Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimal Stopping of Batch Processes with Stochastic Dynamics – A Study of When to Act under Uncertainty

Rafif S. Ramadhan[ab], Luca Grebe[c], Maximilian Maschmeier[c], Johannes Pastoors[c], Eike Cramer[bd]*

[a] Davidson School of Chemical Engineering, Purdue University, West Lafayette, IN 47907, USA
[b] RWTH Aachen University, Process Systems Engineering (AVT.SVT), Aachen 52074, Germany
[c] RWTH Aachen University, Biochemical Process Engineering (AVT.BioVT), Aachen 52074, Germany
[d] University College London, Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, London WC1E 7JE, United Kingdom
* Corresponding Author: e.cramer@ucl.ac.uk

ABSTRACT

Mathematical models in process systems engineering (PSE) are widely used to support decision-making in design and operation, but they are mostly limited to deterministic models. For biochemical systems, the biological variability can give rise to stochastic dynamics. This work addresses the question of when to act in such processes, as the stochastic dynamics affect the timing of important events. We consider the case of batch production of malic acid using *Ustilago trichophora*. The goal is to predict when the substrate concentration falls below a predefined threshold. We extend an existing deterministic model of the process to a stochastic differential equation (SDE) formulation by introducing a Monod-like noise term. Simulations of the SDE model reveal a distribution of substrate depletion times and a deviation between the mean of the stochastic trajectory and the deterministic solution due to nonlinear effects. To determine optimal intervention times under uncertainty, we formulate a finite-horizon optimal stopping problem and solve it using the Longstaff–Schwartz algorithm, also known as the Least-Squares Monte Carlo (LSMC). The resulting distribution of optimal stopping times from the LSMC algorithm is shown to closely match the actual first threshold hitting times obtained from *a posteriori* analysis of the stochastic simulations, as confirmed by a two-sample Kolmogorov–Smirnov test. The results demonstrate that optimal stopping provides a framework for decision-making in stochastic biochemical processes, enabling risk-aware operational strategies beyond deterministic optimisation.

Keywords: Stochastic differential equations (SDEs), optimal stopping, decision-making under uncertainty

INTRODUCTION

Biochemical processes play an increasingly important role in the transition toward more sustainable chemical production. By exploiting the metabolic activity of microorganisms, bioprocesses enable the conversion of renewable feedstocks into fuels and platform chemicals that are essential for downstream manufacturing [1]. As these processes move from laboratory to industrial scale, the need for reliable modeling and decision-support tools becomes increasingly critical.

In process systems engineering, mathematical models are routinely used to support design, optimization, and operational decision-making [2]. For biochemical systems, such models are traditionally formulated as ordinary differential equations (ODEs). While such deterministic models are valuable for capturing average system behavior, biological processes are inherently variable due to intracellular heterogeneity, environmental fluctuations, and nonlinear reaction kinetics [3]. As a result, deterministic predictions may fail to adequately represent the range of possible process trajectories, particularly when decisions depend on threshold-based events such as substrate depletion or intervention timing.

Although uncertainty is often addressed in process control frameworks, stochastic effects are rarely incorporated explicitly into optimization and planning problems for biochemical batch processes [2]. Existing studies on biological noise primarily focus on stochastic simulation and analysis [3], while the implications of

stochastic dynamics for decision-making and optimization remain underexplored. In particular, there is a lack of systematic approaches for determining intervention times in stochastic bioprocesses that account for the full distribution of possible system trajectories rather than a single deterministic prediction.

Stochastic decision-making problems have been extensively studied in financial mathematics, where uncertainty is intrinsic to system dynamics [4, 5]. Here, optimal stopping theory provides a framework for determining when to take an action in a stochastic environment, e.g., to maximize an expected reward. The Longstaff–Schwartz algorithm, also known as Least-Squares Monte Carlo (LSMC), enables the solution of optimal stopping problems using simulated trajectories and regression-based approximation of conditional expectations [6].

In this work, we formulate a stochastic model of a biochemical batch reactor using stochastic differential equations (SDEs) and embed this model within an optimal stopping framework. The stochastic formulation explicitly captures process variability through state-dependent noise, while the optimal stopping problem is defined in terms of substrate depletion. The Longstaff–Schwartz algorithm is then applied to learn an optimal intervention time directly from simulated stochastic trajectories.

In this work, we formulate a stochastic model of a batch bioreactor to produce malic acid via *Ustilago trichophora* [7]. We extend an existing ODE model to an SDE formulation to represent the stochastic process. Based on the stochastic model, we analyse an optimal stopping problem using the Longstaff-Schwartz algorithm. The results from the algorithm show that the stochastic system exhibits behaviours that lead to different intervention strategies compared to the deterministic system. The findings of this work highlight the importance of stochastic modelling for decision-making problems.

STOCHASTIC PROCESS MODEL

This work considers the batch production of the platform chemical malic acid using *U. trichophora* as biomass. For this process, the substrates sucrose, fructose, glucose, and ammonium (NH_4^+), as the source of nitrogen, are supplied. During the initial growth phase, sucrose is enzymatically hydrolysed to the monosaccharides glucose and fructose and utilised mainly for biomass formation. As nitrogen becomes depleted, the culture shifts from growth to production. Under these conditions, malic acid is formed as the predominant product. The deterministic ODE model is taken from Maschmeier [7].

$$\frac{d}{dt}\begin{bmatrix} X_a \\ X_i \\ Suc \\ FruGlu \\ N \\ P \end{bmatrix} = \begin{bmatrix} \mu_a \\ \mu_i \\ -q_{split} \\ q_{split} - \frac{\mu_a}{Y_{XS,a}} - \frac{\mu_i}{Y_{XS,i}} - \frac{q_P}{Y_{PS}} \\ -\frac{\mu_a}{Y_{XN}} \\ q_P \end{bmatrix} X_a \tag{1}$$

Here, the specific growth and product formation rates are:

$$\mu_a = \mu_{a,max} \frac{FruGlu}{FruGlu+K_{FG}} \frac{N}{N+K_N} \tag{2}$$

$$\mu_i = \mu_{i,max} \frac{FruGlu}{FruGlu+K_{FG,2}} \frac{K_{IN}}{N+K_{IN}} \left(1 - e^{\frac{X_i/X_a - \phi_X}{\chi_{acc}}}\right) \tag{3}$$

$$q_{split} = q_{split,max} \frac{Suc}{Suc+K_S} \tag{4}$$

$$q_P = q_{P,max} \frac{FruGlu}{FruGlu+K_{PFG}} \frac{K_{IN}}{N+K_{IN}} \frac{K_{IP}}{K_{IP}+\frac{w_N N}{X_a+X_i}} \tag{5}$$

For the definition of the variables in the deterministic model, see Table 1.

Table 1: State variables in the model.

Symbol	Description
X_a	Active biomass concentration
X_i	Inactive biomass concentration
Suc	Sucrose concentration
$FruGlu$	Sum of Fructose and Glucose concentrations
N	Nitrogen concentration
P	Malic acid concentration

We extend the deterministic model to describe the stochastic dynamics by adding a differential noise term to the ODE, turning the model equation into an SDE. In general, SDEs have the form

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dB_t, \tag{6}$$

where $\mu(X_t)$ is the drift term, $\sigma(X_t)$ the diffusion term, and B_t the Brownian motion (standard Wiener process) [8]. The deterministic model Eq. 1 serves as the drift term. In the following, we will define the noise term.

In biochemical processes, biomass activities are dependent on the availability of substrates. Therefore, the random fluctuations are modelled to follow the same phenomenon. The biomass activities exhibit minimal fluctuations when the substrate is depleted and reach a saturation point when the substrate is abundant. To capture the substrate dependence, this work proposes modelling the noise term σ_i in component *i* using a Monod-like equation. The Monod model is chosen because it relates microbial activities to substrate availability [9]. We note that Eq. 7 does not itself constitute the stochastic noise term. Rather, it defines the substrate-dependent

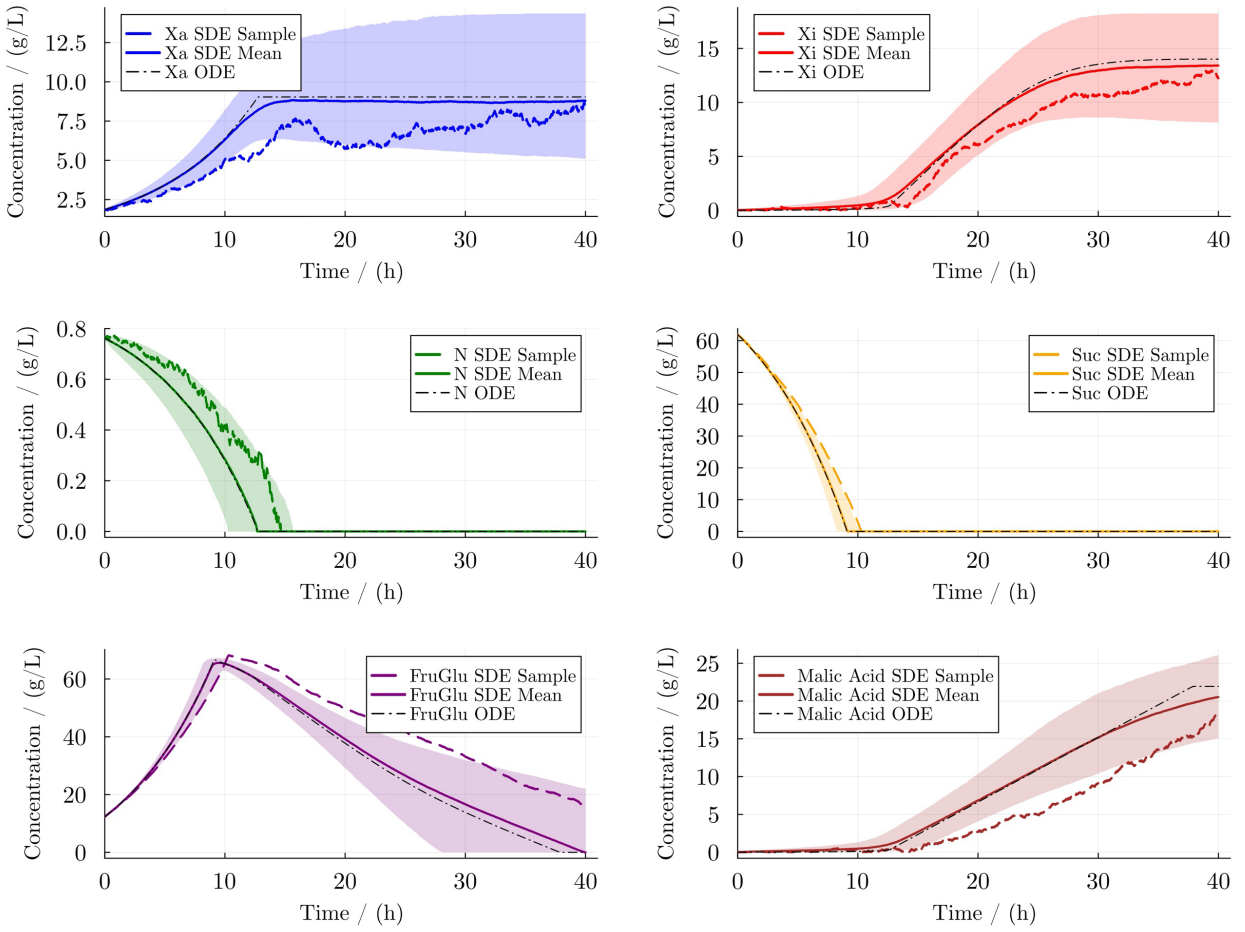

Figure 1. Sample stochastic scenario (dashed), deterministic profile (dash-dot), and SDE summary of several simulated scenarios (Set 1) consisting of the averaged concentration profile (solid line) and the 5th to 95th percentile of the scenario distribution (shaded region). From top to bottom, left to right: active biomass, inactive

diffusion coefficient σ_i, which modulates the stochastic fluctuations in a physically meaningful way. The noise itself enters the SDE through the Wiener increment dB_t.

$$\sigma_i = \sigma_{i,max} \frac{S}{K_{\sigma_i}+S} \tag{7}$$

Here, $\sigma_{i,max}$ is the maximum noise intensity and K_{σ_i} is the associated half-saturation constant. The variable S represents the substrate, which can be taken up by either $Suc, FruGlu,$ or N, depending on which equation the noise term is in. The full noise equation is given in Eq. 8. Each row represents the diffusion coefficient of the corresponding state variable. Thus, together with Equation (1), the mole balance of each species is given as a stochastic differential equation in vector form, with the drift term given by Equation (1) and the diffusion term by Equation (8). The half-saturation constant in the noise term K_{σ_i} is set to the value of the reaction kinetics.

This work aims to discuss the optimal stopping algorithm, and no measurements of the variance parameters are available. Instead, the parameter values in the noise terms are selected arbitrarily to be $\sigma_{i,max} = 0.05\ 1/\sqrt{\mathrm{h}}$ for all components. The validation of this proposed noise model and its parameters based on experimental data falls outside the scope of this work and, therefore, can be conducted as a separate study, where other noise models can also be considered.

$$\sigma(X_t) = \begin{bmatrix} \sigma_{X_a,max} \frac{FruGlu}{K_{FG}+FruGlu} \\ \sigma_{X_i,max} \frac{FruGlu}{K_{FG,2}+FruGlu} \\ \sigma_{Suc,max} \frac{Suc}{K_S+Suc} \\ \sigma_{FG,max} \frac{FruGlu}{K_{FG}+FruGlu} \\ \sigma_{N,max} \frac{N}{K_N+N} \\ \sigma_{P,max} \frac{FruGlu}{K_{PFG}+FruGlu} \end{bmatrix} X_a \tag{8}$$

In this work, we simulated two sets of stochastic sample paths, which will be referred to as Sets 1 and 2 in the following. Each set contains five thousand scenarios. Figure 1 shows a sample stochastic scenario from Set 1

and the summary of Set 1 obtained via the Euler-Maruyama scheme. The Euler-Maruyama method is the stochastic extension of Euler's method to numerically solve SDEs [8]. The ensemble summary consists of an averaged profile and the 95% confidence interval at every time point. The initial values are provided in Table 2, and the model parameters in Table 3.

Table 2: Initial values

Variable	Concentration (g/L)
X_a	1.85
X_i	0
Suc	61.96
$FruGlu$	12.3
N	0.762
P	0

Figure 1 shows that the average concentration profile of the SDE does not align with the deterministic ODE model (shown in dashed lines) and that the concentrations from the stochastic SDE model are distributed over a range of values. This observation suggests that different results may be obtained when the models are applied to further analysis, such as optimisation or process control. The difference between the stochastic scenarios and the deterministic model can be attributed to the interplay between the model's nonlinearities and the noise. This interplay can be described using Jensen's inequality, which states that the mean of a nonlinear function of a random variable is not equal to the function of the expected value itself [8, 10].

OPTIMAL STOPPING

Optimal stopping is a problem commonly encountered in derivatives finance, specifically in the context of American Options. An optimal stopping problem poses the question: "When should one exercise an American Option to maximise the payoff?" For more background information on derivatives finance, the reader is referred to [4, 5].

Least-Squares Monte Carlo

LSMC was originally developed for the valuation of American Options [6]**.** However, the algorithm itself only requires simulated trajectories and a reward function. Thus, the algorithm is model-agnostic, and the method is applicable to problems outside of finance if the underlying model is a stochastic process. This work investigates the application of LSMC to solve an optimal stopping problem, where the underlying stochastic process is the stochastic reactor model.

Consider the SDE given in Eq. 6. The objective of an optimal stopping problem is to determine a time τ that maximises the expected value of a certain reward function G_t. The reward function encodes the performance criterion associated with stopping the process at time t.

$$V^* = \sup_{\tau \in [0,T]} \mathbb{E}[G_\tau] \tag{9}$$

Table 3: Model parameters

Symbol	Description	Value	Unit
$\mu_{a,max}$	Max. growth rate of active biomass	0.125	1/h
$\mu_{i,max}$	Max. growth rate of inactive biomass	0.125	1/h
$q_{split,max}$	Max. split rate	1.985	1/h
$q_{P,max}$	Max. product formation rate	0.095	1/h
K_S	Half-saturation constant of sucrose	0.00321	g/L
K_{FG}	Half-saturation constant of active biomass	0.147	g/L
$K_{FG,2}$	Half-saturation constant of inactive biomass	3.277	g/L
K_N	Half-saturation constant of ammonia	3.8E-5	g/L
K_{PFG}	Half-saturation constant of product	0.0175	g/L
K_{IN}	Nitrogen inhibition constant	0.0147	g/L
K_{IP}	Product inhibition constant	0.147	g/L
w_N	Internal nitrogen mass fraction	0.08	-
χ_{acc}	Acceleration factor	0.3	-
ϕ_X	Max. ratio of inactive to active biomass	1.56	-
$Y_{XS,a}$	Yield factor FruGlu to active biomass	0.531	-
$Y_{XS,i}$	Yield factor FruGlu to inactive biomass	0.799	-
Y_{PS}	Yield factor FruGlu to product	0.508	-
Y_{XN}	Yield factor ammonia to biomass	9.428	-

This formulation constitutes a finite-horizon optimal stopping problem, i.e., the process must be stopped before or at time T. Furthermore, the stopping decision must be adapted to the filtration generated by the stochastic process $\{\mathcal{F}_t\}$. In other words, the stopping decision must be made based on the observed process up to time t. The first step of the algorithm is to approximate the continuous time horizon on a discrete time grid $\{t_n\}_{n=0}^{N}$. Here, G_n is the reward from stopping at time t_n. The algorithm proceeds with backward induction to determine the optimal stopping time. The backward induction corresponds to the dynamic programming recursion,

Eq. 10, with the terminal condition $S_N = G_N$.

$$S_n = \max\{G_n, \mathbb{E}[S_{n+1}|\mathcal{F}_n]\} \tag{10}$$

At each time step t_n, the decision is obtained by comparing the immediate reward G_n with the continuation value, which is defined as the conditional expectation of the future rewards given the current information $\mathbb{E}[S_{n+1}|\mathcal{F}_n]$. The optimal stopping time is thus the earliest time at which stopping is preferable to continuation, i.e., $G_n > \mathbb{E}[S_{n+1}|\mathcal{F}_n]$.

The main component of the LSMC algorithm is to approximate the continuation value, which is essential in the backward induction step. Longstaff and Schwartz proposed approximating the continuation value via simulation and regression [6]. First, a number of sample paths of the stochastic process are generated via Monte Carlo simulation. At each time step, the continuation value is approximated by a regression of realised future rewards onto a set of basis functions ϕ_i:

$$\mathbb{E}[\, G_{n+1} \mid \mathcal{F}_n \,] \approx \sum_{i=0}^{K} \beta_i \, \phi_i(X_n) \tag{11}$$

Here, the regression coefficients β_i are obtained via least-squares minimisation over simulated paths where stopping is admissible. The original method uses the first three weighted Laguerre polynomials. However, numerical tests indicate that the choice of basis polynomials, e.g., unweighted Laguerre, Hermite, Legendre, or Chebyshev, is trivial [6].

Optimal Stopping of Stochastic Batch Reactor

We now formulate a finite-horizon optimal stopping problem based on the stochastic batch reactor model. The objective is to determine the intervention time at which the substrate concentration approaches a predefined depletion threshold ϵ. We define the reward function to be

$$G_t = -(FruGlu(t) - \epsilon)^2, \tag{12}$$

which is maximised when $FruGlu(t) = \epsilon$. The depletion threshold is defined as 5% of the initial concentration of sucrose, fructose, and glucose $\epsilon = 0.05 \cdot (FruGlu(t_0) + Suc(t_0))$. The optimal stopping problem is given by Eq. 9, and the continuation value approximation is

$$\mathbb{E}[\, G_{n+1} \mid \mathcal{F}_n \,] \approx \sum_{i=0}^{3} \beta_i \, \phi_i(FruGlu(t_n)). \tag{13}$$

In this work, we use the first four unweighted Laguerre polynomials as basis functions:

$$\phi_n(x) = \frac{e^x}{n!}\frac{d^n}{dx^n}(x^n e^{-x}) \text{ for } n \in \{0, 1, 2, 3\}. \tag{14}$$

RESULTS AND DISCUSSION

The Longstaff-Schwartz algorithm is applied to the formulated optimal stopping problem to obtain a distribution of optimal intervention time. Figure 2 shows the distribution of stopping times obtained from the LSMC algorithm along with a kernel density estimate (KDE) of the distribution. The resulting distribution is unimodal with a mean of 32.91 h and a mode at 33.46 h, which is significantly earlier than the time the deterministic model reaches the depletion threshold, which is at 35.68 h.

To assess whether the LSMC algorithm correctly identifies the relevant stopping point, the distribution of the optimal stopping times is compared to the distribution of the first hitting times. We define the first hitting time as the time when the stochastic path first crosses the depletion threshold, which is obtained from *a posteriori* analysis of the generated stochastic sample paths. The distribution of the first hitting times is shown in Figure 3. We found the average first hitting time to be 32.91 h and the mode of the distribution is 33.47 h, which matches the statistics of the optimal stopping times.

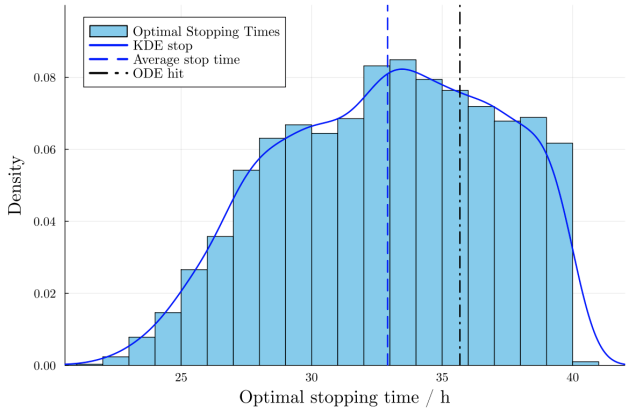

Figure 2. Distribution of the optimal stopping times from LSMC and the corresponding KDE.

The similarity between the two distributions is further quantified using a two-sample Kolmogorov-Smirnov test. The test fails to reject the null hypothesis that the two samples are drawn from the same distribution, with a negligible maximum deviation between the empirical cumulative distribution functions obtained from the time distributions, which is shown in Figure 4. This statistical agreement supports the conclusion that the LSMC algorithm yields consistent results.

To assess the robustness of the LSMC method, the analysis was repeated on Simulation Set 2. Set 2 differs in that the generated scenarios now exhibit increased variability and multimodality. However, the average optimal stopping time of 32.96 h differs only slightly from the average time found in the first simulation set (Set 1). This similarity is shown in Figure 5. The Kolmogorov-Smirnov test again indicates statistical similarity between the distributions of the optimal stopping times and the actual first hitting times for Set 2. This finding suggests that the

LSMC algorithm can reliably identify stopping times across different stochastic realisations, considering the nonlinearities of the stochastic process.

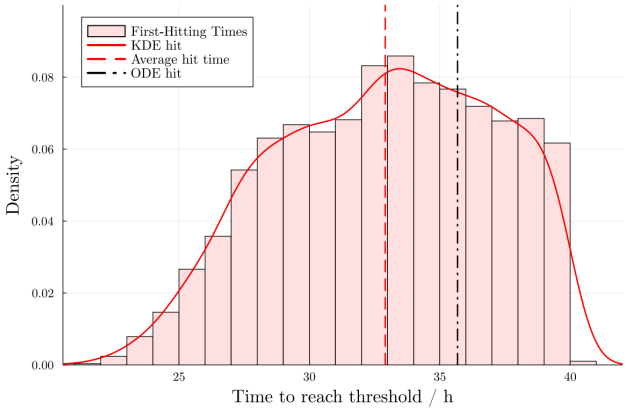

Figure 3. Distribution of the first hitting times (actual times when the substrate reaches the threshold) and the corresponding KDE.

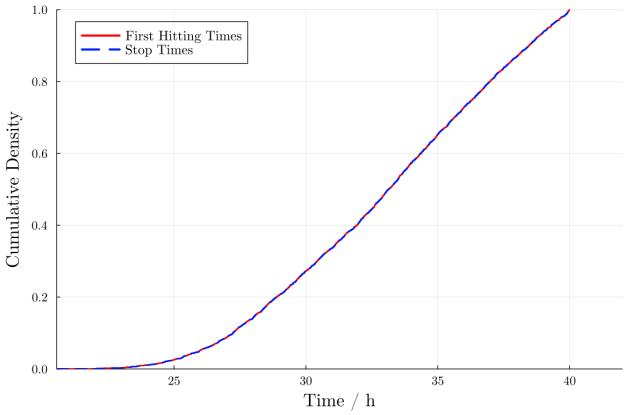

Figure 4. Cumulative distribution function (CDF) of the first hitting times and the optimal stopping times

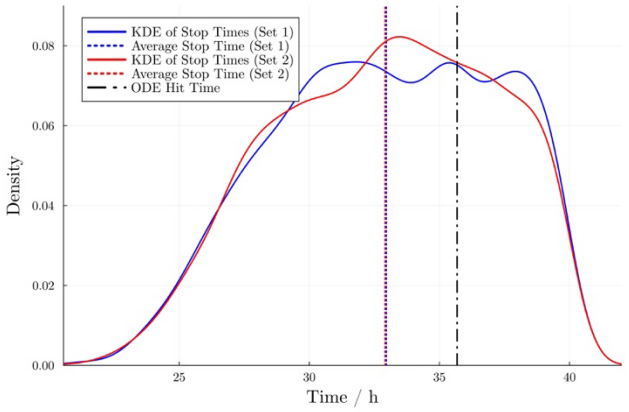

Figure 5. Optimal stopping time distributions from Simulation Sets 1 (red) and 2 (blue).

CONCLUSION

This work demonstrates that stochastic optimal stopping provides a viable and interpretable framework for intervention timing in biochemical batch processes, e.g., by adding new substrate or stopping the production process. By extending a deterministic bioprocess model to a stochastic formulation and coupling it with a regression-based optimal stopping algorithm, the study shows that uncertainty-aware decision-making can lead to systematically earlier interventions than deterministic predictions.

The results demonstrate that stochastic dynamics have a measurable impact on the timing of interventions in biochemical batch processes. Across both simulation sets, the obtained optimal stopping times consistently precede the time the deterministic model reaches the depletion threshold. This shift reflects the influence of stochastic fluctuations on process trajectories and highlights the limitations of deterministic models in decision-making problems under uncertainty.

We have also shown through the close agreement between the optimal stopping time distribution and the distribution of the first-hitting times that the LSMC algorithm accurately captures the stochastic dynamics of the system with little sensitivity across different sets of stochastic realisations. This finding provides empirical validation of the optimal stopping formulation and highlights the broader applicability of methods originating in financial mathematics to process systems engineering problems involving stochastic dynamics. The effectiveness of the optimal stopping framework depends critically on the fidelity of the underlying stochastic model. Future work should therefore prioritise validation and refinement of stochastic models using experimental data, including the exploration of species-specific or correlated stochastic perturbations.

A notable downside of the SDE model in this work is that it allows for physically unrealistic behaviour, e.g., the small increase in nitrogen concentration observed in Figure 1**.** Future work will focus on modelling with alternative stochastic models with an emphasis on physical constraints, specifically mass and energy balances as well as nonequilibrium reactions. Further extensions may include the application of the proposed framework to fed-batch or continuous systems, where real-time measurements enable dynamic intervention strategies. Integrating stochastic optimal stopping with techno-economic analyses also represents a promising direction, allowing intervention decisions to be evaluated directly in terms of operational cost or economic performance. Together, these developments would move stochastic optimal stopping from a methodological concept toward a practical decision-support tool for industrial bioprocess operations.

AUTHOR IDENTIFIERS

Author ORCIDs:
Ramadhan, R.S.: 0009-0007-6861-0066
Grebe, L.: 0009-0009-5532-5797
Pastoors, J.: 0000-0003-2654-0593
Cramer, E.: 0000-0002-6882-5469

REFERENCES

1. Almquist J, Cvijovic M, Hatzimanikatis V, Nielsen J, Jirstrand M. Kinetic models in industrial biotechnology – improving cell factory performance. Metabolic Engineering 24:38-60 (2014). https://doi.org/10.1016/j.ymben.2014.03.007
2. M. D. Graham and J. B. Rawlings. Modeling and Analysis Principles for Chemical and Biological Engineers. Santa Barbara: Nob Hill Publishing, 2022. isbn:9780975937761.
3. Pischel D, Sundmacher K, Flassig RJ. Efficient simulation of intrinsic, extrinsic and external noise in biochemical systems. Bioinformatics 33:i319-i324 (2017). https://doi.org/10.1093/bioinformatics/btx253
4. Privault N. Introduction to stochastic finance with market examples. Chapman and Hall/CRC (2022). https://doi.org/10.1201/9781003298670
5. Glasserman P. Monte carlo methods in financial engineering. Springer New York (2003). https://doi.org/10.1007/978-0-387-21617-1
6. Longstaff FA, Schwartz ES. Valuing american options by simulation: a simple least-squares approach. Rev. Financ. Stud. 14:113-147 (2001). https://doi.org/10.1093/rfs/14.1.113
7. , Diaz F, Sommerfeld M, , Hovestadt G, , Latacz D, , Friedrich B, . Lessons learned from attempts at minimising CO2 emissions in process metallurgy – pyrolysed secondary raw materials, bio-coke, and hydrogen as alternative reducing agents. 12th International Conference of Molten Slags, Fluxes and Salts (MOLTEN 2024) Proceedings :1519-1532 (2024). https://doi.org/10.62053/zkqg2762
8. Thygesen UH. Stochastic differential equations for science and engineering. Chapman and Hall/CRC (2023). https://doi.org/10.1201/9781003277569
9. Monod J. THE GROWTH OF BACTERIAL CULTURES. Annu. Rev. Microbiol. 3:371-394 (1949). https://doi.org/10.1146/annurev.mi.03.100149.002103
10. A. Nobel. Convex Functions and Jensen's Inequality. Sept. 2024. url: https://nobel.web.unc.edu/teaching/stor- 555-theoretical- statistics/stor- 555-lecture-notes/