

RESEARCH ARTICLE OPEN ACCESS

Estimating Mean Viral Load Trajectory From Intermittent Longitudinal Data and Unknown Time Origins

Yonatan Woodbridge^{1,2}  | Micha Mandel³ | Yair Goldberg⁴ | Amit Huppert^{1,5}

¹The Gertner Institute for Epidemiology & Health Policy Research, Sheba Medical Center, Ramat Gan, Israel | ²Department of Computer Science, Holon Institute of Technology, Holon, Israel | ³Department of Statistics and Data Science, The Hebrew University of Jerusalem, Jerusalem, Israel | ⁴Faculty of Industrial Engineering and Management, Technion - Israel Institute of Technology, Haifa, Israel | ⁵Department of Epidemiology and Preventive Medicine, School of Public Health, Faculty of Medical and Health Sciences, Tel-Aviv University, Tel-Aviv, Israel

Correspondence: Yonatan Woodbridge (yonatan.woodbridge@mail.huji.ac.il)

Received: 19 April 2024 | **Revised:** 26 January 2025 | **Accepted:** 7 February 2025

Funding: This work was supported by the Israel Science Foundation, Grant No. 3663/19.

Keywords: Ct-value | EM algorithm | multivariate normal distribution | SARS-Cov-2

ABSTRACT

Viral load (VL) in the respiratory tract is the leading proxy for assessing infectiousness potential. Understanding the dynamics of disease-related VL within the host is of great importance, as it helps to determine different policies and health recommendations. However, normally the VL is measured on individuals only once, in order to confirm infection, and furthermore, the infection date is unknown. It is therefore necessary to develop statistical approaches to estimate the typical VL trajectory. We show here that, under plausible parametric assumptions, two measures of VL on infected individuals can be used to accurately estimate the VL mean function. Specifically, we consider a discrete-time likelihood-based approach to modeling and estimating partial observed longitudinal samples. We study a multivariate normal model for a function of the VL that accounts for possible correlation between measurements within individuals. We derive an expectation-maximization (EM) algorithm which treats the unknown time origins and the missing measurements as latent variables. Our main motivation is the reconstruction of the daily mean VL, given measurements on patients whose VLs were measured multiple times on different days. Such data should and can be obtained at the beginning of a pandemic with the specific goal of estimating the VL dynamics. For demonstration purposes, the method is applied to SARS-Cov-2 cycle-threshold-value data collected in Israel.

1 | Introduction

The Viral load (VL) is the amount of viral nucleic acid within the host, expressed as the number of viral particles in a given volume. The typical measure unit is the cycle-threshold-value (Ct-value). In SARS-Cov-2, the typical Ct-value for a positive sample is between 15 and 40, representing the number of duplications required for the amount of viral genetic material to reach a certain detectable fluorescence threshold. It is inversely correlated to the VL, with a lower Ct-value indicating a higher

VL. Sampling is done by taking a nasal swab conducted in testing centers or medical clinics, after which the sample is transferred to one of several special labs for analysis.

VL typically increases exponentially after infection, until reaching a peak after which it starts to decline exponentially as a result of the host's immune response [1]. Understanding the VL trajectory is of great importance, as it affects the rate of infectiousness, generation time and disease duration [2]. Real-time estimates of quantitative viral shedding dynamics

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *Statistics in Medicine* published by John Wiley & Sons Ltd.

will enable better evidence-based public health interventions, such as lock-downs, length of quarantine, mask-wearing and other health-related policies. Most VL studies are based on longitudinal data, making them difficult and expensive to conduct [3–5]. A second difficulty stems from the fact that for most cases the exact day of infection is unknown or uncertain, due to the study design which follows participants only after diagnosis [3, 5]. As a practical solution, the time-origin can be defined as the study-entry time, which mostly coincides with symptom appearance. However, this approach is biased as the time-origin clearly precedes the symptoms' appearance.

Our main goal is to provide basic practical guidelines for data collection and VL curve reconstruction that can be used in future outbreaks. Specifically, the approach can be applied to a semi-controlled experiment where a sample of infected individuals is followed and provides a second VL measurement several days after the first positive test. Such data are logistically much simpler to collect compared to a full longitudinal study, yet, as we show here, can provide important information about the mean VL trajectory, even when the infection day is unknown.

Several studies addressed the problem of unknown time-origin in the context of HIV infections, since initial infection time is typically unknown (the time scale since infection until detection in HIV is months or years while for pandemics, such as SARS-Cov-2, it takes only several days). Berman [6] studied the T4 level trajectory from infection, using a stochastic process with exponential damping function, treating the time from infection to diagnosis as a latent variable. A similar approach was used in [7] to estimate the time from infection. Other models use empirical Bayes approaches [8], or biology-based dynamical models [9] to investigate the progression of various biomarkers from initial infection. A recent paper [10] proposed a likelihood-based estimation method for longitudinal trajectory estimation, which can be incorporated into survival models. This model can be applied to a wider range of datasets, as demonstrated on cervical dilation and medfly data.

In this paper, we study a discrete-time longitudinal trajectory estimation method. Similar to [10], our method is based on maximum likelihood estimation (MLE). However, here we consider the discrete-time scenario with only a few measurements given for each individual (we mainly consider the case of two measurements per individual). Our main goal is to make use of data collected in a designed experiment and estimate the mean population trajectory; that is, the mean Ct at each time point, if all individual trajectories were aligned up, starting from a common time-origin representing the infection time. We incorporate in our model a covariance matrix which accounts for possible within-individual correlation, as in longitudinal data analysis or in linear mixed-effects models [11]. Overall, the unknown parameters included in the likelihood function are the daily means, the covariance matrix, and the discrete distribution of the onset time. Using a Gaussian model for the measured values, we derive an EM algorithm for estimation, which treats the unknown time-origin and the unknown measurements as latent variables. Owing to the discrete-time approach and the multivariate normal model, the EM procedure mostly involves closed form calculations and standard quadratic function

minimizations. The model is flexible and various constraints can be imposed on the parameters, such as unimodality and structured covariance matrices. We show that under mild conditions, the model is identifiable in all of its parameters.

Our model formulation is motivated by the SARS-Cov-2 pandemic. In order to apply the approach on real data, we use a dataset on numerous PCR tests and VL measurements conducted by a major Israeli lab between January to June 2022. Many individuals had their VL measured twice or more on different days. We show that although the day of infection is unknown, the large number of samples enables reconstructing the population's VL trajectory using our model formulation and assumptions. We compare our results to known results in the literature.

The paper is organized as follows. Section 2 formulates the model, the missing data-generating process, and derives the likelihood function. Section 3 proves that under certain mild conditions, the parameters are identifiable. Section 4 develops an EM algorithm with and without structural constraints on the model's parameters. Section 5 reports results of numerical simulations, while Section 6 applies the method on SARS-Cov-2 data. Section 7 summarizes the paper and discusses limitations, and future directions.

2 | Model

Let x be an integer denoting the day following the time-origin, and let y_x be the outcome of interest as measured on day x . We assume a model of the form $y_x = \theta_x + \epsilon_x$, where θ_x ($x = 1, \dots, 2d - 1$) are fixed parameters, and $(\epsilon_1, \dots, \epsilon_{2d-1}) \sim F$ are zero-mean random variables which are possibly correlated; the use of vectors of length $2d - 1$ will be explained in the sequel. We concentrate on the multivariate normal case $F = \mathcal{N}(0, \Sigma)$, where $\Sigma \in \mathbb{R}^{2d-1 \times 2d-1}$ is a covariance matrix. The parameters $\theta = (\theta_1, \dots, \theta_{2d-1})$ and Σ are unknown, and should be estimated using incomplete data, as described below.

For n statistically independent samples $(y_{i,1}, \dots, y_{i,2d-1}) \sim \mathcal{N}(\theta, \Sigma)$, parameter estimation by maximum likelihood is straightforward. Suppose now that only $m_i \in \{1, \dots, d\}$ components from each vector are observed, but their indexing is unknown. Concretely, the data comprise of n vectors of observations $\{(\tilde{y}_{i,1}, \dots, \tilde{y}_{i,m_i})\}_{i=1}^n$ where $\tilde{y}_{i,j} = y_{i,x_{i,j}}$. However, instead of fully observing the indices $(x_{i,1}, \dots, x_{i,m_i})$, we assume that $x_{i,1} \in \{1, \dots, d\}$ and that only the difference between the indices $\Delta_{i,j}$ are available:

$$\Delta_{i,j} = x_{i,j+1} - x_{i,j}$$

for $j = 1, \dots, m_i - 1$. For our motivating example, the assumptions state that the first observation is during the “infection” period that lasts for about d days, but the exact day post infection is unknown, therefore only the differences between measurement times are exactly observed. If $m_i = 1$ we formally define $\Delta_{i,j} = 0$. Otherwise, since the $x_{i,j}$'s are integers, $\Delta_{i,j}$'s are also positive integers satisfying $\sum_j \Delta_{i,j} = x_{i,m_i} - x_{i,1} < d$. The goal is to estimate θ and Σ based on $\{(\tilde{y}_{i,1}, \dots, \tilde{y}_{i,m_i}, \Delta_{i,1}, \dots, \Delta_{i,m_i-1})\}_{i=1, \dots, n}$.

The model considers d possible time points during which the infection is “active”. However, in practice d is often unknown, and the latest observation time x_{i,m_i} might be larger than d . In fact, as $x_{i,1} = d$ and $x_{i,m_i} - x_{i,1} = d - 1$ are both possible, x_{i,m_i} can be as large as $2d - 1$. In typical settings, such as that of VL, it is reasonable to assume a “steady state” after time d , so we propose the following convention: The means and standard deviations from time d onward are all equal, that is, $\theta_j = \theta_d$ and $\sigma_j^2 = \sigma_d^2$ for $j = d, d + 1, \dots, 2d - 1$. The assumption of a complete viral clearance after d days is also used by [12] and serves as an approximation for small VL levels. As we will see in the numerical simulations, selecting a proper d is important for the estimate’s accuracy.

We illustrate the settings described above with a simple example. Consider three persons with daily Ct-values shown in the following table, starting in their infection day indexed as 1. Here $d = 5$, so that $2d - 1 = 9$.

day	1	2	3	4	5	6	7	8	9
Person 1	38	30	25	27	40	40	40	40	40
Person 2	36	29	28	30	35	40	40	40	40
Person 3	36	29	28	30	35	40	40	40	40

The observed Ct-values are marked with squares, so these are on days 3 and 5 post infection for the first subject, days 4 and 7 for the second subject, and days 1, 4 and 8 for the third subject. While we observe the time differences between samples, $\Delta_{i,j}$, their times relative to infection (relative to day 1) are unknown. For person 1, for instance, we observe a pair of Ct-values $(\tilde{y}_{1,1}, \tilde{y}_{1,2}) = (25, 40)$, and the time difference $\Delta_{1,1} = 2$ days between samples, but we do not know that these are on days 3 and 5 from infection. All $d = 5$ possible infection days should be considered as candidates for the first observed Ct-value.

To construct the likelihood, let $q_j = p(x_{i,1} = j)$ ($j = 1, \dots, d$) be the probability function of the first observed time $x_{i,1}$. Assume a non-informative selection mechanism in a sense that given $(x_{i,1}, \dots, x_{i,m_i})$, the observed data $(\tilde{y}_{i,1}, \dots, \tilde{y}_{i,m_i})$ follows the model $\tilde{y}_{i,j} = \theta_{x_{i,j}} + \epsilon_{x_{i,j}}$ as described above (i.e., there is no selection bias). While this assumption might be considered strong, it holds in a controlled experiment, which should be conducted during a pandemic in order to support policy decisions.

For ease of notation, we use a single index j to denote the variance parameters (i.e., the diagonal values of Σ), such that:

$$\text{diag}(\Sigma) = (\sigma_1^2, \dots, \sigma_{2d-1}^2) \quad (1)$$

For the off-diagonal covariance terms, we use two indices: $[\Sigma]_{k,\ell} = \sigma_{k,\ell}$. Given $x_{i,1} = j$, the vector $\tilde{\mathbf{y}}_i = (\tilde{y}_{i,1}, \dots, \tilde{y}_{i,m_i})$ has an m_i -variate normal distribution with mean $\theta_{i|j} = (\theta_j, \theta_{j+\Delta_{i,1}}, \dots, \theta_{j+\sum_{\ell=1}^{m_i-1} \Delta_{i,\ell}})$ and covariance

$$\Sigma_{i|j} = \begin{pmatrix} \sigma_j^2 & \cdots & \sigma_{j+\sum_{\ell=1}^{m_i-1} \Delta_{i,\ell}} \\ \vdots & \ddots & \vdots \\ \sigma_{j+\sum_{\ell=1}^{m_i-1} \Delta_{i,\ell}} & \cdots & \sigma_{j+\sum_{\ell=1}^{m_i-1} \Delta_{i,\ell}} \end{pmatrix} \quad (2)$$

As $x_{i,1}$ is not observed, the likelihood is a mixture of normal densities:

$$\mathcal{L}_m = \prod_{i=1}^n \sum_{j=1}^d \frac{q_j}{(2\pi)^{m_i/2} \sqrt{|\Sigma_{i|j}|}} \exp\left(-\frac{1}{2}(\tilde{\mathbf{y}}_i - \theta_{i|j})^T \Sigma_{i|j}^{-1} (\tilde{\mathbf{y}}_i - \theta_{i|j})\right) \quad (3)$$

Direct maximization of (3) is difficult. Treating $x_{i,1}$ and the unobserved components within $(y_{i,1}, \dots, y_{i,2d-1})$ as latent variables, we derive an EM algorithm [13] that maximizes the likelihood.

3 | Model Identification

In this section, we focus on the case where only pairs of samples $(\tilde{y}_{i,1}, \tilde{y}_{i,2})$ are given, together with the difference $\Delta_{i,1}$. We prove that under certain uniqueness conditions, the model is identifiable even for paired data. As a matter of convenience, we consider $i = 1$ and omit the indices, and use the notation $\mathcal{N}(\cdot; \cdot)$ to mark the marginal bivariate Normal p.d.f with given parameters. Thus, for a given Δ , the joint mixture distribution is

$$\mathcal{L}_{\theta, \Sigma, q}(\tilde{y}_1, \tilde{y}_2, \Delta) = \sum_{j=1}^d q_j \mathcal{N}(\tilde{y}_1, \tilde{y}_2; \theta_{1|j}, \Sigma_{1|j}) \quad (4)$$

The mean and covariance $\theta_{1|j}, \Sigma_{1|j}$ are defined in Section 2 (see (2) and the discussion above it). To prove model identifiability, we show that if $\mathcal{L}_{\theta, \Sigma, q}(\tilde{y}_1, \tilde{y}_2, \Delta) = \mathcal{L}_{\theta', \Sigma', q'}(\tilde{y}_1, \tilde{y}_2, \Delta)$ for all $\tilde{y}_1, \tilde{y}_2 \in \mathbb{R}$ and $\Delta \in \{1, \dots, d\}$, then $\theta = \theta'$, $\Sigma = \Sigma'$, and $q = q'$.

Without any constraints, the model is not identifiable. For example, if $q_1 = 0$ then clearly θ_1 and σ_1^2 are not identifiable. Similarly, if with probability 1 $\Delta \neq j$ for some $1 \leq j \leq d - 1$, then the $(k, k + j)$ th ($k = 1, \dots, d - j$) elements of Σ are not identifiable. However, the model can be partially identifiable if we assume the following reasonable conditions:

- **Condition 1.** The pairs (θ_j, σ_j^2) $j = 1, \dots, d$ are unique.
- **Condition 2.** The mixing probabilities q_j are positive for all $j = 1, \dots, d$.

The proof uses results on identifiability of finite mixture models [14, 15]. Specifically, [15] proved that an n -dimensional Gaussian mixture model (GMM) is identifiable up to permutations of indices. Taking the special 2-dimensional case, [15] shows that if a GMM p.d.f with d different Gaussian components is identical to another, such that

$$\sum_{j=1}^d \alpha_j \mathcal{N}(x; \mu_j, \mathbf{Q}_j) = \sum_{j=1}^d \alpha'_j \mathcal{N}(x; \mu'_j, \mathbf{Q}'_j), \quad \forall x \in \mathbb{R}^2$$

then a permutation π exists satisfying $\pi(\alpha'_1, \dots, \alpha'_d) = (\alpha_1, \dots, \alpha_d)$, $\pi(\mu'_1, \dots, \mu'_d) = (\mu_1, \dots, \mu_d)$ and $\pi(\mathbf{Q}'_1, \dots, \mathbf{Q}'_d) = (\mathbf{Q}_1, \dots, \mathbf{Q}_d)$.

Our model does not fall into the GMM framework, as only bivariate marginals of the $2d - 1$ dimensional normal distribution are observed. Nevertheless, we can use the identifiability result of GMM to prove identification in our setting.

Proposition 1. *Suppose that $P(\Delta = j) > 0$ for $j = 1, \dots, d - 1$. Under Conditions 1 and 2, the model (4) is identifiable in θ, Σ and q .*

Proof. For $\Delta = 1$, we observe the mixture

$$\sum_{j=1}^d q_j \mathcal{N}\left(\tilde{y}_1, \tilde{y}_2; \begin{pmatrix} \theta_j \\ \theta_{j+1} \end{pmatrix}, \begin{pmatrix} \sigma_j^2 & \sigma_{j,j+1} \\ \sigma_{j,j+1} & \sigma_{j+1}^2 \end{pmatrix}\right) \quad (5)$$

By identifiability of GMM models, the bivariate distributions in (5) and the mixing probabilities q_1, \dots, q_d are all identifiable up to an indices permutation. By the first condition, the pair (θ_1, σ_1^2) appears only in one of the mixtures, which identifies the first index. This also identifies q_1 , (θ_2, σ_2^2) , and $\sigma_{1,2}$ via the identification of $\mathcal{N}\left(\tilde{y}_1, \tilde{y}_2; \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \sigma_{1,2} \\ \sigma_{1,2} & \sigma_2^2 \end{pmatrix}\right)$. Once (θ_2, σ_2^2) is identified, (θ_3, σ_3^2) and $\sigma_{2,3}$ are identified via the identification of $\mathcal{N}\left(\tilde{y}_1, \tilde{y}_2; \begin{pmatrix} \theta_2 \\ \theta_3 \end{pmatrix}, \begin{pmatrix} \sigma_2^2 & \sigma_{2,3} \\ \sigma_{2,3} & \sigma_3^2 \end{pmatrix}\right)$. Continuing with the same reasoning shows identifiability of the mixing probabilities and all means, variances and single-lag correlations. Repeating the same arguments for $\Delta = k$ ($2 \leq k \leq d - 1$) establishes identification of the k -lag correlations. \square

For a single Ct-value, the model is not identifiable, as there is no information for estimating correlations. However, under Conditions 1 and 2 above, the means and variances (μ_i, σ_i^2) $i = 1, \dots, d$ are identifiable up to permutation of indices. While this is of limited use for most practical purposes, single observations can still identify interesting quantities such as the largest and smallest mean values.

4 | An EM Algorithm

4.1 | The Unconstrained Model

In terms of an EM algorithm, we define the complete likelihood to be of the time-origin $x_{i,1}$ and all measurements $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,2d-1})^T$ of individual i . Since $\mathbf{y}_i \sim \mathcal{N}(\theta, \Sigma)$ $i = 1, \dots, n$ are independent, the complete likelihood simply becomes:

$$\mathcal{L}_c = \prod_{i=1}^n \prod_{j=1}^d \left[\frac{q_j}{(2\pi)^{(2d-1)/2} \sqrt{|\Sigma|}} \times \exp\left(-\frac{1}{2}(\mathbf{y}_i - \theta)^T \Sigma^{-1} (\mathbf{y}_i - \theta)\right) \right]^{\mathbb{1}_{x_{i,1}=j}} \quad (6)$$

where $\mathbb{1}_{x_{i,1}=j}$ indicates the event $\{x_{i,1} = j\}$. The conditional expectation of $-\log \mathcal{L}_c$ is calculated in the E-step, which is then maximized over the unknown parameters in the M-step. These steps are repeated until convergence.

E-step. Denote by $\Theta^{(t)}$ the set of estimated parameters obtained at the t -th iteration, which contains the mean vector $\theta^{(t)}$, the covariance $\Sigma^{(t)}$, and the prior probability parameters $q^{(t)} = \{q_1^{(t)}, \dots, q_d^{(t)}\}$. For ease of notation, we use D_i to mark the observed data on subject i , that is, $\{D_i\}_{i=1, \dots, n} \equiv \{(\tilde{y}_{i,1}, \dots, \tilde{y}_{i,m_i}, \Delta_{i,1}, \dots, \Delta_{i,m_i-1})\}_{i=1, \dots, n}$, and the boldface letter $\mathbf{D} = \cup_i D_i$ to mark the entire observed data.

The negative log of (6), without constant terms, can be written as:

$$-\log(\mathcal{L}_c) = \sum_{i=1}^n \sum_{j=1}^d \mathbb{1}_{x_{i,1}=j} \times \left\{ \frac{1}{2} \log(|\Sigma|) + \frac{1}{2} (\mathbf{y}_i - \theta)^T \Sigma^{-1} (\mathbf{y}_i - \theta) - \log(q_j) \right\} \quad (7)$$

Computing the conditional expectation of (7) requires the calculation of $E_{\Theta^{(t)}}(1_{x_{i,1}=j} | \mathbf{D})$ and $E_{\Theta^{(t)}}(1_{x_{i,1}=j} f(\mathbf{y}_i) | \mathbf{D})$ for a linear and a quadratic function $f(\cdot)$. Denoting the former as $E_{ij}^{(t)}$, we have:

$$E_{ij}^{(t)} = p_{\Theta^{(t)}}(x_{i,1} = j | \mathbf{D}) = \frac{p_{\Theta^{(t)}}(D_i | x_{i,1} = j) q_j^{(t)}}{\sum_{k=1}^d p_{\Theta^{(t)}}(D_i | x_{i,1} = k) q_k^{(t)}}$$

The expectation $E_{ij}^{(t)}$ is readily obtained by recalling that $D_i | \{x_{i,1} = k\}$ has a normal distribution with mean $\theta_{i|k}$ and covariance matrix $\Sigma_{i|k}$; see (2) and the discussion above it.

For the second expectation, $\mathbb{E}_{\Theta^{(t)}}(1_{x_{i,1}=j} f(\mathbf{y}_i) | \mathbf{D})$, we have:

$$\mathbb{E}_{\Theta^{(t)}}(1_{x_{i,1}=j} f(\mathbf{y}_i) | \mathbf{D}) = \mathbb{E}_{\Theta^{(t)}}(f(\mathbf{y}_i) | x_{i,1} = j, D_i) \times E_{ij}^{(t)}$$

The expectation $\mathbb{E}_{\Theta^{(t)}}(f(\mathbf{y}_i) | x_{i,1} = j, D_i)$ involves calculation of the first two moments of \mathbf{y}_i conditionally on the event $\{D_i, x_{i,1} = j\}$, which can be obtained using known properties of multivariate normal distributions. The details are deferred to Appendix I, see Supporting Information.

M-step. The M-step should minimize for θ, Σ , and q_1, \dots, q_d the expression

$$-\mathbb{E}_{\Theta^{(t)}}(\log(\mathcal{L}_c) | \mathbf{D}) = \sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \left(\frac{1}{2} \log(|\Sigma|) - \log(q_j) + \frac{1}{2} \text{trace}(\Sigma^{-1} \mathbf{Y}_{\theta, i|j}^{(t)}) \right) \quad (8)$$

where

$$\mathbf{Y}_{\theta, i|j}^{(t)} = \theta \theta^T - \mathbf{y}_{i|j}^{(t)} \theta^T - \theta (\mathbf{y}_{i|j}^{(t)})^T + \mathbf{C}_{i|j}^{(t)} \quad (9)$$

and $\mathbf{y}_{i|j}^{(t)}$ and $\mathbf{C}_{i|j}^{(t)}$ are defined in Appendix I, see Supporting Information. Without any constraint, the minimum points $\boldsymbol{\theta}^{(t+1)}$ and $\boldsymbol{\Sigma}^{(t+1)}$ of (8) over $\boldsymbol{\theta}$ and $\boldsymbol{\Sigma}$ are similar to the sample mean and covariance. To see that, note that the ℓ -th partial derivative of (8) over $\boldsymbol{\theta}$, in a trace form, is:

$$\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \text{trace} \left(\boldsymbol{\Sigma}^{-1} \left(\boldsymbol{\theta} \mathbf{e}_\ell^T + \mathbf{e}_\ell \boldsymbol{\theta}^T - \mathbf{y}_{i|j}^{(t)} \mathbf{e}_\ell^T - \mathbf{e}_\ell (\mathbf{y}_{i|j}^{(t)})^T \right) \right)$$

where \mathbf{e}_ℓ is the ℓ -th unit basis vector. It can be shown that the gradient becomes zero at the point

$$\boldsymbol{\theta}^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \mathbf{y}_{i|j}^{(t)} \quad (10)$$

As for the covariance, the derivative of (8) over $\boldsymbol{\Sigma}^{-1}$ is:

$$\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \left(\boldsymbol{\Sigma} - \mathbf{Y}_{\boldsymbol{\theta}, i|j}^{(t)} \right)$$

which becomes zero at

$$\boldsymbol{\Sigma}^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \mathbf{Y}_{\boldsymbol{\theta}^{(t+1)}, i|j}^{(t)}$$

Lastly, the updated probability parameters are the sample proportions:

$$q_k^{(t+1)} = \frac{\sum_{i=1}^n E_{ik}^{(t)}}{\sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)}} \quad k = 1, \dots, d$$

4.2 | Model Constraints

Since the model has many unknown parameters, it is helpful to specify structural constraints that are based on prior knowledge and are relevant to the problem's domain. We discuss several constraints that are relevant to our motivation problem of VL reconstruction.

Means. First note from (8) and (9) that estimating $\boldsymbol{\theta}$ requires solving the following quadratic program:

$$\min_{\boldsymbol{\theta}} \left\{ n \boldsymbol{\theta}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\theta} - 2 \boldsymbol{\theta}^T \boldsymbol{\Sigma}^{-1} \left(\sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \mathbf{y}_{i|j}^{(t)} \right) \right\} \quad (11)$$

We have already made a structural constraint in the model formulation (Section 2), by assuming all means θ_j for $j \geq d$ are equal to θ_d . This reflects the fact that VL stabilizes after recovery. A natural assumption is a unimodal model, as the VL is expected to increase initially until reaching a maximal point, and then to decrease over time. Concretely, assuming a peak at time $d_{\max} \in \{1, \dots, d\}$, we impose the linear constraints $\theta_i \leq \theta_{i+1} \quad \forall i = 1, \dots, d_{\max} - 1$ and $\theta_i \geq \theta_{i+1} \quad \forall i = d_{\max}, \dots, 2d - 1$. If d_{\max} is unknown, we propose to run the EM-algorithm for every $d_{\max} = 1, \dots, d$ and select the estimate with the largest likelihood value.

To further reduce the number of parameters, a parametric model can be used [12] assumes a piece-wise constant model for the mean trajectory. While it is easy to adapt such a model here, we

prefer a smoother model and focus on a family of unimodal functions specified by three parameters $\alpha_1, \alpha_2, \alpha_3$:

$$\theta_k = \alpha_1 k^{\alpha_2 - 1} e^{-k/\alpha_3}, \quad \alpha_1, \alpha_2, \alpha_3 > 0; \quad k = 1, \dots, d \quad (12)$$

and $\theta_k = \theta_d$ for $k \geq d$. These unimodal functions, which we refer to as the unimodal Gamma (due to their similarity to the Gamma function), can smoothly capture the exponential growth, followed by an exponential decline rate, which characterizes most Ct-value and VL trajectories [4]. In that case, minimization of (11) can be carried out using a grid search over α_1, α_2 and α_3 .

Covariance. Within individual dependence should be taken into account when samples contain longitudinal measurements per individual. It is natural to assume a model of the form $\boldsymbol{\Sigma} = \sigma^2 \mathbf{R}$, where σ^2 is the variance and \mathbf{R} the correlation matrix. Various models for the covariance matrix have been suggested, see for example Chapter 7 of [16]. Since correlation tends to vanish as time difference increases, we choose to focus on a parsimonious first-order autoregressive model $\boldsymbol{\Sigma} = \sigma^2 \mathbf{R}(\rho)$ where $\text{Corr}(Y_{ij}, Y_{ik}) = \rho^{|k-j|}$. The M-step now involves the minimization of (8) over $\boldsymbol{\Sigma}$ under the constraint $\boldsymbol{\Sigma} = \sigma^2 \mathbf{R}(\rho)$, which leads to:

$$\min_{\sigma^2, \rho} (2d - 1)n \log \sigma^2 + n \log R(\rho) + \frac{1}{\sigma^2} \text{trace} \left(\mathbf{R}^{-1}(\rho) \left(\sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \mathbf{Y}_{\boldsymbol{\theta}, i|j}^{(t)} \right) \right)$$

For a fixed $\rho_0 \in (-1, 1)$, simple differentiation shows that

$$\sigma_0^2 = \frac{1}{(2d - 1)n} \text{trace} \left(\mathbf{R}^{-1}(\rho_0) \left(\sum_{i=1}^n \sum_{j=1}^d E_{ij}^{(t)} \mathbf{Y}_{\boldsymbol{\theta}, i|j}^{(t)} \right) \right)$$

minimizes the above function. A line search procedure can be applied in order to find the minimum point by calculating (ρ_0, σ_0) in a grid over $(-1, 1)$, and selecting the point that gives the smallest objective value.

Another family of covariance matrices of interest has the linear form

$$\boldsymbol{\Sigma}(\boldsymbol{\beta}) = \sum_{j=1}^J \beta_j \mathbf{B}_j \quad (13)$$

where \mathbf{B}_j ($j = 1, \dots, J$) are known symmetric matrices. An example is the heteroskedastic model in which $\text{Var}(Y_{ij}) = \sigma_j^2$. Here \mathbf{B}_j for $j = 1, \dots, d$ is a matrix with 1 in the j th entry of the diagonal and 0 in all other entries. If in addition $\text{Cov}(Y_{ij}, Y_{ik}) = \sigma_{|k-j|}$ depends on the lag between observations, matrices having 1 in entries $(j, j \pm k)$ and 0 otherwise can be added to (13). A fixed point procedure for estimating the coefficient vector $\boldsymbol{\beta}$ of a linear covariance model was proposed in [17]; we briefly describe the implementation of this algorithm to our problem in Appendix II, see Supporting Information.

4.3 | Technical Remarks

Since the Ct-value is inversely correlated to the VL, and it is upper bounded by 40, we use the negative linear transformation

$y_{ij} = 40 - Ct_{ij}$, where Ct_{ij} is the Ct-value of subject i measured on day j . The values are highly variable and, together with non-convexity of the likelihood, can lead to poor estimation. Reducing the number of parameters by using prior knowledge on the expected Ct curve, and specifically focusing on the family of unimodal functions (12), helps solving the problem to some extent.

To further deal with the non-convexity of the likelihood, we apply the following procedure. First, we run the algorithm several times, starting from a random realization of all parameters. Second, for each set of initial values, instead of directly estimating the parameters under model (12), we found it better to first estimate θ under a unimodal constraint, as discussed in Section 4.2, and use the results as initial values after fitting the function (12) to the estimated means. Specifically, we generated five sets of estimates of (θ, Σ, q) using the EM-algorithm under a unimodal constraint on θ assuming $d_{max} = m$ for $2 \leq m \leq 6$, which is its expected range, and then use the results as initial values for the EM algorithm to the unimodal Gamma model (12), choosing the estimate that gives the largest likelihood value. The final estimate is the one that maximizes the likelihood over all replications.

5 | Numerical Simulations

5.1 | Performance of the Method

Hay et al. [5] measured the Ct-values daily on a group of NBA players, providing a small sample of real-world Ct data. We use this dataset as a baseline for our simulation study. For our first simulation setting, we approximated the sample mean of the daily Ct-value by fitting (12). We constructed an AR(1) covariance matrix by setting $\sigma^2 = 10$, $\rho = 0.9$ and $\Sigma = \sigma^2 R(\rho)$; see discussion about covariance in Section 4.1. Lastly, we defined a vector of decreasing values q in the simplex. We set these parameters as the ground truth when generating the samples and calculating the performance of the method. The parameters θ and q that were used in the simulation are shown as the solid red lines in Figure 1a,c, respectively.

The choice of Δ is important for the performance of the estimator, and can be determined by policy makers in future pandemic-like scenarios. We assume only paired data ($m_i = 2$ for all i), as logistically it will be easier to design experiments in which infected

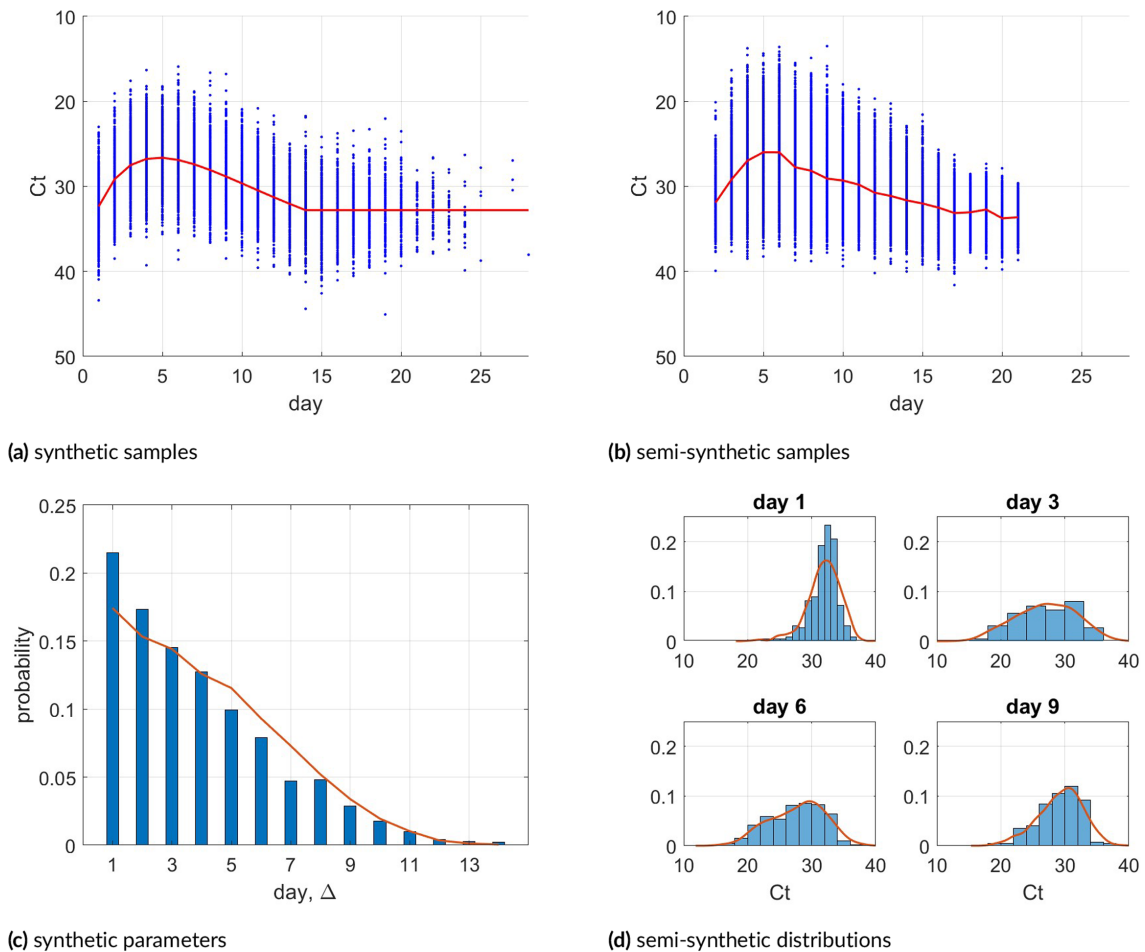


FIGURE 1 | Top row: Ct-value plotted against day after infection. Blue dots are samples generated synthetically (a) and semi-synthetically (b). Solid red lines represent the daily mean Ct-values. (c): Synthetic sample parameters: q (red line) and $p(\Delta)$ (histogram). (d): Histograms of the Ct-values as provided in [5] with the density estimation of the corresponding semi-synthetic Ct-value distributions (solid red lines), on days 1, 3, 6, and 9.

individuals are sampled only twice. We compare three distributions for the Δ values: Uniformly distributed over $\{1, \dots, 14\}$, uniformly distributed over $\{2, 3\}$, and a decreasing distribution as shown by the histogram in Figure 1c. To generate a sample, we randomly select integer $1 \leq x_{i,1} \leq 14$ using q , and sample Δ according to its distribution. We then generate $\tilde{y}_{i,1}, \tilde{y}_{i,2}$ from a bivariate normal distribution parameterized by the corresponding means in θ and the 2×2 sub-covariance of Σ . A scatter plot of the samples $(x_{i,1}, \tilde{y}_{i,1}), (x_{i,2}, \tilde{y}_{i,2})$ is shown in Figure 1a.

In order to better mimic the real world Ct-values and to study the performance of the estimator when the assumptions do not hold, we conducted an additional set of simulations, this time sampling directly from the data provided by [5]. As mentioned before, the viral load of each participant was monitored on a daily basis, so information on the day of infection for each infected participant is quite accurate, and the Ct-values on almost all days following infection are known. As the number of NBA players participating in the study was small, basing the simulation on the individuals' data was impractical. Instead, we sampled independent Ct-values for 20 days using the marginal empirical distributions of the Ct-value data and added a subject-specific $\mathcal{N}(0, 1)$ random number to the whole vector in order to add within individual correlations. To further deal with the small sample size, we added for each individual an additional noise for each coordinate using independent $\mathcal{N}(0, 1)$ random numbers. The empirical distributions of the Ct-values on days 1, 3, 6, and 9, as well as the corresponding marginal distributions used in the simulation, are shown in Figure 1d. After sampling the whole vector, we randomly select two entries from the distribution of Δ presented in Figure 1c. Figure 1b is the equivalence of Figure 1a presenting the mean and the sampled data for this scenario. Of course, when estimating the parameters, only the data $(\tilde{y}_{i,1}, \tilde{y}_{i,2}, \Delta_i = x_{i,2} - x_{i,1})$ are used.

We quantify the estimates' accuracy using the normalized mean-square error (NMSE):

$$\text{NMSE}(\hat{\theta}) = \|\hat{\theta} - \theta\|^2 / \|\theta\|^2$$

We found that the homoscedastic AR(1) structure for the covariance matrix does not hold. Indeed, simulations showed no consistency of the estimate when assuming that structure. We therefore chose to use the linear covariance model (13), where the main diagonal can have different values, thus accounting for heteroscedasticity. To account for within-individual correlation we set the two closest off-diagonal entries in Σ to be non-zero with a unique value for each off-diagonal ($\Sigma_{i,j} = \sigma_i^2$ if $i = j$, $\Sigma_{i,j} = c_1$ if $|i - j| = 1$, $\Sigma_{i,j} = c_2$ if $|i - j| = 2$, and $\Sigma_{i,j} = 0$ if $|i - j| \geq 3$).

The NMSE results versus the number of samples are shown in Figure 2. Each NMSE value is calculated as the average of 200 trials. Figure 2a–c show NMSE results for θ , (σ, ρ) and q respectively, in the three different settings of Δ . Figure 2d displays all estimates of θ from the 200 trials for $n = 100$ and $n = 1000$ (green lines). As expected, the variance is lower when $n = 1000$. The average of 200 estimates coincide with the true θ (solid lines), which suggests the estimate is approximately unbiased. Estimation of q is more challenging as the values have a smaller scale (between 0 and 1). Interestingly, the estimates of all parameters perform better when $\Delta \sim \text{Unif}(1, \dots, 14)$.

Figure 3a shows the semi-synthetic samples' NMSE results for all three parameters. The NMSE results were averaged over 100 trials, a sufficient number to produce smooth curves in this case. Here Σ has a linear covariance structure, so the NMSE is computed via the Frobenius norm. As previously, θ estimate is more accurate in terms of NMSE compared to q , which is more difficult to estimate. Figure 3b shows 200 estimates of θ together with their average and true θ , for $n = 100$ and $n = 1000$.

5.2 | Choice of d

The latest day d that limits the day after infection on which the first measurement was taken is somewhat arbitrary, reflecting a day on which the VL is low and individuals can be regarded as "recovered". In this part, we examine how the choice of d affects the estimate's accuracy. We use the same setting as in the first simulation model, using a uniform distribution of Δ over $\{1, \dots, 14\}$. Generating 100 samples in each trial, we perform parameter estimation assuming various d values: $d = 7, 10, 14$, and 20, where $d = 14$ is the ground-truth.

Panel (a) of Figure 4 shows the average results of the four estimates of θ compared to the ground-truth under model (12), and panel (b) shows the same for a unimodal-constrained model. The estimate coincides with the true θ in the early days, while towards the end the estimates differ. In general, the model is more sensitive to d values smaller than the true d . The value on which the VL is maximal, which is an important parameter, seems quite stable to the value of d , especially under the unimodal constraint.

6 | Viral Load Data

To lower the risk of sampling bias, the approach should be applied to controlled data, where Δ is determined by the researcher. While such data are not available to us, we apply the method to purely observational data and compare the results to that of [5].

The data contains Nucleocapsid gene (N -gene) Ct-values, measured by a major lab in Israel, on swab samples taken from patients who were tested positive for SARS-Cov-2. The samples were taken between January and June, 2022, during which the SARS-Cov-2 infections were most likely caused by the Omicron variant. The data contains records on patients whose Ct-value was measured once, twice, or multiple times on different days. Specifically, out of 222 668 records, about 97% records contain a single Ct-measurement, while about 2.7%, 0.26%, and 0.04% respectively contain pairs, triples, and quadruples of Ct-values measured on different days. These amount to over 6000 pairs, 580 triples, and 89 quadruples of Ct-values.

Since the Ct-values are documented only on patients who were tested positive, the $\hat{\theta}$ estimate should be interpreted as the daily mean Ct-value of all infected persons; more formally, the expectation of the Ct-value conditioned on Ct-value less than 40. This is different from the mean Ct-value trajectory which includes deceased and recovered persons. The latter is usually the focus of viral load studies, such as [3, 5]. Nevertheless, because the probability of recovery (i.e., for the Ct-value to reach 40) is low on the early days following infection, these two trajectories most

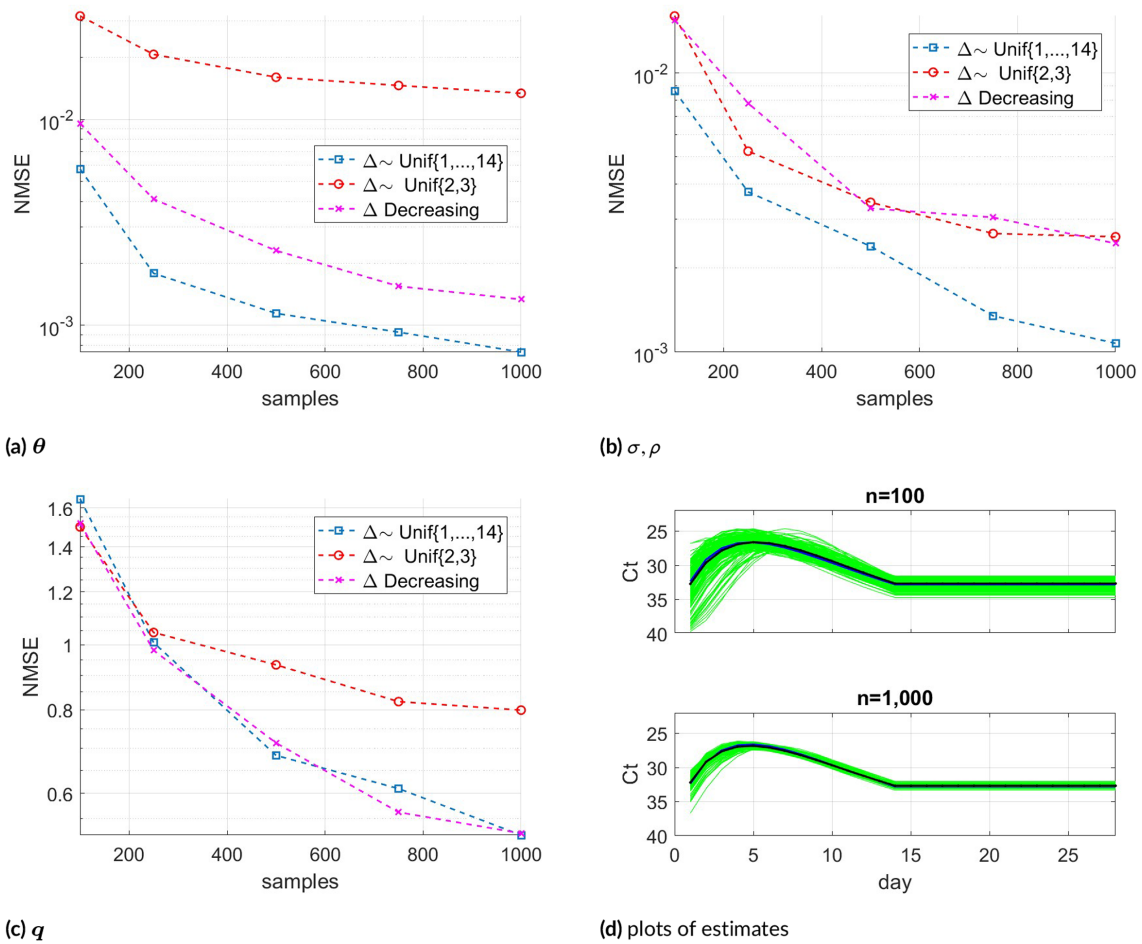


FIGURE 2 | Synthetic sample NMSE for the following estimates: (a) θ , (b) σ, ρ in Σ , (c) q , and (d): Plots of θ estimates of 200 trials (green lines) for $n = 100$ and $n = 1000$, together with their averages and true θ (solid lines).

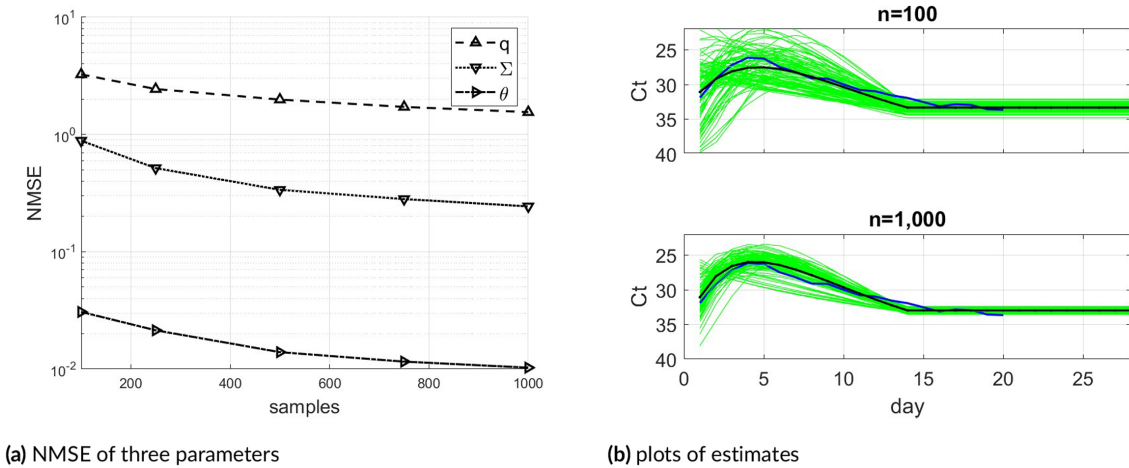
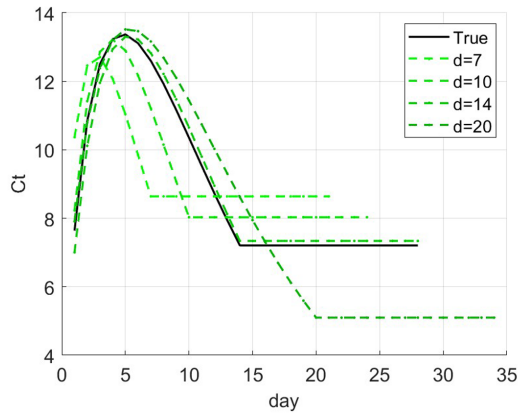


FIGURE 3 | (a): Semi-synthetic sample NMSE for the following θ, σ, ρ in Σ and q . (b): Plots of θ estimates of 200 trials (green lines) for $n = 100$ and $n = 1000$, together with their averages and true θ (solid lines).

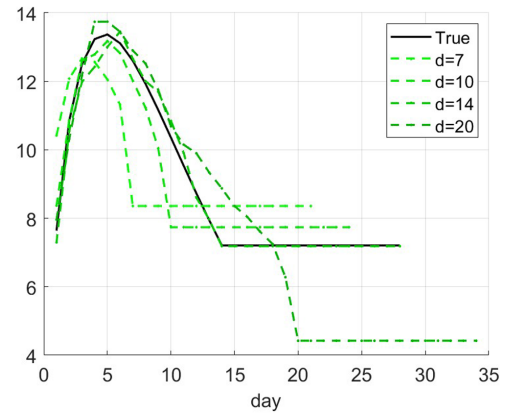
likely coincide on these days. In any case, the last day d should be chosen carefully and the estimate should be interpreted accordingly. Regarding deaths, during our follow-up period the proportion of deceased patients among all confirmed cases is negligible (0.09%, [18]). Thus, we believe that excluding the

Ct-value documents of deceased patients will introduce a very little bias when considering the population Ct-value.

Figure 5a shows a histogram of the first Ct-value of individuals whose Ct-value was measured multiple times on different days

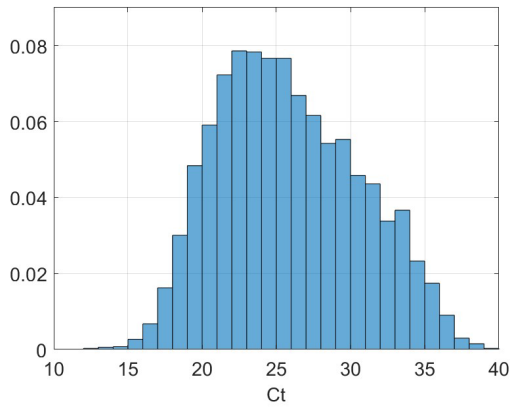


(a) Gamma

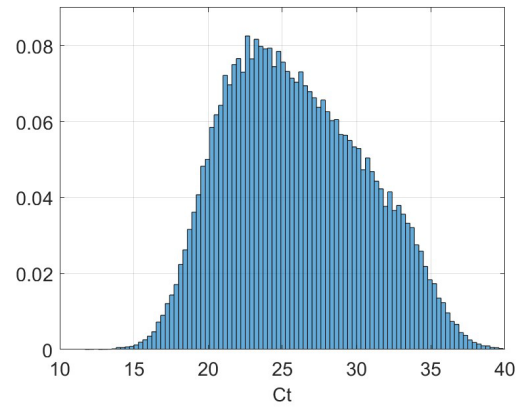


(b) Unimodal

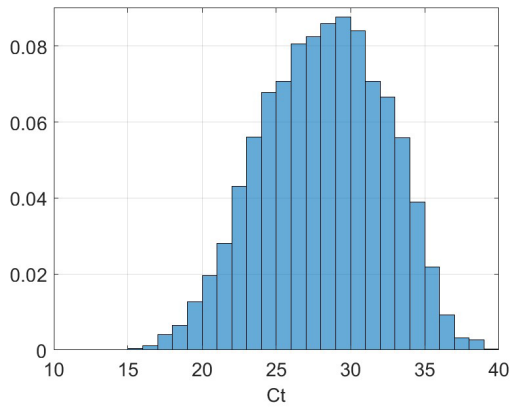
FIGURE 4 | (a): Plots of θ estimates for $d = 7, 10, 14$, and 20 assuming the unimodal Gamma function in (12). (b): Plots of θ estimates assuming a general unimodal function.



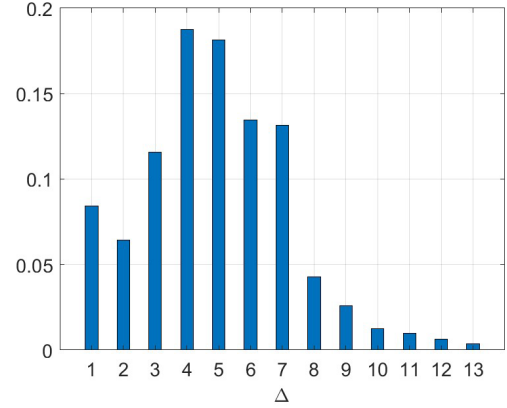
(a) $n = 6,653$



(b) $n = 216,015$



(c) $n = 6,645$



(d) Δ

FIGURE 5 | Top row: Distribution of first Ct-value measurement among: Persons who were tested multiple times (a), and persons who were tested once (b). Number of samples is given at the bottom. (c): Distribution of second Ct-value measurement. (d): Histogram of time difference Δ in days between the first and second measurement.

(6653 samples). This histogram is similar to that of the Ct-value of individuals who were measured only once (Figure 5b, with 216015 samples). We also notice that although samples with Ct-value of 40 are not included, the tails decrease smoothly, without a sudden truncation at Ct-value of 40. Figure 5c shows a

histogram of the second Ct-value of individuals whose Ct-value was measured multiple times. The distribution is shifted more to the right, showing an average increase in Ct-value between the tests. As a larger Ct-value represents a smaller VL, this suggests that the second measurements were typically taken after

the infection's peak. Figure 5d shows the distribution of Δ —the difference in days between the first and second measurements (samples with $\Delta > 13$ were excluded; the number of such samples is negligible accounting for 1% of all samples). Most of the Δ values are between 1 and 7 days, which constitute about 90% of all samples.

We included all pairs, triples, and quadruples of Ct-values in estimation. We used the structural assumptions set in the semi-synthetic simulation, namely the Gamma parametric class of θ and the linear covariance structure (due to possible heteroscedasticity). Although the maximum value of Δ is 13, the maximum number of days after infection, d , that best fits the data is unknown. A natural choice is $d = 14$; however, we also examine $d = 7, 10$ and 20 . For each d , we follow the same procedure as described previously, running the EM algorithm several times using different random initializations, and selecting the one that gives the largest likelihood value. Since the number of samples is quite large, we use 1000 iterations as convergence might possibly be slower.

Figure 6a shows the daily Ct-value estimation results for $d = 7, 10, 14$, and 20 . Since $d = 10$ gives the largest likelihood value we regard this as the estimate of choice and it is shown as a black solid line. The curve shows that the viral load reaches its peak during day 4 after infection (with $\text{Ct} \approx 23$), a result that is consistent with current knowledge [5]. The Ct-value on day 10 reaches a value of 30, and it represents the average Ct-value on day 10 and onward. Panel (b) of Figure 6 shows the daily Ct-value estimates under unimodality constraints for various d , which are quite similar to the parametric Gamma curves shown in (a).

7 | Discussion

In the midst of the COVID-19 pandemic, crucial policy decisions have been made under the pressure of limited evidence. One such example is determining the appropriate length of quarantine for infected individuals – a decision that carries substantial economic and ethical consequences. In order to ensure

responsible and effective policies, it is imperative to base these decisions on rapid and accurate information. Estimating viral load (VL) trajectories – a key factor in understanding infectiousness – traditionally requires resource-intensive and complex longitudinal studies. These studies are not only time-consuming but often involve substantial costs and logistical challenges. These led to delays in obtaining critical data needed for informed policy adjustments.

The method presented here requires much less resource (performing a second measure for a sample of infected individuals), is logistically much simpler than traditional methods, and can provide fast and important estimates for the parameters of interest. Importantly, the method's reliance on routine surveillance data makes it more adaptable and feasible for rapid implementation. In the face of emerging variants of concern, the ability to promptly evaluate data and adjust policies becomes even more vital. The current approach represents a significant advancement, offering the promise of faster decision-making processes while maintaining scientific rigor. Although the method is based on strong modeling assumptions, it can provide good initial estimates that can be updated when more complete data are collected. From VL trajectory estimates one can derive other parameters of special interest to epidemiologists, such as the duration of viral shedding, given some VL threshold.

Using available observations data is often subject to selection bias that may affect the statistical results. As a referee commented, individuals who die, especially those who die immediately after infection, are under-represented. We recommend conducting a semi-controlled design in which a second test is performed to a sample of individuals who were found positive to the virus. The EM algorithm can deal with individuals who are included in the sample, but do not show up for their scheduled second test by using $m_i = 1$ for these individuals. While this does not solve the selection bias problem, it may somewhat reduce it and can be used for sensitivity analysis under various assumptions regarding the expected VL trajectory for these partially-observed individuals. For a pandemic with a low mortality rate, such as COVID-19, this bias is expected to be small.

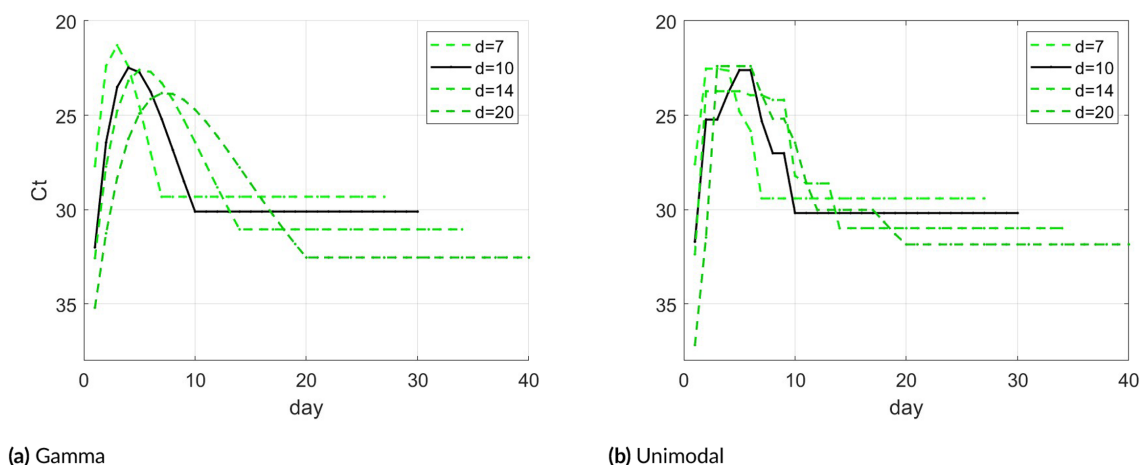


FIGURE 6 | (a): Daily Ct-value estimates for $d = 7, 10, 14$, and 20 based on the parametric model (12). (b): Daily Ct-value estimates for $d = 7, 10, 14$, and 20 based on a unimodal constraint.

In the synthetic simulation, we examined how the distribution of the time difference Δ between pairs of samples affects the accuracy of the estimate. We found that the uniformly distributed Δ gave a stable and accurate estimate, which suggests using this approach when planning the design. A detailed analysis for finding the optimal Δ distribution is a possible future research direction. Another important issue is the case where the samples greatly depart from the normal distribution. This includes, for instance, the bimodal distribution that can be relevant if recovered persons, whose VL becomes a fixed zero, are included in the samples.

Software

Matlab code for the EM-algorithm and numerical simulations can be accessed at <https://github.com/yonatan123/VL-trajectory>.

Acknowledgments

The authors would like to thank Arnona Ziv for providing the data.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Due to Israel Ministry of Health's regulations, individual Ct-level data cannot be publicly shared.

References

1. S. M. Kissler, J. R. Fauver, C. Mack, et al., "Viral Dynamics of Acute SARS-CoV-2 Infection and Applications to Diagnostic and Public Health Strategies," *PLoS Biology* 19, no. 7 (2021): e3001333.
2. M. Marks, P. Millat-Martinez, D. Ouchi, et al., "Transmission of COVID-19 in 282 Clusters in Catalonia, Spain: A Cohort Study," *Lancet Infectious Diseases* 21, no. 5 (2021): 629–636.
3. P. Y. Chia, S. W. X. Ong, C. J. Chiew, et al., "Virological and Serological Kinetics of SARS-CoV-2 Delta Variant Vaccine Breakthrough Infections: A Multicentre Cohort Study," *Clinical Microbiology and Infection* 28, no. 4 (2022): 612.
4. R. Ke, P. P. Martinez, R. L. Smith, et al., "Daily Longitudinal Sampling of SARS-CoV-2 Infection Reveals Substantial Heterogeneity in Infectiousness," *Nature Microbiology* 7, no. 5 (2022): 640–652.
5. J. A. Hay, S. M. Kissler, J. R. Fauver, et al., "Quantifying the Impact of Immune History and Variant on SARS-CoV-2 Viral Kinetics and Infection Rebound: A Retrospective Cohort Study," *eLife* 11 (2022): e81849.
6. S. M. Berman, "A Stochastic Model for the Distribution of HIV Latency Time Based on T4 Counts," *Biometrika* 77, no. 4 (1990): 733–741.
7. N. Dubin, S. Berman, M. Marmor, B. Tindall, D. D. Jarlais, and M. Kim, "Estimation of Time Since Infection Using Longitudinal Disease-Marker Data," *Statistics in Medicine* 13, no. 3 (1994): 231–244.
8. V. Degruittola, N. Lange, and U. Dafni, "Modeling the Progression of HIV Infection," *Journal of the American Statistical Association* 86, no. 415 (1991): 569–577.
9. J. Drylewicz, J. Guedj, D. Commenges, and R. Thiébaud, "Modeling the Dynamics of Biomarkers During Primary HIV Infection Taking Into Account the Uncertainty of Infection Date," *Annals of Applied Statistics* 4, no. 4 (2010): 1847–1870.

10. T. Wang, S. J. Ratcliffe, and W. Guo, "Time-To-Event Analysis With Unknown Time Origins via Longitudinal Biomarker Registration," *Journal of the American Statistical Association* 118, no. 543 (2022): 1–16.
11. R. H. Myers, D. C. Montgomery, G. G. Vining, and T. J. Robinson, *Generalized Linear Models: With Applications in Engineering and the Sciences* (John Wiley & Sons, 2012).
12. H. C. S. Karita, T. Q. Dong, C. Johnston, et al., "Trajectory of Viral RNA Load Among Persons With Incident SARS-CoV-2 G614 Infection (Wuhan Strain) in Association With COVID-19 Symptom Onset and Severity," *JAMA Network Open* 5, no. 1 (2022): e2142796.
13. A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum Likelihood From Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society: Series B* 39, no. 1 (1977): 1–22.
14. H. Teicher, "Identifiability of Finite Mixtures," *Annals of Mathematical Statistics* 34, no. 4 (1963): 1265–1269.
15. S. J. Yakowitz and J. D. Spragins, "On the Identifiability of Finite Mixtures," *Annals of Mathematical Statistics* 39, no. 1 (1968): 209–214.
16. G. M. Fitzmaurice, N. M. Laird, and J. H. Ware, *Applied Longitudinal Analysis* (John Wiley & Sons, 2012).
17. T. W. Anderson, "Asymptotically Efficient Estimation of Covariance Matrices With Linear Structure," *Annals of Statistics* 1, no. 1 (1973): 135–141.
18. Israel Ministry of Health Dashboard (Corona), <https://datadashboard.health.gov.il/portal/dashboard/health>.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.