

Control-Polygon Epidemiology: A Bernstein-Bezier Framework for Epidemic Reconstruction, Forecasting, and Control

Jomar Rabajante

jfrabajante@up.edu.ph

University of the Philippines Los Banos <https://orcid.org/0000-0002-0655-0893>

Method Article

Keywords: Bernstein polynomial, Bezier curve, effective reproduction number, epidemic forecasting, inverse control, decision support

Posted Date: October 6th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11261513/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Control-Polygon Epidemiology: A Bernstein–Bézier Framework for Epidemic Reconstruction, Forecasting, and Control

Jomar F. Rabajante^{*1}

¹Institute of Mathematical Sciences, University of the Philippines Los Baños,
Laguna, Philippines

Abstract

Many data-driven biological systems must be reconstructed from observations, forecast under uncertainty, and steered toward desired outcomes, yet these tasks are often assigned different model representations. In computational epidemiology, this fragmentation obscures parameter meaning and makes model outputs harder to compare and communicate. We introduce ‘control-polygon epidemiology’, a unified Bernstein–Bézier framework in which an effective reproduction-number trajectory is encoded by a small set of interpretable control points. Historical polygons are estimated by constrained least squares or Bayesian negative-binomial renewal inference; forecasts attach new, smoothly joined Bernstein segments instead of extrapolating the historical polynomial; and inverse control maps a desired epidemic trajectory back to feasible intervention intensity, implementation speed, uncertainty, and cost. The representation yields global convex-hull threshold certificates, finite-difference momentum and curvature, coefficient-level uncertainty propagation, spectral calibration to mechanistic models, and structural lower bounds: C transversal threshold crossings require at least $C + 1$ control points, while E isolated interior extrema require at least $E + 2$. Synthetic experiments span renewal, SIR, SEIR, SIRS, vector-borne, age-structured, vaccination, isolation/hospitalization, and overdispersed observation models. In Philippine COVID-19 rolling-origin tests, a history-only peak-aware continuation improved on the original joined-Bernstein forecast in 87.5% of origin–horizon comparisons and on a constant- R_t baseline in 75%, while stress tests correctly exposed the impossibility of forecasting genuinely unobserved structural shocks from case history alone. The framework provides a low-dimensional bridge between statistical learning, mechanistic transmission models, probabilistic forecasting, and interactive decision science.

Keywords: Bernstein polynomial; Bézier curve; effective reproduction number; epidemic forecasting; inverse control; decision support

^{*}Corresponding author: jfrabajante@up.edu.ph

Author summary

Epidemic decisions often begin with three questions: what is happening now, what may happen next, and what trajectory do we want instead? We represent all three with the same small set of Bernstein–Bézier control points. Historical control points are inferred from surveillance data and summarize the effective reproduction number. Future control points encode a transparent forecast continuation and can incorporate external information such as an emerging variant. For intervention planning, decision makers can move target epidemic or R_t points in an interactive display; a mechanistic model then translates the chosen target into required intervention intensity, feasibility, expected health outcomes, and cost. The graphical handles have mathematical meaning: their levels bound the complete trajectory, their differences determine direction and curvature, and their number has structural lower bounds determined by the number of threshold crossings or extrema. The framework does not replace causal intervention-effect estimates or economic data; instead, it provides a common, interpretable interface through which those inputs can be used consistently.

1 Introduction

The effective or instantaneous reproduction number, R_t , is a central summary of epidemic transmission and underpins extensive work on real-time estimation, uncertainty quantification, and decision support [1–6]. In renewal models, expected incidence is current transmissibility multiplied by infectiousness generated by recent cases, so R_t is a natural latent forcing function for both retrospective inference and forecasting [3, 7, 8]. Yet reconstruction, forecasting, and intervention design are typically represented with different parameterizations.

Flexible smooth models of incidence or R_t include sliding-window estimators, penalized splines, Bayesian P-splines, generalized additive models, Gaussian processes, and state-space or semi-mechanistic formulations [3, 9–13]. These methods are powerful, but flexible coefficients are not always easy to interpret globally. Forecasting adds a further difficulty: a fitted historical curve does not specify what should happen after the forecast boundary, and structural shocks such as variant replacement can invalidate purely history-based continuation [14–16].

Bernstein polynomials and their Bézier interpretation provide positivity, partition of unity, convex-hull containment, finite-difference derivatives, and variation-diminishing properties [17–19]. Bernstein and Bézier methods are not new to epidemics: they have been used as numerical bases for stochastic SIR systems [20], to parameterize epidemic optimal controls [21], to represent time-varying transmission [22], and to fit segmented epidemic curves [23]. Likewise, peak-aware and growth-geometry forecasting have clear precedents [24, 25]. Our claim is therefore not novelty of the basis itself.

We propose ‘control-polygon epidemiology’: the same epidemiologically interpretable Bernstein polygon is used for three linked tasks. First, reconstruction estimates the historical R_t polygon from incidence using constrained least squares (LS) or Bayesian renewal inference. Second, forecasting attaches a new future polygon, with optional growth- and event-informed adaptation, rather than extrapolating the historical polynomial outside its domain. Third,

inverse control starts from a desired epidemic or R_t trajectory and solves backward for intervention intensity, feasibility, and cost. This last direction is intentionally compatible with interactive high-level decision making: a modeller or policymaker can move target control points and immediately see the mechanistically implied intervention and uncertainty.

The framework contributes five main ideas. (i) It unifies reconstruction, forecasting, and inverse control in one representation. (ii) It makes interpretability mathematical through convex-hull threshold certificates, finite-difference momentum/curvature, and coefficient-level uncertainty propagation. (iii) It gives structural lower bounds on the minimum number of control points needed to express specified threshold crossings or extrema. (iv) It provides both LS and Bayesian implementations for reconstruction and forecasting, and parallel point/posterior formulations for inverse control. (v) It connects the same polygon to renewal, compartmental, vector-borne, age-structured, vaccination, and hospitalization/isolation mechanisms. Detailed proofs, derivations, full worked examples, extended sensitivity analyses, and the Shiny user manual are provided in the Supplementary Material.

2 Framework

2.1 Bernstein representation and renewal reconstruction

For $u = (t - t_0)/T \in [0, 1]$, define

$$B_{j,n}(u) = \binom{n}{j} u^j (1 - u)^{n-j}, \quad j = 0, \dots, n, \quad (1)$$

and represent the latent reproduction-number trajectory as

$$R(u) = \sum_{j=0}^n \rho_j B_{j,n}(u), \quad K = n + 1. \quad (2)$$

The coefficients ρ_j are the epidemiological control points. For incidence Y_t and generation-interval weights w_s ,

$$\Lambda_t = \sum_{s=1}^S w_s Y_{t-s}, \quad \mu_t = R_t \Lambda_t, \quad (3)$$

with an overdispersed observation model

$$Y_t \mid \mu_t, \phi \sim \text{NB}(\mu_t, \phi), \quad \text{Var}(Y_t \mid \mu_t) = \mu_t + \mu_t^2 / \phi. \quad (4)$$

Two complementary inferential engines are used. Given empirical renewal ratios $\tilde{R}_t$, constrained LS estimates

$$\hat{\boldsymbol{\rho}} = \arg \min_{\boldsymbol{\rho} > 0} \|W^{1/2}(\tilde{\mathbf{R}} - B\boldsymbol{\rho})\|_2^2 + \lambda \|D_2 \boldsymbol{\rho}\|_2^2, \quad (5)$$

with optional shape constraints. The Bayesian formulation uses positive coefficients $\rho_j = \exp(\theta_j)$, a smoothness prior on first or second differences of θ_j , and the negative-binomial

renewal likelihood. Posterior draws yield credible bands for R_t and predictive intervals for incidence. The LS implementation is fast and transparent; the Bayesian version propagates parameter and observation uncertainty. Full priors, bootstrap procedures, and posterior diagnostics are in Supplementary Sections S1–S3. For the table of symbols and notation, refer to Table 3.

2.2 Theoretical guarantees and minimum complexity

The usefulness of Eq. (2) comes from properties that hold for the complete curve, not only at the displayed handles. Table 1 summarizes the results; proofs are given in Supplementary Section S2.

Table 1: Main geometric guarantees of the epidemiological control polygon.

Property	Epidemiological consequence
Convex hull	$\min_j \rho_j \leq R(u) \leq \max_j \rho_j$. If every control point is below one, the complete fitted R_t trajectory is below one; analogously above one.
First differences	$R'(t) = \frac{n}{T} \sum_{j=0}^{n-1} (\rho_{j+1} - \rho_j) B_{j,n-1}(u)$, so neighboring differences encode transmission momentum.
Second differences	The sign and magnitude of $\Delta^2 \rho_j$ encode curvature/acceleration and give sufficient convexity/concavity certificates.
Threshold complexity	If the curve has C transversal crossings of $R_t = 1$, then the number of control points $K \geq C + 1$.
Extremum complexity	If it has E isolated interior extrema, then the number of control points $K \geq E + 2$.
Uncertainty propagation	Coefficient-wise lower/upper limits propagate to pointwise lower/upper curves because Bernstein weights are nonnegative and sum to one.
Spectral calibration	If a next-generation operator is $\mathcal{K}(t; \beta) = \beta(t)A(t)$, choosing $\beta(t) = R(t)/\rho(A(t))$ reproduces the Bernstein target exactly whenever scalar calibration is valid.

The crossing and extremum results are structural lower bounds, not automatic statistical choices. We use them to define a smallest admissible search region for K ; predictive validation, residual adequacy, and Bayesian predictive criteria then determine whether larger K is supported. Figure 1 illustrates why the coefficients are more than generic smoothing parameters: their vertical positions bound the complete R_t curve, while their first differences determine the direction and magnitude of transmission momentum.

2.3 Forecasting with a joined future polygon

Directly evaluating the historical polynomial beyond $u = 1$ loses the convex-hull guarantee and is not our forecast. Instead, define a future segment over horizon H ,

$$R_F(v) = \sum_{j=0}^m q_j B_{j,m}(v), \quad v \in [0, 1]. \quad (6)$$




Figure 1: Interpretability of the Bernstein control polygon. Control-point levels bound the full R_t curve, while neighboring differences determine its derivative.

Continuity sets $q_0 = \rho_n$. Matching the derivative at the boundary gives

$$q_1 = q_0 + \frac{nH}{mT_H}(\rho_n - \rho_{n-1}), \quad (7)$$

where T_H is the historical time scale used for the derivative. Remaining future points are generated from a declared continuation process or event model.

When a wave signal is already visible at the origin, long-history smoothing can delay the forecast. We therefore estimate recent log-growth from only pre-origin observations, convert it to a short-horizon reproduction-number anchor through the generation interval, and compress the resulting continuation into joined Bernstein segments. Forecast priorities can be declared through

$$D_{\mathbf{w}} = w_c d_{\text{curve}} + w_t \frac{|\hat{t}_p - t_p|}{\tau_t} + w_h \left| \frac{\hat{Y}_p - Y_p}{Y_p} \right| + w_d \frac{|\hat{t}_{\text{onset}} - t_{\text{onset}}|}{\tau_d}, \quad (8)$$

with balanced, peak-priority, timing-priority, or user-defined weights. Parameters are selected only on earlier rolling origins. If external surveillance supplies an expected peak window or event magnitude, it is included explicitly as a scenario assumption rather than hidden future-data tuning.

For a newly detected variant with frequency $p(t)$ and effective transmission multiplier a , one simple event map is

$$q_j^{(V)} = q_j^{(0)} \{1 + (a - 1)p_j\}. \quad (9)$$

Bayesian forecasting draws historical and future control points, event parameters, and dispersion, then propagates them through the renewal model. Outputs include predictive incidence intervals and probabilities such as

$$P \left(\max_{1 \leq h \leq H} R_{T+h} > 1 \mid Y, \mathcal{E} \right). \quad (10)$$

2.4 Inverse control and interactive policy design

Forecasting asks what may happen; inverse control asks what intervention is required to obtain a chosen target. Represent a standardized intervention trajectory by

$$u(t) = \sum_{j=0}^n a_j B_{j,n}(t/T), \quad 0 \leq u(t) \leq 1. \quad (11)$$

Under a multiplicative transmission-reduction mechanism,

$$R_t = R_t^{\text{free}}[1 - \varepsilon u(t)], \quad u^*(t) = \frac{1}{\varepsilon} \left(1 - \frac{R_t^*}{R_t^{\text{free}}} \right), \quad (12)$$

subject to $0 \leq u^*(t) \leq u_{\text{max}}$. If $R_t^{\text{free}} = R_0$ is constant and target and control use the same degree,

$$a_j = \frac{1 - \rho_j^*/R_0}{\varepsilon}. \quad (13)$$

For vaccination, isolation, hospitalization, vector control, or age-structured contact interventions, the same target is mapped through the appropriate state, removal, contact, or next-generation mechanism rather than treating all policies as equivalent scalar suppression.

The policy polygon has its own exact geometry:

$$\frac{1}{T} \int_0^T u(t) dt = \frac{1}{n+1} \sum_{j=0}^n a_j, \quad (14)$$

so mean control-point height equals mean intervention intensity, while

$$\sup_t |\dot{u}(t)| \leq \frac{n}{T} \max_j |a_{j+1} - a_j| \quad (15)$$

bounds the implementation ramp rate. Quadratic burden and smoothness costs are exact quadratic forms in the coefficient vector. Thus an interactive app can let decision makers move a small number of target points while the model updates feasibility, intervention intensity, health outcomes, and costs in real time. LS produces a point control conditional on estimated parameters; Bayesian propagation produces distributions for required control, cost, and constraint satisfaction. Detailed derivations and intervention-specific mappings are in Supplementary Section S4.

3 Results

3.1 Structural complexity and reconstruction

The first result is that the theoretical complexity floor is recoverable from data when the signal-to-noise ratio is adequate. A controlled theorem-verification experiment constructed trajectories with one, two, and four transversal crossings of $R_t = 1$. Their degree-minimal representations required $K = 2, 3, 5$, exactly the lower bound $K = C + 1$. With 120 low-noise

Table 2: Examples of intervention mechanisms in the control layer.

Intervention	Mechanistic entry	Interpretation of the control
Lockdown/contact NPIs	contact rate or contact-matrix elements	reduction in contact opportunities; not automatically a literal closure percentage
Masking/ventilation	per-contact transmission probability	reduction in effective transmission conditional on uptake/effectiveness
Testing/isolation	additional infectious removal rate	shorter effective infectious duration
Vaccination	protected state and/or susceptibility/infectiousness modifiers	cumulative state-dependent effect; cannot generally be inverted independently at each time
Hospitalization/isolation	removal plus residual infectiousness and capacity constraints	transmission effect only when admission isolates infectious individuals; occupancy is also a system constraint
Vector control	biting, competence, recruitment, or mortality	mechanism-specific reduction in vector-mediated transmission

replicates per scenario, the first adequate fitted model recovered that floor in every replicate. Figure 2 shows that adequacy first appears at the theorem-implied K , rather than at an arbitrary model size selected after inspection.

The robustness experiment also separates a mathematical requirement from statistical detectability. For one- and two-crossing trajectories, the lower bound was recovered in 100% of replicates across all tested noise levels. For four crossings, recovery remained 98% at noise SD 0.05 but fell to 76% at SD 0.08, even though the theorem itself remained valid. Thus the lower bound can be used as a principled starting point for model selection, while data quality determines whether all structurally possible features are empirically resolvable (Supplementary Figure S3 and Supplementary Table S3).

The second reconstruction experiment tested whether the control polygon can be inferred with uncertainty rather than only drawn from known coefficients. In the worked negative-binomial renewal experiment, historical R_t RMSE was 0.0310 for constrained LS and 0.0272 for the Bayesian posterior median. The LS bootstrap covered the true historical curve on 94.3% of days and the Bayesian 95% credible band on 100% in this realization. The Bayesian sampler also recovered the observation dispersion ($\tilde{\phi} = 32.5$ for true $\phi = 35$), with maximum classical $\hat{R} = 1.0017$ and minimum approximate effective sample size above 2,096. Figure 3 shows that the two engines recover the same broad transmission geometry while providing complementary uncertainty summaries. These are worked-example diagnostics, not universal frequentist coverage guarantees or evidence that Bayesian inference must dominate LS.

3.2 Event-informed and peak-aware forecasting

The joined future polygon avoids a common ambiguity of smooth historical fits: the historical Bernstein polynomial is never evaluated outside its fitted domain. Instead, a new segment





Figure 2: Controlled verification of the minimum-control-point result. In the low-noise experiment, the first adequate K coincides with the structural threshold-crossing lower bound.



(a) Historical LS/Bayesian reconstruction.

(b) Bayesian future R_t forecast.

Figure 3: Parallel LS and Bayesian inference in the synthetic renewal experiment.

carries explicit assumptions about continuation, growth, or a detected event. In 100 stochastic futures per variant multiplier, event information had little value when the effective multiplier was one but increasing value as the future departed from the no-event continuation. At a 42-day horizon, mean WIS for event-informed versus no-event forecasts was 20.20 versus 20.31 at multiplier 1.0, 26.51 versus 35.33 at 1.4, and 34.35 versus 99.77 at 1.8. At multiplier 1.6 and a 60-day horizon, mean WIS was 54.75 for the event-informed forecast versus 142.09 for the no-event continuation. Figure 4 therefore demonstrates a useful separation: the Bernstein geometry does not invent an event, but it provides a controlled, interpretable channel through which external biological or surveillance information can alter the future trajectory.

The held-out synthetic forecast also illustrates uncertainty propagation. Future R_t RMSE was 0.116 for the LS point forecast and 0.0693 for the Bayesian posterior median; the Bayesian 95% R_t credible band covered all held-out latent values in this realization, and the 95% incidence posterior predictive interval covered 98.3% of held-out days. The posterior probability of any future $R_t > 1$ was 0.977. Supplementary Figures S7–S9 show the historical posterior, future R_t distribution, and time-varying resurgence probability in detail.

A different forecasting failure occurs when a wave is already accelerating at the forecast origin but a long-history smoother reacts too slowly. The peak-aware continuation addresses this by using only pre-origin recent growth to anchor the first future segment and by allowing peak timing, height, onset delay, and whole-curve error to receive different weights. This mechanism is evaluated on historical Philippine origins below and in Supplementary Figures S17–S21; its role is not to foresee unobserved shocks, but to reduce avoidable phase lag once a new growth regime is already visible.




Figure 4: Future control polygons separate history-based continuation from externally informed structural change.

3.3 Mechanistic portability and inverse control

The portability experiment deliberately held the same Bernstein R_t target fixed while changing the biological mechanism underneath it. Six mechanistic systems—SEIR, SIRS with waning immunity, a Ross–Macdonald-type human–vector model, a three-group age-structured model, vaccination with breakthrough infection, and SEIHR hospitalization/isolation with residual infectiousness—all reproduced the target next-generation reproduction number to

machine precision (maximum absolute calibration error $< 5 \times 10^{-16}$). A seventh example added negative-binomial observation overdispersion without altering the latent transmission geometry. These examples are listed and derived in Supplementary Section S5, shown in Supplementary Figures S10–S11, and summarized in Supplementary Table S5.

The examples are intentionally heterogeneous because equal R_t trajectories need not imply equal biological states. The SEIR model generated a peak of 4.64 daily exposures per 100,000; the vector-borne and age-structured examples generated peak aggregate incidences of 9.83 and 8.34 per 100,000, respectively; the vaccination example reached 37.7% in the vaccinated-susceptible state; and the hospitalization/isolation example reached 34.0 isolated/hospitalized individuals per 100,000. Thus the control polygon can serve as a common transmission-level coordinate system while latent-stage delays, waning immunity, host structure, protection, isolation, and observation noise remain mechanistically distinct. The result is portability of the representation, not identifiability of those mechanisms from aggregate R_t alone.

Inverse-control experiments then reversed the direction of calculation: rather than asking what epidemic follows from a policy, they asked what policy polygon is needed to track a chosen epidemic target. Normalized target-tracking errors were 5.6×10^{-5} (SIR), 5.2×10^{-5} (SIRS), 2.1×10^{-4} (SEIR), and 3.6×10^{-4} (vector-borne); the age-structured multilever optimization achieved aggregate- R_t RMSE 5.9×10^{-5} . In the SIR illustration, the uncontrolled prevalence peak exceeded 215,000 infectious individuals whereas the target peak was approximately 1,354. Figure 5 shows the intended meeting-room workflow: decision makers manipulate a small target polygon on a readable scale, while a separate full-scale panel preserves the magnitude of the uncontrolled epidemic and the model returns the intervention polygon required to achieve the target.

The supplementary control suite makes the same point under different mechanisms. Supplementary Figure S13 gives SIR, SIRS, SEIR, and vector-borne inverse tracking; Supplementary Figure S14 solves a multilever age-structured problem across school, work, and community contacts; and Supplementary Figure S15 converts policy intensity into a health–burden frontier. In the illustrative SIR cost scan, no added control produced 875,261 cumulative infections, whereas the reference inverse-control polygon produced 13,762, averting approximately 861,499 infections while having mean standardized control 0.567. For the illustrative burden weight $\lambda = 0.15$, this reference policy was near the minimum of the scanned normalized loss. These numbers are synthetic and dimensionless on the policy side; their purpose is to show that epidemiological targets, implementation constraints, and local cost functions can be optimized in the same coefficient space once real application-specific costs are supplied.

3.4 Philippine COVID-19 illustration

We analyzed a reporting-date series with 1,178 daily observations from 30 January 2020 through 21 April 2023 and 4,087,964 reported cases. The series is used only as an aggregated methodological illustration; reported cases are not equated with infections. Context on Philippine epidemic waves and response is consistent with Bansilan and Rabajante [26].

The historical reconstruction used $K = 36$ and achieved log-NRMSE 0.0287 against the 7-day reporting-date case series. The extended analyses in Supplementary Section S8 show the reconstructed R_t , generation-interval/ K sensitivity, terminal pseudo-forecasting, nine-origin




Figure 5: Interactive inverse-control concept. A target epidemic/control polygon is manipulated directly; the mechanistic model returns the implied intervention and a separate full-scale panel preserves the magnitude of the uncontrolled epidemic without covering the target-scale plot.

baseline evaluation, historical-origin stress tests, peak-aware tuning, and the Omicron event scenarios (Supplementary Figures S16–S23). These analyses are deliberately broader than a single favorable forecast because they expose where the representation succeeds and where information is absent.

A stringent 2020-only origin illustrates that boundary: none of the cases-only continuations could anticipate the April 2021 wave before its driver was visible. Once acceleration was present, however, peak-aware continuation materially improved the forecast. With a 20 March 2021 origin, the original joined Bernstein forecast missed the April peak by 31 days and had 180% peak-height error, whereas the peak-aware cases-only continuation reduced the timing error to 7 days and the height error to 4.9% (Supplementary Figure S18). After selecting history-only hyperparameters on three earlier pre-Omicron origins and freezing them, the adaptive method was evaluated on six later origins: it had lower MAE than the original joined-Bernstein forecast in 87.5% of origin–horizon comparisons and lower MAE than a constant- R_t baseline in 75%. Median improvements were clearest over 28–60 days, while the constant- R_t baseline remained competitive at 14 days.

The end-2021 stress test provides the complementary lesson. A cases-only continuation underestimated the unprecedented Omicron surge, so the paper does not present history-only geometry as a substitute for variant surveillance. Once an external high-transmission event assumption is declared, the future polygon can represent that information explicitly. In the peak-priority Omicron illustration, both the LS point trajectory and scenario-centered Bayesian central trajectory peak on 18 January 2022 at approximately 35,193 cases/day versus an observed 34,808, or about 1.1% peak-height error. Figure 6 places that conditional scenario beside the original epidemic curve and separates it visually from cases-only continuations. The near match is therefore a demonstration of event-conditioned representation and decision targeting, not prospective proof that cases alone predicted Omicron.

Finally, the Philippine example closes the reconstruction–forecast–control loop. Figure 7 shows LS/Bayesian reconstruction, LS/Bayesian event-conditioned forecasting, and a hypothetical inverse-control calculation asking what additional effective suppression would be required to keep $R_t \leq 0.90$. The maximum additional suppression required under the illustrated target was 0.808 and the mean Bayesian requirement was 0.219; an illustrative minimum-burden decomposition distributed that requirement across contact NPIs, testing/isolation, and vaccination-derived protection. These quantities are standardized decision indices rather than estimates of the causal effects of historical Philippine policies. The empirical example therefore demonstrates how one observed dataset can move through all three uses of the framework without conflating inference, conditional forecasting, and policy counterfactuals.

4 Discussion

Control-polygon epidemiology is best viewed as a bridge between flexible statistical smoothing, mechanistic epidemic models, and decision-facing control. Its principal advantage is not that Bernstein polynomials fit curves that other bases cannot fit; rather, the same low-dimensional coefficients admit exact global interpretation, a principled forecast continuation, and an




Figure 6: Philippine Omicron-period forecast comparison. Forecast curves are plotted against the original observed epicurve and a separate full-scale panel is used where needed. Event-conditioned forecasts are distinguished from cases-only continuations.




Figure 7: Philippine COVID-19 demonstration of the three goals: LS/Bayesian reconstruction, LS/Bayesian forecasting, and hypothetical LS/Bayesian inverse control. Companion full-scale panels are used where necessary so extreme trajectories do not obscure the decision-relevant scale.

inverse mapping from target trajectories to policy controls.

The framework also makes its limitations explicit. R_t remains a latent effective quantity and can be sensitive to the generation interval, reporting delay, ascertainment, and model specification [5, 8]. Intervention controls are only mechanistic when the dose-response mapping is specified and calibrated; a generic 30% effective suppression cannot automatically be interpreted as a 30% lockdown. Likewise, monetary cost-effectiveness requires local economic inputs. Forecast geometry cannot infer an unobserved future shock. The Philippine 2020-only and cases-only Omicron experiments are retained precisely to demonstrate this information boundary.

Statistically, the Bayesian reference implementation is useful for uncertainty propagation but should be stress-tested for posterior geometry and identifiability in application-specific settings. The companion Shiny app uses a faster Laplace/posterior-draw approximation for interaction; high-stakes final analyses should confirm non-Gaussian posteriors with full MCMC when necessary. Operational forecast validation should also use many origins, nested model selection, reporting-delay adjustments, strong baselines, and prospective evaluation [27, 28].

There are several natural extensions: hierarchical polygons across regions, multi-strain control polygons, event covariates learned rather than fixed, partial pooling of policy effects, and multi-objective MPC with explicit health-system constraints. The coefficient-space view may also simplify communication of uncertainty: decision makers can see not only forecast bands but which control points and intervention levers generate them.

5 Conclusion

Control-polygon epidemiology contributes a single Bernstein–Bézier representation for five linked objectives. First, it reconstructs existing epidemic dynamics with interpretable control points using constrained LS or Bayesian renewal inference. Second, it forecasts by attaching a new future polygon rather than extrapolating the historical polynomial, with both LS and Bayesian uncertainty-aware implementations and optional peak/timing priorities or event information. Third, it reverses the problem for control: modeller-policymakers can move target epidemic or R_t control points interactively, after which a mechanistic model returns required intervention intensity, feasibility, uncertainty, and cost. Fourth, it provides structural minimum-complexity results: C resolvable crossings of $R_t = 1$ require at least $C + 1$ control points and E isolated interior extrema require at least $E + 2$. Fifth, it makes interpretability a mathematical property: convex-hull bounds certify thresholds, neighboring differences encode transmission or policy momentum, second differences encode curvature, uncertainty propagates across the full curve, and spectral calibration connects the same polygon to multiple mechanistic models.

The framework does not replace causal intervention estimates, economic data, or prospective forecast validation. Its contribution is a transparent common language through which these ingredients can be integrated. Observed data become an interpretable current polygon; current data plus external evidence become a probabilistic future polygon; and public-health targets become feasible, costed control polygons. The same mathematical object there-

fore supports retrospective understanding, prospective forecasting, and actionable epidemic control.

6 Table of symbols and notations

Table 3: Definition of symbols and notation used in the framework.

Symbol/notation	Definition
t	Calendar or epidemic time.
t_0	Beginning of the historical fitting interval.
T	Length of the time interval over which a Bernstein–Bézier trajectory or control is defined.
$u = (t - t_0)/T$	Rescaled historical time, with $u \in [0, 1]$.
v	Rescaled time for a future forecast segment, with $v \in [0, 1]$.
n	Degree of the historical Bernstein polynomial or policy-control polynomial.
m	Degree of a joined future Bernstein forecast segment.
$K = n + 1$	Number of Bernstein control points in a degree- n representation.
$B_{j,n}(u)$	Bernstein basis polynomial, $B_{j,n}(u) = \binom{n}{j} u^j (1 - u)^{n-j}$.
$\boldsymbol{\rho} = (\rho_0, \dots, \rho_n)^\top$	Vector of epidemiological Bernstein control-point coefficients.
ρ_j	j th epidemiological control point determining the latent reproduction-number trajectory.
$R(u)$	Bernstein representation of the reproduction-number trajectory, $R(u) = \sum_{j=0}^n \rho_j B_{j,n}(u).$
R_t	Effective or instantaneous reproduction number at time t .
$\tilde{R}_t$	Empirical or preliminary renewal-ratio estimate of R_t used in constrained least-squares fitting.
$\hat{\boldsymbol{\rho}}$	Estimated vector of Bernstein control points.
θ_j	Log-scale control-point parameter, $\theta_j = \log \rho_j$, used to enforce positivity in Bayesian inference.
$\Delta \rho_j = \rho_{j+1} - \rho_j$	First difference between adjacent control points; determines local transmission momentum through the derivative of the Bernstein curve.
$\Delta^2 \rho_j$	Second finite difference of the control points; describes curvature or acceleration of the trajectory.
Y_t	Observed incidence or reported case count at time t .
$\mathbf{Y}$	Vector of observed incidence values.
w_s	Generation-interval or infectiousness weight at lag s .

Symbol/notation	Definition
S	Maximum lag retained in the discrete generation-interval distribution.
Λ_t	Total infectiousness generated by previous incidence, $\Lambda_t = \sum_{s=1}^S w_s Y_{t-s}$.
μ_t	Conditional mean incidence at time t , with $\mu_t = R_t \Lambda_t$.
ϕ	Negative-binomial dispersion parameter; larger ϕ implies less observation overdispersion.
$\text{NB}(\mu_t, \phi)$	Negative-binomial observation distribution parameterized by mean μ_t and dispersion ϕ .
$\text{Var}(Y_t \mid \mu_t)$	Conditional incidence variance, $\mu_t + \mu_t^2/\phi$.
B	Bernstein design matrix evaluated at the observation times.
$\tilde{\mathbf{R}}$	Vector of empirical renewal-ratio estimates used in least-squares reconstruction.
W	Weight matrix used in constrained weighted least-squares estimation.
λ	Smoothness-penalty weight in Bernstein reconstruction; elsewhere, when explicitly stated in a cost analysis, it may denote the chosen burden/cost trade-off weight.
D_2	Second-difference operator applied to the vector of Bernstein coefficients.
$\ \cdot\ _2$	Euclidean norm.
C	Number of transversal crossings of the epidemic threshold $R_t = 1$.
E	Number of isolated interior extrema of a Bernstein trajectory when used in the structural-complexity result.
$R'(t)$	Time derivative of the reproduction-number trajectory.
$\rho(A)$	Spectral radius, or dominant eigenvalue magnitude, of matrix/operator A .
$\mathcal{K}(t; \beta)$	Time-dependent next-generation operator.
$A(t)$	Mechanistic next-generation structure excluding the scalar transmission multiplier.
$\beta(t)$	Scalar transmission-rate multiplier used in spectral calibration.
H	Forecast horizon.
$R_F(v)$	Joined future Bernstein reproduction-number trajectory, $R_F(v) = \sum_{j=0}^m q_j B_{j,m}(v)$.
q_j	j th control point of the joined future forecast polygon.
$q_j^{(0)}$	Baseline or no-event future control point.
$q_j^{(V)}$	Event- or variant-informed future control point.
T_H	Historical time scale used when matching the derivative of the historical and future Bernstein segments.

Symbol/notation	Definition
$D_{\mathbf{w}}$	Weighted forecast-priority loss combining whole-curve, peak-timing, peak-height, and onset-delay errors.
w_c	Weight assigned to whole-curve forecast discrepancy.
w_t	Weight assigned to peak-timing error.
w_h	Weight assigned to peak-height error.
w_d	Weight assigned to onset-delay error.
d_{curve}	Whole-curve forecast discrepancy used in the forecast-priority loss.
t_p	Observed or reference peak time.
$\widehat{t}_p$	Forecast peak time.
Y_p	Observed or reference peak incidence.
$\widehat{Y}_p$	Forecast peak incidence.
t_{onset}	Observed or reference epidemic-wave onset time.
$\widehat{t}_{\text{onset}}$	Forecast epidemic-wave onset time.
τ_t	Scaling constant for peak-timing error.
τ_d	Scaling constant for onset-delay error.
$p(t)$	Frequency or prevalence of an emerging variant or event-associated population fraction at time t .
p_j	Variant frequency or event level associated with future control point j .
a	Effective transmission multiplier for an emerging variant in the event-informed forecast map.
$\mathcal{E}$	Declared external event or surveillance information used to condition a forecast.
$u(t)$	Standardized time-varying intervention or policy-control trajectory, constrained to $0 \leq u(t) \leq 1$.
a_j	j th Bernstein coefficient of the intervention-control polygon, $u(t) = \sum_{j=0}^n a_j B_{j,n}(t/T)$.
$\mathbf{a} = (a_0, \dots, a_n)^\top$	Vector of policy-control coefficients.
R_t^{free}	Reproduction number in the absence of the additional intervention being optimized.
R_t^*	Desired or target reproduction-number trajectory.
$u^*(t)$	Intervention intensity required by the inverse-control map to attain the target trajectory.
ε	Maximum or calibrated effectiveness of the standardized intervention in the multiplicative transmission-reduction model.
u_{max}	Maximum feasible intervention intensity imposed by policy or implementation constraints.

Symbol/notation	Definition
R_0	Baseline reproduction number when the uncontrolled reproduction number is assumed constant.
$\dot{u}(t)$	Time derivative of the intervention trajectory; represents implementation or relaxation speed.
$\frac{1}{T} \int_0^T u(t) dt$	Mean standardized intervention intensity over the policy interval; for a Bernstein control polygon it equals $(n+1)^{-1} \sum_{j=0}^n a_j$.
$\sup_t \dot{u}(t) $	Maximum absolute implementation ramp rate over the control interval.
WIS	Weighted interval score used to evaluate probabilistic forecasts.
RMSE	Root-mean-square error.
MAE	Mean absolute error.
NRMSE	Normalized root-mean-square error.
log - NRMSE	NRMSE computed after the stated logarithmic transformation of the incidence series.
$\mathbb{E}[\cdot]$	Expectation operator.
$\text{Var}(\cdot)$	Variance operator.
$\hat{\cdot}$	Estimated, fitted, or forecast quantity.
$(\cdot)^*$	Target or inverse-control quantity.
$(\cdot)^{\text{free}}$	Quantity evaluated under the uncontrolled or no-additional-control condition.
$(\cdot)^{(0)}$	Baseline/no-event forecast quantity.
$(\cdot)^{(V)}$	Variant- or event-informed forecast quantity.

Online version of the Shiny applications

Visit the link for the Shiny app accompanying this paper: <https://jfrabajanteuplb.shinyapps.io/CONPOLYEPI>

For the sample Decision Support Shiny app: <https://jfrabajanteuplb.shinyapps.io/EPIDSS>

Reproducibility and data availability

The reproducibility package contains all simulation scripts, numerical summaries, figure files, the Philippine COVID-19 reporting-date series, the complete Supplementary Material, and the companion Shiny app. The Philippine COVID-19 data from the Department of Health contains 1,178 daily observations from 30 January 2020 through 21 April 2023 and 4,087,964 total reported cases. No individual-level data are used.

Code availability

Reference code implementations reproduce the reconstruction, complexity, Bayesian forecasting, mechanistic embedding, inverse-control/cost, and Philippine empirical analyses. A companion R Shiny app presents the reconstruction–forecast–control workflow interactively, including custom data upload, editable control points, editable model parameters, and LS and interactive Bayesian/Laplace modes. The complete Shiny user manual is provided as Supplementary Section S13.

Ethics statement

This methodological study uses synthetic simulations and an aggregated Philippine COVID-19 reporting-date case series. No individual-level human-subject records, personal identifiers, or private health information are analyzed.

Competing interests

The author declares no competing interests.

Declaration of Generative AI Use

During the preparation of this manuscript, the author used GPT-5.6 Sol (OpenAI) for grammar checking, code debugging, and code optimization. All AI-assisted text and code were carefully reviewed, edited as needed, and verified for correctness by the author, who takes full responsibility for the content and accuracy of the publication.

References

- [1] Jacco Wallinga and Peter Teunis. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of Epidemiology*, 160(6):509–516, 2004. doi: 10.1093/aje/kwh255.
- [2] Christophe Fraser. Estimating individual and household reproduction numbers in an emerging epidemic. *PLoS ONE*, 2(8):e758, 2007. doi: 10.1371/journal.pone.0000758.
- [3] Anne Cori, Neil M. Ferguson, Christophe Fraser, and Simon Cauchemez. A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology*, 178(9):1505–1512, 2013. doi: 10.1093/aje/kwt133.
- [4] Robin N. Thompson, Jessica E. Stockwin, Ryan D. van Gaalen, Jonathan A. Polonsky, Zhian N. Kamvar, Paul A. Demarsh, Elisabeth Dahlgvist, Siyu Li, Eva Miguel, Thibaut Jombart, Justin Lessler, Simon Cauchemez, and Anne Cori. Improved inference of

- time-varying reproduction numbers during infectious disease outbreaks. *Epidemics*, 29:100356, 2019. doi: 10.1016/j.epidem.2019.100356.
- [5] Katelyn M. Gostic, Lauren McGough, Edward B. Baskerville, Sam Abbott, Keya Joshi, Christine Tedijanto, Rebecca Kahn, Rene Niehus, James A. Hay, Pablo M. De Salazar, Joel Hellewell, Sophie Meakin, James D. Munday, Nikos I. Bosse, Katharine Sherratt, Robin N. Thompson, Laura F. White, Jana S. Huisman, Jérémie Scire, Sebastian Bonhoeffer, Tanja Stadler, Jacco Wallinga, Sebastian Funk, Marc Lipsitch, and Sarah Cobey. Practical considerations for measuring the effective reproductive number, r_t . *PLoS Computational Biology*, 16(12):e1008409, 2020. doi: 10.1371/journal.pcbi.1008409.
 - [6] Kris V. Parag. Improved estimation of time-varying reproduction numbers at low case incidence and between epidemic waves. *PLoS Computational Biology*, 17(9):e1009347, 2021. doi: 10.1371/journal.pcbi.1009347.
 - [7] Jacco Wallinga and Marc Lipsitch. How generation intervals shape the relationship between growth rates and reproductive numbers. *Proceedings of the Royal Society B*, 274(1609):599–604, 2007. doi: 10.1098/rspb.2006.3754.
 - [8] Sang Woo Park, David Champredon, Joshua S. Weitz, and Jonathan Dushoff. A practical generation-interval-based approach to inferring the strength of epidemics from their speed. *Epidemics*, 27:12–18, 2019. doi: 10.1016/j.epidem.2018.12.002.
 - [9] Oswaldo Gressani, Jacco Wallinga, Christian L. Althaus, Niel Hens, and Christel Faes. Epilps: A fast and flexible bayesian tool for estimation of the time-varying reproduction number. *PLoS Computational Biology*, 18(10):e1010618, 2022. doi: 10.1371/journal.pcbi.1010618.
 - [10] R. Pandey and H. Tolani. Penalized splines model to estimate time-varying reproduction number for covid-19 in india: A bayesian semi-parametric approach. *Clinical Epidemiology and Global Health*, 18:101176, 2022. doi: 10.1016/j.cegh.2022.101176.
 - [11] Oliver Eales, Kylie E. C. Ainslie, Caroline E. Walters, Haowei Wang, Christina Atchison, Deborah Ashby, Christl A. Donnelly, Graham Cooke, Wendy Barclay, Helen Ward, Ara Darzi, Paul Elliott, and Steven Riley. Appropriately smoothing prevalence data to inform estimates of growth rate and reproduction number. *Epidemics*, 40:100604, 2022. doi: 10.1016/j.epidem.2022.100604.
 - [12] Eugen Piricalabelu. A spline-based time-varying reproduction number for modelling epidemiological outbreaks. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 72(3):688–702, 2023. doi: 10.1093/jrssc/qlad027.
 - [13] Xiaoxi Pang, Yang Han, Elise Tressier, Nur Abdul Aziz, Lorenzo Pellis, Thomas House, and Ian Hall. Time-varying reproduction number estimation: Fusing compartmental models with generalized additive models. *Journal of the Royal Society Interface*, 22(222):20240518, 2025. doi: 10.1098/rsif.2024.0518.

- [14] Nicholas G. Davies, Sam Abbott, Rosanna C. Barnard, Christopher I. Jarvis, Adam J. Kucharski, James D. Munday, Carl A. B. Pearson, Timothy W. Russell, Damien C. Tully, Alex D. Washburne, et al. Estimated transmissibility and impact of sars-cov-2 lineage b.1.1.7 in england. *Science*, 372(6538):eabg3055, 2021. doi: 10.1126/science.abg3055.
- [15] Erik Volz, Swapnil Mishra, Meera Chand, Jeffrey C. Barrett, Robert Johnson, Lily Geidelberg, Wes R. Hinsley, Daniel J. Laydon, Gavin Dabrera, Àine O’Toole, et al. Assessing transmissibility of sars-cov-2 lineage b.1.1.7 in england. *Nature*, 593:266–269, 2021. doi: 10.1038/s41586-021-03470-x.
- [16] Fritz Obermeyer, Martin Jankowiak, Nikolaos Barkas, Stephen F. Schaffner, Jesse D. Pyle, Leonid Yurkovetskiy, Matteo Bosso, Daniel J. Park, Mehrtash Babadi, Bronwyn L. MacInnis, Jeremy Luban, Pardis C. Sabeti, and Jacob E. Lemieux. Analysis of 6.4 million sars-cov-2 genomes identifies mutations associated with fitness. *Science*, 376(6599):1327–1332, 2022. doi: 10.1126/science.abm1208.
- [17] G. G. Lorentz. *Bernstein Polynomials*. Chelsea Publishing Company, New York, 2 edition, 1986.
- [18] Rida T. Farouki and Tim N. T. Goodman. On the optimal stability of the bernstein basis. *Mathematics of Computation*, 65(216):1553–1566, 1996. doi: 10.1090/S0025-5718-96-00759-2.
- [19] Gerald Farin. *Curves and Surfaces for CAGD: A Practical Guide*. Morgan Kaufmann, San Francisco, 5 edition, 2002.
- [20] Neda Rahmani, M. Khodabin, and E. Hashemizadeh. Numerical solution of stochastic sir model by bernstein polynomials. *Journal of Interpolation and Approximation in Scientific Computing*, 2016(Special Issue 1):19–25, 2016. doi: 10.5899/2016/jiasc-00102.
- [21] Tibor Kmet and Maria Kmetova. Bézier curve parametrisation and echo state network methods for solving optimal control problems of sir model. *Biosystems*, 186:104029, 2019. doi: 10.1016/j.biosystems.2019.104029.
- [22] Qing Liu and Longbing Cao. Modeling time evolving covid-19 uncertainties with density dependent asymptomatic infections and social reinforcement. *Scientific Reports*, 12:5891, 2022. doi: 10.1038/s41598-022-09879-2.
- [23] Emre Eroglu. Accurate estimation of covid-19 active cases using bézier curve-based mathematical modeling. *Journal of Advances in Applied & Computational Mathematics*, 12:217–230, 2025. doi: 10.15377/2409-5761.2025.12.13.
- [24] Graham C. Gibson, Spencer J. Fox, Emily Javan, Susan E. Ptak, Oluwasegun M. Ibrahim, Michael Lachmann, and Lauren Ancel Meyers. Improving outbreak forecasts through model augmentation. *Proceedings of the National Academy of Sciences of the United States of America*, 122(43):e2508575122, 2025. doi: 10.1073/pnas.2508575122.

- [25] Matthew J. Y. Shin, Juliette Paireau, and Simon Cauchemez. Leveraging regularity in covid-19 growth rate dynamics for epidemic wave forecasting. *Epidemics*, 55:100908, 2026. doi: 10.1016/j.epidem.2026.100908.
- [26] Norvin P. Bansilan and Jomar F. Rabajante. A descriptive analytics of the covid-19 pandemic in a middle-income country with forward-looking insights. *Healthcare Analytics*, 5:100320, 2024. doi: 10.1016/j.health.2024.100320.
- [27] Nicholas G. Reich, Logan C. Brooks, Spencer J. Fox, Sasikiran Kandula, Craig J. McGowan, Evan Moore, Dave Osthus, Evan L. Ray, Abhinav Tushar, Teresa K. Yamana, Matthew Biggerstaff, Michael A. Johansson, Roni Rosenfeld, and Jeffrey Shaman. A collaborative multiyear, multimodel assessment of seasonal influenza forecasting in the united states. *Proceedings of the National Academy of Sciences USA*, 116(8):3146–3154, 2019. doi: 10.1073/pnas.1812594116.
- [28] Estee Y. Cramer, Evan L. Ray, Velma K. Lopez, Johannes Bracher, Andrea Brennen, Alvaro J. Castro Rivadeneira, Aaron Gerding, Tilmann Gneiting, Katie H. House, Yuxin Huang, Dasuni Jayawardena, Abdul-Hafiz Kanji, Amit Khandelwal, Khoa Le, Anja Mühlemann, Elaine O. Nsoesie, S. Shah, Nicholas G. Reich, et al. Evaluation of individual and ensemble probabilistic forecasts of covid-19 mortality in the united states. *Proceedings of the National Academy of Sciences USA*, 119(15):e2113561119, 2022. doi: 10.1073/pnas.2113561119.

Supplementary Material for Control-Polygon Epidemiology: A Bernstein–Bézier Framework for Epidemic Reconstruction, Forecasting, and Control

Jomar F. Rabajante^{*1}

¹University of the Philippines Los Baños, Laguna, Philippines

Overview

This document contains the technical detail as supplement to the main article. It presents the complete model formulation, derivations and proofs, forecasting/control algorithms, all examples, extended Philippine COVID-19 experiments, sensitivity analyses, simulation settings, forecast scoring details, and a user manual for the companion Shiny application.

Shiny app accompanying this paper: <https://jfrabajanteuplb.shinyapps.io/CONPOLYEPI>

Sample Decision Support Shiny app: <https://jfrabajanteuplb.shinyapps.io/EPIDSS>

Main-article to Supplementary-Material crosswalk.

Main-article topic	Detailed location in this supplement
Historical reconstruction and LS/Bayesian inference	S1
Mathematical guarantees, minimum-complexity proofs, uncertainty propagation	S2 and S11
Joined, peak-aware, priority-weighted, and event-informed forecasting	S3 and S12
Inverse control, intervention mechanisms, cost identities, decision support meeting workflow	S4
Complete synthetic experiments and mechanistic examples	S5
Practical K selection and growth-rate interpretation	S6–S7
Extended Philippine COVID-19 analyses and sensitivity tests	S8 and S10
Complete model/observation equation compendium	S9
Complete simulation and Philippine-example settings	S9–S10
Shiny installation, tab-by-tab workflow, interpretation, and cautions	S13
Applied reproducibility/reporting checklist	S14

^{*}Corresponding author: jfrabajante@up.edu.ph

S1 Detailed reconstruction model and inference

S1.1 Bernstein representation of the effective reproduction number

Let the historical observation interval be $t \in [t_0, t_1]$ with duration $T = t_1 - t_0 > 0$, and define normalized time

$$u(t) = \frac{t - t_0}{T} \in [0, 1]. \quad (\text{S1})$$

For degree n , define the Bernstein basis

$$B_{j,n}(u) = \binom{n}{j} u^j (1 - u)^{n-j}, \quad j = 0, \dots, n. \quad (\text{S2})$$

We model

$$R(t) = \sum_{j=0}^n \rho_j B_{j,n}\{u(t)\}, \quad \rho_j > 0, \quad (\text{S3})$$

where $K = n + 1$ is the number of control points. The pairs

$$\left(t_0 + \frac{j}{n}T, \rho_j\right), \quad j = 0, \dots, n, \quad (\text{S4})$$

form the displayed control polygon. Except at the endpoints, these are generally not interpolation points on the effective or time-varying reproduction number $R(t)$; they are coefficients whose convex combinations generate the curve $[1, 2]$.

To prevent a common geometric misinterpretation: a control point acts directly on R_t and only indirectly on the epidemic curve. That is, when the coefficients represent R_t itself, the position of the control points relative to one (1) and their successive differences have exact implications for the latent transmission trajectory.

S1.2 Renewal observation model

Let Y_t be observed incident infections or a reporting-adjusted incidence proxy. Let $w_s \geq 0$ be a discrete generation-interval distribution with $\sum_{s=1}^S w_s = 1$. Total infectiousness is

$$\Lambda_t = \sum_{s=1}^S w_s Y_{t-s}. \quad (\text{S5})$$

The conditional mean incidence is

$$\mu_t = R_t \Lambda_t, \quad (\text{S6})$$

and we use an overdispersed observation model

$$Y_t \mid \mu_t, \phi \sim \text{NB}(\mu_t, \phi), \quad (\text{S7})$$

parameterized so that

$$\mathbb{E}(Y_t \mid \mu_t) = \mu_t, \quad \text{Var}(Y_t \mid \mu_t) = \mu_t + \frac{\mu_t^2}{\phi}. \quad (\text{S8})$$

This is deliberately close to established renewal estimators [3–5]; the methodological novelty is concentrated in the structured R_t representation and forecast continuation rather than in redefining renewal transmission.

For reported-case applications, Eq. (S6) should be embedded in a reporting-delay, ascertainment, or latent-infection layer when those processes are material. The control polygon is not a remedy for observation bias [4, 6].

S1.3 Two inference engines for the historical control polygon

The geometric representation does not prescribe a single inferential method. We use two complementary engines: a fast constrained least-squares estimator and a fully generative Bayesian negative-binomial renewal model. The former is useful for initialization, screening candidate K , and rapid operational fitting; the latter is the primary route to coherent latent-trajectory uncertainty and probabilistic forecasting.

S1.3.1 Constrained least squares on empirical renewal ratios

For days with $\Lambda_t > 0$, define the empirical renewal ratio

$$\tilde{R}_t = \frac{Y_t}{\Lambda_t}. \quad (\text{S9})$$

Let $\mathbf{B}$ be the Bernstein design matrix with entries $B_{tj} = B_{j,K-1}(u_t)$, let $\mathbf{W}$ be a nonnegative diagonal weight matrix, and let D_2 be the second-difference matrix. A penalized constrained estimator is

$$\hat{\boldsymbol{\rho}}_{\text{LS}} = \arg \min_{\boldsymbol{\rho}} \left\| \mathbf{W}^{1/2} (\tilde{\mathbf{R}} - \mathbf{B}\boldsymbol{\rho}) \right\|_2^2 + \lambda \|D_2 \boldsymbol{\rho}\|_2^2, \quad A\boldsymbol{\rho} \geq \mathbf{b}. \quad (\text{S10})$$

Positivity uses $\rho_j \geq \varepsilon$; monotonic or convex segments can be imposed through linear constraints on first or second differences. Such Bernstein shape restrictions reduce to linear-inequality constrained least squares and can be solved as a quadratic program [7, 8]. Without active constraints,

$$\hat{\boldsymbol{\rho}}_{\text{LS}} = \left(\mathbf{B}^\top \mathbf{W} \mathbf{B} + \lambda D_2^\top D_2 \right)^{-1} \mathbf{B}^\top \mathbf{W} \tilde{\mathbf{R}}. \quad (\text{S11})$$

Because $\tilde{R}_t$ is heteroscedastic and uses observed incidence on both sides of the renewal ratio, we treat Eq. (S10) as a computational benchmark rather than the preferred generative likelihood. Confidence bands are obtained by parametric renewal bootstrap: simulate incidence from the fitted model, refit the control polygon, and take pointwise bootstrap quantiles.

S1.3.2 Bayesian negative-binomial renewal inference

For the primary inferential model, retain Eqs. (S5)–(S7) and parameterize

$$\rho_j = \exp(\theta_j), \quad (\text{S12})$$

which enforces positivity without truncating the R_t curve itself. One smoothness prior used in the simulations is

$$\theta_j \sim N(0, \sigma_0^2), \quad (\text{S13})$$

$$\Delta^2 \theta_j \sim N(0, \tau_2^2), \quad (\text{S14})$$

$$\log \phi \sim N(m_\phi, s_\phi^2), \quad (\text{S15})$$

with hyperparameters declared in the simulation settings. The posterior is

$$p(\boldsymbol{\theta}, \phi \mid \mathbf{Y}) \propto \prod_t p_{\text{NB}} \left(Y_t; \Lambda_t \sum_j e^{\theta_j} B_{j,K-1}(u_t), \phi \right) p(\boldsymbol{\theta}) p(\phi). \quad (\text{S16})$$

The prior can be replaced by truncated-normal or general linear-inequality constructions when strict shape restrictions are scientifically justified [9, 10]. Posterior samples are transformed to ρ , then through the Bernstein basis to obtain pointwise credible intervals for latent R_t . Drawing negative-binomial observations conditional on posterior R_t samples produces posterior predictive intervals for incidence; these intervals include observation/process variability and are therefore conceptually distinct from credible intervals for the latent transmission curve.

For a future horizon $h = 1, \dots, H$, each posterior historical draw is joined to a new future polygon by theorem 13. Remaining future coefficients are sampled from Eq. (S40) and, when external event information is available, from uncertain event parameters such as variant multiplier a and takeover rate s . This yields posterior quantities directly relevant to surveillance,

$$P(R_{T+h} > 1 \mid \mathbf{Y}, \mathcal{E}), \quad P\left(\max_{1 \leq h \leq H} R_{T+h} > 1 \mid \mathbf{Y}, \mathcal{E}\right), \quad (\text{S17})$$

where $\mathcal{E}$ denotes available event information.

S1.3.3 Selecting K after the structural floor

The theorems below provide a necessary floor K_{floor} , not a complete model-selection rule. Candidate fitting therefore begins at that floor. For constrained least squares, blocked or rolling-origin cross-validation and forecast error can be combined with residual adequacy. For Bayesian models, leave-one-out predictive criteria, WAIC, or rolling-origin posterior predictive scores are natural choices. In both cases we recommend the smallest $K \geq K_{\text{floor}}$ whose threshold features are stable and whose predictive adequacy is acceptable, rather than maximizing in-sample fit.

S2 Proofs and mathematical properties of the epidemiological control polygon

The results below are elementary consequences of the Bernstein basis but become useful when the coefficients have epidemiological meaning. They distinguish global guarantees from heuristic visual interpretation.

Theorem 1 (Convex-hull and transmission-threshold certificate). *For $u \in [0, 1]$, let*

$$R(u) = \sum_{j=0}^n \rho_j B_{j,n}(u).$$

Then

$$\min_j \rho_j \leq R(u) \leq \max_j \rho_j. \quad (\text{S18})$$

Consequently: (i) if $\rho_j > 0$ for all j , then $R(u) > 0$; (ii) if $\rho_j \leq 1$ for all j , then $R(u) \leq 1$ for all u ; and (iii) if $\rho_j \geq 1$ for all j , then $R(u) \geq 1$ for all u .

Proof. For $u \in [0, 1]$, each $B_{j,n}(u) \geq 0$ and

$$\sum_{j=0}^n B_{j,n}(u) = (u + 1 - u)^n = 1.$$

Thus $R(u)$ is a convex combination of the coefficients ρ_j . Equation (S18) and the threshold consequences follow immediately. $\square$

Remark 2. The theorem provides a sufficient global certificate. Mixed coefficients above and below one do not imply that every corresponding level is attained, because the control polygon is not an interpolant. The threshold-crossing theorem below gives the appropriate upper bound.

Proposition 3 (Spectral calibration into compartmental models). *Let a compartmental epidemic model have an instantaneous next-generation matrix of the form*

$$\mathcal{K}(t; \beta) = \beta(t)A(t), \quad (\text{S19})$$

where $A(t)$ is a nonnegative matrix with spectral radius $\rho\{A(t)\} > 0$. Let $R_B(t)$ be any positive Bernstein reproduction-number curve of the form in Eq. (S3). If

$$\beta(t) = \frac{R_B(t)}{\rho\{A(t)\}}, \quad (\text{S20})$$

then the model's instantaneous effective reproduction number is exactly

$$\rho\{\mathcal{K}(t; \beta)\} = R_B(t). \quad (\text{S21})$$

Proof. The spectral radius is positively homogeneous: for scalar $c \geq 0$, $\rho(cA) = c\rho(A)$. Therefore

$$\rho\{\mathcal{K}(t; \beta)\} = \beta(t)\rho\{A(t)\} = R_B(t).$$

□

Remark 4. Proposition 3 is the bridge used in the SEIR, SIRS, age-structured, vaccination, and hospitalization/isolation examples below. It does not imply that every mechanistic model admits a single scalar calibration; the required next-generation structure must first be derived. Vector-borne transmission, for example, contains two directional transmission terms, but if both are multiplied by a common positive scalar then the next-generation spectral radius scales by that scalar and an analogous calibration is available [11, 12].

Theorem 5 (Finite differences determine derivatives). *For integer $r \in \{1, \dots, n\}$,*

$$\frac{d^r R}{du^r} = \frac{n!}{(n-r)!} \sum_{j=0}^{n-r} \Delta^r \rho_j B_{j,n-r}(u). \quad (\text{S22})$$

In calendar time,

$$\frac{dR}{dt} = \frac{n}{T} \sum_{j=0}^{n-1} \Delta \rho_j B_{j,n-1}(u). \quad (\text{S23})$$

Proof. Using

$$\frac{d}{du} B_{j,n}(u) = n\{B_{j-1,n-1}(u) - B_{j,n-1}(u)\},$$

with out-of-range basis terms set to zero, collect adjacent coefficients to obtain

$$R'(u) = n \sum_{j=0}^{n-1} (\rho_{j+1} - \rho_j) B_{j,n-1}(u).$$

Repeated differentiation gives Eq. (S22). Since $du/dt = 1/T$, Eq. (S23) follows by the chain rule. □

Corollary 6 (Monotonicity and curvature certificates). *If $\Delta\rho_j \geq 0$ for all j , then $R(t)$ is nondecreasing; if $\Delta\rho_j \leq 0$ for all j , it is nonincreasing. If $\Delta^2\rho_j \geq 0$ for all j , $R(t)$ is convex; if $\Delta^2\rho_j \leq 0$ for all j , it is concave.*

Proof. Apply theorem 5 and the nonnegativity of every Bernstein basis function on $[0, 1]$. $\square$

Equation (S23) motivates the term transmission momentum for $\Delta\rho_j$. Positive differences push the derivative upward and negative differences push it downward. The second differences describe changes in that momentum. This is a geometric description of the latent R_t process, not a claim that R_t itself is a physical momentum variable.

Theorem 7 (Variation-diminishing threshold-crossing bound). *Let $c \in \mathbb{R}$ and define the threshold-centered coefficient sequence*

$$a_j = \rho_j - c, \quad j = 0, \dots, n.$$

Let $V(a_0, \dots, a_n)$ be the number of sign variations after zeros are removed. The number N_c of distinct transversal roots of $R(u) - c$ in $(0, 1)$ satisfies

$$N_c \leq V(a_0, \dots, a_n) \leq n. \quad (\text{S24})$$

More generally, the classical Bernstein variation-diminishing result bounds real roots with the usual multiplicity qualifications.

Proof. Write

$$R(u) - c = \sum_{j=0}^n (\rho_j - c) B_{j,n}(u).$$

The Bernstein basis is a normalized totally positive basis on $(0, 1)$. The variation-diminishing property of totally positive bases states that the number of sign changes (and, with the standard formulation, zeros counted with appropriate multiplicity) of a linear combination cannot exceed the sign variation of its coefficient sequence [1, 13, 14]. Hence Eq. (S24) follows. $\square$

Corollary 8 (Structural lower bound from epidemic threshold crossings). *If a candidate R_t trajectory must contain at least C transversal crossings of $R_t = 1$, then any Bernstein representation requires*

$$K \geq C + 1. \quad (\text{S25})$$

Furthermore, the threshold-centered sequence $(\rho_0 - 1, \dots, \rho_{K-1} - 1)$ must have at least C sign variations.

Proof. Set $c = 1$ in theorem 7. Since $V \leq n = K - 1$, $C \leq K - 1$. $\square$

Theorem 9 (Structural lower bound from interior extrema). *Let E be the number of isolated interior extrema of $R(u)$ whose derivative changes sign. Then*

$$E \leq V(\Delta\rho_0, \dots, \Delta\rho_{n-1}) \leq n - 1 = K - 2. \quad (\text{S26})$$

Therefore a trajectory requiring E such extrema must satisfy

$$K \geq E + 2. \quad (\text{S27})$$

Proof. By theorem 5, $R'(u)/n$ is a Bernstein polynomial of degree $n - 1$ with coefficients $\Delta\rho_j$. Apply theorem 7 to $R'(u)$ at level zero. Each isolated extremum with a sign-changing derivative contributes a distinct interior zero of R' . $\square$

Remark 10 (What can be known in advance). Equations (S25) and (S27) are necessary structural lower bounds, not formulas for the statistically optimal K . Noise, weakly separated transitions, asymmetric waves, reporting artifacts, and rapid curvature may require K above the theorem floor. Conversely, apparent crossings induced by noise should not be used as structural requirements. In applications we therefore begin candidate fitting at

$$K_{\text{floor}} = \max(C_{\text{req}} + 1, E_{\text{req}} + 2) \quad (\text{S28})$$

and select the smallest K that also satisfies predictive and residual adequacy criteria.

Theorem 11 (Coefficient-wise uncertainty envelope). *Suppose each control point obeys $\ell_j \leq \rho_j \leq u_j$. Define*

$$L(u) = \sum_{j=0}^n \ell_j B_{j,n}(u), \quad U(u) = \sum_{j=0}^n u_j B_{j,n}(u). \quad (\text{S29})$$

Then for every $u \in [0, 1]$,

$$L(u) \leq R(u) \leq U(u). \quad (\text{S30})$$

Proof. Because $B_{j,n}(u) \geq 0$, multiplying the coefficient inequalities by $B_{j,n}(u)$ preserves order. Summing over j yields Eq. (S30). $\square$

This envelope is exact for the hyper-rectangle of coefficient intervals in the sense that it is the pointwise minimum and maximum induced by independently varying coefficients within their intervals. Posterior coefficient dependence can yield narrower credible bands; those should be computed by joint posterior propagation rather than by marginal intervals alone.

Proposition 12 (Posterior pushforward and probabilistic threshold certificates). *Let $\mathbf{b}(u) = (B_{0,n}(u), \dots, B_{n,n}(u))^T$ and suppose the posterior distribution of $\boldsymbol{\rho}$ has finite second moments with mean $\mathbf{m}$ and covariance Σ . Then*

$$\mathbb{E}\{R(u) \mid \mathbf{Y}\} = \mathbf{b}(u)^T \mathbf{m}, \quad (\text{S31})$$

$$\text{Cov}\{R(u), R(v) \mid \mathbf{Y}\} = \mathbf{b}(u)^T \Sigma \mathbf{b}(v). \quad (\text{S32})$$

Moreover,

$$P\left(\sup_{u \in [0,1]} R(u) \leq 1 \mid \mathbf{Y}\right) \geq P\left(\max_j \rho_j \leq 1 \mid \mathbf{Y}\right), \quad (\text{S33})$$

$$P\left(\inf_{u \in [0,1]} R(u) \geq 1 \mid \mathbf{Y}\right) \geq P\left(\min_j \rho_j \geq 1 \mid \mathbf{Y}\right). \quad (\text{S34})$$

Proof. The first two identities follow from $R(u) = \mathbf{b}(u)^T \boldsymbol{\rho}$ and linearity of conditional expectation and covariance. For the probability bounds, theorem 1 shows that the coefficient events $\{\max_j \rho_j \leq 1\}$ and $\{\min_j \rho_j \geq 1\}$ are sufficient, respectively, for the corresponding global curve events. Taking posterior probabilities preserves set inclusion. $\square$

The bounds in Eq. (S34) are conservative because mixed control points can still generate an entirely subcritical or supercritical curve. Direct posterior evaluation of $\sup R(u)$ on a sufficiently dense grid gives a sharper numerical probability, while the coefficient event remains a transparent global certificate.

S3 Detailed forecasting derivations and algorithms

S3.1 Why direct polynomial extrapolation is not the proposed forecast

The convex-hull interpretation in theorem 1 relies on $u \in [0, 1]$. Outside that domain the Bernstein basis functions no longer all form nonnegative convex weights. Consequently, evaluating a fitted historical polynomial at $u > 1$ can produce unstable or implausible behavior and loses the control-polygon threshold certificate. Our forecast therefore starts a new Bernstein segment at the forecast origin.

Let the historical segment be

$$R_H(t) = \sum_{j=0}^n \rho_j B_{j,n} \left(\frac{t - t_0}{T_H} \right), \quad t \in [t_0, T_0], \quad (\text{S35})$$

where $T_H = T_0 - t_0$. Let the forecast horizon have length H and define

$$R_F(t) = \sum_{j=0}^m q_j B_{j,m} \left(\frac{t - T_0}{H} \right), \quad t \in [T_0, T_0 + H]. \quad (\text{S36})$$

Theorem 13 (C^1 joining condition). *The two segments are continuous at T_0 if*

$$q_0 = \rho_n. \quad (\text{S37})$$

They are continuously differentiable at T_0 if, in addition,

$$\frac{m}{H}(q_1 - q_0) = \frac{n}{T_H}(\rho_n - \rho_{n-1}), \quad (\text{S38})$$

or equivalently

$$q_1 = q_0 + \frac{nH}{mT_H}(\rho_n - \rho_{n-1}). \quad (\text{S39})$$

Proof. The endpoint interpolation property of Bernstein curves gives $R_H(T_0) = \rho_n$ and $R_F(T_0) = q_0$, proving Eq. (S37). By theorem 5, the historical terminal derivative is $n(\rho_n - \rho_{n-1})/T_H$ and the future initial derivative is $m(q_1 - q_0)/H$. Equality gives Eqs. (S38)–(S39). $\square$

The remaining future coefficients $q_2, \dots, q_m$ encode the forecast hypothesis. They may follow momentum continuation, mean reversion, covariate regression, posterior dynamics learned from past waves, or event-informed priors. A simple stochastic prior is

$$q_j = q_{j-1} + \kappa(q_{j-1} - q_{j-2}) + \eta(\bar{R} - q_{j-1}) + \epsilon_j, \quad \epsilon_j \sim N(0, \sigma_q^2), \quad (\text{S40})$$

with positivity enforced by truncation or a positive link. Here κ controls momentum, η controls mean reversion, and $\bar{R}$ is a long-run reference level.

S3.2 Peak-aware adaptive continuation when a wave signal is already visible

A long historical fitting window can suppress a rapidly changing terminal signal. This is particularly problematic near the beginning of a new wave: the global Bernstein reconstruction may be smooth and stable while the final few observations are already growing exponentially. We therefore add an

optional peak-aware adaptive continuation that remains within the control-polygon framework but gives recent growth a separate role.

Let $Y_t^{(7)}$ denote the 7-day mean and estimate the terminal exponential growth rate from the final w observations by

$$(\hat{a}, \hat{g}_T) = \arg \min_{a, g} \sum_{r=0}^{w-1} \left[\log \{ Y_{T_0-w+1+r}^{(7)} + 1 \} - a - gr \right]^2. \quad (\text{S41})$$

If the generation interval is Gamma with shape k and scale θ , the Euler–Lotka relation gives the local growth-implied reproduction number

$$R_g = (1 + \theta \hat{g}_T)^k, \quad (\text{S42})$$

provided $1 + \theta \hat{g}_T > 0$ [15, 16]. This quantity is not substituted blindly for the historical polygon. Instead it anchors a short future template only when $\hat{g}_T$ exceeds a declared growth trigger g_c .

For the history-only mode used in the Philippine backtest, the template is

$$\tilde{R}(h) = \begin{cases} R_0, & 1 \leq h \leq d, \\ R_{\text{post}} + (R_0 - R_{\text{post}}) \exp\{-(h - d)/\tau\}, & h > d, \end{cases} \quad (\text{S43})$$

where $R_0 = \min(R_g, R_{\text{max}})$. If $\hat{g}_T \leq g_c$, the continuation instead mean-reverts smoothly from the recent median R_t . The daily template is then compressed into two positive Bernstein segments, with the first segment intentionally short so that a sharp terminal surge is not smeared over the full forecast horizon. The split is a numerical representation choice, not a claim that transmission truly changes discontinuously.

The history-only hyperparameters are selected by a peak-aware rolling criterion using origins strictly before the evaluation period,

$$D_{\text{peak}} = \frac{|\hat{t}_p - t_p|}{14} + 2 \left| \frac{\hat{Y}_p - Y_p}{Y_p} \right| + \frac{\text{MAE}(\hat{Y}, Y)}{\max(\bar{Y}, 100)}. \quad (\text{S44})$$

For the Philippine illustration, a grid search over $w \in \{4, 5, 7\}$, $d \in \{4, 6, 8, 10\}$ d, $\tau \in \{3, 4, 5, 6, 8\}$ d, and $R_{\text{post}} \in \{0.6, 0.7, 0.8, 0.9\}$ was performed only at three pre-Omicron origins (24 March, 19 June, and 14 September 2021). The selected history-only rule was $w = 4$, $d = 10$ d, $\tau = 8$ d, $R_{\text{post}} = 0.9$, with $g_c = 0.02 \text{ d}^{-1}$ and $R_{\text{max}} = 2.5$. Six later origins were reserved for out-of-selection evaluation.

An event-informed high-growth mode removes the history-only cap when an external surge or variant alert is active. We expose two declared event profiles in the Philippine illustration. The event-balanced profile uses a short high-growth phase followed by approximately one-generation decay toward $R_{\text{post}} = 0.8$. The event peak-priority profile places greater decision weight on Eq. (S45)’s timing and height terms and uses a slightly longer high-growth phase with a more subcritical post-wave anchor. In the Bayesian version, recent log-growth is sampled from its local regression posterior and the event scale, high-growth duration, decay time, post-wave anchor, and phase adjustment are assigned uncertainty before propagation through the renewal model. To keep the probabilistic forecast centered on the event scenario actually declared by decision makers, the simulated $\log R_t$ deviations are pointwise recentered so that their posterior median equals the deterministic event curve before incidence propagation. This scenario centering uses no post-origin cases; it prevents nonlinear uncertainty propagation from shifting the Bayesian central scenario away from the LS/event specification. Predictive intervals still reflect the full centered distribution. These are conditional scenario assumptions, not estimates obtained from post-origin case outcomes; the cases-only and event-conditioned versions remain distinct throughout the results.

S3.3 User-weighted forecast priorities, onset delay, and phase adjustment

Peak date and peak height are often direct decision targets rather than secondary summaries. Validation studies of epidemic models therefore score peak date and magnitude explicitly [17], and experimental work shows that additional information about epidemic timing or intensity can materially improve peak forecasts [18]. We consequently generalize Eq. (S44) to a user-weighted objective

$$D_{\mathbf{w}} = w_c d_{\text{curve}} + w_t \frac{|\hat{t}_p - t_p|}{\tau_t} + w_h \left| \frac{\hat{Y}_p - Y_p}{Y_p} \right| + w_d \frac{|\hat{t}_{\text{onset}} - t_{\text{onset}}|}{\tau_d}, \quad (\text{S45})$$

where d_{curve} may be MAE, log-NRMSE, WIS, or another declared whole-trajectory score; τ_t and τ_d are scale constants. The default balanced profile uses all four terms; a peak-priority profile increases w_t and w_h ; a timing-priority profile increases w_t and w_d ; and advanced users may supply arbitrary nonnegative weights. This is intentionally a decision layer, not a new probability law.

To correct systematic lead/lag behavior without using centered future smoothing, the short future template may also include a phase parameter δ ,

$$\tilde{R}_\delta(h) = \tilde{R}(h + \delta), \quad (\text{S46})$$

with $\delta > 0$ advancing the future template and $\delta < 0$ delaying it. In deployable use, δ and the other shape parameters are selected only from earlier rolling origins using Eq. (S45); they are never chosen from the future being forecast. This directly addresses reporting/smoothing delays while preserving temporal honesty.

Least-squares and Bayesian versions remain parallel. In the LS version, candidate future-polygons are selected by minimizing $D_{\mathbf{w}}$ on training origins. In the Bayesian version, recent log-growth is assigned a posterior distribution, event-shape parameters receive declared priors, and posterior future polygons are propagated through the renewal model. Decision weights may then be used to choose among posterior summaries or actions, while credible and predictive intervals remain probabilistic statements. If a policymaker supplies an external prior on expected peak date or peak magnitude, it enters as explicit information rather than as hidden future outcome data.

No weighting scheme can guarantee zero unknown future peak error. When decision makers possess an external target or prior range for the peak—for example, an expected peak window from genomic growth, hospitalization leading indicators, or a scenario exercise—the algorithm can switch to a target-constrained priority mode. Let t_p^{ext} and Y_p^{ext} denote externally supplied peak targets. Candidate future polygons can minimize

$$D_{\mathbf{w}}^{\text{ext}} = D_{\mathbf{w}}^{\text{history}} + \lambda_t \left(\frac{\hat{t}_p - t_p^{\text{ext}}}{\tau_t} \right)^2 + \lambda_h \left\{ \log \left(\frac{\hat{Y}_p}{Y_p^{\text{ext}}} \right) \right\}^2, \quad (\text{S47})$$

where $D_{\mathbf{w}}^{\text{history}}$ penalizes departure from the history-based continuation. Large λ_t or λ_h makes timing or height dominate; in the limit these can be implemented as approximate equality constraints. This is useful when a meeting explicitly prioritizes “hit this peak window” or “keep the peak below this level.” The resulting error can be made arbitrarily small relative to the supplied target, but not necessarily relative to the unknown future if that target is wrong. External targets must therefore be displayed and archived as assumptions, not hidden tuning parameters.

For method diagnosis we therefore also allow a clearly labeled retrospective oracle mode that minimizes Eq. (S45) against the realized future. It answers “how close could this parameterization come if the future peak were known?” and is useful for studying representational capacity, but it is never counted as prospective or rolling-origin forecast skill. This distinction is especially important for the Omicron example below.

S3.4 Event-informed forecasting for a newly detected variant

A future variant cannot be inferred from pre-emergence incidence alone without an external model of emergence. The framework therefore treats a variant as an observed or hypothesized future event rather than claiming to predict its spontaneous appearance. Let $p(t)$ denote the future population share of the variant. A parsimonious takeover model is

$$p(t) = \text{logit}^{-1}\{\text{logit}(p_0) + s(t - T_0)\}, \quad (\text{S48})$$

where p_0 is its share at the forecast origin and s is its logistic growth rate. Such relative-growth information can be informed by genomic surveillance and lineage-fitness models [19–21].

Let $a > 0$ be an effective relative transmission multiplier. If $q_j^{(0)}$ are the no-event future control points and $p_j = p(t_j)$ is evaluated at their future control times, define the event-informed prior mean

$$q_j^{(V)} = q_j^{(0)}\{1 + (a - 1)p_j\}. \quad (\text{S49})$$

This is intentionally a scenario model: a can summarize transmissibility, immune escape, or other effects only to the extent that they act multiplicatively on the effective reproduction number. More detailed applications can replace Eq. (S49) with a multi-strain renewal system.

Proposition 14 (Coefficient-level variant threshold). *Assume $a > 1$ and $0 < q_j^{(0)} < 1$. The event-informed control point exceeds one if and only if*

$$p_j > \frac{1/q_j^{(0)} - 1}{a - 1}. \quad (\text{S50})$$

Proof. From Eq. (S49), $q_j^{(V)} > 1$ is equivalent to $q_j^{(0)}\{1 + (a - 1)p_j\} > 1$. Rearranging yields Eq. (S50). $\square$

Because interior control points need not lie on the curve, Eq. (S50) is a coefficient-level threshold, not an exact time at which R_t crosses one. The actual crossing is determined from the full Bernstein curve. The advantage is that the future curve remains globally bounded by its event-informed control polygon.

S3.5 Forecast algorithm

S4 Detailed inverse-control derivations and policy mappings

S4.1 From “what will happen?” to “what trajectory do we want?”

Forecasting asks for a distribution over future epidemic trajectories conditional on current information. Policy design asks the inverse question: given a desired trajectory, what intervention is required, is it feasible, and what does it cost? This distinction motivates a second control polygon. Let the standardized intervention intensity be

$$u(t) = \sum_{j=0}^n a_j B_{j,n}(t/T), \quad 0 \leq a_j \leq u_{\max} \leq 1, \quad (\text{S51})$$

where $u = 0$ denotes no additional control and $u = 1$ denotes the maximum intensity under the declared intervention scale. The vector $\mathbf{a} = (a_0, \dots, a_n)^\top$ is the policy control polygon. Its

Algorithm 1 Event-informed Bernstein renewal forecasting

Require: Incidence $Y_{1:T_0}$, generation weights $w_{1:S}$, historical K , future K_F , forecast horizon H , dispersion ϕ , optional event information $\mathcal{E}$.

- 1: Fit the historical polygon by constrained LS or draw $(\boldsymbol{\rho}, \phi)$ from Eq. (S16); Bayesian draws are preferred for probabilistic forecasting.
 - 2: Compute threshold-centered signs, $\Delta\rho_j$, $\Delta^2\rho_j$, and curve-level uncertainty for each draw.
 - 3: **for** each bootstrap or posterior draw **do**
 - 4: If no terminal surge rule is active, set $q_0 = \rho_n$, impose q_1 from Eq. (S39), and sample $q_2^{(0)}, \dots, q_m^{(0)}$ from Eq. (S40) or another declared future-control process.
 - 5: If the peak-aware rule in section S3.2 is active, estimate $\hat{g}_T$ and R_g , construct Eq. (S43), and compress it into short joined Bernstein future segments before renewal propagation.
 - 6: **if** event information $\mathcal{E}$ is available **then**
 - 7: Sample uncertain event parameters and transform future controls; for a variant use Eqs. (S48)–(S49).
 - 8: **end if**
 - 9: Evaluate $R_F(t)$ only on the new future segment $v \in [0, 1]$ and simulate future incidence through Eqs. (S5)–(S7).
 - 10: **end for**
 - 11: Summarize R_F by credible/confidence bands, incidence by predictive intervals, and threshold risk by Eq. (S17).
 - 12: Score forecasts using proper probabilistic metrics such as WIS [22].
-

epidemiological meaning is determined by the mechanistic way in which u enters the transmission model.

For a policymaker-facing implementation, the target itself may also be parameterized by Bernstein points. For example, a positive infectious-prevalence target can be written

$$\log I^*(t) = \sum_{j=0}^{K_I-1} \xi_j B_{j, K_I-1}(t/T). \quad (\text{S52})$$

A decision maker can move the vertical coordinates ξ_j in an interactive display; the model evaluates I^* , its analytic derivatives, the implied R_t^* , and the required intervention. We deliberately keep calendar time monotone and do not interpret free horizontal movement of a general parametric Bézier point as policy frequency. In the standard Bernstein graph representation, the coefficient abscissae are fixed. Operational frequency is better represented by piecewise Bernstein segments whose lengths are the policy-review intervals (for example, 7, 14, or 30 days).

S4.2 Least-squares and Bayesian control as parallel decision engines

The same two inferential philosophies used for reconstruction and forecasting are retained for inverse control. A point-estimate or least-squares control design treats the reconstructed states and parameters as fixed at $\hat{\boldsymbol{\vartheta}}$ and solves

$$\hat{\mathbf{a}}_{\text{LS}} = \arg \min_{\mathbf{a} \in \mathcal{A}} J(\mathbf{a}; \hat{\boldsymbol{\vartheta}}), \quad (\text{S53})$$

subject to the epidemic dynamics and feasibility constraints. Bootstrap refits can be pushed through the same optimizer to obtain frequentist uncertainty on the required control.

Under Bayesian reconstruction, ϑ has posterior $p(\vartheta \mid Y)$. There are then two natural outputs. The first is the induced posterior distribution of draw-specific optimal policies,

$$\mathbf{a}^{*(s)} = \arg \min_{\mathbf{a} \in \mathcal{A}} J(\mathbf{a}; \vartheta^{(s)}), \quad \vartheta^{(s)} \sim p(\vartheta \mid Y), \quad (\text{S54})$$

which yields posterior medians and credible bands for intervention intensity, cost, and health outcomes. The second is a Bayes action that minimizes posterior expected loss,

$$\hat{\mathbf{a}}_{\text{Bayes}} = \arg \min_{\mathbf{a} \in \mathcal{A}} E\{J(\mathbf{a}; \vartheta) \mid Y\}. \quad (\text{S55})$$

Safety requirements can also be written as chance constraints, for example

$$P\{H(t) \leq H_{\max} \ \forall t \in [0, T] \mid Y, \mathbf{a}\} \geq 1 - \alpha. \quad (\text{S56})$$

Thus the policymaker-facing interface should always allow a point/LS view and a Bayesian uncertainty-aware view. Moving a target point changes the same geometric object in either mode; what changes is whether the downstream intervention is a single optimized polygon or a posterior distribution over feasible polygons.

S4.3 Exact policy geometry and cost identities

Theorem 15 (Policy-polygon geometry). *Let $u(t)$ be given by Eq. (S51), with degree n over $t \in [0, T]$, and define $\Delta a_j = a_{j+1} - a_j$. Then:*

$$\min_j a_j \leq u(t) \leq \max_j a_j, \quad (\text{S57})$$

$$\frac{1}{T} \int_0^T u(t) dt = \frac{1}{n+1} \sum_{j=0}^n a_j, \quad (\text{S58})$$

$$\dot{u}(t) = \frac{n}{T} \sum_{j=0}^{n-1} \Delta a_j B_{j,n-1}(t/T), \quad (\text{S59})$$

$$\sup_t |\dot{u}(t)| \leq \frac{n}{T} \max_j |\Delta a_j|, \quad (\text{S60})$$

$$\int_0^T |\dot{u}(t)| dt \leq \sum_{j=0}^{n-1} |\Delta a_j|. \quad (\text{S61})$$

Moreover, if $G^{(n)}$ is the Bernstein Gram matrix

$$G_{ij}^{(n)} = \int_0^1 B_{i,n}(v) B_{j,n}(v) dv = \frac{\binom{n}{i} \binom{n}{j}}{(2n+1) \binom{2n}{i+j}}, \quad (\text{S62})$$

then

$$\frac{1}{T} \int_0^T u(t)^2 dt = \mathbf{a}^\top G^{(n)} \mathbf{a}, \quad (\text{S63})$$

$$\int_0^T \dot{u}(t)^2 dt = \frac{n^2}{T} (\Delta \mathbf{a})^\top G^{(n-1)} (\Delta \mathbf{a}). \quad (\text{S64})$$

Proof. Equation (S57) is the convex-hull property. Because $\int_0^1 B_{j,n}(v) dv = 1/(n+1)$, changing variables $v = t/T$ gives Eq. (S58). Equation (S59) follows from the standard Bernstein derivative identity. The derivative is a scaled convex combination of the first differences, giving Eq. (S60); applying the triangle inequality before integration and using $\int_0^1 B_{j,n-1} = 1/n$ gives Eq. (S61). Expanding u^2 and $\dot{u}^2$ in the Bernstein basis gives Eqs. (S63)–(S64). The product integral in Eq. (S62) follows from the beta integral. $\square$

Proposition 16 (Meaning of moving one graphical handle). *If only policy coefficient a_k is moved by δ , then*

$$u_{\text{new}}(t) - u_{\text{old}}(t) = \delta B_{k,n}(t/T), \quad (\text{S65})$$

so $\sup_t |u_{\text{new}} - u_{\text{old}}| \leq |\delta|$ and the change in mean intervention intensity is exactly $\delta/(n+1)$. Under Eq. (S69) with fixed R_{free} , the corresponding reproduction-number change is

$$R_{\text{new}}(t) - R_{\text{old}}(t) = -\varepsilon R_{\text{free}}(t) \delta B_{k,n}(t/T). \quad (\text{S66})$$

If instead one log-epidemic target coefficient ξ_k in Eq. (S52) is moved by δ , then

$$\frac{I_{\text{new}}^*(t)}{I_{\text{old}}^*(t)} = \exp\{\delta B_{k,K_I-1}(t/T)\}. \quad (\text{S67})$$

Thus a control-point edit has an exact, localized-in-weight mathematical effect even though a Bernstein basis function has global support.

Proof. Linearity of the Bernstein expansion gives the first identity. Nonnegativity and $B_{k,n} \leq 1$ give the supremum bound, while Eq. (S58) gives the mean shift. Substitution into Eq. (S69) gives the reproduction-number effect. The final identity follows by exponentiating the additive change $\delta B_{k,K_I-1}$ in log prevalence. $\square$

Thus the height of policy points can be mechanistic after an intervention scale is defined: their arithmetic mean is exactly the mean intervention intensity. Their successive differences encode tightening or relaxation, their maximum difference limits implementation ramp speed, and their total variation bounds switching burden. These properties are useful in interactive decision support because they are computable directly from the few displayed coefficients, without numerically integrating a finely discretized control. Bernstein direct-transcription methods exploit related geometric properties in general optimal-control problems [23, 24].

S4.4 Exact mapping under multiplicative transmission reduction

Suppose a policy acts on effective transmission through

$$\beta(t) = \beta_{\text{free}}(t)\{1 - \varepsilon u(t)\}, \quad 0 < \varepsilon \leq 1, \quad (\text{S68})$$

where ε converts the standardized policy scale into effective transmission reduction. If all other state variables are held at the values used to define $R_{\text{free}}(t)$, then

$$R_t = R_{\text{free}}(t)\{1 - \varepsilon u(t)\}. \quad (\text{S69})$$

Hence a target R_t^* requires

$$u^*(t) = \frac{1}{\varepsilon} \left\{ 1 - \frac{R_t^*}{R_{\text{free}}(t)} \right\}. \quad (\text{S70})$$

For suppression-only control, a necessary pointwise feasibility condition is

$$R_{\text{free}}(t)(1 - \varepsilon u_{\text{max}}) \leq R_t^* \leq R_{\text{free}}(t). \quad (\text{S71})$$

If $R_{\text{free}} = R_0$ is constant and R_t^* and u^* use the same Bernstein degree, then the coefficient mapping is exact:

$$\rho_j^* = R_0(1 - \varepsilon a_j), \quad a_j = \varepsilon^{-1}(1 - \rho_j^*/R_0). \quad (\text{S72})$$

This is the cleanest setting in which control-point height has an immediate intervention interpretation. For more complex interventions—vaccination, isolation, age-specific measures, vector control, or combinations—the map must be derived from the model rather than assumed.

One-point numerical sanity check. Suppose the no-new-intervention reproduction number is $R_{\text{free}} = 1.60$, the meeting target is $R_t^* = 0.90$, and a unit of standardized control delivers $\varepsilon = 0.80$ proportional transmission reduction. Equation (S70) gives

$$u^* = \frac{1 - 0.90/1.60}{0.80} = 0.546875.$$

Thus the model requires about 0.547 of the declared policy scale at that instant. This does not mean “54.7% school closure” unless the policy scale has been calibrated that way. If the same R_{free} applies to an entire Bernstein segment, each target coefficient ρ_j^* can be converted point-by-point using Eq. (S72). This arithmetic is intentionally simple enough to verify during an interactive meeting.

S4.5 Exact inverse tracking for SIR, SIRS, and SEIR targets

The target curve can be translated directly into a control when the state equations permit inversion. Consider the controlled SIR model

$$\dot{S} = -\beta_0(1 - \varepsilon u)SI/N, \quad (\text{S73})$$

$$\dot{I} = \beta_0(1 - \varepsilon u)SI/N - \gamma I, \quad (\text{S74})$$

$$\dot{R} = \gamma I. \quad (\text{S75})$$

Theorem 17 (Exact SIR target tracking). *Let $I^*(t) > 0$ be continuously differentiable and define*

$$q^*(t) = \dot{I}^*(t) + \gamma I^*(t), \quad (\text{S76})$$

$$S^*(t) = S_0 - \int_0^t q^*(s) ds. \quad (\text{S77})$$

If $q^* \geq 0$, $S^* > 0$, and

$$u^*(t) = \frac{1}{\varepsilon} \left[1 - \frac{Nq^*(t)}{\beta_0 S^*(t) I^*(t)} \right] \quad (\text{S78})$$

lies in $[0, u_{\text{max}}]$, then the controlled SIR system initialized at $(S_0, I^*(0), R_0)$ exactly follows $I(t) = I^*(t)$, with the consistent recovered state.

Proof. Substituting Eq. (S78) into the infection flow gives $\beta_0(1 - \varepsilon u^*)S^*I^*/N = q^*$. Therefore $\dot{S}^* = -q^*$ and $\dot{I}^* = q^* - \gamma I^*$, exactly matching the first two SIR equations; the recovered equation follows by conservation. $\square$

For SIRS with waning rate ω , first solve

$$\dot{R}^* = \gamma I^* - \omega R^*, \quad S^* = N - I^* - R^*, \quad (\text{S79})$$

and substitute these states into Eq. (S78). For SEIR with progression rate σ , a twice-differentiable infectious target yields

$$E^* = \frac{\dot{I}^* + \gamma I^*}{\sigma}, \quad (\text{S80})$$

$$q_E^* = \dot{E}^* + \sigma E^*, \quad (\text{S81})$$

$$S^*(t) = S_0 - \int_0^t q_E^*(s) ds, \quad (\text{S82})$$

$$u^*(t) = \frac{1}{\varepsilon} \left[1 - \frac{N q_E^*(t)}{\beta_0 S^*(t) I^*(t)} \right]. \quad (\text{S83})$$

Because a Bernstein target has analytic derivatives, Eqs. (S78)–(S83) can be recomputed essentially immediately after a target control point is moved.

S4.6 Vector-borne and age-structured inverse controls

For a simple human–vector system in which intervention u reduces vector-to-human transmission,

$$\dot{S}_h = -b_{vh,0}(1-u)S_h I_v / N_v, \quad (\text{S84})$$

$$\dot{I}_h = b_{vh,0}(1-u)S_h I_v / N_v - \gamma_h I_h, \quad (\text{S85})$$

$$\dot{I}_v = b_{hv}S_v I_h / N_h - \mu_v I_v, \quad (\text{S86})$$

with the corresponding susceptible-vector equation. Given I_h^* , solve the vector subsystem driven by that human target, define $q_h^* = \dot{I}_h^* + \gamma_h I_h^*$, and obtain

$$u^*(t) = 1 - \frac{q_h^*(t) N_v}{b_{vh,0} S_h^*(t) I_v^*(t)}, \quad (\text{S87})$$

provided the denominator is positive and the resulting control is feasible.

In an age-structured model, several policy levers can produce the same aggregate R_t . Decompose the contact matrix by setting,

$$C(\mathbf{u}) = C_H + (1 - u_S)C_S + (1 - u_W)C_W + (1 - u_O)C_O, \quad (\text{S88})$$

for home, school, work, and other/community contact matrices. The effective reproduction number is the spectral radius of the corresponding next-generation matrix. Given R_t^* , one can choose among nonunique policy combinations by solving

$$\mathbf{u}^*(t) = \arg \min_{0 \leq \mathbf{u} \leq \mathbf{u}_{\max}} \left\{ \sum_k c_k u_k^2 + \eta \|\mathbf{u} - \mathbf{u}_{\text{prev}}\|_2^2 \right\} \quad \text{s.t.} \quad \rho\{\mathcal{K}(t, \mathbf{u})\} = R_t^*. \quad (\text{S89})$$

The c_k values encode relative implementation or socioeconomic burdens. This makes nonidentifiability operationally useful: many mechanisms may achieve the epidemiological target, and optimization selects a feasible mixture according to declared policy preferences rather than pretending that an aggregate R_t uniquely identifies the intervention.

S4.7 Intervention-specific mechanistic controls: what a policy point means

A single scalar “intervention intensity” is useful for geometry but is not a universal biological quantity. Different interventions enter different terms of an epidemic model. We therefore recommend a mechanism-first translation: each displayed policy polygon $u_k(t) = \sum_j a_{kj} B_{j,n}$ controls a declared epidemiological quantity, and its coefficients are interpreted only after that quantity and its dose-response scale are specified. Lockdowns and contact NPIs primarily reduce effective contact or per-contact transmission; vaccination changes the susceptible or protected state over time; testing and isolation shorten community infectious duration; hospitalization is usually a health-system state and capacity constraint, but admission to a low-infectivity isolation state can also reduce onward transmission. These mechanisms should not be collapsed into one coefficient unless the purpose is explicitly an aggregate effective-control index.

Table S1: Mechanistic interpretation of common intervention polygons. $u_k(t) \in [0, 1]$ is a standardized decision variable; it is not automatically a percentage of a real policy. Calibration parameters must come from intervention, mobility, clinical, or program data.

Intervention	Mechanistic entry	Example mapping to transmission or states	Natural cost/burden input
Lockdown / contact NPI	$\beta(t) = \beta_0 \{1 - \varepsilon_L u_L(t)\}$ or setting-specific contact matrices	$R_t = R_{\text{free}}(t) \{1 - \varepsilon_L u_L(t)\}$ in the scalar multiplicative case	lost activity, mobility/contact reduction, compliance burden
Masking / ventilation / hygiene	per-contact transmission multiplier	$\beta = \beta_0 \prod_k \{1 - \varepsilon_k u_k\}$ or a calibrated log-linear alternative	program, equipment, adherence costs
Vaccination	$\dot{V} = \nu_{\text{max}} u_V(t) S - \dots$	with leaky efficacy e_V , $S_{\text{eff}} = S + (1 - e_V)V$ and $R_t \propto S_{\text{eff}}$	doses, delivery, adverse-event/program costs
Testing / isolation	additional removal $\alpha(t) = \alpha_0 + \alpha_{\text{max}} u_I(t)$	simple SIR-type $R_t = \beta S / [N \{\gamma + \alpha(t)\}]$	tests, tracing, isolation-days, compliance
Hospital admission / isolation	flow $h(t)I$ from I to low-infectivity H	$R_t = \frac{\beta S}{N} \frac{1 + \kappa h / \gamma_H}{\gamma_I + h}$ when H has relative infectiousness κ	beds, staffing, admission/isolation-days
Hospital capacity	state constraint rather than transmission control	$H(t) \leq H_{\text{max}}$ enters the feasible set or objective	occupancy, surge beds, staff, opportunity cost
Vector control	biting/transmission or vector mortality	$b(t) = b_0 \{1 - \varepsilon_b u_b\}$ and/or $\mu_v(t) = \mu_{v0} + \eta_v u_v$	spraying, source reduction, coverage

Pointwise inverse examples. Suppose a reference state has $R_{\text{free}} = 1.60$ and the meeting target is $R_t^* = 0.90$. For a contact NPI with $\varepsilon_L = 0.80$, Eq. (S70) requires $u_L = 0.547$. If instead transmission is reduced only by faster isolation, take $\gamma = 1/7 \text{ d}^{-1}$ and $\alpha_{\text{max}} = 0.25 \text{ d}^{-1}$. Since the uncontrolled infection pressure is $A = R_{\text{free}}\gamma$, the required added isolation rate is

$$\alpha^* = \frac{A}{R_t^*} - \gamma = 0.111 \text{ d}^{-1}, \quad u_I^* = \alpha^* / \alpha_{\text{max}} = 0.444. \quad (\text{S90})$$

For admission to an isolation/hospital state with $\gamma_I = 1/7$, $\gamma_H = 1/8$, relative infectiousness $\kappa = 0.05$, and $h_{\max} = 0.30 \text{ d}^{-1}$, write $r = R_t^*/A$. Solving

$$r = \frac{1 + \kappa h / \gamma_H}{\gamma_I + h} \quad (\text{S91})$$

gives

$$h^* = \frac{1 - r\gamma_I}{r - \kappa/\gamma_H} = 0.124 \text{ d}^{-1}, \quad u_H^* = 0.412. \quad (\text{S92})$$

For a leaky vaccine with efficacy against infection $e_V = 0.80$, an instantaneous effective-susceptible approximation would require protected susceptible fraction

$$c_V^* = \frac{1 - R_t^*/R_{\text{free}}}{e_V} = 0.547. \quad (\text{S93})$$

Vaccination is not truly instantaneous, however. If protection is accumulated over a 30-day campaign with $\nu(t) = 0.05u_V \text{ d}^{-1}$ and no depletion/waning during this short calculation, $c_V = 1 - \exp(-0.05u_V 30)$ gives $u_V \approx 0.528$. In the full model, vaccination must be optimized dynamically because today's doses change tomorrow's susceptible state.

Combined portfolios and nonidentifiability. Multiple interventions can produce the same R_t^* . A transparent composite model is

$$R_t(\mathbf{u}) = R_{\text{free}}(t) \{1 - \varepsilon_L u_L(t)\} \frac{S_{\text{eff}}(t; u_V)}{S_{\text{eff, free}}(t)} \frac{\gamma + \alpha_0}{\gamma + \alpha_0 + \alpha_{\max} u_I(t)}, \quad (\text{S94})$$

with additional factors when hospitalization/isolation or setting-specific controls are modeled explicitly. Because many $\mathbf{u}$ can satisfy $R_t(\mathbf{u}) = R_t^*$, the controller should select a portfolio by minimizing declared cost, inequity, implementation-speed, or feasibility penalties. This is not an identifiability failure to hide: it is the policy choice that the optimization layer must expose. Epidemic optimal-control studies have long considered vaccination, quarantine/isolation, and distancing controls [25, 26]; the control-polygon contribution is to make their time profiles low-dimensional, visually editable, and algebraically linked to bounds and costs.

S4.8 Cost, cost–benefit, and cost-effectiveness in coefficient space

A general objective can combine target tracking, health-system burden, and intervention burden:

$$J(\mathbf{a}) = w_Y \int_0^T \{Y(t) - Y^*(t)\}^2 dt + w_H \int_0^T [H(t) - H_{\max}]_+^2 dt + C_P(\mathbf{a}) + C_H(\mathbf{x}). \quad (\text{S95})$$

For a scalar policy, an interpretable policy-cost model is

$$C_P = c_1 \int_0^T u(t) dt + c_2 \int_0^T u(t)^2 dt + c_3 \int_0^T \dot{u}(t)^2 dt + c_4 \text{TV}(u), \quad (\text{S96})$$

where the first three terms are exactly or directly computable from theorem 15 and the final term is bounded by coefficient differences. Monetary health costs may include case, hospitalization, and mortality costs when defensible unit-cost data are available. Standard outputs can then include infections or hospitalizations averted, incremental cost-effectiveness ratios, benefit–cost ratios, or net monetary benefit. Cost-effectiveness analysis is already well established in epidemic optimal-control studies [26]; the proposed contribution is that the policy geometry exposed to decision makers is the same low-dimensional object used by the mathematical controller. No monetary interpretation is implied without locally calibrated cost inputs.

S4.9 Interactive high-level meetings and receding-horizon control

The geometric representation suggests a direct decision workflow. At a policymaker/decision maker review or decision support meeting: (1) update the reconstruction in Supplementary Section S1 using the newest surveillance data; (2) display the forecasting analysis in Supplementary Section S3 and its uncertainty; (3) ask decision makers to specify a target incidence, prevalence, hospitalization, or R_t curve by moving a small number of vertical control points; (4) map that target through the mechanistic model to the policy polygon; (5) flag infeasible targets immediately; (6) display intervention intensity, ramp rate, duration above a chosen threshold, projected cases/hospitalizations, and cost or cost-effectiveness metrics; and (7) compare alternative target polygons or policy mixtures on a Pareto frontier. A graphical interface can therefore make the control points genuine decision handles, while the equations remain underneath the visualization.

Only the first policy segment should normally be implemented. At the next review, states and parameters are re-estimated and the remaining target/control polygon is optimized again. This is receding-horizon or model predictive control, a strategy already used in epidemic intervention planning [27, 28]. In a piecewise-Bernstein implementation, the segment length is also an operational decision: a 7-day segment encodes weekly review, a 14-day segment biweekly review, and so forth. The framework therefore separates three concepts that can otherwise be conflated: coefficient height (intensity), coefficient differences (ramp/switching behavior), and segment duration (review frequency).

Algorithm 2 Interactive meeting-room decision loop

- 1: Update surveillance data and reconstruct the current control polygon by LS and Bayesian inference (Supplementary Section S1).
 - 2: Display both point/LS and Bayesian probabilistic forecasts, including important event scenarios (Supplementary Section S3).
 - 3: Let decision makers edit a small number of target control points for incidence, hospitalization, or R_t .
 - 4: Recompute Y^* , R_t^* , and analytic derivatives immediately.
 - 5: Compute the LS inverse policy and, from posterior draws, the Bayesian distribution of required policies/costs.
 - 6: Compress each required control to policy points $\mathbf{a}$ and display point estimates plus uncertainty bands.
 - 7: Flag violations of $0 \leq u \leq u_{\max}$, ramp limits, health-system constraints, or budget constraints.
 - 8: Display cases/hospitalizations averted, policy burden, cost, cost-effectiveness/net benefit, and Pareto alternatives.
 - 9: Save the chosen target and assumptions; implement only the first review segment, then repeat with new data.
-

S5 Complete synthetic numerical experiments

All simulations are reproducible from the accompanying code. Experiments 1–7 test reconstruction, geometry, inference, forecasting, and mechanistic embedding using synthetic data. Experiment 8 tests inverse control and policy geometry across SIR, SIRS, SEIR, vector-borne, and age-structured systems and includes normalized health–policy trade-offs. Experiment 9 uses a frozen Philippine COVID-19 reporting-date series as an empirical illustration; because that series contains cases but not intervention-effect or economic data, its policy-control panel is explicitly non-causal.




Figure S1: Bernstein transmission geometry. Control-point levels determine a convex-hull envelope, while neighboring differences determine the derivative.

S5.1 Experiment 1: control-polygon geometry

We used seven control points

$$\boldsymbol{\rho} = (0.78, 0.72, 0.95, 1.38, 1.32, 0.96, 0.82) \quad (\text{S97})$$

to illustrate threshold geometry and transmission momentum. The R_t curve is a convex combination of these values, while its derivative is a degree-five Bernstein curve whose coefficients are scaled first differences. Figure S1 shows the direct relation between the displayed control polygon and the derivative.

S5.2 Experiment 2: theorem-driven minimum control-point complexity

To verify the lower bound in corollary 8, we constructed degree-minimal Bernstein curves with exactly one, two, and four transversal crossings of $R_t = 1$. The corresponding true coefficient vectors were

$$(0.415, 1.715), \quad (1.325, 0.425, 1.525), \quad (1.258, 0.457, 1.746, 0.327, 1.398). \quad (\text{S98})$$

They therefore have $K = 2, 3, 5$ control points, exactly matching the theoretical floors $C + 1$. Each curve was observed at 121 equally spaced times with independent Gaussian noise of standard deviation 0.012 on the R_t scale. For each of 120 replicates, we fitted positive Bernstein curves with $K = 2, \dots, 8$ by constrained least squares. The selected model was the smallest K having the correct threshold-crossing count and residual RMSE ≤ 0.025 .

All 120 replicates recovered the theorem floor in each of the three scenarios (table S2 and figure S2). This experiment is a controlled verification, not evidence that the lower bound will always be selected from epidemic case data. A noise sensitivity experiment with 200 replicates per scenario and noise level showed that four-crossing recovery at the theorem floor declined from 1.00 at SD 0.025 to 0.98 at SD 0.05 and 0.76 at SD 0.08, while the simpler one- and two-crossing curves remained fully recoverable over the simulated range (figure S3).

S5.3 Experiment 3: renewal epidemic with an emerging variant

We next simulated a stochastic renewal epidemic. The generation interval was a discretized gamma distribution with mean 5 days and standard deviation 2 days, truncated at 20 days and renormalized.

Table S2: Controlled recovery of the structural threshold-crossing lower bound. Each scenario used 120 replicates at noise SD 0.012.

Scenario	True crossings C	True K	Theorem floor $C + 1$	Floor recovery
One crossing	1	2	2	1.00
Two crossings	2	3	3	1.00
Four crossings	4	5	5	1.00




Figure S2: Probability that a candidate number of control points achieves adequate geometric recovery in the theorem-verification experiment. The first adequate K coincides with the threshold-crossing lower bound for the low-noise design.

Historical R_t over 120 days used seven control points

$$(1.25, 1.31, 1.16, 0.93, 0.79, 0.84, 0.82). \quad (\text{S99})$$

Incidence followed Eq. (S7) with negative-binomial size $\phi = 50$. A six-control-point future baseline was attached over 60 days with the first two points satisfying the C^1 condition. The baseline then moved gradually toward approximately 0.86. A variant began at frequency $p_0 = 0.02$ with logistic growth rate $s = 0.12$ per day. In the main illustration its effective relative transmission multiplier was $a = 1.55$.

The historical segment was re-estimated with $K = 7$ and a second-difference penalty. Three operational forecasts were compared: (i) an event-informed Bernstein forecast using the variant signal; (ii) the same future polygon without the event adjustment; and (iii) a constant- R_t baseline based on recent transmission. We also display direct polynomial extrapolation of the historical Bernstein curve only as a diagnostic counterexample; it is not part of the proposed method because its convex-hull guarantee is lost outside the historical interval.

Figure S4 shows the latent R_t forecast and the resulting renewal-incidence forecasts. The example demonstrates how external evidence can cause the future control polygon to cross from a subcritical baseline into a supercritical region even when the historical endpoint is below one.




Figure S3: Noise sensitivity of selecting the theorem lower bound. Structural bounds remain mathematically valid at all noise levels, but data-driven recovery can fail when observation noise obscures threshold geometry.




Figure S4: Event-informed forecasting under a simulated emerging variant. The future forecast is a newly joined Bernstein segment, not an extrapolation beyond the historical Bernstein domain.

S5.4 Experiment 4: probabilistic forecast value across variant advantages

Forecast value was quantified over effective variant multipliers

$$a \in \{1.0, 1.2, 1.4, 1.6, 1.8\}. \quad (\text{S100})$$

For each multiplier, 100 stochastic future epidemics were generated conditional on the same historical trajectory. For every replicate and model, 250 predictive renewal trajectories were generated. We evaluated 14-, 28-, 42-, and 60-day forecasts using a WIS constructed from the predictive median and central 50%, 80%, and 90% intervals, following the proper interval-score framework of Bracher et al. [22].

At 42 days, mean WIS for event-informed, no-event, and constant- R_t forecasts was respectively 20.20, 20.31, and 20.14 when $a = 1.0$; there is correctly little value to an event adjustment when no event effect exists. At $a = 1.4$, the scores were 26.51, 35.33, and 46.51; at $a = 1.6$, 30.16, 57.58, and




Figure S5: Mean 42-day WIS versus effective variant transmission multiplier. Event information provides little gain when the multiplier is one, but becomes increasingly valuable as the shock departs from the no-event continuation.

71.62; and at $a = 1.8$, 34.35, 99.77, and 115.51. The advantage widened with horizon: for $a = 1.6$ at 60 days, mean WIS was 54.75 for the event-informed forecast versus 142.09 for the no-event forecast and 164.64 for the constant baseline (table S3 and figure S5). These experiments demonstrate conditional value when the event specification is informative; they do not show that the framework can detect an unknown future variant before surveillance evidence exists.

Table S3: Mean WIS over 100 stochastic futures at the 42-day horizon. Lower is better. Each forecast distribution used 250 predictive trajectories.

Variant multiplier a	Event-informed	No-event	Constant R_t	Gain vs no-event
1.0	20.20	20.31	20.14	0.5%
1.2	23.51	24.25	30.39	3.1%
1.4	26.51	35.33	46.51	25.0%
1.6	30.16	57.58	71.62	47.6%
1.8	34.35	99.77	115.51	65.6%

S5.5 Experiment 5: coefficient uncertainty propagated to the whole trajectory

To illustrate theorem 11, coefficient means were set to

$$(0.85, 0.78, 0.88, 1.20, 1.32, 1.06, 0.90) \quad (\text{S101})$$

with coefficient standard errors

$$(0.05, 0.06, 0.07, 0.10, 0.11, 0.08, 0.06). \quad (\text{S102})$$

Marginal 95% coefficient intervals were propagated using Eq. (S30). The resulting band contains every curve generated by independent coefficient choices within those intervals (figure S6). For




Figure S6: Coefficient-wise uncertainty envelope. Nonnegative Bernstein weights propagate lower and upper coefficient limits into a global pointwise bound for R_t over the complete fitted interval.

posterior inference, joint posterior simulation is preferable when coefficient dependence is substantial, but the deterministic envelope remains useful as a conservative visual certificate.

S5.6 Experiment 6: constrained least squares versus Bayesian renewal inference

We next tested whether the same Bernstein control polygon can be estimated with conventional least squares and with a fully probabilistic renewal likelihood. A 140-day historical trajectory used $K = 7$ true control points

$$\boldsymbol{\rho} = (1.27, 1.33, 1.15, 0.91, 0.76, 0.80, 0.84), \quad (\text{S103})$$

a gamma generation interval with mean 5 d and standard deviation 2 d discretized to 20 d, negative-binomial dispersion $\phi = 35$, and initial incidence 95. A 60-day held-out future was generated from a $K_F = 6$ C^1 -joined polygon with the same momentum/mean-reversion rule used in Eq. (S40); an emerging variant had true multiplier $a = 1.55$, takeover slope $s = 0.11 \text{ d}^{-1}$, and initial share $p_0 = 0.02$.

For constrained LS, Eq. (S10) used positivity bounds $0.05 \leq \rho_j \leq 4$, a second-difference penalty $\lambda = 3$, infectiousness-based weights, and 300 parametric bootstrap refits for 95% confidence bands. Bayesian inference used Eq. (S16) with $\sigma_0 = 0.55$, $\tau_2 = 0.17$, and $\log \phi \sim N(\log 30, 0.70^2)$. Four adaptive Metropolis-within-Gibbs chains each ran for 8,000 iterations with 3,000 burn-in iterations and thinning by two. All $\hat{R}$ values were below 1.002 and the minimum approximate effective sample size across sampled parameters exceeded 2,096. For forecasting, 3,000 future control polygons propagated historical posterior uncertainty, $N(0, 0.025^2)$ future-control innovations, $\log a \sim N(\log 1.50, 0.10^2)$, and $s \sim N(0.105, 0.018^2)$ truncated below 0.02. A subset of 1,200 posterior draws generated negative-binomial predictive incidence trajectories.

On the historical interval, R_t RMSE was 0.0310 for constrained LS and 0.0272 for the Bayesian posterior median. The LS 95% bootstrap band covered 94.3% of the true historical R_t values, while the pointwise Bayesian 95% credible band contained all historical R_t values in this realization (figure S7). The posterior median dispersion was 32.48, close to the simulated value 35. These are single-realization calibration diagnostics rather than general coverage guarantees.




Figure S7: Inference for the historical Bernstein R_t trajectory. Constrained least squares is accompanied by a parametric-bootstrap confidence band; Bayesian negative-binomial renewal inference is accompanied by a posterior credible band. The distinction is inferential: the LS band describes repeated-sampling uncertainty under the bootstrap model, while the Bayesian band is a posterior probability statement conditional on the specified prior and likelihood.


For the 60-day held-out future, the LS event-informed point forecast had R_t RMSE 0.116, whereas the Bayesian median forecast had RMSE 0.0693. The pointwise Bayesian 95% credible band contained all 60 simulated future R_t values in this realization, and the 95% posterior predictive interval contained 98.3% of held-out incidence days (figure S8). The posterior probability that R_t exceeded one at least once during the 60-day horizon was 0.977. Day-specific probabilities $P(R_{T+h} > 1 \mid Y, \mathcal{E})$ are shown in figure S9; unlike a binary threshold declaration from a single estimated curve, these probabilities retain uncertainty in historical control points, future-control innovations, variant advantage, variant takeover, and negative-binomial dispersion.

S5.7 Experiment 7: embedding the same control polygon in seven epidemic model classes


The renewal model is the primary forecasting formulation because it maps R_t directly to incidence through a generation-interval distribution. The control polygon is not restricted to renewal models, however. To test portability, we embedded the same positive Bernstein target

$$R_B(t) = \sum_{j=0}^7 \rho_j B_{j,7}(t/210), \quad \boldsymbol{\rho} = (1.16, 1.10, 0.91, 0.82, 1.09, 1.17, 0.96, 0.86), \quad (\text{S104})$$

over 210 days into seven distinct model/observation structures. For the six mechanistic models, the transmission coefficient or common transmission scale was recalibrated at each time using the appropriate next-generation expression, so the numerical effective reproduction number matched $R_B(t)$ to machine precision (maximum absolute error $< 5 \times 10^{-16}$ in the supplied verification code). This exercise illustrates mathematical compatibility, not identifiability: fitting a population-level



(a) Latent future R_t .



(b) Posterior predictive incidence.

Figure S8: Bayesian event-informed forecast from posterior control polygons. The R_t band is a 95% credible interval for the latent transmission trajectory; the incidence band is a 95% posterior predictive interval and additionally contains stochastic observation/process variability.




Figure S9: Posterior probability of supercritical transmission over the forecast horizon. The curve estimates $P(R_{T+h} > 1 | Y, \mathcal{E})$ from joint posterior/future-control draws rather than thresholding a single point forecast.

R_t does not by itself identify latent periods, waning rates, contact matrices, vaccine efficacy, or isolation effects.

Plotting convention for unequal compartment scales. Raw susceptible or uncontrolled compartments can be orders of magnitude larger than incidence, infectious prevalence, or a policy target. Overlaying them on one linear axis makes the scientifically relevant smaller trajectories effectively invisible. Throughout the examples we therefore report epidemic flows/prevalences per 100,000 when appropriate, plot population fractions separately, use a secondary axis only for a genuinely different unit such as R_t or vector prevalence, and use a separate full-scale companion panel when an uncontrolled epidemic is much larger than the target. No trajectory is rescaled independently within the same quantitative comparison.

S5.7.1 Example 7A: SEIR latent-stage dynamics

For population size N , the SEIR model is

$$\dot{S} = -\beta(t)SI/N, \quad \dot{E} = \beta(t)SI/N - \sigma E, \quad (\text{S105})$$

$$\dot{I} = \sigma E - \gamma I, \quad \dot{R} = \gamma I. \quad (\text{S106})$$

With complete progression from E to I , the instantaneous next-generation reproduction number is

$$R_e(t) = \frac{\beta(t)S(t)}{\gamma N}. \quad (\text{S107})$$

Hence Eq. (S20) becomes

$$\beta(t) = \frac{\gamma N R_B(t)}{S(t)}. \quad (\text{S108})$$

We used $N = 10^6$, $(S_0, E_0, I_0, R_0) = (999740, 160, 100, 0)$, mean latent period $1/\sigma = 3$ d, and mean infectious period $1/\gamma = 6$ d. The resulting incidence and infectious prevalence are shown in panel (a) of figure S10. The latent stage delays the incidence response relative to the control-polygon changes even though the instantaneous next-generation number is calibrated exactly [11, 29].

S5.7.2 Example 7B: SIRS with waning immunity

The SIRS equations are

$$\dot{S} = -\beta(t)SI/N + \omega R, \quad (\text{S109})$$

$$\dot{I} = \beta(t)SI/N - \gamma I, \quad (\text{S110})$$

$$\dot{R} = \gamma I - \omega R. \quad (\text{S111})$$

Again $R_e(t) = \beta(t)S(t)/(\gamma N)$, so we use the calibration in Eq. (S108). Parameters were $(S_0, I_0, R_0) = (799800, 200, 200000)$, $1/\gamma = 6$ d, and mean immunity duration $1/\omega = 120$ d. Because $R_B(t)$ is imposed as the effective reproduction number, waning immunity changes the susceptible pool and therefore the transmission coefficient required to realize the same R_t trajectory; it does not create an additional unmodeled threshold crossing. This distinction is useful when embedding an inferred R_t in a mechanistic model. In a causal analysis where $\beta(t)$ rather than R_t is held fixed, waning susceptibility would itself change $R_e(t)$ [29].

S5.7.3 Example 7C: vector-borne transmission

We used a Ross–Macdonald-type human–vector system [12]:

$$\dot{S}_h = -b_{vh}(t)S_h I_v / N_v, \quad (\text{S112})$$

$$\dot{I}_h = b_{vh}(t)S_h I_v / N_v - \gamma_h I_h, \quad (\text{S113})$$

$$\dot{R}_h = \gamma_h I_h, \quad (\text{S114})$$

$$\dot{S}_v = \mu_v N_v - b_{hv}(t)S_v I_h / N_h - \mu_v S_v, \quad (\text{S115})$$

$$\dot{I}_v = b_{hv}(t)S_v I_h / N_h - \mu_v I_v. \quad (\text{S116})$$

For infectious states (I_h, I_v) , the next-generation spectral radius is

$$R_e(t) = \left[\frac{b_{vh}(t)b_{hv}(t)S_h(t)S_v(t)}{N_v N_h \gamma_h \mu_v} \right]^{1/2}. \quad (\text{S117})$$

Let $b_{vh}(t) = c(t)b_{vh}^0$ and $b_{hv}(t) = c(t)b_{hv}^0$. Then

$$c(t) = \frac{R_B(t)}{\sqrt{b_{vh}^0 b_{hv}^0 S_h S_v / (N_v N_h \gamma_h \mu_v)}}. \quad (\text{S118})$$

We used $N_h = 10^6$, $N_v = 2 \times 10^6$, $(I_{h0}, I_{v0}) = (180, 1200)$, $\gamma_h = 1/6 \text{ d}^{-1}$, $\mu_v = 1/14 \text{ d}^{-1}$, $b_{vh}^0 = 0.35 \text{ d}^{-1}$, and $b_{hv}^0 = 0.25 \text{ d}^{-1}$. Human incidence and vector prevalence are shown in panel (c) of figure S10. This example demonstrates that the control polygon may calibrate a next-generation spectral radius even when transmission alternates between two host types rather than occurring through a single infectious compartment.

S5.7.4 Example 7D: age-structured transmission

For age groups $a = 1, \dots, G$,

$$\dot{S}_a = -S_a \lambda_a(t), \quad (\text{S119})$$

$$\dot{I}_a = S_a \lambda_a(t) - \gamma_a I_a, \quad (\text{S120})$$

$$\dot{R}_a = \gamma_a I_a, \quad (\text{S121})$$

$$\lambda_a(t) = \beta(t) \sum_{b=1}^G C_{ab} \frac{I_b}{N_b}. \quad (\text{S122})$$


The next-generation matrix at unit β has entries

$$A_{ab}(t) = \frac{S_a(t)C_{ab}}{N_b \gamma_b}, \quad (\text{S123})$$


so $\beta(t) = R_B(t)/\rho\{A(t)\}$ by proposition 3. We used three groups (0–19, 20–64, 65+ years) with population sizes (250000, 550000, 200000), common infectious period 6 d, initial infectious counts (80, 120, 30), and contact matrix

$$C = \begin{pmatrix} 3.5 & 1.2 & 0.4 \\ 0.9 & 3.2 & 0.8 \\ 0.4 & 1.0 & 2.0 \end{pmatrix}. \quad (\text{S124})$$


These values are illustrative rather than estimates for a particular country; empirical work can substitute measured or synthetic contact matrices such as those developed by Prem et al. [30]. The age-specific incidence curves differ substantially even though the population-level next-generation spectral radius follows the same $R_B(t)$ (panel (d) of figure S10).




(a) SEIR.



(b) SIRS with waning immunity.



(c) Vector-borne transmission.



(d) Age-structured transmission.

Figure S10: Mechanistic embeddings of the common Bernstein target in Eq. (S104). The same effective reproduction-number geometry can be mapped into latent-stage, waning-immunity, vector-borne, and multigroup systems once the corresponding next-generation structure is specified.

S5.7.5 Example 7E: negative-binomial overdispersed observations

The control polygon describes latent transmission and need not force observations to be Poisson. Using the SEIR expected daily incidence μ_t from Example 7A, we simulated

$$Y_t \mid \mu_t \sim \text{NB}(\mu_t, \phi), \quad \phi = 12, \quad (\text{S125})$$

with

$$\text{Var}(Y_t \mid \mu_t) = \mu_t + \frac{\mu_t^2}{12}. \quad (\text{S126})$$

Thus the latent R_t and mechanistic mean remain unchanged while reported counts can show substantial day-to-day scatter (panel (a) of figure S11). This separates transmission geometry from observation overdispersion, matching the role of the negative-binomial layer in Eq. (S7).

S5.7.6 Example 7F: vaccination with imperfect protection

Let V denote vaccinated individuals who remain susceptible to breakthrough infection with relative susceptibility $1 - \varepsilon$. The model is

$$\dot{S} = -\beta SI/N - \nu(t)S, \quad (\text{S127})$$

$$\dot{V} = \nu(t)S - (1 - \varepsilon)\beta VI/N, \quad (\text{S128})$$

$$\dot{I} = \beta I\{S + (1 - \varepsilon)V\}/N - \gamma I, \quad (\text{S129})$$

$$\dot{R} = \gamma I. \quad (\text{S130})$$

The effective reproduction number is

$$R_e(t) = \frac{\beta(t)}{\gamma N} \{S(t) + (1 - \varepsilon)V(t)\}, \quad (\text{S131})$$

and therefore

$$\beta(t) = \frac{\gamma N R_B(t)}{S(t) + (1 - \varepsilon)V(t)}. \quad (\text{S132})$$

We used $(S_0, V_0, I_0, R_0) = (999780, 0, 220, 0)$, $1/\gamma = 6$ d, vaccine efficacy against infection $\varepsilon = 0.75$, and vaccination rate

$$\nu(t) = \begin{cases} 0, & t < 45, \\ 0.004 \text{ d}^{-1}, & 45 \leq t < 150, \\ 0.001 \text{ d}^{-1}, & t \geq 150. \end{cases} \quad (\text{S133})$$

The simulation reports total and breakthrough incidence together with the vaccinated-susceptible fraction (panel (b) of figure S11). In a future-event forecast, a vaccination campaign can alternatively be translated directly into downward shifts of future R_t control-point priors; the explicit compartmental representation shown here is useful when coverage and breakthrough composition are themselves outcomes of interest.

S5.7.7 Example 7G: hospitalization or isolation with residual infectiousness

We considered an SEIHR model in which individuals leave community infectiousness either by recovery or by hospitalization/isolation:

$$\dot{S} = -\beta S(I + \kappa H)/N, \quad (\text{S134})$$

$$\dot{E} = \beta S(I + \kappa H)/N - \sigma E, \quad (\text{S135})$$

$$\dot{I} = \sigma E - (\gamma_I + \alpha)I, \quad (\text{S136})$$

$$\dot{H} = \alpha I - \gamma_H H, \quad (\text{S137})$$

$$\dot{R} = \gamma_I I + \gamma_H H. \quad (\text{S138})$$

Here $0 \leq \kappa \leq 1$ is infectiousness while hospitalized/isolated relative to community infectiousness. A newly infected individual contributes an expected infectious-time equivalent

$$D_I = \frac{1}{\gamma_I + \alpha} + \kappa \frac{\alpha}{(\gamma_I + \alpha)\gamma_H}. \quad (\text{S139})$$

Therefore


$$R_e(t) = \frac{\beta(t)S(t)}{N} D_I, \quad \beta(t) = \frac{NR_B(t)}{S(t)D_I}. \quad (\text{S140})$$

We used initial state $(S_0, E_0, I_0, H_0, R_0) = (999740, 140, 120, 0, 0)$. The rates were $\sigma = 1/3$, $\gamma_I = 1/8$, $\alpha = 1/5$, and $\gamma_H = 1/8 \text{ d}^{-1}$, with residual infectiousness $\kappa = 0.08$. The model generates a hospital/isolation occupancy trajectory in addition to infection incidence (panel (c) of figure S11), demonstrating how the same control-polygon R_t can be connected to a health-system outcome.


Table S4: Summary of the seven additional worked examples. Every mechanistic example uses the common control polygon in Eq. (S104); the negative-binomial example adds an observation layer to the SEIR mean.

Example	How the Bernstein curve enters	Primary additional scientific quantity
SEIR	$\beta = \gamma NR_B/S$	Latent-stage delay and infectious prevalence
SIRS	$\beta = \gamma NR_B/S$	Waning immunity and implied transmission rate
Vector-borne	Common scale on human–vector and vector–human transmission	Human incidence and vector prevalence
Age-structured	$\beta = R_B/\rho(A)$	Age-specific incidence under a contact matrix
Negative binomial	$Y_t \sim \text{NB}(\mu_t, 12)$	Reporting/observation overdispersion
Vaccination	$\beta = \gamma NR_B/[S + (1 - \varepsilon)V]$	Coverage and breakthrough incidence
Hospital/isolation	$\beta = NR_B/(SD_I)$	Hospital/isolation occupancy and residual infectiousness


These examples clarify the scope of the geometric theorems. The threshold, finite-difference, and complexity results apply directly whenever the Bernstein curve itself is the effective reproduction number. If instead one places a Bernstein polynomial on a primitive parameter such as $\beta(t)$, mobility, or contact intensity and derives R_t afterward, the control points retain their numerical-basis role but their direct relation to the epidemic threshold $R_t = 1$ no longer follows from theorem 1; the next-generation mapping must then be analyzed explicitly.



(a) Negative-binomial overdispersion.



(b) Vaccination with breakthrough infection.



(c) Hospitalization/isolation with residual infectiousness.

Figure S11: Observation and intervention extensions. Overdispersion is added without changing the latent transmission curve, whereas vaccination and isolation enter mechanistically through the susceptible-equivalent pool or the expected infectious-time contribution.

S5.8 Experiment 8: step-by-step inverse control, interactive targets, and cost trade-offs

We next reverse the usual modeling direction. The common planning horizon is $T = 120$ days and $N = 10^6$. A hypothetical decision meeting specifies the log-infectious-prevalence target by seven control-polygon ordinates corresponding to (200, 350, 1000, 2500, 1800, 800, 300) infectious persons. Because a Bernstein curve is not an interpolant, these are control ordinates rather than required point values; the resulting target reaches a maximum of 1,354 infectious persons. The steps are: (i) evaluate the Bernstein target and its analytic derivatives; (ii) invert the mechanistic equations to obtain the exact continuous-time control $u^*(t)$; (iii) compress u^* into a small Bernstein policy polygon by bounded least squares; (iv) resimulate the mechanistic system using the fitted policy polygon; and (v) compute policy-geometry and health-burden summaries.

SIR. We use $\gamma = 1/7 \text{ d}^{-1}$, $\beta_0 = 0.34 \text{ d}^{-1}$, $S_0 = N - I^*(0)$, and $\varepsilon = 1$. Equation (S78) requires $u^* \in [0.464, 0.720]$. A seven-point policy polygon has coefficients

$$(0.4975, 0.4366, 0.4384, 0.5325, 0.6460, 0.6999, 0.7200), \quad (\text{S141})$$

with control-function RMSE 2.8×10^{-5} and normalized target-tracking RMSE 5.6×10^{-5} . Its exact mean control from Eq. (S58) is 0.567. Without control, infectious prevalence peaks near 2.16×10^5 in this illustrative parameterization.

SIRS. We use $\gamma = 1/7 \text{ d}^{-1}$, mean immunity duration $1/\omega = 120 \text{ d}$, $R(0) = 100,000$, and $\beta_0 = 0.40 \text{ d}^{-1}$. Solving the target-consistent recovered state before inversion gives $u^* \in [0.502, 0.754]$. The fitted seven-point polygon has mean intensity 0.606, control-function RMSE 2.4×10^{-5} , and target-tracking normalized RMSE 5.2×10^{-5} .

SEIR. We use latent period $1/\sigma = 3 \text{ d}$, infectious period $1/\gamma = 7 \text{ d}$, and $\beta_0 = 0.42 \text{ d}^{-1}$. The target derivatives determine E^* and the required exposure flow through Eq. (S83). The exact control range is $[0.515, 0.809]$; the seven-point approximation has mean intensity 0.644 and target-tracking normalized RMSE 2.1×10^{-4} .

Vector-borne. Human population is $N_h = 10^6$, vector population $N_v = 2 \times 10^6$, $\gamma_h = 1/7 \text{ d}^{-1}$, vector mortality $\mu_v = 1/14 \text{ d}^{-1}$, fixed human-to-vector coefficient $b_{hv} = 0.32$, baseline vector-to-human coefficient $b_{vh,0} = 0.12$, and $I_v(0) = 1000$. The human target is 0.7 times the common prevalence target. Solving the vector dynamics conditional on that target and applying Eq. (S87) gives $u^* \in [0.489, 0.908]$. An eight-point policy polygon tracks the target with normalized RMSE 3.7×10^{-4} .

Age-structured control. Three age groups have populations (250,000, 550,000, 200,000), infectious period 6 d, and initial infections (80, 120, 30). We use

$$\begin{aligned} C_H &= \begin{pmatrix} 1.3 & 0.3 & 0.2 \\ 0.25 & 1.1 & 0.25 \\ 0.2 & 0.3 & 1.0 \end{pmatrix}, & C_S &= \begin{pmatrix} 3.2 & 0.7 & 0.05 \\ 0.5 & 0.15 & 0.03 \\ 0.02 & 0.03 & 0 \end{pmatrix}, \\ C_W &= \begin{pmatrix} 0.05 & 0.2 & 0.02 \\ 0.1 & 2.8 & 0.35 \\ 0.02 & 0.25 & 0.2 \end{pmatrix}, & C_O &= \begin{pmatrix} 0.5 & 0.5 & 0.15 \\ 0.35 & 0.8 & 0.35 \\ 0.15 & 0.4 & 0.5 \end{pmatrix}. \end{aligned}$$

The transmission scale is calibrated to initial uncontrolled $R_t = 2$. The desired reproduction-number polygon has six points (1.35, 0.95, 0.72, 0.82, 1.02, 0.88). At each day Eq. (S89) is solved using relative squared-cost weights (1.3, 1.8, 0.7) for school, work, and other/community reductions, with a small penalty on changes from the previous day's policy. The optimized setting-specific controls are then compressed into seven-point Bernstein polygons. The achieved R_t has RMSE 5.9×10^{-5} relative to the target. The school, work, and other controls range approximately over $[0.460, 0.763]$, $[0.311, 0.596]$, and $[0.539, 0.950]$, respectively. The values are illustrative relative burdens, not empirical estimates of real policy costs.




Figure S12: Policymaker-facing inverse design. A decision maker moves a small number of target-epidemic control points (A). Because the uncontrolled SIR peak is over two orders of magnitude larger than the target, a separate full-scale panel (B) preserves that context without flattening the target-scale comparison. The mechanistic model translates the target into the required effective transmission-reduction policy and its Bernstein control points (C). Alternative policy intensities are then displayed against health benefit (D). The plotted costs are normalized illustrative burdens, not monetary Philippine estimates.

To illustrate costing without inventing currency values, we scale the reference SIR policy polygon from 0 to 1.2 and define the normalized objective

$$J_\lambda = \frac{\text{cumulative infections}}{N} + \lambda \frac{1}{T} \int_0^T u(t)^2 dt, \quad (\text{S142})$$




Figure S13: Step-by-step inverse-control examples for SIR, SIRS, SEIR, and vector-borne systems. Each panel overlays the prescribed infectious-prevalence target, the trajectory achieved after compressing the exact inverse control into a small Bernstein policy polygon, and the resulting policy intensity.




Figure S14: Age-structured multilevel control. The left panel compares desired and achieved aggregate R_t . The right panel shows the minimum-cost mixture of school, work, and other/community contact reductions and their fitted control points. Home contacts are not controlled in this example.

for $\lambda \in \{0.05, 0.15, 0.40\}$. The uncontrolled simulation generates 875,261 cumulative infections over the horizon. Increasing intervention intensity averts infections but increases policy burden; the preferred point changes with λ (figure S15). In an applied analysis, Eq. (S142) should be replaced or supplemented by locally estimated economic costs, DALYs/QALYs, hospital costs, or net monetary benefit rather than treating λ as known.




Figure S15: Health-policy trade-off in the illustrative SIR control experiment. Left: mean-square intervention burden versus infections averted. Right: normalized total loss for three policy-burden weights. This figure demonstrates how moving or scaling control points can update a decision frontier immediately; it is not an empirical cost-effectiveness estimate.

Table S5: Tangible SIR cost-benefit examples for $\lambda = 0.15$. Policy scale 1.0 is the reference Bernstein polygon obtained from inverse target tracking. Values are deterministic outputs of the synthetic experiment and are dimensionless except for infection counts.

Policy scale	Mean control	Mean-square burden	Cumulative infections	Infections averted	$J_{0.15}$
0.0	0.000	0.000	875,261	0	0.8753
0.5	0.284	0.082	672,079	203,182	0.6845
1.0	0.567	0.330	13,762	861,499	0.0632
1.2	0.681	0.475	2,161	873,100	0.0734

The table also illustrates why “stronger” is not automatically “better” once policy burden is priced: moving from scale 1.0 to 1.2 averts additional infections but increases $J_{0.15}$ because the marginal health gain is smaller than the added quadratic policy burden under that illustrative preference weight.

S6 Additional guidance for selecting the number of control points

The theorem floor and the statistical optimum solve different problems. A practical selection procedure should therefore have two stages.

First, define only the geometric features that are scientifically resolvable: the minimum required number of $R_t = 1$ crossings C_{req} and/or interior extrema E_{req} . These may come from a preliminary robust estimate, prior knowledge about distinct transmission regimes, or a feature-detection algorithm

with uncertainty. Begin at

$$K_{\text{floor}} = \max\{2, C_{\text{req}} + 1, E_{\text{req}} + 2\}. \quad (\text{S143})$$

Second, fit increasing K and choose the smallest model that achieves prespecified adequacy: acceptable residual structure, threshold-feature stability, and out-of-sample forecast score. For LS fits, blocked or rolling-origin cross-validation is preferred to ordinary random-fold validation because time ordering must be respected. For Bayesian fits, leave-one-out/WAIC calculations can be useful for retrospective comparison, but rolling-origin posterior predictive scoring is more directly aligned with forecasting. Information criteria should supplement rather than replace feature and predictive adequacy. This is analogous in spirit to selecting spline complexity, but the theorem floor prevents searching models that are structurally incapable of the required geometry.

Degree elevation creates nonminimal Bernstein representations of the same polynomial, so a large K is not intrinsically more mechanistic. The meaningful object is the smallest control-point complexity needed under a declared basis, time window, data resolution, penalty/prior, and adequacy criterion. As with any smoothing problem, this operational minimum is design dependent.

S7 Additional interpretation and relationship to growth rates

If infectious prevalence $I(t)$ followed the simple continuous-time equation

$$\frac{dI}{dt} = \gamma I \{R_e(t) - 1\}, \quad (\text{S144})$$

then

$$\frac{d \log I}{dt} = \gamma \{R_e(t) - 1\}. \quad (\text{S145})$$

For incidence data, however, growth rate and R_t are linked through the generation-interval distribution rather than this simple identity [15, 16]. The renewal framework is therefore preferable for translating a Bernstein R_t trajectory into incidence. Control-point momentum $\Delta \rho_j$ should be interpreted as the change in the latent reproduction-number curve, while the implied incidence growth rate depends on recent infectiousness and the generation interval.

This distinction also clarifies the relation to Shin et al. [31]. Their method learns regularity in epidemic growth-rate trajectories to form forecasting priors. Our framework instead parameterizes R_t as a control polygon, uses exact Bernstein shape constraints, and propagates it through a renewal mechanism. The approaches are complementary: empirical priors learned from historical growth geometry could be used to inform the future control-point prior in Eq. (S40).

S8 Extended Philippine COVID-19 analyses, 2020–2023

We finally illustrate the framework using the frozen reporting-date series supplied with the reproducibility package. It contains 1,178 daily observations from 30 January 2020 through 21 April 2023 and sums to 4,087,964 reported cases; the cumulative endpoint agrees with the contemporaneous ASEAN BioDiaspora situational report [32]. Bansilan and Rabajante provide a complementary descriptive analysis and timeline of the Philippine pandemic, including transmission, testing, mobility, distancing, lockdown, quarantine, isolation, vaccination, and other response dimensions [33]. Reported cases are not treated as infections, and this example does not adjust for changes in ascertainment, testing, delays, immunity, or variant-specific generation intervals. It is therefore an illustration of the computational workflow rather than a causal reconstruction of Philippine transmission.

Historical reconstruction. We use a 7-day trailing mean of reported cases, a discretized gamma generation interval with mean 5 d and SD 2 d truncated at 20 d, and the renewal ratio $\tilde{R}_t = Y_t/\Lambda_t$ after the first 20 days. The LS reconstruction uses a 36-point positive penalized Bernstein curve with second-difference penalty 18 and square-root incidence weights; a 14-day moving-block residual bootstrap provides a descriptive 95% confidence band. In parallel, a Bayesian negative-binomial renewal fit uses positive Bernstein controls, a second-difference prior on log controls, and a Laplace posterior approximation for the long empirical series; the full MCMC formulation is given in Supplementary Section S1 and the inference experiment in Supplementary Section S5. The two reconstructions are intentionally shown together in Fig. S22. The large control-point count is an expressive multiyear representation, not a universal optimum.




Figure S16: Philippine COVID-19 illustration, 30 January 2020–21 April 2023. Top: daily reporting-date cases, 7-day trailing mean, and the mean reconstructed through the fitted Bernstein renewal model. Bottom: empirical renewal ratios, the 36-point Bernstein R_t curve, control points, and a moving-block bootstrap band. The model is descriptive and inherits reporting and generation-interval assumptions.

Terminal 60-day pseudo-forecast with the adaptive extension. We hold out 21 February–21 April 2023. The original joined-Bernstein forecast has MAE 121 cases/day against the held-out 7-day series and the constant- R_t comparator has MAE 110. The peak-aware history-only continuation reduces MAE to 75 cases/day. However, all three point forecasts still place their within-horizon maximum at the beginning of the holdout, whereas the observed holdout maximum occurs on 21 April. This is an important distinction: the adaptive continuation can improve curve error when

local dynamics are informative, but it cannot time a later resurgence that is not yet visible at the forecast origin.

Peak-aware rolling selection and out-of-selection evaluation. The adaptive history-only rule was selected using only three pre-Omicron rolling origins (24 March, 19 June, and 14 September 2021) and the peak-aware criterion in Eq. (S44). We then froze its hyperparameters and evaluated six later origins from December 2021 through February 2023 at 14-, 28-, 42-, and 60-day horizons. Across these 24 out-of-selection origin–horizon comparisons, the adaptive Bernstein forecast had lower MAE than the original joined-Bernstein forecast in 87.5% of comparisons and lower MAE than the constant- R_t comparator in 75%. The overall median MAE ratio was 0.697 relative to the joined forecast and 0.647 relative to the constant- R_t forecast. At 14 days the median adaptive/constant ratio remained above one, so the improvement is not universal; the largest gains appear at 28–60 days. Figure S17 and Table S6 report the reserved evaluation origins only.




Figure S17: Philippine COVID-19 out-of-selection rolling evaluation after peak-aware hyperparameter selection on three earlier pre-Omicron origins. Left: median MAE for peak-aware adaptive Bernstein, original joined Bernstein, and constant- R_t forecasts across six later origins. Right: median MAE ratios of the adaptive method to each comparator; values below one favor the adaptive method. The adaptive rule is therefore evaluated on origins not used to choose its history-only shape parameters.

Table S6: Six-origin out-of-selection Philippine rolling evaluation. Win fractions are the fractions of the six evaluation origins at which peak-aware adaptive Bernstein had lower MAE than the named comparator.

Horizon (d)	Adaptive MAE	Joined MAE	Constant MAE	A/J ratio	A/C ratio	Win vs J	Win vs C
14	47.6	138.3	65.1	0.484	1.281	0.833	0.500
28	168.1	354.4	181.6	0.584	0.718	1.000	0.833
42	229.2	480.1	446.1	0.777	0.527	1.000	0.833
60	270.4	765.8	708.7	0.940	0.647	0.667	0.833

Sensitivity. We also varied the generation-interval mean over 4, 5, and 6 days and historical complexity over $K \in \{24, 36, 48\}$. The full reconstruction sensitivity is reported in Supplementary

Table S13. These analyses reinforce the intended interpretation: Bernstein geometry is a transparent forecasting scaffold, while skill depends on how future control points are informed. Local growth, event-conditioned priors, richer surveillance streams, and formal rolling-origin model selection are substantive parts of an operational implementation rather than cosmetic refinements.

Historical-origin and peak-aware wave forecasts. We next separate three questions that are often conflated. First, can a curve observed only through 31 December 2020 anticipate the large April 2021 wave? The answer remains no. The observed 7-day series peaks on 15 April 2021 at 10,815 cases/day. The original joined forecast peaks immediately at about 1,398 cases/day, the peak-aware cases-only rule peaks immediately at about 1,233 cases/day, and the constant- R_t comparator also fails. No rearrangement of the already-observed curve creates information about a future structural shock that has not yet left a signal.

Second, once the 2021 wave is visibly growing, does preserving the terminal signal reduce peak error? Using an origin of 20 March 2021 and no post-origin case data, the peak-aware cases-only forecast predicts a peak on 8 April at 11,348 cases/day, compared with the observed 15 April peak of 10,815. Its peak-date error is 7 days and peak-height MARE is 4.9%, with MAE 575 cases/day. The original joined forecast has a 31-day peak-date error, peak-height MARE 180%, and MAE 11,005; the constant- R_t comparator is still poorer. This experiment illustrates the purpose of Eq. (S41): when a wave signal exists, it should not be erased by a long fitting window.

Third, using all observations through 31 December 2021, can the January 2022 Omicron-era peak be forecast more faithfully? The observed 7-day series peaks on 18 January at 34,808 cases/day. The original joined forecast still peaks on 4 January at only 413 cases/day, while the cases-only adaptive rule still underestimates the unprecedented surge. This is precisely where external event information matters. Omicron had been detected in the Philippines by December 2021, and genomic analyses identify BA.1 and BA.2 before the January surge, with sustained BA.2 transmission [34]. The event-balanced LS scenario peaks one day late with 14.4% peak-height error. When the user explicitly switches to the peak-priority objective in Eq. (S45), the illustrative event-conditioned LS profile peaks on 18 January at 35,193 cases/day—zero timing error and 1.1% peak-height error. Its scenario-centered Bayesian counterpart also peaks on 18 January at about 35,193 cases/day (1.1% peak-height error), while the 95% posterior predictive interval retains uncertainty around that declared event scenario. Scenario centering is performed on the latent future- R_t distribution using the deterministic event curve as the prior median; it does not use post-origin case outcomes. These peak-priority settings are presented as an event-conditioned scenario/decision profile, not as evidence that cases alone identified Omicron biology or that zero future peak error can be guaranteed. The separate retrospective-oracle diagnostic exists only to quantify representational capacity and is excluded from forecast-skill summaries.

Figure S18 overlays the competing forecasts so the gain in peak timing and height is directly visible. Figure S19 retains the broader scenario-envelope stress test to show uncertainty in event amplitude and duration.

Flexible priority profiles and full epicurve overlays. Figure S20 shows the event-balanced and event peak-priority forecasts directly against the observed Omicron-era epicurve and compares peak-date and peak-height errors. The peak-priority profile changes only declared forecast priorities/scenario assumptions; it does not alter the historical observations. Figure S21 then overlays forecasts against the original held-out epicurve at four representative origins so that timing errors, amplitude errors, and phase-specific successes/failures are visually inspectable rather than summarized only by scalar scores.

Table S7: Philippine peak-focused retrospective forecasts. Peak values use the 7-day reporting-date series. Event-informed Omicron rows are explicitly conditional scenarios rather than biological estimates or prospective validation.

Origin / target	Forecast	MAE	Obs. peak	Pred. peak	Date error	Height	MARE
31 Dec 2020 / Apr 2021	joined Bernstein	3784	10,815	1,398	104 d		0.871
31 Dec 2020 / Apr 2021	peak-aware cases-only	4073	10,815	1,233	104 d		0.886
20 Mar 2021 / Apr 2021	joined Bernstein	11,005	10,815	30,279	31 d		1.800
20 Mar 2021 / Apr 2021	peak-aware cases-only	575	10,815	11,348	7 d		0.049
31 Dec 2021 / Jan 2022	joined Bernstein	13,472	34,808	413	14 d		0.988
31 Dec 2021 / Jan 2022	peak-aware cases-only	9892	34,808	10,482	11 d		0.699
31 Dec 2021 / Jan 2022	event-balanced LS	–	34,808	29,793	1 d		0.144
31 Dec 2021 / Jan 2022	event peak-priority LS	–	34,808	35,193	0 d		0.011
31 Dec 2021 / Jan 2022	event peak-priority Bayesian central	–	34,808	35,193	0 d		0.011

All three goals with LS and Bayesian inference. The empirical example is also run end-to-end with both inferential philosophies. Figure S22 shows (A) LS and Bayesian reconstruction, (B) LS and Bayesian Omicron-era forecasting, and (C) LS and Bayesian hypothetical control. For the control panel, the meeting target is $R_t \leq 0.90$. The LS curve is the pointwise additional effective suppression implied by the LS event forecast. The Bayesian curve and 95% credible band push posterior future- R_t draws through the same inverse map. A purely illustrative three-lever portfolio then decomposes the required effect into standardized contact-NPI, isolation, and vaccination-protection indices by minimizing quadratic burden with declared relative cost weights. These component curves are not estimates of historical Philippine policy effects; they demonstrate how intervention-specific mappings can be attached once credible local effect and cost data are supplied.

Illustrative target-control translation. For the same forecast origin, we add a non-causal meeting scenario: keep the future effective reproduction number at or below 0.90. Relative to the model’s no-new-intervention event forecast, Eq. (S70) with $\varepsilon = 1$ implies an additional effective-transmission reduction reaching about 0.808 and averaging about 0.219 over the 60-day horizon in this deliberately stringent $R_t \leq 0.90$ scenario. The scenario-centered Bayesian median gives essentially the same central requirement, with posterior uncertainty shown in Figure S22. These numbers are not estimates of the effect of any Philippine policy: the case series contains neither intervention exposure nor an identified dose–response function. In the hypothetical portfolio demonstration, the same required suppression can be allocated among contact NPIs/lockdown, faster testing–isolation, vaccination-derived protection, or admission to a low-infectivity isolation state according to the mechanistic mappings in Table S1 and declared cost weights. Hospital occupancy itself remains a state/capacity constraint unless hospitalization also changes infectious contact. The panel therefore demonstrates how a target chosen in a meeting can be translated first into required effective suppression and then, only after calibration, into intervention-specific actions.




Figure S18: Philippine forecast comparisons emphasizing peak timing and height. A: data through 31 December 2020 do not contain the future driver of the April 2021 wave, so all cases-only point forecasts fail. B: with a 20 March 2021 origin, the target-scale panel shows how the peak-aware cases-only rule preserves the already-visible growth signal. C: a separate full-scale panel retains the exploding comparator without covering B. D: at 31 December 2021, the cases-only adaptive rule improves on the globally smoothed joined forecast but still underestimates the Omicron-era amplitude; the explicitly variant-informed high-growth scenario substantially reduces timing and height error. E: peak-date and peak-height errors for the forecastable wave origins. No post-origin case data are used to construct the forecasts shown in A–C.




Figure S19: Historical-origin Philippine stress tests and event-scenario envelope. A–B: data through 31 December 2020 are insufficient to anticipate the April 2021 wave. C–D: a case-history-only forecast through 31 December 2021 misses the January 2022 Omicron-era surge, while a family of temporary event-conditioned R_t shocks makes rapid resurgence possible. The shaded region is a scenario envelope, not a confidence or credible interval.




Figure S20: Philippine Omicron-era priority forecasts. A: LS event-balanced, LS event peak-priority, and Bayesian event peak-priority forecasts on a policy-relevant scale. B: the same Omicron forecast on the full predictive scale, placed in a separate panel rather than an inset. C–D: peak-date and peak-height errors. The event peak-priority profile explicitly gives larger decision weight to timing and height. It is a conditional scenario profile, not a claim that the peak can be known from cases alone.




Figure S21: Forecasts overlaid with the original Philippine epicurve, including 60 days of pre-origin context, at multiple historical origins. The 2020-only panel shows the information boundary before the 2021 wave driver appears; the March 2021 panel shows a wave once terminal growth is visible; the December 2021 panel includes the Bayesian event peak-priority Omicron scenario; and the terminal 2023 panel illustrates a low-incidence phase. A fifth companion panel displays the Omicron forecast on the full predictive scale, outside the main comparison. These overlays are intended to make delays and amplitude errors visible rather than relying only on aggregate metrics.




Figure S22: Philippine COVID-19 demonstration of the paper's three goals with both LS and Bayesian engines. A: reconstruct the observed epicurve. B: forecast the Omicron-era epicurve on a policy-relevant scale. C: display the same forecast on the full predictive scale without an inset. D: translate a hypothetical target $R_t \leq 0.90$ into LS and Bayesian distributions of required effective suppression and one illustrative minimum-burden portfolio across contact NPIs, isolation, and vaccination-protection levels. Control components are hypothetical standardized indices, not causal effect estimates for actual Philippine policies.




Figure S23: Philippine 60-day pseudo-forecast and illustrative target-control translation. Top: held-out 7-day case series, joined-Bernstein forecast, bootstrap band, and constant- R_t baseline. Bottom: forecast R_t , the meeting target $R_t \leq 0.90$, and the additional effective suppression required relative to the no-new-intervention forecast. The lower policy curve is a mathematical requirement, not a causal estimate for a particular intervention.

S9 Complete simulation settings

S9.1 Equation compendium used in all worked examples

For reproducibility, this subsection collects in one place the complete dynamical and observation equations used by the worked examples. The same positive Bernstein curve

$$R_B(t) = \sum_{j=0}^{K-1} \rho_j B_{j,K-1}(t/T), \quad \rho_j > 0, \quad (\text{S146})$$

is interpreted as the effective reproduction-number trajectory unless explicitly stated otherwise. The model-specific transmission coefficient is recalibrated from $R_B(t)$ through the appropriate next-generation expression. Thus the geometric driver is shared, while latent states and biological interpretation remain model specific.

Renewal reconstruction and forecasting. The incidence examples use

$$\Lambda_t = \sum_{s=1}^L w_s Y_{t-s}, \quad (\text{S147})$$

$$\mu_t = R_t \Lambda_t, \quad (\text{S148})$$

$$Y_t \mid R_t, \phi \sim \text{NB}(\mu_t, \phi), \quad \text{Var}(Y_t \mid R_t, \phi) = \mu_t + \mu_t^2 / \phi, \quad (\text{S149})$$

where w_s is the discretized generation-interval mass function. A deterministic LS forecast recursively applies

$$\hat{Y}_{T+h} = R_{T+h}^{\text{future}} \sum_{s=1}^L w_s \tilde{Y}_{T+h-s}, \quad (\text{S150})$$

with $\tilde{Y}$ equal to observed incidence before the forecast origin and recursively forecast incidence afterward. Bayesian forecasts repeat the same recursion over posterior/future draws of R_{T+h} and then sample the negative-binomial observation layer.

SIR. The baseline SIR example uses

$$\dot{S} = -\beta(t)SI/N, \quad (\text{S151})$$

$$\dot{I} = \beta(t)SI/N - \gamma I, \quad (\text{S152})$$

$$\dot{R} = \gamma I, \quad (\text{S153})$$

with

$$R_e(t) = \frac{\beta(t)S(t)}{\gamma N}, \quad \beta(t) = \frac{\gamma N R_B(t)}{S(t)}. \quad (\text{S154})$$

For the transmission-reduction control example,

$$\beta(t) = \beta_0 \{1 - \varepsilon u(t)\}, \quad (\text{S155})$$

so the same three ODEs apply after substituting Eq. (S155).

SEIR. The latent-stage example uses

$$\dot{S} = -\beta(t)SI/N, \quad (\text{S156})$$

$$\dot{E} = \beta(t)SI/N - \sigma E, \quad (\text{S157})$$

$$\dot{I} = \sigma E - \gamma I, \quad (\text{S158})$$

$$\dot{R} = \gamma I, \quad (\text{S159})$$

with

$$R_e(t) = \frac{\beta(t)S(t)}{\gamma N}, \quad \beta(t) = \frac{\gamma N R_B(t)}{S(t)}. \quad (\text{S160})$$

The controlled SEIR example replaces the last expression by $\beta(t) = \beta_0\{1 - \varepsilon u(t)\}$ and retains the same progression equations.

SIRS with waning immunity. The SIRS example uses

$$\dot{S} = -\beta(t)SI/N + \omega R, \quad (\text{S161})$$

$$\dot{I} = \beta(t)SI/N - \gamma I, \quad (\text{S162})$$

$$\dot{R} = \gamma I - \omega R, \quad (\text{S163})$$

with $R_e(t) = \beta(t)S(t)/(\gamma N)$ and therefore $\beta(t) = \gamma N R_B(t)/S(t)$. The inverse-control example instead uses $\beta(t) = \beta_0\{1 - \varepsilon u(t)\}$.

Vector-borne transmission. The Ross–Macdonald-type example is

$$\dot{S}_h = -b_{vh}(t)S_h I_v/N_v, \quad (\text{S164})$$

$$\dot{I}_h = b_{vh}(t)S_h I_v/N_v - \gamma_h I_h, \quad (\text{S165})$$

$$\dot{R}_h = \gamma_h I_h, \quad (\text{S166})$$

$$\dot{S}_v = \mu_v N_v - b_{hv}(t)S_v I_h/N_h - \mu_v S_v, \quad (\text{S167})$$

$$\dot{I}_v = b_{hv}(t)S_v I_h/N_h - \mu_v I_v. \quad (\text{S168})$$

Its instantaneous next-generation spectral radius is

$$R_e(t) = \left[\frac{b_{vh}(t)b_{hv}(t)S_h(t)S_v(t)}{N_v N_h \gamma_h \mu_v} \right]^{1/2}. \quad (\text{S169})$$

For the shared-scaling example, $b_{vh} = c(t)b_{vh}^0$ and $b_{hv} = c(t)b_{hv}^0$ with $c(t)$ chosen so that Eq. (S169) equals $R_B(t)$. In the inverse-control example only vector-to-human transmission is controlled, $b_{vh}(t) = b_{vh,0}\{1 - u(t)\}$, while b_{hv} is held fixed.

Age-structured transmission. For groups $a = 1, \dots, G$,

$$\dot{S}_a = -S_a \lambda_a(t), \quad (\text{S170})$$

$$\dot{I}_a = S_a \lambda_a(t) - \gamma_a I_a, \quad (\text{S171})$$

$$\dot{R}_a = \gamma_a I_a, \quad (\text{S172})$$

$$\lambda_a(t) = \beta(t) \sum_{b=1}^G C_{ab}(t) \frac{I_b}{N_b}. \quad (\text{S173})$$

At unit β , the next-generation matrix has elements

$$A_{ab}(t) = \frac{S_a(t)C_{ab}(t)}{N_b\gamma_b}, \quad \beta(t) = \frac{R_B(t)}{\rho\{A(t)\}}. \quad (\text{S174})$$

The multilever control example decomposes

$$C(t) = C_H + \{1 - u_S(t)\}C_S + \{1 - u_W(t)\}C_W + \{1 - u_O(t)\}C_O, \quad (\text{S175})$$

where home contacts are left uncontrolled in the worked example and the school, work, and other/community controls are selected subject to the desired next-generation spectral radius.

Negative-binomial observation example. The overdispersed-count example does not change the SEIR ODEs. If μ_t is the deterministic SEIR daily infection flow, then

$$Y_t \mid \mu_t \sim \text{NB}(\mu_t, \phi), \quad \phi = 12, \quad \text{Var}(Y_t \mid \mu_t) = \mu_t + \frac{\mu_t^2}{\phi}. \quad (\text{S176})$$

This separates latent transmission dynamics from the observation process.

Vaccination with breakthrough infection. The vaccination example uses

$$\dot{S} = -\beta SI/N - \nu(t)S, \quad (\text{S177})$$

$$\dot{V} = \nu(t)S - (1 - \varepsilon)\beta VI/N, \quad (\text{S178})$$

$$\dot{I} = \beta I\{S + (1 - \varepsilon)V\}/N - \gamma I, \quad (\text{S179})$$

$$\dot{R} = \gamma I, \quad (\text{S180})$$

with

$$R_e(t) = \frac{\beta(t)}{\gamma N} \{S(t) + (1 - \varepsilon)V(t)\}, \quad \beta(t) = \frac{\gamma N R_B(t)}{S(t) + (1 - \varepsilon)V(t)}. \quad (\text{S181})$$

The time-varying vaccination rate $\nu(t)$ is the piecewise schedule reported in Example 7F and Table S9.

Hospitalization/isolation with residual infectiousness. The SEIHR example uses

$$\dot{S} = -\beta S(I + \kappa H)/N, \quad (\text{S182})$$

$$\dot{E} = \beta S(I + \kappa H)/N - \sigma E, \quad (\text{S183})$$

$$\dot{I} = \sigma E - (\gamma_I + \alpha)I, \quad (\text{S184})$$

$$\dot{H} = \alpha I - \gamma_H H, \quad (\text{S185})$$

$$\dot{R} = \gamma_I I + \gamma_H H. \quad (\text{S186})$$

The expected infectious-time equivalent is

$$D_I = \frac{1}{\gamma_I + \alpha} + \kappa \frac{\alpha}{(\gamma_I + \alpha)\gamma_H}, \quad (\text{S187})$$

and hence

$$R_e(t) = \frac{\beta(t)S(t)}{N} D_I, \quad \beta(t) = \frac{N R_B(t)}{S(t)D_I}. \quad (\text{S188})$$

If hospitalization/isolation is used as a control lever, its causal effect must enter through α , κ , another removal parameter, or a calibrated combination; hospital occupancy by itself is an outcome, not automatically a transmission intervention.

Inverse-control target equations. For the SIR target-prevalence example, a positive differentiable target $I^\star(t)$ implies

$$q^\star(t) = \dot{I}^\star(t) + \gamma I^\star(t), \quad S^\star(t) = S_0 - \int_0^t q^\star(s) ds, \quad (\text{S189})$$

and the exact transmission-reduction control is

$$u^\star(t) = \frac{1}{\varepsilon} \left[1 - \frac{Nq^\star(t)}{\beta_0 S^\star(t) I^\star(t)} \right]. \quad (\text{S190})$$

For SEIR,

$$E^\star(t) = \frac{\dot{I}^\star(t) + \gamma I^\star(t)}{\sigma}, \quad q_E^\star(t) = \dot{E}^\star(t) + \sigma E^\star(t), \quad (\text{S191})$$

with $S^\star(t) = S_0 - \int_0^t q_E^\star(s) ds$ and Eq. (S190) after replacing $q^\star$ by $q_E^\star$. SIRS additionally solves $\dot{R}^\star = \gamma I^\star - \omega R^\star$ and sets $S^\star = N - I^\star - R^\star$. These equations are the exact continuous-time targets that are subsequently compressed into small Bernstein policy polygons in Experiment 8.

Equation-to-example map. Experiments based only on Bernstein geometry use Eq. (S146); reconstruction and forecasting use Eqs. (S149)–(S150); mechanistic portability uses Eqs. (S153)–(S188); and inverse-control examples use the corresponding controlled ODE together with Eq. (S190) or its model-specific analogue.

S10 Inverse-control and Philippine-example settings

S11 Additional proof detail: degree elevation and minimality

A polynomial of degree at most n can be expressed in any Bernstein basis of degree $m \geq n$ by degree elevation. Therefore the existence of a representation with K control points never implies that K is minimal. The bounds in corollary 8 and theorem 9 work in the opposite direction: specified geometric features make some low degrees impossible. In exact arithmetic, the minimal Bernstein degree equals the ordinary polynomial degree after cancellation. In statistical applications, however, the target curve is not known exactly, and “minimal” should mean the smallest candidate K satisfying declared geometric and predictive adequacy criteria. This operational definition avoids confusing degree elevation with evidence for additional epidemic complexity.

S12 Forecast scoring

For a central $(1 - \alpha)$ prediction interval $[\ell, u]$ and observation y , the interval score is

$$\text{IS}_\alpha(y; \ell, u) = (u - \ell) + \frac{2}{\alpha}(\ell - y)\mathbf{1}(y < \ell) + \frac{2}{\alpha}(y - u)\mathbf{1}(y > u). \quad (\text{S192})$$

The simulation code combines absolute median error with weighted interval scores using $\alpha = 0.1, 0.2$, and 0.5 , following the structure described by Bracher et al. [22]. The same quantile levels and weighting should be fixed before empirical rolling-origin evaluation.

Table S8: Inputs used in the reproducible simulations.

Component	Setting
Random seed	20260908
Generation interval	Discretized gamma, mean 5 d, SD 2 d, maximum lag 20 d, renormalized
Observation model	Negative binomial with mean $R_t \Lambda_t$
Historical variant experiment	120 days, seven Bernstein control points
Inference experiment	140 historical + 60 held-out days; $K = 7$, $K_F = 6$, $\phi = 35$
LS inference	$\lambda = 3$; bounds $[0.05, 4]$; 300 parametric bootstrap refits
Bayesian inference	$\sigma_0 = 0.55$, $\tau_2 = 0.17$, $\log \phi \sim N(\log 30, 0.70^2)$; 4 chains $\times$ 8000 iterations; burn-in 3000; thin 2
Bayesian future uncertainty	3000 future polygons; future innovation SD 0.025; $\log a \sim N(\log 1.50, 0.10^2)$; $s \sim N(0.105, 0.018^2)$ truncated below 0.02; 1200 incidence predictive draws
Future forecast segment	60 days, six Bernstein control points
Historical true controls	(1.25, 1.31, 1.16, 0.93, 0.79, 0.84, 0.82)
Variant start share	$p_0 = 0.02$
Variant logistic growth	$s = 0.12$ per day
Variant multipliers	1.0, 1.2, 1.4, 1.6, 1.8
Renewal dispersion	$\phi = 50$ in variant experiment
Monte Carlo futures	100 per variant multiplier
Predictive samples	250 per model per future
Forecast horizons	14, 28, 42, 60 days
Complexity verification grid	121 time points; candidates $K = 2, \dots, 8$
Complexity verification noise	Gaussian SD 0.012 on R_t scale; 120 replicates
Robustness noise SDs	0.006, 0.012, 0.025, 0.05, 0.08; 200 replicates per scenario/SD

Table S9: Complete settings for the seven additional model-class examples. Rates are per day unless stated otherwise.

Example	Declared settings
Common target	210 days; $K = 8$; $\boldsymbol{\rho} = (1.16, 1.10, 0.91, 0.82, 1.09, 1.17, 0.96, 0.86)$; daily output; ODE maximum step 0.25 d.
SEIR	$N = 10^6$; $(S_0, E_0, I_0, R_0) = (999740, 160, 100, 0)$; $\sigma = 1/3$; $\gamma = 1/6$.
SIRS	$N = 10^6$; $(S_0, I_0, R_0) = (799800, 200, 200000)$; $\gamma = 1/6$; $\omega = 1/120$.
Vector-borne	$N_h = 10^6$; $N_v = 2 \times 10^6$; $(I_{h0}, I_{v0}) = (180, 1200)$; $\gamma_h = 1/6$; $\mu_v = 1/14$; $b_{vh}^0 = 0.35$; $b_{hv}^0 = 0.25$; vectors have births at rate $\mu_v N_v$.
Age-structured	$N = (250000, 550000, 200000)$; $I_0 = (80, 120, 30)$; $\gamma_a = 1/6$ for all groups; contact matrix given in Eq. (S124).
Negative binomial	Mean daily incidence equals the deterministic SEIR exposure flow; size/dispersion $\phi = 12$; variance $\mu + \mu^2/\phi$; random seed 20261608.
Vaccination	$N = 10^6$; $(S_0, V_0, I_0, R_0) = (999780, 0, 220, 0)$; $\gamma = 1/6$; $\varepsilon = 0.75$; $\nu = 0$ before day 45, 0.004 on days 45–149, and 0.001 thereafter.
Hospital/isolation	$N = 10^6$; $(S_0, E_0, I_0, H_0, R_0) = (999740, 140, 120, 0, 0)$; $\sigma = 1/3$; $\gamma_I = 1/8$; $\alpha = 1/5$; $\gamma_H = 1/8$; $\kappa = 0.08$.

Table S10: Fully declared settings for the inverse-control examples. All values are synthetic and illustrative.

Example	Inputs and policy mechanism
Common target	$T = 120$ d, $N = 10^6$, seven log-prevalence control ordinates corresponding to (200, 350, 1000, 2500, 1800, 800, 300)
SIR	$\gamma = 1/7$, $\beta_0 = 0.34$, multiplicative transmission reduction $\beta = \beta_0(1 - u)$
SIRS	$\gamma = 1/7$, $\omega = 1/120$, $R(0) = 100,000$, $\beta_0 = 0.40$, $\beta = \beta_0(1 - u)$
SEIR	$\sigma = 1/3$, $\gamma = 1/7$, $\beta_0 = 0.42$, $\beta = \beta_0(1 - u)$
Vector-borne	$N_h = 10^6$, $N_v = 2 \times 10^6$, $\gamma_h = 1/7$, $\mu_v = 1/14$, $b_{hv} = 0.32$, $b_{vh,0} = 0.12$, $I_v(0) = 1000$; u reduces vector-to-human transmission
Age-structured	populations (250000, 550000, 200000), infectious period 6 d, initial infections (80, 120, 30); explicit C_H, C_S, C_W, C_O in text; initial uncontrolled $R_t = 2$; relative costs (1.3, 1.8, 0.7)
Cost scan	reference SIR policy scaled from 0 to 1.2; normalized burden weights $\lambda = 0.05, 0.15, 0.40$

Table S11: Philippine COVID-19 illustration settings.

Component	Setting
Data	1,178 daily reporting-date case counts, 30 Jan 2020–21 Apr 2023; total 4,087,964
Preprocessing	7-day trailing mean; no deconvolution for reporting delay or ascertainment
Generation interval	Gamma-discretized, mean 5 d, SD 2 d, truncated at 20 d and renormalized
Historical reconstruction	$K = 36$, positive penalized LS, second-difference penalty 18, square-root incidence weights
Historical uncertainty	160 moving-block residual bootstrap refits, block length 14 d
Pseudo-forecast	final 60 d held out; original joined future segment plus peak-aware cases-only continuation
Forecast uncertainty	240 moving-block refits for the original joined forecast; peak-aware comparison is a point-forecast evaluation
Peak-aware selection	three pre-Omicron origins (24 Mar, 19 Jun, 14 Sep 2021); grid $w \in \{4, 5, 7\}$, hold $\in \{4, 6, 8, 10\}$ d, $\tau \in \{3, 4, 5, 6, 8\}$ d, $R_{\text{post}} \in \{0.6, 0.7, 0.8, 0.9\}$; selected (4, 10, 8, 0.9)
Out-of-selection backtest	six later origins from Dec 2021–Feb 2023; horizons 14, 28, 42, and 60 d; compare peak-aware adaptive, joined Bernstein, and constant- R_t
Historical-origin stress test 1	train through 31 Dec 2020; 150-d forecast of first large 2021 wave; additional early-wave origin on 20 Mar 2021
Historical-origin stress test 2	train through 31 Dec 2021; compare joined, constant- R_t , peak-aware cases-only, and variant-informed peak-aware forecasts; broader 16-scenario event envelope retained
Sensitivity	Generation-interval mean 4, 5, 6 d (SD 2 d) crossed with $K = 24, 36, 48$ for descriptive reconstruction sensitivity
Comparator	constant R_t from median of final 14 training days
Illustrative meeting target	$R_t \leq 0.90$; required additional effective suppression relative to the model’s no-new-intervention forecast; no specific policy-effect or monetary-cost claim

Table S12: Forecast-priority, Bayesian event, and Philippine hypothetical-control settings used in the reproducible examples. Priority weights define Eq. (S45); they are decision preferences, not fitted biological parameters.

Component	Declared setting
Balanced forecast loss	$(w_c, w_t, w_h, w_d) = (1, 1, 2, 0.75)$ for whole-curve, peak-time, peak-height, and onset-delay components, respectively.
Peak-priority loss	$(w_c, w_t, w_h, w_d) = (0.25, 5, 8, 2)$.
Timing-priority loss	$(w_c, w_t, w_h, w_d) = (0.25, 10, 3, 4)$.
Deployable history-only profile	Selected by pre-Omicron rolling origins; $(\text{scale}, \text{hold}, \tau, R_{\text{post}}, \delta) = (0.95, 10 \text{ d}, 10 \text{ d}, 0.90, -2 \text{ d})$ in the reported run.
Omicron event-balanced profile	$(1.00, 8 \text{ d}, 5 \text{ d}, 0.80, 0 \text{ d})$.
Omicron event peak-priority profile	$(0.99, 9 \text{ d}, 5 \text{ d}, 0.70, 0 \text{ d})$ for the forward event-conditioned scenario. A separate retrospective-oracle diagnostic is excluded from forecast-skill claims.
Bayesian event uncertainty	Recent log-growth slope prior SD 0.35 d^{-1} ; event scale SD 0.05; hold SD 1.5 d; τ SD 1 d; post-event R_t SD 0.08; phase-shift SD 1 d; illustrative negative-binomial predictive dispersion $\phi = 35$.
Philippine Bayesian reconstruction	$K = 20$ positive control points on the log scale; coefficient-level prior SD 0.70; second-difference prior SD 0.25; $\log \phi \sim N(\log 20, 0.80^2)$; MAP with Laplace approximation and 900 posterior draws.
Hypothetical Philippine control	Target $R_t \leq 0.90$; standardized contact-NPI efficacy 0.80; baseline removal $\gamma = 1/7 \text{ d}^{-1}$; maximum additional isolation/removal 0.25 d^{-1} ; vaccination-protection efficacy index 0.80; relative quadratic burden weights $(1.00, 0.65, 0.85)$; SLSQP allocation.

Table S13: Philippine reconstruction sensitivity to the assumed generation-interval mean and Bernstein complexity. Log-NRMSE is calculated against the 7-day reporting-date case series. The variation in the fitted R_t range illustrates that reproduction-number levels depend on epidemiological assumptions even when reconstruction error changes only modestly.

GI mean (d)	K	log-NRMSE	min R_t	max R_t
4	24	0.02362	0.953	1.390
4	36	0.02363	0.933	1.379
4	48	0.02362	0.915	1.318
5	24	0.02865	0.944	1.531
5	36	0.02865	0.919	1.526
5	48	0.02863	0.898	1.455
6	24	0.03355	0.937	1.651
6	36	0.03353	0.905	1.644
6	48	0.03350	0.881	1.562

S13 Shiny application user manual

The companion Control-Polygon Epidemiology Studio is an interactive R Shiny implementation of the paper’s reconstruction–forecast–control workflow. The app is designed both to reproduce paper examples and to support user-supplied incidence data. The interactive implementation uses the same Bernstein representation but prioritizes responsiveness; its Bayesian mode uses a fast Laplace/posterior-draw approximation and should be confirmed with full MCMC for high-stakes inference when posterior non-Gaussianity is material.

S13.1 Installation

Open R in the `shiny_app/` folder and run

```
source("install_packages.R")
shiny::runApp()
```

Required packages are `shiny`, `bslib`, `plotly`, `DT`, `ggplot2`, `deSolve`, `MASS`, and `htmltools`.

Note: The online Shiny app does not need to install the packages.

S13.2 Recommended workflow

1. **Activate data.** Use a built-in example or upload a CSV. A downloadable two-column template (date, cases) is provided. Select a date/time column if available and a nonnegative incidence/case-count column. If no date column is supplied, the app creates a sequential daily index. Every tab displays a dataset/source banner stating whether it is using the active uploaded/preset dataset, a stored forecast derived from that dataset, a mechanistic paper preset, or the fixed Philippine paper dataset.
2. **Reconstruct.** Choose K , generation-interval assumptions, smoothing/penalty settings, and LS or Bayesian mode. Inspect the historical R_t polygon, uncertainty, inferred threshold crossings, residual diagnostics, and the implied minimum structural complexity.
3. **Forecast.** Choose the forecast horizon and objective profile. Balanced, peak-priority, timing-priority, and custom weights correspond to the loss in Eq. (S45). Event-conditioned scenarios can be activated only when such information would genuinely be available at the forecast origin. Inspect forecast epicurves against historical context, R_t trajectories, uncertainty intervals, peak date/height, and threshold risk.
4. **Control.** Specify a target R_t cap or edit target Bernstein-point values in the table. The app computes point and Bayesian inverse-control requirements, tests feasibility, and can allocate generic effective suppression across illustrative contact-NPI, testing/isolation, and vaccination/protection levers. To keep the tab responsive, Bayesian control uncertainty reuses a bounded deterministic subset of the forecast posterior draws rather than generating a second large posterior array.
5. **Complexity, cost and sensitivity.** Explore the minimum- K theorems and normalized health-policy cost trade-offs. Replace normalized costs with calibrated monetary or utility functions for applied work. Use Sensitivity Analysis to vary one control input at a time and identify which outputs are most responsive.

6. **Mechanistic laboratory.** Reproduce the SIR, SEIR, SIRS, vector-borne, age-structured, vaccination, SEIHR, and negative-binomial examples. Edit scalar and vector parameters, contact matrices, initial conditions, and target control points.
7. **Export.** Download plots and tables from each workflow before changing the active data or assumptions. Each analytical figure has an explicit PNG button and a figure-data CSV button; displayed tables have CSV and PNG export buttons. Record the exact settings used in a research analysis.

S13.3 Editing Bernstein control points

Control-point editing is table-based. In every polygon panel, edit the **Value** column in the table below the Plotly graph; the curve updates immediately. Plotly is for zoom, pan, hover, and visual inspection. Table-based editing leaves an exact auditable record that can be exported as CSV. This is suitable for meeting-room use: a modeller or policymaker changes a small number of interpretable ordinates and immediately sees the implied R_t , incidence, intervention, and cost outputs.

S13.4 Tabs and interpretation

Tab	Main inputs/process	Interpretation
Home	Workflow, assumptions, quick start	Read before analysis; distinguishes inference, scenarios, and control.
Data & Reconstruction	CSV data, K , generation interval, LS/Bayesian settings	Historical R_t , control points, uncertainty, threshold geometry, residuals.
Forecasting	Horizon, future K , objective weights, event assumptions	Future epicurve/ R_t , prediction intervals, peak timing/height, $P(R_t > 1)$.
Control & Policy	Target curve/cap, intervention bounds, policy costs	Required suppression, feasibility, lever allocation, posterior control/cost uncertainty.
Cost-Benefit	Cost weights, control options	Normalized or calibrated trade-off frontier between epidemic and policy burden.
Model Laboratory	Model class, editable parameters, contact matrix, target polygon	Reproduces mechanistic examples and reveals how the same R_t target maps to different mechanisms.
Philippine COVID-19	Historical-origin presets and Omicron/event scenarios	Reproduces the empirical examples while separating cases-only and event-conditioned information.
Complexity & K	Requested crossings/extrema and candidate K	Shows the structural floor $K \geq C+1$ and $K \geq E+2$; data may require larger K .
Methods & Downloads	Reporting checklist and export tools	Archive assumptions, outputs, and reproducibility metadata.

S13.5 Plot controls

Users can select color palettes and set explicit axis ranges. Mechanistic plots support raw values, per-100,000 scaling, and $\log_{10}(1+x)$ display to prevent very large epidemic trajectories from making smaller trajectories unreadable. Each analytical figure has a dedicated PNG download button plus a CSV button for the plotted data; Plotly's own toolbar is available for interactive inspection. Each table has CSV export buttons.

S13.6 Uploading data for independent research

The active dataset is shared across reconstruction and forecasting, while Control, Cost-Benefit and Sensitivity Analysis explicitly report the stored forecast/control source they are using. A downloadable CSV template is available before upload. Dataset/source banners appear at the top of every tab so the user can distinguish uploaded data, built-in presets, stored derived results, and synthetic/model-laboratory examples. Users should upload nonnegative incidence at a consistent time step and document reporting-date versus onset-date interpretation. Generation-interval parameters, reporting delays, imported cases, ascertainment, and under-reporting are not automatically identifiable from a case column; applied users should adjust or extend the core functions when these features matter. A fitted R_t curve is therefore conditional on the stated observation model.

S13.7 Cautions for LS and Bayesian modes

LS provides fast point estimates and bootstrap-style uncertainty when requested. The interactive Bayesian mode is intended for exploration and approximates the posterior around a fitted mode. Separate codes were used for the more explicit Bayesian workflow shown in the manuscript figures. Event-conditioned forecasts are conditional on declared event information. In the app, an event-informed profile means a future R_t scenario modified by external evidence available at the forecast origin (for example, documented variant emergence, a scheduled intervention, school reopening, or another known event), not by post-origin case counts. The built-in Event peak-priority template multiplies the terminal growth-derived anchor by 1.35, holds it for approximately 8 days, decays with a 5-day time constant toward $R = 0.70$, and applies a 1-day phase lead; these values are editable scenario assumptions rather than universally estimated constants. The app must not use post-origin outcomes as hidden forecast inputs. Effective suppression is a mathematical control quantity until a calibrated intervention-effect function maps it to vaccination, isolation, NPI, or other real policy units.

S13.8 Project structure

- `app.R`: Shiny user interface and server logic.
- `R/core.R`: Bernstein, renewal, LS/Bayesian inference, forecasting, and inverse-control utilities.
- `R/models.R`: mechanistic paper examples.
- `R/ui_helpers.R`: reusable user-interface and plotting helpers.
- `data/philippine_covid.csv`: frozen paper dataset.
- `data/example_incidence.csv`: worked example dataset.

- `data/upload_template.csv`: minimal date/cases upload template exposed by the app download button.
- `www/styles.css`: visual design.
- `install_packages.R`: package installer.

S14 Reproducibility checklist for applied users

For an external analysis, report at minimum: data source and retrieval date; date definition and time step; preprocessing and smoothing; generation-interval distribution and sensitivity analysis; historical and future K ; LS penalties/constraints or Bayesian priors; forecast-origin information set; objective weights for peak/timing/curve/delay; event assumptions; model parameters and initial conditions; intervention-effect mappings; policy bounds; cost functions; uncertainty method; validation origins and horizons; scoring rules; and software/package versions. Distinguish cases-only forecasts, event-conditioned forecasts, retrospective oracle diagnostics, and policymaker target scenarios in every table and figure.

References

- [1] Rida T. Farouki. The bernstein polynomial basis: A centennial retrospective. *Computer Aided Geometric Design*, 29(6):379–419, 2012. doi: 10.1016/j.cagd.2012.03.001.
- [2] Gerald Farin. *Curves and Surfaces for CAGD: A Practical Guide*. Morgan Kaufmann, San Francisco, 5 edition, 2002.
- [3] Anne Cori, Neil M. Ferguson, Christophe Fraser, and Simon Cauchemez. A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology*, 178(9):1505–1512, 2013. doi: 10.1093/aje/kwt133.
- [4] Katelyn M. Gostic, Lauren McGough, Edward B. Baskerville, Sam Abbott, Keya Joshi, Christine Tedijanto, Rebecca Kahn, Rene Niehus, James A. Hay, Pablo M. De Salazar, Joel Hellewell, Sophie Meakin, James D. Munday, Nikos I. Bosse, Katharine Sherratt, Robin N. Thompson, Laura F. White, Jana S. Huisman, Jérémie Scire, Sebastian Bonhoeffer, Tanja Stadler, Jacco Wallinga, Sebastian Funk, Marc Lipsitch, and Sarah Cobey. Practical considerations for measuring the effective reproductive number, r_t . *PLoS Computational Biology*, 16(12):e1008409, 2020. doi: 10.1371/journal.pcbi.1008409.
- [5] Oswaldo Gressani, Jacco Wallinga, Christian L. Althaus, Niel Hens, and Christel Faes. Epilps: A fast and flexible bayesian tool for estimation of the time-varying reproduction number. *PLoS Computational Biology*, 18(10):e1010618, 2022. doi: 10.1371/journal.pcbi.1010618.
- [6] Tom Britton and Gianpaolo Scalia Tomba. Estimation in emerging epidemics: Biases and remedies. *Journal of the Royal Society Interface*, 16(150):20180670, 2019. doi: 10.1098/rsif.2018.0670.
- [7] Jiangdian Wang and Sujit K. Ghosh. Shape restricted nonparametric regression with bernstein polynomials. *Computational Statistics & Data Analysis*, 56(9):2729–2741, 2012. doi: 10.1016/j.csda.2012.02.018.

- [8] Rahul Ghosal, Sujit Ghosh, Jacek Urbanek, Jennifer A. Schrack, and Vadim Zipunnikov. Shape-constrained estimation in functional regression with bernstein polynomials. *Computational Statistics & Data Analysis*, 178:107614, 2023. doi: 10.1016/j.csda.2022.107614.
- [9] Jianhua Ding and Zhongzhan Zhang. Bayesian regression on non-parametric mixed-effect models with shape-restricted bernstein polynomials. *Journal of Applied Statistics*, 43(14): 2524–2537, 2016. doi: 10.1080/02664763.2016.1142940.
- [10] Rahul Ghosal and Sujit K. Ghosh. Bayesian inference for generalized linear model with linear inequality constraints. *Computational Statistics & Data Analysis*, 166:107335, 2022. doi: 10.1016/j.csda.2021.107335.
- [11] Odo Diekmann, J. A. P. Heesterbeek, and Michael G. Roberts. The construction of next-generation matrices for compartmental epidemic models. *Journal of the Royal Society Interface*, 7(47):873–885, 2010. doi: 10.1098/rsif.2009.0386.
- [12] David L. Smith, Katherine E. Battle, Simon I. Hay, Christopher M. Barker, Thomas W. Scott, and F. Ellis McKenzie. Ross, macdonald, and a theory for the dynamics and control of mosquito-transmitted pathogens. *PLoS Pathogens*, 8(4):e1002588, 2012. doi: 10.1371/journal.ppat.1002588.
- [13] Samuel Karlin. *Total Positivity, Volume I*. Stanford University Press, Stanford, California, 1968.
- [14] G. G. Lorentz. *Bernstein Polynomials*. Chelsea Publishing Company, New York, 2 edition, 1986.
- [15] Jacco Wallinga and Marc Lipsitch. How generation intervals shape the relationship between growth rates and reproductive numbers. *Proceedings of the Royal Society B*, 274(1609):599–604, 2007. doi: 10.1098/rspb.2006.3754.
- [16] Sang Woo Park, David Champredon, Joshua S. Weitz, and Jonathan Dushoff. A practical generation-interval-based approach to inferring the strength of epidemics from their speed. *Epidemics*, 27:12–18, 2019. doi: 10.1016/j.epidem.2018.12.002.
- [17] Kimberly A. Dautel, Ephraim Agyingi, and Pras Pathmanathan. Validation framework for epidemiological models with application to covid-19 models. *PLOS Computational Biology*, 19(3):e1010968, 2023. doi: 10.1371/journal.pcbi.1010968.
- [18] Thomas McAndrew, Graham C. Gibson, David Braun, Abhishek Srivastava, and Kate Brown. Chimeric forecasting: An experiment to leverage human judgment to improve forecasts of infectious disease using simulated surveillance data. *Epidemics*, 47:100756, 2024. doi: 10.1016/j.epidem.2024.100756.
- [19] Nicholas G. Davies, Sam Abbott, Rosanna C. Barnard, Christopher I. Jarvis, Adam J. Kucharski, James D. Munday, Carl A. B. Pearson, Timothy W. Russell, Damien C. Tully, Alex D. Washburne, et al. Estimated transmissibility and impact of sars-cov-2 lineage b.1.1.7 in england. *Science*, 372(6538):eabg3055, 2021. doi: 10.1126/science.abg3055.
- [20] Erik Volz, Swapnil Mishra, Meera Chand, Jeffrey C. Barrett, Robert Johnson, Lily Geidelberg, Wes R. Hinsley, Daniel J. Laydon, Gavin Dabrera, Àine O’Toole, et al. Assessing transmissibility of sars-cov-2 lineage b.1.1.7 in england. *Nature*, 593:266–269, 2021. doi: 10.1038/s41586-021-03470-x.

- [21] Fritz Obermeyer, Martin Jankowiak, Nikolaos Barkas, Stephen F. Schaffner, Jesse D. Pyle, Leonid Yurkovetskiy, Matteo Bosso, Daniel J. Park, Mehrtash Babadi, Bronwyn L. MacInnis, Jeremy Luban, Pardis C. Sabeti, and Jacob E. Lemieux. Analysis of 6.4 million sars-cov-2 genomes identifies mutations associated with fitness. *Science*, 376(6599):1327–1332, 2022. doi: 10.1126/science.abm1208.
- [22] Johannes Bracher, Evan L. Ray, Tilmann Gneiting, and Nicholas G. Reich. Evaluating epidemic forecasts in an interval format. *PLoS Computational Biology*, 17(2):e1008618, 2021. doi: 10.1371/journal.pcbi.1008618.
- [23] Lorenzo A. Ricciardi and Massimiliano Vasile. Direct transcription of optimal control problems with finite elements on bernstein basis. *Journal of Guidance, Control, and Dynamics*, 42(2): 229–243, 2019. doi: 10.2514/1.G003753.
- [24] Calvin Kielas-Jensen, Venanzio Cichella, Thomas Berry, Isaac Kaminer, Claire Walton, and Antonio Pascoal. Bernstein polynomial-based method for solving optimal trajectory generation problems. *Sensors*, 22(5):1869, 2022. doi: 10.3390/s22051869.
- [25] Rocio Balderrama, Javier Peressutti, Juan Pablo Pinasco, Federico Vazquez, and Constanza Sanchez de la Vega. Optimal control for a sir epidemic model with limited quarantine. *Scientific Reports*, 12:12583, 2022. doi: 10.1038/s41598-022-16619-z.
- [26] Md Abdul Kuddus, Anip Kumar Paul, and Thitiya Theparod. Cost-effectiveness analysis of covid-19 intervention policies using a mathematical model: An optimal control approach. *Scientific Reports*, 14:494, 2024. doi: 10.1038/s41598-023-50799-6.
- [27] Raffaele Carli, Graziana Cavone, Nicola Epicoco, Paolo Scarabaggio, and Mariagrazia Dotoli. Model predictive control to mitigate the covid-19 outbreak in a multi-region scenario. *Annual Reviews in Control*, 50:373–393, 2020. doi: 10.1016/j.arcontrol.2020.09.005.
- [28] Marcelo M. Morato, Saulo B. Bastos, Daniel O. Cajueiro, and Julio E. Normey-Rico. An optimal predictive control strategy for covid-19 (sars-cov-2) social distancing policies in brazil. *Annual Reviews in Control*, 50:417–431, 2020. doi: 10.1016/j.arcontrol.2020.07.001.
- [29] Herbert W. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000. doi: 10.1137/S0036144500371907.
- [30] Kiesha Prem, Alex R. Cook, and Mark Jit. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLoS Computational Biology*, 13(9):e1005697, 2017. doi: 10.1371/journal.pcbi.1005697.
- [31] Matthew J. Y. Shin, Juliette Paireau, and Simon Cauchemez. Leveraging regularity in covid-19 growth rate dynamics for epidemic wave forecasting. *Epidemics*, 55:100908, 2026. doi: 10.1016/j.epidem.2026.100908.
- [32] ASEAN BioDiaspora Virtual Center. Covid-19 and mpox situational report, 24 april 2023, issue no. 475, 2023. Reports 4,087,964 cumulative confirmed COVID-19 cases in the Philippines as of 21 April 2023.
- [33] Norvin P. Bansilan and Jomar F. Rabajante. A descriptive analytics of the covid-19 pandemic in a middle-income country with forward-looking insights. *Healthcare Analytics*, 5:100320, 2024. doi: 10.1016/j.health.2024.100320.

- [34] Yao-Tsun Li, Francisco Gerardo M. Polotan, Gerald Ivan S. Sotelo, Anne Pauline A. Alpino, Ardiane Ysabelle M. Dolor, Ma. Angelica A. Tujan, Ma. Ricci R. Gomez, Othoniel Jan T. Onza, Angela Kae T. Chang, Criselda T. Bautista, June C. Carandang, Dodge R. Lim, Lei Lanna M. Dancel, Mayan Uy-Lumandas, Timothy John R. Dizon, Katie Hampson, Simon Daldry, Joseph Hughes, and Kirstyn Brunker. Lineage ba.2 dominated the omicron sars-cov-2 epidemic wave in the philippines. *Virus Evolution*, 8(2):veac078, 2022. doi: 10.1093/ve/veac078.