



HAL
open science

Optimal control of bioproduction in the presence of population heterogeneity

Davin Lunz, Joseph Frédéric Bonnans, Jakob Ruess

► **To cite this version:**

Davin Lunz, Joseph Frédéric Bonnans, Jakob Ruess. Optimal control of bioproduction in the presence of population heterogeneity. 2021. hal-03445175v2

HAL Id: hal-03445175

<https://inria.hal.science/hal-03445175v2>

Preprint submitted on 18 Jan 2023 (v2), last revised 19 Jan 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal control of bioproduction in the presence of population heterogeneity

Davin Lunz^{*1,2}, J. Frédéric Bonnans³, and Jakob Ruess^{1,2}

¹Inria Paris, 2 rue Simone Iff, 75012 Paris, France

²Institut Pasteur, 28 rue du Docteur Roux, 75015 Paris, France

³Université Paris-Saclay, CNRS, CentraleSupélec, Inria, Laboratory of signals and systems, 91190, Gif-sur-Yvette, France

Abstract

Cell-to-cell variability, born of stochastic chemical kinetics, persists even in large isogenic populations. In the study of single-cell dynamics this is typically accounted for. However, on the population level this source of heterogeneity is often sidelined to avoid the inevitable complexity it introduces. The homogeneous models used instead are more tractable but risk disagreeing with their heterogeneous counterparts and may thus lead to severely suboptimal control of bioproduction. In this work, we introduce a comprehensive mathematical framework for solving bioproduction optimal control problems in the presence of heterogeneity. We study population-level models in which such heterogeneity is retained, and propose order-reduction approximation techniques. The reduced-order models take forms typical of homogeneous bioproduction models, making them a useful benchmark by which to study the importance of heterogeneity. Moreover, the derivation from the heterogeneous setting sheds light on parameter selection in ways a direct homogeneous outlook cannot, and reveals the source of approximation error. With view to optimally controlling bioproduction in microbial communities, we ask the question: when does optimising the reduced-order models produce strategies that work well in the presence of population heterogeneity? We show that, in some cases, homogeneous approximations provide remarkably accurate surrogate models. Nevertheless, we also demonstrate that this is not uniformly true: overlooking the heterogeneity can lead to significantly suboptimal control strategies. In these cases, the heterogeneous tools and perspective are crucial to optimise bioproduction.

Contents

1	Introduction	3
2	Heterogeneous bioproduction models and their optimal control and approximation	3
2.1	The influence of heterogeneity in bioproduction	3
2.2	Modular model class	4
2.2.1	Population growth	5
2.2.2	Bioreactor removal	6
2.2.3	Discrete dynamics	7
2.2.4	Continuum dynamics	7
2.2.5	Optimal control of the PDE	8
2.3	Moment approximations of model components	8
2.3.1	Population growth	9
2.3.2	Bioreactor removal	9
2.3.3	Discrete dynamics	10
2.3.4	Continuum dynamics	10
2.3.5	Initial conditions	10
2.3.6	Optimal control of the ODE	11
3	Bioproduction by direct induction	11
3.1	PDE formulation	11
3.1.1	Concentration paradigm	11
3.1.2	Copy-number paradigm	12

*davin.lunz@inria.fr

3.2	A moment closure approximation	13
3.2.1	Optimal control of the ODE model	13
3.3	Numerical results: comparing the PDE and ODE models	14
4	Inducing a transcription factor that modulates production	15
4.1	PDE copy-number formulation	15
4.2	A moment closure approximation	16
4.2.1	Optimal control of the ODE model	17
4.3	Numerical results	17
4.3.1	Parametrised oscillatory controls	17
4.3.2	Optimal controls	18
4.3.3	Comparing the optimal controls	20
4.4	Optimising the period	21
5	Inducing differentiation in bioproducing consortia	22
5.1	PDE concentration formulation	22
5.2	Moment closure approximations	23
5.2.1	Zeroth-order closure	23
5.2.2	First-order closure	24
5.2.3	Optimal control of the ODE model	25
5.3	Numerical results: comparing the PDE and ODE models	25
6	Differentiation in consortia with modulated production	26
6.1	PDE concentration formulation	26
6.2	Moment closure approximations	28
6.3	Numerical results: comparing the PDE and ODE models	29
7	Conclusions	32
	References	32
	Appendix	36
A	Optimal control of the PDE	36
A.1	The abstract setting	36
A.2	The modular model class	37
A.3	PDE discretisation	38
A.3.1	Population growth	38
A.3.2	Bioreactor removal	39
A.3.3	Discrete dynamics	39
A.3.4	Continuum dynamics	39
A.4	Optimal control of the discretised PDE	40
B	Details of the moment closure calculations	40
B.1	Copy-number population growth	41
B.2	Continuum dynamics	42
C	Singular arcs in the single-species ODE model	42
D	Division ratios	45
E	Numerical artefact	45
F	Increasing the resolution	46
G	MPC with partial observability	48
H	Optimal control near the terminal time	48
I	Zeroth-order approximation and fitting parameters	52

1 Introduction

Stochastic gene expression leads to cell-to-cell variability within a population of genetically identical cells, which can manifest in heterogeneous behaviour, such as response to stimuli or cell-fate decisions [3, 7, 18, 42]. Our ability to image and track single-cell dynamics and model their statistics has led to a thorough mechanistic understanding of the behaviour of individual cells and communities [43]. Nevertheless, when it comes to bioengineering, bioproduction models typically neglect this heterogeneity due to the complexity it introduces [16], especially when it comes to process control [17]. When this nonuniformity is ignored, as is the case in homogeneous models, the dynamics can be heavily skewed. For example, small subpopulations that are weak bioproducers may be fast growers, having alleviated the onerous burden of production [9, 20, 33, 51]. These cells may dominate the population and exhibit vastly different behaviour (such as gene expression dynamics, metabolic activity, growth rate, etc.) than the mean dynamics predict. This casts doubt over whether homogeneous models can be used to optimally control bioproduction in the presence of population heterogeneity.

The aim of this work is twofold. First, we pose a collection of heterogeneous multiscale models that cover a variety of behaviour and features that routinely appear in bioproduction applications. The models take the form of systems of partial differential equations and partial integro-differential equations (abbreviated simply as PDEs) and may be derived from first principles [14] or from the chemical master equation — the central governing equation in chemical kinetics modeling [26, 27]. Crucially, our population-level models incorporate sources of heterogeneity born of single-cell stochastic gene expression that encompass a class of models of diverse biological relevance [11, 27, 44]. We study the optimal control of these multiscale models and thereby attempt to unite the heterogeneous population models arising from the important stochastic chemical kinetics with optimal control theory and practice.

Second, we study order-reduction techniques based on a moment closure approximation of the heterogeneous models, which yield homogeneous surrogates. Deriving these reduced models from the heterogeneous setting provides several key advantages over more common ODE approaches that ignore heterogeneity from the outset. For example, many parameters in the reduced-order ODE models may be expressed as functions of the underlying heterogeneous single-cell dynamics, allowing closed-form expressions for parameters that would otherwise have to be fitted [5]. Moreover, the derivation allows us to clearly identify the source of the discrepancy between the reduced-order models and their heterogeneous counterparts. This allows us to pin down the precise assumptions that would otherwise remain tacit.

The heterogeneous models are more challenging to analyse and more computationally demanding to optimise numerically. This begs the question of how important the heterogeneity is, particularly with regard to the control. Studying the homogeneous approximations and their optimal control allows us to directly address this question. We find that the homogeneous models can certainly be quantitatively reliable surrogates. However, this is not always the case, and the heterogeneity may influence the control to the extent that the reduced-order models perform relatively poorly. In the final case we study, for example, this poor performance manifests in significantly reduced bioproduction.

After developing the theory for a class of bioproduction models, we demonstrate the optimal control results on a range of examples with a wide array of biological features. The underlying stochastic chemical kinetics include gene expression in both birth–death as well as bursts [12, 24]. We consider gene expression from the concentration perspective, and separately from the copy-number perspective [34]. We show how these processes are modeled both in the PDE setting, as well as how these differences manifest in the moment closure approximations. We consider models controlled by light-induced responses in both monocultures and consortia [6, 23, 38, 53].

The paper is organised as follows. In section 2 we outline the class of bioproduction models we study, summarise the optimal control framework, and derive the homogeneous moment closure approximation. The following four sections, sections 3 to 6, contain example systems of interest. In each, a heterogeneous model is presented, followed by a homogeneous approximation, and then numerical results comparing the optimal control of both models. A diverse choice of models allows us to observe different model features and how their constitutive ingredients interact with view to bioproduction control.

2 Heterogeneous bioproduction models and their optimal control and approximation

2.1 The influence of heterogeneity in bioproduction

It is instructive to motivate the study of heterogeneous models and their homogeneous counterparts using a sketch of a simple, idealised scenario. We consider a protein of interest that is present in different levels across a population, as typically arises due to stochastic gene expression. There is some selection process occurring within the population, say growth, that is modulated by the protein via a step function, (fig. 1a). This may be caused

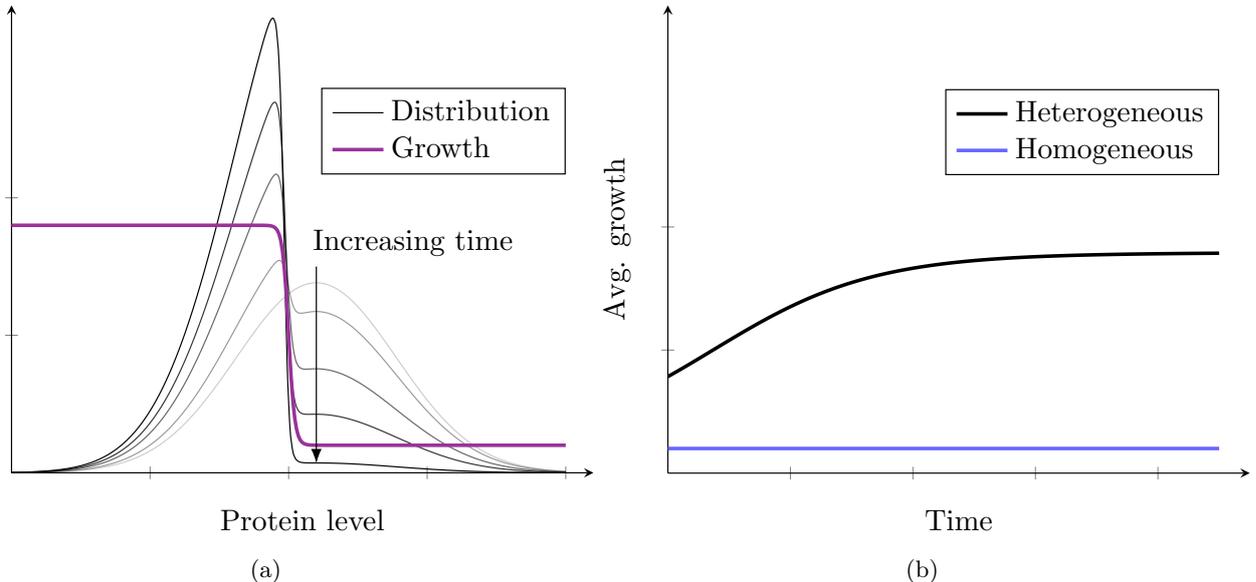


Figure 1: (a) Sketch of how the influence of protein-modulated growth on the distribution of protein levels can lead to selection. (b) A homogeneous model does not capture the time-varying growth as it does not account for the selection.

by the protein being toxic or burdensome to the cell [9, 20, 33, 51]. We then consider how the average selection evolves with time (fig. 1b) and ask: how does a homogeneous model compare with its heterogeneous counterpart?

In fig. 1, we assume that stochastic gene expression gave rise to an initial population distribution in the absence of selection. We then induce a protein-dependent growth process that is assumed, for the purposes of demonstration, to occur on a timescale much faster than the gene expression. Homogeneous models effectively represent the full underlying distribution exclusively by its mean with no account of variability. Therefore, the homogeneous perspective sees no change to the mean, retaining a low and time-invariant average growth rate. On the other hand, in the fully heterogeneous account, the state-dependent growth selects for cells of lower protein levels thereby skewing the distribution in the heterogeneous model and leading to a higher average growth rate (but lower average protein levels). This raises the question: do the different dynamics captured by the models manifest in vastly different optimal controls that maximise bioproduction? We will demonstrate that, in certain circumstances, homogeneous approximations provide sufficiently good surrogate models. In other circumstances, the discrepancy between the full transient behaviour and the approximations is definitive in describing the optimal control, and the heterogeneity must be accounted for.

2.2 Modular model class

Here we sketch the structure of a general class of models by delineating different constitutive components of the bioproduction process. We then provide concrete instances of each component under different modeling assumptions. This approach enables us to construct models in a modular fashion by plugging in the modeling description relevant for different bioproduction configurations.

We begin by considering a collection of *continuum* quantities within each cell, such as levels of a protein of interest or a transcription factor. We denote the j th quantity x_j and the vector of the entire collection $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}_+^d$. We then consider distinct *discrete* types of cells, with cells of a given type being indistinguishable. Our interest is in the dynamics of populations of large numbers of cells, as is typically the case in bioproduction applications. Importantly, heterogeneity at the single-cell level depends on the stochastic chemical kinetics, not the population size, and therefore persists in large populations [27]. The primary object we study is the (expected) population density of each class over the continuum quantities. This description allows us to capture the population mass in any given state, and thus allows a description that governs the evolution based on the multiscale dynamics comprising the single-cell kinetics and the population-level processes. Ultimately, it also allows us to quantify the bioproduction, which is typically the quantity we seek to optimise.

We denote the population density of cells of class k over the continuum quantities \mathbf{x} at time t by $p_k(\mathbf{x}, t)$.

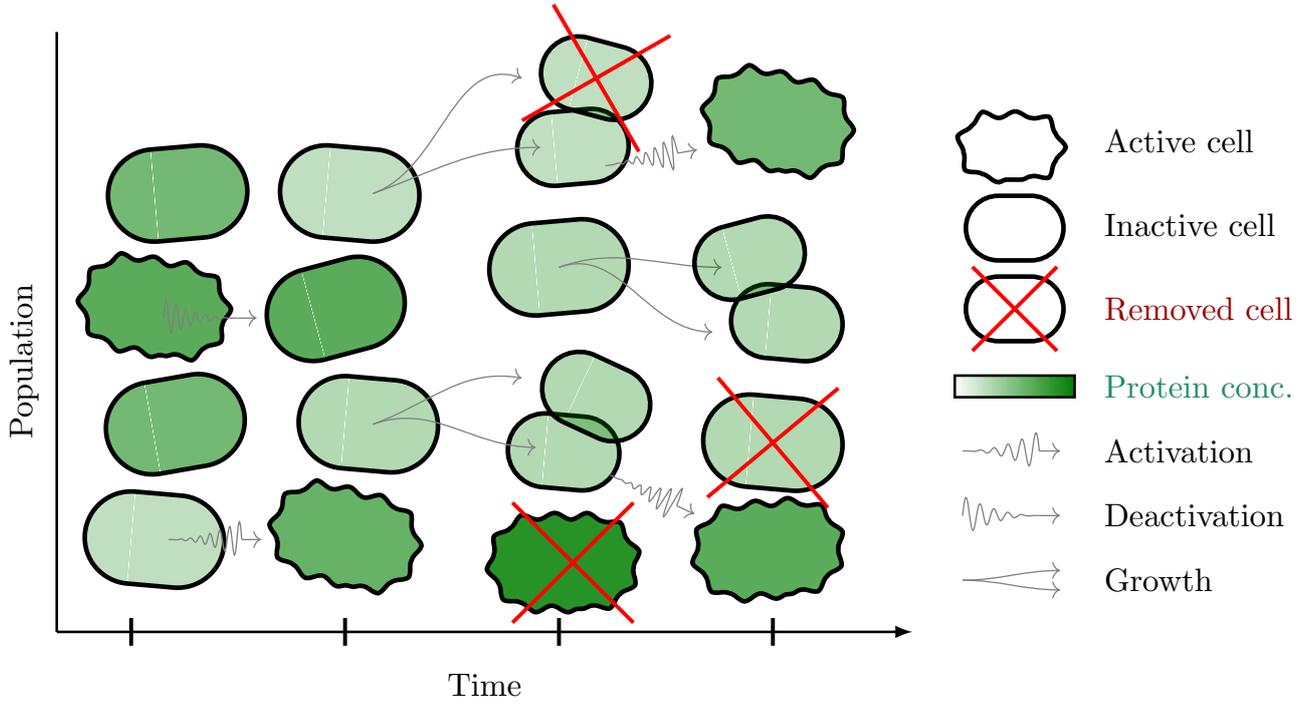


Figure 2: Schematic of the multiscale evolution of a population of cells. Each cell has an evolving internal configuration comprising a discrete state whether the cell is active or inactive (section 2.2.3), and a continuous level of a protein of interest (section 2.2.4). The population changes composition via selective state-dependent mechanisms such as growth as cells proliferate (here, cells with lower protein levels tend to have duplicated more, see section 2.2.1). In turbidostat mode, the total population mass is maintained by dilution (indicated by red crosses, see section 2.2.2).

Models governing the evolution of the collection of population densities in all states, $\{p_k(\mathbf{x}, t)\}_k$, take the form

$$\begin{aligned} \frac{\partial}{\partial t} p_k(\mathbf{x}, t) = & \text{Population growth} + \text{Bioreactor removal} \\ & + \text{Discrete dynamics} + \text{Continuum dynamics}, \end{aligned} \quad (1a)$$

$$p_k(\mathbf{x}, 0) = p_{0k}(\mathbf{x}), \quad (1b)$$

for $k = 1, \dots, \mathcal{K}$. Each term on the right-hand side of (1) contributes to the rate of change of the density of cells of class k with internal state \mathbf{x} , and thus each may depend on both k and \mathbf{x} , as well as time t . A schematic illustration is provided in fig. 2. We now detail a collection of examples for each component on the right-hand side of (1), from which we compose all of the case studies we subsequently explore.

2.2.1 Population growth

The population growth refers to the processes by which the cells proliferate. A simple and classical example is exponential growth. If the continuum quantities \mathbf{x} represent *concentrations* (that is, copy numbers divided by cell volume) that remain essentially constant over cell division events [34], then the rate of change of cell density due to growth may be governed by a law of the form

$$\frac{\partial}{\partial t} p_k(\mathbf{x}, t) = g_k(\mathbf{x}, t) p_k(\mathbf{x}, t) + (\text{Non-growth sources}), \quad (2)$$

where the growth rate $g_k(\mathbf{x}, t)$ exhibits dependence on the states k (e.g. different cell types grow at different rates) and \mathbf{x} (e.g. enhanced or inhibited cell division for different levels represented in \mathbf{x}) as well as time t (e.g. fluctuating ambient environmental conditions).

An alternative is to consider the continuum states \mathbf{x} to represent (scaled) copy numbers. In this case, cell division events cause the state \mathbf{x} to be partitioned and shared between the two daughter cells, which is described by a fragmentation process. We define the set of states in \mathbb{R}_+^d with each component no greater than each corresponding component of $\hat{\mathbf{x}}$, which we denote $\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}} := [0, \hat{x}_1] \times \dots \times [0, \hat{x}_d]$ where $\mathbf{0} = (0, \dots, 0) \in \mathbb{N}_0^d$ is the d -dimensional vector of zeros. Similarly $\hat{\mathbf{x}} \geq \mathbf{x}$ denotes the set where each component is no less than its

counterpart in \mathbf{x} . The rate of change of cell density is then governed by

$$\begin{aligned} \frac{\partial}{\partial t} p_k(\mathbf{x}, t) = & -g_k(\mathbf{x}, t)p_k(\mathbf{x}, t) + 2 \int_{\hat{\mathbf{x}} \geq \mathbf{x}} g_k(\hat{\mathbf{x}}, t)p_k(\hat{\mathbf{x}}, t)\theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\hat{\mathbf{x}} \\ & + (\text{Non-growth sources}), \end{aligned} \quad (3)$$

where the growth rate is still $g_k(\mathbf{x}, t)$ and $\theta_k(\mathbf{x} | \hat{\mathbf{x}})$ represents the probability density that, for a given daughter cell, a mother cell in state $\hat{\mathbf{x}}$ divides to produce a daughter in state \mathbf{x} . The factor of two reflects the cell division of the mother into two daughter cells. We require that θ_k be a symmetric probability density:

$$\int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} = 1, \quad \theta_k(\mathbf{x} | \hat{\mathbf{x}}) = \theta_k(\hat{\mathbf{x}} - \mathbf{x} | \hat{\mathbf{x}}). \quad (4)$$

As we will demonstrate in deriving the moment contributions of the terms in (3), assumptions (4) guarantee that the fragmentation contributes an increase in the total number of cells per unit time while conserving the total protein copy number.

2.2.2 Bioreactor removal

Bioreactors operate in a variety of modes. Batch operation describes a bioreactor where the initial colony is placed into the media and left to develop without breaking containment. Typically the contents are stirred to ensure a homogeneous solution. Batch operation requires no ‘‘Bioreactor removal’’ term in equation (1). While operation is straightforward, one major disadvantage of batch operation is that once the nutrients in the original media are exhausted the desired production slows to a halt. This then requires a work-intensive sterilisation between short-lived batches, and complicates the modeling since the culture transitions through several phenomenologically different phases (lag, exponential growth, stationary, death, long-term stationary [39]). Various modes of continuous operation allow fresh media to be introduced into the bioreactor, and the existing solution diluted to remove toxins and byproducts as well as maintain an equilibrium population density. This allows longer operation, reducing the burden of frequently resetting the bioreactor, and simplifies the modeling as it allows the colony and nutrients to remain in a single phase.

Turbidostat operation is a particularly useful mode, where the dilution is tuned in real time so as to maintain a constant optical density. Taking the optical density as a proxy for the population density, we may model this contribution via the dilution rate, $\Lambda(t)$,

$$\frac{\partial}{\partial t} p_k(\mathbf{x}, t) = -\Lambda(t)p_k(\mathbf{x}, t) + (\text{Non-bioreactor sources}), \quad (5)$$

where $\Lambda(t)$ does not depend on k or \mathbf{x} , but does depend on the population density as follows. Integrating (1) over the state space $\mathbf{x} \in \mathbb{R}_+^d$ and summing over all discrete states $k = 1, \dots, \mathcal{K}$, the discrete and continuum dynamics cancel as they are *conservative*; describing how members of the population change their states k and \mathbf{x} without changing the total population mass. In this case, where all changes in population mass stem from the ‘‘Population growth’’ terms, it follows that

$$\Lambda(t) = \frac{\sum_{k=1}^{\mathcal{K}} \int_{\mathbb{R}_+^d} [\text{Population growth}]_k(\mathbf{x}, t) d\mathbf{x}}{\sum_{k=1}^{\mathcal{K}} \int_{\mathbb{R}_+^d} p_k(\mathbf{x}, t) d\mathbf{x}}, \quad (6)$$

where we have shown the explicit dependence of the population growth contributions on the states k and \mathbf{x} and time t . The turbidostat dilution rate (6) is simply the mean growth rate, which is precisely the balance required to preserve a constant total population mass. Moreover, since the denominator in (6) is constant, we may assume, without loss of generality, that it is unity. This is equivalent to rescaling the population densities $p_k(\mathbf{x}, t)$ by this constant value. This gives a simpler form for the dilution rate, namely

$$\Lambda(t) = \sum_{k=1}^{\mathcal{K}} \int_{\mathbb{R}_+^d} [\text{Population growth}]_k(\mathbf{x}, t) d\mathbf{x}. \quad (7)$$

2.2.3 Discrete dynamics

Discrete dynamics refers to the change of state of a cell in class k to some other class ℓ , at a rate denoted by $c_{k\ell}(\mathbf{x}, t)$, whereby

$$\begin{aligned} \frac{\partial}{\partial t} p_k(\mathbf{x}, t) = & -p_k(\mathbf{x}, t) \sum_{\ell=1}^{\mathcal{K}} c_{k\ell}(\mathbf{x}, t) + \sum_{\ell=1}^{\mathcal{K}} c_{\ell k}(\mathbf{x}, t) p_\ell(\mathbf{x}, t) \\ & + (\text{Non-discrete sources}). \end{aligned} \quad (8)$$

In this work, we use the discrete states to model different cell classes, as we will detail. Nevertheless, the discrete states could, in principle, also describe the stochastic single-cell chemical kinetics by representing the copy number of a species produced in numbers below what would justify a continuum approximation [26, 27], or a binary internal state such as governed by promoter switching or binding kinetics. The copy-number paradigm described above pertains to continuum species. It may be straightforwardly extended to account for division of discrete copy numbers.

2.2.4 Continuum dynamics

Continuum dynamics refers to the change of the internal states of cells from, say, \mathbf{x} to \mathbf{y} . It is common, as is exclusively the case in this work, that these internal continuum states represent levels of different proteins, whose dynamics are governed by gene expression models. For example, the evolution of the continuum state \mathbf{x} under a birth–death process with rate $r_k(\mathbf{x}, t) \geq 0$ is described by the Fokker–Planck operator

$$\begin{aligned} \frac{\partial}{\partial t} p_k(\mathbf{x}, t) = & -\mathbf{e}_k \cdot \nabla [r_k(\mathbf{x}, t) p_k(\mathbf{x}, t)] + \frac{1}{2\Omega} \mathbf{e}_k^\top \nabla^2 [r_k(\mathbf{x}, t) p_k(\mathbf{x}, t)] \mathbf{e}_k \\ & + (\text{other sources}), \end{aligned} \quad (9)$$

where $\mathbf{e}_k \in \mathbb{Z}^d$ is a vector representing the number of reactant molecules minus the number of product molecules of each continuum species [27], ∇^2 denotes the Hessian, and $\Omega \gg 1$ is the system size parameter, representing the characteristically large copy numbers that justify a continuum formulation [26, 15, 14]. The operator is associated with the zero normal-flux boundary condition

$$(\mathbf{e}_k \cdot \mathbf{n}) \left\{ r_k(\mathbf{x}, t) p_k(\mathbf{x}, t) - \frac{1}{2\Omega} \mathbf{e}_k \cdot \nabla [r_k(\mathbf{x}, t) p_k(\mathbf{x}, t)] \right\} = 0, \quad \mathbf{x} \in \partial \mathbb{R}_+^d, \quad (10)$$

where \mathbf{n} denotes the outward-facing normal to the boundary. There may be many such processes occurring simultaneously on the state k , in which case we sum over all the associated Fokker–Planck operators and boundary conditions.

In many biologically relevant scenarios, production in bursts (where multiple gene expression events occur in rapid succession) provides a more realistic description of the underlying production mechanism [12, 24], whereby the continuum dynamics take the form

$$\begin{aligned} \frac{\partial}{\partial t} p_k(\mathbf{x}, t) = & -f_k(\mathbf{x}, t) p_k(\mathbf{x}, t) \\ & + \|\mathbf{e}_k\| \int_{\mathbf{x} - z\mathbf{e}_k \in \mathbb{R}_+^d} f_k(\mathbf{x} - z\mathbf{e}_k, t) p_k(\mathbf{x} - z\mathbf{e}_k, t) Q_k(\mathbf{x} - z\mathbf{e}_k, z \|\mathbf{e}_k\|) dz \\ & + (\text{other sources}), \end{aligned} \quad (11)$$

where $\mathbf{e}_k \in \mathbb{Z}^d$ again represents the number of molecules of each continuum species in the reactants and products of a single gene production event, $f_k(\mathbf{x}, t)$ denotes the rate of production bursts, and $Q_k(\mathbf{x}, y)$ denotes the probability density, characteristically of exponential form, that the bursty production process in a cell of state \mathbf{x} yields a cell state $\mathbf{x} + y\mathbf{e}_k / \|\mathbf{e}_k\|$. We thus require that, for all \mathbf{x} ,

$$\int_0^\infty Q_k(\mathbf{x}, y) dy = 1. \quad (12)$$

Here too, if several such processes occur simultaneously these are to be summed.

There are no boundary conditions to be imposed with the burst operator, however, the formulation (11) implicitly imposes boundary conditions. The present formulation assumes that all components of the vector \mathbf{e}_k are non-negative. Were this not the case, the first term on the right-hand side of (11) would need to be modulated by a coefficient representing the fraction of valid jumps: $v_k(\mathbf{x}) := \|\mathbf{e}_k\| \int_{\mathbf{x} + z\mathbf{e}_k \in \mathbb{R}_+^d} Q_k(z \|\mathbf{e}_k\|) dz$. This

would disallow jumps that use “exhausted” continuum species (that is, jumps to states \mathbf{x} with at least one negative component) and instead leave the state unchanged. This treatment can be handled in the present formulation by simply absorbing the coefficient $v_k(\mathbf{x})$ into the rate $f_k(\mathbf{x}, t)$. A more physically relevant model would place the fraction $1 - v_k(\mathbf{x})$ of invalid jumps at the boundary, however, for the sake of simplicity, we will not account for this.

We have presented different example components for each of the terms in equation (1). This allows us to consider the broad class of PDEs where these components are chosen independently and pieced together to model a range of bioproduction systems, as we subsequently demonstrate in the examples. Given this class of systems, we proceed now to discuss the optimal control of the PDE models.

2.2.5 Optimal control of the PDE

Our primary focus is on optimally controlling bioproduction models of the form (1), which requires two additional ingredients. First, we must specify how the rates in each contributing term depend on the control input, which we denote $u(t)$ throughout. Second, we must specify some objective function J_{PDE} which we seek to optimise. We rewrite the dynamics (1) in vector form as

$$\frac{\partial}{\partial t} \mathbf{p}(\mathbf{x}, t) = \mathbf{M}(\mathbf{p}(\cdot, t), t, u(t)) \mathbf{p}(\mathbf{x}, t), \quad \mathbf{p}(\mathbf{x}, 0) = \mathbf{p}_0(\mathbf{x}), \quad (13a)$$

where the operator of the right-hand side \mathbf{M} now exhibits explicit (non-local) dependence on the system state $\mathbf{p}(\cdot, t)$ and the control $u(t)$. We consider optimising objectives of the form

$$J_{\text{PDE}} = \int_{\mathbb{R}_+^d} \psi(\mathbf{x}, \mathbf{p}(\mathbf{x}, T)) \, d\mathbf{x} + \int_0^T \int_{\mathbb{R}_+^d} \varphi(\mathbf{x}, t, u(t), \mathbf{p}(\cdot, t)) \, d\mathbf{x} \, dt, \quad (13b)$$

where ψ and φ map to \mathbb{R} . Typically, as is the case in this study, the objective measures the amount of bioproduct. The control problem (13) is to maximise the objective (13b) subject to the dynamics (13a) and a bounded control, say, $u(t) \in [0, 1]$.

How do we optimally control this family of PDE models? While full detail is provided in Appendix A, it is instructive to outline the structure of the approach. The principal aim is to calculate the “derivative” of the objective with respect to the control. If we can accomplish this, we have both an analytical characterisation of the optimal control — where the derivative vanishes — as well as a numerical recipe for computing it — given a control, we may use the derivative to iteratively update the control in a direction which better optimises the objective. In practice, we use L-BFGS-B, a limited-memory, control-bounded implementation of the BFGS routine [31], which uses the difference between iterates to approximate second-order information (the Hessian) and thereby precondition the derivative to provide a curvature-informed search direction. A stable and consistent discretisation of the PDE model is detailed in Appendix A.

So how do we calculate the derivative? We take the adjoint approach: in brief, we construct the Lagrangian by adjoining the system dynamics to the objective function:

$$\begin{aligned} \mathcal{L} := J_{\text{PDE}} - \int_0^T \int_{\mathbb{R}_+^d} \boldsymbol{\lambda}(\mathbf{x}, t) \left[\frac{\partial}{\partial t} \mathbf{p}(\mathbf{x}, t) - \mathbf{M}(\mathbf{p}(\cdot, t), t, u(t)) \mathbf{p}(\mathbf{x}, t) \right] \, d\mathbf{x} \, dt \\ - \int_{\mathbb{R}_+^d} \boldsymbol{\lambda}(\mathbf{x}, 0) [\mathbf{p}(\mathbf{x}, 0) - \mathbf{p}_0(\mathbf{x})] \, d\mathbf{x}. \end{aligned} \quad (14)$$

Equating the derivative of the Lagrangian \mathcal{L} with respect to the system state $\mathbf{p}(\mathbf{x}, t)$ to zero gives a terminal condition and a backward (in time) equation governing the costate $\boldsymbol{\lambda}(\mathbf{x}, t)$. Then, the derivative of the objective with respect to the control — the quantity we sought to compute — is found to be an explicit function of the state and the costate. This recipe may be followed for the discretised formulation identically to yield a discretised costate equation and discretised gradient of the objective with respect to the discretised control.

2.3 Moment approximations of model components

The analysis of the PDE models outlined in section 2.2, as well as their optimal control, is formidably challenging. The computational cost of their numerical simulation and optimisation quickly becomes prohibitive with increasing resolution and dimension, motivating dimension reduction techniques. Many order-reduction (and discretisation) techniques, such as spectral methods, Galerkin methods, projection and decomposition methods, are based on expressing the solution $\mathbf{p}(\mathbf{x}, t)$ as a combination of basis functions approximately spanning the state space. In this work, bearing in mind that our primary goal is to determine the optimal control but not necessarily keep track of the full distribution $\mathbf{p}(\mathbf{x}, t)$, we consider dimension reduction by tracking only the moments of the distribution, thereby removing the independent state variable \mathbf{x} altogether.

In the general case, where the rate and burst functions introduced in section 2.2 are state-dependent, the moments of the distribution are governed by an infinite hierarchy of coupled equations, which require a closure approximation. We focus on low-order moment closure approximations, up to and including first order. Zeroth-order approximations essentially assume that the underlying heterogeneous quantities have a stationary distribution. First-order approximations preserve the time-varying nature of the heterogeneous quantities, however, they do not preserve the cell-to-cell variability. Thus, the comparison between the heterogeneous PDE models and their homogeneous ODE approximations provide great insight into the influence of the cell-to-cell variability on the control. Moreover, the resulting homogeneous ODE models are of the type typically found in the bioproduction literature, making any deviation in performance particularly pertinent.

Taking the components introduced in section 2.2, we derive the associated low-order contributions to the moment closure approximations. We introduce the moments

$$X_{j,k}^i(t) := \int_{\mathbb{R}_+^d} x_j^i p_k(\mathbf{x}, t) d\mathbf{x}. \quad (15)$$

When $i = 0$, the index $j \in \{1, \dots, d\}$ is of no consequence and we write $X_{.,k}^0(t)$. Since we are only interested in moments up to first order, we only multiply by a single variable x_j , but no mixed moments.

The governing equation of each moment $X_{j,k}^i(t)$ is found by multiplying the PDE governing $p_k(\mathbf{x}, t)$ by x_j^i and integrating over the state space. We close the moment hierarchy at first order by extracting all state-dependent rate functions from under the integral and evaluating them at the associated (approximate) mean, that is, for an arbitrary function f ,

$$\int_{\mathbb{R}_+^d} x_j^i f(\mathbf{x}, t) p_k(\mathbf{x}, t) d\mathbf{x} \approx f(\mathbf{X}_k(t), t) X_{j,k}^i(t), \quad (16a)$$

where

$$\mathbf{X}_k(t) := \left(\frac{X_{1,k}^1(t)}{X_{.,k}^0(t)}, \dots, \frac{X_{d,k}^1(t)}{X_{.,k}^0(t)} \right). \quad (16b)$$

This approximation is justified when the distribution $p_k(\mathbf{x}, t)$ is well approximated by its mean [21] provided the function f is sufficiently slowly varying with respect to the state. Naturally, this condition is sufficient but not necessary: there could be scenarios where the distributions possess significant heterogeneity and, nevertheless, this approximation faithfully captures the system behaviour and control. Returning to the earlier comment regarding techniques that represent the solution in a basis, we identify the closure (16) as being equivalent to the representation $p_k(\mathbf{x}, t) = X_{.,k}^0(t) \delta(\mathbf{x} - \mathbf{X}_k(t))$ for the d -dimensional Dirac mass δ .

We now use the moment closure (16) on each aforementioned PDE component. This generates the associated ODE forms, which may analogously be used to compose homogeneous ODE approximations of the underlying heterogeneous PDE model. Details of the calculations appear in Appendix B.

2.3.1 Population growth

The growth law in the concentration paradigm (2) contributes to the moments via

$$\frac{d}{dt} X_{j,k}^i(t) \approx g_k(\mathbf{X}_k(t), t) X_{j,k}^i(t) + (\text{Non-growth sources}). \quad (17)$$

In the copy-number paradigm (3), we find (see Appendix B) that

$$\frac{d}{dt} X_{j,k}^0(t) = g_k(\mathbf{X}_k(t), t) X_{.,k}^0(t) + (\text{Non-growth sources}), \quad (18)$$

$$\frac{d}{dt} X_{j,k}^1(t) = (\text{Non-growth sources}). \quad (19)$$

In other words, the copy-number growth increases the population mass, the zeroth moment, by the average growth rate but does not affect the first moment.

2.3.2 Bioreactor removal

The batch mode of operation has no contribution to the dynamics. The turbidostat contribution (5) is approximated by

$$\frac{d}{dt} X_{j,k}^i(t) = -\Lambda(t) X_{j,k}^i(t) + (\text{Non-bioreactor sources}), \quad (20a)$$

where the turbidostat dilution rate $\Lambda(t)$ given in (7) may be expressed self-consistently as the sum of the growth contributions of zeroth order. In both the concentration (17) and copy-number (18) paradigms the zeroth-order contributions are $g_k(\mathbf{X}_k(t), t)X_{\cdot, k}^0(t)$, therefore, in both cases,

$$\Lambda(t) = \sum_{k=1}^{\mathcal{K}} g_k(\mathbf{X}_k(t), t)X_{\cdot, k}^0(t). \quad (20b)$$

It follows that, in turbidostat mode, the total population mass $\sum_{k=1}^{\mathcal{K}} X_{\cdot, k}^0(t)$ remains constant in the moment closure approximation.

2.3.3 Discrete dynamics

The discrete dynamics (8) lend the contributions

$$\begin{aligned} \frac{d}{dt} X_{j, k}^i(t) &= -X_{j, k}^i(t) \sum_{\ell=1}^{\mathcal{K}} c_{k\ell}(\mathbf{X}_k(t), t) + \sum_{\ell=1}^{\mathcal{K}} c_{\ell k}(\mathbf{X}_\ell(t), t) X_{j, \ell}^i(t) \\ &+ (\text{Non-discrete sources}). \end{aligned} \quad (21)$$

Here too the conservation property of the discrete dynamics is preserved: summing (21) over k we find that $\sum_{k=1}^{\mathcal{K}} X_{j, k}^i(t)$ is changed only by non-discrete sources.

2.3.4 Continuum dynamics

It is shown in Appendix B, by neglecting the second-order terms, that the Fokker–Planck contributions (9) give rise to

$$\frac{d}{dt} X_{\cdot, k}^0(t) = (\text{other sources}), \quad (22)$$

$$\frac{d}{dt} X_{j, k}^1(t) \approx -e_{k, j} r_k(\mathbf{X}_k(t), t) X_{\cdot, k}^0(t) + (\text{other sources}), \quad (23)$$

where $e_{k, j}$ is the j th component of the vector \mathbf{e}_k . The moment structure (22) and (23) preserves the conservation property that the subpopulation mass of index k , $X_{\cdot, k}^0(t)$, is unchanged by the Fokker–Planck contributions.

For the bursty production (11), we find that

$$\frac{d}{dt} X_{\cdot, k}^0(t) \approx (\text{other sources}), \quad (24)$$

$$\frac{d}{dt} X_{j, k}^1(t) \approx \frac{e_{k, j}}{\|\mathbf{e}_k\|} f_k(\mathbf{X}_k(t), t) Q_{k, 1}(\mathbf{X}_k(t)) X_{\cdot, k}^0(t) + (\text{other sources}), \quad (25)$$

where we have defined

$$Q_{k, 1}(\mathbf{x}) := \int_0^\infty y Q_k(\mathbf{x}, y) dy, \quad (26)$$

also guaranteeing the PDE feature that $X_{\cdot, k}^0(t)$ is unchanged by these continuum sources.

For example, the typical jump kernel is an exponential of the form

$$Q_k(\mathbf{x}, y) = \frac{e^{-y/q_k(\mathbf{x})}}{q_k(\mathbf{x})}, \quad (27)$$

for a (possibly state-dependent) mean burst size q . This satisfies (12), and gives $Q_{k, 1}(\mathbf{x}) = q_k(\mathbf{x})$.

2.3.5 Initial conditions

The moments inherit their initial conditions from the moments of the initial density (1b), that is,

$$X_{j, k}^i(0) = \int_{\mathbb{R}_+^d} x_j^i p_{0k}(\mathbf{x}) d\mathbf{x}. \quad (28)$$

This concludes the presentation of the homogeneous ODE analogues of the underlying heterogeneous PDE components. These may be similarly pieced together to model a wide range of bioproduction systems, as will be demonstrated in the examples of the subsequent sections. We finish with a brief overview of the numerical and analytical strategies we apply in optimally controlling these composite ODE models.

2.3.6 Optimal control of the ODE

Given a reduced-order ODE system, we apply Pontryagin’s maximum principle to study the optimal control analytically. For the numerical computation of the optimal control, we discretise the ODE and objective function via a backward Euler scheme, and implement a multiple-shooting direct numerical approach by encoding the nonlinear program in CasADi [4], a framework that performs automatic differentiation to equip the IPOPT solver [49] with gradient information. To make the comparison between the PDE and ODE reasonable, we use the same time discretisation: equispaced points with the time step small enough to guarantee that the PDE scheme is monotone (see Appendix A.3).

Beginning with a generic PDE model (1), we have delineated various forms for the different components, each of which may be plugged in to compose the optimal control problem (13). For each component form, we have derived the associated first-order moment approximation, which provides a homogeneous reduced-order approximation. In the following four sections, we consider four specific bioproduction models of different composition and study the optimal control of the heterogeneous PDE and homogeneous ODE approximation.

3 Bioproduction by direct induction

3.1 PDE formulation

3.1.1 Concentration paradigm

The first scenario we seek to model is of an isogenic community of photoreceptive cells (see e.g. Refs. [8, 22, 30, 40, 47, 52]). Shining light on the culture induces the production of a protein of interest, whose *concentration* in a cell we denote by x . We assume that the protein production is described by a birth–death process, which, via stochastic gene expression, gives rise to cell-to-cell variability in its concentration. We also assume that the protein of interest is disadvantageous to the cell such that a higher concentration results in a slower growth rate $g(x)$. Finally, we consider the cell colony to be hosted in a nutrient-rich bioreactor running in turbidostat mode.

The population density with protein of interest at level x at time t is denoted by $p(x, t)$. We consider the governing equation

$$\begin{aligned} \frac{\partial p}{\partial t}(x, t) &= g(x)p(x, t) - \Lambda(t)p(x, t) \\ &\quad - \frac{\partial}{\partial x}[(\alpha u(t) - \beta x - g(x)x)p(x, t)] \\ &\quad + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2}[(\alpha u(t) + \beta x + g(x)x)p(x, t)], \end{aligned} \quad (29a)$$

$$p(x, 0) = p_0(x), \quad (29b)$$

subject to the zero normal-flux boundary conditions

$$(\alpha u(t) - \beta x - g(x)x)p(x, t) - \frac{1}{2\Omega} \frac{\partial}{\partial x}[(\alpha u(t) + \beta x + g(x)x)p(x, t)] = 0, \quad (29c)$$

at $x = 0$ and as $x \rightarrow \infty$.

The first- and second-order differential terms in (29a) are the Fokker–Planck contributions capturing the stochastic expression of the protein of interest as well as cell growth. These terms stem from a birth–death process with a maximal birth rate of α , modulated by the (dimensionless) light signal $u(t) \in [0, 1]$, and a linear death rate with coefficient $\beta + g(x)$.

To understand the birth and death rates in the concentration paradigm, we write the number of protein molecules (the copy number) as $\#_{\text{protein}}$. A birth event represents the production of the protein of interest, which happens at a constant rate (for a given illumination) given a fixed amount of genetic material to transcribe: over an infinitesimal time step Δt we write $\#_{\text{protein}} \mapsto \#_{\text{protein}} + \alpha u(t)\Delta t$. A death event due to protein degradation can occur to any of the existing protein molecules, and therefore occurs at a linear rate: $\#_{\text{protein}} \mapsto [1 - \beta\Delta t]\#_{\text{protein}}$. Since cell growth is well described by exponential growth [41, 19, 10], we expect a birth event of cellular material to be happening at a constant rate and to be proportional to the existing quantity of material. Denoting the cellular material $\#_{\text{cell}}$, we can write that the cellular materials changes via $\#_{\text{cell}} \mapsto [1 + g(x)\Delta t]\#_{\text{cell}}$ where $x = \#_{\text{protein}}/\#_{\text{cell}}$ is the concentration of proteins that determines the cellular growth rate. The change in concentration brought about by this cellular birth event is equivalent (to leading order in Δt) to the change in protein: $\#_{\text{protein}} \mapsto (1 + g(x)\Delta t)\#_{\text{protein}}$. This explains why the concentration dilution due to cell growth takes the same form as the protein degradation term.

There is a structural difference between the birth and death events described above: the birth event results in an additive change to the protein copy number, while the death event and cell growth result in a multiplicative

change. The latter is preserved when we consider the concentration, however, the former is not. This means that, strictly speaking, the birth event induces a change in concentration that depends on knowledge of either $\#_{\text{protein}}$ or $\#_{\text{cell}}$, since $x \mapsto x + \alpha u(t)\Delta t/\#_{\text{cell}}$. To capture this properly would require us to reinstate either of these two quantities. This modelling and computational costs are heavy, and instead, we assume that a constant rate is an adequate approximation. The multiscale interaction between scales is manifest: the single-cell dynamics are influenced by the population-level cell growth, while the cell growth depends on the concentration dictated at the single-cell level. Thus, phenomenologically we have captured the underlying processes and their mutual interaction even if the coefficient α ought to depend on the internal state more intricately.

The population growth rate $g(x)$ reflects the fact that the protein of interest is toxic to the cell, preferentially selecting for lower levels of x via

$$g(x) = g_\infty + \frac{g_0 - g_\infty}{1 + (x/K)^n}. \quad (29d)$$

In other words, we assume that $g_0 > g_\infty \geq 0$ and $n > 0$.

We make a brief technical remark regarding the linear dependence on the light intensity input $u(t)$. The linear form comes with no loss of generality: were the light intensity to have a nonlinear response on the birth rate, say $W(u(t))$ for some (normalised) nonlinear $W : [0, 1] \rightarrow [0, 1]$ that is onto $[0, 1]$, we could replace any light profile $u(t)$ in the current model by $W^{-1}(u(t))$ in the nonlinear model and achieve equivalent dynamics. Here, W^{-1} denotes any right inverse of W . In particular, if W is invertible then this is the only such mapping. With this in mind, it is more generic to refer to the light control $u(t)$ in the linear model as the light influence, describing its effect (which might depend nonlinearly on the light input) rather than the physical input, but we will not be strict with this phrasing.

The dilution rate $\Lambda(t)$ of the turbidostat mode, given in (7), takes the form

$$\Lambda(t) = \int_0^\infty g(x)p(x, t) dx =: \langle g, p(t) \rangle, \quad (29e)$$

where we have adopted the inner-product notation.

We seek to maximise the bioproduct, which is the amount of protein in the bioreactor at the terminal time $t = T$ combined with the amount we have extracted from the system via dilution throughout the process, namely

$$J_{\text{PDE}} = \langle x, p(T) \rangle + \int_0^T \Lambda(t) \langle x, p(t) \rangle dt. \quad (29f)$$

Our strategy to solve the optimal control problem (29) is grounded in the adjoint approach [13] detailed in Appendix A. We briefly note that some analytical progress may be made in the optimal control of the PDE. The terminal condition of the costate (97) takes the form $\lambda(x, T) = x$. The derivative of the pay-off with respect to the control (99) (in the direction of a Dirac mass at time t) takes the L^2 inner-product form $\delta_u J_{\text{PDE}} = \alpha \langle \frac{\partial}{\partial x} \lambda(x, t) + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2} \lambda(x, t), p(x, t) \rangle$. In the vicinity of the terminal time, we expect, by continuity, the costate to approximate its terminal condition, whereby $\delta_u J_{\text{PDE}} \approx \alpha \langle 1, p \rangle = \alpha > 0$. We deduce that, in a vicinity of the terminal time, the objective can be increased by a larger control. Therefore, in this neighbourhood, the terminal pay-off dominates the control, which is thus maximal $u(t) = 1$.

We now proceed to consider analogous dynamics for the copy-number paradigm.

3.1.2 Copy-number paradigm

The copy-number paradigm retains a similar structure. The independent variable x no longer represents a protein concentration, but rather the (scaled) copy number of the protein of interest. The population density $p(x, t)$ of copy number x at time t is governed by

$$\begin{aligned} \frac{\partial p}{\partial t}(x, t) &= -g(x)p(x, t) + 2 \int_x^\infty g(z)p(z, t)\theta(x | z) dz \\ &\quad - \Lambda(t)p(x, t) \\ &\quad - \frac{\partial}{\partial x} [(\alpha u(t) - \beta x)p(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2} [(\alpha u(t) + \beta x)p(x, t)], \end{aligned} \quad (30)$$

subject to the same initial condition (29b) and boundary conditions (29c). The quantity $\theta(x | z)$ represents the probability density that a cell of copy number z divides to produce, for a given daughter cell, a copy number x . In this copy-number paradigm, the degradation proportional to β represents exclusively degradation of the protein molecules. The turbidostat mode maintains a dilution rate $\Lambda(t)$ of the same form (29e).

Here too we seek to maximise the bioproduct as given in (29f). The argument in section 3.1.1 that the terminal pay-off induces a final arc of maximal optimal control is preserved in this paradigm.

3.2 A moment closure approximation

Using the components presented in section 2.3, we seek to derive ODE models approximating the low-order moments of the full distribution $p(x, t)$ governed by the models (29) and (30). Since we are in dimension $d = 1$ with $\mathcal{K} = 1$ state, we ignore the j and k subscripts, which can only take the values $j = k = 1$, by defining the shorthands

$$X^i(t) := X_{1,1}^i(t). \quad (31)$$

For the turbidostat operation, in which the total population mass is kept constant, the zeroth-order description is obtained by summing (17), (20) and (22) to yield $\dot{X}^0(t) = g(X^0(t))(1 - X^0(t))$ and $X^0(0) = 1$ (as discussed in section 2.2.2). It follows that $X^0(t) = 1$ for all t .

For the concentration model (29), combining the contributions of growth, turbidostat mode, and continuum dynamics (there are no discrete states in this example) via (17), (20) and (23), respectively, we find that

$$\dot{X}^1(t) = \alpha u(t) - [\beta + g(X^1(t))]X^1(t), \quad X^1(0) = \int_0^\infty x p_0(x) dx =: X_0^1, \quad (32)$$

where the dot denotes the time derivative.

For the copy-number model (30), the zeroth-order dynamics remain identical. At first order, using (20) and (23) as above but (19) instead of (17), we find that the ODE approximation is also given by (32). That is, the homogeneous approximation is not able to distinguish between the concentration and copy-number paradigms.

In both concentration and copy-number formulations, the objective (29f) takes the moment form $X^1(T) + \int_0^T \Lambda(t)X^1(t) dt$, where the dilution rate (29e) is approximated via (20b), which takes the form $\Lambda(t) = g(X^1(t))$.

Ultimately, the ODE version of the problem is to maximise

$$J_{\text{ODE}} = X^1(T) + \int_0^T g(X^1(t))X^1(t) dt, \quad (33a)$$

subject to the dynamics

$$\dot{X}^1(t) = \alpha u(t) - [\beta + g(X^1(t))]X^1(t), \quad X^1(0) = X_0^1. \quad (33b)$$

3.2.1 Optimal control of the ODE model

We next apply Pontryagin's principle [36] to the optimal control problem (33). Since the initial state is fixed and the final state is unconstrained, we may discard so-called abnormal extremals [48, Prop. 6.4.4]. We introduce the Hamiltonian

$$H = \lambda(t) \{ \alpha u(t) - [\beta + g(X^1(t))]X^1(t) \} - g(X^1(t))X^1(t), \quad (34)$$

where $\lambda(t)$ is the costate, governed by the costate equation and terminal condition

$$\dot{\lambda}(t) = \lambda(t)[\beta + g(X^1(t)) + g'(X^1(t))X^1(t)] + g'(X^1(t))X^1(t) + g(X^1(t)), \quad (35a)$$

$$\lambda(T) = -1, \quad (35b)$$

where the prime denotes differentiation with respect to the argument.

Pontryagin's maximum principle dictates that the optimal control minimises the Hamiltonian from among the admissible controls. Since the Hamiltonian is linear with respect to the control, it may be expressed as $H = H_0 + (\partial H / \partial u)u(t)$, where $\partial H / \partial u$ is called the switching function and is given by

$$\frac{\partial H}{\partial u} = \alpha \lambda(t). \quad (36)$$

It follows from Pontryagin's maximum principle that the optimal control is of bang–bang form when the switching function is not zero, that is, it takes either its minimum or maximum bound:

$$u(t) = \begin{cases} 1, & \lambda(t) < 0, \\ 0, & \lambda(t) > 0. \end{cases} \quad (37)$$

On intervals where the switching function vanishes identically, $\lambda(t) = 0$, called singular arcs, the maximum principle no longer provides the direct form of the optimal control. Instead, we may use the state and costate

equations to derive the optimal control, in combination with the fact that $\lambda(t)$ and its derivative vanish identically. From the costate equation, we know that, on a singular arc,

$$g(X^1(t)) + X^1(t)g'(X^1(t)) = 0, \quad (38)$$

for $g(x; g_0, g_\infty, K, n)$ defined in (29d). In Appendix C we study the equation (38) and summarise the results here. There can only be singular arcs if $n > 1$ and either (i) $g_\infty = 0$, or (ii) $g_\infty > 0$ and $4g_\infty/(g_0 - g_\infty) < (n - 1)^2/n$. If this condition is satisfied, we have explicit expressions for the unique state $X^1(t)$ satisfying the equation (38).

The condition (38) intuitively characterises the singular arc: it is nothing more than $X^1(t)$ obtaining a value such that the running pay-off $g(X^1(t))X^1(t)$ is extremal, that is, $X^1(t) = x$ where $d[g(x)x]/dx = 0$. If such a state may be reached then it is ideal to maintain it until some vicinity of the terminal time, where the terminal pay-off dominates the optimal control, as we will discuss.

Furthermore, for any constant $X^1(t)$ we can deduce the control by solving (33b) to give

$$u(t) = \frac{[\beta + g(X^1(t))]X^1(t)}{\alpha}. \quad (39)$$

Thus we have derived a closed-form expression for the control on singular arcs. In general, such calculations only guarantee the control as a function of the state and costate, however, in our case, the control depends only on the state, which is particularly useful for the practitioner.

Before we present numerical results, we conclude with an observation analogous to that made for the PDE model: the optimal control will be at its maximum value in some vicinity of the terminal time $t = T$. This follows immediately from the terminal condition (35b) and the bang-bang form of the optimal control (37).

3.3 Numerical results: comparing the PDE and ODE models

To gauge the performance of the ODE approximation (33), we first compute the optimal control for the ODE model $u_{\text{ODE}}(t)$ and the PDE model $u_{\text{PDE}}(t)$. Then, considering the PDE model to be the ground truth, we compute the objective function for the PDE model using each control, $J_{\text{PDE}}(u_{\text{ODE}}(t))$ and $J_{\text{PDE}}(u_{\text{PDE}}(t))$. Finally, we assess the relative performance by considering their ratio

$$\mathcal{J}_{\text{ODE}} := \frac{J_{\text{PDE}}(u_{\text{ODE}}(t))}{J_{\text{PDE}}(u_{\text{PDE}}(t))}. \quad (40)$$

We add a superscript ‘conc’ or ‘copy’ corresponding to the underlying PDE model paradigm.

In fig. 3 we display results for both the concentration and copy-number models, (29) and (30), respectively, alongside the ODE approximation (33) using the parameters

$$g_0 = 1, \quad g_\infty = 0, \quad n = 5, \quad \alpha = 0.5, \quad \beta = 0.7, \quad \Omega = 1000, \\ \hat{\theta}(x | y) = \frac{\delta(x/y - c) + \delta(x/y - 1 + c)}{2}, \quad c = 0.5, \quad p_0(x) = \delta(x), \quad (41)$$

where δ denotes a Dirac delta function, giving the ODE an initial condition $X_0^1 = 0$. In the copy-number model, this parametrisation describes cell division from a mother cell into two daughter cells of relative sizes c and $1 - c$, that is, neglecting stochasticity in daughter sizes (see Appendix B). The choice of $c = 0.5$ is symmetric division, however, the results presented here are preserved surprisingly closely for asymmetric division (see Appendix D). Moreover, we used 101 discrete nodes in the discretisation of the state space, with a time-step of $\Delta t \approx 3.78 \times 10^{-3}$, to guarantee numerical stability of the diffusion scheme (via a monotone scheme, as detailed in Appendix A.3).

We see from fig. 3 that the ODE provides a good control approximation. As to be expected, some of the transient timing present in the PDE control, especially transitioning between singular and nonsingular arcs, are not perfectly captured by the ODE optimal control. Nevertheless, the relative objective differs by less than 1%. Both concentration and copy-number PDE models admit qualitatively similar optimal controls.

With a concrete example of a concentration model (29) and its copy-number analogue (30), we may reflect on the similarities and distinctions between these two paradigms. Firstly, even though the state-dependent growth in these examples takes the same functional form, it represents something biologically distinct. Is the growth inhibited by a large fraction of proteins of interest, or a large absolute number of them? In practice, the answer may depend on the underlying biological mechanism of growth inhibition. One could, of course, retain information on both the cell size and copy numbers, thereby capturing both concentrations and copy-numbers in a single formulation. However, this comes at the significant cost of increased model dimensionality. We restrict our focus, therefore, to the concentration and copy-number reduced-order models.

Secondly, the concentration model captures cell growth, via an effective dilution reaction, whereas the copy-number model describes cell division events. The latter can be more descriptive, for example, it can account for

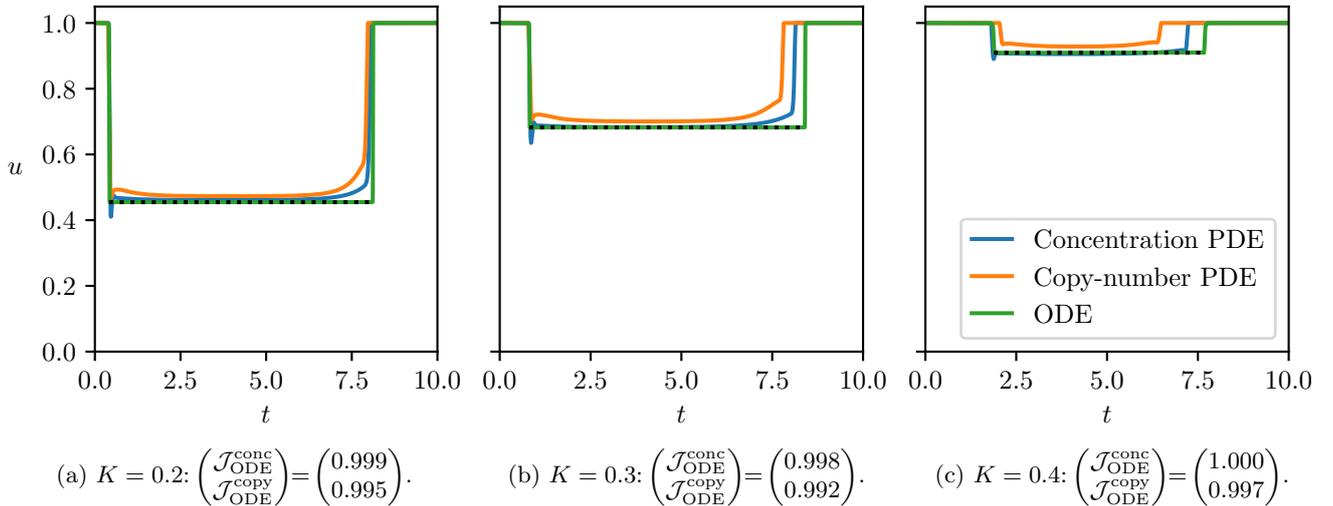


Figure 3: Optimal controls of the PDE problems (29) and (30) and the ODE problem (33) using parameters in (41). Black dotted lines show the singular arc (39).

asymmetric division, but it also requires non-local dynamics and additional parameters describing the cell division statistics. In the general case, the copy-number framework is capable of capturing stochastic division ratios. A binomial partition is one example of a common stochastic discrete division kernel [35], summing the number of molecules that ended up in one of the two daughter cells. In the continuum limit, this takes a Gaussian form, as the sample mean of independent random variables tends to the normal distribution in the limit of large system size (Central Limit Theorem).

Finally, the two models produce the same approximate zeroth- and first-order dynamics. One must proceed to higher order to find a difference in the moment approximations [29]. It follows that, as long as the first two moments provide a sufficiently good surrogate model, the discrepancy between these two model formulations may be neglected.

In the following section, we turn our attention to a more involved model with more prominent transience, and again study how the moment closure approximation performs.

4 Inducing a transcription factor that modulates production

4.1 PDE copy-number formulation

In this section, we propose a biologically richer variant of the model studied in section 3. Physically, the induction mechanism requires a photoreceptive transcription factor (TF), such as EL222 [32, 40]. Upon shining the light, the TF is activated, which then modulates the production of the protein of interest x . In contrast to the previous model, we now recognise that the level of activated TF varies between cells and fluctuates in time [2], motivating us to track the level of the activated TF, which we denote by a new independent variable y . We assume that the inactivated TF is expressed constitutively and present in non-limiting abundance, while the activation is governed by a birth–death process. In this model, the growth is inhibited by the presence of the TF, occurring at a rate $g(y)$, where the form of g from (29d) is retained but it is now evaluated at y . The production of the protein of interest is considered to be positively modulated by the TF at a rate $\alpha(y)$, for strictly increasing α of the Hill form

$$\alpha(y) = \alpha_\infty + \frac{\alpha_0 - \alpha_\infty}{1 + (y/L)^m}, \quad (42)$$

for constants $m > 0$ and $\alpha_\infty > \alpha_0 \geq 0$.

The population density of cells with (scaled) protein and TF copy-numbers x and y , respectively, at time t is

denoted $p(x, y, t)$, and is governed by

$$\begin{aligned} \frac{\partial p}{\partial t}(x, y, t) &= -g(y)p(x, y, t) + 2 \int_x^\infty \int_y^\infty g(\hat{y})p(\hat{x}, \hat{y}, t)\theta(x, y | \hat{x}, \hat{y}) d\hat{y} d\hat{x} \\ &\quad - \Lambda(t)p(x, y, t) \\ &\quad - \frac{\partial}{\partial x}[(\alpha(y) - \beta x)p(x, y, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2}[(\alpha(y) + \beta x)p(x, y, t)] \\ &\quad - \frac{\partial}{\partial y}[(au(t) - by)p(x, y, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial y^2}[(au(t) + by)p(x, y, t)], \end{aligned} \quad (43a)$$

$$p(x, y, 0) = p_0(x, y). \quad (43b)$$

subject to the zero normal-flux boundary conditions

$$(\alpha(y) - \beta x)p(x, y, t) - \frac{1}{2\Omega} \frac{\partial}{\partial x}[(\alpha(y) + \beta x)p(x, y, t)] = 0, \quad (43c)$$

at $x = 0$ and as $x \rightarrow \infty$, and

$$(au(t) - by)p(x, y, t) - \frac{1}{2\Omega} \frac{\partial}{\partial y}[(au(t) + by)p(x, y, t)] = 0, \quad (43d)$$

at $y = 0$ and as $y \rightarrow \infty$.

Here, $\theta(x, y | \hat{x}, \hat{y})$ denotes the probability density that a dividing mother cell in the state (\hat{x}, \hat{y}) produces, for a given daughter cell, a cell in the state (x, y) .

In turbidostat operation the dilution rate (7) becomes

$$\Lambda(t) = \int_0^\infty \int_0^\infty g(y)p(x, y, t) dx dy. \quad (43e)$$

We seek to maximise the protein production, given by

$$J_{\text{PDE}} = \int_0^\infty \int_0^\infty xp(x, y, T) dx dy + \int_0^T \Lambda(t) \int_0^\infty \int_0^\infty xp(x, y, t) dx dy dt. \quad (43f)$$

It is interesting to consider the argument from section 3.1.1 regarding the PDE control near the terminal time. In this case, the terminal costate condition is again $\lambda(x, y, T) = x$, therefore, $M^\top \lambda(x, y, t) \approx 0$ which gives $\delta_u \mathcal{L} \approx 0$ in the vicinity of $t = T$, and the sign is undetermined. Further analysis is required to confirm whether the control is necessarily maximal near the terminal time.

Before turning to the optimal control of problem (43), we derive a moment-closure approximation.

4.2 A moment closure approximation

The PDE formulation (43) is of dimension $d = 2$ with a single state $\mathcal{K} = 1$. To simplify the notation, we define the shorthands

$$X(t) := X_{1,1}^1(t), \quad Y(t) := X_{2,1}^1(t), \quad (44)$$

for which we have the ODE problem of maximising

$$J_{\text{ODE}} = X(T) + \int_0^T g(Y(t))X(t) dt, \quad (45a)$$

subject to the dynamics

$$\dot{X}(t) = \alpha(Y(t)) - [\beta + g(Y(t))]X(t), \quad X(0) = X_0, \quad (45b)$$

$$\dot{Y}(t) = au(t) - [b + g(Y(t))]Y(t), \quad Y(0) = Y_0, \quad (45c)$$

for initial conditions

$$X_0 := \int_0^\infty \int_0^\infty xp_0(x, y) dx dy, \quad Y_0 := \int_0^\infty \int_0^\infty yp_0(x, y) dx dy. \quad (46)$$

4.2.1 Optimal control of the ODE model

In this section, we briefly study the optimal control of the ODE problem (45). Again with the aim of leveraging Pontryagin's maximum principle, we introduce the Hamiltonian

$$H = \lambda_1(t) \{ \alpha(Y(t)) - [\beta + g(Y(t))]X(t) \} + \lambda_2(t) \{ au(t) - [b + g(Y(t))]Y(t) - g(Y(t))X(t) \}. \quad (47)$$

The costate equations and terminal conditions are given by

$$\dot{\lambda}_1(t) = \lambda_1(t)[\beta + g(Y(t))] + g(Y(t)), \quad (48a)$$

$$\lambda_1(T) = -1, \quad (48b)$$

$$\dot{\lambda}_2(t) = \lambda_1(t)[- \alpha'(Y(t)) + g'(Y(t))X(t)] + \lambda_2(t)[b + g(Y(t)) + g'(Y(t))Y(t)] + g'(Y(t))X(t), \quad (48c)$$

$$\lambda_2(T) = 0. \quad (48d)$$

The switching function is $\partial H / \partial u = a\lambda_2(t)$. Therefore, the optimal control is of bang–bang form:

$$u(t) = \begin{cases} 1, & \lambda_2(t) < 0, \\ 0, & \lambda_2(t) > 0. \end{cases} \quad (49)$$

Singular arcs are characterised by $\lambda_2(t) = \dot{\lambda}_2(t) = 0$ whence

$$\lambda_1(t) = \frac{g'(y(t))X(t)}{\alpha'(Y(t)) - g'(Y(t))X(t)}. \quad (50)$$

Differentiating again and solving for $u(t)$ gives the optimal control on singular arcs, namely,

$$u(t) = \frac{R(X(t), Y(t))}{D(X(t), Y(t))}, \quad (51a)$$

with R and D taking the forms

$$R(x, y) = xy [b + g(y)] [\alpha''(y)g'(y) - \alpha'(y)g''(y)] + \beta x g'(y) [xg'(y) - 2\alpha'(y)] + \alpha'(y) [\alpha(y)g'(y) - g(y)\alpha'(y)], \quad (51b)$$

$$D(x, y) = ax [\alpha''(y)g'(y) - \alpha'(y)g''(y)]. \quad (51c)$$

Conveniently, the form (51) is again a function of the state only (and not the costate).

The terminal condition $\lambda_2(T) = 0$ in (48) does not directly grant that the optimal control will be maximal in a vicinity of the terminal time. Nevertheless, we see from (48) that

$$\dot{\lambda}_2(T) = \alpha'(Y(T)). \quad (52)$$

Therefore, $\dot{\lambda}_2(T) = \alpha'(Y(T)) > 0$ for all $Y(T) > 0$ and thus, in a vicinity of the terminal time, $\lambda_2(t) < 0$ from which it follows, by (49), that the optimal control is maximal.

4.3 Numerical results

As in the one-dimensional case, we consider deterministic fragmentation ratio $c \in [0, 1]$, that is, when there is no stochasticity in the daughter copy-number fraction. In this case, $\hat{\theta}$ takes the form

$$\hat{\theta}(y | z) = \frac{\delta(y/z - c) + \delta(y/z - 1 + c)}{2}, \quad (53)$$

satisfying conditions (122) and (124) (see Appendix B).

4.3.1 Parametrised oscillatory controls

Before turning our attention to the control of problems (43) and (45), it is instructive to begin with preliminary remarks. Physically, our goal is to maximise protein production, which involves producing both more protein factories (that is, cells, via growth) and producing more protein in each factory. Application of the control $u(t)$ induces the production/activation of the TF, which, in turn, has two effects that influence our objective in opposite

ways: increased protein production and decreased growth. Therefore, to maximise overall protein production, it might be best to find some middle ground where these competing interests are optimally balanced, reminiscent of the controls in section 3. However, if the range of TF levels y for which $g(y)$ is non-negligible does not overlap the range for which $\alpha(y)$ is non-negligible, then all such compromises produce poor protein production. In this case, it may be optimal to perform an oscillatory induction: induce so as to raise TF levels and protein production, then stop induction so that growth may be resumed (without too large an effect on protein levels), and repeat. This reasoning finds grounds in several experimental studies [20, 52].

Unless stated otherwise, for all simulations in this section we use the parameters

$$\begin{aligned} g_0 &= 1, & g_\infty &= 0, & K &= 0.3, & n &= 10, \\ \alpha_0 &= 0, & \alpha_\infty &= 0.8, & L &= 0.4, & m &= 20, \\ a &= 3, & b &= 5, & \beta &= 1, & \Omega &= 1000, \\ p_0(x, y) &= \delta(x)\delta(y), \end{aligned} \tag{54}$$

with δ denoting a Dirac delta function, producing the initial conditions $X_0 = Y_0 = 0$.

We may explore the foregoing intuition on a parametrised family of controls $v(t)$. For reasons that will become clear, we set the control to its maximum value at the beginning and end of the interval $v(t) = 1$ for $t \in [0, t_0] \cup (T - t_0, T]$. We split the remaining interval $[t_0, T - t_0]$ into N identical pieces, with the control taking its maximum $v = 1$ during the middle $\phi \in [0, 1]$ fraction of the interval enclosed symmetrically by $v = 0$. This allows us to consider N pulses comprising a fraction ϕ of the interval (excluding the first and last pulses). This may be expressed via the formula

$$v(t; \phi, N, t_0) = \begin{cases} 1, & t_0 \notin [t_0, T - t_0], \\ V\left(\left[\frac{t-t_0}{T-2t_0} \bmod \frac{1}{N}\right] N\right) & t_0 \in [t_0, T - t_0], \end{cases} \tag{55a}$$

where,

$$V(z) = \begin{cases} 0, & z \notin \left[\frac{1-\phi}{2}, \frac{1+\phi}{2}\right], \\ 1, & z \in \left[\frac{1-\phi}{2}, \frac{1+\phi}{2}\right], \end{cases} \tag{55b}$$

and $a \bmod b := a - b \lfloor a/b \rfloor \in [0, b)$ gives the remainder after division of $a \geq 0$ by $b > 0$. Example members of this family of controls are depicted in fig. 4.

We emphasise that we are not arguing that the family of controls $v(t; \phi, N, t_0)$ is necessarily optimal, instead, it is a useful parametrised family to explore the model behaviour. Principally, the projection of the control onto this two-dimensional parameter space (we will keep t_0 fixed) will grant a deeper understanding of the objective function.

In fig. 5 we show simulation results for the dynamics and evaluate the objectives in problems (43) and (45) on the (N, ϕ) parameter space in both the PDE and ODE formulations. It is interesting to note that the ODE plots captures some, but not all, of the qualitative features of their PDE counterparts. For example, the peak J values are skewed to the right (beyond $\phi = 0.5$) in both the PDE and ODE formulations. On the other hand, the PDE model exhibits the global property that $J|_{\phi \approx 0} < J|_{\phi \approx 1}$, which the ODE model does not capture. Most strikingly, both PDE and ODE models exhibit non-monotonicity with respect to both the fraction of full intensity ϕ and the number of pulses N . This demonstrates that there is a nontrivial balance of timescales.

4.3.2 Optimal controls

With this intuition of oscillatory controls in mind, we proceed to compute optimal controls. Since the parametrised controls $v(t; \phi, N, 1.5)$ achieved similar objectives for similar values of ϕ but different values of N , despite these controls being markedly different from one another, we foresee that our optimisation routine may converge to local minima. To this end, we choose a family of parametrised initial control guesses $u_0(t; N)$, that, on the one hand, provide a hint of oscillation, but on the other hand, do not prescribe excessive structure to the control from the outset (having a maximal range of only $[1/4, 3/4]$):

$$u_0(t; 0) = \frac{1}{2}, \quad \text{and} \quad u_0(t; N) = \frac{1}{4} \cos\left(2\pi N \frac{t}{T}\right) + \frac{1}{2}, \quad \text{for} \quad N \geq 1. \tag{56}$$

As expected, different initial guesses lead to convergence towards different local optima (fig. 6). Moreover, many iterations were required to make significant changes in the control, suggesting that the objective landscape is particularly flat. This combination of challenging optimisation features motivates an additional step to refine the optimal controls. We see from fig. 6 that, for this choice of parameters, the optimal controls converge

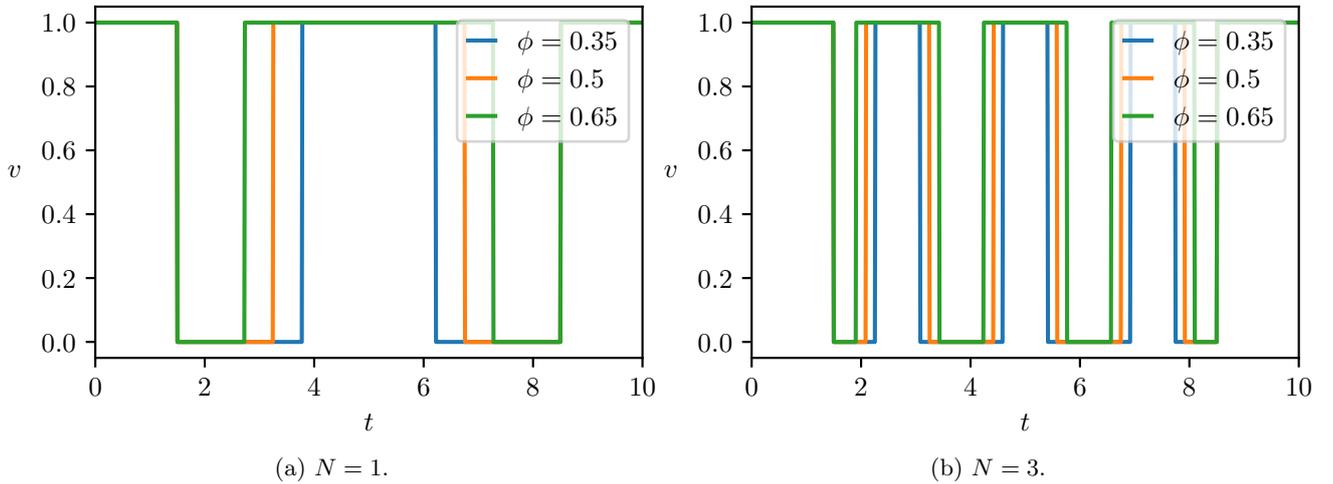


Figure 4: Parametrised controls $v(t; \phi, N, t_0)$ given in (55) for various values of N and ϕ with $t_0 = 1.5$.

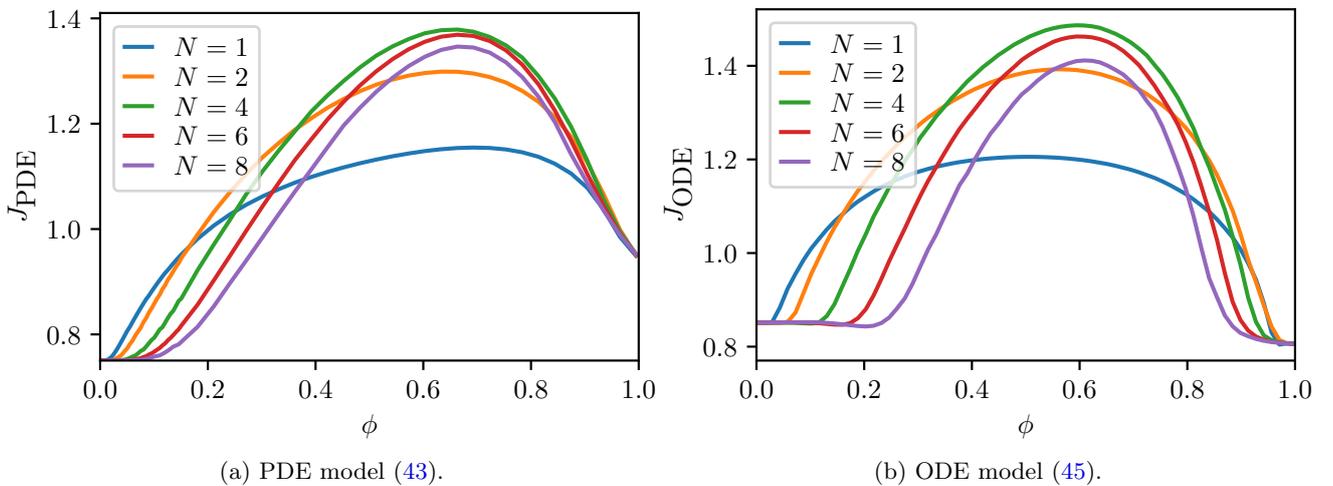


Figure 5: The protein production objective function J for both (a) PDE and (b) ODE formulations, using the parametrised family of controls $v(t; \phi, N, t_0)$ in (55) where $t_0 = 1.5$ and the remaining parameters are given in (54).

towards purely bang–bang forms, comprising a finite number of switches between the minimal and maximal control bounds. Thus, we can parametrise such controls by their switching structure, the times at which the control switches from one bound to the other. This formulation leads to the reduced finite-dimensional setting of maximising the objective for a fixed number of switches. We choose the naive approach of minimising each switching location one at a time, and repeating this until no improvement can be made on any location (more sophisticated methods are available, see [37], for example). For a fixed time discretisation of the PDE, this simple approach is convenient (although not particularly efficient) because we may choose the switching structure as integer indices. This additional parametrised optimisation step achieves two outcomes. First, it unifies the switching structures: controls with the same number of switches in fig. 7 can be slightly different, however, after the parametrised step they converge while improving the objective value. Second, by keeping the number of switches constant throughout the optimisation, this step allows us to study the space of locally optimal controls with a given number of switches, something not guaranteed by the gradient-based optimisation, where the switching structure is not fixed.

With the optimal controls for the PDE models (43) in hand, we proceed to optimise the ODE model (45). In fig. 8 we show the optimal controls for different initial control guesses. Here too we observe local optima to which the optimisation converges, although not for the simple initial conditions in fig. 8.

One prominent structural discrepancy between the ODE optimal controls and the PDE optimal controls is the deviation of the control from the bang–bang structure. As discussed in the analysis of section 4.2.1, these deviations must be singular arcs. Taking the solution $x(t)$ and $y(t)$ corresponding to the controls obtained in fig. 8, we calculate the singular control as by (51). We find that the singular arcs agree very well with the singular control calculation for the larger values of the control (fig. 8g black dotted curves). However, the smaller control values that deviate from $u(t) = 0$ do not match the singular control in (51). These regions are highly insensitive

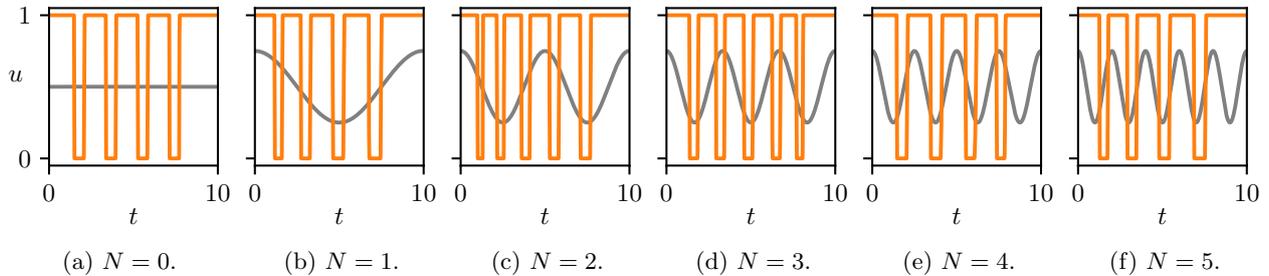


Figure 6: Optimal controls (orange curves) for the PDE model (43), using initial control guesses $u_0(t; N)$ defined in (56) for different values of N (grey curves). The optimisation was terminated after 1500 objective evaluations.

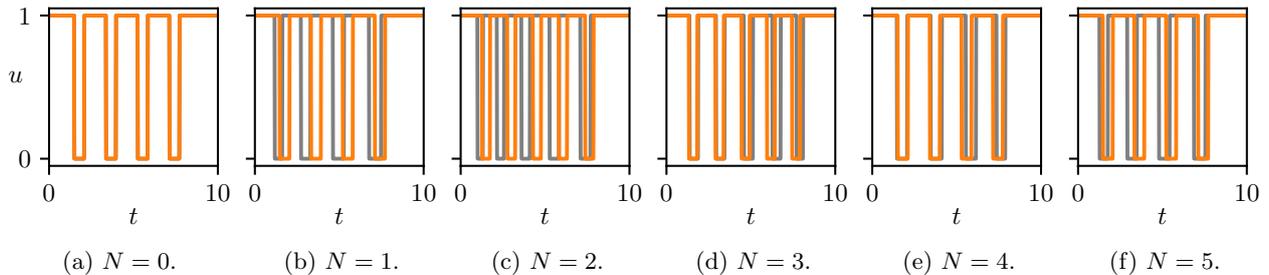


Figure 7: Optimal controls (orange curves) for the PDE model (43), using the results displayed in fig. 6 to parametrise an initial switching structure (grey curves) which is optimised iteratively for each switch in turn. The optimisation was terminated once no switch could be moved to improve the objective.

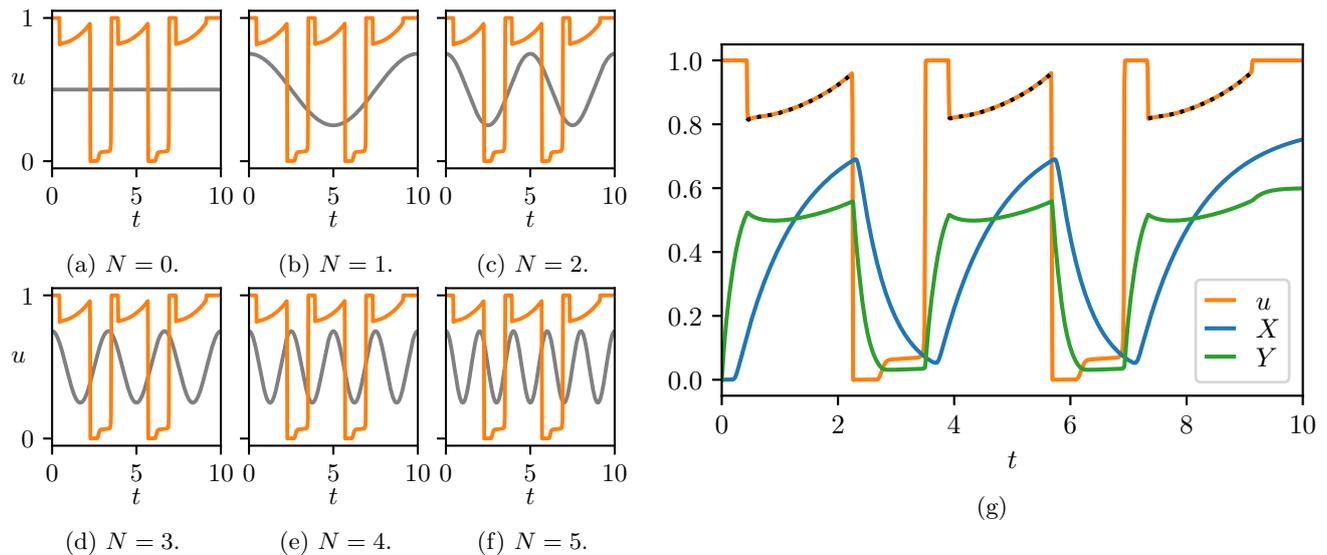


Figure 8: (a–f) Optimal controls (orange curves) for the ODE model (45), using initial control guesses $u_0(t; N)$ defined in (56) for different values of N (grey curves). The optimisation was terminated upon reaching a relative convergence tolerance of 10^{-11} . (g) Optimal control $u(t)$ from (a–f) alongside state variables $X(t)$ and $Y(t)$ (solid curves) overlaid with the singular control (black dotted curves) derived analytically in (51).

to small changes in the state and control (setting the control to $u(t) = 0$ instead of the singular arc changes the ODE objective function by a factor of $\mathcal{O}(10^{-5})$ and the PDE objective function by a factor of $\mathcal{O}(10^{-3})$), and we thus do not consider such anomalies to be the source of important error. Nevertheless, for the sake of completeness, we study the source of the error in Appendix E and show that it is an artefact of the discretisation.

4.3.3 Comparing the optimal controls

To test the performance of the ODE approximation, we again consider the PDE as the ground truth, and use the ODE controls in the PDE dynamics. There are multiple locally optimal PDE and ODE controls, which we distinguish by their M valleys, counting one valley as consecutive transitions from $u(t) = 1$ to $u(t) = 0$ and back

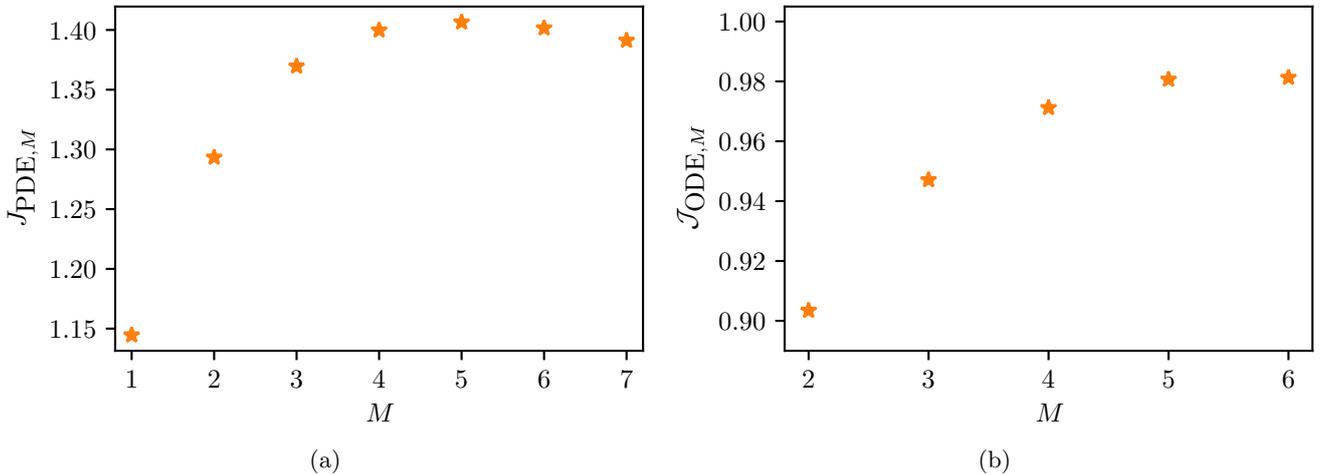


Figure 9: (a) Objective value of the PDE optimal controls $J_{\text{PDE},M}$ with M valleys using PDE model (43). (b) Relative performance of the ODE model (45) compared to the PDE model, as defined in (58), for different locally optimal ODE controls with M valleys. Other parameters are given in (54).

to $u(t) = 1$. Denoting such a control by $u_{\text{PDE},M}(t)$ or $u_{\text{ODE},M}(t)$, we calculate the objective with respect to the PDE dynamics, which we denote $J_{\text{PDE}}(u_{\text{PDE},M}(t))$ or $J_{\text{PDE}}(u_{\text{ODE},M}(t))$, respectively. We employ the shorthand

$$J_{\text{PDE},M} := J_{\text{PDE}}(u_{\text{PDE},M}(t)), \quad (57)$$

to denote the objective value of the locally optimal PDE control with M valleys. The measure of ODE performance is then taken as the ratio of this objective function to the best PDE objective among all (locally) optimal PDE controls, which may be written as

$$\mathcal{J}_{\text{ODE},M} := \frac{J_{\text{PDE}}(u_{\text{ODE},M}(t))}{\max_M J_{\text{PDE},M}}. \quad (58)$$

First, we study the PDE controls $J_{\text{PDE},M}$. In fig. 9a, we show the objective function values corresponding to the PDE optimal controls for different numbers of valleys M . We observe local maxima (with respect to M) for $M = 5$ valleys. The non-monotonicity with respect to M mirrors that seen for the parametrised controls in fig. 5a.

We may now calculate the relative ODE performance for different numbers of valleys M , as illustrated in fig. 9b. We observe what appears to lead to a non-monotonicity with respect to M (we could not find initial conditions that converged to controls with larger values of M so were unable to extend the plot to the region where we expect the non-monotonicity to manifest). However, at higher resolution the performance for this range of M values is non-monotonic (see Appendix F). This mirrors the non-monotonicity in fig. 5b.

Ultimately, we have demonstrated that the challenging optimal landscape may be tamed both in the PDE and ODE settings. The best ODE performance is a mere 2% shy of the optimal PDE result, suggesting that the ODE is an accurate surrogate model for the PDE control problem, even in the context of transient oscillatory behaviour.

4.4 Optimising the period

The optimal controls discussed until now have been computed with respect to objective functions that encapsulate total production by summing the running protein extracted from the system via turbidostat dilution with the protein remaining in the bioreactor at the final time $t = T$. For long-run production environments, the terminal cost is expected to be significant only towards the end of the time horizon. Motivated by long-term dynamics, where we are primarily concerned with the running objective, along with our observation of oscillatory optimal controls, we use the surrogate ODE model (45) to study the periodic free-endpoint problem of determining the optimal control and the optimal oscillation period τ^* that maximises the average running protein production:

$$\max_{\tau, 0 \leq u(t) \leq 1} \frac{1}{\tau} \int_0^\tau g(y(t))x(t) dt, \quad (59)$$

subject to the dynamics

$$\dot{x}(t) = \alpha(y(t)) - (\beta + g(y(t)))x(t), \quad x(0) = x(\tau), \quad (60a)$$

$$\dot{y}(t) = au(t) - (b + g(y(t)))y(t), \quad y(0) = y(\tau). \quad (60b)$$

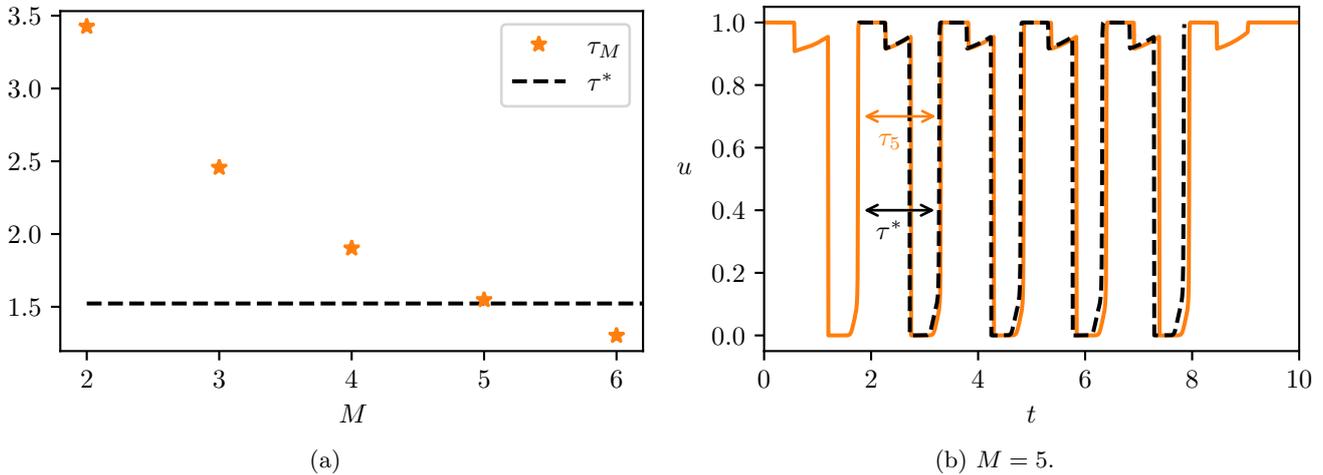


Figure 10: (a) Periods τ_M of the optimal controls $u_{\text{ODE},M}$ of the full ODE problem (45) excluding initial and terminal intervals (stars). The dashed line shows the optimal period τ^* as calculated for the periodic problem (59) and (60). (b,c) Optimal controls $u_{\text{ODE},M}$ of the full problem (45) (coloured curves), overlaid by $M - 1$ copies of the periodic optimal controls (59) and (60) (dashed black curves). Periods of each control are highlighted.

The associated optimal control of the periodic problem (59) and (60) is of a similar form to the problem posed on a fixed time horizon. The optimal oscillation period is $\tau^* \approx 1.523$. In fig. 10a we show these values against the locally optimal periods τ_M , denoting the period of the optimal controls of the full problem for controls with M valleys. The periods are calculated by ignoring the initial and terminal transients, where initial and terminal conditions and objectives have a dominant influence, and finding the average period. Naturally, the periods decrease for increasing M , and we find fairly close agreement between τ^* and τ_5 , the best performing ODE controls for each model.

To see the agreement more clearly, we plot the optimal control $u_{\text{ODE},5}$ of the full problem (45) and (45a) (coloured curves), and on each overlay $M - 1$ copies of the periodic optimal controls (59) and (60) (dashed black curves). We see that the optimal period is very close to, but slightly shorter than, the period of the aperiodic optimal control.

5 Inducing differentiation in bioproducing consortia

5.1 PDE concentration formulation

In this section we study an example of a microbial consortium (inspired by experiments such as in Ref. [1]), thereby relaxing the assumption that the cell colony is isogenic throughout the bioproduction. We consider a construct in which a photoreceptive TF is constitutively produced. Upon light induction, the TF is recruited in the production of recombinase which leads to recombination [1]. Only in the differentiated cell construct is the protein of interest produced. Ultimately, the light induction leads to a community of two genetically distinct subpopulations. We assume that the undifferentiated cells grow at some rate A , while the differentiated cells have an inhibited growth rate $0 \leq a < A$, as a result of the protein production. We thus call the undifferentiated cells the “growers” and the differentiated cells the “producers”. We further assume that the TF recruitment and cell recombination is faster than the other transient timescales and may be neglected.

Ultimately, we arrive at the population dynamics model

$$\begin{aligned} \frac{\partial g}{\partial t}(x, t) &= Ag(x, t) - \Lambda(t)g(x, t) - u(t)h(x)g(x, t) \\ &\quad - \frac{\partial}{\partial x}[(\alpha - \beta x - Ax)g(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2}[(\alpha + \beta x + Ax)g(x, t)], \end{aligned} \quad (61a)$$

$$\dot{p}(t) = ap(t) - \Lambda(t)p(t) + u(t) \int_0^\infty h(x)g(x, t) dx, \quad (61b)$$

where the TF concentration is distributed over the x domain, $\Lambda(t)$ is the total population growth rate, $h(x)$ is the differentiation rate upon (maximal) light induction for TF x , where the light intensity is $u(t)$, α and β are the birth and death rates of the TF concentration, respectively, and $g(x, t)$ and $p(t)$ represent the population density and population mass of the growers and producers, respectively. The dot represent differentiation with respect to time.

We assume that $0 \leq u(t) \leq 1$ is a dimensionless (normalised) light intensity, having absorbed the (dimensional) maximum light intensity into the modulation function $h(x)$ to give a rate. We employ modulation described by a Hill functional form

$$h(x) = f \frac{x^n}{K^n + x^n}, \quad (61c)$$

for a maximal differentiation rate f , a characteristic concentration K and a Hill coefficient n .

The Fokker–Planck components describing the TF expression are supplemented with the zero-flux boundary conditions

$$(\alpha - \beta x - Ax)g(x, t) - \frac{1}{2\Omega} \frac{\partial}{\partial x} [(\alpha + \beta x + Ax)g(x, t)] = 0, \quad (61d)$$

at $x = 0$ and as $x \rightarrow \infty$.

The turbidostat mode gives a dilution rate of

$$\Lambda(t) = ap(t) + A \int_0^\infty g(x, t) dx. \quad (61e)$$

We also impose initial conditions

$$g(x, 0) = g_0(x), \quad p(0) = 0, \quad (61f)$$

where $g_0(x)$ is the steady-state solution of (61a) to (61e) with no light input, $u(t) = 0$. This choice reproduces typical lab conditions where the optogenetic strains are cultivated in the dark for several cell generations.

Finally, we seek to maximise total production of the protein of interest, which we assume is proportional to the mass of producers:

$$J_{\text{PDE}} = p(T) + \int_0^T \Lambda(t)p(t) dt. \quad (61g)$$

We have thus arrived at the complete formulation of the optimal control problem: $\max_{0 \leq u(t) \leq 1} J_{\text{PDE}}$ subject to the population dynamics (61a) to (61f). The maximal optimal control in the vicinity of the terminal time, discussed in section 3.1.1, is not straightforwardly obtainable in this problem. Equipped with the optimal control problem we proceed to pose a homogeneous model, dispensing with the TF heterogeneity arising from cell-to-cell variability. In this way, we can probe the importance of including cell-to-cell variability for the purposes of modeling and optimal control.

5.2 Moment closure approximations

5.2.1 Zeroth-order closure

The PDE model (61) encapsulates heterogeneous expression of the transcription factor x that modulates the light-induced differentiation, which is lost in the zeroth-order approximation when integrating over the state x . It was for this reason that the moment closures performed in sections 3 and 4 include first-order dynamics. However, in the case of model (61) there is additional heterogeneity in the population structure as some cells are growers g and others producers p . In this case, the zeroth-order moment closure yields nontrivial dynamics since this second source of heterogeneity is preserved when integrating over the state space.

This example is of dimension $d = 1$ with $\mathcal{K} = 2$ states. To simplify the notation, we define

$$G(t) := X_{:,1}^0(t). \quad (62)$$

The zeroth-order approximation is

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - u(t)h\left(\frac{X_{:,1}^1(t)}{G(t)}\right)G(t), \quad G(0) = 1, \quad (63a)$$

$$\dot{p}(t) = ap(t) - \Lambda(t)p(t) + u(t)h\left(\frac{X_{:,1}^1(t)}{G(t)}\right)G(t), \quad p(0) = 0. \quad (63b)$$

The system (63) closes only at first order. However, since it is remarkably similar to homogeneous models typically used for such consortia (see, for example, Refs. [1, 2, 50]), it is worth seeking a zeroth-order closure.

To this end, we decompose the grower population $g(x, t)$ into its total mass $G(t)$, and its heterogeneous distribution $\hat{w}(x, t) := g(x, t)/G(t)$ which satisfies $\int_0^\infty \hat{w}(x, t) dx = 1$:

$$g(x, t) = G(t)\hat{w}(x, t) \quad \text{where} \quad \int_0^\infty \hat{w}(x, t) dx = 1. \quad (64)$$

Substituting (64) into (61), we obtain

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - Fu(t)G(t), \quad G(0) = 1, \quad (65a)$$

$$\dot{p}(t) = ap(t) - \Lambda(t)p(t) + Fu(t)G(t), \quad p(0) = 0, \quad (65b)$$

where the differentiation coefficient $F = \int_0^\infty h(x)\hat{w}(x, t) dx$ is a function of time. The turbidostat mode gives the dilution rate

$$\Lambda(t) = AG(t) + ap(t). \quad (65c)$$

The objective function (61g) retains the same form, namely,

$$J_{\text{ODE}} = p(T) + \int_0^T \Lambda(t)p(t) dt. \quad (65d)$$

We consider the zeroth-order moment closure where we approximate the time-dependent quantity F by some constant value. This is tantamount to assuming that the distribution $\hat{w}(x, t)$ is time invariant, which we expect to be a good approximation under steady conditions. Since typical bioproduction applications involve timescales of several cell generations, longer than the typical timescale of cellular processes, the grower distribution may be well approximated by a quasi-stationary distribution $\hat{w}(x, t) \approx w(x)$. Of course, high-frequency control inputs may cause the distribution to deviate from quasi-stationarity. Here, as introduced in section 2.1, the state-dependent selection, in this case the differentiation represented by $h(x)$, vanishes in the homogeneous model (65).

It is important to note that the ODE model (65) may be arrived at, a-priori, based on the physical description of the system provided at the beginning of section 5.1 if cell-to-cell variability is ignored and only the total mass of each subpopulation is modeled. It seems that this reasoning is the modus operandus within the field (e.g., Refs. [1, 2, 16, 17, 50]). What the above reduction shows is that, under an assumption of quasi-stationarity, it is an eminently reasonable approximation of the PDE model, which explains some of the good agreement and performance demonstrated in section 5.3. However, the derivation highlights why we do not expect the initial or terminal transient intervals — during which the quasi-stationarity assumption is manifestly violated — to be well approximated. Moreover, a new problem arises: how should we determine the parameter F ? One common approach is to perform characterisation experiments with some given light induction profile(s) $u(t)$, and fit the parameter F so that experimental results match as close as possible with theory.

We may leverage the density conservation ($G(t) + p(t) = 1$ for all time t) to eliminate $p(t)$ and thus reduce (65) to the one-dimensional problem

$$\dot{G}(t) = (A - a - Fu(t))G(t) - (A - a)G(t)^2, \quad G(0) = 1, \quad (66a)$$

where we aim to maximise the objective function

$$J_{\text{ODE}} = 1 - G(T) + \int_0^T a + (A - 2a)G(t) - (A - a)G(t)^2 dt. \quad (66b)$$

We call the ODE model (66) “homogeneous” with reference to the fact that the heterogeneity in TF level is not captured. However, the ODE model is not homogeneous in the broadest sense of the word: it preserves two subpopulations, thereby reflecting the genetic source of heterogeneity in the population. It is specifically the importance of the heterogeneity born of cell-to-cell variability within each subpopulation that we explore. This is the feature that distinguishes the PDE and ODE models.

5.2.2 First-order closure

A first-order moment closure approximation of model (61) allows us to capture additional features of the model and its x -distribution. Defining the shorthand

$$X(t) = X_{\cdot,1}^1(t), \quad (67)$$

we arrive at the ODE approximation

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - u(t)h\left(\frac{X(t)}{G(t)}\right)G(t), \quad (68a)$$

$$G(0) = 1, \quad (68b)$$

$$\dot{p}(t) = ap(t) - \Lambda(t)p(t) + u(t)h\left(\frac{X(t)}{G(t)}\right)G(t), \quad (68c)$$

$$p(0) = 0, \quad (68d)$$

$$\dot{X}(t) = -\Lambda(t)X(t) - u(t)h\left(\frac{X(t)}{G(t)}\right)X(t) + \alpha G(t) - \beta X(t), \quad (68e)$$

$$X(0) = \alpha/(\beta + A), \quad (68f)$$

where $\Lambda(t)$ is again given by (65c). The initial conditions are consistent with the PDE model (61f) (since we have $G(t) + p(t) = 1$ and $p(0) = 0$). The objective retains the form (65d).

The model (68) seems more involved than the zeroth-order system (65), however, we may ignore the p -dynamics via the conservation property $G(t) + p(t) = 1$. Similarly, we note that the ratio $X(t)/G(t)$ remains constant, since it satisfies

$$\frac{d}{dt}\left(\frac{X(t)}{G(t)}\right) = (\beta + A)\left(\frac{\alpha}{\beta + A} - \frac{X(t)}{G(t)}\right), \quad \frac{X(0)}{G(0)} = \frac{\alpha}{\beta + A}. \quad (69)$$

It follows that $X(t)/G(t) = \alpha/(\beta + A)$ for all $t > 0$. Therefore, the first-order dynamics (68) reduce to the one-dimensional

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - u(t)h\left(\frac{\alpha}{\beta + A}\right)G(t), \quad G(0) = 1. \quad (70)$$

along with the objective (66b). In other words, we obtain the zeroth-order system (65) where the differentiation parameter is given by

$$F = h\left(\frac{\alpha}{\beta + A}\right). \quad (71)$$

This allows us to characterise the ODE model by appealing to the full heterogeneous model, without additional experiments (although assumes that we possess knowledge of these PDE parameters). Again at first order, the state-dependent selection is aggregated via a constant.

5.2.3 Optimal control of the ODE model

A complete account of the dynamics and optimal control of (66) is provided in Ref. [28]. In contrast to the PDE case, the ODE case admits a closed-form solution by means of Pontryagin's maximum principle. Here, we suffice with a summary of the main results: If $A \leq 2a$ or $A \geq 2F$, the optimal control is simply $u(t) = 1$ identically. If $2a < A < 2F$, then there exists a critical time horizon $T_{\text{crit}} = t_s + T_s$ such that for $T \leq T_{\text{crit}}$ the optimal control is $u(t) = 1$. However, if $2a < A < 2F$ and $T > T_{\text{crit}}$, then the optimal control is $u(t) = 1$ on $t \in [0, t_s)$, followed by an interval on which $u(t) = A/(2F)$ for $t \in [t_s, T - T_s]$, and then $u(t) = 1$ for the remaining $t \in (T - T_s, T]$. Explicit expressions are available for the durations t_s and T_s .

Physically, when the growth rates of the two populations are not sufficiently separated, $A \leq 2a$, the optimal control is to simply differentiate at the maximum rate. When the growth rates are sufficiently separated, $A > 2a$, but the differentiation rate is not sufficiently faster than the growth rate, $A \geq 2F$, the optimal control is again to differentiate maximally. When the growth rates are sufficiently separated and the differentiation is sufficiently fast, $A > 2a$ and $A < 2F$, the optimal control may include an arc on which the control is not bang-bang if the time horizon is long enough. This critical time is required to allow the state to evolve to a balance between production and growth that maximises the running protein production, while allowing a final period sufficient to obtain an optimal contribution from the protein present at the terminal time.

In all cases, it transpires that the optimal control is maximal near the terminal time.

5.3 Numerical results: comparing the PDE and ODE models

Equipped with the heterogeneous PDE problem formulation (61) and the homogeneous ODE problem formulation (66), we now compare the solutions.

We first note that the ODE problem is significantly more tractable: we have deduced an explicit solution of the optimal control. Notwithstanding, explicit solutions or fast calculations may not be of great utility if the results

are inaccurate, and thus we must compare the optimal controls and their performance on the underlying ground-truth model. We are assuming that the PDE dynamics are the true system and measurements are noiseless, continuous in time and fully observable (although we will discuss partial observability momentarily). We choose parameter values of the ground-truth PDE model and assume exact knowledge of these in the ODE model.

We then take optimal controls of the ODE model, u_{ODE} , and use them in the ground-truth PDE model, denoting the associated objective pay-off by $J_{\text{PDE}}(u_{\text{ODE}})$. In practice, it is more common to employ a closed-loop control to allow corrections due to noisy observations and dynamics as well as model mismatch. While we suppress noise in this work, model mismatch (between the ground-truth heterogeneous PDE model and its homogeneous approximations) will be a source of deviation. Therefore, we also implement a commonly used control algorithm that closes the loop in an attempt to address this source of error. The procedure is called receding-horizon control, or model-predictive control (MPC), as we will refer to it here, and involves periodically observing state information and recalculating the open-loop solution given the revised state, as follows. We solve the open-loop ODE problem and apply this control in the ground-truth PDE simulation over some time interval $t \in [0, t_1]$. After this first time horizon, we observe the state of the PDE system from which we derive the state of the ODE system $G(t_1)$. We then use this value as an initial state and solve the ODE optimal control over the remaining time horizon $[t_1, T]$ given this revised initial state. We repeat this process at several intervals $t_1 < t_2 < \dots < T$. We choose the state updates to occur hourly to mimic common laboratory capabilities. For this reason we solve the optimal control problem for an arbitrary initial condition in Ref. [28]. We denote the associated objective pay-off by $J_{\text{PDE}}(u_{\text{MPC}})$.

To compare with the objective pay-off of the heterogeneous model with its optimal control u_{PDE} , we consider the ratios

$$\mathcal{J}_{\text{ODE}} := \frac{J_{\text{PDE}}(u_{\text{ODE}})}{J_{\text{PDE}}(u_{\text{PDE}})}, \quad \mathcal{J}_{\text{MPC}} := \frac{J_{\text{PDE}}(u_{\text{MPC}})}{J_{\text{PDE}}(u_{\text{PDE}})}, \quad (72)$$

representing the relative performance of the ODE controls, evaluated in the PDE model, compared with the PDE controls u_{PDE} evaluated in the PDE model.

For all numerical experiments, unless otherwise stated, we used the default values

$$A = \alpha = 1, \quad a = \beta = 0, \quad n = 10, \quad K = 0.5, \quad f = 4, \quad T = 10, \quad \Omega = 1000. \quad (73)$$

There are two implementations for the MPC loop, depending on the observability of the state. If only the grower mass $G(t)$ is observable, then we may use the one-dimensional system (70) which has enveloped the higher order variable $X(t)$. If full distributional information is observable in real time, then we can also update the state $X(t)$ and solve the optimal control for system (68).

In fig. 11 we show the optimal controls corresponding to the full PDE model, the open-loop ODE model, and the MPC control (assuming full observability thus updating both $G(t)$ and $X(t)$). We observe a similar general structure for all values of K and both the PDE and ODE controls: initial and terminal intervals on which the control is maximum $u(t) = 1$ with perhaps an intermediate interval during which the control is not at its maximum, conforming with the open-loop ODE solution. The closed-loop MPC control imposes some minor corrections but remains broadly alike. The MPC control when updating only $G(t)$ looks similar and achieves almost identical performance (see Appendix G).

For $K = 0.3$ the fit is excellent and the ODE controls are barely distinguishable from the PDE control, capturing both the initial and terminal transient intervals as well as the intermediate singular arc, with performance metrics of near unity. As K increases, the fits become increasingly poor. For $K = 0.4$ and $K = 0.5$, the open-loop control deviates from the PDE control in underestimating the initial and terminal transient durations as well as the control value along the singular arc. The cost in performance compared to the full PDE optimisation, $1 - \mathcal{J}_{\text{ODE}}$, is less than 4%. The MPC control is able to alleviate much of this mismatch and improves the performance ($\mathcal{J}_{\text{MPC}} > \mathcal{J}_{\text{ODE}}$), incurring performance costs of less than 0.5%. We conclude that the homogeneous ODE approximation (65) provides a good surrogate model to replace the optimal control of the heterogeneous PDE problem (61) for this choice of parameters.

In light of the excellent agreement obtained by the ODE model it is tempting to draw broad conclusions; perhaps the neglected heterogeneity is of little importance in a wide range of such PDE models. However, this inference turns out to be premature, and the homogeneous approach is less robust than these initial results suggest, as we proceed to demonstrate.

6 Differentiation in consortia with modulated production

6.1 PDE concentration formulation

In this section we extend model (61) from section 5 by introducing a new source of heterogeneity. Beginning with the physical description introduced in section 5.1, we revise the source of the growth inhibition in the producer

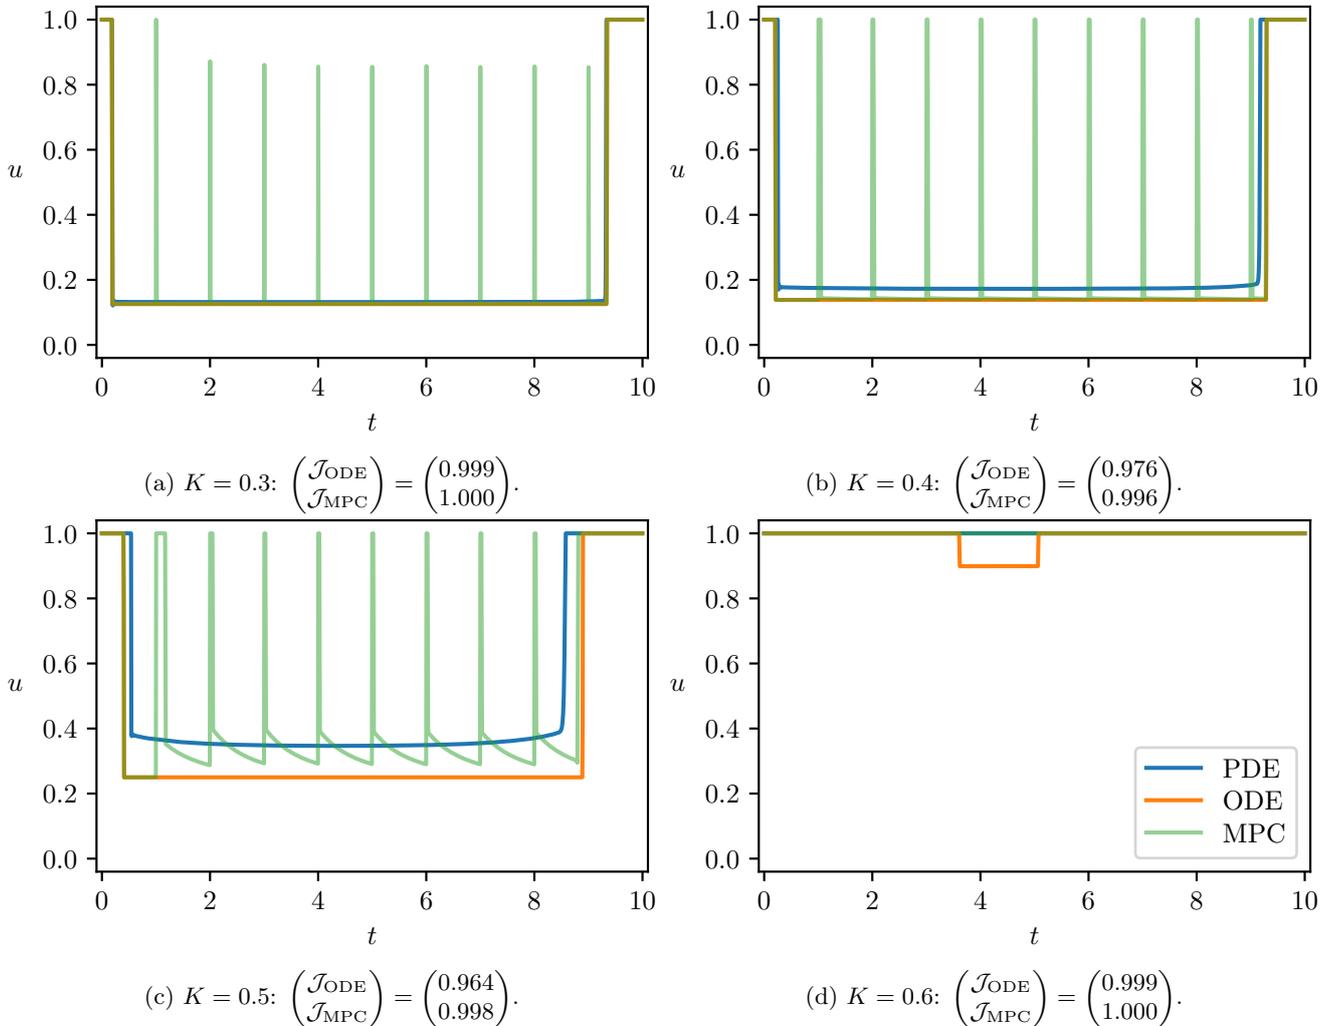


Figure 11: Optimal controls from the PDE model (61), the ODE model (68), and the same ODE model in closed-loop MPC (updating both $G(t)$ and $X(t)$), for various values of the Hill parameter K . Other parameters are given in (73).

cells. We previously modeled the burden of protein production via a uniformly diminished growth rate, representing the fact that the burden is largely independent of protein concentration. We now consider a case in which the growth is inhibited not by protein production burden but instead by the presence of the protein produced, which is toxic to the cell. In fact, in such cases, microbial consortia are considered particularly useful, since a healthy and growing subpopulation can be maintained alongside the subpopulation that is effectively poisoning itself to produce the protein of interest, thereby retaining both a growing population and protein production [38]. In this case, the ground-truth model is revised to include heterogeneity in the protein concentrations, on which the diminished growth rate depends. To this end, we must incorporate in the model the mechanism of protein production, which we assume is governed by an underlying birth–death process.

The protein level is distributed in an additional independent variable y , and the growth rate $\hat{a}(y)$ is modulated by the presence of the protein: $\hat{a}(0) = A$, while $\hat{a}(\infty) = a_\infty \in [0, A)$. We consider the monotonically decreasing Hill form

$$\hat{a}(y) = a_\infty + \frac{A - a_\infty}{1 + (y/L)^m}, \quad (74)$$

where the growth rate $\hat{a}(y) \in [a_\infty, A]$ interpolates nonlinearly between the minimum and maximum growth rates, a_∞ and A , respectively, with the characteristic protein concentration associated with diminished growth is L .

The model then takes the form

$$\begin{aligned} \frac{\partial g}{\partial t}(x, t) &= Ag(x, t) - \Lambda(t)g(x, t) - u(t)h(x)g(x, t) \\ &\quad - \frac{\partial}{\partial x}[(\alpha - \beta x - Ax)g(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2}[(\alpha + \beta x + Ax)g(x, t)], \end{aligned} \quad (75a)$$

$$\begin{aligned} \frac{\partial p}{\partial t}(y, t) &= \hat{a}(y)p(y, t) - \Lambda(t)p(y, t) \\ &\quad - \frac{\partial}{\partial y}[(\mu - \nu y - \hat{a}(y)y)p(y, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial y^2}[(\mu + \nu y + \hat{a}(y)y)p(y, t)]. \end{aligned} \quad (75b)$$

In addition to the TF boundary conditions (61d), we impose the normal flux boundary conditions associated with the protein production, namely,

$$\begin{aligned} (\mu - \nu y - \hat{a}(y)y)p(y, t) - \frac{1}{2\Omega} \frac{\partial}{\partial y}[(\mu + \nu y + \hat{a}(y)y)p(y, t)] \\ = \begin{cases} u(t) \int_0^\infty h(x)g(x, t) dx, & y = 0, \\ 0, & y \rightarrow \infty. \end{cases} \end{aligned} \quad (75c)$$

In turbidostat mode, the dilution rate is given by

$$\Lambda(t) = \int_0^\infty Ag(x, t) dx + \int_0^\infty \hat{a}(y)p(y, t) dy. \quad (75d)$$

We impose initial conditions

$$g(x, 0) = g_0(x), \quad p(0) = 0, \quad (75e)$$

where $g_0(x)$ is again the steady-state solution of (75a) to (75d) with no light input, $u(t) = 0$.

We seek to maximise total production of the protein of interest, which is given by

$$J_{\text{PDE}} = \int_0^\infty yp(y, T) dy + \int_0^T \Lambda(t) \int_0^\infty yp(y, t) dy dt. \quad (75f)$$

In this case too, the optimal control near the terminal time cannot be straightforwardly determined as in section 3.1.1.

6.2 Moment closure approximations

Armed with the extended PDE model (75) encapsulating the heterogeneity of both TF and protein levels, we now turn our attention to a homogeneous approximation. Adopting the same approach as in section 5.2.1, we derive the zeroth-order approximation by integrating over the state space and assuming quasi-stationary distributions. The resulting ODE remains identical in form to system (65), with the constant parameter a approximating the time-dependent average diminished growth rate given by

$$a \approx \frac{\int_0^\infty \hat{a}(y)p(y, t) dy}{\int_0^\infty p(y, t) dy}. \quad (76)$$

This parameter a would, in principle, have to be fitted based on characterisation experiments. Since it is an average $\hat{a}(y)$ defined in (74), we know that it must lie in the range of the function: $[a_\infty, A]$. In this case, the first-order closure will not provide a closed-form expression for this average, but will shed more light on its nature.

We now turn to the first-order moment closure approximation of the extended model (75). Defining

$$G(t) = X_{\cdot,1}^0(t), \quad X(t) = X_{1,1}^1(t), \quad P(t) = X_{\cdot,2}^0(t), \quad Y(t) = X_{2,2}^1(t), \quad (77)$$

we obtain the first-order system

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - u(t)h\left(\frac{X(t)}{G(t)}\right)G(t), \quad (78a)$$

$$G(0) = 1, \quad (78b)$$

$$\dot{P}(t) = \hat{a}\left(\frac{Y(t)}{P(t)}\right)P(t) - \Lambda(t)P(t) + u(t)h\left(\frac{X(t)}{G(t)}\right)G(t), \quad (78c)$$

$$P(0) = 0, \quad (78d)$$

$$\dot{X}(t) = -\Lambda(t)X(t) - u(t)h\left(\frac{X(t)}{G(t)}\right)X(t) + \alpha G(t) - \beta X(t), \quad (78e)$$

$$X(0) = \alpha/(\beta + A), \quad (78f)$$

$$\dot{Y}(t) = -\Lambda(t)Y(t) + \mu P(t) - \nu Y(t), \quad (78g)$$

$$Y(0) = 0, \quad (78h)$$

where the turbidostat dilution rate is given by

$$\Lambda(t) = AG(t) + \hat{a}\left(\frac{Y(t)}{P(t)}\right)P(t). \quad (78i)$$

The initial conditions are consistent with the PDE model and the above approximations. The objective is given by

$$J_{\text{ODE}} = Y(T) + \int_0^T \Lambda(t)Y(t) dt. \quad (79)$$

Strictly speaking, the model (78) is ill posed because the terms $Y(t)/P(t)$ are undefined at $t = 0$ with the initial condition $P(0) = 0$. This is easily remedied by considering the small-time asymptotics of the evolution near the first light impulse, which we assume, without loss of generality, is at the initial time $t = 0$. Taking the scaled time $t = \epsilon \hat{t}$ for an infinitesimally small $0 < \epsilon \ll 1$ and $\hat{t} = \mathcal{O}(1)$, we scale the states via $P(t) = \epsilon \hat{P}(\hat{t})$ and $Y(t) = \epsilon^2 \hat{Y}(\hat{t})$. Dropping the hats for notational convenience, we see that, to leading order with respect to ϵ , $G(t) \sim 1$, and the dynamics of $Y(t)$ and $P(t)$ reduce to

$$\dot{P}(t) \sim u(t)h\left(\frac{\alpha}{\beta + A}\right), \quad \dot{Y}(t) \sim \mu P(t). \quad (80)$$

Therefore, the ratio $Y(t)/P(t) = \mathcal{O}(\epsilon)$ for any $u(t) = \mathcal{O}(1)$, in other words, it remains vanishingly small near the initial time and thus we may consider $\hat{a}(Y(0)/P(0))$ to take the value $\hat{a}(0)$ for $t = 0$.

In the first-order moment closure (78), the grower population dynamics remain identical to those of (68) in section 5.2.2, and, therefore, the mean TF concentration $X(t)/G(t)$ remains constant, as per (69). In contrast, the mean protein concentration $Y(t)/P(t)$, and thus the diminished growth rate, $\hat{a}(Y(t)/P(t))$, remain time-varying. Nevertheless, we note that the ratio $Y(t)/P(t)$ satisfies

$$\begin{aligned} \frac{d}{dt} \left(\frac{Y(t)}{P(t)} \right) &= \mu - \left[\nu + \hat{a}\left(\frac{Y(t)}{P(t)}\right) + u(t)h\left(\frac{X(t)}{G(t)}\right) \frac{G(t)}{P(t)} \right] \frac{Y(t)}{P(t)} \\ &\leq \mu - (\nu + a_\infty) \frac{Y(t)}{P(t)}, \end{aligned} \quad (81)$$

with an initial value of zero, based on the preceding discussion of the small-time asymptotics. It follows that $Y(t)/P(t) \leq \mu/(\nu + a_\infty)$, and thus the approximate diminished growth rate is bounded by $\hat{a}(Y(t)/P(t)) \geq \hat{a}(\mu/[\nu + a_\infty])$. When $\mu/(\nu + a_\infty)$ is significantly smaller than the characteristic value L , this bound suggests that the zeroth-order parameter a may be approximated by $a \approx A$.

The optimal control is again maximal in a vicinity of the terminal time as we demonstrate in Appendix H.

6.3 Numerical results: comparing the PDE and ODE models

In addition to the defaults in (73), numerical solutions in this section also employ the default values

$$a_\infty = 0, \quad m = 10, \quad \mu = 0.5, \quad \nu = 0.7. \quad (82)$$

Our central measure of performance is the ratio of the PDE objective function when applying the ODE or MPC control compared to the PDE control, denoted \mathcal{J}_{ODE} and \mathcal{J}_{MPC} , respectively, and defined in (72).

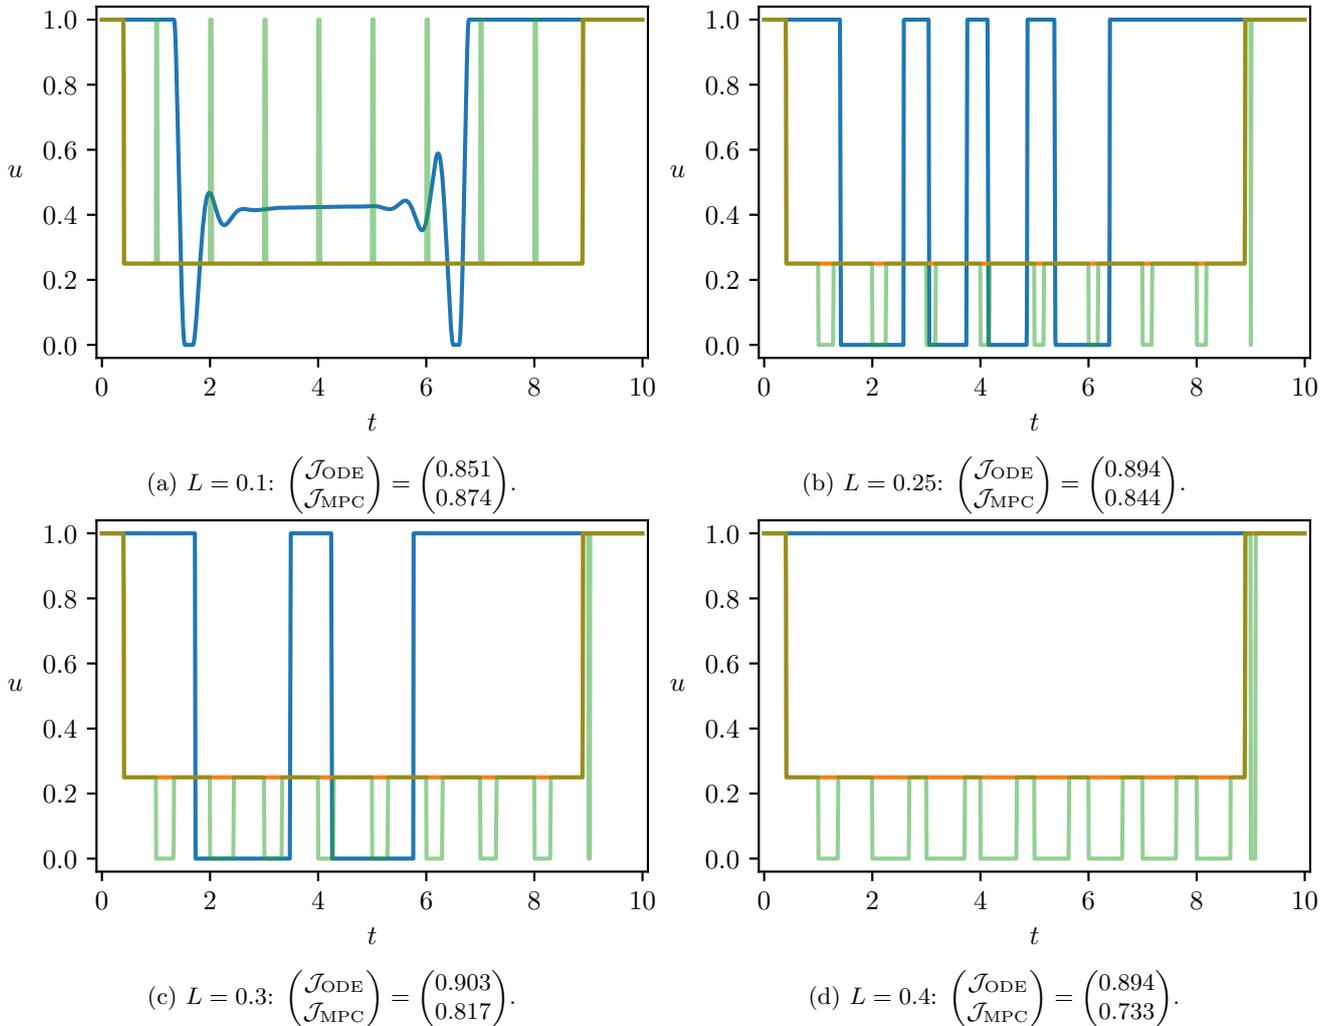


Figure 12: Optimal controls from the PDE model (75), the zeroth-order ODE model (65), and the same ODE model in closed-loop MPC, for various values of the Hill parameter K . For the ODE model, the parameter F is given by (71) while the parameter a is given by fitting. Other parameters are given in (73) and (82). The PDE control in (a) is obtained with an additional regularisation term $\epsilon \int_0^T (du/dt)^2 dt$ in the cost function with $\epsilon = 10^{-4}$ to avoid high-frequency oscillations. This regularisation comes at a negligible cost: the objective pay-off, calculated without any regularisation, is reduced by a factor of only $\mathcal{O}(10^{-6})$.

In section 6.2 we derived both zeroth- and first-order moment closure approximations. While the first-order approximation (78) produces an ODE system characterised by the underlying PDE model parameter, the zeroth-order approximation (producing the ODE system (65)) yields the composite parameter a which estimates the quantity in (76). For the numerical simulations performed in this section, we fit the parameter a by performing an *in silico* characterisation experiment: we choose the control profile $u(t) = 1$ and run the full PDE model, which gives the state $G(t)$ over a time interval, which we fit to the model by minimising the square residual. This mimics common practice, where models are matched to data to identify parameters.

In fig. 12 we show the optimal controls corresponding to the full PDE model (75), the open-loop zeroth-order ODE model (65), and the corresponding MPC control as we vary the Hill parameter L . In contrast to the controls of section 5, the PDE optimal control (blue curves) has a richer structure: there are singular arcs and (damped) switching (fig. 12a), transient oscillatory controls (figs. 12b and 12c), and, for large enough K , a transition to the full-light control (fig. 12d).

The ODE model is dependent on L only via the fitting procedure. That the open-loop ODE controls (orange curves) in fig. 12 are indistinguishable for all L is a result of the fact that the fitting gives $a \approx 0$ for all values of L . Nevertheless, the ODE performance is always above 85% of the full PDE optimal control. Considering the closed-loop MPC control, it is surprising to observe that, other than for $L = 0.1$, the performance is significantly degraded, reaching a low of 73% of the full PDE optimal control. When the surrogate model does not faithfully capture the system dynamics, closing the loop via MPC can amplify the mismatch. Another prominent discrepancy is the duration of the initial and terminal intervals of full light intensity, which the ODE models consistently

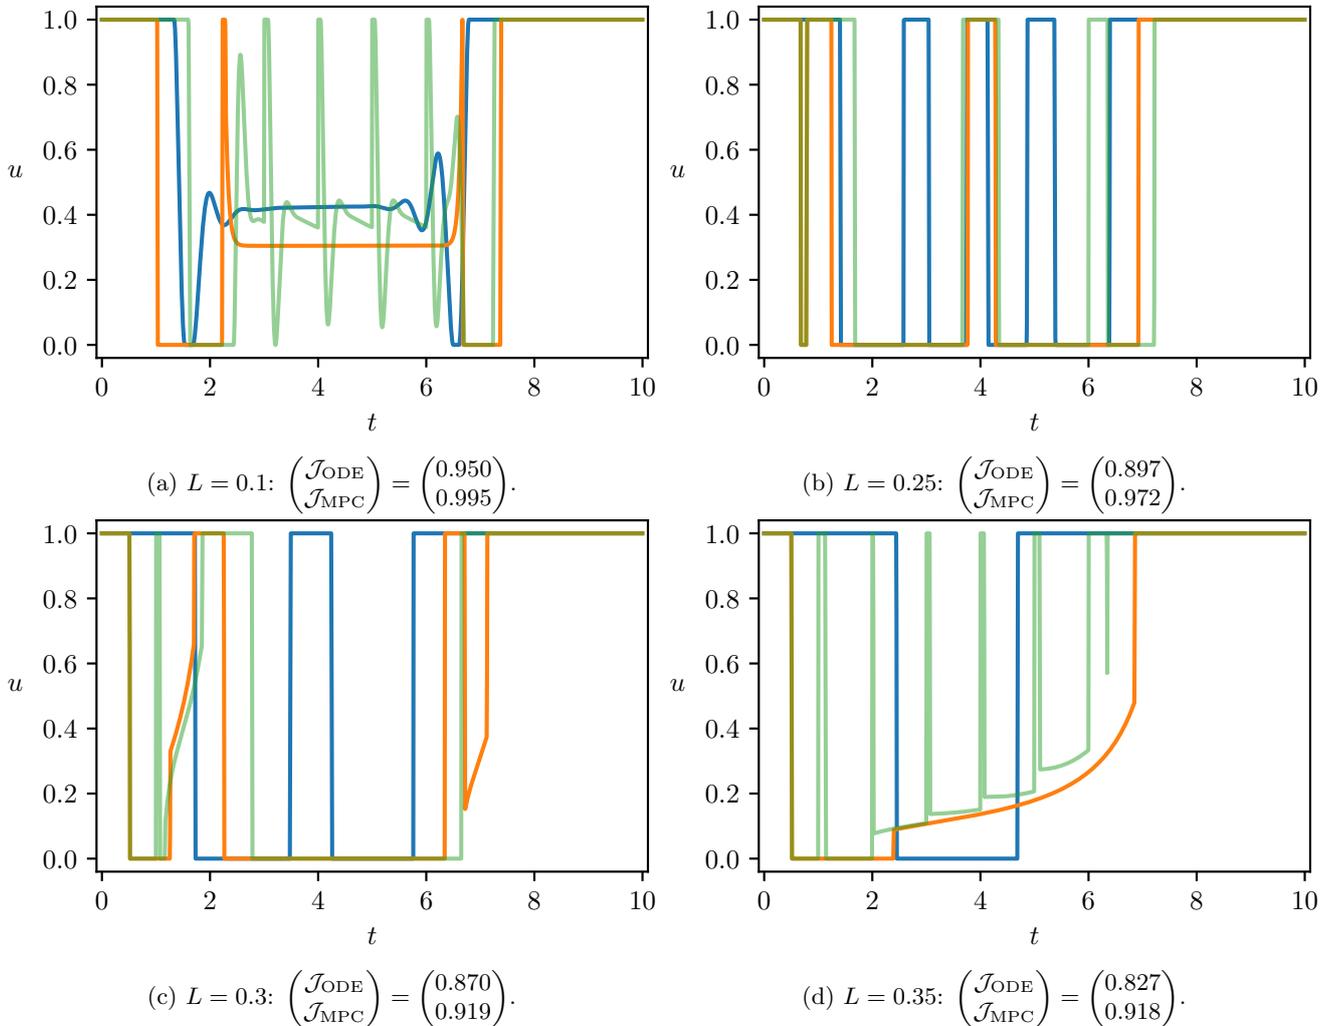


Figure 13: Optimal controls from the PDE model (75), the first-order ODE model (78), and the same ODE model in closed-loop MPC, for various values of the Hill parameter K . Other parameters are given in (73) and (82). The controls in (a) is obtained with an additional regularisation term as detailed in fig. 12.

underestimate.

It might be tempting to think that the parameter value for F derived by first-order consideration lends a hand to the poor performance. In Appendix I, we reproduce the ODE results by fitting F (rather than prescribing F via (71)). While this improves both the open- and closed-loop performance, particularly for small values of L (in particular, the singular arc is well approximated), the most severely degraded performance remains essentially unimproved. In practice, since fitting involves matching noisy and sparse experimental data, it is not clear how much of the improved performance (observed here under noiseless *in silico* conditions with dense measurements) could be attained.

We deduce that the suboptimal performance is a result of the zeroth-order model not sufficiently capturing the underlying dynamics, and we turn our attention to the first-order moment closure approximation.

The open-loop ODE controls in fig. 13 succeed in qualitatively reproducing more of the PDE optimal control structure. The initial and terminal transients of full light are longer than their zeroth-order counterparts, although still underestimating the PDE transients. The singular ODE control in fig. 13a captures not only the nonsingular and singular arcs, but also the switching transitions between the two. The MPC corrections to the singular arc in fig. 13a do a strikingly good job of approximating the PDE control, performing only 0.5% worse than the full PDE control. For higher values of L the singular arc is replaced by pulse forms reminiscent of the PDE control (figs. 13b to 13d). These pulse forms do not match the PDE pulses closely enough to give the open-loop control noticeably improved performance over the zeroth-order controls (in fact, the performance is often noticeably reduced). However, the MPC performance *uniformly* improves upon its open-loop counterpart, performing at above 91% of the PDE control in all cases studied. The case of $L = 0.4$ (not shown in fig. 13), where the zeroth-order performance is 73% (fig. 12d), increases to 100% as the optimal control matches the full-light profile.

We conclude that the zeroth-order moment closure approximation admits some shortcomings, in particular, the

requirement of parameter fitting and the poor closed-loop performance. These are significantly improved by the first-order approximation, which do not introduce new parameters and greatly improve close-loop performance. Nevertheless, this example demonstrates cases where the first-order approximation, despite reasonable closed-loop performance, admits an optimal control that is qualitatively mismatched to the full PDE optimal control: higher-order dynamics have important influence on the dynamics and control [29].

In this study, we have intentionally chosen a simple ground-truth model of microbial consortia: there are just two subpopulations, and, at first, just a single additional heterogeneous level stemming from cell-to-cell variability. The extended model had two heterogeneous levels, but these sources of variability did not interact. One purpose of this simplicity is to retain tractability, which allows us to glean deep insight as to how the physical parameters alter the dynamics and the optimal control. A second motivation for studying these simple models is to reveal the minimal set of ingredients required for the influence of the heterogeneity to prove critical. Surprisingly, even a small excursion from the simplest heterogeneous setting begins to challenge the reduced-order surrogates.

7 Conclusions

We study a class of population-level models that capture heterogeneous gene expression born of underlying stochastic chemical kinetics at the single-cell level. By considering generic influences in bioproduction processes independently, each with several modeling realisations, we allow a modular composition that spans a particularly wide class of bioproduction systems. This makes our approach applicable to different bioproduction infrastructure (e.g. batch vs. turbidostat) as well as different modeling paradigms (e.g. concentration vs. copy number). Naturally, the heterogeneous models are cumbersome both analytically and numerically. We thus introduce low-order moment-closure techniques that significantly reduce the model complexity while retaining the same modular construction. We demonstrate that these reduced-order models are often excellent surrogate models, and their optimal control provides good control performance on the heterogeneous system. Moreover, their analysis is significantly more tractable, and can provide deep physical insight into the dynamics and controls. The numerical optimisation of the ODE models typically took on the order of seconds, a considerable improvement over the PDE optimisation that took typically on the order of several minutes. Nevertheless, we also demonstrate that, in certain circumstances, the homogeneous ODE approximations can fall dramatically short, providing suboptimal controls when the underlying approximation assumptions are violated. This demonstrates the potential of the heterogeneity to significantly impact the optimal control.

Given that both approaches have their advantages and disadvantages, our primary focus is to study both simultaneously. This is particularly fruitful as the two approaches complement each other: the ODE models can provide informed initial guesses for the optimal control, saving valuable computational effort and sometimes being the difference between convergence and divergence, while the PDE models can inform parameter choice for their ODE counterparts. This mutually beneficial interplay suggests that simultaneously developing both approaches makes for a powerful toolkit for optimal control of bioproduction.

The models here were chosen to exhibit a variety of useful paradigms in bioproduction modeling and are intended to be representative but not exhaustive. It could prove important to study models where both the copy-number and concentration (by tracking cell size) are accounted for [45]. The model lends itself to many immediate biologically relevant extensions, such as copy-number models with division of discrete species.

More sophisticated heterogeneity-aware order-reduction techniques need to be developed for cases where the reduced-order methods presented here prove insufficient. We expect this to be particularly important for more complicated heterogeneity structure, such as systems with multimodal distribution. In these cases, moment closures of orders higher than the first are capable of capturing the additional distribution structure, however, these come with several associated analytical and numerical challenges.

References

- [1] C. Aditya, F. Bertaux, G. Batt, and J. Ruess. A light tunable differentiation system for the creation and control of consortia in yeast. *Nature communications*, 12(1):1–10, 2021.
- [2] C. Aditya, F. Bertaux, G. Batt, and J. Ruess. Using single-cell models to predict the functionality of synthetic circuits at the population scale. *P. Natl. Acad. Sci. USA*, 119(11):e2114438119, 2022.
- [3] S. J. Altschuler and L. F. Wu. Cellular heterogeneity: do differences make a difference? *Cell*, 141(4):559–563, 2010.
- [4] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl. CasADi: a software framework for nonlinear optimization and optimal control. *Math. Program. Comput.*, 11(1):1–36, 2019.

- [5] C. Barajas and D. Del Vecchio. Effects of spatial heterogeneity on bacterial genetic circuits. *PLOS Comput. Biol.*, 16(9):1–25, 09 2020.
- [6] K. Brenner, L. You, and F. H. Arnold. Engineering microbial consortia: a new frontier in synthetic biology. *Trends Biotechnol.*, 26(9):483–489, 2008.
- [7] A. Brock, H. Chang, and S. Huang. Non-genetic heterogeneity—a mutation-independent driving force for the somatic evolution of tumours. *Nat. Rev. Genet.*, 10(5):336–342, 2009.
- [8] C. Carrasco-López, S. A. García-Echauri, T. Kichuk, and J. L. Avalos. Optogenetics and biosensors set the stage for metabolic cybergenetics. *Curr. Opin. Biotech.*, 65:296–309, 2020.
- [9] F. Ceroni, A. Boo, S. Furini, T. E. Gorochoowski, O. Borkowski, Y. N. Ladak, A. R. Awan, C. Gilbert, G.-B. Stan, and T. Ellis. Burden-driven feedback control of gene expression. *Nat. Methods*, 15(5):387–393, 2018.
- [10] S. Cooper. Distinguishing between linear and exponential cell growth during the division cycle: single-cell studies, cell-culture studies, and the object of cell-cycle research. *Theor. Biol. Med. Model.*, 3(1):1–15, 2006.
- [11] L. Duso and C. Zechner. Stochastic reaction networks in dynamic compartment populations. *Proc. Natl. Acad. Sci. U.S.A.*, 2020.
- [12] N. Friedman, L. Cai, and X. S. Xie. Linking stochastic dynamics to population distribution: An analytical framework of gene expression. *Phys. Rev. Lett.*, 97:168302, Oct 2006.
- [13] M. B. Giles and N. A. Pierce. An introduction to the adjoint approach to design. *Flow Turbul. Combust.*, 65(3):393–415, 2000.
- [14] D. T. Gillespie. The chemical Langevin equation. *J. Chem. Phys.*, 113(1):297–306, 2000.
- [15] R. Grima, P. Thomas, and A. V. Straube. How accurate are the nonlinear chemical Fokker-Planck and chemical Langevin equations? *J. Chem. Phys.*, 135(8):084103, 2011.
- [16] J. Harmand, C. Lobry, A. Rapaport, and T. Sari. *The chemostat: Mathematical theory of microorganism cultures*. John Wiley & Sons, USA, 2017.
- [17] J. Harmand, C. Lobry, A. Rapaport, and T. Sari. *Optimal Control in Bioprocesses: Pontryagin’s Maximum Principle in Practice*. John Wiley & Sons, USA, 2019.
- [18] S. Huang. Non-genetic heterogeneity of cells in development: more than just noise. *Development*, 136(23):3853–3862, 2009.
- [19] S. Iyer-Biswas, C. S. Wright, J. T. Henry, K. Lo, S. Burov, Y. Lin, G. E. Crooks, S. Crosson, A. R. Dinner, and N. F. Scherer. Scaling laws governing stochastic growth and division of single bacterial cells. *P. Natl. Acad. Sci. USA*, 111(45):15912–15917, 2014.
- [20] J. Izard, C. D. C. Gomez Balderas, D. Ropers, S. Lacour, X. Song, Yi. Yang, A. B. Lindner, J. Geiselmann, and H. de Jong. A synthetic growth switch based on controlled expression of rna polymerase. *Mol. Syst. Biol.*, 11(11):840, 2015.
- [21] R. A. Khan. Approximation for the expectation of a function of the sample mean. *Statistics*, 38(2):117–122, 2004.
- [22] M. A. Lalwani, S. S. Ip, C. Carrasco-López, C. Day, E. M. Zhao, H. Kawabe, and J. L. Avalos. Optogenetic control of the lac operon for bacterial chemical and protein production. *Nat. Chem. Biol.*, 17(1):71–79, 2021.
- [23] Z. Li, X. Wang, and H. Zhang. Balancing the non-linear rosmarinic acid biosynthetic pathway by modular co-culture engineering. *Metab. Eng.*, 54:1–11, 2019.
- [24] Y. T. Lin and C. R. Doering. Gene expression dynamics with stochastic bursts: Construction and exact results for a coarse-grained model. *Phys. Rev. E*, 93:022409, Feb 2016.
- [25] D. Luenberger. An introduction to observers. *IEEE T. Automat. Contr.*, 16(6):596–602, 1971.
- [26] D. Lunz. On continuum approximations of discrete-state markov processes of large system size. *Multiscale Model. Sim.*, 19(1):294–319, 2021.
- [27] D. Lunz, G. Batt, J. Ruess, and J. F. Bonnans. Beyond the chemical master equation: Stochastic chemical kinetics coupled with auxiliary processes. *PLOS Comput. Biol.*, 17(7):1–24, 07 2021.

- [28] D. Lunz and J. F. Bonnans. Modelling and optimal control of a two-species bioproducing microbial consortium. *SIAM J. Appl. Math.*, 2022. Accepted. <https://hal.inria.fr/hal-03479385v2>.
- [29] D. Lunz, J. F. Bonnans, and J. Ruess. Revisiting moment-closure methods with heterogeneous multiscale population models. *Math. Biosci.*, 350:108866, 2022.
- [30] A. Miliás-Argeitis, M. Rullan, S. K. Aoki, P. Buchmann, and M. Khammash. Automated optogenetic feedback control for precise and robust regulation of gene expression and cell growth. *Nat. Commun.*, 7(1):1–11, 2016.
- [31] J. L. Morales and J. Nocedal. Remark on “Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization”. *ACM T. Math. Software*, 38(1):1–4, 2011.
- [32] L. B. Motta-Mena, A. Reade, M. J. Mallory, S. Glantz, O. D. Weiner, K. W. Lynch, and K. H. Gardner. An optogenetic gene expression system with rapid activation and deactivation kinetics. *Nat. Chem. Biol.*, 10(3):196–202, 2014.
- [33] E.-M. Nikolados, A. Y. Weiße, F. Ceroni, and D. A. Oyarzún. Growth defects and loss-of-function in synthetic gene circuits. *ACS Synth. Biol.*, 8(6):1231–1240, 2019.
- [34] J. Paijmans, M. Bosman, P. R. ten Wolde, and D. K. Lubensky. Discrete gene replication events drive coupling between the cell cycle and circadian clocks. *P. Natl. Acad. Sci. USA*, 113(15):4063–4068, 2016.
- [35] G. Peruzzi, M. Miotto, R. Maggio, G. Ruocco, and G. Gosti. Asymmetric binomial statistics explains organelle partitioning variance in cancer cell proliferation. *Commun. Phys.*, 4(1):1–12, 2021.
- [36] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The mathematical theory of optimal processes*. Gordon & Breach Science Publishers, New York, 1986. Reprint of the 1962 English translation.
- [37] M. J. D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Comput. J.*, 7(2):155–162, 1964.
- [38] K. M. Rapp, J. P. Jenkins, and M. J. Betenbaugh. Partners for life: building microbial consortia for the future. *Curr. Opin. Biotech.*, 66:292–300, 2020.
- [39] M. D. Rolfe, C. J. Rice, S. Lucchini, C. Pin, A. Thompson, A. D. S. Cameron, M. Alston, M. F. Stringer, R. P. Betts, J. Baranyi, M. W. Peck, and J. C. D. Hinton. Lag phase is a distinct growth phase that prepares bacteria for exponential growth and involves transient metal accumulation. *J. bacteriol.*, 194(3):686–701, 2012.
- [40] M. Rullan, D. Benzinger, G. W. Schmidt, A. Miliás-Argeitis, and M. Khammash. An optogenetic platform for real-time, single-cell interrogation of stochastic transcriptional regulation. *Mol. Cell*, 70(4):745–756, 2018.
- [41] I. Soifer, L. Robert, and A. Amir. Single-cell analysis of growth in budding yeast and bacteria reveals a common size regulation strategy. *Curr. Biol.*, 26(3):356–361, 2016.
- [42] S. L. Spencer, S. Gaudet, J. G. Albeck, J. M. Burke, and P. K. Sorger. Non-genetic origins of cell-to-cell variability in trail-induced apoptosis. *Nature*, 459(7245):428–432, 2009.
- [43] D. G. Spiller, C. D. Wood, D. A. Rand, and M. R. H. White. Measurement of single-cell dynamics. *Nature*, 465(7299):736–745, 2010.
- [44] P. Thomas. Intrinsic and extrinsic noise of gene expression in lineage trees. *Sci. Rep.*, 9(1):1–16, 2019.
- [45] P. Thomas and V. Shahrezaei. Coordination of gene expression noise with cell size: analytical results for agent-based models of growing cell populations. *J. Roy. Soc. Interface*, 18(178):20210274, 2021.
- [46] V. Thomée. Finite difference methods for linear parabolic equations. *Handbook of numerical analysis*, 1:5–196, 1990.
- [47] J. E. Toettcher, D. Gong, W. A. Lim, and O. D. Weiner. Light-based feedback for controlling intracellular signaling dynamics. *Nat. Methods*, 8(10):837–839, 2011.
- [48] R. Vinter. *Optimal control*. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, 2010.
- [49] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1):25–57, March 2006.

- [50] E. Weill, V. Andréani, C. Aditya, P. Martinon, J. Ruess, G. Batt, and F. Bonnans. Optimal control of an artificial microbial differentiation system for protein bioproduction. In *18th European Control Conference (ECC)*, pages 2663–2668. IEEE, 2019.
- [51] A. Y. Weiße, D. A. Oyarzún, V. Danos, and P. S. Swain. Mechanistic links between cellular trade-offs, gene expression, and growth. *P. Natl. Acad. Sci. USA*, 112(9):E1038–E1047, 2015.
- [52] E. M. Zhao, Y. Zhang, J. Mehl, H. Park, M. A. Lalwani, J. E. Toettcher, and J. L. Avalos. Optogenetic regulation of engineered cellular metabolism for microbial chemical production. *Nature*, 555(7698):683–687, 2018.
- [53] K. Zhou, K. Qiao, S. Edgar, and G. Stephanopoulos. Distributing a metabolic pathway among a microbial consortium enhances production of natural products. *Nat. Biotechnol.*, 33(4):377–383, 2015.

Appendix

A Optimal control of the PDE

In this appendix, we set forth the adjoint approach for optimal control of PDEs in a general abstract setting (Appendix A.1). We then demonstrate how it applies to the modular class of models introduced in section 2.2 (Appendix A.2). This is followed by the discretisation of the model components (Appendix A.3) and the calculation of the discrete gradient for numerical optimisation (Appendix A.4).

A.1 The abstract setting

First, we generalise the optimal control problem (13) in an abstract Banach setting, writing

$$\max_{u \in K_U} J(u, p), \quad (83)$$

such that

$$A(u, p) = 0, \quad (84)$$

where $J : U \times V \rightarrow \mathbb{R}$ and $A : V \times U \rightarrow W$ are both C^1 for Banach spaces U , V , and W , with $u \in K_U \subseteq U$ and $p \in K_V \subseteq V$. This notation preserves the control as u , the state as p , and the objective as J , with the operator A representing the state constraints, that is, the state equation as well as relevant boundary and initial conditions.

Let (u_0, p_0) be a zero of A , that is, $A(u_0, p_0) = 0$. Further, assume that the derivative with respect to the state p , which we denote $D_p A(u_0, p_0)$, a continuous linear map from $V \rightarrow W$, is invertible. It follows, by the implicit function theorem, that there exist open neighborhoods \mathcal{U} of u_0 and \mathcal{P} of p_0 , and a C^1 function $\phi : U \rightarrow P$, such that, any $(u, p) \in \mathcal{U} \times \mathcal{P}$ is a zero of A iff $p = \phi(u)$.

Differentiating the state constraint for any $u \in \mathcal{U}$, we see, by the chain rule, that

$$D_p A(u, \phi(u)) D_u \phi(u) + D_u A(u, \phi(u)) = 0. \quad (85)$$

Since the set of invertible mappings is open, for u sufficiently close to u_0 , $D_p A(u, \phi(u))$ is close to $D_p A(u_0, \phi(u_0))$ and thus invertible. It then follows that

$$D_u \phi(u) = -[D_p A(u, \phi(u))]^{-1} D_u A(u, \phi(u)). \quad (86)$$

We introduce the reduced pay-off

$$\hat{J}(u) = J(u, \phi(u)). \quad (87)$$

The derivative of the reduced pay-off is the primary object of interest, giving the first-order optimality conditions. By application of the chain rule, we determine that the derivative in the direction $v \in U$ satisfies

$$D_u \hat{J}(u)v = \langle D_u J(u, \phi(u)), v \rangle + \langle D_p J(u, \phi(u)), D_u \phi(u)v \rangle = \left\langle D_u J(u, \phi(u)) + \{D_u \phi(u)\}^\top D_p J(u, \phi(u)), v \right\rangle. \quad (88)$$

With the aim of replacing the term involving the difficult-to-compute derivative $D_u \phi(u)$, we use (86) to obtain

$$\{D_u \phi(u)\}^\top D_p J(u, \phi(u)) = - \left\{ [D_p A(u, \phi(u))]^{-1} D_u A(u, \phi(u)) \right\}^\top D_p J(u, \phi(u)) = [D_u A(u, \phi(u))]^\top \lambda, \quad (89)$$

for the costate λ that solves the costate equation, namely,

$$- [D_p A(u, \phi(u))]^\top \lambda = D_p J(u, \phi(u)). \quad (90)$$

Combining (88) and (89), it follows that the derivative of the reduced pay-off takes the form

$$D_u \hat{J}(u) = D_u J(u, \phi(u)) + [D_u A(u, \phi(u))]^\top \lambda. \quad (91)$$

The optimal control has a vanishing pay-off derivative. Therefore, the first-order optimality conditions comprise the state equation (84), the costate equation (90), and a vanishing pay-off derivative $D_u \hat{J}(u) = 0$ given in (91).

The aim of this section was to describe the adjoint approach in a totally generic setting, to demonstrate that the technique is not limited to the concrete model class we study for bioproduction applications. Nevertheless, it is instructive to outline the formulation for the class of models described in section 2.2, as we do in the next section.

A.2 The modular model class

Here, we rewrite the optimal control problem (13) to make explicit the terms nonlinear in state, which stem only from the turbidostat dilution rate (7). As pointed out in section 2.3.2, in both the concentration and copy-number paradigms, the rate of change of $\int_{\mathbb{R}_+^d} p_k(\mathbf{x}, t) d\mathbf{x}$ driven by population growth is of the form $\int_{\mathbb{R}_+^d} g_k(\mathbf{x}, t) p_k(\mathbf{x}, t) d\mathbf{x}$ (see Appendix B.1 for the calculation in the copy-number case). Therefore, in both cases, the turbidostat dilution rate (7) is given by

$$\Lambda(t) = \int_{\mathbb{R}_+^d} \mathbf{g}(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x}, \quad (92)$$

where we have introduced the vector of growth rates $\mathbf{g} = (g_1, \dots, g_d)^\top$ with explicit dependence on the control. We may thus write the operator in (13a) as a linear operator combined with the nonlinear turbidostat contribution,

$$\mathbf{M}(\mathbf{p}(\cdot, t), t, u(t)) \mathbf{p}(\mathbf{x}, t) = \mathbf{M}_0(t, u(t)) \mathbf{p}(\mathbf{x}, t) - \left(\int_{\mathbb{R}_+^d} \mathbf{g}(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x} \right) \mathbf{p}(\mathbf{x}, t). \quad (93a)$$

For the objective (13b), we assume that the terminal pay-off integrand ψ is a linear function of the state, and the running pay-off integrand φ combines a part not dependent on the state, a part linear in the state, and a product of the turbidostat dilution rate by a linear function of the state. This takes the form

$$\begin{aligned} J_{\text{PDE}} &= \int_{\mathbb{R}_+^d} \psi(\mathbf{x})^\top \mathbf{p}(\mathbf{x}, T) d\mathbf{x} + \int_0^T \varphi_u(t, u(t)) dt + \int_0^T \int_{\mathbb{R}_+^d} \varphi_0(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x} dt \\ &\quad + \int_0^T \left(\int_{\mathbb{R}_+^d} \mathbf{g}(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x} \right) \int_{\mathbb{R}_+^d} \varphi_\Lambda(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x} dt. \end{aligned} \quad (93b)$$

In batch mode, there are no nonlinear contributions, and these may be discarded in the above formulation.

Given this concrete setting, it is intuitive to derive the adjoint approach in a Lagrangian framework, as follows. First, we introduce the Lagrangian

$$\mathcal{L} = J_{\text{PDE}} - \int_0^T \left\langle \boldsymbol{\lambda}(t), \frac{\partial \mathbf{p}}{\partial t}(t) - \mathbf{M}(\mathbf{p}(\cdot, t), t, u(t)) \mathbf{p}(t) \right\rangle dt - \langle \boldsymbol{\lambda}(0), \mathbf{p}(0) - \mathbf{p}_0 \rangle, \quad (94)$$

where $\boldsymbol{\lambda}(\mathbf{x}, t)$ denotes the costate and adopting the L^2 inner-product notation (suppressing state dependence)

$$\langle \mathbf{f}(t), \mathbf{g}(t) \rangle := \int_{\mathbb{R}_+^d} \mathbf{f}(\mathbf{x}, t) \cdot \mathbf{g}(\mathbf{x}, t) d\mathbf{x}. \quad (95)$$

After integrating by parts in time we rewrite the Lagrangian as

$$\mathcal{L} = J_{\text{PDE}} - \langle \boldsymbol{\lambda}(T), \mathbf{p}(T) \rangle + \langle \boldsymbol{\lambda}(0), \mathbf{p}_0 \rangle + \int_0^T \left\langle \frac{\partial \boldsymbol{\lambda}}{\partial t}(t) + \mathbf{M}(\mathbf{p}(\cdot, t), t, u(t))^\top \boldsymbol{\lambda}(t), \mathbf{p}(t) \right\rangle dt. \quad (96)$$

We now seek the costate $\boldsymbol{\lambda}(\mathbf{x}, t)$ guaranteeing a vanishing first variation of the Lagrangian with respect to the state $\mathbf{p}(\mathbf{x}, t)$. This will allow us to calculate the derivative of the pay-off without having to compute how the derivative of the state $\mathbf{p}(\mathbf{x}, t)$ with respect to the control $u(t)$. Ultimately, this provides the costate equation, which takes the form

$$\begin{aligned} \frac{\partial \boldsymbol{\lambda}}{\partial t}(\mathbf{x}, t) &= -\mathbf{M}(\mathbf{p}(\cdot, t), t, u(t))^\top \boldsymbol{\lambda}(\mathbf{x}, t) + \mathbf{g}(\mathbf{x}, t) \langle \boldsymbol{\lambda}(t), \mathbf{p}(t) \rangle - \varphi_0(\mathbf{x}, t, u(t)) \\ &\quad - \varphi_\Lambda(\mathbf{x}, t, u(t)) \int_{\mathbb{R}_+^d} \mathbf{g}(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x} - \mathbf{g}(\mathbf{x}, t, u(t)) \int_{\mathbb{R}_+^d} \varphi_\Lambda(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x}, \end{aligned} \quad (97a)$$

$$\boldsymbol{\lambda}(\mathbf{x}, T) = \psi(\mathbf{x}). \quad (97b)$$

Substituting the form (93a) of the operator \mathbf{M} into dynamics (97a), we obtain the factorised form

$$\begin{aligned} \frac{\partial \boldsymbol{\lambda}}{\partial t}(\mathbf{x}, t) &= -\mathbf{M}_0(t, u(t))^\top \boldsymbol{\lambda}(\mathbf{x}, t) + [\boldsymbol{\lambda}(\mathbf{x}, t) - \varphi_\Lambda(\mathbf{x}, t, u(t))] \int_{\mathbb{R}_+^d} \mathbf{g}(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x} \\ &\quad + \mathbf{g}(\mathbf{x}, t) \langle \boldsymbol{\lambda}(t), \mathbf{p}(t) \rangle - \varphi_0(\mathbf{x}, t, u(t)) - \mathbf{g}(\mathbf{x}, t, u(t)) \int_{\mathbb{R}_+^d} \varphi_\Lambda(\mathbf{x}, t, u(t))^\top \mathbf{p}(\mathbf{x}, t) d\mathbf{x}. \end{aligned} \quad (98)$$

The costate ensures a vanishing Lagrangian derivative with respect to the state, therefore, after application of the chain rule, we deduce that the derivative of the objective with respect to the control takes the form

$$\delta_u J_{\text{PDE}}(u(t), \mathbf{p}(\mathbf{x}, t; u(t))) = \delta_u J_{\text{PDE}}(u(t), \mathbf{p}(\mathbf{x}, t)) + \int_0^T \langle \boldsymbol{\lambda}(t), \delta_u \mathbf{M}(\mathbf{p}(\cdot, t), u(t)) \mathbf{p}(t) \rangle dt. \quad (99)$$

The form (99) is clearly an instance of its abstract analogue (91). The first-order optimality condition is a vanishing first variation with respect to the control $u(t)$.

This encapsulates the adjoint approach: the state equation (13a), costate equation (97), and the vanishing objective derivative (99) characterise optimal controls. A discrete analogue of the adjoint derivation allows the problem to be solved numerically, as we now detail.

A.3 PDE discretisation

In this section, we detail the finite-difference discretisation of each component of the operator in problem (13) used in the examples. We truncate and discretise the state space \mathbb{R}_+^d to give the uniform grid $\{0, N_1 \Delta x\} \times \dots \times \{0, N_d \Delta x\}$, where Δx is the state resolution. We index the numerical scheme by integer vectors, where $\mathbf{x}_{\mathbf{J}} := \mathbf{J} \Delta \mathbf{x} = (J_1 \Delta x, \dots, J_d \Delta x)$, and write the discrete probability mass as $\mathbf{p}_{\mathbf{J}}^n = \mathbf{p}(\mathbf{x}_{\mathbf{J}}, t^n)$, for a sequence of time steps $\{t^1, \dots, t^{N+1}\}$. We similarly use the \mathbf{J} subscript and n superscript to denote evaluation of any function at the point $(\mathbf{x}_{\mathbf{J}}, t^n)$. Since the PDE components in section 2 are written in terms of their scalar components $p_k(\mathbf{x}, t)$, for any $k = 1, \dots, \mathcal{K}$, here too we write the scheme for the scalar component $(p_k)_{\mathbf{J}}^n$ while dropping the subscript k , where possible, to simplify the notation. Similarly, we drop the subscript k from the associated rate functions. We write the scheme in explicit form, which is what we used for the numerical results in this work, calling the n th time step $\Delta t^n := t^{n+1} - t^n$.

A.3.1 Population growth

The growth law in the concentration paradigm (2) takes the form

$$\frac{p_{\mathbf{J}}^{n+1} - p_{\mathbf{J}}^n}{\Delta t^n} = g_{\mathbf{J}}^n p_{\mathbf{J}}^n + (\text{Non-growth sources}). \quad (100)$$

For the copy-number paradigm (3), consider any discrete state within the grid $\mathbf{x}_{\mathbf{J}}$ dividing into two cells with states of fractions $c \in (0, 1)$ and $1 - c$ of the original (see Appendix B for a further discussion of this particular division kernel). The daughter cell of fraction c will have the state $c\mathbf{x}$, which lies in the lattice rectangle with opposite corners at $\mathbf{x}_{\lfloor c\mathbf{J} \rfloor}$ and $\mathbf{x}_{\lfloor c\mathbf{J} \rfloor + \mathbf{1}}$, where $\lfloor c\mathbf{J} \rfloor$ denotes the vector $(\lfloor cJ_1 \rfloor, \dots, \lfloor cJ_d \rfloor)$ and $\mathbf{1}$ denotes a d -dimensional vector of ones. We split the density between the 2^d nodes at the corners of the rectangle via bilinear interpolation. For the sake of concreteness, we present the discretisation in dimension $d = 2$, noting that the approach generalises to arbitrary dimension. Writing $\mathbf{I} = (I, J)$, this takes the form

$$\Delta p_{I,J} = -P_{I,J} + (\text{other sources}), \quad (101a)$$

$$\Delta p_{\lfloor cI \rfloor, \lfloor cJ \rfloor} = (1 - \psi_{cI})(1 - \psi_{cJ})P_{I,J} + (\text{other sources}), \quad (101b)$$

$$\Delta p_{\lfloor cI \rfloor + 1, \lfloor cJ \rfloor} = \psi_{cI}(1 - \psi_{cJ})P_{I,J} + (\text{other sources}), \quad (101c)$$

$$\Delta p_{\lfloor cI \rfloor, \lfloor cJ \rfloor + 1} = (1 - \psi_{cI})\psi_{cJ}P_{I,J} + (\text{other sources}), \quad (101d)$$

$$\Delta p_{\lfloor cI \rfloor + 1, \lfloor cJ \rfloor + 1} = \psi_{cI}\psi_{cJ}P_{I,J} + (\text{other sources}), \quad (101e)$$

where the parenthetical (other sources) includes any change in population density not resulting from the division of cells of state $\mathbf{x}_{\mathbf{J}}$ into a daughter of fraction c , and adopting the shorthands

$$\Delta p_{\mathbf{J}} = \frac{p_{\mathbf{J}}^{n+1} - p_{\mathbf{J}}^n}{\Delta t^n}, \quad P_{\mathbf{J}} = g_{\mathbf{J}}^n p_{\mathbf{J}}^n, \quad \psi_z = z - \lfloor z \rfloor. \quad (101f)$$

The forms (101b) to (101e) are duplicated for $1 - c$. Summing these contributions, and neglecting the other sources, we recover the (discretised) contributions to the rates of change of the zeroth and first moments (the

population density and total copy numbers):

$$\begin{aligned}
\sum_{I,J} \Delta p_{I,J} &= \sum_{I,J} P_{I,J} = \sum_{I,J} g(y_J) p(x_I, y_J, t), \tag{102} \\
\sum_{I,J} x_I \Delta p_{I,J} &= -\Delta x \sum_{I,J} I P_{I,J} + \Delta x \sum_{I,J} \sum_{c,1-c} \left\{ [cI] (1 - \psi_{cI})(1 - \psi_{cJ}) + ([cI] + 1) \psi_{cI}(1 - \psi_{cJ}) \right. \\
&\quad \left. + [cI] (1 - \psi_{cI}) \psi_{cJ} + ([cI] + 1) \psi_{cI} \psi_{cJ} \right\} P_{I,J} \\
&= \Delta x \sum_{I,J} \left\{ -I + [cI] + \psi_{cI} + [(1-c)I] + \psi_{(1-c)I} \right\} P_{I,J} \\
&= 0, \tag{103}
\end{aligned}$$

and similarly $\sum_{I,J} y_J \Delta p_{I,J} = 0$. These are the discrete analogues of the continuous (18) and (19).

The above discretisations preserve the PDE and ODE overall growth rate structure:

$$\sum_{\mathbf{J}} \frac{p_{\mathbf{J}}^{n+1} - p_{\mathbf{J}}^n}{\Delta t^n} = \sum_{\mathbf{J}} g_{\mathbf{J}}^n p_{\mathbf{J}}^n + (\text{Non-growth sources}). \tag{104}$$

A.3.2 Bioreactor removal

The turbidostat contribution (5) is discretised as in (100), with the dilution rate (7) approximated by the discrete form

$$\Lambda(t) \approx \sum_{k=1}^{\mathcal{K}} \sum_{\mathbf{J}} (g_k^n)_{\mathbf{J}}^n (p_k^n)_{\mathbf{J}}^n. \tag{105}$$

Subscripts k have been momentarily reintroduced in order to sum over all components. Combining (104) and (105), we find that this discretisation preserves the turbidostat feature of constant population mass: $\sum_{k=1}^{\mathcal{K}} \sum_{\mathbf{J}} (p_k^n)_{\mathbf{J}}^n$ remains constant for all time points n .

A.3.3 Discrete dynamics

The discrete dynamics (8) are discretised identically to the growth terms in (100), whereby the conservation property of the discrete dynamics is preserved: summing over k we find that all discrete terms cancel.

A.3.4 Continuum dynamics

The Fokker–Planck contributions (9) are, in divergence form (see Appendix B.2 for a continuum analogue), discretised as:

$$\frac{p_{\mathbf{J}}^{n+1} - p_{\mathbf{J}}^n}{\Delta t^n} = -\frac{\Phi_{\mathbf{J}+e}^n - \Phi_{\mathbf{J}}^n}{\Delta x} + (\text{other sources}), \tag{106a}$$

where the numerical flux is defined as

$$\Phi_{\mathbf{J}}^n := r_{\mathbf{J}-e}^n p_{\mathbf{J}-e}^n - \frac{1}{2\Omega} \frac{r_{\mathbf{J}}^n p_{\mathbf{J}}^n - r_{\mathbf{J}-e}^n p_{\mathbf{J}-e}^n}{\Delta x}. \tag{106b}$$

Expanding the divergence form, we readily see that the scheme is consistent with the continuum form, with central differences in the second-order term and an upwinded stencil (the rates are non-negative) in the first-order term:

$$\frac{p_{\mathbf{J}}^{n+1} - p_{\mathbf{J}}^n}{\Delta t^n} = -\frac{r_{\mathbf{J}}^n p_{\mathbf{J}}^n - r_{\mathbf{J}-e}^n p_{\mathbf{J}-e}^n}{\Delta x} + \frac{1}{2\Omega} \frac{r_{\mathbf{J}+e}^n p_{\mathbf{J}+e}^n - 2r_{\mathbf{J}}^n p_{\mathbf{J}}^n + r_{\mathbf{J}-e}^n p_{\mathbf{J}-e}^n}{(\Delta x)^2} + (\text{other sources}), \tag{107}$$

which preserves the conservation property that $X_{\cdot,k}^0(t)$ is not changed by the Fokker–Planck contributions.

The zero-flux boundary conditions (10) are enforced by setting to zero all fluxes (106b) whose support is beyond the lattice.

By summing over all state indices \mathbf{J} , it is evident from the divergence form of the discretisation (106) and the aforementioned boundary conditions that the continuum terms on the right-hand side vanish. Therefore, the discretisation preserves the feature that $\sum_{\mathbf{J}} p_{\mathbf{J}}^n$ is not altered by the continuum Fokker–Planck dynamics.

For desirable properties including numerical stability [46], we seek to ensure that the scheme (107) is monotone. A scheme is monotone if, when writing $p_{\mathbf{J}}^{n+1}$ as a linear combination of the state variables $\sum_{\mathbf{J}} \alpha_{\mathbf{J}} p_{\mathbf{J}}^n$, the coefficients $\alpha_{\mathbf{J}} \geq 0$ are non-negative. Rearranging (107) we see that

$$p_{\mathbf{J}}^{n+1} = \left[1 - \Delta t^n \left(\frac{1}{\Delta x} + \frac{r_{\mathbf{J}}^n}{\Omega(\Delta x)^2} \right) \right] p_{\mathbf{J}}^n + \left[\frac{\Delta t^n r_{\mathbf{J}+\mathbf{e}}^n}{2\Omega(\Delta x)^2} \right] p_{\mathbf{J}+\mathbf{e}}^n + \left[\frac{\Delta t^n}{\Delta x} + \frac{\Delta t^n r_{\mathbf{J}-\mathbf{e}}^n}{2\Omega(\Delta x)^2} \right] p_{\mathbf{J}-\mathbf{e}}^n + (\text{other sources}). \quad (108)$$

The coefficients of the off-diagonal terms in (108) are strictly positive, and the coefficient of all diagonal term will be non-negative if and only if

$$\Delta t^n \leq \min_{\mathbf{J}} \left(\frac{1}{\Delta x} + \frac{r_{\mathbf{J}}^n}{\Omega(\Delta x)^2} \right)^{-1}. \quad (109)$$

In practice, when we combine the various contributions, including multiple continuum reactions of Fokker–Planck form as well as terms stemming from other dynamics, we must ensure the scheme is monotone. We write the combined discrete system in vector form as

$$\frac{\mathbf{p}^{n+1} - \mathbf{p}^n}{\Delta t^n} = \mathbf{M}^n \mathbf{p}^n, \quad \mathbf{p}^0 = \mathbf{p}_0, \quad (110)$$

for all $n = 1, \dots, N$, where we have placed all entries of $(p_k)_{\mathbf{J}}^n$ (for all $k = 1, \dots, \mathcal{K}$ and all \mathbf{J} on the lattice) into a vector. We emphasise that the operator \mathbf{M}^n depends, in general, on the state \mathbf{p}^n and a discrete control u^n .

To ensure that the diagonal terms have non-negative coefficients, we require

$$\Delta t^n \leq \frac{1}{-\min \text{diag}(\mathbf{M}^n)}, \quad (111)$$

where diag represents the vector of diagonal entries in the matrix operator \mathbf{M}^n , and assuming that the minimal diagonal entry is negative (otherwise the scheme is monotone regardless of the time step Δt^n).

A.4 Optimal control of the discretised PDE

The aim of this section is to show how we calculate the derivative of the objective function with respect to the control using the adjoint method. We introduce the discrete Lagrangian

$$\mathcal{L} = J - \sum_{n=0}^N (\boldsymbol{\lambda}^{n+1})^\top [\mathbf{p}^{n+1} - (\mathbf{I} + \Delta t^n \mathbf{M}^n) \mathbf{p}^n] - (\boldsymbol{\lambda}^0)^\top [\mathbf{p}^0 - \mathbf{p}_0], \quad (112)$$

from which we extract the costate equation and terminal conditions by ensuring that $\partial \mathcal{L} / \partial p_\ell^n = 0$ for all $n = 0, \dots, N+1$ and vector elements ℓ , which takes the vector form

$$\mathbf{0} = \left(\frac{\partial J}{\partial \mathbf{p}^n} \right)^\top - \boldsymbol{\lambda}^n + [\mathbf{I} + \Delta t^n (\mathbf{M}^n)^\top] \boldsymbol{\lambda}^{n+1} + \Delta t^n (\mathbf{p}^n)^\top \frac{\partial (\mathbf{M}^n)^\top}{\partial \mathbf{p}^n} \boldsymbol{\lambda}^{n+1}, \quad \boldsymbol{\lambda}^{N+1} = \left(\frac{\partial J}{\partial \mathbf{p}^{N+1}} \right)^\top. \quad (113)$$

The discrete costate equation (113) is identifiable as consistent with its continuum analogue (97).

The discrete adjoint approach proceeds as follows. For a given discrete control $\mathbf{u} = (u^0, \dots, u^N)$, we solve the PDE model (110) forwards to give \mathbf{p}^n for all $n = 0, \dots, N+1$. Then we solve the costate equation (113) backwards to give the costate $\boldsymbol{\lambda}^n$ for all $n = 0, \dots, N+1$. Finally, the gradient of the objective with respect to the control is given by

$$\nabla_{\mathbf{u}} J = \frac{\partial \mathcal{L}}{\partial \mathbf{u}} = \frac{\partial J}{\partial \mathbf{u}} + \Delta t^n \sum_{n=0}^N (\boldsymbol{\lambda}^{n+1})^\top \frac{\partial \mathbf{M}^n}{\partial \mathbf{u}} \mathbf{p}^n. \quad (114)$$

The discrete gradient (114) is analogous to its continuum counterpart (99), with $\partial J / \partial \mathbf{u} = 0$ in this work.

Thus the discrete adjoint formulation provides us with gradient information that may be supplied to standard gradient-based numerical optimisation routines. In this work, we employ the L-BFGS-B optimisation algorithm, a limited-memory, control-bounded implementation of the BFGS routine [31]. Ultimately, this provides a numerical recipe to converge towards local control optima.

B Details of the moment closure calculations

A more extensive treatment of the moments (of arbitrary order) is provided in Ref. [29]. We reproduce here the results pertinent to zeroth and first orders.

B.1 Copy-number population growth

In the copy-number paradigm, we multiply (3) by x_j^i and integrate over the state space to obtain

$$\frac{d}{dt} X_{j,k}^i(t) = -g_k(\mathbf{X}_k(t), t) X_{j,k}^i(t) + 2 \int_{\mathbb{R}_+^d} x_j^i \int_{\hat{\mathbf{x}} \geq \mathbf{x}} g_k(\hat{\mathbf{x}}, t) p_k(\hat{\mathbf{x}}, t) \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\hat{\mathbf{x}} d\mathbf{x} + (\text{Non-growth sources}), \quad (115)$$

where, changing the order of integration, we see that

$$\int_{\mathbb{R}_+^d} x_j^i \int_{\hat{\mathbf{x}} \geq \mathbf{x}} g_k(\hat{\mathbf{x}}, t) p_k(\hat{\mathbf{x}}, t) \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\hat{\mathbf{x}} d\mathbf{x} = \int_{\mathbb{R}_+^d} g_k(\hat{\mathbf{x}}, t) p_k(\hat{\mathbf{x}}, t) \int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} x_j^i \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} d\hat{\mathbf{x}}. \quad (116)$$

This motivates a study of the integral $\int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} x_j^i \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x}$ leveraging assumptions (4). When $i = 0$, the integral is unity. When $i = 1$, we obtain the first-order relation

$$\begin{aligned} \int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} x_j \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} &= \int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} x_j \theta_k(\hat{\mathbf{x}} - \mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} = \int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} (\hat{x}_j - x_j) \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} \\ &= \hat{x}_j - \int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} x_j \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x}, \end{aligned} \quad (117)$$

from which it follows that

$$\int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} x_j \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} = \frac{\hat{x}_j}{2}. \quad (118)$$

The forms (18) and (19) follow.

We briefly derive a simpler representation for a particular case employed in this study. In the copy-number paradigm, we focus on cell division that preserves a fixed copy-number ratio between mother and daughter cells for all continuum species. That is, when a mother cell in state $\hat{\mathbf{x}}$ divides into two daughters with states \mathbf{x} and $\hat{\mathbf{x}} - \mathbf{x}$, we consider the particular case where $\mathbf{x} = c\hat{\mathbf{x}}$ (and thus $\hat{\mathbf{x}} - \mathbf{x} = (1 - c)\hat{\mathbf{x}}$), and the daughters inherit copy number fractions $c \in [0, 1]$ and $1 - c$ of the mother cell's copy number for each species. This is realised with θ of the form

$$\theta_k(\mathbf{x} | \hat{\mathbf{x}}) = \frac{1}{\hat{x}_1} \hat{\theta}_k(x_1 | \hat{x}_1) \prod_{j=2}^d \frac{1}{\hat{x}_j} \delta\left(\frac{x_j}{\hat{x}_j} - \frac{x_1}{\hat{x}_1}\right), \quad (119)$$

where δ is a Dirac delta function and $\hat{\theta}$ is a univariate probability density. The class of cell-division laws in (119) allows us to rewrite the high-dimensional copy-number contribution (3) via the change of variables

$$\hat{z}_1 = \frac{\hat{x}_1}{x_1}, \quad \hat{z}_j = \frac{\hat{x}_j}{x_j} - \frac{\hat{x}_1}{x_1}, \quad (120)$$

for $j = 2, \dots, d$, whereupon we obtain the one-dimensional representation

$$\int_{\hat{\mathbf{x}} \geq \mathbf{x}} g_k(\hat{\mathbf{x}}, t) p_k(\hat{\mathbf{x}}, t) \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\hat{\mathbf{x}} = \int_1^\infty g_k(z\mathbf{x}, t) p_k(z\mathbf{x}, t) z^{d-2} \hat{\theta}_k(x_1 | zx_1) dz. \quad (121)$$

Moreover, we may express conditions (4) as conditions on $\hat{\theta}$. The symmetry condition is equivalent to symmetry in $\hat{\theta}$, namely,

$$\hat{\theta}_k(x | \hat{\mathbf{x}}) = \hat{\theta}_k(\hat{\mathbf{x}} - x | \hat{\mathbf{x}}). \quad (122)$$

Via the change of variables

$$z_1 = \frac{x_1}{\hat{x}_1}, \quad z_j = \frac{x_j}{\hat{x}_j} - \frac{x_1}{\hat{x}_1}, \quad (123)$$

for $j = 2, \dots, d$, we obtain the density condition

$$1 = \int_{\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}}} \theta_k(\mathbf{x} | \hat{\mathbf{x}}) d\mathbf{x} = \int_{z^{-1}(\mathbf{0} \leq \mathbf{x} \leq \hat{\mathbf{x}})} \theta_k(z_1 \hat{x}_1 | \hat{x}_1) \prod_{j=2}^d \delta(z_j) dz = \int_0^1 \hat{\theta}_k(z \hat{x}_1 | \hat{x}_1) dz. \quad (124)$$

Condition (124) demonstrates the sense in which $\hat{\theta}$ is a probability density. The transformed conditions (122) and (124) hold for all $d \geq 1$.

B.2 Continuum dynamics

The contributions stemming from the Fokker–Planck terms (9) can be derived by expressing the operator in divergence form, namely,

$$\frac{\partial}{\partial t} p_k(\mathbf{x}, t) = -\nabla \cdot \Phi_k(\mathbf{x}, t) + (\text{other sources}), \quad (125)$$

for the flux $\Phi_k(\mathbf{x}, t)$ given by

$$\Phi_k(\mathbf{x}, t) = \mathbf{e}_k \left\{ r_k(\mathbf{x}, t) p_k(\mathbf{x}, t) - \frac{1}{2\Omega} \mathbf{e}_k \cdot \nabla [r_k(\mathbf{x}, t) p_k(\mathbf{x}, t)] \right\}. \quad (126)$$

By application of the divergence theorem, and using the zero normal flux boundary condition (10), it follows that there are no zeroth-order contributions ($i = 0$). This reflects the fact that the process describes a change in the continuum species, but not the population mass. By application of Green's identities, we find, for $i = 1$, that

$$\begin{aligned} - \int_{\mathbb{R}_+^d} x_j \nabla \cdot \Phi_k(\mathbf{x}, t) d\mathbf{x} &= \int_{\mathbb{R}_+^d} \nabla x_j \cdot \Phi_k(\mathbf{x}, t) d\mathbf{x} = \int_{\mathbb{R}_+^d} e_{k,j} \left\{ r_k(\mathbf{x}, t) p_k(\mathbf{x}, t) - \frac{1}{2\Omega} \mathbf{e}_k \cdot \nabla [r_k(\mathbf{x}, t) p_k(\mathbf{x}, t)] \right\} d\mathbf{x} \\ &\approx e_{k,j} r_k(\mathbf{X}_k(t), t) X_{\cdot,k}^0(\mathbf{x}, t), \end{aligned} \quad (127)$$

where $e_{k,j}$ denotes the j th element of the vector \mathbf{e}_k . The approximation in (127), in addition to the closure technique of approximating the nonlinear functions at the mean, stems from ignoring terms of order $\mathcal{O}(1/\Omega)$. Since these terms admit a divergence form, they may be written as boundary terms, and thus provide a way to verify that the boundary contributions (multiplied by $1/\Omega$) are negligible (if we can solve or approximate the full solution, or if we have good reason to expect the distribution to be supported away from the boundary). It follows that the first-order contributions are approximated by (23).

The bursty production (11) similarly has no zeroth-order contribution. To see this, note that the zeroth-order terms sum all possible jumps — for each \mathbf{x} add all the valid jumps *to* \mathbf{x} — which may be expressed from the dual perspective — for each \mathbf{x} add all the valid jumps *from* \mathbf{x} . Then, utilising (12), we find that

$$\begin{aligned} \|\mathbf{e}_k\| \int_{\mathbb{R}_+^d} \int_{\mathbf{x} - z\mathbf{e}_k \in \mathbb{R}_+^d} f_k(\mathbf{x} - z\mathbf{e}_k, t) p_k(\mathbf{x} - z\mathbf{e}_k, t) Q_k(\mathbf{x} - z\mathbf{e}_k, z \|\mathbf{e}_k\|) dz d\mathbf{x} \\ = \int_{\mathbb{R}_+^d} f_k(\mathbf{x}, t) p_k(\mathbf{x}, t) \int_0^\infty \|\mathbf{e}_k\| Q_k(\mathbf{x}, z \|\mathbf{e}_k\|) dz d\mathbf{x} \\ = \int_{\mathbb{R}_+^d} f_k(\mathbf{x}, t) p_k(\mathbf{x}, t) d\mathbf{x}. \end{aligned} \quad (128)$$

At first order, we again appeal to the dual perspective to evaluate the integral:

$$\begin{aligned} \int_{\mathbb{R}_+^d} x_j \|\mathbf{e}_k\| \int_{\mathbf{x} - z\mathbf{e}_k \in \mathbb{R}_+^d} f_k(\mathbf{x} - z\mathbf{e}_k, t) p_k(\mathbf{x} - z\mathbf{e}_k, t) Q_k(\mathbf{x} - z\mathbf{e}_k, z \|\mathbf{e}_k\|) dz d\mathbf{x} \\ = \int_{\mathbb{R}_+^d} \int_{\mathbf{x} - z\mathbf{e}_k \in \mathbb{R}_+^d} (x_j - z e_{k,j} + z e_{k,j}) \|\mathbf{e}_k\| f_k(\mathbf{x} - z\mathbf{e}_k, t) p_k(\mathbf{x} - z\mathbf{e}_k, t) Q_k(\mathbf{x} - z\mathbf{e}_k, z \|\mathbf{e}_k\|) dz d\mathbf{x} \\ = \int_{\mathbb{R}_+^d} x_j f_k(\mathbf{x}, t) p_k(\mathbf{x}, t) d\mathbf{x} + e_{k,j} \int_{\mathbb{R}_+^d} f_k(\mathbf{x}, t) p_k(\mathbf{x}, t) \int_0^\infty z \|\mathbf{e}_k\| Q_k(\mathbf{x}, z \|\mathbf{e}_k\|) dz d\mathbf{x} \\ \approx f_k(\mathbf{X}_k(t), t) X_{j,k}^1(t) + \frac{e_{k,j}}{\|\mathbf{e}_k\|} f_k(\mathbf{X}_k(t), t) Q_1(\mathbf{X}_k(t)) X_{\cdot,k}^0(t), \end{aligned} \quad (129)$$

where we have defined

$$Q_{k,1}(\mathbf{x}) := \int_0^\infty y Q_k(\mathbf{x}, y) dy. \quad (130)$$

The result (24) follows.

C Singular arcs in the single-species ODE model

Our goal in this appendix is to study the function $g'(x)x + g(x)$, where $g(x)$ is defined in (29d) which we reproduce here

$$g(x) = g_\infty + \frac{g_0 - g_\infty}{1 + (x/k)^n}. \quad (131)$$

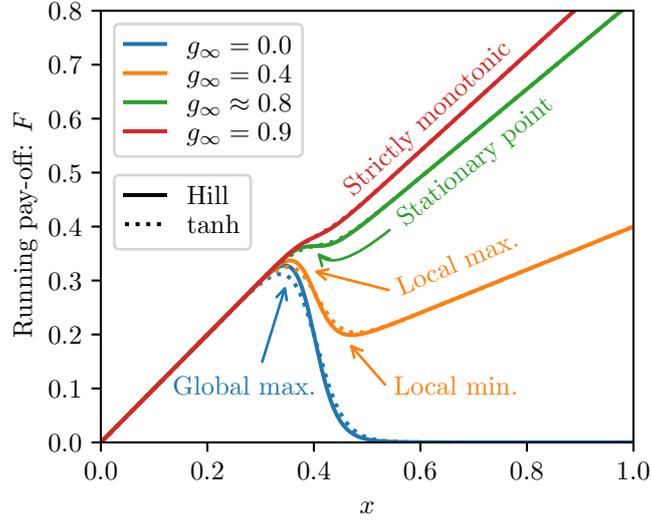


Figure 14: The running pay-off $F(x) = g(x)x$ for different parametrisations of $g(x)$. Solid curves show the Hill form (29d) while dashed curves show the (nearly identical) hyperbolic tangent form (136) sharing the same parameter values: $g_0 = 1$, $K = 0.4$, $n = 20$. When $g_\infty = 0$ we observe the global maximum. When $g_\infty \in (0, g_c)$, where $g_c \approx 0.8186$ is such that $4cn = (n-1)^2$, there are two roots of f corresponding to the local maximum and the local minimum. When $g_\infty = g_c$ there is a unique double root of f corresponding to a stationary point. For $g_\infty > g_c$ the F curve is strictly monotonic.

Recall from the discussion leading up to equation (38) that the singular arcs are characterised by states $X^1(t)$ satisfying $g'(X^1(t))X^1(t) + g(X^1(t)) = 0$. Defining

$$\begin{aligned} f(z) &:= g'(z)z + g(z) \\ &= -n \left(\frac{z}{K}\right)^n \frac{g_0 - g_\infty}{[1 + (z/K)^n]^2} + g_\infty + \frac{g_0 - g_\infty}{1 + (z/K)^n} \\ &= g_\infty + \frac{g_0 - g_\infty}{1 + (z/K)^n} \left(1 - n \left(\frac{z}{K}\right)^n \frac{1}{1 + (z/K)^n}\right), \end{aligned} \quad (132)$$

we thus set out to study the roots of $f(z)$.

If $g_\infty = 0$ then $f(z) = 0 \iff (z/K)^n = 1/(n-1)$ in which case there is exactly one real solution, and it is positive iff $n > 1$.

If $g_\infty > 0$, we may consider $f(z)$ as a function of $1 + (z/K)^n$ via $f(z) = \hat{f}(1 + (z/K)^n)$ where

$$\hat{f}(y) := g_\infty + \frac{g_0 - g_\infty}{y} \left(1 - n \frac{y-1}{y}\right). \quad (133)$$

The roots of $\hat{f}(y)$ solve the quadratic equation $cy^2 + (1-n)y + n = 0$ for $c := g_\infty/(g_0 - g_\infty)$ where $c > 0$ as we assume that g is strictly decreasing. It follows that $f(z) = 0 \iff \hat{f}(1 + (z/K)^n) = 0$, which is solved by

$$\left(\frac{z}{K}\right)^n = -1 + \frac{(n-1) \pm \sqrt{(n-1)^2 - 4cn}}{2c}. \quad (134)$$

It thus suffices to consider the value of $(z/K)^n$ as appears in (134).

If $4cn > (n-1)^2$ there are no real solutions. If $4cn = (n-1)^2$ then the discriminant is zero and there is exactly one real solution for $(z/K)^n$. Thus there is exactly one non-negative solution for z iff $n-1 \geq 2c = (n-1)^2/2n$. If $n \in [0, 1]$, then $n-1 \leq 0$ while $c > 0$ and there are no non-negative solutions, while if $n > 1$ the inequality holds strictly and there is exactly one positive solution. If $4cn < (n-1)^2$, then the discriminant is strictly positive and there are two positive roots for $(z/K)^n$ if $n > 1$, and no non-negative roots if $n \in [0, 1]$. To see why both roots are positive in the first case, note that the smaller root is positive iff

$$\begin{aligned} \frac{(n-1) - \sqrt{(n-1)^2 - 4cn}}{2c} > 1 &\iff (n-1) - 2c > \sqrt{(n-1)^2 - 4cn} \\ &\iff ((n-1) - 2c)^2 > (n-1)^2 - 4cn \\ &\iff 4c(1+c) > 0, \end{aligned} \quad (135)$$

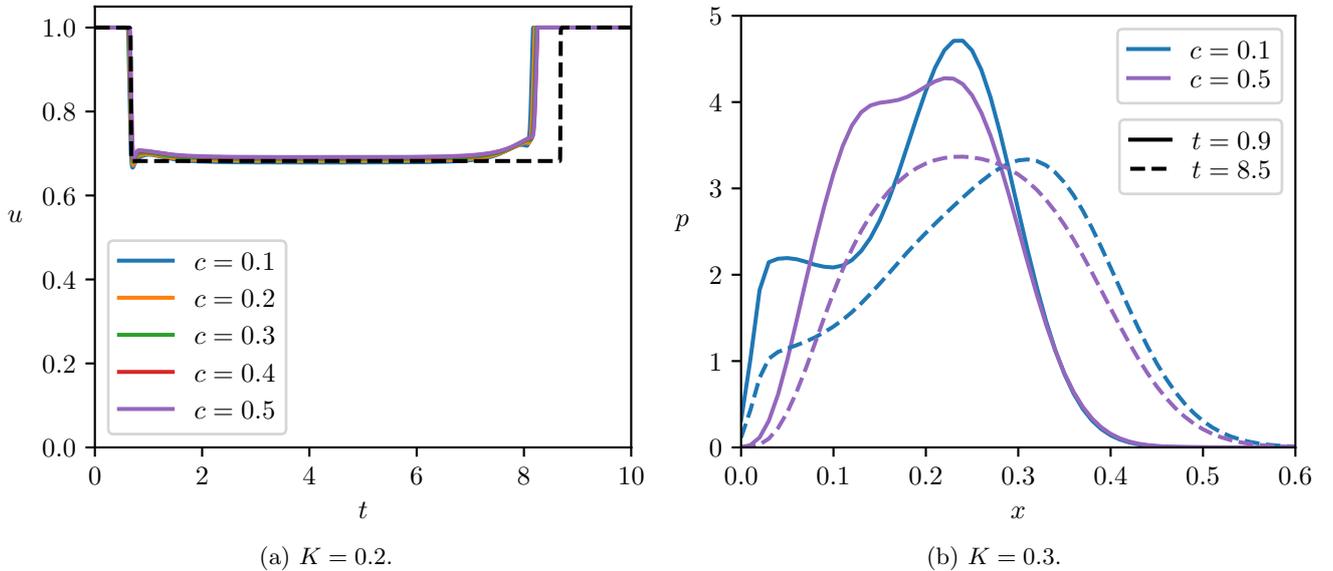


Figure 15: (a) Optimal controls $u(t)$ for the PDE copy-number problem (30) (coloured curves) for various division ratios $c \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ overlaid with the optimal control of the ODE model (33) (black dashed curve). Other parameters are given in (41). (b) Population distributions $p(x, t)$ associated with the optimal controls in (a) for various division ratios c at specific times t .

where we use the fact that $(n-1) - 2c > (n-1)(n+1)/2n > 0$ when $n > 1$ to justify squaring both sides. Thus, there are two positive solutions for z . To see why the second case admits no non-negative roots, note that the roots are bounded from above by $-1 + (n-1)/2c$, so when $n \leq 1$, this bound is strictly negative.

To recap the case of $g_\infty > 0$: If $4c > (n-1)^2/n$, there are no singular arcs. If $4c \leq (n-1)^2/n$ and $n > 1$, there is at least one and at most two values that $X^1(t)$ can take that satisfy the singular-arc condition (38), otherwise there are no singular arcs.

To gain further insight, we identify that the singular arc corresponds to maximising the running pay-off $F(x) := g(x)x$ in (33a), whose derivative with respect to the state x is precisely $f(x)$. Since $F(0) = 0$ and $F'(0) = f(0) = g(0) = g_0 > g_\infty \geq 0$, it follows that F is initially increasing. Therefore, when f has no roots, F must continue increasing. When f has a unique simple root, it must have a turning point, at which point F obtains its global maximum, after which it decreases. When f has a unique double root, it must have a stationary point and continue increasing. When f has two roots, it must have two turning points, first a local maximum, then a local minimum (at which a singular arc will not form, as this minimises the running pay-off), after which it continues increasing. These cases are all illustrated in fig. 14.

Therefore, when f has no roots, or when it has a double root, no singular arc is formed since the running pay-off F is monotonically increasing, and thus maximised for ever larger $X^1(t)$. The optimal control is thus $u(t) = 1$ for all time. When f has a unique simple root F is maximised at the global maximum, and we expect a singular arc to form if this global maximum is at a reachable state and there is sufficient time. When f has two simple roots a singular arc may form depending on whether the value of F at the local maximum surpasses all reachable pay-off values $F(X^1(t))$ and the time horizon.

While the algebra relied heavily on the particular form (29d) of $g(x)$, it is important to highlight that the results are remain robust to changes in the sigmoidal form. For example, in fig. 14 we show the running pay-off $F(x) = g(x)x$ for $g(x)$ defined by the hyperbolic tangent

$$g(x) = g_0 + (g_\infty - g_0) \frac{\tanh[n(x - K)] + 1}{2}, \quad (136)$$

where all parameters g_0 , g_∞ , K , and n remain the same. The F curves are almost indistinguishable preserving the features of the Hill case, even though the algebra in the hyperbolic tangent case becomes analytically intractable as the equations become transcendental.

Ultimately, we have uncovered conditions sufficient to suppress singular arc formation. When a singular arc does form, we have deduced that it can occur at a unique state value $X^1(t)$.

D Division ratios

The ODE approximation is independent of θ (and thus c), since it only accounts for the changes in copy number, which are unaffected by division. This is not as dire as it might seem for the ODE approximation, because the PDE optimal control proves insensitive to all values of c we employed, even though the copy-number distributions are noticeably different (see fig. 15). We observe in fig. 15b that the distributions for $c = 0.1$ have a larger contribution near $x = 0$ and are most spread than the distributions for $c = 0.5$, both on the singular arc ($t = 0.9$) and on the terminal nonsingular arc ($t = 8.5$). Nevertheless, the corresponding optimal controls in fig. 15a are remarkably similar.

E Numerical artefact

In this appendix, we study the source of the discrepancy between the analytical and numerical optimal controls of problem (45). The analytical singular control is given by (51), however, this does not coincide with the numerical singular control in the vicinity of $u(t) = 0$, as illustrated in fig. 8.

To see concretely that the analytical form does not match the numerical solution, note that intervals of small control values $u(t)$ correspond to intervals where $Y(t) \approx 0.03$, which is significantly smaller than both K and L . When $y \ll L$, the function $\alpha(y)$ and its derivatives have the leading-order forms

$$\alpha(y) \sim \alpha_0, \quad \alpha'(y) \sim -(\alpha_0 - \alpha_\infty) \frac{my^{m-1}}{L^m}, \quad \alpha''(y) \sim -(\alpha_0 - \alpha_\infty) \frac{m(m-1)y^{m-2}}{L^m}. \quad (137)$$

Analogous forms are true for $g(y)$ and its derivatives when $y \ll K$. It follows from (51) that the singular control is of the form

$$\frac{R(x, y)}{D(x, y)} = \frac{\beta g'(y)[xg'(y) - 2\alpha'(y)]}{a[\alpha''(y)g'(y) - \alpha'(y)g''(y)]} + \frac{y[b + g(y)]}{a} \sim \frac{\beta y}{a(m-n)} \left[\frac{(g_0 - g_\infty)nL^m}{(\alpha_0 - \alpha_\infty)mK^n} xy^{n-m} - 2 \right] + \frac{y(b + g_0)}{a}. \quad (138)$$

For the parameters in (54) and $y \approx 0.03$ as in fig. 8, the leading-order approximation in (138) evaluates, approximately, to

$$\frac{R(x, y)}{D(x, y)} \approx 10^{-3} [-2x \times 10^{12} + 2] + 6 \times 10^{-2}. \quad (139)$$

The first term inside the brackets is negative and of exceedingly large magnitude, dominating the rest of the contributions. Importantly, this is not a valid value for the control and hence cannot coincide with the positive singular arc of the numerical solution.

Note that this discrepancy does not always arise. For example, when we change the sign of the approximation (138) by choosing $m = 9 < 10 = n$ the numerical solution agrees more closely with the analytical solutions (see fig. 16).

To explore the source of the numerical deviation, we perform an analysis of the discrete adjoint problem. Recall that the optimal control problem is to maximise (45a) subject to the dynamics (45). The objective may be discretised by the Riemann sum

$$J = X^N + \Delta t \sum_{k=0}^{N-1} g(Y^k)X^k, \quad (140)$$

alongside a forward-Euler discretisation of the dynamics, of the form

$$\frac{X^{k+1} - X^k}{\Delta t} = \alpha(Y^k) - (\beta + g(Y^k))X^k =: M_1^k, \quad X^0 = X_0, \quad (141a)$$

$$\frac{Y^{k+1} - Y^k}{\Delta t} = aY^k - (b + g(Y^k))Y^k =: M_2^k, \quad Y^0 = Y_0, \quad (141b)$$

for $k = 0, \dots, N-1$. Although the results in the main section were obtained using a backward-Euler scheme, the anomalous singular arcs that do not match the continuum theory are present in solutions using both the forward and backward discretisations, and so we choose the simplest possible scheme to analyse.

We introduce the Lagrangian

$$\mathcal{L} = J + \sum_{k=0}^{N-1} \lambda_1^{k+1} (X^{k+1} - X^k - \Delta t M_1^k) + \sum_{k=0}^{N-1} \lambda_2^{k+1} (Y^{k+1} - Y^k - \Delta t M_2^k) + \lambda_1^0 (X^0 - X_0) + \lambda_2^0 (Y^0 - Y_0). \quad (142)$$

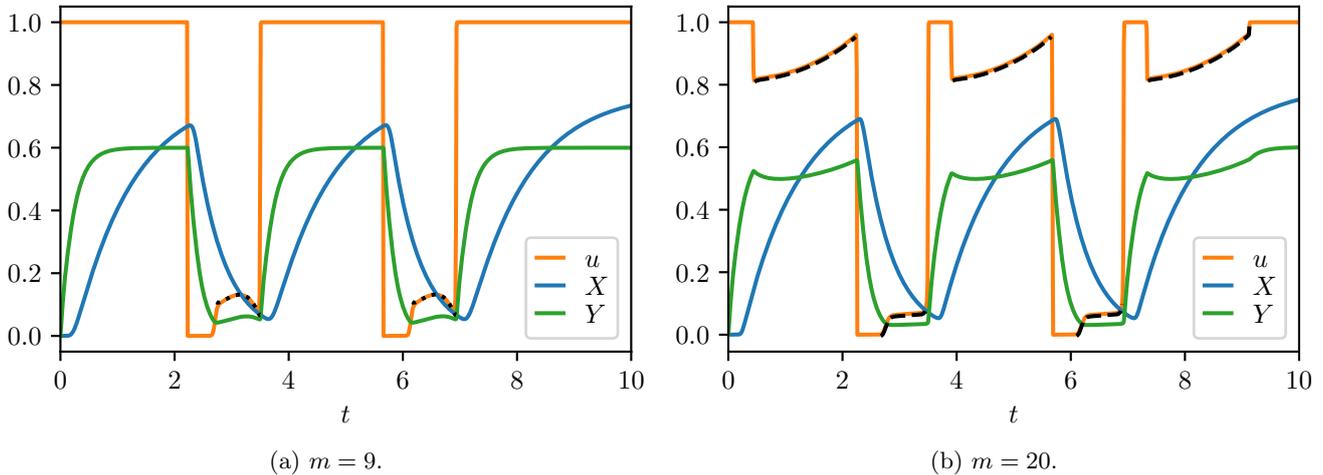


Figure 16: Optimal controls $u(t)$ and state variables $X(t)$ and $Y(t)$ (solid curves) for the ODE model (45) overlaid with the singular control (black broken curves) of the (a) continuum formulation (51), and of the (b) discrete formulation (146), for different values of m and other parameters defined in (54).

The costate equations and terminal conditions are derived by ensuring that $\partial\mathcal{L}/\partial X^k = \partial\mathcal{L}/\partial Y^k = 0$ for all k , from which we deduce that, for $k = 0, \dots, N-1$,

$$\frac{\lambda_1^{k+1} - \lambda_1^k}{\Delta t} = \lambda_1^{k+1} [\beta + g(Y^k)] + g(Y^k), \quad \lambda_1^N = -1, \quad (143a)$$

$$\frac{\lambda_2^{k+1} - \lambda_2^k}{\Delta t} = \lambda_1^{k+1} [-\alpha'(Y^k) + g'(Y^k)X^k] + \lambda_2^k [b + g(Y^k) + g'(Y^k)Y^k] + g'(Y^k)X^k, \quad \lambda_2^N = 0. \quad (143b)$$

Note that the discrete adjoint formulation is consistent with the continuum formulation (48). The derivative of the discrete objective $J(u^0, \dots, u^N)$ with respect to the discrete control u^k is then given by the partial derivative of the Lagrangian (considered as a function of the state, costate, and control $\mathcal{L}(\{X^k\}, \{Y^k\}, \{\lambda_1^k\}, \{\lambda_2^k\}, \{u^k\})$)

$$\frac{\partial J}{\partial u^k} = \frac{\partial \mathcal{L}}{\partial u^k} = -a\Delta t \lambda_2^{k+1}, \quad (144)$$

for $k = 0, \dots, N-1$. In the discrete setting, as in the continuum setting, singular arcs may arise where, for the optimal solution, $\lambda_2^{k+1} = 0$ over some contiguous range of k . In this case, the optimal control may be determined analogously to the continuum case. Since $\lambda_2^{k+1} = 0$ for some range of k , we deduce from (143b) that

$$\lambda_1^{k+1} = \frac{g'(Y^k)X^k}{\alpha'(Y^k) - g'(Y^k)X^k}. \quad (145)$$

Substituting (145) into (143a) we find that

$$\frac{1}{\Delta t} \left(\frac{g'(Y^k)X^k}{g'(Y^k)X^k - \alpha'(Y^k)} - \frac{g'(Y^{k-1})X^{k-1}}{g'(Y^{k-1})X^{k-1} - \alpha'(Y^{k-1})} \right) = \frac{g'(Y^k)X^k}{g'(Y^k)X^k - \alpha'(Y^k)} [\beta + g(Y^k)] - g(Y^k). \quad (146)$$

Using (141) we now replace X^k and Y^k by functions of X^{k-1} , Y^{k-1} , and u^{k-1} . If we consider the limit as $\Delta t \rightarrow 0$, we may Taylor with respect to Δt and retain the leading-order contribution. Ultimately, this produces a linear equation for u^{k-1} whose solution coincides with the continuum singular control (51). If, however, we retain a nonzero time step $\Delta t > 0$, and (numerically) find the root of equation (146), we recover the (numerical) singular arcs of the solution (see fig. 16). We thus infer that it is precisely an artefact of the discretisation that causes spurious singular arcs.

F Increasing the resolution

In this appendix, we study the improvements when we increase the discretisation resolution of the (x, y) -state space for the models in section 4. Recall that the models in section 3 are one-dimensional, allowing us to use a resolution of $N_{\text{state}} = 100$ nodes in the state variable with acceptable performance on a standard laptop (on the order of one second to solve the state and costate equation, with good optimisation convergence in a few hundred BFGS iterations). The models in section 4 are two-dimensional, and we chose to use a coarse resolution

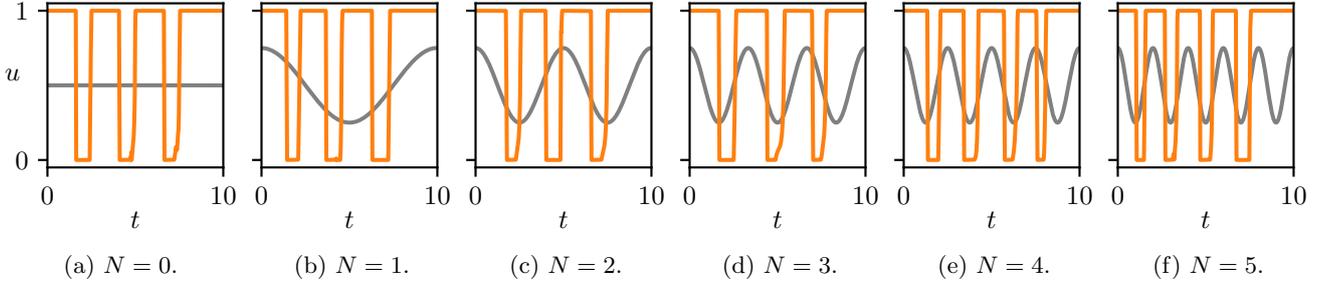


Figure 17: Optimal controls (orange curves) for the PDE copy-number model (43), using initial control guesses $u_0(t; N)$ defined in (56) for different values of N (grey curves). The optimisation was terminated after 1500 objective evaluations with $N_{\text{state}} = 40$ discrete nodes in each state dimension.

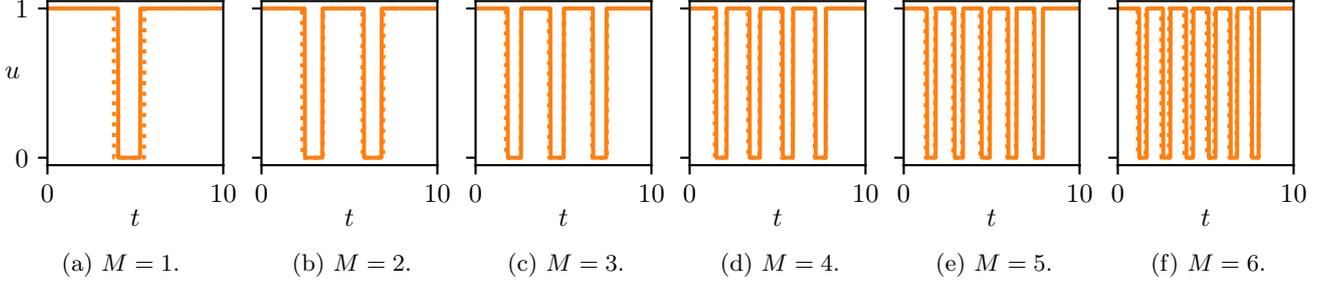


Figure 18: Locally optimal switching structures for the PDE model (43), for resolutions $N_{\text{state}} = 20$ (solid curves) and $N_{\text{state}} = 40$ (dotted curves). Each switch is optimised in turn, terminating once no switch could be moved to improve the objective.

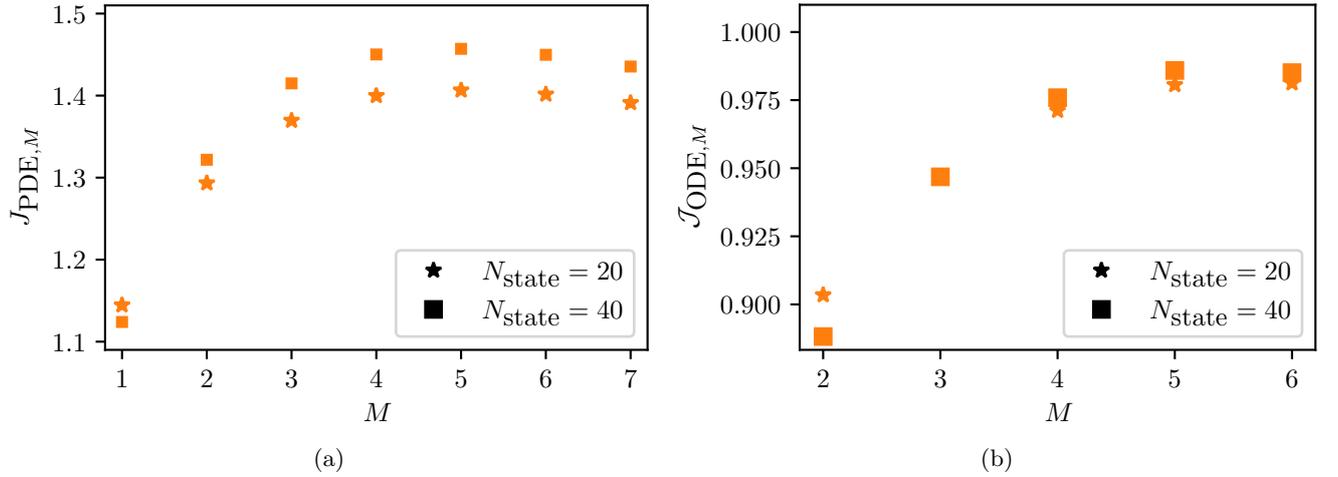


Figure 19: (a) Objective value for the PDE optimal controls $J_{\text{PDE},M}$ with M valleys using N_{state} nodes in the state discretisation of model (43). (b) Relative performance of the ODE model (45) compared to the PDE model, as defined in (58), for different locally optimal ODE controls with M valleys, using N_{state} nodes in the state discretisation. Other parameters are given in (54).

of $N_{\text{state}} = 20$ nodes in each state variable (for a total state space of 400 nodes) which gave similar computational times for solving the state and costate equations (on the order of one second), however, the objective landscape is more intricate and thus more BFGS iterations were required, as well as many more initial conditions since there were multiple local optima. To examine whether this is too coarse, we subsequently performed the optimisation using $N_{\text{state}} = 40$ nodes in each state variable, which resulted in computational times of approximately 30 seconds per iteration (that is, to solve the state and costate equations and compute the derivative). Note that, for a finer state discretisation, the time step must be smaller in order to retain a monotone scheme, which also contributes to the increased execution time. To retain a fair comparison between the PDE and the ODE results, here too we discretised the time in the ODE optimisation problem to match the PDE time discretisation.

The resulting optimal PDE controls were qualitatively similar (fig. 17), and the parametrised optimisation step produced almost identical switching structures (fig. 18). While the objective function values for PDE optimal

controls were up to 7% larger (fig. 19a), the ODE optimal controls were similarly larger, resulting in an almost unchanged relative ODE performance (fig. 19b). We thus deduce the accuracy of the results even for the coarse resolution.

G MPC with partial observability

In the main text, we present results for MPC controls for first-order moment closure approximations of systems (68) and (78) that assumed complete state observability. For the sake of completeness, we present here the numerical simulations where only the zeroth-order information is observable (the total grower mass, $G(t)$, from which the total producer mass is given by $1 - G(t)$) but the first-order information is not ($X(t)$ for system (68) and $X(t)$ and $Y(t)$ for system (78)). In real-world scenarios, the input is expected to be noisy, in which case an appropriate fusion between the noisy data and model predictions must be decided upon by some filtering algorithm. Similarly, missing state variables are typically estimated by constructing observers [25]. Here, we consider the simpler case of noiseless data and no observer construction, in order to directly study the influence of partially missing state information without mitigating efforts. Having showed analytically in section 5.2.2 that $X(t)/G(t)$ is constant, we have two options for updating the unobservable state $X(t)$ given new $G(t)$ data in the control loop: (i) use the constant ratio to update the value of $X(t)$, or (ii) update only the observable states, retaining the previous value of $X(t)$. The $Y(t)$ value is retained in both cases (we do not treat the case where some of the first-order information is observable, although this can be done similarly). We present numerical results for both algorithms here.

For system (68) with optimal controls appearing in fig. 11 for full observability, we add the case where $X(t)$ is unobservable. $X(t)$ can be updated to satisfy $X(t)/G(t) = \alpha/(\beta + A)$ (fig. 20) or not updated in the MPC loop (fig. 21). Only the MPC control differs between these figures. Overall, the performance with partial observability is slightly reduced. The discrepancy is not more than 2% and the performance in the partially observable case remains high.

For system (78) with optimal controls for the first-order moment closure appearing in fig. 13 for full observability, we add the case where both $X(t)$ and $Y(t)$ are unobservable. $X(t)$ can be updated to satisfy $X(t)/G(t) = \alpha/(\beta + A)$ (fig. 22) or not updated in the MPC loop (fig. 23). Only the MPC varies between these three figures. Here too, overall, the performance with partial observability is slightly reduced. The discrepancy is not more than 2% and the performance in the partially observable case remains above 90%.

H Optimal control near the terminal time

We study the optimal control of the ODE problem (78) and (79) after removing the $P(t)$ and $X(t)$ dynamics as described in section 5.2.2, to give the ODE system

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - u(t)h\left(\frac{\alpha}{\beta + A}\right)G(t), \quad G(0) = 1, \quad (147a)$$

$$\dot{Y}(t) = -\Lambda(t)Y(t) + \mu[1 - G(t)] - \nu Y(t), \quad Y(0) = 0, \quad (147b)$$

where the turbidostat dilution rate is given by

$$\Lambda(t) = AG(t) + \hat{a}\left(\frac{Y(t)}{1 - G(t)}\right)[1 - G(t)]. \quad (147c)$$

We introduce the Hamiltonian

$$H = \lambda_1(t) \left[AG(t) - \Lambda(t)G(t) - u(t)h\left(\frac{\alpha}{\beta + A}\right)G(t) \right] + \lambda_2(t) \{ -\Lambda(t)Y(t) + \mu[1 - G(t)] - \nu Y(t) \} - \Lambda(t)Y(t), \quad (148)$$

and the costate equations

$$\dot{\lambda}_1(t) = -\lambda_1(t) \left[A - \Lambda(t) - \frac{\partial \Lambda(t)}{\partial G(t)}G(t) - u(t)h\left(\frac{\alpha}{\beta + A}\right) \right] - \lambda_2(t) \left[-\frac{\partial \Lambda(t)}{\partial G(t)}Y(t) - \mu \right] + \frac{\partial \Lambda(t)}{\partial G(t)}Y(t), \quad (149a)$$

$$\lambda_1(T) = 0, \quad (149b)$$

$$\dot{\lambda}_2(t) = -\lambda_1(t) \left[\frac{\partial \Lambda(t)}{\partial Y(t)}G(t) \right] - \lambda_2(t) \left[-\Lambda(t) - \frac{\partial \Lambda(t)}{\partial Y(t)}Y(t) - \nu \right] + \Lambda(t) + \frac{\partial \Lambda(t)}{\partial Y(t)}Y(t), \quad (149c)$$

$$\lambda_2(T) = -1. \quad (149d)$$

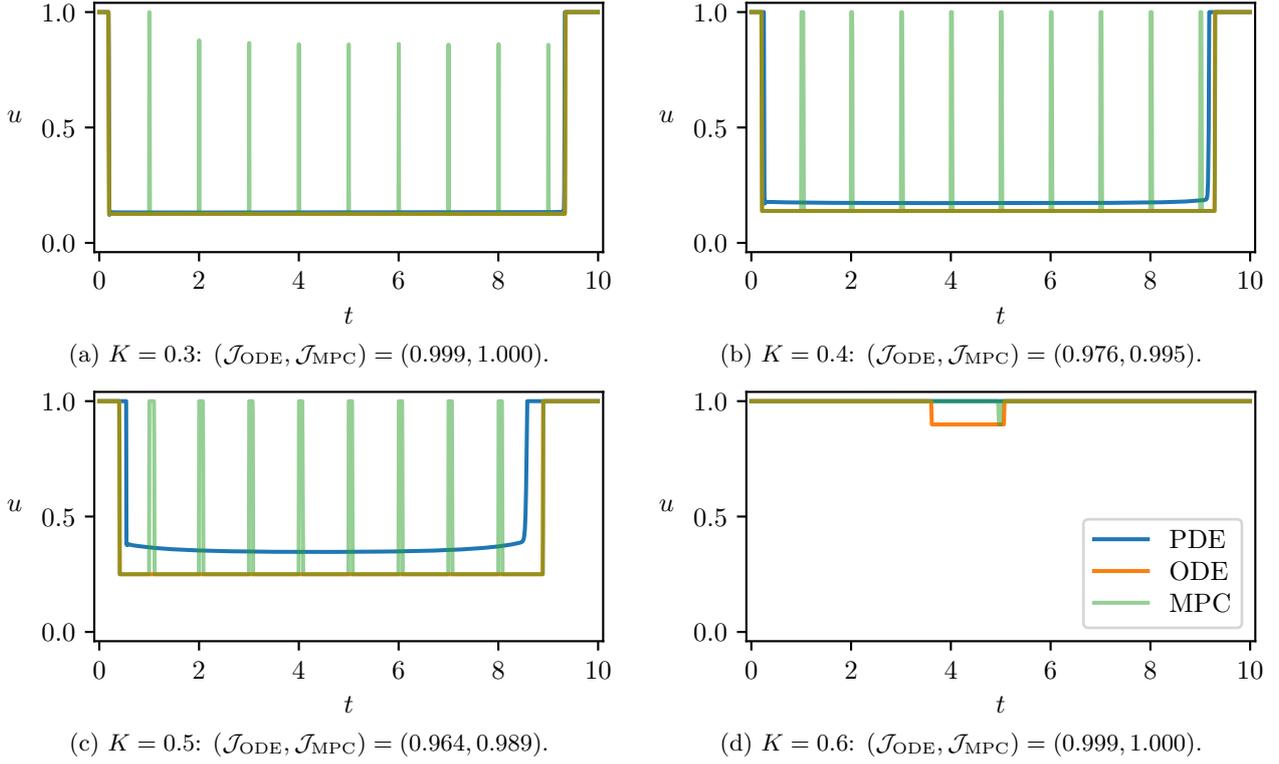


Figure 20: Optimal controls from the PDE model (61), the ODE model (65), and the same ODE model in closed-loop MPC (for partial observability, where the state $G(t)$ is observable and the state $X(t)$ is updated to satisfy $X(t)/G(t) = \alpha/(\beta + A)$), illustrated as blue, orange, and dashed green curves, respectively, for various values of the Hill parameter K . Other parameters are given in (73).

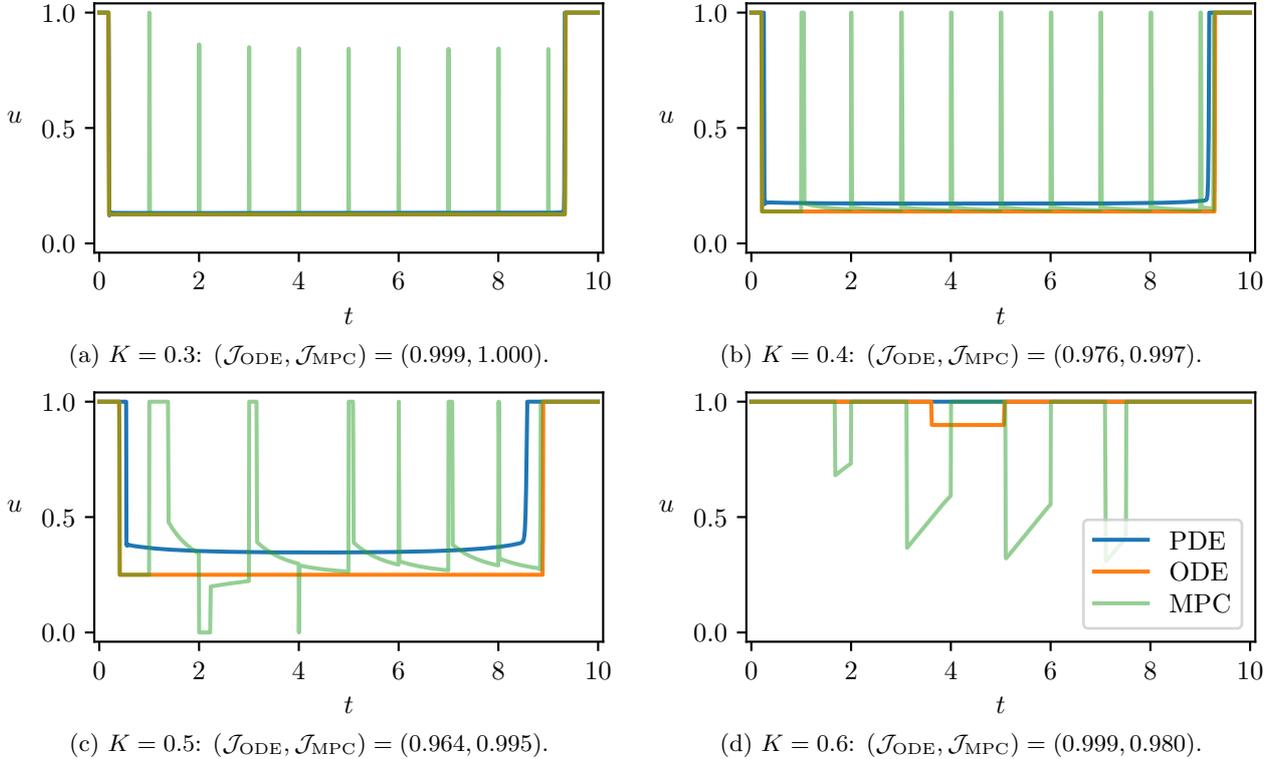


Figure 21: Optimal controls from the PDE model (61), the ODE model (65), and the same ODE model in closed-loop MPC (for partial observability, where the state $G(t)$ is observable and the state $X(t)$ is not updated on new data input), illustrated as blue, orange, and dashed green curves, respectively, for various values of the Hill parameter K . Other parameters are given in (73).

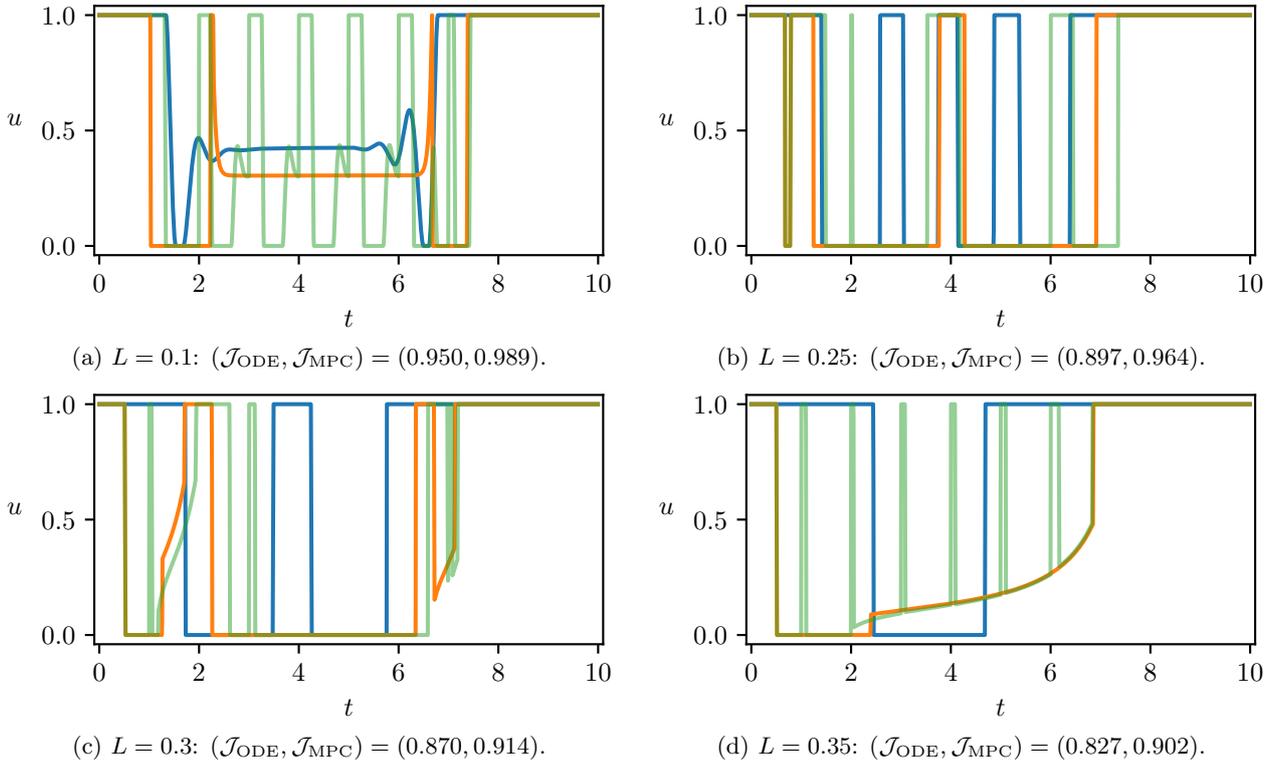


Figure 22: Optimal controls from the PDE model (75), the first-order ODE model (78) where states $P(t)$ and $X(t)$ have been excluded analytically as described in section 5.2.2, and the same ODE model in closed-loop MPC (updating the states $G(t)$ and preserving the state $Y(t)$), illustrated as blue, orange, and dashed green curves for various values of the Hill parameter K . Other parameters are given in (73) and (82). The PDE control in (a) is obtained with an additional regularisation term as detailed in fig. 12.

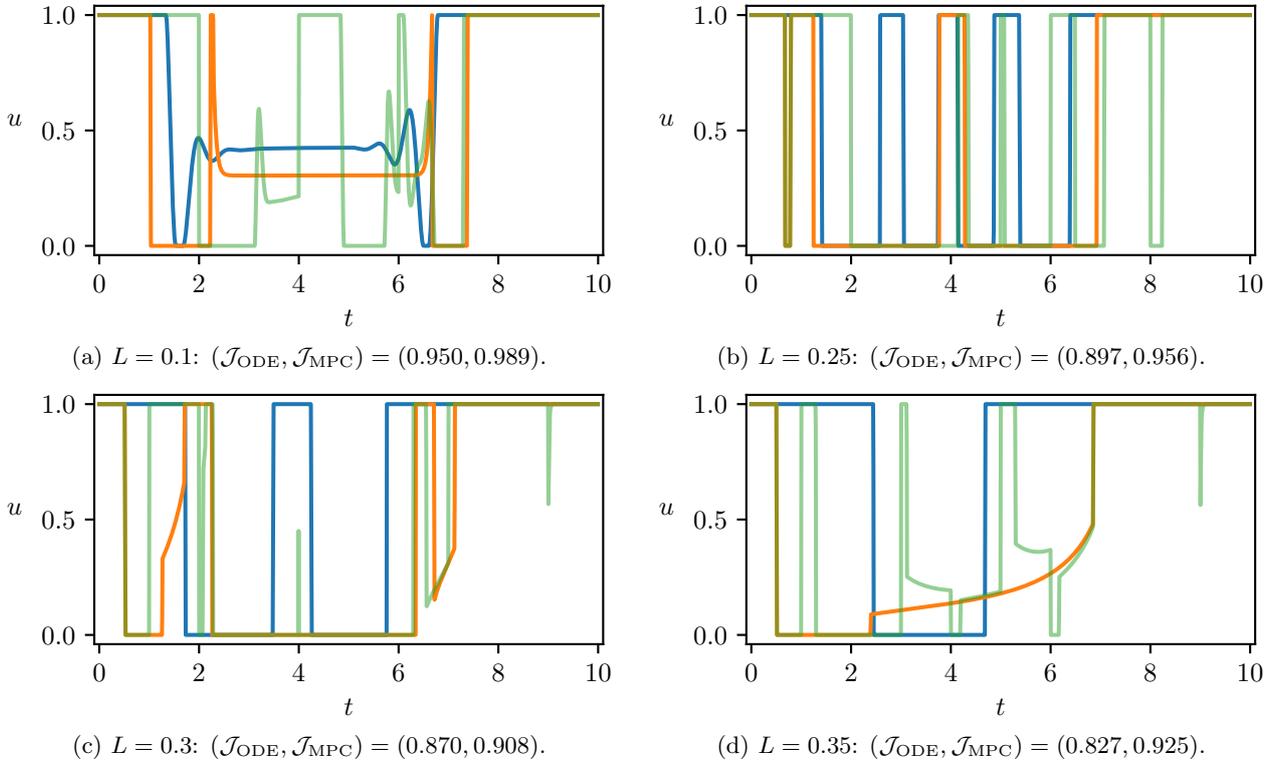


Figure 23: Optimal controls from the PDE model (75), the first-order ODE model (78), and the same ODE model in closed-loop MPC (updating the states $G(t)$, updating $X(t)$ to satisfy $X(t)/G(t) = \alpha/(\beta + A)$ and preserving the state $Y(t)$), illustrated as blue, orange, and dashed green curves for various values of the Hill parameter K . Other parameters are given in (73) and (82). The PDE control in (a) is obtained with an additional regularisation term as detailed in fig. 12.

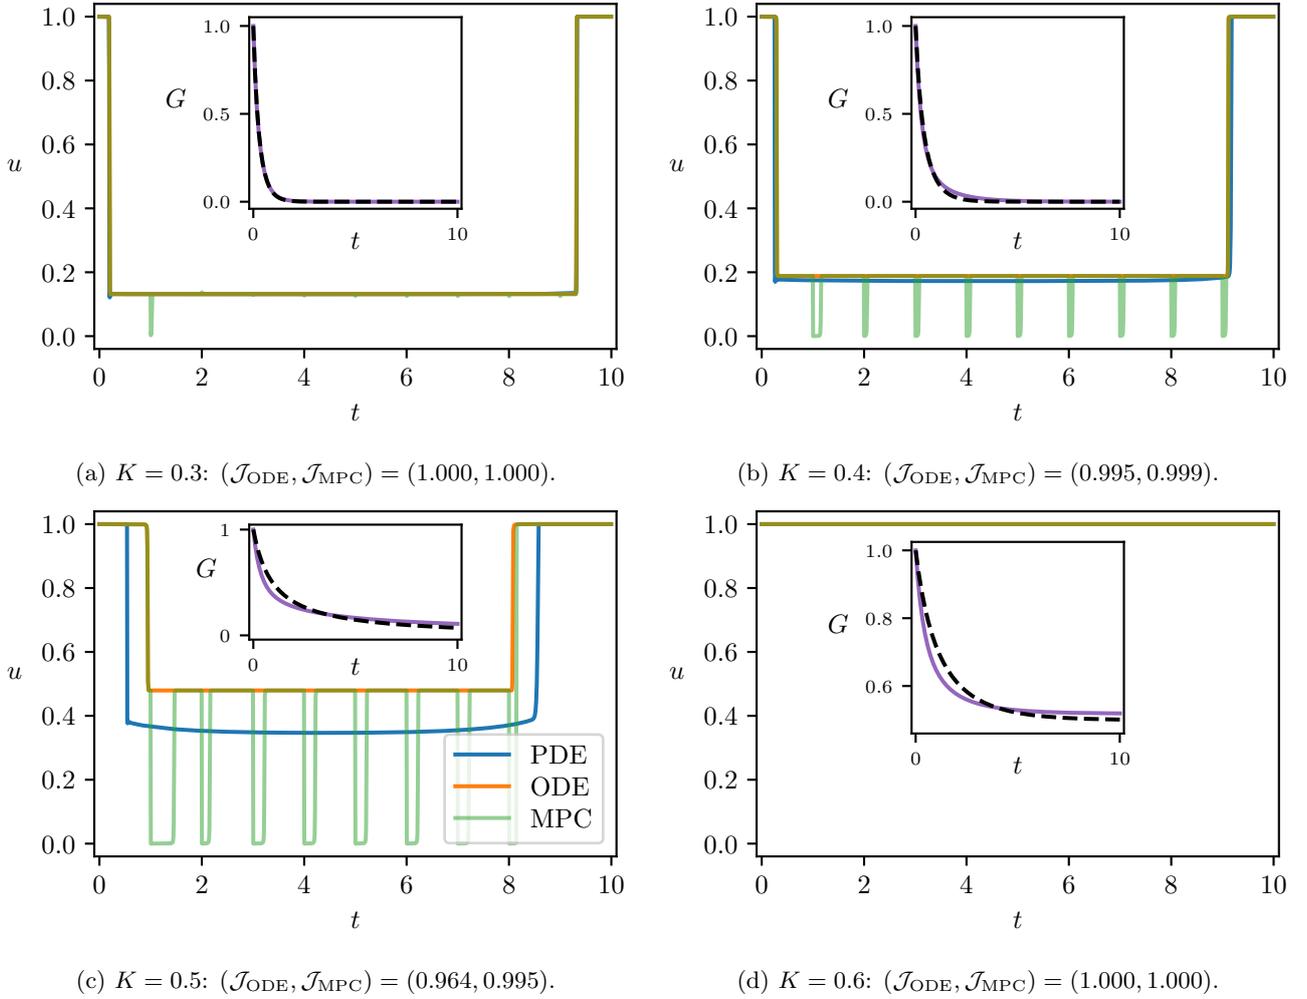


Figure 24: Optimal controls from the PDE model (61), the zeroth-order ODE model (65) where the parameter F is fitted, and the same ODE model in closed-loop MPC, illustrated as blue, orange, and dashed green curves, respectively, for various values of the Hill parameter K . Other parameters are given in (73). Inset plots show the simulated $G(t)$ dynamics (purple curves), overlaid by the fitted match (black dotted curves).

The switching function is given by

$$\frac{\partial H}{\partial u} = -\lambda_1(t)h\left(\frac{\alpha}{\beta + A}\right)G(t). \quad (150)$$

Since $\Lambda(t)$ is an interpolation between A and some evaluation of \hat{a} , which is itself an interpolation between a_∞ and A (see (74)), it holds that $\Lambda(t) \in [a_\infty, A]$, and it follows that $G(t)$ decreases at most exponentially, and thus remains positive for all time. Therefore, the switching function vanishes if and only if $\lambda_1(t) = 0$, from which it follows that, on non-singular arcs (where $\lambda_1(t) \neq 0$), the optimal control takes the bang-bang form

$$u(t) = \begin{cases} 1, & \lambda_1(t) > 0, \\ 0, & \lambda_1(t) < 0. \end{cases} \quad (151)$$

At the terminal time $\lambda_1(T) = 0$, and so the bang-bang form (151) is not informative. Instead, we look at the derivative of the costate at the terminal time,

$$\dot{\lambda}_1(T) = -\mu < 0, \quad (152)$$

which is strictly negative since we assume a nontrivial protein production rate $\mu > 0$. It follows, by the terminal condition $\lambda_1(T) = 0$ and continuity, that $\lambda_1(t) > 0$ in some vicinity of the terminal time $(T - \epsilon, T)$ for some $\epsilon > 0$. Therefore, in this neighbourhood, the optimal control is bang-bang and, from (151), we see that it is maximal.

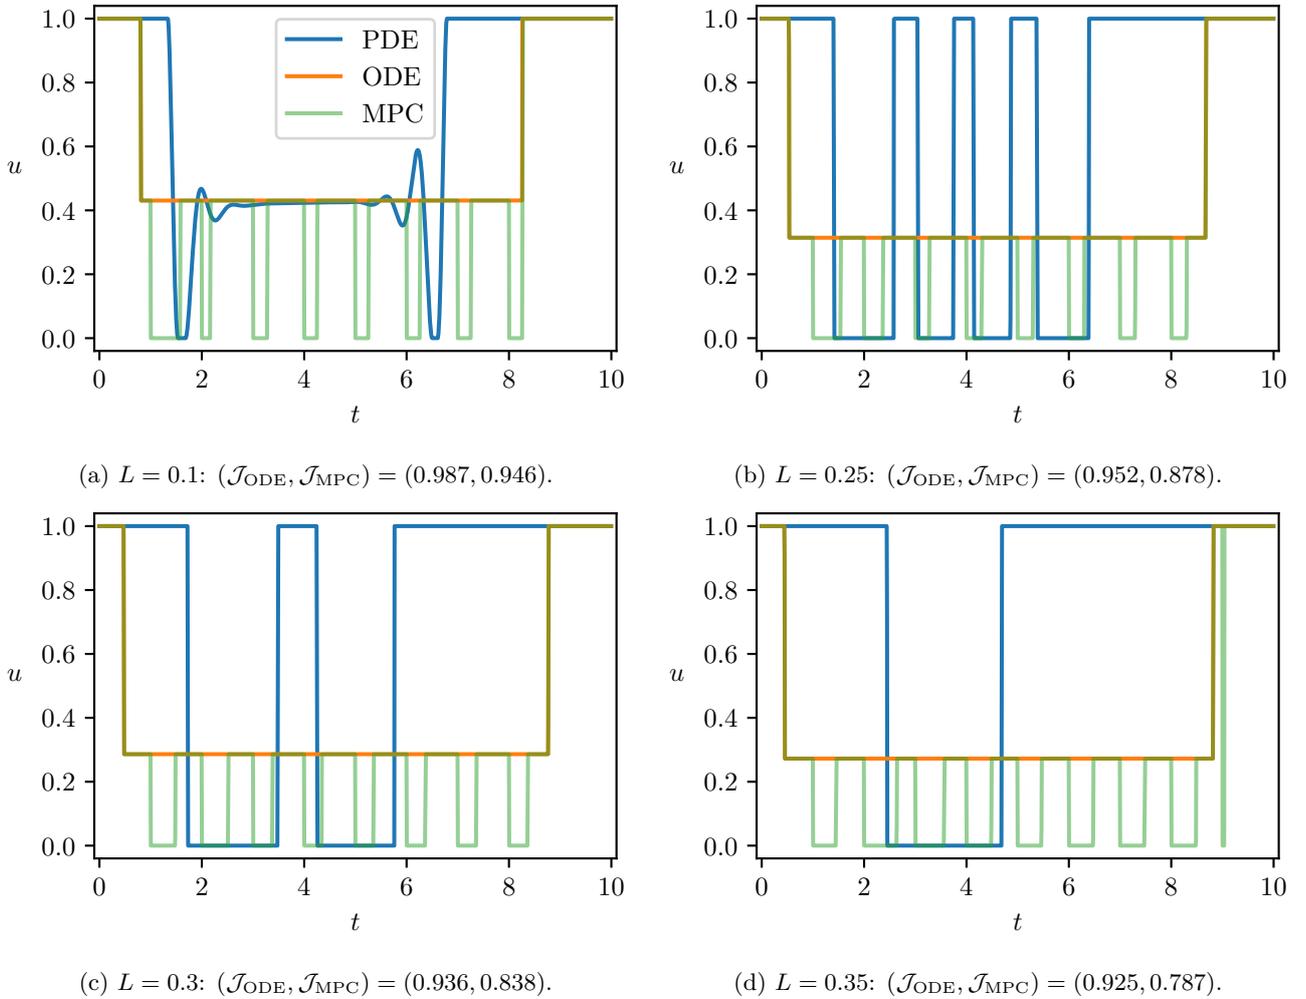


Figure 25: Optimal controls from the PDE model (75), the zeroth-order ODE model (65), and the same ODE model in closed-loop MPC, illustrated as blue, orange, and dashed green curves for various values of the Hill parameter K . For the ODE model, both parameters a and F are given by fitting. Other parameters are given in (73) and (82). The PDE control in (a) is obtained with an additional regularisation term as detailed in fig. 12.

I Zeroth-order approximation and fitting parameters

In the main text, we used the first-order form (71) to provide the value for the parameter F in the zeroth-order system (65). One might propose not to use the associated first-order approximations but instead to fit the parameter F (as we do later for the parameter a). We argue here that this does not qualitatively change the results.

For the simpler equation (61) of section 5 the results in fig. 11 are compared with the results in fig. 24 where the parameter F has been obtained by fitting: a characterisation “experiment” is run using $u(t) = 1$ and matching the dynamics of $G(t)$. We see that the results are very similar. The main difference is that, whereas the open-loop ODE control previously underestimated the PDE optimal control when using the first-order F value, in this case it is overestimated. This discrepancy is almost entirely corrected by the closed-loop MPC. The insets illustrate the “experimental data” $G(t)$ as purple curves, overlaid by the fitted match as black dotted curves.

For equation (75) in section 6 the results in fig. 12 (using first-order information for F) are to be compared with the results in fig. 25 (fitted F). The results obtained by fitting the parameter F show a noticeable quantitative improvement, in particular, the singular arc is much better approximated (figs. 12a and 25a). Nevertheless, the controls appear qualitatively very similar. Moreover, as L increases, the improvements diminish: by $L = 0.5$ the open-loop improvement is less than 2% and the MPC improvement is less than 1%. This is because fitting the parameter F allows for a slightly more favourable parametrisation, however, no structural improvement to the approximation has been made. For this reason, the poorest performances, which occur at the higher values of L , remain largely unimproved.

This procedure is idealised in several ways. Firstly, there is no noise. Secondly, we take data points at

the PDE discretisation, which is on the order of one measurement every twenty seconds. For sampling and flow cytometry measurements this cannot be achieved. In realistic settings where only noisy and sparse measurements are available, it is likely that the utility of fitting will be significantly impacted. For this reason, we treat this case only as a footnote, despite the somewhat improved performance. Moreover, the main performance degradation in fig. 25d remains effectively the same, highlighting the fact that the primary source of error remains unaddressed by even this idealised improvement, and more fundamental changes are needed, such as those brought about by the first-order approximation.