation with only two scalar parameters per state: mode location $\lambda^{s,h} > 0$ and temperature $\nu^{s,h} > 0$:

$$[\pi_1^{s,h}, \dots, \pi_D^{s,h}] \leftarrow \text{softmax} \left(\frac{\log \text{PoiPMF}(1|\lambda^{s,h})}{\nu^{s,h}}, \dots, \frac{\log \text{PoiPMF}(D|\lambda^{s,h})}{\nu^{s,h}} \right). \quad (4)$$

where PoiPMF denotes the probability mass function of the Poisson distribution. Fig. 2 shows how possible distributions over D days depend on these two parameters. Intuitively, $\lambda^{s,h}$ controls the location of the most likely duration (mode). Large temperatures $\nu^{s,h} \rightarrow \infty$ makes the distribution $\pi^{s,h}$ more uniform, while as $\nu^{s,h} \rightarrow 0$ the distribution becomes peaked at the mode. This parameterization encodes two valuable inductive biases: similar stay-lengths have similar probabilities and the overall distribution is either unimodal or flat. This 2-parameter formulation has been previously used for supervised modeling of ordinal outcomes (Beckham and Pal, 2017), but not to our knowledge for durations in a Markov model. We compare our formulation to other possibilities (truncated Poisson) in App. B.9.

Joint Probability Distribution. Each patient’s entire hospital stay can be captured by a sequence of stage indicators, health indicators, and durations, with joint probability:

$$p(s_{1:L_n}, h_{1:L_n}, \Delta_{1:L_n}) = p_\tau(s_1, h_1) \prod_{\ell=2}^{L_n} p_\tau(s_\ell, h_\ell | s_{\ell-1}, h_{\ell-1}) \cdot \prod_{\ell=1}^{L_n} p_\xi(\Delta_\ell | s_\ell, h_\ell) \quad (5)$$

Here we’ve gathered all transition probabilities into τ , such that $\tau = \{d, \rho\}$ and all duration parameters are denoted $\xi = \{\lambda, \nu\}$ for simplicity. Notably, this model structure can be mapped to the semi-Markov model known as an explicit-duration (hidden) Markov model (Ferguson, 1980; Mitchell and Jamieson, 1993; Yu, 2010), by mapping the health and stage indicator variable sequences $h_{1:L_n}, s_{1:L_n}$ to a discrete state sequence $z_{1:L_n}$ with 9 possible states: the start state **S**, the two absorbing states **T**, **R**, and 6 intermediate states **G0**, **G1**, **I0**, **I1**, **V0**, **V1** where the 0 (1) suffix in the state name represents the declining (recovering) variant of each stage, respectively. The transition probability parameters for this HMM are specified in Supplementary Table C.1. Throughout the rest of the paper, we use the compact notation $\mathbf{z}_n = z_{n,1:L_n}$ below to represent the state sequence for patient n , which can be easily mapped back to the stage sequence $s_{n,1:L_n}$ and health sequence $h_{n,1:L_n}$.

3.2. Prior beliefs about transition and duration parameters

Priors on transition parameters. We can use published statistics to inform our prior beliefs about transition probability parameters $\tau = \{\rho, d\}$. Published probability statistics from the U.S. Centers for Disease Control (CDC) (see Table 2 of CDC (2020)) indicate the fraction of hospitalized COVID-19 patients who are admitted to ICU, placed on the ventilator, and eventually die, for each of three age groups (18-49, 50-64, and 65+). We translate this information into Beta priors for each ρ_k , so that each prior distribution has a mean matching these statistics while covering the overall range across age groups with non-trivial mass. See Appendix C.2 for details.

We next set the early death probabilities. Based on conversations with clinical experts, we expect the chance of early death in the general ward is very low ($\mathbb{E}[d_{\mathbf{G}}] \approx 1\%$) and the chance of early death in the ICU is only slightly larger ($\mathbb{E}[d_{\mathbf{I}}] \approx 2\%$). We again set




Figure 2: Visualization of possible duration distributions under our tractable parameterization (Eq. 4). Parameter $\lambda > 0$ (increasing across columns) controls the location of the mode, while temperature $\nu > 0$ (increasing top-to-bottom across rows) controls entropy around the mode.

Beta priors to match these means with little uncertainty around them (such that the 95th percentile is below 0.05) to reflect our confidence that early death is possible but rare.

Priors on duration parameters. To determine priors over the durations spent in each state (indexed by s, h), we draw relevant parameters independently from a common prior:

$$\lambda^{s,h} \sim \text{TruncNormal}(\mu^{s,h}, (\sigma^{s,h})^2, \text{lower} = 0, \text{higher} = D), \quad \log_{10} \nu^{s,h} \sim \mathcal{N}(0.5, 0.5^2). \quad (6)$$

Visuals of our prior’s imposed distribution on durations in each state can be seen in Fig. 5.

Our truncated normal prior for the mode location $\lambda^{s,h}$ of each state reflects that these quantities are difficult to know without reliable data about previous patient experience in the region of interest. We select a mean of $\mu^{s,h} = 8$ days with a standard deviation $\sigma^{s,h} = 3$. Under this Normal prior, the mode of each segment’s duration distribution could plausibly range from 1 to 17 days (99% of the probability mass lies within the interval $\mu^{s,h} \pm 3\sigma^{s,h}$).

Our prior for each state’s temperature $\nu^{s,h}$ assumes the \log of this value will be normally distributed. By using a prior for $\log_{10} \nu^{s,h}$ with mean 0.5 and standard deviation 0.5, we allow plausible temperatures to span both high and low entropy distributions (see Fig. 2).

4. Estimating the Posterior Distribution over Parameters using ABC

In order to train our model so its parameters reasonably explain hospital trajectories in a specific region or site of interest, we assume we have access to observed daily counts $\mathbf{y}_{1:T} = \{y_{1:T}^k\}_{k \in \mathcal{K}}$ during the training period of interest. Here the set of observed stages \mathcal{K} could include any of the intermediate or terminal stages; typically $\mathcal{K} = \{\mathbf{G}, \mathbf{I}, \mathbf{V}, \mathbf{T}\}$. We also assume knowledge of the admission counts $a_{1:T}$ for the training period, which represent a total of N patients who enter the hospital system, with $N = \sum_{t=1}^T a_t$.

Given observed counts $\mathbf{y}_{1:T}, a_{1:T}$, we wish to be able to draw samples from the posterior $p(\tau, \xi, \mathbf{z}_{1:N}, \mathbf{\Delta}_{1:N} | \mathbf{y}_{1:T}, a_{1:T})$. Samples from this distribution allow us to properly understand the range of possible values for the model’s transition parameters τ and duration parameters $\xi = \{\lambda, \nu\}$. Samples of parameters can further allow us to produce *forecasts*, as we can draw samples of counts $\mathbf{y}_{T+1:T+F}$ for a “testing” period of F days $-T+1, \dots, T+F$ – by taking some assumed admissions $a_{T+1:T+F}$ for this period and sampling forward from the ACED-HMM.

For our model, it is easy to *simulate* patient care sequences and to produce aggregate counts from these sequences (simply counting the total number of patients simulated in each stage at each day). However, it is cumbersome to define an explicit probability mass function for both the individual patient care sequences $\mathbf{z}_n, \mathbf{\Delta}_n$ and the observed aggregate counts $\mathbf{y}_{1:T}$. This would only become more difficult if we explore more complex extensions of our model (e.g. add age-dependent transition or duration probabilities). We would like to simply define a *simulation* that produces samples of aggregate counts, and then estimate a *posterior distribution* over the parameters of this simulation that fits observed counts.

4.1. Approximate Bayesian Computation

To fit posterior distributions, we turn to Approximate Bayesian Computation (ABC) (Marjoram et al., 2003; Marin et al., 2012; Cranmer et al., 2020), a method for fitting models that are naturally expressed as simulations. ABC has found applications in ecology (Beaumont, 2010), population genetics (Beaumont et al., 2002), and epidemiology (Blum and Tran, 2010). To use standard Markov chain Monte Carlo methods in our application would require an explicit ability to compute the likelihood of observed counts arising from our model’s simulations. ABC offers a simpler approach: we need only the ability to simulate counts and to decide if simulated counts are “close enough” to observed data.

Our approach is based on the ABC MCMC algorithm presented by Marjoram et al. (2003). We iteratively propose new candidate values for parameters τ, ξ , and then decide to accept or reject these based on an acceptance criterion designed to ensure the overall sequence of sampled parameters converges to the target posterior distribution. This looks superficially similar to conventional Metropolis-Hastings MCMC (Metropolis et al., 1953; Andrieu et al., 2003). However, the test for acceptance uses an evaluation of a distance function d between observed data $\mathbf{y}_{1:T}$ and our simulation’s sampled counts $\tilde{\mathbf{y}}_{1:T}$. This distance threshold check replaces the evaluation of an explicit likelihood (which would require specification of a probability mass function and the ability to evaluate it efficiently).

An advantage of our ABC approach is that we do not need to compute or even define an explicit probability distribution for the likelihood $p(\mathbf{y}_{1:T} | \mathbf{z}_{1:N}, \mathbf{\Delta}_{1:N})$. This allows us to avoid a potentially detrimental *misspecified* parametric form chosen for computational convenience. Another advantage of ABC is that we can customize the distance function to meet application-specific needs; specifying a likelihood that achieves the same goals and is also a valid probability mass function might be more challenging. For example, our chosen distance below emphasizes fitting the most recent daily counts in the training period (so our forecasts will generalize to the future well) and fitting later stages of the model well (since those stages are harder to estimate).

Distance computation. Let \mathcal{K} be the set of observed stages, and K the number of stages ($K = |\mathcal{K}|$). Let y_t^k denote the observed count at time t in stage k , and \tilde{y}_t^k denote the

simulation’s sampled count. We define the distance as a weighted mean absolute error:

$$d(\mathbf{y}_{1:T}, \tilde{\mathbf{y}}_{1:T}) = \frac{1}{K \cdot T} \sum_{t=1}^T \sum_{k=1}^K w_{tk} \frac{|y_t^k - \tilde{y}_t^k|}{\max(y_t^k, \tilde{y}_t^k)} \quad (7)$$

where w_{tk} is a scalar weight that determines the relative “importance” of matching the counts at time t and stage k . We define w_{tk} as a product of a time-specific weight v_t and a stage-specific weight u_k . For timestep weight v_t , we apply a recency bias: we linearly interpolate between $v_1 = 0.5$ at day 1 and $v_T = 1.5$ at day T , so that our model is better at making forecasts after day T (and the average weight is 1). We use stage-specific weights u_k because preliminary experiments suggested the *later* stages of the model are harder to fit (while all N patients enter the general ward, only a smaller fraction will reach the ICU or the ventilator). Thus, we set u_k values so later stages are worth more. See App. D for further details and justification of our chosen distance, and App. B.8 for ablation studies.

Proposal distributions q . Our ABC algorithm requires a procedure to draw a new plausible parameter value given a current value. For each type of parameter (transition probabilities, duration modes, duration temperatures), we define a standard “random walk” proposal distribution q with the mean at the previous value and a variance that encourages modest exploration. See Appendix D for parameter-specific details.

ABC algorithm. Initial values for the transition parameters τ and duration parameters ξ are samples from our chosen priors $p(\tau)$ and $p(\xi)$ defined above. Until convergence, the algorithm visits each individual parameter in turn, either a specific transition probability parameter ρ_s or d_s within τ , or a state-specific duration mode $\lambda^{s,h}$ or temperature $\nu^{s,h}$ within ξ . At each visit, the algorithm draws a candidate value from the proposal distribution q defined above. Then, we decide to “accept” or “reject” the candidate. Denote the current parameters as τ, ξ , and new candidate parameters (with one entry updated) as τ^*, ξ^* .

There are two stages of acceptance. First, counts are simulated from our ACED-HMM progression model for the entire training period using the proposed parameters τ^*, ξ^* , generating $\tilde{\mathbf{y}}_{1:T}$. Then, we compute the distance d between the samples $\tilde{\mathbf{y}}_{1:T}$ and true counts $\mathbf{y}_{1:T}$. If the distance d is below a chosen tolerance threshold ε , the algorithm moves to the second stage of acceptance; otherwise the proposal is rejected. In stage two, the proposal is accepted if a random uniform draw between 0 and 1 is less than the acceptance ratio $\alpha = \frac{p(\tau^*, \xi^*)q(\tau^*, \xi^* \rightarrow \tau, \xi)}{p(\tau, \xi)q(\tau, \xi \rightarrow \tau^*, \xi^*)}$. The first stage of acceptance checks the fit to the data (replacing the likelihood used in conventional MCMC), whereas the second stage further filters the acceptance via information from the prior and from the proposal, ensuring the Markov chain will converge to the target posterior distribution (Marjoram et al., 2003).

Scheduling the distance tolerance threshold. Because random initial parameters are unlikely to explain observed data well, we find it useful to start with a lenient value of the data-agreement threshold ε , i.e. one that accepts all proposals. We then decay ε across iterations toward a smaller value that better ensures fit to observed data. In App. D, we describe two phases for our ABC algorithm: a *burn-in* phase and a *sampling* phase. During *burn-in*, we gradually decay the value of ε over many iterations to reach a good explanation of the data. Then, during *sampling* we hold ε fixed and collect parameter samples, which are treated as representative of the target posterior distribution. During burn-in, we find it

helpful to periodically *increase* ε sharply and then decay gradually to escape local optima. See Fig. D.1 for a visual representation of escaping a local optima using this idea.

Ensembling. To obtain robust estimates of the posterior, we combine samples from multiple independent runs of our ABC algorithm into an ensemble. For any given dataset, we expect there might be several different parameter configurations that explain the training counts well (due to the flexibility of the model and possible identifiability issues). Ensembling allows us to capture more of these configurations than a single run of ABC might recover. By aggregating, we gain statistical strength in the forecasts and a deeper understanding of the uncertainty in learned parameter values.

Synthetic validation. We conducted several experiments on synthetic “toy” datasets (generated by our model) to verify that our ABC procedures can (approximately) recover plausible posterior distributions. See Appendix B.7 for a thorough description.

5. Data for Case Studies

We fit our proposed model to pursue two case study applications over a range of national contexts and geographic scales. First, as a primary use case we pursued *regional* forecasts for several states in the United States. State-level modeling allows us to make forecasts that interest policy makers and public health officials within that state. We also show how such regional forecasts might be used to assess societal value. Second, we performed *site-level* forecasts for specific hospitals in the United Kingdom (UK). These experiments allow us to assess potential utility to hospital administrators for planning purposes.

Below, we describe the data acquired for training and evaluating our forecasting model in both cases. Each dataset is assumed to provide census counts for the training period, denoted as days $1, 2, \dots, T$, for some subset or additive combination of the stages of our model (general ward \mathbf{G} , off-ventilator in ICU \mathbf{I} , on-ventilator in ICU \mathbf{V} , terminal death \mathbf{T} , recovered \mathbf{R}). To evaluate forecasts against a “ground truth”, we further require the same counts for the future period of interest, starting on day $T + 1$ and ending on $T + F$.

In order to assess our model, during both training and forecasting we assume access to true daily *admissions* counts to the general ward. For a real forecast after day T , these counts will not be known. We leave this complex scenario to future work. We could use hand-designed projections representing plausible future admission scenarios from clinical experts, or develop a forecasting model to project admissions (as in Qian et al. (2020)).

5.1. U.S. state hospital data for regional forecasts

We obtain hospital utilization counts for four US states: Massachusetts (MA), Utah (UT), South Dakota (SD), and California (CA). We use counts from November 11th, 2020 to January 11th, 2021 as training data, and forecast one month into the future. We selected these states to represent diverse geographic regions as well as diverse hospital utilization dynamics and policy responses to the pandemic during the selected period. Hospitalizations in SD peaked in late November, in MA in early January, and in UT in late January.

For all states, we use data from the COVID Tracking Project ([The Atlantic Monthly Group, 2021](#))³ and the [U. S. Dept. of Health and Human Services \(2021\)](#)⁴. For each state we have access to a daily time series of General Ward (\mathbf{G}) counts, counts in the ICU off the ventilator (\mathbf{I}), counts on the ventilator (\mathbf{V}) and terminal death counts (\mathbf{T}). For Utah and California in particular, we have access only to aggregate ICU counts ($\mathbf{I} + \mathbf{V}$) that do not distinguish between on and off the ventilator. For all states, we found that *smoothing* the terminal counts produced more sensible data (in the raw data, some weekends and holidays report implausibly low counts for deaths that reflect imperfect reporting processes). See Appendix A for data collection and preprocessing details.

5.2. United Kingdom hospital data for site-level forecasts

We obtain daily occupancy counts at different care stages from two U.K. hospitals: South Tees and Oxford University, made available by the UK National Health Service⁵. We use counts from November 3rd to January 3rd as training data, and forecast one month into the future (counts past Feb. 3rd were unavailable when we ran the experiments). We selected these two hospitals because they represented large daily volumes with fewer data irregularities than other alternatives. After data collection (see details in Appendix A), for each site we obtained daily counts for total beds used ($\mathbf{G} + \mathbf{I} + \mathbf{V}$), ventilators used (\mathbf{V}), and recovered patients discharged from the hospital (stage \mathbf{R} of our model).

6. Results of Case Studies

We now describe several experimental results that assess our model’s ability to forecast accurately as well as its scalability, sensitivity to available data, and ability to answer what-if questions.

6.1. Forecast assessment: Can this model capture qualitative trends?

We show our method’s forecasted distribution over daily counts for Massachusetts in Fig. 3. Importantly, our ACED-HMM mechanistic model (blue lines), trained during a period of *rising* counts, can reasonably forecast the *downward* trend in hospital counts seen in MA after mid-January 2021. The figure also shows the published forecasts for MA from IHME ([Reiner et al., 2021](#)) made on Jan. 15th. The IHME model is primarily trained to match death counts from the training period and to our knowledge is not informed by state-specific hospital census counts. Consequently, we see large systematic overestimates. We do stress that our model incorporates actual admissions data (which IHME does not apparently use), so this is not quite a fair comparison (see App. F for thorough discussion), but it is still a compelling demonstration of benefits of our mechanistic model.

We further examine a forecast for a specific hospital site in the UK (“South Tees”) in Fig. 4. Here, we see how our method can capture *rising* trends during the evaluation period for both overall occupied beds (the sum of several stages: $\mathbf{G} + \mathbf{I} + \mathbf{V}$) as well as the number

3. <https://covidtracking.com/data>

4. <https://healthdata.gov/Hospital/COVID-19-Reported-Patient-Impact-and-Hospital-Capa/g62h-syeh>

5. <https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/>

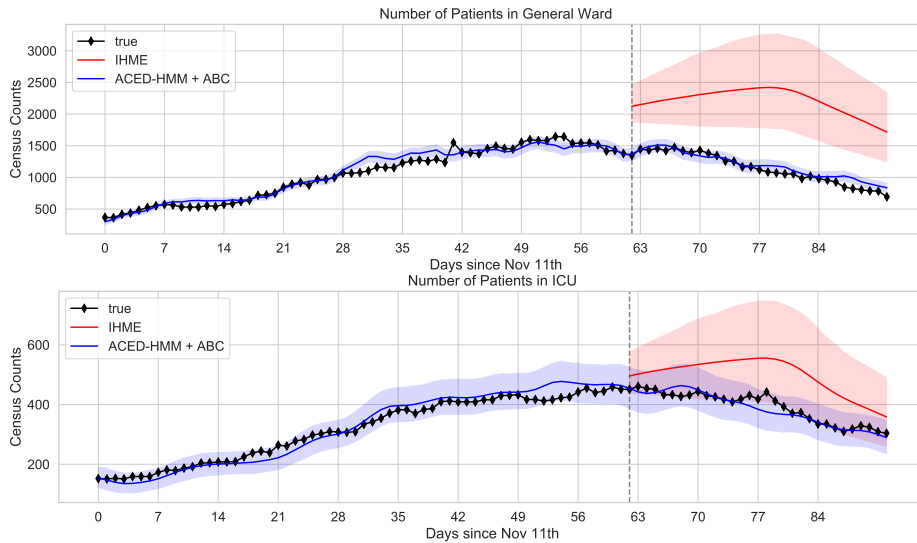


Figure 3: **Fit and forecasts on General Ward (G) and Total ICU (I + V) counts for Massachusetts.** Using ABC, we estimate a posterior over model parameters from 2 months of training counts (Nov. 11 2020 - Jan. 11, 2021, left of the dashed line). We stress that our ABC method is provided true daily admission counts for all days (0 - 92). We also include forecasts from IHME released on Jan. 15 2021. See App. F for details about comparing IHME’s forecasts to ours. For our method, shaded intervals show the 2.5th and 97.5th percentiles of 2000 posterior samples. For IHME, shaded intervals show their published “upper” and “lower” forecasts.

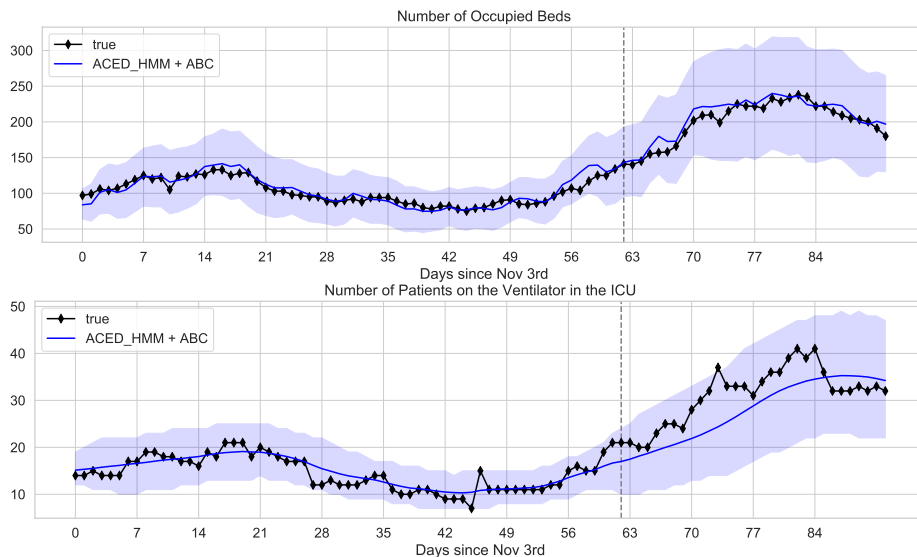


Figure 4: **Fit and forecasts on Occupied Beds (G + I + V) and Ventilator (V) counts for South Tees hospital in the UK.** Parameters were selected via ABC to fit 2 months of counts (Nov. 3 '20 - Jan. 3, '21, left of the dashed line). Shaded intervals show the 2.5th and 97.5th percentiles of 2000 posterior samples.

of ventilators in use (stage **V**), despite the fact that during the training period the counts were mostly flat and had only recently begun to rise.

	Method	G	InGeneralWard	I + V	InICU	V	OnVentInICU	T	Death (smooth)
		MAE	lower - upper	MAE	lower - upper	MAE	lower - upper	MAE	lower - upper
MA	Median of train set	211.2	NA	73.5	NA	80.1	NA	17.0	NA
	Bayes LR (day)	931.3	927.2 - 935.3	187.2	186.5 - 188.0	103.0	102.7 - 103.3	26.9	26.7 - 27.2
	Bayes LR (day + adm.)	277.8	275.6 - 280.5	24.3	23.8 - 24.8	24.3	23.9 - 24.8	8.0	7.8 - 8.2
	ACED-HMM + ABC	65.0	61.5 - 68.9	15.7	14.5 - 16.9	34.0	32.6 - 35.7	8.0	7.8 - 8.3
	IHME	1066.1	*543.4 - 1754.9	104.0	*15.6 - 257.9	25.7	*41.4 - 106.5	21.6	*8.7 - 49.7
	AR-Poisson	498.0	438.3 - 553.6	162.2	148.6 - 173.5	128.3	118.5 - 138.6	34.8	28.4 - 42.0
	ACED-HMM + Prior	868.6	779.6 - 945.8	293.6	272.6 - 313.3	119.6	113.8 - 124.8	40.7	39.1 - 42.2
	Mean test-set y	1141.6		392.5		249.1		66.5	
SD	Median of train set	234.0	NA	55.1	NA	23.9	NA	8.2	NA
	Bayes LR (day)	55.0	54.3 - 55.8	5.8	5.5 - 6.1	7.0	6.7 - 7.2	2.6	2.4 - 2.8
	Bayes LR (day + adm.)	53.1	52.6 - 53.9	4.3	4.2 - 4.5	5.1	4.9 - 5.3	2.8	2.7 - 2.9
	ACED-HMM + ABC	9.2	8.5 - 10.1	4.5	4.3 - 4.8	4.8	4.5 - 5.1	2.8	2.7 - 2.9
	IHME	71.8	*6.6 - 166.4	17.7	*4.2 - 46.0	5.9	*6.3 - 20.6	2.5	*3.6 - 5.0
	Mean test-set y	102.5		34.9		23.6		7.8	
	Mean test-set y	102.5		34.9		23.6		7.8	
UT	Median of train set	75.9	NA	42.1	NA	NA	NA	2.7	NA
	Bayes LR (day)	68.7	67.9 - 69.5	33.6	33.1 - 34.2	NA	NA	2.7	2.7 - 2.8
	Bayes LR (day + adm.)	63.4	62.8 - 63.9	32.6	32.1 - 33.1	NA	NA	2.6	2.5 - 2.7
	ACED-HMM + ABC	20.8	20.2 - 21.3	18.2	17.4 - 19.2	NA	NA	2.4	2.3 - 2.5
	IHME	332.9	*142.6 - 578.4	105.0	*35.7 - 213.9	NA	NA	7.4	*2.7 - 15.7
	Mean test-set y	272.9		164.3		NA		11.7	

Table 1: **Quantitative error assessment on 3 U.S. states.** We measure error with respect to the observed counts during the test period (Jan. 12 - Feb. 11, 2021) at each possible care stage: general ward **G**, in ICU (including on and off ventilator) **I+V**, on ventilator in ICU **V**, and *smoothed* death counts **T**). For all states and stages, we report the mean count y across the testing period to give a sense of scale. Cells marked NA were not available for that state (see discussion in App. A). AR-Poisson baseline is described in App. E, IHME baseline in App. F. We communicate uncertainty by reporting lower and upper estimates of the MAE. To obtain these intervals for our ACED-HMM and other probabilistic baselines, we repeat all MAE computations across 100 separate batches of posterior samples and then report the 2.5th and 97.5th percentiles across these batches. For the IHME baseline, a distribution of forecasted samples is not published. Instead, we evaluate and report the MAE of the provided “lower” and “upper” forecasts. This second kind of interval (marked with *) should not be directly compared to the first, as they capture different aspects of uncertainty.

Together, these qualitative results indicate that our model’s learned transition and duration distributions can successfully *translate* several weeks into the future in some cases. Again, we stress that our model incorporates the true admission counts for the testing period. However, the accurate ventilator demand prediction shown in Fig. 3-4 suggests that our ACED-HMM’s transition and duration modeling is reasonably robust.

6.2. Quantitative assessment: How accurate are forecasts?

We quantitatively evaluate our forecasts for 3 US states and 2 UK hospitals. For all U.S. states, we compare to forecasts from the IHME model (Reiner et al., 2021). We further compare to non-mechanistic Bayesian linear regression models that treat observed counts in each stage as linear functions of the day index or the day plus the last 3 weeks of admissions counts. Finally, we tried a non-mechanistic autoregressive model with Poisson likelihood designed for single hospital forecasting (Lee et al., 2021), and an ACED-HMM using our prior (not a posterior fit to the data). Because these last two models performed poorly on two sites, we did not apply them to others.

First, we calculate the Mean Absolute Error (MAE) between each method’s mean prediction and the observed counts during the testing period. See MAE results for the US in Table 1. For Massachusetts, our ACED-HMM achieves a noticeably better MAE of 65 for general ward (stage **G**), while all other methods have MAE values above 211. In addition to

this relative gain, we find that on a typical day, our prediction is within 4-10% of the actual general ward count (on average during the test period MA hospitals had 1141.6 patients in the general ward (range 694-1474)). Thus, our general ward predictions are sufficiently accurate to be operationally useful for planning. Similar MAE improvements for our method can be found for ICU beds in MA; for ventilator and death counts our method is competitive but not superior to others. For the states of SD and UT, each much smaller than MA, our mean absolute error (MAE) across all stages remained below 21, while the MAE values for other methods are considerably worse (e.g. the IHME baseline exceeds 100 MAE for the ICU in UT and exceeds 71 MAE for the general ward in SD). These results illustrate that our method can produce accurate forecasts across several diverse scenarios. Further results for the UK showing similar advantages for our ACED-HMM can be found in App. B.1.

Second, we assess the *coverage* of the learned posterior distribution across the 50%, 80% and 95% centered intervals of the posterior. For an ideal, well-calibrated posterior, we should observe that X% of the actual counts in the testing period fall within an interval chosen to cover X% of the distribution’s mass. See results for the US and the UK in Appendix B.3. We find the coverage of all methods could be improved substantially. For example for MA the interval our ACED-HMM predicts will cover 95% of data for stages **G** and **I** actually covers only 68% and 65% of observed counts. Improving coverage is challenging and an important area for future work.

6.3. Interpretation of learned posteriors over parameters

In Fig. 5, we show the learned posterior distributions over *all* model parameters after fitting to the MA data, revealing several insights about hospital dynamics during the training period. For instance, the learned recovery probability in **I** is lower than that indicated by our CDC-informed prior, indicating a potential deviation in hospital dynamics or in the patient age distribution in MA compared to the US as a whole. The learned posterior indicates, with high confidence, that the average length of stay in **I** (ICU *not* on the ventilator) is very low, both for recovering ($h = 1$) and especially for declining ($h = 0$) patients, indicating that a patient who is on a declining path in the ICU is very quickly moved to the ventilator. Last, patients on the ventilator seem to stay there for shorter periods of time when *declining* rather than when *recovering*. However, we note that the duration posterior for recovering on the ventilator closely matches the prior and is bimodal, which may suggest some chains are producing poor posterior estimates. This finding is not surprising: only a few patients recover from the ventilator, and thus with little data it is a difficult parameter to recover.

In Fig. B.1 in the Appendix, we further illustrate the learned posteriors for the South Tees hospital site in the UK. As expected, we observe much higher variance in these learned posteriors for a single hospital than the posteriors for the (much larger) Massachusetts dataset above. However, we notice some potentially valuable differences compared to MA. Duration distributions appear to be longer, and patients appear more likely to recover both in the General Ward and in the ICU.

6.4. Data granularity ablation: Are results sensitive to missing some stages?

Public data sources may differ in terms of the granularity of available counts. For example, for U.S. states we can usually gather separate counts of for each of the **G**, **I**, and **V** interme-

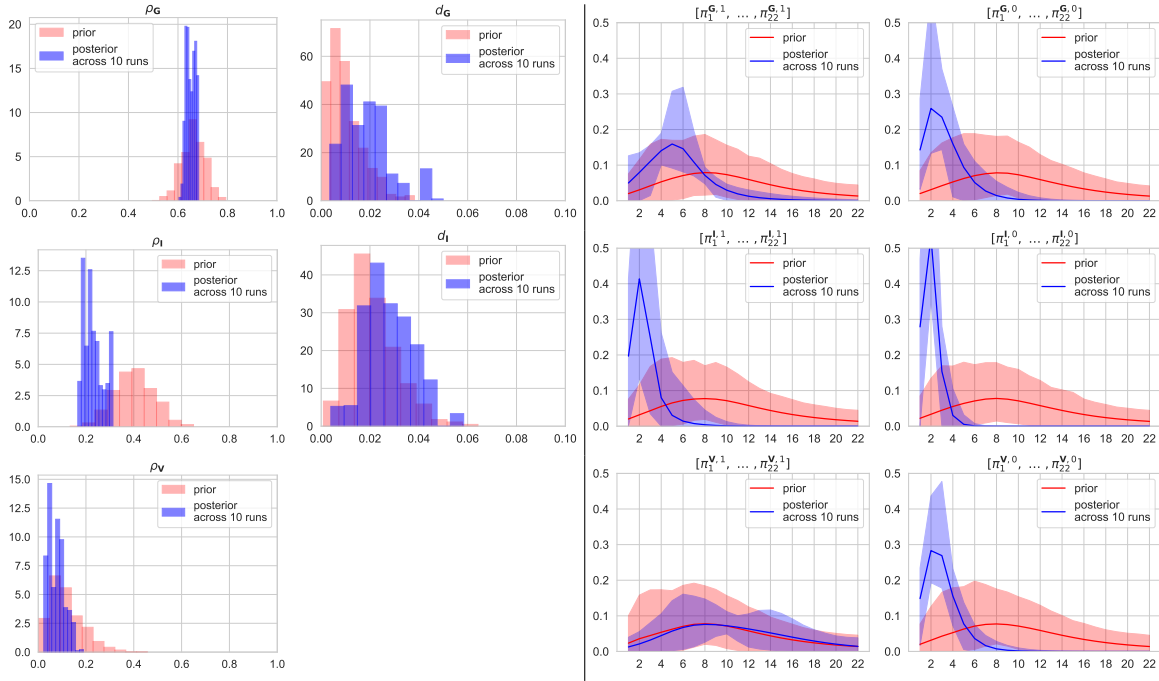


Figure 5: **Posterior distributions over parameters for Massachusetts (trained on Nov. 11 - Jan. 11).** We show transition parameters (left) and duration parameters (right) after fitting on 2 months of counts, where each day we used available census counts for \mathbf{G} , \mathbf{I} , \mathbf{V} , and \mathbf{T} . The colored interval of duration plots shows the 2.5 - 97.5th percentile intervals of 2000 samples (10 runs, each with 200 samples). The prior is also shown for comparison.

diagnose stages as well as the terminal stage \mathbf{T} . In contrast, for the U.K. we only have counts of all beds ($\mathbf{G} + \mathbf{I} + \mathbf{V}$), patients on ventilation (\mathbf{V}), and recovered (\mathbf{R}). Other sites may yet offer other different subsets of counts $\{\mathbf{G}, \mathbf{I}, \mathbf{V}, \mathbf{T}, \mathbf{R}\}$.

To assess sensitivity to data granularity, we trained the ACED-HMM on MA data after an intentional preprocessing that aggregated all intermediate stages. We performed a fit to this reduced data – only counts of all beds ($\mathbf{G} + \mathbf{I} + \mathbf{V}$) and death (\mathbf{T}) – and to the original data. App. B.5 describes detailed results. As expected, forecasts for model stages where training data was not provided are generally of lesser quality, unless closely-related counts informed training (e.g. \mathbf{V} forecasts were reasonable given \mathbf{T} data).

6.5. Scalability assessment: Can our ABC scale to large populations?

A key design requirement for our ABC procedure is the ability to fit posteriors to data from high-population regions or even entire countries. We propose a simple but effective way to *scale up*: given a scaling factor $r > 1$, we simulate only a fraction $\frac{1}{r}$ of all admissions (at a fraction of the computational cost). Then we rescale the resulting aggregate counts produced by the simulation appropriately: $r\tilde{y}_t^k$. In expectation, this approach should produce an unbiased estimate of simulated counts from the original admitted population. In Fig. B.4, we show that our proposed scaling procedure is effective. Simulating only 20% of admissions in MA and rescaling, we see a qualitatively similar forecast with only modestly increased

uncertainty. With exact simulation of Massachusetts already feasible, by applying 10-x or 20-x scaling we can easily forecast for any U.S. state of interest.

6.6. What-if scenario assessment: Can we assess the impact of interventions?

We now consider using our model to assess the societal value of possible *interventions*. Our mechanistic ACED-HMM model can consider changes to its parameters or assumptions about admissions. We stress this analysis is exploratory, not a proper causal model.

Below, we study two possible scenarios for the state of California (CA), which chose to enact significant *restrictions* in late November 2020 in light of rising hospital utilization and especially rising ICU utilization. We imagine each intervention is rolled out on Nov. 19th (when CA’s Limited Stay at Home Order was announced). We used estimated posteriors trained on CA count data from November 11th to January 11th (using 30-x scaling). Because the training period overlaps with the intervention period, this analysis should be interpreted as a *retrospective* characterization of what might have happened. Our method can forecast into the future if desired, but we have not done so here for simplicity.

What-if Scenario 1: Intervention that lowers hospitalization rates. A recent Phase 3 trial suggests that two monoclonal antibody drugs (bamlanivimab and etesevimab) together reduced hospitalizations and death by 87% when given to severely ill COVID-19 patients before they needed the hospital (Eli Lilly and Company, 2021). Suppose that policy makers could have aggressively deployed this treatment across the state of California after Nov. 19th, 2020. In Fig. 6, we show our model’s forecasts for several scenarios: what actually happened (no reduction in admissions), an ideal retrospective rollout (a linear ramp up to a full 87% reduction in hospital admissions after 30 days), and two more realistic scenarios that assume a 50% and 25% reduction in hospital admissions, respectively. Given Fig. 6 plus knowledge of capacity limits, policy makers could use such quantitative forecasts to decide which rollout strategies to prioritize.

What-if Scenario 2: Pharmaceutical that lowers recovery times in the hospital. We next consider an intervention that changes the *number of days* spent when recovering in the hospital. Our Scenario 2 supposes that a new treatment reduces in-hospital recovery time by 25%. This is a conservative translation of a recent clinical trial described by Beigel et al. (2020), suggesting that the drug remdesivir reduced median time to recovery from 15 days (placebo) to 10 days (treatment) for hospitalized patients with COVID-19.

To operationalize this scenario, for each segment ℓ of patient n ’s trajectory we impose a 25% reduction in the sampled duration $\Delta_{n\ell}$ when the patient is recovering (when $h_{n\ell} = 1$). Durations when health is declining are not modified. When we reduce durations, we round to an integer number of days in a way that preserves expected value (e.g. if the down-scaled duration is 9.75 days, we return a value of 10 days 75% of the time and 9 days 25% of the time). We also enforce that each segment’s duration cannot be reduced below 1 day.

Fig. 7 shows the difference in utilization between the standard and 25% reduction in recovery durations. As expected, we see noticeable differences in the populations in the general ward. However, the model suggests the ICU counts would not change much, probably because the ICU is unfortunately dominated by *declining* rather than *recovering* patients.

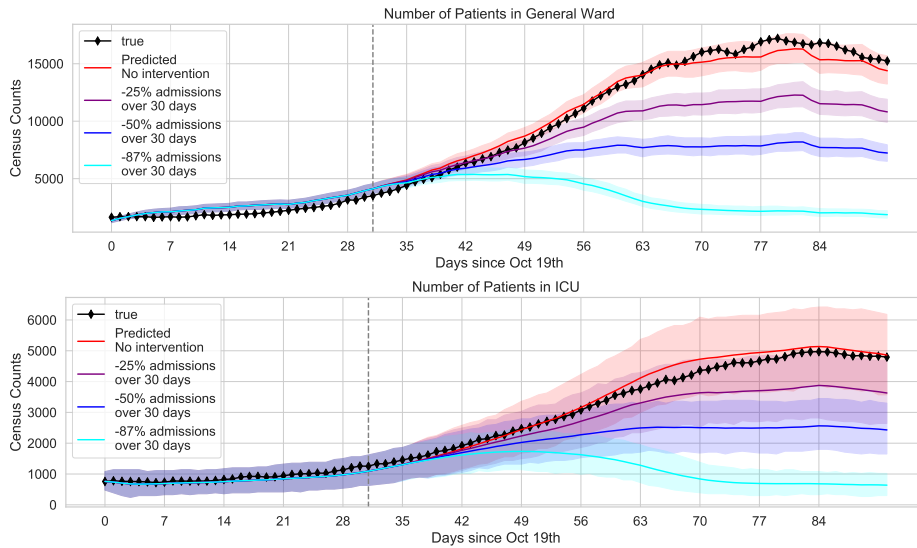


Figure 6: Forecasted demand for General Ward (top) and ICU (bottom) beds in California under Scenario 1 (Sec. 6.6).

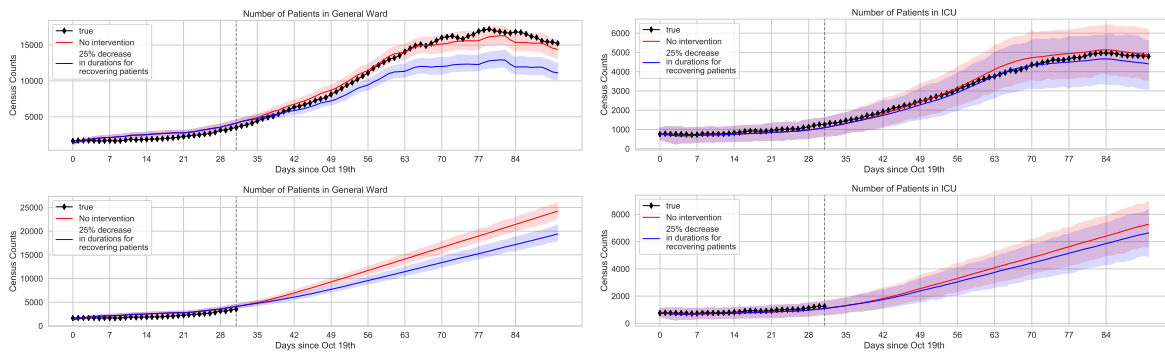


Figure 7: Forecasted demand for General Ward beds (left) and ICU beds (right) using actual admission counts (top) and a worst-case projection of admission counts (bottom) in California under Scenario 2 (Sec. 6.6).

7. Discussion

We have contributed a new probabilistic model, developed a scalable ABC algorithm to draw samples from the (approximate) posterior distribution over model parameters, and validated these contributions on several datasets relevant to hospital demand modeling.

Limitations. Our proposed ACED-HMM model has several limitations. A major limitation is the assumed availability of admission counts when making forecasts, when of course for a real forecast into the future these counts are unknown. We hope in future work to explore integrating models for forecasting future admissions, such as the model by Qian et al. (2020). Another major limitation is that we assume that the values of most parameters are static over time. In reality, we know that many factors change over time, including properties of the disease (e.g. variants with different infection dynamics), characteristics of the infected population (e.g. age distributions and vaccination status), and hospital care

practices. In the language of statistics, we know the real progression of patients through the hospital is neither *homogenous* nor *Markov*, so we should not expect our forecasts to be accurate far into the future. We could allow the model to *learn* how recovery probabilities or duration distributions might change over time. Finally, our model would likely be improved by capturing age as a key prognostic covariate that informs both the recovery probabilities and the duration probabilities. Naturally this would require more information, such as estimates of the age distribution on each day in the target region and the effect of mandated non-pharmaceutical interventions or vaccination on the evolving age distribution. However, if available this information could increase the model’s power as an explanatory tool.

Further limitations can be ascribed to our ABC methodology. Though our ABC procedure is simple and scalable, it is inefficient in the way it proposes samples via a random walk, with rather high rejection rates. Future work may consider Hamiltonian ABC (Meeds et al., 2015), using *gradient* information to improve proposals and accelerate posterior exploration.

Advantages. Towards our twin goals of being able to provide reasonable forecasts and assess societal value, the *extreme portability* of our approach is a key advantage. The model requires only daily count data that are available from public health agencies throughout the world. There is no requirement for patient-level data. We further emphasize that the model and learning procedures are easily extended. Changing the model (such as adding age-dependent duration probabilities or allowing recovering patients to later decline) would simply require editing the simulation code that samples each patient’s trajectory and specifying relevant prior and proposal distributions for any additional parameters.

Overall, we hope this study and our open-source code release leads to an improved ability to make hospital demand forecasts that are portable (by requiring only aggregate count data), interpretable and extrapolate-able (by using a mechanistic model), accurate (by fitting model parameters to local conditions via our ABC algorithm), and appropriately uncertain (by producing posterior distributions rather than point estimates). Although further work is needed to properly assess the societal value of possible interventions with this model, we believe our work represents a promising first step.

Acknowledgements

Authors GMV, CN, JTC, and MCH gratefully acknowledge funding support from a consortium of industry partners for this project – AstraZeneca, Bristol-Myers Squibb, Eli Lilly, Gilead Sciences, Janssen, Merck, Pfizer, and Regeneron – via an agreement with the Center for the Evaluation of Value and Risk in Health (CEVR) at Tufts Medical Center.

Authors MCH and DMK thank the Office of the Vice Provost for Research at Tufts University for support for this project under a “Tufts Springboard” award.

We thank several graduate students in the Department of Computer Science at Tufts University who assisted with various early aspects of this project: Preetish Rath, Zhe Huang, and Panos Lymperopoulos (early model prototyping efforts) as well as Gabriel Appleby, Diana Eastman, and Ab Mosca (visualization efforts). We thank Jinny Park (RA at Tufts Medical Center) for assistance with literature surveys and project management.

References

- Shakir Ahamad, Scotty Branch, Shea Harrelson, Mohd Kamil Hussain, Mohammad Saquib, and Saeed Khan. Primed for global coronavirus pandemic: Emerging research and clinical outcome. *European Journal of Medicinal Chemistry*, 209:112862, 2021.
- Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50(1-2):5–43, 2003.
- Ruth E. Baker, Jose-Maria Peña, Jayaratnam Jayamohan, and Antoine Jérusalem. Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biology Letters*, 14(5):20170660, 2018.
- Mark A. Beaumont. Approximate Bayesian Computation in Evolution and Ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41(1):379–406, 2010. URL <https://doi.org/10.1146/annurev-ecolsys-102209-144621>.
- Mark A. Beaumont, Wenyang Zhang, and David J. Balding. Approximate Bayesian Computation in Population Genetics. *Genetics*, 162(4), 2002. URL <https://www.genetics.org/content/162/4/2025>.
- Christopher Beckham and Christopher Pal. Unimodal probability distributions for deep ordinal classification. In *International Conference on Machine Learning*, 2017. URL <http://proceedings.mlr.press/v70/beckham17a.html>.
- John H. Beigel, Kay M. Tomashek, Lori E. Dodd, Aneesh K. Mehta, Barry S. Zingman, Andre C. Kalil, Elizabeth Hohmann, Helen Y. Chu, Annie Luetkemeyer, et al. Remdesivir for the Treatment of Covid-19 — Final Report. *New England Journal of Medicine*, 383(19):1813–1826, 2020.
- Michael G. B. Blum and Viet Chi Tran. HIV with contact tracing: A case study in approximate Bayesian computation. *Biostatistics*, 11(4):644–660, 2010.
- U. S. CDC. COVID-19 Pandemic Planning Scenarios. Technical report, U. S. Centers for Disease Control, 2020. URL <https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html#table-2>.
- Peter Chen, Ajay Nirula, Barry Heller, Robert L. Gottlieb, Joseph Boscia, Jason Morris, Gregory Huhn, Jose Cardona, Bharat Mocherla, et al. SARS-CoV-2 Neutralizing Antibody LY-CoV555 in Outpatients with Covid-19. *New England Journal of Medicine*, 384(3):229–237, 2021.
- Stephen E. Congly, Rhea A. Varughese, Crystal E. Brown, Fiona M. Clement, and Lynora Saxinger. Treatment of Moderate to Severe Respiratory COVID-19—A Cost-Utility Analysis. *medRxiv*, 2020. URL <https://www.medrxiv.org/content/10.1101/2020.09.21.20199182v2>.
- Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48), 2020. URL <https://www.pnas.org/content/117/48/30055>.

- Eli Lilly and Company. Lilly’s bamlanivimab and etesevimab together reduced hospitalizations and death in Phase 3 trial for early COVID-19. 2021. URL <https://investor.lilly.com/news-releases/news-release-details/lillys-bamlanivimab-and-etesevimab-together-reduced>.
- Richard H Epstein and Franklin Dexter. A Predictive Model for Patient Census and Ventilator Requirements at Individual Hospitals During the Coronavirus Disease 2019 (COVID-19) Pandemic: A Preliminary Technical Report. *Cureus*, 12(6), 2020.
- J. D. Ferguson. Variable duration models for speech. In *Proceedings of the Symposium on the Application of Hidden Markov Models to Text and Speech*, Princeton, NJ, USA, 1980.
- Seth Flaxman, Swapnil Mishra, Axel Gandy, H. Juliette T. Unwin, Thomas A. Mellan, Helen Coupland, Charles Whittaker, Harrison Zhu, Tresnia Berah, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, 584(7820): 257–261, 2020.
- Joseph Friedman, Patrick Liu, Emmanuela Gakidou, and IHME COVID19 Model Comparison Team. Predictive performance of international COVID-19 mortality forecasting models. *medRxiv*, page 2020.07.13.20151233, 2020.
- Afschin Gandjour. How Many Intensive Care Beds are Justifiable for Hospital Pandemic Preparedness? A Cost-effectiveness Analysis for COVID-19 in Germany. *Applied Health Economics and Health Policy*, 19(2):181–190, 2021.
- Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian Data Analysis*. CRC Press, 2013.
- Youyang Gu. Concerns with the IHME Model, 2020. URL <https://covid19-projections.com/about/#concerns-with-the-ihme-model>.
- Matthew D Hoffman and Andrew Gelman. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, page 31, 2014. URL <http://jmlr.org/papers/volume15/hoffman14a/hoffman14a.pdf>.
- IHME COVID-19 Forecasting Team. Supplementary Information: Modeling COVID-19 scenarios for the United States. Supplementary Material, *Nature Medicine*, 2020. URL <https://doi.org/10.1038/s41591-020-1132-9>.
- John P.A. Ioannidis, Sally Cripps, and Martin A. Tanner. Forecasting for COVID-19 has failed. *International Journal of Forecasting*, 2020.
- Nicholas P. Jewell, Joseph A. Lewnard, and Britta L. Jewell. Predictive Mathematical Models of the COVID-19 Pandemic: Underlying Principles and Value of Projections. *JAMA*, 323(19):1893–1894, 2020.
- Youngji Jo, Lise Jamieson, Ijeoma Edoaka, Lawrence Long, Sheetal Silal, Juliet R. C. Pulliam, Harry Moultrie, Ian Sanne, Gesine Meyer-Rath, et al. Cost-effectiveness of remdesivir and dexamethasone for COVID-19 treatment in South Africa. *medRxiv: The Preprint Server for Health Sciences*, 2020.

- David Koeckerling, Joseph Barker, Nadeesha L. Mudalige, Oluwatobiloba Oyefeso, Daniel Pan, Manish Pareek, Jonathan P. Thompson, and G. Andre Ng. Awake prone positioning in COVID-19. *Thorax*, 75(10):833–834, 2020.
- Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M. Blei. Automatic Differentiation Variational Inference. *Journal of Machine Learning Research*, 18(14):1–45, 2017. URL <http://jmlr.org/papers/v18/16-107.html>.
- Alexandra Hope Lee, Panagiotis Lymperopoulos, Joshua T Cohen, John B Wong, and Michael C Hughes. Forecasting COVID-19 Counts at a Single Hospital: A Hierarchical Bayesian Approach. In *ICLR 2021 Workshop on Machine Learning for Preventing and Combating Pandemics*, 2021. URL <https://arxiv.org/abs/2104.09327>.
- Michael Lingzhi Li, Hamza Tazi Bouardi, Omar Skali Lami, Nikolaos Trichakis, Mohammad Fazel Zarandi, and Dimitris Bertsimas. Overview of DELPHI Model V3. Technical report, COVIDAnalytics, 2020. URL https://covidanalytics.io/DELPHI_documentation_pdf.
- Jean-Michel Marin, Pierre Pudlo, Christian P. Robert, and Robin J. Ryder. Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012. URL <https://doi.org/10.1007/s11222-011-9288-2>.
- Paul Marjoram, John Molitor, Vincent Plagnol, and Simon Tavaré. Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100(26), 2003. URL <https://www.pnas.org/content/100/26/15324>.
- Edward Meeds, Robert Leenders, and Max Welling. Hamiltonian ABC. *arXiv:1503.01916 [cs, q-bio, stat]*, 2015. URL <http://arxiv.org/abs/1503.01916>.
- Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- C. D. Mitchell and L. H. Jamieson. Modeling duration in a hidden Markov model with the exponential family. In *1993 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 2, pages 331–334, 1993.
- Christopher JL Murray, others, and the IHME COVID-19 health service utilization forecasting team. Forecasting the impact of the first wave of the COVID-19 pandemic on hospital demand and deaths for the USA and European Economic Area countries. *medRxiv*, 2020. URL <https://www.medrxiv.org/content/10.1101/2020.04.21.20074732v1>.
- Peter J. Neumann, Joshua T. Cohen, David D. Kim, and Daniel A. Ollendorf. Consideration Of Value-Based Pricing For Treatments And Vaccines Is Important, Even In The COVID-19 Pandemic. *Health Affairs*, 40(1):53–61, 2020.
- Zhaozhi Qian, Ahmed M. Alaa, and Mihaela van der Schaar. CPAS: The UK’s national machine learning-based hospital capacity planning system for COVID-19. *Machine Learning*, pages 1–21, 2020.

- Robert C. Reiner, Ryan M. Barber, James K. Collins, Peng Zheng, Christopher Adolph, James Albright, Catherine M. Antony, Aleksandr Y. Aravkin, Steven D. Bachmeier, et al. Modeling COVID-19 scenarios for the United States. *Nature Medicine*, 27(1):94–105, 2021.
- Michael Roimi, Rom Gutman, Jonathan Somer, Asaf Ben Arie, Ido Calman, Yaron Bar-Lavie, Udi Gelbshtein, Sigal Liverant-Taub, Arnona Ziv, et al. Development and validation of a machine learning model predicting illness trajectory and hospital utilization of COVID-19 patients—a nationwide study. *Journal of the American Medical Informatics Association: JAMIA*, 2021.
- John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. Probabilistic programming in Python using PyMC3. *PeerJ Computer Science*, 2:e55, 2016.
- The Atlantic Monthly Group. The COVID tracking project - Data API, 2021. URL <https://covidtracking.com/data/api>.
- U. S. Dept. of Health and Human Services. COVID-19 reported patient impact and hospital capacity by state timeseries, 2021. URL <https://healthdata.gov/Hospital/COVID-19-Reported-Patient-Impact-and-Hospital-Capa/g62h-syeh>.
- Martin J. Wainwright and Michael I. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- Gary E. Weissman, Andrew Crane-Droesch, Corey Chivers, ThaiBinh Luong, Asaf Hanish, Michael Z. Levy, Jason Lubken, Michael Becker, Michael E. Draugelis, et al. Locally Informed Simulation to Predict Hospital Capacity Needs During the COVID-19 Pandemic. *Annals of Internal Medicine*, 2020.
- Shun-Zheng Yu. Hidden semi-Markov models. *Artificial Intelligence*, 174(2):215–243, 2010.
- Jiancheng Zhang, Bing Xie, and Kenji Hashimoto. Current status of potential therapeutic candidates for the COVID-19 crisis. *Brain, Behavior, and Immunity*, 87:59–73, 2020.
- Difan Zou, Lingxiao Wang, Pan Xu, Jinghui Chen, Weitong Zhang, and Quanquan Gu. Epidemic Model Guided Machine Learning for COVID-19 Forecasts in the United States. *medRxiv*, page 2020.05.24.20111989, 2020.

Appendix A. Datasets

In our public code release, we include CSV files of the datasets we used for both the US and UK. Below, we describe how we obtained and preprocessed this data from original sources.

A.1. U.S. data.

Data representing state-level counts came from the [U. S. Dept. of Health and Human Services \(2021\)](#) and the Covid Tracking Project (CTP, [The Atlantic Monthly Group \(2021\)](#))⁶. We accessed the data in February 2021. These sources provide data for all 50 states, we selected a subset of states of interest: California (CA), Massachusetts (MA), South Dakota (SD), and Utah (UT). HHS data provided daily counts of hospital admissions and general ward occupancy counts (stage **G**), while CTP data provided ICU counts (**I** and **V**), as well as deaths (stage **T**).

As much as possible, we validated the consistency of these 2 public data sources by assuring that these counts matched with available corresponding data provided directly by state government public health agencies. For example, we verified the MA data matched counts released by the Massachusetts Department of Public Health⁷.

In summary, our US predictive models all use standardized counts representing adult hospitalizations for *confirmed* COVID-19 cases. We are certain that our data for the general ward (**G**) accounts for the adult population while excluding pediatric cases. Through data-consistency checks across the 2 data sources, we convinced ourselves that CTP’s **I**, **V**, and **T** data can be assumed as a good representation of adult populations (even though CTP reports uncertainty about whether their numbers include pediatric cases).

HHS fields used. HHS data descriptions are found on the clickable columns of the data web page⁸. We used the following data fields, with quoted definitions extracted from the data source.

- **previous day admission adult covid confirmed** : “Number of patients who were admitted to an adult inpatient bed on the previous calendar day who had confirmed COVID-19 at the time of admission in this state” (We considered previous day admissions as a feasible current day admission because admissions rates were not abruptly changing in orders of magnitudes between days. We recommend considering making the 1 day adjustment in future use of this data field.)
- **total adult patients hospitalized confirmed covid** : “Reported patients currently hospitalized in an adult inpatient bed who have laboratory-confirmed COVID-19. This include those in observation beds”
- **staffed icu adult patients confirmed covid** : “Reported patients currently hospitalized in an adult ICU bed who have confirmed COVID-19 in this state”

6. <https://api.covidtracking.com/v1/states/daily.csv>

7. <https://www.mass.gov/info-details/archive-of-covid-19-cases-in-massachusetts>

8. <https://healthdata.gov/Hospital/COVID-19-Reported-Patient-Impact-and-Hospital-Capacity/g62h-syeh/data>

CTP fields used. CTP data descriptions are found in the 'historic values for all states' section of the data API webpage page⁹. We used the following data fields, with quoted definitions extracted from the data source.

- **inIcuCurrently** : "Individuals who are currently hospitalized in the Intensive Care Unit with COVID-19. Definitions vary by state / territory, and it is not always clear whether pediatric patients are included in this metric. Where possible, we report patients in the ICU with confirmed or suspected COVID-19 cases"
- **onVentilatorCurrently** : "Individuals who are currently hospitalized under advanced ventilation with COVID-19. Definitions vary by state / territory, and it is not always clear whether pediatric patients are included in this metric. Where possible, we report patients on ventilation with confirmed or suspected COVID-19 cases."
- **deathIncrease** : "Daily increase in death, calculated from the previous day's value."

We note that as of 2021-04-14, the CTP website relates the following message: "*The COVID Tracking Project has ended all data collection as of March 7, 2021. These files are still available, but will only include data up to March 7, 2021. We'll be publishing research through May, and then we will—fully and accessibly—archive our work and be done.*"

Obtaining stage-specific counts from raw data fields.

- Counts for daily admissions to the general ward were obtained from field **previous day admission adult covid confirmed**.
- **G** : Daily counts for general ward occupancy were obtained from field **total adult patients hospitalized confirmed covid** subtracting field **staffed icu adult patients confirmed covid**
- **I** : Daily counts for ICU-off-ventilator were obtained from field **inIcuCurrently** subtracting field **onVentilatorCurrently**
- **V** : Daily counts for ICU-on-ventilator were obtained from field **onVentilatorCurrently**
- **T** : Daily counts obtained from field **deathIncrease**

Variable data granularity across states. Not all states report counts from all stages. For example, UT and CA only provide aggregate ICU counts, and thus do not distinguish between on and off the ventilator. We thus take the provided count as the sum of the two ICU stages **I, V** of our model, and adapt our ABC accordingly to evaluate a distance on this combined stage.

Smoothing terminal counts. For all states, we find the daily counts for the terminal death stage **T** to be both noisy and unreliable due to human factors in the reporting process. In particular, the reported counts often drop to unrealistic values at or close to zero during holidays and weekends (e.g. in some states we observed zero deaths on Thanksgiving and Christmas). Thus, we *smooth* these counts by replacing each day's raw value with the centered moving average across 5 days, being careful to not leak test-counts information into the moving average of training counts.

9. <https://covidtracking.com/data/api>

A.2. UK data.

We selected two hospital sites in England, United Kingdom with consistent data availability and a large volume of cases. We used public data sourced from the UK National Health Service (NHS)¹⁰, which provided the count of beds occupied by COVID-19 patients on each day at each hospital. We used the following data fields:

- mechanical ventilators used for covid patients (\mathbf{V})
- total number of beds used for covid patients ($\mathbf{G} + \mathbf{I} + \mathbf{V}$)
- patients discharged from covid hospitalization (\mathbf{R})
- patients admitted with covid
- patients diagnosed in-hospital with covid but who were not admitted as covid patients

In our experiments, we aggregated daily covid admissions and daily in-hospital covid diagnoses to form the daily admissions in \mathbf{G} that are then fed into ACED-HMM. The two hospital sites we selected had the same available counts. Unlike for US data, we did not find the need to smooth any count.

Appendix B. Additional Results

We offer extended commentary on the experimental protocols and results of Sec. 6 here.

B.1. Quantitative Error Assessment: UK Results

Here, in Table B.1 we provide mean absolute error assessments for our method and various baselines for two hospital sites in the U.K. This matches the main paper’s Table 1, which focused on U.S. state data.

Table B.1 shows consistent gains for our proposed ACED-HMM especially in the “all beds” and “on ventilator” stages. For the “recovered” stage, our method is competitive for South Tees, and mildly competitive for the Oxford site (Bayes LR seems to offer slightly better MAE on that site: 10.5 vs. our 13.4).

B.2. Quantitative Error Assessment: Baseline Details

In Tables 1 and B.1, we compared our proposed ACED-HMM to several reasonable baselines. (We did not try all methods on all datasets due to limited computational resources; poor performing baselines such as ACED-HMM+Prior and AR-Poisson were cut from further experiments to save time and energy). Further details about these baselines are below.

ACED-HMM + Prior. First, we consider “ACED-HMM + Prior”, which makes forecasts using the *prior* distribution over our model’s parameters, rather than the posterior fit via ABC. This could indicate how reliable our model would be if we did not adapt parameters to the specific region of interest but instead chose values for transition and duration probabilities from a reasonable literature survey (as our prior is chosen). This could also indicate how important other factors (such as the provided admission counts in testing period) are to the forecast.

10. <https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/>

	Method	G + I + V		All Beds			V OnVentInICU			R Recovered		
		MAE	lower - upper	MAE	lower - upper	MAE	lower - upper	MAE	lower - upper			
South Tees	ACED-HMM + ABC	8.7	7.7 - 10.2	4.6	3.9 - 5.1	5.7	5.5 - 5.8					
	ACED-HMM + Prior	118.4	109.8 - 125.9	10.7	10.0 - 11.4	5.9	5.7 - 6.2					
	AR-Poisson	49.2	40.3 - 60.5	417.3	15.8 - 2209.5	49.2	9.7 - 518.7					
	Median predictor	102.0	NA	17.1	NA	9.9	NA					
	Bayes LR (day)	114.0	113.4 - 114.6	20.5	20.4 - 20.6	13.6	13.4 - 13.7					
	Bayes LR (day+adm.)	81.7	81.3 - 82.1	13.9	13.8 - 14.0	12.4	12.3 - 12.6					
	Mean Test y	199.5		31.1		20.8						
Oxford	ACED-HMM + ABC	18.0	15.4 - 22.4	6.3	5.5 - 6.9	13.4	13.2 - 13.6					
	Median predictor	216.7	NA	34.7	NA	25.1	NA					
	Bayes LR (day)	71.5	70.5 - 72.5	16.2	16.0 - 16.3	12.5	12.3 - 12.7					
	Bayes LR (day+adm.)	21.8	21.4 - 22.3	10.5	10.4 - 10.7	10.5	10.3 - 10.7					
	Mean Test y	274.7		37.7		33.1						

Table B.1: Quantitative error assessment for predictions for two UK hospital sites, evaluated over the test period Jan. 4 - Feb. 3, 2021. We assess mean absolute error (MAE) at each stage where we have observed data: total hospital beds **G + I + V**, on ventilator in ICU **V**, and recovered **R**. For all sites and stages, we report the mean count in the test period to indicate scale. We report lower and upper estimates on the MAE, via repeating all MAE computations across 100 separate batches of posterior samples. We report the 2.5th and 97.5th percentiles across these batches.

AR-Poisson. Second, we compare to “AR-Poisson”, a non-mechanistic probabilistic autoregressive model with a Poisson likelihood that forecasts each univariate count time-series individually, inspired by the COVID-19 single site hospital forecasts of Lee et al. (2021). This is a custom univariate forecasting model fit to each stage separately. It does not use admission counts at all. The AR-Poisson uses modern Hamiltonian MCMC to fit posteriors (rather than our less-efficient-to-explore ABC). See Appendix E for complete details. As documented in Lee et al. (2021), we believe this model represents a reasonable baseline when forecasts need a roughly linear function of the training data. However, when test period looks different than training in trends (as in Fig. 3, we should expect (and indeed see in practice) that performance suffers.

Median of train set. This baseline simply predicts the median count observed in the training set of each stage. We selected median (not mean) because it is the simple summary statistic that provably optimizes mean absolute error (MAE) as a performance metric.

Bayes LR. We used two versions of a Bayesian linear regression model (as implemented in sklearn with default parameters). The first version (“Bayes LR (day)”) used as its only feature the absolute time in days (day 0 is the start of the training set, then the days increase through the entire training period and into the testing period). The second version (“Bayes LR (day+adm.)”) took as input a feature vector containing absolute time in days, as well as the admission count from each of the last 21 days (we set the admissions values to 0 for days before day 0 in the training set). This model has the key property of using recent admissions as its covariates, which are also used by ACED-HMM to make forecasts.

IHME. For all U.S. states we compare against the January 15th public release of IHME model forecasts (Reiner et al., 2021). These forecasts were produced by the IHME team by training on data available before January 12th, thus matching the start of our testing period almost exactly. This model does not make use of admission counts from the testing period. See Appendix F for details.

Uncertainty estimation details. For all probabilistic models (both versions of ACED-HMM, the AR-Poisson model, and the two Bayesian Ridge regression baselines), we compute the MAE and its confidence interval in Tab. 1 as follows. We draw 100 samples from the posterior (prior for ACED-HMM + Prior) distribution over the model parameters, we produce individual forecasts for each sample, and compute the MAE using the mean of these forecasts. We repeat the process 100 times, and we report the average MAE with the 2.5th - 97.5th percentile range across these batches as confidence interval. The public release of forecasts from IHME only provides three daily forecast values - a mean estimate, a lower bound and an upper bound - thus making it impossible to compute a confidence interval analogous to that which we compute for the other three models. Thus, we simply report the MAE computed using each of the three estimates. The IHME uncertainty intervals (marked with *) are given for completeness, but should be recognized as different from the other intervals and not directly compared.

B.3. Quantitative Assessment: Posterior Coverage

In Tab. B.2 below, we evaluate the posterior coverage of our estimated posteriors on heldout data at target fractions of 50%, 80% and 95%, for both U.S. and U.K. data. Ideal probabilistic models that are well-calibrated would have an empirical fraction of observations *exactly equal* to the target fraction.

Across Tab. B.2, we see that all models could be better calibrated. Our proposed ACED-HMM appears “too confident”: it primarily underestimates its posterior mass intervals. For example for MA the interval it predicts will cover 95% of data for **G** and **I** stages actually covers only 68% and 65% of observed counts. One remedy that may help is the scalability approximation that we propose (Sec. B.6) provides more conservative uncertainty estimates.

U.S.								U.K.								
	Coverage (%)	G	I	V	I+V	T	T sm.		Coverage (%)	R	G+I+V	V				
SD	ACED-HMM + ABC	95	87	77	81	100	45	77	S. Tees	AR-Poisson	95	90	100	100		
		80	52	48	71	97	23	68			80	52	100	100		
		50	42	26	32	55	6	32			50	3	58	74		
UT	ACED-HMM + ABC	95	71	NA	NA	100	55	100			ACED-HMM + Prior	95	77	100	84	
		80	39	NA	NA	100	32	87				80	58	13	42	
		50	19	NA	NA	61	23	39				50	29	10	26	
MA	AR-Poisson	95	81	100	45	61	100	100		Oxford	ACED-HMM + ABC	95	55	100	100	
		80	45	90	26	35	68	87					80	19	100	77
		50	29	35	13	13	32	10					50	10	87	35
	ACED-HMM + Prior	95	52	39	61	71	42	23								
		80	6	35	42	39	26	23								
		50	6	19	19	23	23	19								
	ACED-HMM + ABC	95	68	65	42	97	55	84								
		80	48	45	16	94	39	74								
		50	32	16	3	74	16	58								

Table B.2: Coverage assessment for estimated posteriors fit to U.S. data (left) and U.K. data (right). This metric assesses how reliably calibrated the forecasted *distributions* each method produces may be on future unseen data. Given many samples from a predicted distribution, we obtain a centered interval of specified coverage percentage by computing relevant percentiles from posterior samples (e.g. for the 80% percentage our interval is defined by the 10th and 90th percentiles). We report the observed percentage of daily count data within the testing period at each stage that fall within that predicted interval. Ideally, the observed fraction should be close to the intended coverage fraction.

ABC may be an issue too here: using a proper likelihood (rather than a distance function) may lead to improved calibration.

Our ACED-HMM trained with ABC does have a reasonable fit to the $\mathbf{G} + \mathbf{I} + \mathbf{V}$ counts and \mathbf{V} counts on both UK hospitals, as well as a slight over-estimate of the interval for the total ICU bed count ($\mathbf{I} + \mathbf{V}$) on all US states. These bed counts have been commonly used by US policy makers to determine the imposition of social-distancing measures.

We wish to emphasize that coverages need to be put in the context of their MAEs as well. For instance, while ACED-HMM has worse 95% coverage than AR-Poisson on \mathbf{G} counts, that is because AR-Poisson has much worse MAE with a very high uncertainty interval, effectively making AR-Poisson's predictions less useful than those of ACED-HMM.

B.4. Interpretation of Learned Posteriors: UK Results

We visualize the learned distributions for all parameters of our ACED-HMM in Fig. B.1.

As expected, these posteriors (derived from data for a single hospital) show a higher variance (uncertainty) than the posteriors derived from the state of Massachusetts in the main paper's Fig. 5. The distributions shown for this U.K. hospital appear to favor longer durations. Patients appear more likely to recover both in the General Ward and in the ICU than in MA.

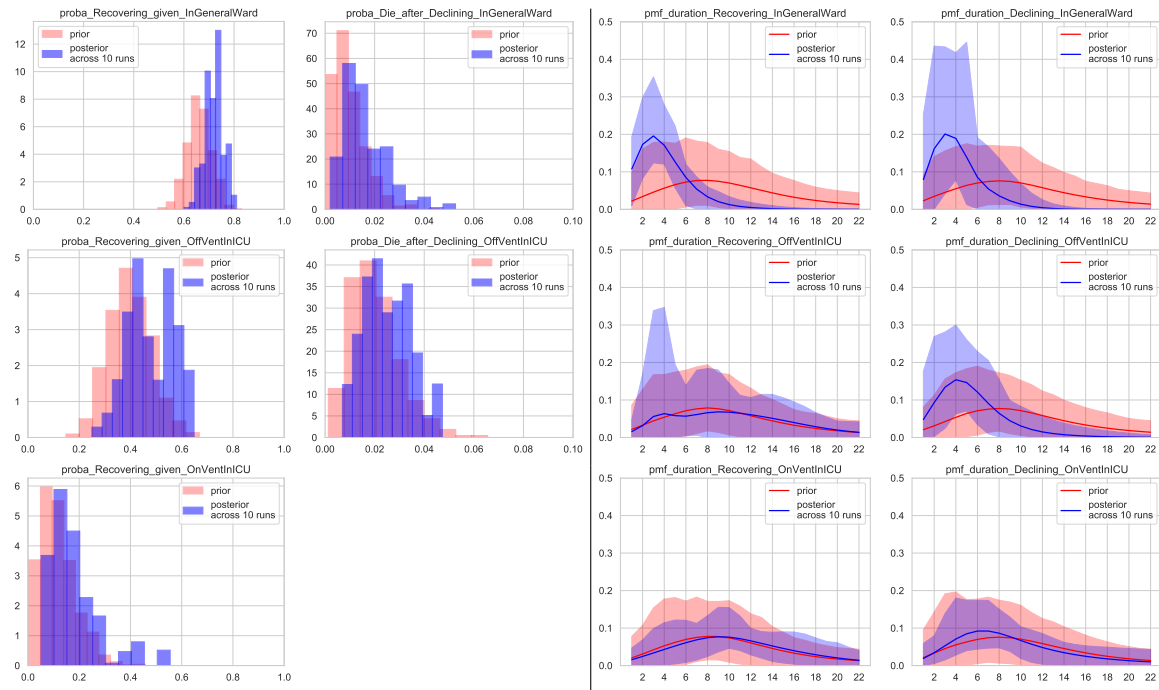


Figure B.1: **Posterior distributions over parameters for South Tees hospital in the UK (trained on Nov. 3rd - Jan. 3rd).** We show transition parameters (left) and duration parameters (right) after fitting on 2 months of counts, where each day we used available census counts for \mathbf{R} , $\mathbf{G} + \mathbf{I} + \mathbf{V}$, and \mathbf{V} . The colored interval of duration plots shows the 2.5 - 97.5th percentile intervals of 2000 samples (10 runs, each with 200 samples). The prior is also shown for comparison.

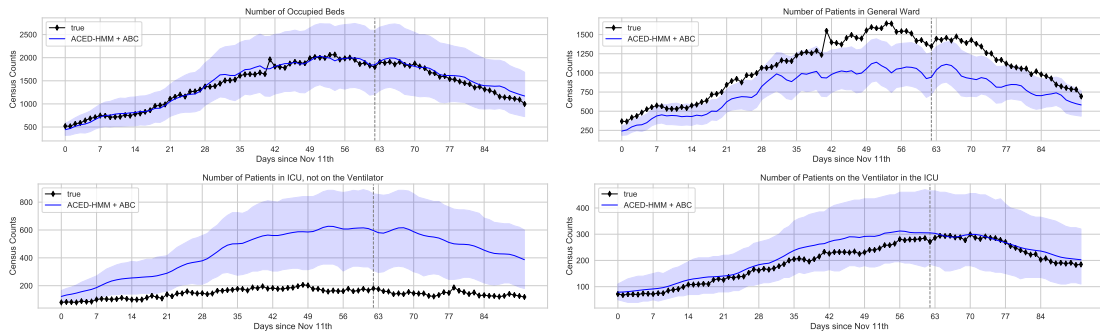


Figure B.2: **Fit and forecasts on MA data for our data granularity ablation study.** Here, the ACED-HMM was trained to match the total number of hospital beds ($\mathbf{G} + \mathbf{I} + \mathbf{V}$) as well as terminal counts (\mathbf{T}) for MA state-level data. (Compare to Fig. 3, which fit on counts for individual stages \mathbf{G} , \mathbf{I} , \mathbf{V} , and \mathbf{T}). The model is able to fit and forecast the total bed counts accurately (upper-left), but did not recover the counts at some individual compartments (\mathbf{G} : lower left, \mathbf{I} upper right) well. The ventilator counts (\mathbf{V} : lower right) are near-optimally recovered as they have a very direct influence over the observed \mathbf{T} counts.

B.5. Data granularity ablation study

To assess the model’s sensitivity to data granularity, we trained the ACED-HMM on the Massachusetts data we used in our other experiments, but with a key manually imposed difference in available data: instead of fitting to the finest granularity of available counts – \mathbf{G} , \mathbf{I} , \mathbf{V} and \mathbf{T} separately – we instead aggregate all intermediate stages and fit the parameters to daily counts of all beds ($\mathbf{G} + \mathbf{I} + \mathbf{V}$) and death (\mathbf{T}) only. See Fig. B.2 for fits and forecasts to \mathbf{G} , \mathbf{I} , \mathbf{V} and $\mathbf{G} + \mathbf{I} + \mathbf{V}$ counts, and Fig. B.3 for the learned model parameters.

We find that, while the model trained with less-granular data can fit the aggregate “all beds” counts well in both training and test periods, the model could not recover individual stage counts from the \mathbf{G} and \mathbf{I} stages well. In particular, the algorithm converged to parameters that undershot \mathbf{G} counts and overshot \mathbf{I} counts. The differing signs of these errors roughly cancel out and thus aggregate counts are predicted well. Interestingly, the model managed better predictions of \mathbf{V} . We believe this is due to the presence of \mathbf{T} counts at training time, as these values provide a more direct signal for what the \mathbf{V} counts should be.

This experiment shows that, as expected, forecasts on fine-grained counts that the model was not trained on are generally not reliable, unless some closely-related count is provided (e.g. \mathbf{T} is provided as a signal for \mathbf{V} or \mathbf{R} is provided as a signal for \mathbf{G}). Stronger, more informative priors about the intermediate stages may also improve performance when fine-grained signals are not available.

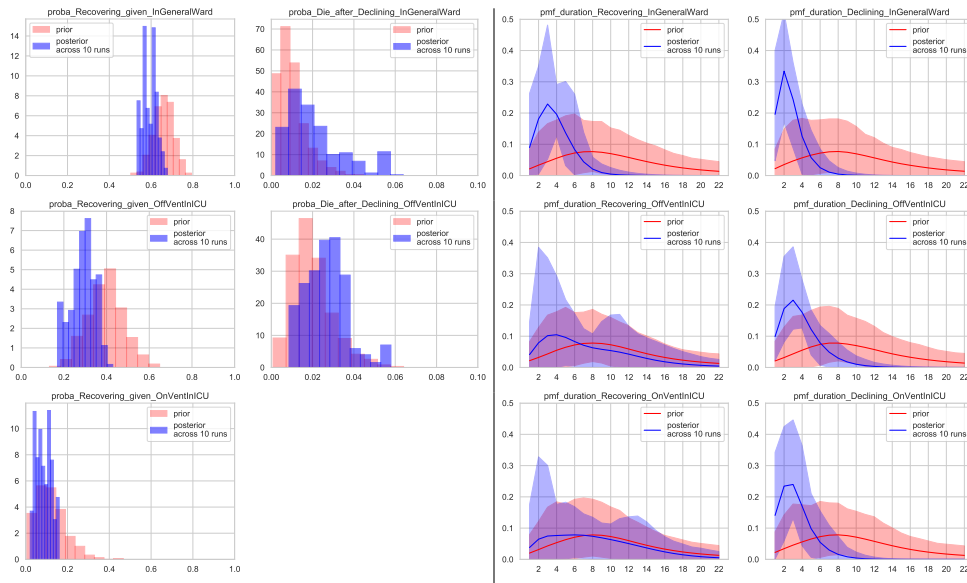


Figure B.3: **Posterior distributions over parameters for MA obtained from our data granularity ablation study.** Here, the ACED-HMM was trained to match the total number of hospital beds ($G + I + V$) as well as terminal counts (T) for MA state-level data. (Compare to Fig. 5, which fit on counts for individual stages G , I , V , and T). The colored interval of duration plots shows the 2.5 - 97.5th percentile intervals of 2000 samples (10 runs, each with 200 samples).

B.6. Scalability Assessment

We visually assess the performance of our procedure for scaling up simulations described in Sec. 6.5 on data from Massachusetts in Fig. B.4. Despite simulating 5x fewer patient trajectories (and thus being substantially cheaper), our scaling leads to reasonably similar overall trends in the forecasts with only modestly increased uncertainty intervals.

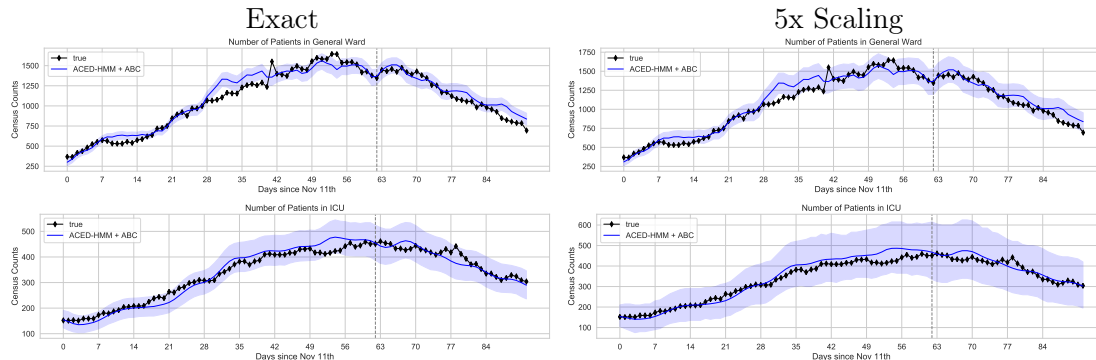


Figure B.4: **Scalability assessment for MA data.** We show the predictions for our ABC method using all admissions recorded in the training period (left), as well as our proposed scaling procedure with a 5x scale factor (right). The scaling simulates only 20% of observed admissions but then scales all counts by 5. The scaling procedure allows simulations to run 5 times faster with the same qualitative trend recovery and only modest increases in uncertainty intervals.

B.7. Synthetic Data Assessment: Can we recover “ideal” model parameters?

We assessed the ability of our procedure to recover “ideal” model parameters under different admissions regimes. We selected a set of parameters that deviated significantly, though not excessively from the prior we derived from CDC data. Then, we simulated sets of counts using these parameters and a fixed random seed under 4 admissions regimes, respectively with 1-x, 3-x, 6-x and 9-x the amounts of true admissions at South Tees Hospitals between November 3rd and February 3rd. We note that the 9-x admissions regime is comparable to that of Massachusetts in terms of total admitted patients, and it thus can be considered a regional-level regime. We then used ABC with the same prior to learn 5 sets of model parameters to fit each of the 5 simulated counts, respectively. We used all possible counts for training: \mathbf{R} , \mathbf{G} , \mathbf{I} , \mathbf{V} and \mathbf{T} .

ABC was able to identify parameters that provided a close-to-optimal fit to the training counts as well as high-quality forecasts on test counts. While the parameters learned were optimal in their ability to explain the training data, in no case were they “ideal”, i.e. the true parameters. In particular, we observed that, the higher the number of total admitted patients, the closer the learned parameters are to the true ones, as can be seen by comparing Fig. B.5 and Fig. B.6, which show the posterior distributions for the 1-x (hospital-level) and 9-x (regional-level) regimes, respectively. Under the regional-level regime, 5 parameters are fully recovered (i.e. mean of posterior matches true parameters), 2 parameters are partially recovered (i.e. posterior shifts from prior towards true parameter), 3 parameters are not recovered (i.e. no shift from prior towards true parameters).

Under our modeling assumptions, higher admissions regimes generate counts which are reproducible by fewer sets of parameters. In other words, the space of “optimal” parameters shrinks. Thus, on higher admissions regimes, our procedure converges closer to the true parameters.

Furthermore, we observed that certain parameters are easier to recover than others. For instance, parameters that have a direct effect on \mathbf{R} and \mathbf{T} counts seem to be among the first to be recovered, this includes the duration probabilities recovering in \mathbf{G} and declining in \mathbf{V} . By contrast, the parameters at the interface between hospital stages seem harder to recover.

We believe another useful experiment using synthetic data to be one that answers the question *How informative does the prior need to be for the posterior to recover all the “ideal” parameters?*. We leave the answering of this question for future work.

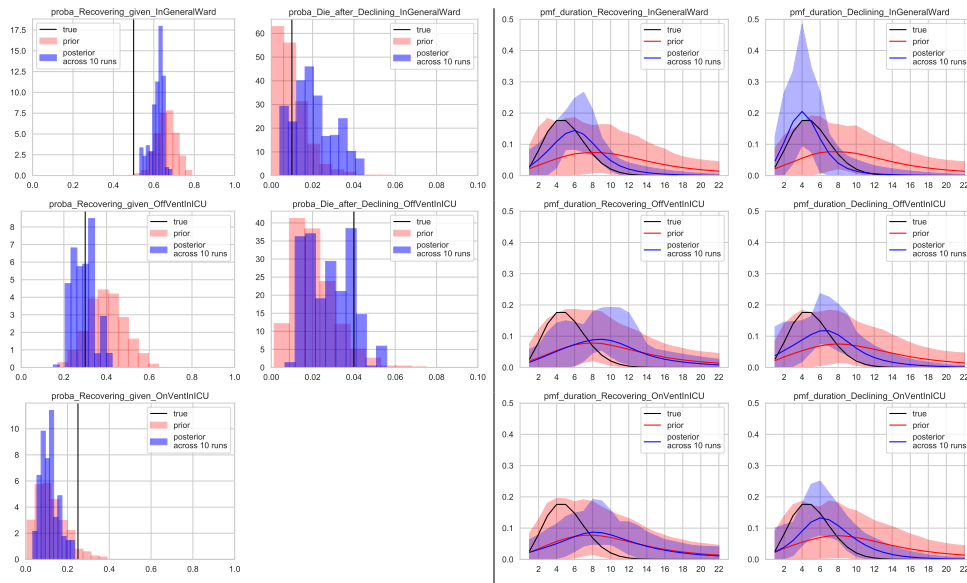


Figure B.5: **Posterior distributions over parameters for synthetic data with hospital-level admissions.** We show transition parameters (left) and duration parameters (right) after fitting on 61 days of counts, where each day we used available simulated census counts for \mathbf{R} , \mathbf{G} , \mathbf{I} , \mathbf{V} , and \mathbf{T} . The colored interval of duration plots shows the 2.5 - 97.5th percentile intervals of 2000 samples (10 runs, each with 200 samples). The prior is also shown for comparison.

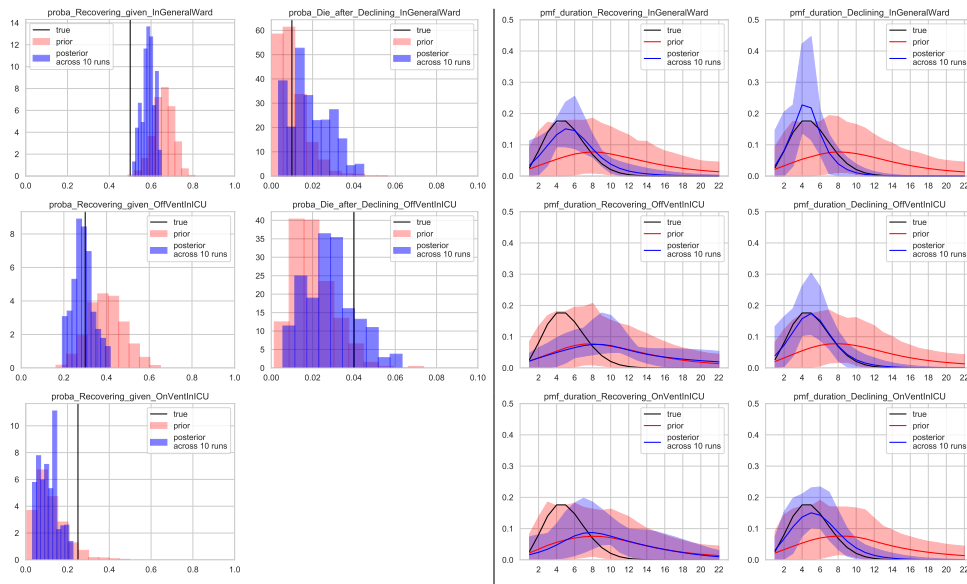


Figure B.6: **Posterior distributions over parameters for synthetic data with regional-level admissions.** We show transition parameters (left) and duration parameters (right) after fitting on 61 days of counts, where each day we used available simulated census counts for \mathbf{R} , \mathbf{G} , \mathbf{I} , \mathbf{V} , and \mathbf{T} . The colored interval of duration plots shows the 2.5 - 97.5th percentile intervals of 2000 samples (10 runs, each with 200 samples). The prior is also shown for comparison.

B.8. Importance of weighting on distance computation

Here, we assess the impact of weighting counts within our distance computation (via the weight parameter w_{tk} in Eq. 7) on forecast quality. To do so, we fit the parameters of ACED-HMM to the Massachusetts and Utah datasets used in our experiments using the same procedure described in Sec. 4, with the only difference that w_{tk} is set to 1.0 for each timestep t and stage k , thus effectively removing any timestep-specific and stage-specific biasing in the distance computation.

Table B.3 reports the test-set MAE scores of posteriors trained with this unweighted distance, and compares them to the test MAE scores of posteriors trained with our weighted distance in Eq. 7. On Utah data, the use of weighted distance does not appear to have a meaningful impact on forecast quality over using the simpler distance: across all stages the estimated MAE interval shows substantial overlap. On Massachusetts data, on the other hand, the use of weights results in slightly better forecasts of **I+V** and **V** counts, with slightly worse forecasts on **G** counts. This effect was sought after in our weight calibration, as, for the purposes of evaluation societal impact, we care more to forecast ICU counts accurately than General Ward counts. Overall, as the differences between weighted and unweighted versions of our ACED-HMM seem to be minor, we do not deem the timestep-specific and stage-specific weighting as leading factors in the success of our method.

Method		G InGeneralWard			I + V InICU			V OnVentInICU			T sm.		Death	
		MAE	lower	- upper	MAE	lower	- upper	MAE	lower	- upper	MAE	lower	- upper	
MA	ACED-HMM + ABC weighted dist.	65.0	61.5	- 68.9	15.7	14.5	- 16.9	34.0	32.6	- 35.7	8.0	7.8	- 8.3	
	ACED-HMM + ABC unweighted dist.	60.1	58.0	- 62.7	16.4	15.0	- 17.6	37.8	35.2	- 40.4	7.8	7.3	- 8.2	
	Mean Test y	1141.6			392.5			249.1			66.5			
UT	ACED-HMM + ABC weighted dist.	20.8	20.2	- 21.3	18.2	17.4	- 19.2	NA			2.4	2.3	- 2.5	
	ACED-HMM + ABC unweighted dist.	20.8	19.9	- 21.7	18.8	16.9	- 20.7	NA			2.5	2.3	- 2.7	
	Mean Test y	272.9			164.3			NA			11.7			

Table B.3: **Quantitative error assessment comparing two versions of the distance function used in ABC: unweighted and time-specific and stage-specific weighting.** We report results for both MA and UT during the testing period (Jan. 12 - Feb. 11, 2021). The method for computing the MAE is analogous to that of Table 1.

B.9. Exploring different duration parameterizations

We perform two sets of experiments exploring different parameterizations of the duration distributions within ACED-HMM, described below.

Alternative Duration Model Experiment 1: Larger truncation limit $D = 44$. In the first experiment, we increase the upper bound of the per-segment duration distributions from 22 to 44, and train the modified ACED-HMM using our procedure on Massachusetts data. Figure B.7 shows the learned posterior duration distributions for the original upper bound ($D = 22$) and less restrictive upper bound ($D = 44$). The respective posterior distributions of both models look similar, indicating that the algorithm has converged to similar parameter values. Furthermore, almost all distributions have no weight after day 22. The only exception to these two observations is the duration distribution while *recovering on the ventilator*, which does change noticeably between $D = 22$ and $D = 44$.

We note two things about the duration distribution for recovering in stage **V**, which changes between $D = 22$ and $D = 44$. First, the estimated posterior here is very similar to the prior even for $D = 44$. Because the prior does put some weight after day 22 when $D = 44$, perhaps this the reason for the visible change in the posterior. Second, this distribution is the hardest to recover, since very few patients go on the ventilator, and, of those, only about 10% recover, according to our learned posterior in Figure 5. In our earlier experiments with synthetic data, we showed it is difficult to recover the true parameters related to recovering on the ventilator (Figure B.6).

Nevertheless, Table B.4 shows that ACED-HMM with a 44-day-per-segment upper bound achieves better MAE on **V** counts for Massachusetts data (MAE reduced from 34 to 27). We

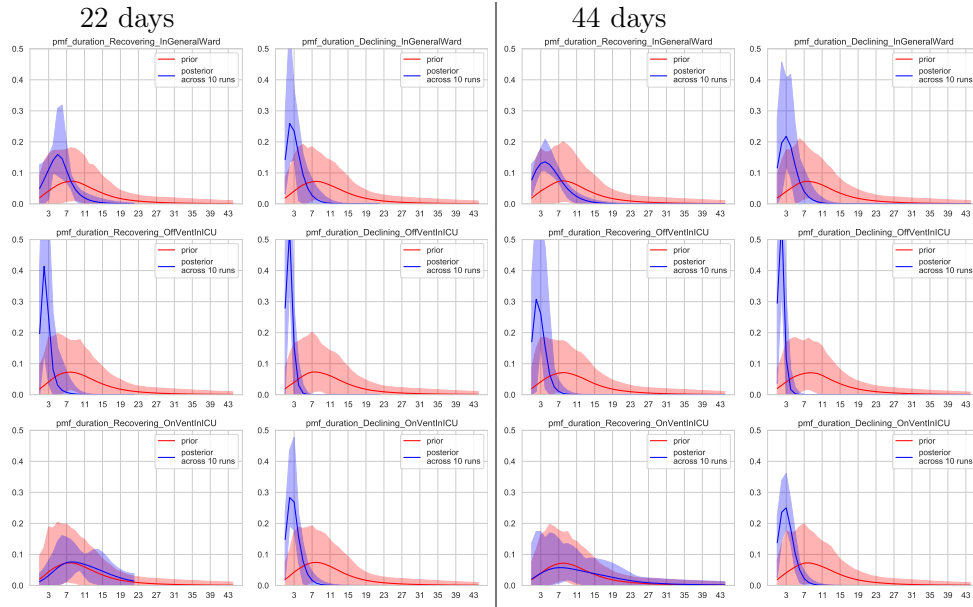


Figure B.7: **Comparison of posterior distributions of durations for ACED-HMM with different duration upper bounds.** The plots on the left are for ACED-HMM with durations capped at 22 days, whereas the plots on the right are for ACED-HMM with durations capped at 44 days.

	Method	G InGeneralWard			I + V			InICU			V OnVentInICU			T sm.		Death	
		MAE	lower	upper	MAE	lower	upper	MAE	lower	upper	MAE	lower	upper	MAE	lower	upper	
MA	ACED-HMM + ABC	65.0	61.5	68.9	15.7	14.5	16.9	34.0	32.6	35.7	8.0	7.8	8.3				
	ACED-HMM + ABC 44 days	68.1	63.2	75.7	15.8	14.2	17.6	27.3	24.2	30.6	7.7	7.3	8.1				
	ACED-HMM + ABC poisson	56.4	54.5	58.5	16.3	15.0	17.8	32.0	29.1	34.4	8.1	7.7	8.5				
	Mean Test y	1141.6			392.5			249.1			66.5						
UT	ACED-HMM + ABC	20.8	20.2	21.3	18.2	17.4	19.2	NA			2.4	2.3	2.5				
	ACED-HMM + ABC 44 days	20.7	19.8	21.4	19.6	18.1	21.4	NA			2.5	2.3	2.7				
	ACED-HMM + ABC poisson	21.0	20.2	21.8	18.4	16.8	19.7	NA			2.5	2.3	2.7				
	Mean Test y	272.9			164.3			NA			11.7						

Table B.4: Quantitative error assessment for predictions for two US states during the testing period (Jan. 11 - Feb. 11, 2021). We compare the performance of three versions of ACED-HMM: 1) using our duration parameterization with each state’s maximum length truncated to 22 days; 2) using our duration parameterization with each state’s maximum length truncated to 44 days; 3) using a poisson parameterization for the durations with each state’s maximum length truncated to 22 days. The method for computing the MAE is analogous to that of Table 1.

believe this improvement may be attributable to the duration distribution while recovering on the ventilator placing weight past day 22, even if, as we discussed above, this is difficult to learn from limited data and may be strongly influenced by the choice of prior. Given this improved MAE result, we recommend setting $D = 44$ or at least higher than 22 days.

Alternative Duration Model Experiment 2: Truncated Poisson. In the second experiment investigating modeling assumptions about durations, we parameterize each duration distribution as a *truncated Poisson*. This is equivalent to fixing the temperature parameter $\nu^{k,h}$ to 1.0 in Eq. (4). This restrictive choice results in a less flexible parameterization that is not able to generate distributions with varying entropy. However, this parameterization also reduces the number of trainable parameters from 17 to 11, thus reducing training time as well as potentially reducing the difficulty of posterior estimation, as there are less possible combinations of parameters with equally good fits to the data.

We fit each version of our model – fixing $\nu^{s,h} = 1$ and estimating its posterior – to the Massachusetts and Utah datasets and performed both a quantitative and a qualitative evaluation.

As a qualitative evaluation, we plot the learned posterior temperature distributions (in log scale), for both MA and UT in Figure B.8. For UT, we see that the learned posteriors over the log temperature parameter often place non-trivial mass outside a narrow interval around zero. Similar findings can be seen in the posteriors for MA. If mass concentrated near $\log_{10} \nu^{k,h} = 0.0$, then we could conclude that $\nu^{k,h} \approx 1.0$ and suggest the posterior was well-approximated by a truncated Poisson). The non-trivial variation around 0.0 in almost all stages suggests that we should continue using our 2-parameter flexible parameterization: at worst, it can match the truncated Poisson as a special case.

As a quantitative evaluation, Table B.4 shows a comparison in test mean absolute error (MAE) between ACED-HMM with our suggested parameterization (learned $\nu^{s,h}$) and with the simpler truncated Poisson parameterization ($\nu^{s,h} = 1.0 \quad \forall s, h$). The truncated Poisson parameterization yields slightly better MAE scores for **G** and **V** counts, and slightly worse ones for **I + V** counts. Our ABC procedure is able to find suitable sets of parameters that fit the data well, even with a less flexible parameterization like the truncated Poisson. We pointed out this issue of parameter identifiability (i.e. multiple distinct sets of parameters having similarly good fits to the data) in our experiment on synthetic data. While we do

not have strong empirical evidence in favor of our suggested 2-parameter formulation, there is also little reason to avoid it given that it is competitive.

Overall, as our purpose is not simply to fit to past data and forecast on future data, but also to learn meaningful posterior distributions, the decision between using our flexible parameterization vs. a poisson one cannot be reduced to a comparison over MAE. Instead, one should use clinical knowledge to determine whether the simplification provided by the Poisson assumption is close enough to the truth for all possible regions and sites. As we lack such knowledge, we recommend using our more flexible parameterization, which may provide better fits and more realistic posterior distributions for regions or sites in which the Poisson assumption is far from the truth.

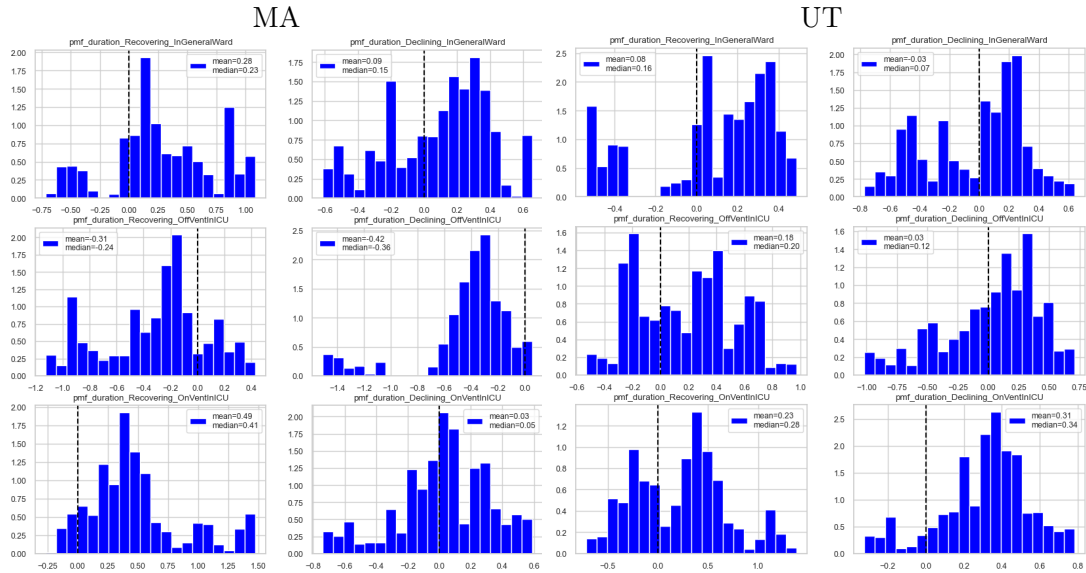


Figure B.8: Posterior distribution of the log temperature for durations in MA (left) and UT (right). For each state, we show samples from the posterior of the log temperature $\log_{10} \nu^{k,h}$ for all stages k and health states h .

Must Durations Be Truncated? Conceptually, there is little need to truncate each segment’s duration distribution. However, in practice truncation may improve efficiency (e.g. by constraining the memory required to represent each segment’s duration distribution). Our public code implementation does not easily allow for an unspecified upper bound, but it should be possible to alter the implementation to achieve this if desired.

Appendix C. Details of our proposed model: ACED-HMM

C.1. View of the ACED-HMM as a Hidden Markov Model

In the main paper, we discuss how our proposed model can be seen as an explicit duration hidden Markov model with 9 states: **S**, **G0**, **G1**, **I0**, **I1**, **V0**, **V1**, **T**, **R**. We provide the transition matrix for this model in Table C.1.

	next state							
	R	G1	I1	V1	G0	I0	V0	T
prev. state					$1 - \rho_G$			
S								
G1	1	ρ_G						
I1	r_I	$1 - r_I$						
V1	r_V		$1 - r_V$					
G0			$\rho_I(1 - d_G)$		$(1 - \rho_I)(1 - d_G)$			d_G
I0				$\rho_V(1 - d_I)$		$(1 - \rho_V)(1 - d_I)$		d_I
V0								1

Table C.1: State transition probabilities assumed by the explicit-duration HMM equivalent of our model: rows correspond to the previous state and columns indicate possible next states. Blank entries all have zero probability, but are kept blank for visual clarity. It is not possible to transition into the start state **S** or out of the terminal states **R** and **T**, so these rows/columns are omitted.

C.2. Prior over recovery probabilities

We first reproduce the relevant numbers from the CDC’s Table 2 (CDC, 2020) in Table C.2.

Age	18-49	50 - 64	65+	Avg.
Percent transfered to ICU	23.8	36.1	35.3	34.3
Percent who receive ventilation	12.0	22.1	21.1	20.4
Percent who die	2.4	10.0	26.6	19.3

Table C.2: Reproduced Table 2 of CDC (2020), accessed Mar. 19, 2021. We added the Avg. column.

We then describe how these estimates are turned into a prior over ρ . We first convert these estimates to single, age-independent estimates by computing the average across age groups weighted by the US country-level age distribution of hospital patient provided by the CDC ¹¹ during the week of November 7th. We show the computed average on the right-most column of Table C.2. Then, assuming fixed values for the probabilities of dying d and given these numbers, we derive the probability of recovering at each stage. Let $p(\mathbf{I})$ be the percent of patients transferred to the ICU, $p(\mathbf{V})$ the percentage of patients who receive ventilation, and $p(\mathbf{T})$ the percent of patients who die (all these numbers can be found in the table above). Then, we compute the mean of each parameter’s prior (denoted ρ_G , ρ_I and ρ_V for convenience) by solving for them in the following equations:

$$\begin{aligned}
 1 - p(\mathbf{I}) &= \rho_G + ((1 - \rho_G)d_G) \\
 p(\mathbf{V}) &= p(\mathbf{I})(1 - \rho_I)(1 - d_I)
 \end{aligned}$$

11. https://gis.cdc.gov/grasp/covidnet/COVID19_3.html

$$p(\mathbf{T}) = p(\mathbf{V})(1 - \rho_{\mathbf{V}}) + p(\mathbf{I})(1 - \rho_{\mathbf{I}})d_{\mathbf{I}} + d_{\mathbf{G}}$$

The found values for $\rho_{\mathbf{G}}$, $\rho_{\mathbf{I}}$ and $\rho_{\mathbf{V}}$ denote the ratios of parameters for the Beta priors. To set the variance for the prior, we set the beta parameters for the prior over each ρ by multiplying ρ and $1 - \rho$ by a scalar value r . For $\rho_{\mathbf{G}}$, we empirically set $r_{\mathbf{G}} = 100$. Then, we set $r_{\mathbf{I}} = r_{\mathbf{G}}(1 - \rho_{\mathbf{G}})$ and $r_{\mathbf{V}} = r_{\mathbf{I}}(1 - \rho_{\mathbf{I}})$. This increases the uncertainty in prior in proportion to the *influx* of patients in the different stages expected under the regime specified by the prior mean probabilities.

For both probabilities of death, we set $r = 200$ to indicate strong confidence in the low value that we set and discourage exploration towards unrealistic values.

Appendix D. Details of ABC Learning Procedure

Biassing the distance computation. For every dataset, we set weight u_k for each count type k such that counts that are generated at stages farther away from admissions are worth more, and such that their average is 1.0, so that the upper bound on the distance is preserved at 1.0. The rationale behind this is that counts generated farther away from admissions are generally harder to recover, as indicated by the greater uncertainty of ABC in predicting those counts that we see in our experiments.

For MA and SD, where $\mathbf{G}, \mathbf{I}, \mathbf{V}$ and \mathbf{T} -smoothed counts are available, we set $u_{\mathbf{G}} = 0.7$, $u_{\mathbf{I}} = 0.9$, $u_{\mathbf{V}} = 1.1$ and $u_{\mathbf{T}} = 1.3$. For UT and CA, where $\mathbf{G}, \mathbf{I} + \mathbf{V}$ and \mathbf{T} -smoothed counts are available, we set $u_{\mathbf{G}} = 0.8$, $u_{\mathbf{I} + \mathbf{V}} = 1.0$ and $u_{\mathbf{T}} = 1.2$. For the data ablation experiment on MA, we set $u_{\mathbf{G} + \mathbf{I} + \mathbf{V}} = 0.8$ and $u_{\mathbf{T}} = 1.2$. For both UK hospitals, we set $u_{\mathbf{R}} = 0.8$, $u_{\mathbf{G} + \mathbf{I} + \mathbf{V}} = 1.0$ and $u_{\mathbf{T}} = 1.2$. For the experiments on synthetic data, we set $u_{\mathbf{G}} = 0.8$, $u_{\mathbf{R}} = 0.9$, $u_{\mathbf{I}} = 1.0$, $u_{\mathbf{V}} = 1.1$ and $u_{\mathbf{T}} = 1.2$.

Justification of the maximum-normalized distance computation. At an individual timestep t , we assess distance via a mean absolute error normalized by the maximum between individual entries (which keeps this value on a unit scale). A side-effect of this normalization is that parameters that overestimate the true counts are preferred over parameters that underestimate the true counts by the same margin. For example, if true count $y_t^k = 20$ and simulated count $\tilde{y}_t^k = 22$, then the relevant distance will be $\frac{|20-22|}{22} = \frac{1}{11}$. Instead, if $\tilde{y}_t^k = 18$, the relevant unweighted portion of the distance will be $\frac{|20-18|}{20} = \frac{1}{10}$. This asymmetry enforces a pessimistic bias that in our experimental circumstances would be appropriate - e.g., when developing forecasts that will be used to determine what resources a hospital will need to meet future demand. In any case, the scheme at most moderately favors over prediction of utilization, compared to under utilization. Alternative distances could be easily considered.

Scheduling the decay. We draw samples from a *burn-in* phase and a *sampling* phase. During the *burn-in* phase, we anneal ε across iterations with the goals of 1) making it converge to an optimal value for sampling, and 2) dragging the parameters to the high-density region of the parameter space. Then, during the *sampling* phase, we stop the annealing, and draw samples in standard ABC MCMC as described in [Marjoram et al. \(2003\)](#), where ε is held fixed across iterations.

Crucially, in the sampling phase, we raise ε by a small value compared to the best value

found during the burn-in phase, so that we accept slightly more diverse samples than it would have otherwise.

In the *burn-in* phase, we anneal ε in the following way. The upper-bound on our distance d allows us to initialize ε to 1.0. In practice, we find that ε can often be safely initialized to a lower value, thus reducing the number of iterations needed for convergence. We anneal ε exponentially after each parameter proposal using hyperparameter γ . Crucially, we never allow ε to take a value below the last accepted distance d_{best} , thus allowing the algorithm to naturally converge to a value of ε that is optimal for sampling, as it is low enough that most proposals get rejected for not being able to provide a better match to the training data. We apply one more change to the annealing schedule: at regular intervals during the *burn-in* phase, we increase ε by a fixed value p . This allows the algorithm to “escape” a local optima that might have been encountered along the way. We found this to be particularly useful for fitting parameters to datasets containing fewer patients, such as single-hospital datasets. Indeed, with fewer patients to model, the variance in the counts generated by a given set of parameters is higher, thus it is more likely that a given set of proposed parameters gets accepted due to a particularly "lucky" simulation, which had assigned an unusually low distance to the proposed parameters.

Specific tuning of the decay schedule. In all our experiments, we initialize ε to 0.7, as we find such value to be always greater than the distance generated using samples from the prior. We fix the number of *burn-in* iterations to 24,000 (where one iteration makes and evaluates a proposal for each of the 17 parameters in turn, for a total of 408,000 total single-parameter iterations), p to 0.05, and f to 0.15. To set the annealing parameter γ , we consider the following. The optimal, convergence value of ε varies by dataset. In particular, due to the design of our distance function, datasets with higher counts have a lower convergence value. An approximate lower bound to the convergence value can be quickly found for any dataset by computing the average distance between multiple sets of counts generated with the dataset’s admissions and a fixed set of parameters, and a single representative of such counts. Though our annealing schedule is designed to be likely to escape local optima, we find it desirable, especially for lower-counts datasets, to set γ such that ε does not excessively undershoot the approximate lower bound by the end of the training iterations. Concretely, this translates to lower-counts datasets having higher γ , and thus slower annealing, in our experiments.

Details of proposal distributions. For each transition probability τ_k , we propose a new value τ_k^* by sampling from a Beta distribution whose mean is the old value and with a scaling parameter $r > 0$ that determines the variance: $\tau_k^* \sim \text{Beta}(r\tau_k, r(1 - \tau_k))$. We set $r = 100$ for the recovering probabilities and to $r = 200$ (lower entropy) for death probabilities.

For the mode of each duration probability $\lambda^{s,h}$, we propose a new value $\lambda^{*,s,h}$ by sampling from a truncated normal distribution with mean at $\lambda^{s,h}$ and a variance of 0.25: $\lambda^{*,s,h} \sim \text{TruncNorm}(\lambda^{s,h}, 0.25, [1, D])$. For the *log* of each temperature parameter $\nu^{s,h}$, we sample a new value from a normal distribution centered at the old value, with a variance of 0.01: $\nu^{*,s,h} \sim \text{Norm}(\nu^{s,h}, 0.01)$.

Warm start. At the start of the simulated training period, it is unrealistic to assume that the initial patients in each hospital stage have only just been admitted to said stage

(i.e. they are all at day zero of their stay). Thus, in an effort to more realistically model the progression through the hospital of the initial population, we apply a simple 'warm start' heuristic, in which we simulate the admission of the initial population in each state by uniformly distributing them during 5 days previous to day zero. We increase the admissions by 3% in both G and V to account for any patients who exit the hospital before the official start of the training period.

Details of ensembling. For all experiment, we collect 200 samples each from 10 different runs of the algorithm, resulting in 2000 total samples. We only use runs which we have verified to have converged to similar values of ε , and thus produce samples that explain the training counts equally well.




Figure D.1: **Trend of accepted distanced on one run of ABC for ACED-HMM on South Tees hospital data.** Each data point indicates a set of parametrs that has surpassed the first stage of acceptance (distance below ε). Around iteration 230,000 the algorithm got stuck in a local optimum, where few proposals were getting accepted. The resetting of ε after iteration 300,000 helped the algorithm get unstuck from the local optimum.

Appendix E. Details of Baseline AR-Poisson forecasting model

We consider a simple probabilistic model for univariate counts, the *latent autoregressive Poisson* or just AR-Poisson for short. This is inspired by previous work on single hospital site forecasting by Lee et al. (2021). The model can be fit separately to each univariate count time series of interest (e.g. separate fit to \mathbf{G} , \mathbf{I} , and \mathbf{V} counts).

Across T days we observe a univariate time series $y_{1:T} = [y_1, y_2, \dots, y_T]$, where $y_t \in \{0, 1, 2, \dots\}$ indicates the count of patients in a particular stage on day t . Our goal is to develop *forecasts* for the next F days, given all previous observations, using a conditional probabilistic model $p(y_{(T+1):(T+F)} \mid y_{1:T})$.

AR-Poisson Model. We consider an order-1 autoregressive process as our latent variable model, where each timestep has a latent real value $f_t \in \mathbb{R}$. We capture dependency across

time in the latent series $f_{1:T}$, and model each count y_t as conditionally independent given f_t :

$$p(f_{1:T}, y_{1:T}) = \prod_{t=1}^T p(f_t | f_{t-1}) \cdot \prod_{t=1}^T p(y_t | f_t). \quad (8)$$

For most timesteps $t > 1$, we generate the latent value f_t using the previous value f_{t-1} :

$$p(f_t | f_{t-1}) = \text{Normal}(\beta_0 + \beta_1 f_{t-1}, \sigma^2), \quad t \in 2, 3, \dots \quad (9)$$

The very first value f_1 is generated with mean β_0 :

$$p(f_1) = \text{Normal}(\beta_0, \sigma^2) \quad (10)$$

The latent-generating parameters are coefficients $\beta = [\beta_0, \beta_1]$ and standard deviation $\sigma > 0$.

We then generate each observed count y_t using a Poisson likelihood, setting the mean parameter by transforming the latent f_t to a positive value via the exponential:

$$p(y_t | f_t) = \text{Poisson}(\exp(f_t)). \quad (11)$$

Forecasting Method. We wish to first sample from the posterior over parameters β , σ , and latent values f : $p(\beta, \sigma, f_{1:T} | y_{1:T})$. We place the follow vague priors over the parameters:

$$\beta_0 \sim \text{Normal}(0, 0.1), \quad (12)$$

$$\beta_1 \sim \text{Normal}(1, 0.1), \quad (13)$$

$$\sigma \sim \text{HalfNormal}(0.1). \quad (14)$$

We then use the No-U-Turn sampler (Hoffman and Gelman, 2014) to perform Markov chain Monte Carlo approximation of the posterior. Our NUTS sampler is implemented using the PyMC3 toolbox (Salvatier et al., 2016).

Given a single posterior sample indexed by s , with parameters β^s, σ^s and latents $f_{1:T}^s$, we can then use the generative model to draw a *forecast* of latents f and counts y for the next F days:

$$f_{T+\tau}^s \sim \text{Normal}(\beta_0^s + \beta_1^s f_{T+\tau-1}^s, \sigma^s), \quad \tau \in 1, \dots, F \quad (15)$$

$$y_{T+\tau}^s \sim \text{Poisson}(\exp(f_{T+\tau}^s)), \quad \tau \in 1, \dots, F \quad (16)$$

We typically draw a forecast of S distinct samples, using $S = 1000$ or more to be sure we're capturing the full distribution.

Results. In Fig. E.1, we compare our AR-Poisson with our mechanistic ACED-HMM. Given the series of counts from the last 28 days of training data, we forecast ahead $F = 31$ days. We draw $S = 1000$ forecast samples and depict the 2.5th, 50th, and 97.5th percentiles.

We can see that our ACED-HMM model has clear advantages over the AR-Poisson, which is limited to capturing linear trends and makes the assumption that the future is like the past.




Figure E.1: Fit and forecasts on General Ward (G) and Total ICU (I + V) counts for MA, comparing our ACED-HMM to the AR-Poisson baseline. Compare to Fig. 3

Appendix F. Details of IHME Baseline forecasts for U.S. States

As a baseline, we consider the public forecasts produced by the IHME COVID-19 forecasting project (Reiner et al., 2021). These forecasts are available online¹² and contain predicted daily occupancy count values for each U.S. state for several stages of hospital care: total beds, ICU beds, ventilators, and deaths.

The core IHME model is targeted at forecasting *deaths* for each U. S. state, using an SEIR compartment model that accounts for a number of local factors. Historical daily death counts represent the primary training data used, though other observed information such as testing and mobility are included. Unlike our approach, the IHME model is not adapted to fit to the observed bed occupancy counts in hospitals in the region of interest. The methodology behind IHME’s forecasts for hospital resource use (total bed usage, ventilator usage, etc.) is described in Section 8 of the Supplementary Information (IHME COVID-19 Forecasting Team, 2020) released alongside their published research article on overall forecasting in the United States (Reiner et al., 2021).

To forecast hospital resource utilization, the IHME team performs a “microsimulation” that starts at deaths and works backwards. For each death, they assume the patient stayed the previous 6 days in the ICU. They further simulate for each death, a predicted age bin using local age distribution data, and simulate a number of individuals in the same age group as having also entering the hospital on the same day but surviving. Most of this “age-hospital-cohort” stay in the general ward (outside the ICU) for 8 days total. A small fraction of this “age-hospital-cohort” (6.3%) are assumed to be admitted to the ICU and stay for 20 days total (the middle 13 in the ICU). 85% of individuals in the ICU are assumed to need the ventilator. These parametric assumptions are reasonably informed by the literature from a specific location (New York state).

Compared to this approach, we argue that our approach of adapting a flexible hospital model with learnable transition probabilities and durations to a region of interest, rather than simply making fixed assumptions for all states, has substantial benefits.

Comparison of IHME forecasts to our methods. Several factors make it difficult to perform a fair head-to-head comparison between our model and the published IHME forecasts. First, the data signals used to train each method differ: while our method uses counts from all available stages, the IHME model relies primarily on daily deaths. Second, our forecasts require an admission count time series for the test period, while IHME does not. Finally, while our model is open source and can be customized to any training period and any region of interest provided the requisite data, we could not find a way to easily get produce customized forecasts from IHME, so we rely on their published numbers. Thus, any comparison between our methods and IHME should be more of a sanity check than a true comparison of equals. We expect that our method, because it is provided more detailed training data for the region of interest and especially given admissions for the testing period, should outperform the IHME baseline.

Acknowledging these challenges, we try to set up as fair a comparison as possible by taking the IHME forecasts released during our target testing period (Jan. 11 - Feb. 11, 2021), which are the ones dated 2021-01-15 on the IHME website. Manual inspection of

12. <https://www.healthdata.org/covid/data-downloads>

the CSV files released on this date indicates that these models had access to death counts until Jan. 11 but not afterwards. We suggest this because before this date values of the “deaths_mean” field in the released spreadsheet are whole numbers, after this date values are fractional indicating an expected value forecast rather than an actual observed count.

Raw forecast information from IHME website. We used the “reference” forecast released by the IHME for 2021-01-15, which provides daily count forecasts of hospital usage for each state of interest. IHME also releases best-case or worst-case projections (imagining different levels of mask usage or movement restrictions), but we did not use these, taking the “reference” forecasts as a reasonable projection of the status quo.

Each field in the released forecast has a “mean”, “lower”, and “upper” value indicating the expected value as well as lower and upper values to suggest an interval of uncertainty.

We used the following fields of IHME “reference” forecasts.

- admis(mean/lower/upper) : hospital admissions by day
- allbed(mean/lower/upper) : total covid beds needed by day
- ICUbed(mean/lower/upper) : ICU covid beds needed by day
- InvVen(mean/lower/upper) : invasive ventilation needed by day
- deaths(mean/lower/upper) : daily covid deaths

We used the PDF file `IHME_COVID_19_Data_Release_Information_Sheet.pdf` included in the ZIP file downloaded from IHME as a “data dictionary” to understand these field names and values.

Standardizing IHME forecasts our data format. We transformed the raw data daily counts to obtain values for our stages by performing the following element-wise subtractions/additions of the above raw data fields:

- General Ward **G** : allbed - ICUbed
- In ICU without ventilator **I**: ICUbed - InvVen
- In ICU on ventilator **V**: InvVen
- In ICU total **I + V**: ICUbed
- Terminal **T**: deaths

MAE computation. To obtain an MAE estimate from IHME at a specific stage (**G, I, V, T**) for a given state, we look in the testing period simply evaluate the MAE for the mean forecast of that stage compared to our ground-truth count values of that stage (as described in App. A).

Intervals To obtain some estimates of uncertainty, we compute the MAE between the true counts and the lower and upper estimates. This is all we can do given only mean/lower/upper values at each day (rather than a proper distribution).

We emphasize the IHME reported interval is not comparable to the intervals produced by our models. Our model's intervals are more easily interpreted in the standard way: they provide a range of plausible values for what the MAE would be given another similar-sized sample from the posterior.