Uncertainty Discourse and Applications

www.uda.reapress.com

Uncert. Disc. Appl. Vol. 3, No. 4 (2026) 307–319.

Paper Type: Original Article

Scenario Quantile Predictive Control of Bioreactor Cultivation under Kinetic Parameter Uncertainty

Tarek Abdelrahman[1], Wolfgang H. Müller[2], Arafa Abdelzaher Nasef[3,*]

[1] Department of Electrical Power and Machines Engineering, Faculty of Engineering, Ain Shams University, Cairo, Egypt.
[2] Department of Computer Science, Technical University of Munich, Munich, Germany.
[3] Department of Engineering, Kafrelsheikh University, Egypt; nasefa50@yahoo.com.

Citation:

Received: 20 December 2025 *Revised: 05 April 2026* *Accepted: 19 June 2026*	Abdelrahman, T., Müller, W. H., & Nasef, A. A. (2026). Scenario quantile predictive control of bioreactor cultivation under kinetic parameter uncertainty. *Uncertainty Discourse and Applications, 3*(4), 307-319.

Abstract

Bioreactor feeding requires a balance between product formation, substrate limits, and uncertain kinetics. This study examines an empirical scenario-quantile model predictive controller that constrains an order statistic of predicted substrate concentrations and adds a quadratic exceedance penalty. The reported simulation campaign uses a constant-volume, three-state continuous-culture surrogate with Monod growth, product inhibition, Luedeking–Piret product formation, and six lognormally dispersed kinetic parameters. Thirty parameter draws crossed with five process-noise seeds give 150 paired runs per controller. The proposed controller has a reported final titer of 5.049 ± 1.593 g/L and a pooled step violation frequency of 4.9%, compared with 2.788 ± 1.130 g/L for Open-Loop (OL) nominal optimization, 4.798 ± 1.516 g/L for Proportional–Integral (PI) control, 4.913 ± 1.506 g/L for fixed-margin Model Predictive Control (MPC), 5.044 ± 1.545 g/L for nominal MPC, and 5.059 ± 1.576 g/L for hard-scenario MPC. The hard-scenario controller has fewer observed violations, while an 80/90/95% order-statistic ablation changes violation frequencies from 7.9% to 2.9% with small mean-titer differences. These observations support a tunable empirical risk mechanism at the tested scale. They establish neither a finite-sample chance-constraint guarantee nor a joint batch safety guarantee. Inference, changing-volume fed-batch dynamics, and structural model error remain outside the executed comparison; the nested simulation design also limits interpretation of the reported run-level significance tests.

Keywords: Empirical quantile control, Model predictive control, Kinetic parameter uncertainty, Continuous culture, Scenario optimization, Bioprocess simulation.

1 | Introduction

Feeding policy links bioreactor kinetics to operational objectives. Monod kinetics describe substrate-limited growth [1], while the Andrews model incorporates substrate inhibition [2]. These mechanisms motivate control of substrate concentration, but a substrate ceiling is a biological inhibition limit only when the

Corresponding Author: nasefa50@yahoo.com
https://doi.org/10.48313/uda.v3i4.117
Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0).

organism and kinetic model support that interpretation. The present study examines uncertain-parameter predictive control in a simulation whose substrate ceiling is an operational guard rather than an identified inhibition threshold.

Dynamic optimization provides a computational basis for selecting cultivation trajectories, including path-constrained nonlinear bioprocess models [3]. A nominal trajectory uses one kinetic parameter vector; changing that vector can alter both productivity and constraint satisfaction. The resulting sensitivity motivates feedback and explicit uncertainty propagation, without implying that every parameter perturbation causes a violation.

Model Predictive Control (MPC) repeatedly solves a finite-horizon problem using updated state information. Established stability and feasibility results require stated model, terminal, and disturbance assumptions [4]. Stochastic MPC instead introduces probability models and violation budgets [5]. The present implementation uses scenario predictions and local nonlinear optimization; its constraint behavior is assessed empirically rather than inferred from a general MPC stability theorem.

Bioprocess estimation can supply kinetic information for a controller [6]. Here the controller receives samples from a specified kinetic distribution, and a paired simulation design compares the influence of uncertainty representation on titer and substrate excursions. This isolates the control layer: Bayesian parameter learning and uncertainty calibration from sparse assays are proposed extensions rather than tested components.

The contribution is an empirical order-statistic controller with an explicit quantile level and exceedance penalty, evaluated against Open-Loop (OL) optimization, Proportional–Integral (PI) control, nominal MPC, hard-scenario MPC, and fixed-margin MPC under common simulated parameter and noise realizations.

The comparison reports final titer, productivity, pooled step violations, runs with any violation, lower-tail titer, and solve time. A quantile-level ablation assesses how the chosen order statistic changes observed constraint behavior.

The implementation uses warm-started Sequential Least Squares Quadratic Programming (SLSQP) single shooting. Its measured solve time describes the stated simulator and hardware campaign, while the distinction between empirical violation frequencies and probabilistic guarantees governs interpretation of the results.

Section 2 places the construction within bioprocess optimization and stochastic control. Sections 3 and 4 define the model, objectives, and implemented scenario mechanism. Section 5 presents the reported paired simulation results, and Sections 6 and 7 discuss their scope and the next validation steps.

2 | Related Work

2.1 | Optimal Control of Cultivation Processes

Bioprocess trajectory optimization combines kinetic models, product-formation laws, and path constraints. Numerical methods for these problems are reviewed by Banga and co-workers [3]. The growth-associated and non-growth-associated terms of the Luedeking–Piret law provide a conventional product model [7]. Their use here specifies a simulator architecture; the nominal coefficients are study settings and are not estimates for a named organism.

2.2 | Model Predictive, Robust, and Stochastic Control

Constrained MPC supplies a receding-horizon optimization framework [4], while stochastic MPC makes the probability model and constraint budget explicit [5]. The scenario approach yields probabilistic results for specified classes of sampled optimization problems, with assumptions and sample requirements tied to those results [8]. Those results do not automatically certify a nonconvex order-statistic program solved locally and reused in closed loop. The construction evaluated here therefore distinguishes its probability target from its empirical planning surrogate.

At the tested scale, tractability depends on horizon length, number of scenarios, and nonlinear-solver behavior. Warm starts reduce repeated optimization effort, but a comparison of one nominal model with ten scenarios is not a controlled scaling experiment.

2.3 | Parameter Uncertainty and Estimation in Bioprocesses

Kinetic estimation and adaptive bioreactor control have an established literature [6]. Bayesian inference provides a distribution over uncertain parameters [9], and sequential Monte Carlo supplies one possible update mechanism [10]. A learned distribution requires an explicit likelihood, assay design, identifiability assessment, and predictive checks. None of these inference steps is exercised in the reported campaign, which samples a known simulation law.

2.4 | Synthesis and Research Gap

The present study connects uncertain kinetics to predictive substrate control through an empirical quantile constraint. Its contribution is the paired comparison and quantile ablation under a specified simulator, rather than a claim that posterior-aware bioprocess control has no precedent. *Table 1* distinguishes the assumptions of established control families from the empirical status of this implementation.

Table 1. Control-design landscape for cultivation under kinetic parameter uncertainty.

Design Family	Uncertainty Representation	Constraint Handling	Guarantee Type	Representative Sources
OL optimal control	None (nominal parameters)	Path constraints in the program	Local or nominal optimum under solver assumptions	[3], [11]
Nominal MPC	Point estimate	Hard constraints	Conditional on terminal and model assumptions	[4], [12]
Min-max robust MPC	Set-based worst case	Tightened constraints	Conditional on uncertainty set and feasibility assumptions	[13]
Tube MPC	Set propagation, affine feedback	Tube-radius tightening	Conditional on tube construction and disturbance assumptions	[14]
Stochastic/chance-constrained MPC	Probability distribution	Chance constraints	Conditional on sampling and formulation assumptions	[5], [8]
Posterior-informed chance-constrained MPC, proposed	Specified kinetic distribution in the experiment	Scenario back-off on state limits	Empirical violation diagnostics only	This design

3 | Problem Formulation and Research Questions

3.1 | Research Questions and Hypotheses

The study asks how scenario-based kinetic uncertainty affects the productivity–violation trade-off and whether the resulting optimization is inexpensive at the tested horizon. The questions below distinguish a design target from an observed result.

RQ1: Can an empirical scenario quantile provide a tunable substrate-risk mechanism in the reported simulator, and how does its observed violation frequency relate to the nominal probability target?

RQ2: How do titer and substrate excursions differ among the six controllers under paired parameter draws and process-noise seeds?

RQ3: What solve-time budget is observed for the stated horizon and scenario count?

The titer hypothesis compares the quantile controller with the fixed-margin controller, while the risk hypothesis examines whether a higher scenario order statistic reduces observed violations. The computation hypothesis concerns the measured budget at the executed scale.

3.2 | System Model and Notation

The general fed-batch model has biomass, substrate, and product concentrations together with culture volume, x = (X, S, P, V) and a clean feed of rate F and substrate concentration S_in. *Eq. (1)* includes the dilution of biomass and product caused by increasing volume. The illustrative inhibited growth law in *Eq. (2)* uses θ = (μ_max, K_S, K_I, K_P, Y, a, b). The executed comparison instead uses the constant-volume system of *Eq. (6)*, a six-parameter law without substrate inhibition. *Table 2* separates the physical quantities from the controller settings.

Table 2. Notation for the robust cultivation control problem.

Symbol	Meaning	Unit
X, S, P	biomass, substrate, and product concentrations	g/L
V	culture volume	L
F	feed rate (control input)	L/h
S_in	feed substrate concentration	g/L
mu_max	maximum specific growth rate	1/h
K_S, K_I	half-saturation and substrate-inhibition constants	g/L
K_P	product-inhibition constant	g/L
Y	biomass yield on substrate	g/g
a, b	growth-associated and non-growth-associated product coefficients	g/g, g/(g h)
theta	kinetic parameter vector	-
Pi(theta | y)	posterior distribution of theta	-
eps	chance-constraint violation budget	-
N, n_s	prediction horizon and scenario count	steps, -

3.3 | Formal Problem Statement

At each control cycle, the ideal stochastic program of *Eq. (4)* balances an expected objective with marginal state constraints of *Eq. (3)*. Its probability must cover the parameter and disturbance laws used to describe the plant. A joint path requirement would instead constrain the event that all inequalities hold simultaneously and would need a separate batch or horizon risk budget. The reported controller approximates the marginal target using deterministic parameter scenarios; it does not solve the ideal program exactly.

$$\begin{aligned} \frac{dX}{dt} &= (\mu - \frac{F}{V})X, \\ \frac{dS}{dt} &= -\frac{\mu X}{Y} + \frac{F}{V}(S_{in} - S), \\ \frac{dP}{dt} &= (a\mu + b)X - \frac{F}{V}P, \\ \frac{dV}{dt} &= F. \end{aligned} \tag{1}$$

$$\mu(S,P) = \mu_{max}\frac{S}{K_S + S + \frac{S^2}{K_I}} \cdot \frac{K_P}{K_P + P}. \tag{2}$$

$$\Pr_{\theta,w}\{g_j(x_{k+\ell}) \le 0\} \ge 1 - \varepsilon_j, \text{ for each l and j.} \tag{3}$$

$$\min_F E_{\theta,w}[J(x, F; \theta, w)], \quad \text{subject to (1), (3), } 0 \le F \le F_{max}, V \le V_{max.} \tag{4}$$

$$\Pi(\theta \mid y) \propto p(y \mid \theta) p_0(\theta). \tag{5}$$

4 | Proposed Methodology

4.1 | Posterior Construction for Kinetic Parameters

A posterior-informed implementation would combine a kinetic likelihood with prior information and update it when assays become available [9], [10]. Such a layer would need parameter identifiability and posterior-predictive checks before its uncertainty could support a control decision. The executed study uses a specified lognormal law directly and contains no assay-driven inference. This distribution is therefore a known simulation input, not an estimated posterior.

4.2 | Scenario-Based Chance-Constrained MPC

At every control cycle, the implemented controller propagates ten deterministic parameter scenarios over the prediction horizon. The scenario set is drawn once per run and reused, as described in Section 5.1. For each horizon step, the ninth ordered substrate prediction must not exceed the ceiling; a quadratic mean-exceedance penalty also enters the objective. This is an empirical nearest-rank 90% planning rule. It permits one of ten scenarios to exceed the ceiling and provides no established 90% population coverage guarantee. The convex scenario results of [8] require assumptions and sample bounds that have not been established for this locally solved nonlinear quantile program.

The scenario order statistic is applied separately at each predicted step. A pointwise scenario constraint does not certify the simultaneous event over the full horizon or batch, and process noise is absent from the planning scenarios. The empirical step and run violation summaries in Section 5 measure these distinct effects. They are diagnostics of the chosen design, not substitutes for a closed-loop probabilistic proof.

4.3 | Comparison Controllers

The nominal MPC benchmark uses one nominal kinetic vector. Hard-scenario MPC enforces the substrate ceiling for every sampled parameter vector. The fixed-margin benchmark, labelled tube-lite in the reported tables and figure, is nominal MPC with a 15% substrate back-off. It contains no propagated tube, affine disturbance feedback, or min-max optimization. Its zero observed violations must therefore be interpreted as a result of this campaign, not as recursive feasibility or worst-case safety.

4.4 | Numerical Implementation

The executed controllers share explicit-Euler prediction and warm-started SLSQP single shooting through SciPy [15]. These are locally solved nonlinear programs [11]; global optimality and recursive feasibility are not established. Previous control blocks initialize the next solve. Reproducible implementation requires the exact objective and terminal surrogate, control-block parameterization, stopping tolerances, unsuccessful-solve policy, and timing protocol. The reported aggregate solve times appear in Section 5.2.

4.5 | Constraint Handling, Complexity, and Failure Modes

The general design distinguishes actuator and vessel limits from recoverable state limits. Its executed instantiation constrains only the dilution-rate actuator and substrate ceiling. Increasing the number of scenario rollouts adds prediction work, but nonlinear-solver iterations also affect total cost. Prior misspecification, poor identifiability, and structural model error are validation concerns for a future learned-parameter implementation; this campaign does not quantify them.

The actuator constraint is $0 \le F \le F_{max}$ in the general fed-batch model and $0 \le u \le 0.5\ h^{-1}$ in the executed constant-volume model. The volume ceiling $V \le V_{max}$ belongs to the general model and is inactive in the experiment.

The substrate ceiling $S \le S_{sup}$ is the general risk constraint. The executed value $S_{max} = 4.0$ g/L is an operational guard; *Eq. (6)* contains no substrate-inhibition term. A product guard $P \le P_{crit}$ or a terminal biomass requirement $X(T) \ge X_{min}$ could be added, but neither is tested here.

Algorithm 1. Empirical scenario quantile MPC used in the simulation.

Input: *Eq. (6)*, current full state, fixed run-level parameter scenarios, horizon H, actuator limits, substrate ceiling, order-statistic level, and penalty coefficient.
Initialize the scenario set and the warm-start control blocks at run start.
At each step, read the simulated full state and propagate every scenario under candidate control blocks.
Enforce actuator bounds and the selected substrate order statistic at each predicted step.
Optimize the mean scenario objective with the quadratic exceedance penalty using warm-started SLSQP.
Apply the first control move, shift the control blocks, and record solve time and realized substrate excursions.
Repeat to T and summarize titer and violations. State estimation, posterior learning, volume constraints, and guarantees of successful or safe fallback are outside the executed algorithm.

5 | Experimental Results and Analysis

5.1 | Experimental Setup

The reported campaign instantiates the control design on a constant-volume, three-state bioreactor surrogate with $x = (X, S, P)$ and manipulated dilution rate $0 \le u \le 0.5\ h^{-1}$. The dilution losses in *Eq. (6)* correspond to a well-mixed continuous-culture model with equal inflow and outflow. They do not describe a changing-volume fed-batch vessel. Growth follows Monod kinetics [1] with multiplicative product inhibition, and product formation uses the Luedeking–Piret structure [7]:

$$\begin{aligned} \frac{dX}{dt} &= (\mu - u)X, \\ \frac{dS}{dt} &= u(S_{in} - S) - \frac{\mu X}{Y}, \\ \frac{dP}{dt} &= (\beta\mu + \gamma)X - uP, \\ \mu &= \mu_{max}\frac{S}{K_S + S}\frac{1}{1 + \frac{P}{K_P}}. \end{aligned} \tag{6}$$

The stated discretization is explicit Euler with $\Delta t = 0.2$ h for 60 steps, giving $T = 12.0$ h. The feed concentration is $S_{in} = 20.0$ g/L and $x_0 = (0.2, 2.0, 0.0)$ g/L. Nominal coefficients are $\mu_{max} = 0.4\ h^{-1}$, $K_S = 0.8$ g/L, $Y = 0.5$ g/g, $\beta = 1.5$ g/g, $\gamma = 0.02$ g/(g h), and $K_P = 6.0$ g/L. The substrate ceiling is 4.0 g/L with a nominal 0.9 marginal probability target. Volume, product, and terminal biomass guards are inactive. Because *Eq. (6)* has no Andrews term [2], the ceiling does not mark a modeled transition into substrate inhibition.

The six parameters follow the study-selected law $\theta = \theta_{nom}\exp(0.2z)$, $z \sim N(0, I_6)$, independently by component. The nominal vector is the median, not the mean; each component has coefficient of variation $\sqrt{(\exp(0.04) - 1)}$, approximately 20.2%. Evaluation crosses 30 parameter draws with five process-noise seeds. Within each of the 150 cells, controllers use the same true parameter vector and noise realization. Additive per-step Gaussian noise has standard deviations (0.01, 0.01, 0.02) g/L, followed by truncation at zero. The two scenario controllers share ten planning parameter draws sampled once per run. A violation is $S > 4.0$ g/L at a sampled step. The campaign reports a total runtime of 168.2 s on a desktop Central Processing Unit (CPU).

OL control solves the nominal single-shooting program once (reported solve time 114 ms) with four 15-step constant-control segments, producing $u = (0.036, 0.019, 0.250, 0.150)\ h^{-1}$. PI control [16] tracks 3.9 g/L with conditional integration; $K_p = 0.2$ and $K_i = 0.005$ were selected from a stated 5×5 nominal-plant search. Nominal MPC uses $H = 10$ steps and two control blocks. Hard-scenario MPC (SMPC) enforces the maximum substrate prediction over ten scenarios. The proposed quantile MPC (CCMPC) uses the ninth order statistic and a quadratic mean squared-exceedance penalty with $\varrho = 2.0$. Fixed-margin MPC (TUBE in the tables) imposes predicted $S \leq 3.4$ g/L. All MPC variants use the stated horizon-end-product and remaining-production surrogate; its exact formula is required to reproduce the comparison. The Chance-Constrained Model Predictive Control (CCMPC)– Stochastic Model Predictive Control (SMPC) comparison changes both the constraint and the penalty, so it is not a pure quantile ablation.

The reported outcomes are final titer $P(T)$, productivity $P(T)/T$, sampled-step violation fraction, peak substrate, and solve time. Tables give mean ± standard deviation, rather than confidence intervals, across the 150 paired cells; lower-decile titer is the mean of the lowest 10% of runs. The reported software versions are Python 3.12.15, NumPy 2.1.3 [17], SciPy 1.14.1 [15], and Matplotlib 3.9.2 [18]. Figures are reported as generated from result files at 300 dpi.

The implemented study isolates the control layer under a known parameter law. It uses full-state feedback and a shared Euler discretization for plant and planner, with parametric mismatch and additive noise. It does not test posterior estimation, structural mismatch, min-max or certified tube control, changing-volume fed-batch operation, or inhibition- and biomass-specific guards. The quantile ablation uses 15 draws crossed with three seeds, a 45-cell subset.

5.2|Main Comparison (RQ1, RQ2)

Table 3 reports the six-method comparison over the 150 paired runs; *Fig. 1* shows substrate trajectories against the chance-constraint bound for the three highest-titer controllers on a representative draw, and *Fig. 2* places every individual run in the violation-titer plane. Three regimes are immediately visible.

Table 3. Closed-loop performance over 150 paired runs (30 parameter draws x 5 process-noise seeds; mean +/- std across runs; violation rate = fraction of the 60 steps with S > S_max = 4.0 g/L; worst-decile titer = mean of the lowest 10% of runs).

Method	Final Titer P(T) (G/L)	Productivity (G/(L H))	Violation Rate (Steps)	Runs With Any Violation	Worst-Decile Titer (G/L)	Compute (Ms/Step)
OL SLSQP (nominal)	2.788 ± 1.130	0.232 ± 0.094	0.491 ± 0.030	150/150 (100.0%)	0.967	0.00 (one-off 114 ms)
PI anti-windup (setpoint 3.9 g/L)	4.798 ± 1.516	0.400 ± 0.126	0.000 ± 0.000	0/150 (0.0%)	2.334	0.00 ± 0.00
Nominal MPC (H = 10)	5.044 ± 1.545	0.420 ± 0.129	0.065 ± 0.110	50/150 (33.3%)	2.490	3.10 ± 0.55
Scenario MPC (10 scenarios, hard)	5.059 ± 1.576	0.422 ± 0.131	0.032 ± 0.094	25/150 (16.7%)	2.508	3.90 ± 0.78
Chance-constrained MPC (proposed)	5.049 ± 1.593	0.421 ± 0.133	0.049 ± 0.108	38/150 (25.3%)	2.451	4.40 ± 0.88
Tube-lite MPC (15% back-off)	4.913 ± 1.506	0.409 ± 0.125	0.000 ± 0.000	0/150 (0.0%)	2.370	3.98 ± 0.69

The nominal OL profile has a reported predicted peak of 4.000 g/L and nominal titer of 4.507 g/L. Applied to the uncertain systems, it gives 2.788 ± 1.130 g/L, or 61.9% of that nominal prediction, and violates the ceiling in 49.1% of sampled steps and all 150 runs. The reported mean peak substrate is 12.22 g/L and the maximum is 14.841 g/L, about 3.7 times the ceiling. The early low feed and mid-course rise expose the sensitivity of an unchanged profile to kinetic mismatch. These results show overfeeding relative to the operational guard; they do not quantify biological damage from substrate inhibition, which is absent from *Eq. (6)*.

PI and fixed-margin MPC have no observed violations in this campaign, with reported peaks of 3.908 and 3.483 g/L. Their mean titers are 4.798 ± 1.516 and 4.913 ± 1.506 g/L. Relative to CCMPC, their differences are −0.251 and −0.136 g/L, corresponding to discounts of approximately 5.0% and 2.7% using CCMPC as denominator. PI has a lower-decile titer of 2.334 g/L but a highest draw-mean titer of 8.467 g/L. Thus the classical loop is competitive within the simulated conditions. Neither zero observed violations nor a fixed planning margin constitutes a safety guarantee.

The mean titers of Nonlinear Model Predictive Control (NMPC), SMPC, and CCMPC differ by at most 0.015 g/L, while pooled step violation fractions are 0.065, 0.032, and 0.049 and runs with any violation are 33.3%, 16.7%, and 25.3%, respectively. NMPC is a point-parameter controller; the other two propagate parameter scenarios. SMPC has the higher reported mean and lower violation frequency than CCMPC, so these data do not establish a performance advantage for the quantile form. Mean solve times are 3.10, 3.90, 4.40, and 3.98 ms for NMPC, SMPC, CCMPC, and fixed-margin MPC. The 4.40-ms value is about 164,000 times shorter than the 720-s sampling interval. This is a measured implementation budget, not a real-time worst-case bound; solver failure and timing variability require separate analysis.

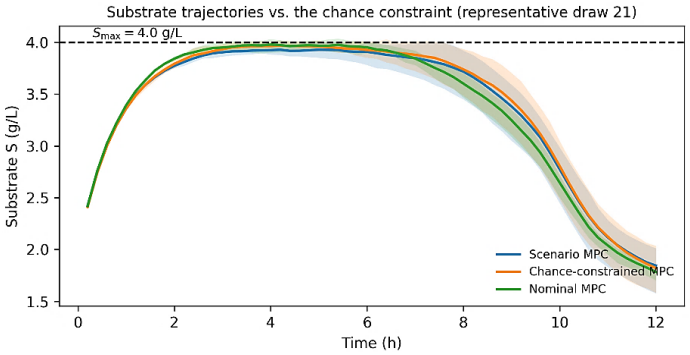


Fig. 1. Substrate trajectories against the chance-constraint bound S_max = 4.0 g/L on the representative parameter draw 21 (the median-performance draw for the proposed controller).

Lines: Seed-mean paths of the three highest-titer controllers; shaded bands: Min-max envelope across the five process-noise seeds; dashed line: S_max (*Fig. 1*).

Fig. 1 shows similar trajectories for the three MPC controllers on draw 21. The paths approach the ceiling, and the seed envelope includes shallow excursions above it. Their similarity is consistent with the small titer differences in *Table 3*; the aggregate violation summaries address behavior beyond this illustrative draw.

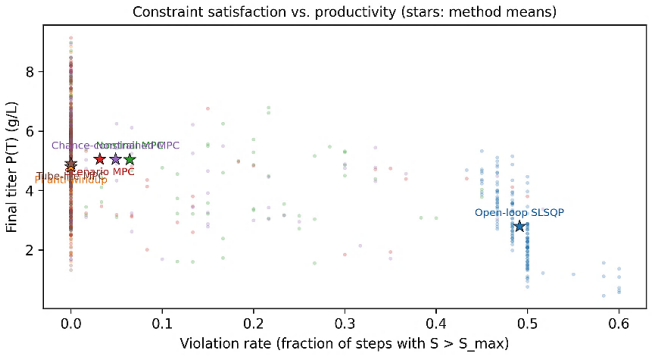


Fig. 2. Per-run substrate-violation fraction versus final titer, with method means marked by stars.

Each method has 150 simulated cells. The tube-lite label denotes fixed-margin nominal MPC (*Fig. 2*).

Fig. 2 places PI and fixed-margin MPC on the observed zero-violation axis, the three MPC variants at higher mean titer with nonzero excursions, and OL control at a substantially larger violation frequency. This comparison is an empirical trade-off across the tested designs, rather than a certified productivity–safety frontier.

5.3 | Ablation: The Quantile Dial

Table 4 compares nearest-rank levels 80%, 90%, and 95% on the stated 45-cell subset. For ten scenarios, the selected ranks are ceil (10q), namely eight, nine, and ten. The 95% rule therefore uses the scenario maximum; its difference from SMPC includes the additional penalty. On the same subset, SMPC has reported titer 5.470 ± 1.610 g/L, violation fraction 0.031, eight violating runs, and lower-decile titer 3.235 g/L.

Table 4. Quantile-level ablation of the chance-constrained controller (45 paired runs: 15 draws x 3 noise seeds; the 90% row is the main configuration of Table 3 restricted to this subset; the last row is scenario MPC on the identical subset).

Constraint Level (Order Statistic Of 10 Scenarios)	Final Titer (G/L)	Violation Rate (Steps)	Runs With Any Violation	Worst-Decile Titer (G/L)
80% (8th of 10)	5.434 ± 1.536	0.079	15/45 (33.3%)	3.231
90% (9th of 10), main configuration	5.482 ± 1.541	0.051	12/45 (26.7%)	3.083
95% (10th of 10) = maximum	5.421 ± 1.530	0.029	8/45 (17.8%)	3.227
Scenario MPC, hard maximum (same subset)	5.470 ± 1.610	0.031	8/45 (17.8%)	3.235

Observed violation fractions decrease from 0.079 to 0.051 to 0.029 as the selected rank increases, and violating-run fractions decrease from 33.3% to 26.7% to 17.8%. Mean titers are 5.434, 5.482, and 5.421 g/L, so the response is not monotone. The reported results include run-level Wilcoxon values $W = 0$, $p = 6.5 \times 10^{-4}$ for 80% versus 95% violations; $W = 435$, $p = 0.358$ and $W = 423$, $p = 0.292$ for 90% versus 80% and 95% titer; and $W = 3$, $p = 2.8 \times 10^{-13}$ for 80% versus 95% titer, with median difference 0.006 g/L and 44 of 45 pairs aligned. The last comparison is statistically different under the reported calculation, although its mean gap is only 0.013 g/L. The three seeds within each of 15 draws limit independent inference. These results support an empirical quantile response, not a universally optimal 90% setting.

The constraint controls an order statistic of scenario substrate predictions at each step. The quadratic exceedance penalty acts on squared positive violations and is not Conditional Value-at-Risk (CVaR), whose optimization construction is different [19]. The 95% row and SMPC have similar violation fractions, 0.029 and 0.031, and a reported titer difference of −0.049 g/L. Without a separate penalty-coefficient ablation and solver diagnostics, the results do not identify the causal contribution or gradient strength of the penalty.

5.4 | Robustness across Parameter Draws

Table 5 reports parameter-draw means over the five noise seeds. CCMPC draw-mean titer spans 1.775–8.269 g/L. The reported analysis identifies draw 29 with $\mu_{max} = 0.243\ h^{-1}$ and draw 13 with $\beta = 2.725$ g/g, and reports titer correlations of +0.84 with μ_{max}, −0.53 with K_S, +0.48 with β, −0.03 with Y, and −0.12 with K_P. These are descriptive associations across 30 draws and do not isolate parameter effects or prove that the same uncertainty controls titer and constraint risk. The substrate guard does not protect final titer against adverse kinetics.

Table 5. Draw-level robustness summary (per-draw means over the 5 noise seeds; 30 draws; per-step satisfaction = 1 - mean violation rate).

Method	Worst-Draw Titer (G/L)	Median-Draw Titer (G/L)	Best-Draw Titer (G/L)	Violation-Free Draws (Of 30)	Worst-Draw Violation Rate	Per-Step Satisfaction
OL	0.626	2.627	5.020	0/30	0.583	0.509
PI anti-windup	1.876	4.880	8.467	30/30	0.000	1.000
Nominal MPC	1.773	4.857	8.193	15/30	0.387	0.935
Scenario MPC	1.957	4.900	8.397	19/30	0.263	0.968
Chance-constrained MPC	1.775	5.026	8.269	15/30	0.300	0.951
Tube-lite MPC	1.717	4.758	7.955	30/30	0.000	1.000

Among runs with any violation, reported mean violating-step fractions are 0.194, 0.189, and 0.193 for NMPC, SMPC, and CCMPC, corresponding to approximately 11–12 sampled steps or 2.3 h. CCMPC violations arise within 15 draws; SMPC has 19 fully violation-free draws. The reported largest MPC peak is 4.058 g/L, approximately 1.5% above the ceiling, compared with 14.841 g/L for OL control. These durations and magnitudes help distinguish excursions, but their mechanisms require trajectory-level examination rather than inference from aggregates alone.

The pooled step satisfaction fractions are 93.5% for NMPC, 95.1% for CCMPC, and 96.8% for SMPC. A pooled average over times and runs does not establish $\Pr\{S(t) \leq S_{max}\} \geq 0.9$ at every time. Fully violation-free runs account for 66.7%, 74.7%, and 83.3%, respectively, and therefore do not meet a 90% batch-level empirical target. On the ablation subset, the 95% order statistic raises violation-free runs from 73.3% to 82.2%. Fixed-margin MPC has zero observed violations, but none of these observations certifies a joint guarantee for future batches.

5.5 | Statistical Analysis

The analysis uses Wilcoxon signed-rank comparisons [20], a Friedman screen [21], and Holm correction across five baseline comparisons [22]. *Table 6* reports those calculations on the 150 paired run cells. The parameter-draw unit is shared by five noise seeds, however, so the cells are nested rather than 150 independent biological draws. The reported p-values are exploratory; confirmatory inference requires a draw-level analysis or a justified clustered procedure.

Table 6. Reported run-level test calculations for final titer. There are 150 paired cells nested within 30 parameter draws; p-values require draw-level or clustered reanalysis before confirmatory interpretation.

Comparison (Final Titer, 150 Paired Runs)	Statistic	P-Value	Holm P	Mean Paired Difference (G/L)
Chance-constrained MPC vs OL SLSQP	W = 0	2.3e-26	1.1e-25	+2.261
Chance-constrained MPC vs PI anti-windup	W = 2667	1.9e-08	7.6e-08	+0.251
Chance-constrained MPC vs nominal MPC	W = 5522	0.792	1.000	+0.005
Chance-constrained MPC vs scenario MPC (primary)	W = 5348	0.555	1.000	-0.010
Chance-constrained MPC vs tube-lite MPC	W = 3615	1.2e-04	3.7e-04	+0.136
Friedman across all six methods	chi-square(5) = 359.1	1.9e-75	-	-

Table 6 reports a Friedman statistic of 359.1, and CCMPC paired mean differences of +2.261, +0.251, +0.005, −0.010, and +0.136 g/L against OL, PI, NMPC, SMPC, and fixed-margin MPC. Its primary CCMPC–SMPC calculation gives W = 5348 and p = 0.555; this absence of rejection does not establish equivalence. The analysis additionally reports a post hoc violation comparison, W = 172.5 and $p = 8.2 \times 10^{-4}$, with 38 versus 25 violating runs. The independent-draw limitation prevents treating them as definitive significance or superiority evidence.

5.6 | Interpretation of the Reported Comparison

The comparison supports three descriptive findings. An unchanged nominal OL profile performs poorly under the stated uncertainty. Feedback controllers reduce substrate excursions substantially, and the scenario representations have similar mean titers but different observed violation frequencies. Within the quantile ablation, a higher selected rank reduces observed violations with small mean-titer changes. The measured optimization cost is modest at this scale. The 26.1% difference between SMPC and NMPC timing does not establish a sublinear scenario-count scaling law, because their constraints and solver paths also differ.

SMPC is ahead of CCMPC in the reported mean titer, violation fraction, and lower-decile titer. Fixed-margin MPC and PI have no observed violations and achieve approximately 97.3% and 95.0% of CCMPC mean titer; PI wins 46 of 150 paired cells. These observations make baseline selection consequential. A practitioner would need application-specific valuations of excursions, computation, and product output before choosing a controller; the simulation does not designate a universally preferable policy.

The known-law scenarios remove estimation error from the experiment, but they do not provide a mathematical optimistic bound on safety under a learned law. Common plant and planner discretization also excludes structural error and integration mismatch. The evidence concerns one simulator, one dispersion scale, one horizon, and the stated paired draws. Its positive result is an empirical mechanism for adjusting observed substrate risk; certified guarantees, posterior learning, and organism-specific validation remain separate tasks.

6 | Discussion

6.1 | Limitations

The study uses a continuous-culture surrogate rather than a volume-explicit fed-batch model. It omits substrate inhibition, dissolved-oxygen and pH interactions, sensor faults, contamination, and estimator delay. Thus the substrate ceiling has an operational role, and final concentration is not a complete measure of harvested mass or economic productivity. Euler step-size convergence and unsuccessful-solve behavior also need verification. The six-component lognormal law is a simulation assumption, not a measured biological variability model.

6.2 | Alternative Viewpoints and Responses

The empirical comparison illustrates when simple policies can be competitive. PI provides direct substrate regulation, and a fixed planning margin reduces observed excursions without scenario propagation. A posterior-aware extension could use estimation tools [6], [9], [10], but credible risk statements would require the resulting distribution to be validated and the control and disturbance assumptions to be stated. Scenario-count, prior-misspecification, structural-mismatch, and penalty-coefficient studies would separate the mechanisms that the present comparison changes together.

7 | Conclusion

This study examines empirical scenario-quantile MPC on a constant-volume continuous-culture surrogate with uncertain kinetics. The reported paired campaign shows similar mean titer across the three MPC controllers, fewer observed violations for hard-scenario MPC than for the ninth-order-statistic controller, and a monotone violation response to the selected order statistic. These results identify a tunable empirical mechanism rather than a certified chance-constrained solution.

At the 90% planning level, CCMPC has reported mean titer 5.049 ± 1.593 g/L, 95.1% pooled step satisfaction, and 74.7% violation-free runs. Its 4.4-ms mean solve time is small relative to the 720-s interval, but neither this timing nor the pooled satisfaction fraction supplies a worst-case or joint guarantee. The reported tests require reanalysis at the parameter-draw unit before inferential claims can be strengthened.

Further validation should use the original executable and raw outputs, specify the terminal objective and failure policy, and analyze uncertainty at the independent-draw level. Subsequent studies can then introduce learned kinetic distributions, explicit disturbance scenarios, and volume-explicit fed-batch dynamics. Formal batch-risk budgeting or CVaR objectives [19] are distinct constructions that would require their own derivation and experiments.

Declarations

Data availability: The study uses simulated bioreactor models. No public access location is specified for the seeds or raw simulation outcomes.

Code availability: The experiment files are identified as exp_27.py, exp_27_results.json, exp_27_f1.png, and exp_27_f2.png. No public access location is specified.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of interest

The authors declare no conflicts of interest.

Author contributions

Tarek Abdelrahman: conceptualization, methodology, software, and writing—original draft. Wolfgang H. Müller: data curation, formal analysis, and visualization. Arafa Abdelzaher Nasef: validation, supervision, and writing—review and editing.

Generative Artificial Intelligence (AI) disclosure: AI-assisted drafting tools were used in preparing the manuscript text. The authors retain responsibility for scientific verification and the accuracy of this disclosure.

References

[1] Monod, J. (1949). The growth of bacterial cultures. *Annual Review of Microbiology*, *3*, 371–394. https://doi.org/10.1146/annurev.mi.03.100149.002103

[2] Andrews, J. F. (1968). A mathematical model for the continuous culture of microorganisms utilizing inhibitory substrates. *Biotechnology and Bioengineering*, *10*(6), 707–723. https://doi.org/10.1002/bit.260100602

[3] Banga, J. R., Balsa-Canto, E., Moles, C. G., & Alonso, A. A. (2005). Dynamic optimization of bioprocesses: Efficient and robust numerical strategies. *Journal of Biotechnology*, *117*(4), 407–419. https://doi.org/10.1016/j.jbiotec.2005.02.013

[4] Mayne, D. Q., Rawlings, J. B., Rao, C. V, & Scokaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, *36*(6), 789–814. https://doi.org/10.1016/S0005-1098(99)00214-9

[5] Mesbah, A. (2016). Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, *36*(6), 30–44. https://doi.org/10.1109/MCS.2016.2602087

[6] Bastin, G., & Dochain, D. (1990). *On line estimation and adaptive control of bioreactors*. Elsevier Amsterdam. 379. https://doi.org/10.1016/C2009-0-12088-3

[7] Luedeking, R., & Piret, E. L. (1959). A kinetic study of the lactic acid fermentation. Batch process at controlled pH. *Journal of Biochemical and Microbiological Technology and Engineering*, *1*(4), 393–412. https://doi.org/10.1002/jbmte.390010406

[8] Calafiore, G. C., & Campi, M. C. (2006). The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, *51*(5), 742–753. https://doi.org/10.1109/TAC.2006.875041

[9] Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and hall cRc. https://doi.org/10.1201/b16018

[10] Doucet, A., De Freitas, N., & Gordon, N. (2001). *Sequential Monte Carlo methods & particle filters resources*. https://doi.org/10.1007/978-1-4757-3437-9

[11] Nocedal, J., & Wright, S. (2006). *Numerical Optimization*. https://doi.org/10.1007/978-0-387-40065-5

[12] J. B. Rawlings, D. Q. Mayne, and M. M. D. (2017). *Model Predictive Control: Theory, Computation, and Design, 2nd ed*. Nob Hill. https://sites.engineering.ucsb.edu/~jbraw/mpc/MPC-book-2nd-edition-1st-printing.pdf

[13] Scokaert, P. O. M., & Mayne, D. Q. (1998). Min max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, *43*(8), 1136–1142. https://doi.org/10.1109/9.704989

[14] Langson, W., Chryssochoos, I., Raković, S. V, & Mayne, D. Q. (2004). Robust model predictive control using tubes. *Automatica*, *40*(1), 125–133. https://doi.org/10.1016/j.automatica.2003.08.009

[15] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., … & van Mulbregt, P. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2

[16] K. J. Åström & T. Hägglund. (1995). *PID controllers: Theory, design, and tuning*. https://portal.research.lu.se/en/publications/pid-controllers-theory-design-and-tuning/

[17] Harris, C. R., Millman, K. J., Van Der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., … & Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357–362. https://doi.org/10.1038/s41586-020-2649-2

[18] Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, *9*(03), 90–95. https://doi.org/10.1109/mcse.2007.55

[19] Rockafellar, R. T., Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*(3), 21–41. https://doi.org/10.21314/JOR.2000.038

[20] Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin*, *1*(6), 80–83. https://doi.org/10.2307/3001968

[21] Friedman, M. (1937). The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, *32*(200), 675–701. https://doi.org/10.1080/01621459.1937.10503522

[22] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, *6*(2), 65–70. https://www.jstor.org/stable/4615733