Modelling and optimal control of a two-species bioproducing microbial consortium

Davin Lunz, Joseph Frédéric Bonnans

HAL Id: hal-03479385

https://inria.hal.science/hal-03479385v2

Submitted on 21 Jun 2022

Modelling and optimal control of a two-species bioproducing microbial consortium

Davin Lunz[*1,2] and J. Frédéric Bonnans[3]

[1]Inria Paris, 2 rue Simone Iff, 75012 Paris, France
[2]Institut Pasteur, 28 rue du Docteur Roux, 75015 Paris, France
[3]Université Paris-Saclay, CNRS, CentraleSupélec, Inria, Laboratory of signals and systems, 91190, Gif-sur-Yvette, France

Abstract

Motivated by recent laboratory experiments, we study microbial populations with light-inducible genetic differentiation that generates a two-species microbial consortium relevant for bioproduction. First, we derive a hierarchy of models describing the evolution of the microbial populations, each with decreasing complexity. This sequential order reduction reveals the connections between several popular classes of models used in this context. Second, we demonstrate the analytical insight the order reduction provides by studying the optimal control of such a reduced-order system of nonlinear ordinary differential equations. Appealing to Pontryagin's maximum principle, we find different optimal control structures within different regions of the parameter space. Explicit solutions are obtained in a subset of parameter space, while, for the remainder of parameter space, closed-form solutions are obtained that depend on a scalar value that solves a particular transcendental equation. We show that a unique solution of the scalar equation exists and lies in a known compact interval, making its numerical approximation particularly easy. The analytical results are verified against direct numerical calculations.

1 Introduction

Microbial consortia are colonies of two or more interacting microorganisms living together. Such colonies offer increased application potential over cultures of single strains since each organism can be specialised for different tasks, allowing a diverse set of mechanisms that need not coexist in a single organism. This promises improvements in bioproduction, bioremediation, and other areas [5, 16, 25, 33]. A primary focus of recent work has been the modelling and control of such consortia [1, 4, 13, 22, 26, 31, 32]. In this work, we tackle both the modelling and control facets. First, we develop a hierarchy of consortia models of increasing coarseness: more coarse-grained models trade off fidelity for tractability. Second, we use the coarsest member of this hierarchy to demonstrate how significant analytical progress can be made in the optimal control of consortia. While these two contributions will require deep dives into mathematically distinct areas, the two outcomes remain intimately connected as the analytical insight of the coarse model informs the higher-fidelity models in the hierarchy.

On the modelling front, we observe that many previous studies (e.g. Refs. [1, 22, 26, 31]) pose ordinary-differential rate equations to model such consortia. This modelling paradigm neglects certain phenomena, such as heterogeneity driven by underlying stochasticity [2, 21, 32] and spatial inhomogeneity driven by non-uniform nutrient availability [4]. Sometimes these simpler models satisfactorily capture the underlying population dynamics. However, it has been demonstrated that there are circumstances where the neglected heterogeneity can render rate-equation-derived controls substantially suboptimal [20, 21]. The modeler is thus faced with the challenge of deciding which model is most appropriate for a given situation. This decision involves balancing a delicate trade-off: retaining a higher-fidelity model typically comes at the cost of increased complexity. Therefore, it is of great utility to the modeler to be equipped with a toolbox containing models of varying fidelity. Indeed, no single model is universally best. Instead, it can be enormously beneficial to combine different models so as to leverage their relative strengths.

---

*davin.lunz@inria.fr

Our first contribution in this work is to develop a hierarchy of models with ranging fidelity (and tractability). We begin by posing a multiscale fully stochastic model comprising a well-mixed population of cells. Stochastic chemical kinetics govern the internal state of each individual cell, while the population evolves via stochastic events of cell division and removal. The complexity of this model renders it all but intractable, motivating order reduction. Taking an expectation over the population stochasticity, we derive a model for the evolution of the expected population density. On the one hand, this reduced-order model preserves the population heterogeneity and conserves the dynamics on both single-cell and population scales. On the other hand, we retain the model complexity of a stochastic single-cell model. A further averaging over the state space demonstrates how we may arrive at the ordinary-differential rate equations. This procession from a stochastic model, via a heterogeneous model, through to a homogeneous model reveals the assumptions required at each stage and thereby crystallises the nature of each model. The end result is a family of models of varying fidelity alongside the underlying assumptions connecting them.

On the control front, synthetically engineering microbial consortia is still a young field, in large part, due to the difficulty of controlling the populations [13]. Armed with the modeling framework described above, our second contribution is an analytical solution to the model-based optimal control problem of maximising bioproduction using a microbial consortium. We study a system of nonlinear ordinary-differential equations obtained via the previously described order-reduction technique; the model is the coarsest member of the model hierarchy. We apply Pontryagin's maximum principle to solve the optimal control problem explicitly, modulo a single scalar value that must be calculated numerically in some regions of parameter space. When this numerical value is required, we show that it is the unique solution of an equation that lies in a known compact interval, making its computation eminently tractable. Having used the coarsest model for the optimal control problem, we discuss how our solution can also inform control problems based on higher-fidelity members of the model hierarchy. This observation highlights one way in which the two contributions of this work are tightly coupled: after posing a biologically faithful model, systematic reduction produces a hierarchy that enables greater tractability, and the associated insights can then in turn inform parameter/control choice in the higher-fidelity models throughout the hierarchy.

The rest of the paper is organised as follows. In section 2, we derive the hierarchy of increasingly reduced-order models. In section 3, using the most tractable model of the hierarchy, we pose and solve an optimal control problem. Finally, in section 4 we summarise our work, tying together the two main contributions, and comment on how the techniques presented may be generalised and extended.

2 Model derivation

In this section, our aim is to derive a sequence of decreasingly complex (and thus increasingly tractable) population models of a microbial consortium. The plan is to start with a high-fidelity stochastic model that is faithful to the underlying biological process. Due to the size of the associated state space, this first stochastic model must be reduced for all practical purposes. Taking an expectation yields a second model described by a system of equations over a large state space, which, while controllable numerically, remain analytically prohibitive. The third model is obtained by averaging over the state space to yield a small system of ODEs. This succession of models provides a blueprint for systematic order reduction of multiscale population models that capture population-level processes while remaining faithful to the underlying stochastic chemical kinetics at the single-cell level. While we will use a particular form for concrete calculations, we try to sketch the derivation generically to demonstrate how it is easily adapted to changes in the underlying biology, both for the single-cell kinetics and the population dynamics.

2.1 Biological setting

We consider an initial population of genetically identical cells constitutively producing a photoreceptive transcription factor. Upon light induction, the transcription factor is recruited in the production of recombinase which leads to genetic recombination and a rewiring of a cell's DNA. Cells that undergo this process have been (irreversibly) differentiated: regions of the genotype not previously used are now available, thereby creating a genetically distinct subpopulation. Only in the differentiated cell construct is the protein of interest produced (to place the burden of production only on the differentiated subpopulation).

We assume that the undifferentiated cells grow and divide at some rate, while the differentiated cells have an inhibited growth rate as a result of the protein production. We thus call the

undifferentiated cells the "growers" and the differentiated cells the "producers". State-dependent growth rates, and their influence on population dynamics, have been widely observed experimentally [2, 6, 8, 14, 23, 27]. We further assume that the transcription factor recruitment and cell recombination is faster than the other transient timescales of the process and may thus be neglected. Previous experimental results show excellent agreement with models adopting both of these assumptions [1, 2]. Nevertheless, we emphasise that our modelling approach is straightforwardly generalisable (at the cost of a larger state space) to the case where we relax these assumptions and track these processes. We assume that the colony is housed in a bioreactor that is run in turbidostat mode, whereby the optical density of the contents (which we assume an accurate proxy for the population density) is kept constant by means of dilution. Finally, we assume that the media is maintained such that nutrients are present in abundance and the culture is well stirred.

2.2 The stochastic heterogeneous population model

To construct the stochastic population process, we begin with the single-cell kinetics described by the chemical master equation [10, 11, 12]. We introduce the three-dimensional state space $\mathcal{X} := \{\mathcal{G}, \mathcal{P}\} \times \mathbb{N}_0 \times \mathbb{N}_0$. The first dimension denotes the cell type of grower or producer. The second and third dimensions correspond to the levels of transcription factor and protein of interest, respectively. We will use the vector notation $\boldsymbol{z} \in \mathcal{X}$ wherever possible for brevity, reverting to the component form $\boldsymbol{z} = (k, x, y)$ where concrete computations are required. The single-cell kinetics are described at time t by the law $Q(\boldsymbol{z}, t) = Q(k, x, y, t)$, whose evolution is governed by the chemical master equation:

$$\frac{\partial}{\partial t} Q(\boldsymbol{z}, t) = -\sum_{E \in \mathcal{F}} E(\boldsymbol{z}, t) Q(\boldsymbol{z}, t) + \sum_{E \in \mathcal{F}} \sum_{\hat{\boldsymbol{z}} \in E^{-1}(\boldsymbol{z})} E(\hat{\boldsymbol{z}}, t) Q(\hat{\boldsymbol{z}}, t), \tag{1}$$

where we have a set of events $E \in \mathcal{F}$ each of which changes the cell state at a rate $E(\boldsymbol{z}, t)$. For a given state $\boldsymbol{z}$ and event $E \in \mathcal{F}$, the set $E^{-1}(\boldsymbol{z})$ denotes all the states $\hat{\boldsymbol{z}}$ for which the event E transforms a cell from state $\hat{\boldsymbol{z}}$ to state $\boldsymbol{z}$. The master equation is of the familiar form: the law evolves by losing all the probability mass associated with changes from the current state to any other, while gaining the probability mass associated with all those states that can transform to the current state.

At the single-cell level we have five events. The differentiation event, as well as four events comprising the birth and death of each of the two quantities accounted for: $\mathcal{F} = \{E_{\text{Diff}}, E_{\text{Birth}}^{\text{TF}}, E_{\text{Death}}^{\text{TF}}, E_{\text{Birth}}^{\text{POI}}, E_{\text{Death}}^{\text{POI}}\}$, where the superscripts TF and POI stand for "transcription factor" and "protein of interest", respectively..

We consider the differentiation rate to be regulated by the light signal $u(t)$ and the transcription factor level x (via some modulation h):

$$E_{\text{Diff}}(k, x, y, t) = u(t) h(x) \delta_{k, \mathcal{G}}, \tag{2}$$

where $\delta_{k, \mathcal{G}}$ denotes the Kronecker delta function, taking the value one when $k = \mathcal{G}$ and zero otherwise. Since differentiation is irreversible, and occurs only to growers, $E_{\text{Diff}}^{-1}(k, x, y)$ is the empty set if $k = \mathcal{G}$, otherwise it is the singleton $\{(\mathcal{G}, x, y)\}$.

We consider birth–death production of the transcription factor (in both growers and producers) and protein of interest (exclusively in the producers) at state-dependent rates:

$$\begin{aligned} E_{\text{Birth}}^{\text{TF}}(k, x, y, t) &= \lambda_k(x), & E_{\text{Death}}^{\text{TF}}(k, x, y, t) &= \mu_k(x), \\ E_{\text{Birth}}^{\text{POI}}(k, x, y, t) &= \lambda(y) \delta_{k, \mathcal{P}}, & E_{\text{Death}}^{\text{POI}}(k, x, y, t) &= \mu(y) \delta_{k, \mathcal{P}}, \end{aligned} \tag{3}$$

where λ and μ represent birth and death rates, respectively, and likewise when the subscript k appears. The corresponding event preimages may be deduced by considering the way each birth and death process changes the cell state. Birth of a transcription factor molecule (in both growers and producers) changes any state from (k, x, y) to $(k, x + 1, y)$. Therefore, the preimage of the TF birth event is the singleton $(E_{\text{Birth}}^{\text{TF}})^{-1}(k, x, y) = \{(k, x - 1, y)\}$. The remaining preimages are calculated analogously, where the signs change for birth and death, and the coordinate of the change becomes y for the POI. While some of these preimages may be outside the state space $\mathcal{X}$, it simplifies the notation substantially to use these preimages and define $Q(\hat{\boldsymbol{z}}, t) = 0$ for all $\hat{\boldsymbol{z}} \in \{\mathcal{G}, \mathcal{P}\} \times \mathbb{Z}^2 \setminus \mathcal{X}$ and $t \geq 0$, rather than introduce indicator functions.

This encapsulates the single-cell process entirely. Now, we introduce a counting process that tracks a population of cells, each independently realising the stochastic dynamics described above [9].
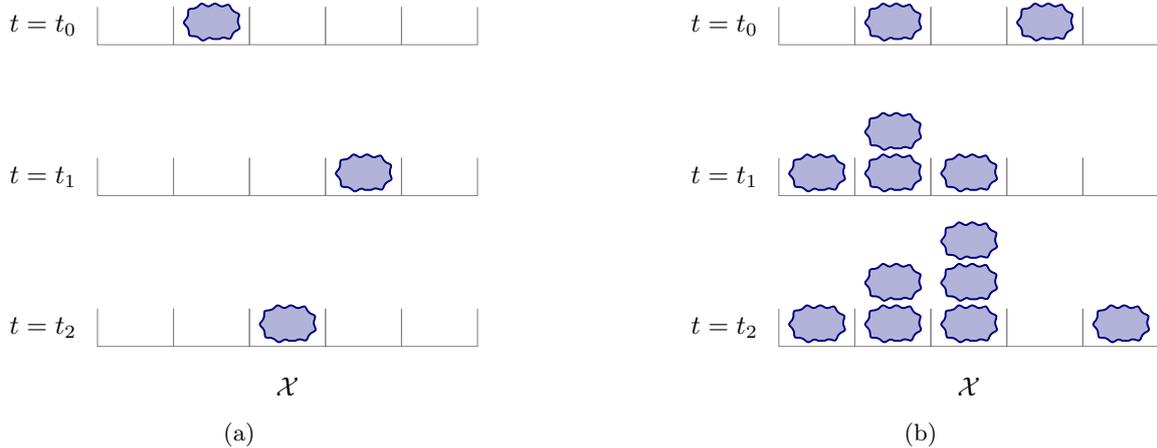
3

Figure 1: Schematic snapshots of (a) a single-cell process; (b) a (population) counting process, at various times. At any time, the single-cell process is described by a member of $\mathcal{X}$ (depicted as a single cell occupying one box in the state space), while the population counting process is described by the number of cells in each state of $\mathcal{X}$ (depicted by some non-negative number of cells in each box of the state space), that is, a mapping belonging to $\mathcal{X} \to \mathbb{N}_0$. For the purposes of illustration, the state space $\mathcal{X}$ is drawn in one dimension.

In addition to the single-cell kinetics, there are population-level dynamics, including the introduction of new cells as existing cells divide, as well as cell removal, to describe continuous bioreactor operation. We assume that these population events are Markovian, just like the single-cell kinetics. It is thus natural to represent the counting process via a master equation, which may be achieved by conceiving of the counting process as a Markov jump process on an extended state space of mappings, namely, $\mathcal{N} := \mathcal{X} \to \mathbb{N}_0$. Just as a single cell, at any time, has a state in $\mathcal{X}$, so too the population described by the counting process is, at any time, in a state within $\mathcal{N}$ (see fig. 1).

For an arbitrary $n \in \mathcal{N}$, the process is described by the law $P(n, t)$, which obeys the master equation

$$\frac{\partial}{\partial t} P(n, t) = -\sum_{E \in \mathcal{E}} \sum_{\boldsymbol{z} \in \mathcal{X}} n(\boldsymbol{z}) E(\boldsymbol{z}, t) P(n, t) + \sum_{E \in \mathcal{E}} \sum_{\boldsymbol{z} \in \mathcal{X}} \sum_{m \in E^{-1}(n, \boldsymbol{z})} m(\boldsymbol{z}) E(\boldsymbol{z}, t) P(m, t), \qquad (4)$$

where $E^{-1}(n, \boldsymbol{z})$ denotes the preimage of states (now a subset of $\mathcal{N}$) that, upon the event $E \in \mathcal{E}$ occurring to a cell in state $\boldsymbol{z}$, yields the state n. Equation (4) is of master-equation form like the more familiar (1). All events, both production at the single-cell level as well as cell division and removal at the population level, occur to individual cells. Therefore, the master equation is the sum of all possible events $E \in \mathcal{E}$ that change the state n, which occurs at a rate of any single event occurring to a single cell $E(\boldsymbol{z}, t)$ multiplied by the number of cells in that state $n(\boldsymbol{z})$.

2.3 The deterministic heterogeneous model

The model (4) is appealing in that it can incorporate the stochasticity of population growth and single-cell kinetics. However, the state space $\mathcal{N}$ is impractically large. Even if we truncate the countably infinite dimensions to some finite size N, this truncated space is still of size N^{2N^2}, which is astronomical for even modest values of N. This explosion of size demands reduction, and motivates considering only the expectation of the counting process at all states $x \in \mathcal{X}$ [28]. We denote this expected population count by

$$q(\boldsymbol{z}, t) := \sum_{n \in \mathcal{N}} n(\boldsymbol{z}) P(n, t), \qquad (5)$$

for all $\boldsymbol{z} \in \mathcal{X}$. We now seek to derive the dynamics of $q(\boldsymbol{z}, t)$ by simplifying the right-hand side of (4) for each event.

First, we lift the single-cell events in $\mathcal{F}$ to the population-level $\mathcal{E}$. The associated rates remain the same and we must identify the associated preimages in the expanded state space. A differentiation event occurring on a cell in state $(\mathcal{G}, x, y)$ leads to that cell becoming a producer in the state $(\mathcal{P}, x, y)$. It follows that the preimage $E_{\mathrm{Diff}}^{-1}(n, k, x, y)$ is empty if $k = \mathcal{P}$ since only growers differentiate. When

$k = \mathcal{G}$, the preimage is equal to the set containing all $m \in \mathcal{N}$ satisfying, for all $(\hat{k}, \hat{x}, \hat{y}) \in \mathcal{X}$,

$$m(\hat{k}, \hat{x}, \hat{y}) = n(\hat{k}, \hat{x}, \hat{y}) + \delta_{(\hat{k},\hat{x},\hat{y}),(\mathcal{G},x,y)} - \delta_{(\hat{k},\hat{x},\hat{y}),(\mathcal{P},x,y)}. \tag{6}$$

For any state $n \in \mathcal{N}$, there can be at most one such state m in the preimage, given by considering (6) as the definition of m. For n with $n(\mathcal{P}, x, y) = 0$ the definition (6) gives a mapping $m \notin \mathcal{N}$ since $m(\mathcal{P}, x, y) = -1$. Therefore, strictly speaking, the preimage state is empty. However, just as in the single-cell case, it is more convenient to consider the preimage to be the singleton containing m from (6) and adopt the convention that $P(m, t) = 0$ for all $m \in (\mathcal{X} \to \mathbb{Z}) \setminus \mathcal{N}$ and all $t \geq 0$. This convention avoids introducing notation to explicitly accommodate these empty preimages. Moreover, since the differentiation rate is zero for $k = \mathcal{P}$ (see (2)), it suffices to take the singleton containing m in (6) as the defining preimage for all $(k, x, y) \in \mathcal{X}$.

All preimages in this model are singletons (using the convention described above), since each event is deterministic. One could also consider events whose outcomes are stochastic, such as production in bursts or copy numbers with stochastic cell division laws, which, while requiring a more involved derivation, is not conceptually more complicated. Since this is the case, we may remove the sum over the preimage set in (4), and replace m with the singleton member, which we denote $m_{n,k,x,y}$ (or, equivalently, $m_{n,\boldsymbol{z}}$) for each event, highlighting the dependence on n and $(k, x, y) = \boldsymbol{z}$ but suppressing the dependence on the event (which will be clear from the context). The master equation (4) then becomes

$$\frac{\partial}{\partial t} P(n, t) = -\sum_{E \in \mathcal{E}} \sum_{\boldsymbol{z} \in \mathcal{X}} n(\boldsymbol{z}) E(\boldsymbol{z}, t) P(n, t) + \sum_{E \in \mathcal{E}} \sum_{\boldsymbol{z} \in \mathcal{X}} m_{n,\boldsymbol{z}}(\boldsymbol{z}) E(\boldsymbol{z}, t) P(m_{n,\boldsymbol{z}}, t). \tag{7}$$

For the event of the birth of a single molecule of transcription factor, the preimage satisfies, for all $(\hat{k}, \hat{x}, \hat{y}) \in \mathcal{X}$,

$$m_{n,k,x,y}(\hat{k}, \hat{x}, \hat{y}) = n(\hat{k}, \hat{x}, \hat{y}) + \delta_{(\hat{k},\hat{x},\hat{y}),(k,x,y)} - \delta_{(\hat{k},\hat{x},\hat{y}),(k,x+1,y)}. \tag{8}$$

The above comment regarding some such m not strictly belonging to $\mathcal{N}$ applies here too, mutatis mutandis, and we do not mention this again.

For the population-level process, we add the birth–death event rates a_k for cell division and Λ for cell removal from the bioreactor:

$$E_{\text{Divide}}(k, x, y, t) = a_k(y), \qquad\qquad E_{\text{Remove}}(k, x, y, t) = \Lambda(t). \tag{9}$$

The associated preimages are also singletons for the same reason as above. We consider a dividing (mother) cell to produce two daughters *of the same state*. If the species we are counting in the x and y variables are copy numbers, we should expect cell division to partition the mother cell's content into the two daughter cells, and thus typically there is to be lower levels of x and y in both daughters (compared to the mother). However, if we the species levels are concentrations, this rule aligns with laboratory observations [21, 24]. Thus, for cell division events, the preimage satisfies

$$m_{n,k,x,y}(\hat{k}, \hat{x}, \hat{y}) = n(\hat{k}, \hat{x}, \hat{y}) - \delta_{(\hat{k},\hat{x},\hat{y}),(k,x,y)}, \tag{10}$$

while for cell removal, the preimage satisfies

$$m_{n,k,x,y}(\hat{k}, \hat{x}, \hat{y}) = n(\hat{k}, \hat{x}, \hat{y}) + \delta_{(\hat{k},\hat{x},\hat{y}),(k,x,y)}, \tag{11}$$

for all $(\hat{k}, \hat{x}, \hat{y}) \in \mathcal{X}$.

By differentiating (5) with respect to time and substituting the master equation (7), we find that the expected cell density satisfies

$$\begin{aligned}
\frac{\partial}{\partial t} q(\boldsymbol{z}, t) &= \sum_{n \in \mathcal{N}} n(\boldsymbol{z}) \frac{\partial}{\partial t} P(n, t) \\
&= -\sum_{n \in \mathcal{N}} n(\boldsymbol{z}) \sum_{E \in \mathcal{E}} \sum_{\hat{\boldsymbol{z}} \in \mathcal{X}} n(\hat{\boldsymbol{z}}) E(\hat{\boldsymbol{z}}, t) P(n, t) + \sum_{n \in \mathcal{N}} n(\boldsymbol{z}) \sum_{E \in \mathcal{E}} \sum_{\hat{\boldsymbol{z}} \in \mathcal{X}} m_{n,\boldsymbol{z}}(\hat{\boldsymbol{z}}) E(\hat{\boldsymbol{z}}, t) P(m_{n,\boldsymbol{z}}, t) \\
&= \sum_{E \in \mathcal{E}} \left(-\sum_{n \in \mathcal{N}} n(\boldsymbol{z}) \sum_{\hat{\boldsymbol{z}} \in \mathcal{X}} n(\hat{\boldsymbol{z}}) E(\hat{\boldsymbol{z}}, t) P(n, t) + \sum_{n \in \mathcal{N}} n(\boldsymbol{z}) \sum_{\hat{\boldsymbol{z}} \in \mathcal{X}} m_{n,\boldsymbol{z}}(\hat{\boldsymbol{z}}) E(\hat{\boldsymbol{z}}, t) P(m_{n,\boldsymbol{z}}, t) \right) \\
&=: \sum_{E \in \mathcal{E}} \Phi_E(\boldsymbol{z}, t).
\end{aligned} \tag{12}$$

We now derive expressions for each event contribution Φ_E in terms of the expected population count. For differentiation, using (2), (6) and (12), we find that

$$
\begin{aligned}
\Phi_{E_{\text{Diff}}}(k,x,y,t) &= -\sum_{n\in\mathcal{N}} n(k,x,y) \sum_{(\hat{k},\hat{x},\hat{y})\in\mathcal{X}} n(\hat{k},\hat{x},\hat{y})u(t)h(\hat{x})\delta_{\hat{k},\mathcal{G}}P(n,t) \\
&\quad + \sum_{n\in\mathcal{N}} n(k,x,y) \sum_{(\hat{k},\hat{x},\hat{y})\in\mathcal{X}} m_{n,\hat{k},\hat{x},\hat{y}}(\hat{k},\hat{x},\hat{y})u(t)h(\hat{x})\delta_{\hat{k},\mathcal{G}}P(m_{n,\hat{k},\hat{x},\hat{y}},t) \\
&= -\sum_{n\in\mathcal{N}}\sum_{(\hat{x},\hat{y})} n(k,x,y)n(\mathcal{G},\hat{x},\hat{y})u(t)h(\hat{x})P(n,t) \\
&\quad + \sum_{n\in\mathcal{N}}\sum_{(\hat{x},\hat{y})} [m_{n,\mathcal{G},\hat{x},\hat{y}}(k,x,y) - \delta_{(k,x,y),(\mathcal{G},\hat{x},\hat{y})} + \delta_{(k,x,y),(\mathcal{P},\hat{x},\hat{y})}] \times \\
&\qquad\qquad\qquad\qquad\qquad m_{n,\mathcal{G},\hat{x},\hat{y}}(\mathcal{G},\hat{x},\hat{y})u(t)h(\hat{x})P(m_{n,\mathcal{G},\hat{x},\hat{y}},t) \\
&= \sum_{n\in\mathcal{N}}\sum_{(\hat{x},\hat{y})} [-\delta_{(k,x,y),(\mathcal{G},\hat{x},\hat{y})} + \delta_{(k,x,y),(\mathcal{P},\hat{x},\hat{y})}]m_{n,\mathcal{G},\hat{x},\hat{y}}(\mathcal{G},\hat{x},\hat{y})u(t)h(\hat{x})P(m_{n,\mathcal{G},\hat{x},\hat{y}},t) \\
&= -u(t)h(x)\delta_{k,\mathcal{G}}\,q(\mathcal{G},x,y,t) + u(t)h(x)\delta_{k,\mathcal{P}}\,q(\mathcal{G},x,y,t).
\end{aligned}
\tag{13}
$$

The last two equalities rely on the fact that, for any functional f, the following two sums are equal:

$$
\begin{aligned}
\sum_{n\in\mathcal{N}}\sum_{(\hat{x},\hat{y})} &f(n,\hat{x},\hat{y})n(\mathcal{G},\hat{x},\hat{y})P(n,t) \\
&= \sum_{n\in\mathcal{N}}\sum_{(\hat{x},\hat{y})} f(m_{n,\mathcal{G},\hat{x},\hat{y}},\hat{x},\hat{y})m_{n,\mathcal{G},\hat{x},\hat{y}}(\mathcal{G},\hat{x},\hat{y})P(m_{n,\mathcal{G},\hat{x},\hat{y}},t).
\end{aligned}
\tag{14}
$$

Equality (14) holds if and only if it holds with $f(n,\hat{x},\hat{y}) = \delta_{(n,\hat{x},\hat{y}),(\hat{n},x,y)}$ for all $\hat{n}\in\mathcal{N}$ and $x,y\in\mathbb{N}_0$. With this choice of f, the left-hand side reduces to $\hat{n}(\mathcal{G},x,y)P(\hat{n},t)$, while the right-hand side reduces to the same thing if the state $\hat{n}$ is the preimage of some state, otherwise it vanishes. Therefore, to prove the equality holds, it suffices to show that $\hat{n}(\mathcal{G},x,y)P(\hat{n},t) = 0$ for all $\hat{n}$ that are not preimages of any state. Intuitively, states are not preimages (with respect to differentiation) only when there are no growers: $\hat{n}(\mathcal{G},x,y) = 0$. Formally, taking such a state $\hat{n}\in\mathcal{N}\setminus\{m_{n,\mathcal{G},\hat{x},\hat{y}} \mid n\in\mathcal{N},\ \hat{x},\hat{y}\in\mathbb{N}_0\}$, we may define the mapping n via (6) by replacing m with $\hat{n}$. If $\hat{n}(\mathcal{G},x,y) > 0$ then $n(\mathcal{G},x,y)\geq 0$ and thus $n\in\mathcal{N}$. It then follows that $m_{n,\mathcal{G},x,y} = \hat{n}$, the preimage of this constructed n coincides with $\hat{n}$, which contradicts the fact that $\hat{n}$ is not the preimage of any state, and the claim follows. Physically, equality (14) reflects the fact that each event is accounted for as a transformation *from* an original state, and as a transformation *to* a different state.

We now proceed to the event contribution for TF birth $\Phi_{E_{\text{Birth}}^{\text{TF}}}$. Using (3), (8) and (12), we find that

$$
\begin{aligned}
\Phi_{E_{\text{Birth}}^{\text{TF}}}(k,x,y,t) &= -\sum_{n\in\mathcal{N}}\sum_{(\hat{k},\hat{x},\hat{y})} n(k,x,y)n(\hat{k},\hat{x},\hat{y})\lambda_{\hat{k}}(\hat{x})P(n,t) \\
&\quad + \sum_{n\in\mathcal{N}}\sum_{(\hat{k},\hat{x},\hat{y})} n(k,x,y)m_{n,\hat{k},\hat{x},\hat{y}}(\hat{k},\hat{x},\hat{y})\lambda_{\hat{k}}(\hat{x})P(m_{n,\hat{k},\hat{x},\hat{y}},t) \\
&= \sum_{n\in\mathcal{N}}\sum_{(\hat{k},\hat{x},\hat{y})} [-\delta_{(k,x,y),(\hat{k},\hat{x},\hat{y})} + \delta_{(k,x,y)(\hat{k},\hat{x}+1,\hat{y})}]m_{n,\hat{k},\hat{x},\hat{y}}(\hat{k},\hat{x},\hat{y})\lambda_{\hat{k}}(\hat{x})P(m_{n,\hat{k},\hat{x},\hat{y}},t) \\
&= -\lambda_k(x)q(k,x,y,t) + \lambda_k(x-1)q(k,x-1,y,t).
\end{aligned}
\tag{15}
$$

Here too the final two equalities rely upon the fact that sums evaluating forms of n or $m_{n,\hat{k},\hat{x},\hat{y}}$ are equal, differing only in states that do not contribute to the sum, as above.

For any $\boldsymbol{z}\in(\{\mathcal{G},\mathcal{P}\}\times\mathbb{Z}^2)\setminus\mathcal{X}$, we adopt the convention that $q(\boldsymbol{z},t) = 0$, which is tantamount to imposing boundary conditions that there is no population density outside the state space.

For the remaining single-cell birth–death events, without repeating the arguments, we find that

$$
\Phi_{E_{\text{Death}}^{\text{TF}}}(k,x,y,t) = -\mu_k(x)q(k,x,y,t) + \mu_k(x+1)q(k,x+1,y,t),
\tag{16}
$$

$$
\Phi_{E_{\text{Birth}}^{\text{POI}}}(k,x,y,t) = \delta_{k,\mathcal{P}}\left[-\lambda(y)q(k,x,y,t) + \lambda(y-1)q(k,x,y-1,t)\right],
\tag{17}
$$

$$
\Phi_{E_{\text{Death}}^{\text{POI}}}(k,x,y,t) = \delta_{k,\mathcal{P}}\left[-\mu(y)q(k,x,y,t) + \mu(y+1)q(k,x+1,y,t)\right].
\tag{18}
$$

For the population-level processes of cell division and removal the calculations are analogous. Using (9), (10) and (12), we see that

$$
\begin{aligned}
\Phi_{E_{\text{Divide}}}(k,x,y,t) = & -\sum_{n\in\mathcal{N}} n(k,x,y) \sum_{(\hat{k},\hat{x},\hat{y})} n(\hat{k},\hat{x},\hat{y}) a_{\hat{k}}(\hat{y}) P(n,t) \\
& + \sum_{n\in\mathcal{N}} n(k,x,y) \sum_{(\hat{k},\hat{x},\hat{y})} m_{n,k,x,y}(\hat{k},\hat{x},\hat{y}) a_{\hat{k}}(\hat{y}) P(m_{n,k,x,y},t) \\
= & -\sum_{n\in\mathcal{N}} \sum_{(\hat{k},\hat{x},\hat{y})} n(k,x,y) n(\hat{k},\hat{x},\hat{y}) a_{\hat{k}}(\hat{y}) P(n,t) \\
& + \sum_{n\in\mathcal{N}} \sum_{(\hat{k},\hat{x},\hat{y})} [m_{n,\hat{k},\hat{x},\hat{y}}(k,x,y) + \delta_{(k,x,y),(\hat{k},\hat{x},\hat{y})}] m_{n,\hat{k},\hat{x},\hat{y}}(\hat{k},\hat{x},\hat{y}) a_{\hat{k}}(\hat{y}) P(m_{n,\hat{k},\hat{x},\hat{y}},t) \\
= & \sum_{n\in\mathcal{N}} \sum_{(\hat{k},\hat{x},\hat{y})} \delta_{(k,x,y),(\hat{k},\hat{x},\hat{y})} m_{n,\hat{k},\hat{x},\hat{y}}(\hat{k},\hat{x},\hat{y}) a_{\hat{k}}(\hat{y}) P(m_{n,\hat{k},\hat{x},\hat{y}},t) \\
= & \; a_k(y) q(k,x,y,t).
\end{aligned}
\tag{19}
$$

Similarly for removal, we find that

$$
\Phi_{E_{\text{Remove}}}(k,x,y,t) = -\Lambda(t) q(k,x,y,t).
\tag{20}
$$

Substituting the forms from (13) to (20) into (12), we obtain the governing dynamics for the expected population count $q(x,y,t)$. To simplify the notation, we define the shorthands

$$
g(x,y,t) := q(\mathcal{G},x,y,t), \qquad\qquad p(x,y,t) := q(\mathcal{P},x,y,t),
\tag{21}
$$

for which we have

$$
\begin{aligned}
\frac{\partial}{\partial t} g(x,y,t) = & \; [a_{\mathcal{G}}(y) - \Lambda(t)] g(x,y,t) - u(t) h(x) g(x,y,t) \\
& - [\lambda_{\mathcal{G}}(x) + \mu_{\mathcal{G}}(x)] g(x,y,t) + \lambda_{\mathcal{G}}(x-1) g(x-1,y,t) + \mu_{\mathcal{G}}(x+1) g(x+1,y,t), \\
\frac{\partial}{\partial t} p(x,y,t) = & \; [a_{\mathcal{P}}(y) - \Lambda(t)] p(x,y,t) + u(t) h(x) g(x,y,t) \\
& - [\lambda_{\mathcal{P}}(x) + \mu_{\mathcal{P}}(x)] p(x,y,t) + \lambda_{\mathcal{P}}(x-1) p(x-1,t) + \mu_{\mathcal{P}}(x+1) p(x+1,y,t) \\
& - [\lambda(y) + \mu(y)] p(x,y,t) + \lambda(y-1) p(x,y-1,t) + \mu(y+1) p(x,y+1,t).
\end{aligned}
\tag{22}
$$

Ultimately, the model (22) recovers the growth and removal forms alongside the single-cell dynamics, all acting on the population density. Thus, by studying the dynamics of the *expected* population density, we have retained heterogeneity driven by stochastic single-cell kinetics alongside contributions from all physical processes, while reducing the model complexity to that of the single-cell dynamics (for example, compare (1) and (22)). This order reduction averages over population-level stochasticity, however, no approximations or additional assumptions are employed. We consolidate this method in the following proposition.

Proposition 1. *For the multiscale stochastic process described by the master equation* (4) *with events as defined in this section, the expected population density* (5) *is governed* exactly *by the law* (22).

The calculations in this section are formal. We have not established the well-posedness of the master equation (4). Similarly, not every solution (even if these exist) has a finite expectation (5). While it is not our aim to deal with these issues in depth, a brief comment is worthwhile. For the master equation (4), a finite truncation of the state space $\mathcal{N}$ (at some maximum number of single-cell states and maximum cell population size, while again adopting the convention of setting probabilities to zero beyond this truncated state space) suffices to reduce the master equation to a finite-dimensional system of linear ODEs guaranteeing well-posedness. Since all physical quantities are expected to remain bounded, the truncated model remains physically sensible. This truncation similarly produces a finite expectation and guarantees that system (5) is finite and linear, thus being well posed.

Truncating only the cell population size (but retaining an infinite single-cell state space $\mathcal{X}$) and assuming the master equation is well posed suffice to guarantee a finite expectation (5). That is, given the solution of the master equation on the truncated state space $\{n \in \mathcal{N} \mid \sum_{z\in\mathcal{X}} n(z) < N\}$,

for some $N < \infty$, the expected population density is well-defined for all time. Assuming that the transition rates are uniformly bounded in space, and noting that all summands are non-negative, we infer that all infinite sums are absolutely convergent, which justifies their rearrangement in deriving the infinite-dimensional system (22). Under these conditions, the operator in system (22) is uniformly Lipschitz, which guarantees its well-posedness.

2.4   The deterministic homogeneous model

The model (22) is numerically tractable for a small enough state space, but quickly becomes cumbersome for a realistically large state space. In this case, by classical methods [18, 19], one may derive a PDE approximation of system (22) which allows for tractable numerical approximations [21]. However, to gain deeper analytical insight it is prudent to further reduce the system. We study the case where growth rates depend only on the cell phenotype:

$$a_{\mathcal{G}}(y) := A, \qquad\qquad a_{\mathcal{P}}(y) := a. \tag{23}$$

This choice, with $A > a$ models uniform growth of each subpopulation, reflecting the growth burden of production of the protein of interest (rather than the toxicity of the protein of interest, which might still depend on the state y).

Defining the zeroth-order moments (recycling the use of the symbol P)

$$G(t) := \sum_{(x,y)} g(x,y,t), \qquad\qquad P(t) := \sum_{(x,y)} p(x,y,t), \tag{24}$$

and formally summing system (22) over all (x, y), we find that

$$\dot{G}(t) = [A - \Lambda(t)]G(t) - u(t) \sum_{(x,y)} h(x)g(x,y,t), \tag{25a}$$

$$\dot{P}(t) = [a - \Lambda(t)]P(t) + u(t) \sum_{(x,y)} h(x)g(x,y,t). \tag{25b}$$

We employ the quasi-stationary closure, whereby we assume that the normalised distribution $g(x,y,t)/G(t)$ is approximately stationary (see, e.g., Ref. [21]), thus the system (25) reduces to

$$\dot{G}(t) = [A - \Lambda(t)]G(t) - Fu(t)G(t), \tag{26a}$$

$$\dot{P}(t) = [a - \Lambda(t)]P(t) + Fu(t)G(t), \tag{26b}$$

where $F \approx \sum_{(x,y)} h(x)g(x,y,t)/\sum_{(x,y)} g(x,y,t)$ represents the global differentiation rate for the maximal light input.

We have thus arrived at a system of two ordinary-differential equations, akin to the dynamical systems models employed previously in the literature [1, 22, 26, 31].

2.5   Synopsis

It is instructive at this stage to take stock. We have derived a hierarchy of models of decreasing complexity (and fidelity): stochastic and heterogeneous, deterministic and heterogeneous, deterministic and homogeneous. The derivation is easily adaptable to similar models with stochastic populations of individual agent, irrespective of the number of microorganisms in the consortium and the number of species tracked in their internal states. The underlying stochastic population model was introduced in Ref. [9]. A deterministic heterogeneous counterpart was derived in Ref. [28], although using different techniques, and the subsequent deterministic homogeneous model was studied in Ref. [20, 21]. To the best of our knowledge, this is the first time these techniques have been combined in an interconnected hierarchy.

Which model to choose depends on the modeler's particular requirements, which motivates a discussion of the relative advantages of each model. As noted, the stochastic model has a state space too large to allow us to calculate the law of the process. If information beyond the expected population density is required, for example, higher-order statistical information, potentially relevant when only a small number of realisations are planned, one could develop the dynamics for higher-order moments in an analogous fashion. Alternatively, one could use stochastic sampling to approximate the law in a Monte Carlo fashion. We refer the reader to Ref. [9], where these approaches are studied in greater depth.

In many practical applications deterministic models suffice. We have labelled these models "heterogeneous" and "homogeneous" but emphasise that in both models we preserve the distinction between different microorganisms that make up the consortium. The label reflects whether the heterogeneity present *within each subpopulation*, driven by stochastic chemical kinetics, is preserved. In some scenarios, ignoring this modelling ingredient results in substantially diminished control performance [20, 21], while in other practical scenarios it may be safely ignored [32].

Even when heterogeneity is a crucial model feature, the coarser homogeneous model can inform parameter choice and initial control guesses for the heterogeneous model, which can save valuable computational effort [21]. Motivated by this salience, we proceed to analyse the deterministic homogeneous system (26).

3 Optimal consortia control

3.1 Problem formulation

For the optimal control of a microbial consortium, we begin with the reduced-order model (26), to which we add arbitrary initial conditions. In practice, the culture is typically grown in darkness so that the initial population is of known composition. Nevertheless, we consider the more general problem of arbitrary initial conditions. This is because we are also interested in using the solution with Model–Predictive Control (MPC), where state updates can be made throughout the bioproduction process, and the optimal control recalculated from the intermediate state. Thus, we have

$$\dot{G}(t) = AG(t) - \Lambda(t)G(t) - Fu(t)G(t), \qquad G(0) = g_0, \qquad (27a)$$

$$\dot{P}(t) = aP(t) - \Lambda(t)P(t) + Fu(t)G(t), \qquad P(0) = p_0, \qquad (27b)$$

where $\Lambda(t)$ is the turbidostat dilution rate, $u(t) \in [0,1]$ is the (dimensionless) light signal that induces differentiation, and $F > 0$ is the differentiation rate for the maximal light signal. If there is significant leakage in the differentiation construct, that is, differentiation occurs even in the absence of light, a basal rate of differentiation should be included in the model. It has been shown experimentally that strains can be engineered with negligible leakage, justifying the present model (see Supplementary Table 2 of Ref. [1]).

We may determine the turbidostat dilution rate by fixing $\Lambda(t)$ such that the total population mass remains constant. Adding (27a) and (27b) and taking $\dot{G}(t) + \dot{P}(t) = 0$ we find that

$$\Lambda(t) = \frac{AG(t) + aP(t)}{G(t) + P(t)} = \frac{AG(t) + aP(t)}{g_0 + p_0}. \qquad (27c)$$

We seek to maximise the protein yield J, which we take to be proportional to the producer population mass from both the diluted media and what remains in the bioreactor at the terminal time $t = T$. The objective function takes the form

$$J = P(T) + \int_0^T \Lambda(t)P(t)\,\mathrm{d}t. \qquad (27d)$$

We can, without loss of generality, take $g_0 + p_0 = 1$, since this amounts to rescaling $G(t)$ and $P(t)$ by the denominator $g_0 + p_0$, whereby the dilution rate is then given by $\Lambda(t) = AG(t) + aP(t)$. We thus obtain the unit conservation property

$$G(t) + P(t) = 1, \qquad (28)$$

from which we reduce the optimal control problem (27) to a single dimension. The problem then reads: maximise

$$J = -G(T) + \int_0^T (A - 2a)G(t) - (A - a)G(t)^2\,\mathrm{d}t, \qquad (29a)$$

for $0 \le u(t) \le 1$, where we have neglected terms that are constant for all controls, subject to the nonlinear dynamics

$$\dot{G}(t) = (A - a - Fu(t))G(t) - (A - a)G(t)^2, \qquad G(0) = g_0 \in [0,1]. \qquad (29b)$$

At the outset, we may make predictions of the optimal control under certain assumptions based on the turnpike property. Since the running component of the objective is a quadratic (thus convex) function of the state, we expect the optimal control to be of turnpike form [29] assuming the time horizon is long enough and the system is controllable. That is, for most of the time horizon, the optimal control should be that which guarantees the state is maximising the quadratic form of the running pay-off. The only deviations from this may be an initial transient where this steady state is established, and a terminal transient where the terminal component of the objective is maximised.

The subsequent analysis agrees with this characterisation, however, it refines and extends it in important ways. First, we deal with the cases of insufficient growth arrest (where the steady state maximising the running pay-off is not reachable) and weak optogenetic induction (where the system is not sufficiently controllable). Second, we deal with arbitrary time horizons, in particular, those that are insufficiently long for the turnpike structure to form. Finally, in all cases we are able to pin down the optimal control structure quantitatively. For example, in the case of strong induction, we have an explicit expression for the critical time horizon at which the turnpike structure manifests.

3.2 Optimal control solution

We begin by studying the nonlinear dynamics of the state $G(t)$. This helps us develop an intuition that can inform our expectations of the control. First, we clarify the bounds of the state dynamics, which provides guarantees on the value of $G(t)$. It will also prove useful in subsequent proofs.

Proposition 2. *The sets $\{0\}$ and $(0, 1]$ are invariant sets of the dynamics of $G(t)$.*

Proof. If $G(t) = 0$, then $\dot{G}(t) = 0$. If $G(t) = 1$, then $\dot{G}(t) \leq 0$. It thus suffices to prove that zero cannot be reached if $0 < G(t) \leq 1$. Since $G(t) \in [0, 1]$, the dynamics (29b) may be written as $\dot{G}(t) = -\gamma(t)G(t)$ where $\gamma(t) \leq F$. An application of Grönwall's inequality suffices to show that the solution is bounded from below by an exponentially decreasing function and thus must remain positive for all time. □

We conclude from Proposition 2 that, for $g_0 > 0$, $G(t) > 0$ for all time t. Moreover, the case of $g_0 = 0$ is trivial, and is not subsequently considered.

We now consider the dynamics (29b) for a constant control $u(t) = u$. In this case, the dynamics are homogeneous, and the quadratic form on the right-hand side of (29b) reveals two equilibria, $G(t) = G_0$ and $G(t) = G_1(u)$, given by

$$G_0 = 0, \qquad\qquad G_1(u) = 1 - \frac{Fu}{A - a}. \qquad\qquad (30)$$

The stability of these equilibria is governed by the sign of $\dot{G}(t)$ in their vicinity. The quadratic form (29b) dictates that, when these two roots are distinct, the lesser is unstable while the greater is stable. The roots coincide for

$$u = \frac{A - a}{F} =: u_0. \qquad\qquad (31)$$

When $u > u_0$, $G_1(u) < G_0$ and the stable root is G_0, while when $u < u_0$ the sign of $G_1(u)$ switches and it becomes the stable root. When G_0 is a stable equilibrium, the grower population $G(t)$ cannot vanish in finite time (Proposition 2).

For the degenerate control value $u = u_0$ the roots coincide and the unique equilibrium $G(t) = G_0 = G_1(u_0) = 0$ is unstable to arbitrary perturbations. Nevertheless, being a population mass, $G(t)$ and its perturbations must be non-negative, and the zero equilibrium is stable to positive perturbations.

It is helpful to develop a physical interpretation of these conditions to complement the analysis. We found a critical control value u_0 at (and beyond) which the grower population tends to zero (whereby the population comprises exclusively producers) as the differentiation, which depletes the grower population, surpasses the influence of the growth. For subcritical controls $u < u_0$, the population tends to a strictly positive equilibrium $G(t) = G_1(u)$ where the differentiation and growth balance to maintain a steady growers population.

Before tackling the optimal control problem, we present a result regarding the nonlinear dynamics with an arbitrary time-varying control input that will prove useful in a later proof.

Proposition 3. *When $G_1(1) > 0$, the state $G(t)$ increases monotonically in $(0, G_1(1))$, while the interval $[G_1(1), 1]$ is an invariant set of the dynamics of $G(t)$ for any $u(t) \in [0, 1]$.*

Proof. If $G(t) \in (0, G_1(1))$ then $\dot{G}(t) \geq (A - a - F)G(t) - (A - a)G(t)^2 > 0$, where the final inequality comes from minimising the quadratic over the domain $(0, G_1(1))$.

If $G(t) = G_1(1)$, then $\dot{G}(t) \geq 0$, and the claim follows from Proposition 2. □

We now study the optimal control of problem (29) via Pontryagin's maximum principle (for a concise and accessible introduction we refer the reader to Ref. [17]). We introduce the Hamiltonian (with the convention of minimising $-J$),

$$H = \lambda(t) \left[(A - a - Fu(t))G(t) - (A - a)G(t)^2 \right] - (A - 2a)G(t) + (A - a)G(t)^2, \qquad (32)$$

from which we determine that the costate equation is given by

$$\dot{\lambda}(t) = -\lambda(t)[A - a - Fu(t) - 2(A - a)G(t)] + (A - 2a) - 2(A - a)G(t), \qquad \lambda(T) = 1. \qquad (33)$$

The Hamiltonian is linear with respect to the control, and therefore the optimal control is given by the bang–bang form when the switching function is nonzero [15]. The switching function is given by

$$\frac{\partial H}{\partial u} = -F\lambda(t)G(t), \qquad (34)$$

and since $F > 0$ and $G(t) > 0$ for all time, the optimal control takes the bang–bang form

$$u(t) = \begin{cases} 1, & \lambda(t) > 0, \\ 0, & \lambda(t) < 0. \end{cases} \qquad (35)$$

It remains to determine the optimal control on singular arcs, that is, nontrivial intervals on which $\lambda(t) = 0$. We denote the bang–bang arcs of control (35) by $\mathcal{B}_-$ and $\mathcal{B}_+$ arcs, for $u(t) = 0$ and $u(t) = 1$, respectively, while we denote a singular arc by $\mathcal{S}$.

Further analysis of the adjoint equation (33) allows us to determine the structure of the optimal control.

Proposition 4. *There can be at most one $\mathcal{B}_-$ arc. This arc, if it exists, must be adjacent to $t = 0$, that is, it must be the time interval $[0, t_0)$ for some $t_0 > 0$.*

Proof. Assume that there exists a nontrivial arc (t_0, t_1), where $0 < t_0 < t_1 < T$, on which $\lambda(t) < 0$ and $\lambda(t_0) = \lambda(t_1) = 0$. It must be that $\dot{\lambda}(t_0) \leq 0 \leq \dot{\lambda}(t_1)$, and it follows from the costate equation (33) that $\dot{\lambda}(t_i) = 2(A - a)(G_s - G(t_i))$ for both $i = 0, 1$ where

$$G_s := \frac{A - 2a}{2(A - a)} < 1. \qquad (36)$$

Therefore, $G(t_0) \geq G(t_1)$. However, the state must be monotonically non-decreasing over the interval since $\dot{G}(t) \geq 0$ when $u(t) = 0$ with equality only for $G(t) = 1$ (since we are not interested in the $G(t) = 0$ case). It follows that $G(t) = 1$ over this entire interval, and thus $\dot{\lambda}(t_1) = 2(A - a)(G_s - 1) < 0$, which is a contradiction.

This demonstrates that there cannot exist such an arc for $0 < t_0 < t_1 < T$. We cannot set $t_1 = T$, since the terminal condition in (33) will not be satisfied. Thus, the only way an arc of $\lambda(t) > 0$ can exist is if $t_0 = 0$. This arc, if it exists, must therefore be unique. □

It is instructive to characterise the singular arcs, as in the following result.

Proposition 5. *On singular arcs, the optimal control is $u(t) = A/(2F)$, and satisfies the strengthened Generalised Legendre-Clebsch (GLC) condition.*

Proof. Assume that the switching function (34) vanishes on some interval. This can occur when either $G(t) = 0$ or $\lambda(t) = 0$. We proved in Proposition 2 that the former cannot happen in finite time (since $g_0 > 0$), and thus focus on the latter. On such intervals, it must hold that $\dot{\lambda}(t) = 0$ and thus from (33) we find that $G(t) = G_s$ defined in (36). Differentiating, $\dot{G}(t) = 0$, from which it follows via (29b) that the singular optimal control is given by

$$u(t) = \frac{A}{2F} =: u_s. \qquad (37)$$

This allows us to calculate the GLC criterion

$$-\frac{\partial}{\partial u} \left(\frac{d^2}{dt^2} \frac{\partial H}{\partial u} \right) = \frac{F^2(A - 2a)^2}{2(A - a)} \geq 0. \qquad (38)$$

Since the state $G(t)$ is positive for all time (Proposition 2), a singular arc, where $G(t) = G_s$, can only exist if $G_s > 0$, that is, $A > 2a$, in which case the expression in (38) will be strictly positive. This is none other than the strengthened GLC condition, also known as the Kelley condition. □

We have ascertained that there is a unique initial (possibly empty) $\mathcal{B}_-$ arc, on $[0, t_0)$, say. There is also a terminal $\mathcal{B}_+$ arc due to the bang–bang form (35) and the terminal condition (33), on $(T - T_s, T]$, say. By continuity of the costate we may assume that this final arc is maximal. When $T - T_s > t_0$, we are interested in pinning down the optimal control structure in the intervening time. We now prove that there can be no additional $\mathcal{B}_+$ arc spanning (t', t''), for some $t_0 < t' < t'' \leq T - T_s$. This demonstrates that, between these initial $\mathcal{B}_-$ and terminal $\mathcal{B}_+$ arcs, there can only be another $\mathcal{B}_+$ arc if it is adjacent to the initial $\mathcal{B}_-$ arc, that is, $t' = t_0$. Since there can be no more than the initial $\mathcal{B}_-$ arc, and no more than these two potential $\mathcal{B}_+$ arcs, the optimal control must of the form $(\mathcal{B}_-)(\mathcal{B}_+)(\mathcal{S})\mathcal{B}_+$, where parentheses enclose arc forms that may be empty. It follows that the singular arc, if it exists, is unique.

Proposition 6. *Considering the initial $\mathcal{B}_-$ arc to span $[0, t_0)$ for some $t_0 \geq 0$, and the terminal $\mathcal{B}_+$ arc to maximally span $(T - T_s, T]$ for some $T_s > 0$, there can be no $\mathcal{B}_+$ arc maximally spanning (t', t''), for any $t_0 < t' < t'' \leq T - T_s$.*

Proof. By way of contradiction, consider such a $\mathcal{B}_+$ arc. The maximality of the arc guarantees that the times t' and t'' are critical times, that is, $\lambda(t') = \lambda(t'') = 0$. Moreover, since $t' > t_0$ by assumption, and there cannot be additional $\mathcal{B}_-$ arcs, it must also hold, by (35), that $\dot{\lambda}(t') = \dot{\lambda}(t'') = 0$. It then follows from the costate equation (33) that $G(t') = G(t'') = G_s$.

On a $\mathcal{B}_+$ arc the control is maximal $u(t) = 1$. In this case, the state evolves strictly monotonically towards the stable equilibrium, G_0 or $G_1(1)$ defined in (30), unless the state $G(t') = G_s$ coincides with one of the equilibria. Since $G(t') = G(t'')$, it must be the latter case, whereby $G(t) = G_s$ on $t \in [t', t'']$. It follows from (33) and (36) that the costate remains zero throughout $[t', t'']$, which contradicts the assumption that the arc is $\mathcal{B}_+$. $\qquad\square$

We formulate the final argument in the proof of Proposition 5 as a separate result forming a condition for the existence of singular arcs.

Proposition 7. *If $A \leq 2a$, there are no singular arcs.*

In the absence of singular arcs, the optimal control is of the form $(\mathcal{B}_-)\mathcal{B}_+$. In this case of $A \leq 2a$, the singular state is negative, and thus unreachable. Intuitively, the $\mathcal{B}_+$ arc will steer the state to a smaller value than a $\mathcal{B}_-$ arc, which, in this case, is advantageous for both the running and terminal pay-offs. Therefore, we expect the optimal control to be a $\mathcal{B}_+$ arc, which we prove in the following result.

Proposition 8. *If $A \leq 2a$, the optimal control is a $\mathcal{B}_+$ arc.*

Proof. If $\lambda(t) = 0$ at any t, it follows from the costate equation (33) that $\dot{\lambda}(t) = (A - 2a) - 2(A - a)G(t) \leq -2(A - a)G(t) < 0$ when $A \leq 2a$ and thus $\lambda(t)$ must remain strictly negative for all time after t. Since $\lambda(T) = 1$, we deduce that $\lambda(t)$ never vanishes and must remain positive for all t. The claim then follows from (35). $\qquad\square$

Having determined the optimal control when $A \leq 2a$, we now turn our attention to the case when $A > 2a$, where there may be singular arcs. We begin by highlighting some useful inequalities that may be verified directly from the definitions in (31), (36) and (37)

$$
\begin{align}
A > 2F \quad &\Longleftrightarrow \quad u_s > 1 \quad &\Longleftrightarrow \quad G_1(1) > G_s, \tag{39a} \\
A > 2a \quad &\Longleftrightarrow \quad u_s < u_0 \quad &\Longleftrightarrow \quad G_s > 0. \tag{39b}
\end{align}
$$

This reveals a useful dichotomy in parameter space and we split the analysis based on the case of weak induction $A \geq 2F$, in which case the singular control u_s is no less than the maximal control, and strong induction $A < 2F$, where the singular control is strictly less than the maximal control.

3.2.1 The case of weak induction $A \geq 2F$

In addition to Proposition 7, another necessary condition for the existence of singular arcs is that the singular control (37) respects the control bounds: $u_s \leq 1$. If $u_s > 1$, the grower mass $G(t)$ will never reach the level (36) that characterises a singular arc.

Proposition 9. *If $A \geq 2F$, there are no singular arcs.*

Proof. If $A \leq 2a$, there are no singular arcs (Proposition 7), therefore, we focus on the case of $A > 2a$. In the case where $A > 2F$, we see from (39) that $G_1(1) > G_s > 0$. Therefore, from Proposition 3, the state cannot remain constant at G_s, and the claim follows.

When $A = 2F$, we look to the optimal control structure. Since every singular arc must be followed by a $\mathcal{B}_+$ arc, it suffices to prove that the arc combination $\mathcal{S}\mathcal{B}_+$ cannot form. The argument is identical to that in the proof of Proposition 6: the singular control is the maximal control $u_s = 1$ when $A = 2F$, therefore, the state remains constant on a $\mathcal{S}\mathcal{B}_+$ arc combination $G(t) = G_s$, and thus so too does the costate (33), contradicting the fact that the second arc is $\mathcal{B}_+$. $\qquad\square$

We deduce, in the case of weak induction, by combining the discussion before Proposition 6, with Propositions 7 to 9, that the optimal control is of the form $(\mathcal{B}_-)\mathcal{B}_+$. We now proceed to narrow down the control structure more concretely. Intuitively, for weak induction, the maximal control does not suffice to maintain the singular state (see (39)), but will drive the state as close as possible to maximising both the running and terminal pay-offs. Therefore, we expect the optimal to be a $\mathcal{B}_+$ arc, as long as the initial state g_0 is not too much smaller than the singular state G_s, as we establish below.

Proposition 10. *If $A \geq 2F$ and $g_0 \geq G_s$, the optimal control is a $\mathcal{B}_+$ arc.*

Proof. Assume, by way of contradiction, that the optimal control is of the form $\mathcal{B}_-\mathcal{B}_+$, where the $\mathcal{B}_-$ arc is on $t \in [0, t_0]$ for some $t_0 > 0$. On this interval $\dot{G}(t) = (A - a)G(t)(1 - G(t)) \geq 0$ with equality only for $G(t) = 1$. Therefore, $G(t) > G_s$ for all $t > 0$. By assumption, $\lambda(t_0) = 0 \leq \dot{\lambda}(t_0)$, but $\dot{\lambda}(t_0) = 2(A - a)(G_s - G(t_0)) < 0$, which is a contradiction. $\qquad\square$

Only when $A > 2a$, $A \geq 2F$ and $g_0 < G_s$, can there exist a $\mathcal{B}_-$ arc. When such an arc exists, we may characterise the state and costate as functions of the switching time $t_0 > 0$ when the costate changes sign and the control switches. Equation (29b) is separable with constant coefficients when the control is constant. For constant control $u(t) = u$ and initial condition g_0 the solution is found to be

$$G(t; u, g_0) = \Sigma(u) \left[1 + \left(\frac{\Sigma(u)}{g_0} - 1 \right) e^{-(A-a-Fu)t} \right]^{-1}, \qquad \Sigma(u) = \frac{A - a - Fu}{A - a}. \qquad (40)$$

It follows that the state $G(t)$ is given by

$$G(t) = \begin{cases} G(t; 0, g_0), & t \in [0, t_0), \\ G(t - t_0; 1, g_1(t_0)), & t \in [t_0, T], \end{cases} \qquad (41)$$

where

$$g_1(t_0) := G(t_0; 0, g_0) = \left[1 + \left(\frac{1}{g_0} - 1 \right) e^{-(A-a)t_0} \right]^{-1}. \qquad (42)$$

This calculation allows us to bound the switching time by observing that $g_1(t_0) < G_s$, otherwise it follows from Proposition 10 that the control is not optimal. Requiring $g_1(t_0) \leq G_s$ yields the bound

$$t_0 < \frac{1}{A - a} \log \left[\left(1 - \frac{2a}{A} \right) \left(\frac{1}{g_0} - 1 \right) \right]. \qquad (43)$$

The bound (43) is not tight, and we seek a more exact characterisation of t_0 by resolving the optimality conditions. The costate may be obtained on $[t_0, T]$, where $u(t) = 1$, by using the integrating factor

$$\Lambda(t) := \exp \left(\int^t A - a - F - 2(A - a)G(t - t_0; 1, g_1(t_0)) \, dt \right)$$

$$= e^{(A-a-F)(t-t_0)} \left(e^{(A-a-F)(t-t_0)} + \frac{\Sigma(1)}{g_1(t_0)} - 1 \right)^{-2}, \qquad (44)$$

up to a positive multiplicative constant, which may be neglected without loss of generality. We thus find that, on $t \in [t_0, T]$, the costate is given by

$$\lambda(t) = \frac{\Lambda(T)}{\Lambda(t)} + \int_T^t [A - 2a - 2(A - a)G(s)] \frac{\Lambda(s)}{\Lambda(t)} \, ds. \qquad (45a)$$

Using the shorthands $c := A - a - F$, $d(t_0) := -1 + \Sigma(1)/g_1(t_0)$, and $z(t; t_0) := e^{c(t-t_0)}$, we may rewrite the costate explicitly via

$$\Lambda(t) = \frac{z(t; t_0)}{(z(t; t_0) + d(t_0))^2},$$

$$\int^t \Lambda(t)\, \mathrm{d}t = -\frac{1}{c}\frac{1}{z(t; t_0) + d(t_0)}, \tag{45b}$$

$$-\int^t 2(A - a)G(t)\Lambda(t)\, \mathrm{d}t = \frac{z(t; t_0)}{(z(t; t_0) + d(t_0))^2} + \frac{1}{z(t; t_0) + d(t_0)}.$$

Note that, since $A > 2a$, $A \geq 2F$, and $g_1(t_0) \leq G_s$, it follows that $c \neq 0$, $d(t_0) \geq 0$ and thus $z(t; t_0) + d(t_0) \neq 0$ for all $t > t_0$. Substituting (45b) into (45a), we find that the costate takes the form

$$\lambda(t) = 1 + \frac{z(t; t_0) + d(t_0)}{z(t; t_0)}\frac{a - F}{A - a - F}\left(1 - \frac{z(t; t_0) + d(t_0)}{z(T; t_0) + d(t_0)}\right). \tag{45c}$$

Therefore, the costate vanishes at the switching point, $\lambda(t_0) = 0$, if and only if

$$R(t_0) := \frac{e^{(A-a-F)(T-t_0)} - 1}{A - a - F}\left[(F - a) - (A - a)g_1(t_0)\right] - 1 = 0. \tag{46}$$

The switching time t_0 must satisfy equation (46). This equation does not appear to admit a closed-form solution, in general. Consider, for instance, any $F > a = 0$, and $A = F + \epsilon$, for a vanishingly small $\epsilon \ll 1$. Upon the change of variables $x = F(T - t_0) - 1$, the leading-order contributions give rise to the equation

$$x e^x = y, \qquad \text{where} \qquad y = \frac{e^{FT-1}}{-1 + 1/g_0} \in (0, \infty), \tag{47}$$

which is transcendental [7].

Nevertheless, we may still say a great deal about solutions. The state $g_1(t_0)$ is a positive and strictly increasing function of t_0, and so the residual $R(t_0)$ is a strictly decreasing function in t_0 that is free from singularities. The residual satisfies $R(T) = -1 < 0$ and thus there is an admissible solution $t_0 \in (0, T)$ if and only if $R(0) > 0$. A necessary condition for $R(0) > 0$ to hold is that the term in square brackets in (46) is positive, namely,

$$g_0 < \frac{F - a}{A - a}. \tag{48}$$

In particular, since $g_0 \in (0, 1]$ it must be that $F > a$. Physically, the condition (48) requires the initial condition to be sufficiently small (note that $(F - a)/(A - a) \leq G_s$) to make the $\mathcal{B}_-\mathcal{B}_+$ strategy worthwhile. When the condition (48) is satisfied, the time horizon T must be sufficiently large so that $R(0)$ can be positive. Ultimately, we arrive at the following result.

Proposition 11. *If $A \geq 2F$, the optimal control is a $\mathcal{B}_+$ arc unless $R(0) > 0$, in which case there is exactly one admissible solution of equation (46) characterising the switching time of the $\mathcal{B}_-\mathcal{B}_+$ control.*

We may also consider approximate solutions. Bearing in mind that we have found an upper bound (43) on t_0 and a lower bound on T, it is reasonable to consider the case of moderately large $T - t_0$ (assuming $g_0 = \mathcal{O}(1)$). In this case we see from (46) that the exponentially large term (multiplying the square brackets) must be balanced by the term in square brackets, which has to be exponentially small. Therefore, up to exponentially small corrections, we expect the solution to be given by t_0 that makes the term in square brackets vanish, namely,

$$t_0 = \frac{1}{A - a}\log\left[\frac{F - a}{A - F}\left(\frac{1}{g_0} - 1\right)\right] \qquad \text{when} \qquad F > a, \tag{49}$$

while $t_0 = 0$ when $F \leq a$, that is, the optimal control is a $\mathcal{B}_+$ arc.

To verify the analytical solution, we implement a multiple-shooting direct numerical approach with a backward-Euler discretisation encoded in CasADi [3], a framework that performs automatic differentiation to produce a nonlinear program equipped with gradient information and solved by the IPOPT solver [30]. The code is included in the Supplementary Materials.
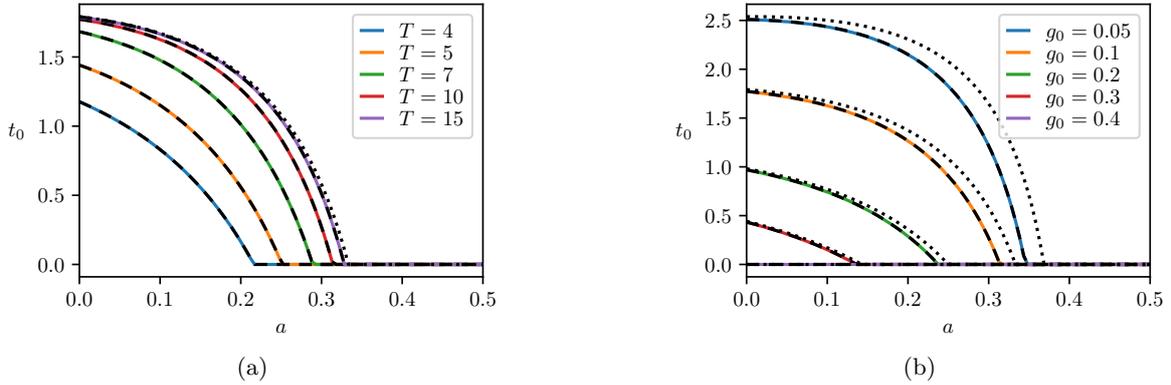
Figure 2: $A \geq 2F$. Switching times t_0 of the optimal control for various parameter values found using optimisation of the full ODE problem (coloured curves), numerical solution of equation (46) $R(t_0) = 0$ (dashed curves), and the approximate solution (49) (dotted curves). Parameters, when not specified, are $A = 1$, $F = 0.4$, $g_0 = 0.1$, and $T = 10$.

In fig. 2 we illustrate the switching times t_0 for the optimal control with weak induction $A \geq 2F$. We see perfect agreement between the full ODE problem (coloured curves) and the numerical solution of equation (46) $R(t_0) = 0$ (dashed curves). For $T \gtrsim 10$, the switching times are in close agreement with the large-T approximation (49) (dotted curved) to which they converge (fig. 2a). As the initial condition g_0 is decreased (fewer growers initially) the optimal switching time increases (fig. 2b), as more time is required to establish an optimal balance between the two population masses. In both cases, we plot results for parameter values of a up to $A/2 = 0.5$, and observe how $t_0 = 0$ for $a \geq F$ (and even before), as there is no solution to equation (46).

When $t_0 > 0$, the switching time is monotonically decreasing with a up until a critical value of a beyond which the switching time remains $t_0 = 0$ as the optimal control is just a $\mathcal{B}_+$ arc. We can find an implicit expression for the critical value of a at which this bifurcation occurs by noting that $R(0) = 0$, therefore,

$$\frac{e^{(A-a-F)T} - 1}{A - a - F} \left[(F - a) - (A - a)g_0 \right] = 1. \tag{50}$$

Equation (50) is transcendental with respect to a but may be solved explicitly for g_0.

We have found *necessary* conditions for the existence of a $\mathcal{B}_-\mathcal{B}_+$ extremal: when these conditions are not met the optimal control must be a $\mathcal{B}_+$ arc. When these conditions are met we have an extremal control but we have not ruled out the possibility that the $\mathcal{B}_+$ control is also extremal (though, in practice, we did not encounter a case where the $\mathcal{B}_+$ arc was optimal when a $\mathcal{B}_-\mathcal{B}_+$ extremal exists). Explicit formulae for the objective function are available for these solutions (though not reproduced here) and thus this too reduces to checking explicit formulae. Alternatively, sufficiency conditions or a uniqueness argument might rule out one of the extremals, although we do not pursue this here. Thus, when $A \geq 2F$, the entire solution is given in closed form, modulo one numerical root-finding procedure relevant in a subset of cases.

3.2.2 The case of strong induction $A < 2F$

In the case of strong induction, $A < 2F$, the control is capable of steering the state to G_s for any initial condition (given enough time and assuming that $A > 2a$ so that $G_s > 0$). Therefore, as opposed to the case of weak induction, a singular arc may arise.

Recall that the optimal control is of the form $(\mathcal{B}_-)(\mathcal{B}_+)(\mathcal{S})\mathcal{B}_+$. We proceed to refine this by proving that also the combination $\mathcal{B}_-\mathcal{B}_+\mathcal{S}$ cannot appear. Enumerating all remaining possible optimal control forms, we then find that the optimal control must be of one of the forms $\mathcal{B}_-\mathcal{S}\mathcal{B}_+$, $\mathcal{B}_+\mathcal{S}\mathcal{B}_+$, $\mathcal{S}\mathcal{B}_+$, $\mathcal{B}_-\mathcal{B}_+$, or $\mathcal{B}_+$. As we will show, all of these formations are obtainable: the first three pertain to situations where there is sufficient time to optimally establish a singular arc (that is, while also optimally accounting for the terminal pay-off), and differ based on the initial conditions, $g_0 < G_s$, $g_0 > G_s$, and $g_0 = G_s$, respectively. The final two formations pertain to situations where there is insufficient time to optimally establish a singular arc, and again distinguished based on the initial condition, $g_0 < G_s$ and $g_0 \geq G_s$, respectively.

Proposition 12. *If $2a < A < 2F$, the optimal control cannot contain the arc combination $\mathcal{B}_-\mathcal{B}_+\mathcal{S}$.*

Proof. Assume by contradiction that such an arc form exists, where the $\mathcal{B}_+$ arc spans the interval (t_0, t_1) for $0 < t_0 < t_1 < T$. It follows that $\lambda(t) > 0$ on $t \in (t_0, t_1)$ and $\lambda(t_0) = \lambda(t_1) = 0$. It must also hold that $\dot{\lambda}(t_0) \geq 0 \geq \dot{\lambda}(t_1)$, and thus it follows from the costate equation (33) that $\dot{\lambda}(t_i) = 2(A - a)(G_s - G(t_i))$ for both $i = 0, 1$. Therefore, $G(t_0) \leq G(t_1)$.

From the state equation (29b) we see that, on a $\mathcal{B}_+$ arc, the state satisfies

$$\dot{G}(t) = (A - a)G(t)\left(1 - \frac{F}{A - a} - G(t)\right) = (A - a)G(t)(G_1(1) - G(t)). \tag{51}$$

Therefore, $G(t)$ is monotonically increasing on the $\mathcal{B}_+$ arc if $G(t_0) < G_1(1)$, monotonically decreasing on the $\mathcal{B}_+$ arc if $G(t_0) > G_1(1)$, and constant on the $\mathcal{B}_+$ arc if $G(t_0) = G_1(1)$. Since the arc following $\mathcal{B}_+$ is the singular arc, we know that $G(t_1) = G_s$. It follows from $2a < A < 2F$ that $G_1(1) < G_s$, and, therefore, the state cannot be increasing or constant on the $\mathcal{B}_+$ arc, since these trajectories will not reach G_s. We deduce that $G(t)$ must be monotonically decreasing along the $\mathcal{B}_+$ arc, in contradiction to earlier. □

We now determine the critical time T_{crit} required to optimally establish a singular arc (and give the optimal control the turnpike structure discussed at the end of section 3.1). For $T \leq T_{\text{crit}}$, no singular arc exists, and we are in a similar situation to the case of weak induction where the optimal control is of the form $(\mathcal{B}_-)\mathcal{B}_+$. For $T > T_{\text{crit}}$, one singular arc is established. The duration T_{crit} comprises the time it takes for $G(t)$ to evolve from $G(0) = g_0$ to $G(t) = G_s$, which we denote t_s, and the time it takes to evolve from $\lambda(t) = 0$ with $G(t) = G_s$ to $\lambda(t) = 1$ with $u(t) = 1$, which we denote T_s.

If $g_0 = G_s$ then $t_s = 0$. If $g_0 < G_s$ then the initial arc will be a $\mathcal{B}_-$ arc and to calculate t_s we take the constant control $u(t) = 0$. Using the solution (41), we obtain t_s by solving $G(t_s; 0, g_0) = G_s$ to find that

$$t_s = \frac{1}{A - a}\log\left[\left(1 - \frac{2a}{A}\right)\left(\frac{1}{g_0} - 1\right)\right] \qquad \text{when} \qquad g_0 \leq G_s. \tag{52a}$$

This is none other than the bound derived in (43). Similarly, for $g_0 > G_s$, solving $G(t_s; 1, g_0) = G_s$ yields

$$t_s = \frac{1}{A - a - F}\log\left\{\left[1 - \frac{1}{g_0}\left(1 - \frac{F}{A - a}\right)\right]\frac{A - 2a}{2F - A}\right\} \qquad \text{when} \qquad \begin{array}{l} g_0 > G_s \text{ and} \\ A - a - F \neq 0. \end{array} \tag{52b}$$

When $c = A - a - F = 0$, the solution (52b) is to be understood in the limit of $c \to 0$, whereby,

$$t_s = \frac{2}{F - a} - \frac{1}{g_0 F} \qquad \text{when} \qquad g_0 > G_s \qquad \text{and} \qquad A - a - F = 0. \tag{52c}$$

The terminal nonsingular arc, on $[T - T_s, T]$, of duration T_s, is characterised by the costate evolving from $\lambda(T - T_s) = 0$ to $\lambda(T) = 1$. Since the problem is homogeneous in time, it will be convenient to translate the time interval to $[0, T_s]$, and consider the state to solve (29b) subject to the initial condition $G(0) = G_s$ and the costate to solve (33) with the conditions $\lambda(0) = 0$ and $\lambda(T_s) = 1$, where $u(t) = 1$ throughout the interval. This solution may be expressed via (45) with $t_0 = 0$, $g_0 = G_s$, and $T = T_s$. Upon enforcing the initial condition $\lambda(0) = 0$, we obtain an equation for the terminal duration T_s, which may be expressed as $R(0)|_{g_0 = G_s, T = T_s} = 0$:

$$e^{(A - a - F)T_s} - \frac{A - 2a}{2F - A} = 0. \tag{53}$$

When $2a < A < 2F$, there exists a unique solution $T_s > 0$ to equation (53), namely,

$$T_s = \frac{1}{A - a - F}\log\left(\frac{A - 2a}{2F - A}\right), \qquad \text{when} \qquad A - a - F \neq 0. \tag{54a}$$

When $c = A - a - F = 0$, the duration T_s may be found by considering the expression (54a) in the limit $c \to 0$, to obtain

$$T_s = \frac{2}{F - a}, \qquad \text{when} \qquad A - a - F = 0. \tag{54b}$$
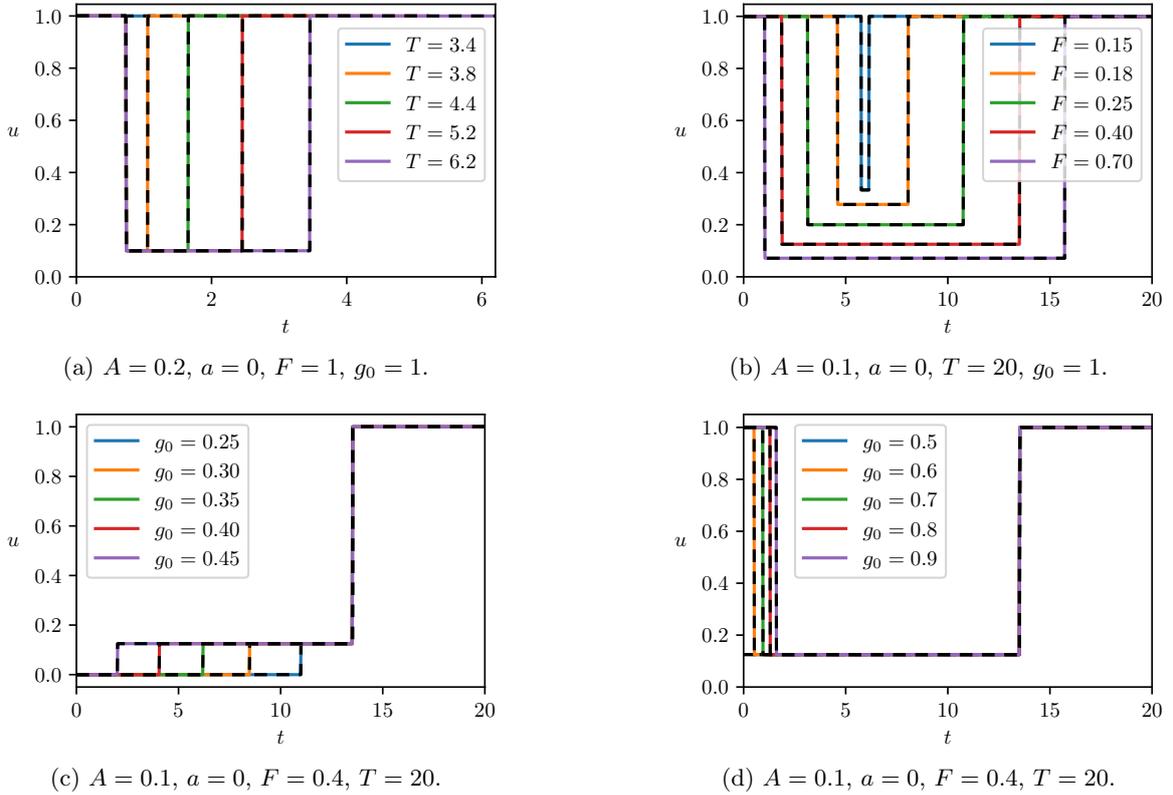
We consolidate the calculations in the following result.

16

(a) $A = 0.2$, $a = 0$, $F = 1$, $g_0 = 1$.

(b) $A = 0.1$, $a = 0$, $T = 20$, $g_0 = 1$.

(c) $A = 0.1$, $a = 0$, $F = 0.4$, $T = 20$.

(d) $A = 0.1$, $a = 0$, $F = 0.4$, $T = 20$.

Figure 3: **$A < 2F$**. Optimal controls for various parameters A, a, F, and time horizons T. Coloured lines show numerical solutions (for 5000 discrete control points), black dotted lines show analytical solutions.

Proposition 13. *If $2a < A < 2F$ and $T \leq T_{crit} := t_s + T_s$, there is no singular arc.*

In fig. 3 we show optimal controls for the case of strong induction ($A < 2F$) with singular arcs of the forms $\mathcal{B}_+\mathcal{S}\mathcal{B}_+$, $\mathcal{B}_-\mathcal{S}\mathcal{B}_+$, and $\mathcal{S}\mathcal{B}_+$. The agreement between the numerical and analytical results remains perfect. We see how the singular arc emerges as T increases from subcritical $T < T_{crit}$ to supercritical $T > T_{crit}$ (fig. 3a where $T_{crit} \approx 3.5$). As the induction becomes stronger with increasing F, the singular arc increases in duration (fig. 3b): the induction is able to faster establish the state $G(t) = G_s$, optimal for the running pay-off, and requires less time of maximum induction for the terminal pay-off near the terminal time $t = T$. In figs. 3c and 3d, we explore the impact of initial conditions $G(0) = g_0$ where all other parameters are kept fixed. When $a = 0$, the optimal running pay-off state is an even split between growers and producers: $G_s = 1/2$. For initial conditions below the singular level, $g_0 < G_s$, the optimal control arc first has a $\mathcal{B}_-$ arc with minimal induction (fig. 3c). This is sustained for longer as the initial grower mass g_0 decreases (reminiscent of fig. 2b). When $g_0 = G_s$ the initial arc is singular $\mathcal{S}$, while for $g_0 > G_s$, the initial arc is maximal induction $\mathcal{B}_+$ (fig. 3d). For all of the controls in figs. 3c and 3d the terminal phase is identical: once the singular arc is established, the discrepancy due to different initial conditions is eliminated.

When $T \leq T_{crit}$, no singular arc is present and the control is of the form $(\mathcal{B}_-)\mathcal{B}_+$. We may again leverage the solutions (41) and (45) for a constant control $u(t) = u$ and apply the same approach of seeking an admissible solution to the equation $R(t_0) = 0$ for the residual $R(t_0)$ defined in (46). The only difference is that, in this case, the value $c = A - a - F$ is not necessarily positive, but can also be negative or zero. When $c < 0$ it suffices to note that $R(t_0)$ remains a monotonically decreasing function of t_0 with no singularities, and, therefore, the previous condition is preserved: an admissible solution $t_0 \in (0, T)$ exists if and only if $R(0) > 0$.

Some care is required in the case of $c = 0$, which is removably singular (for example, see the solution (40) for $u = 1$) and must be tackled in the limit as $c \to 0$ as $R(t_0)$ is not strictly defined at $c = 0$. We find that a costate solution vanishing at t_0 exists if and only if

$$S(t_0) := \lim_{c \to 0} R(t_0) = (T - t_0)\left[(F - a) - Fh_1(t_0)\right] - 1 = 0, \tag{55a}$$

17

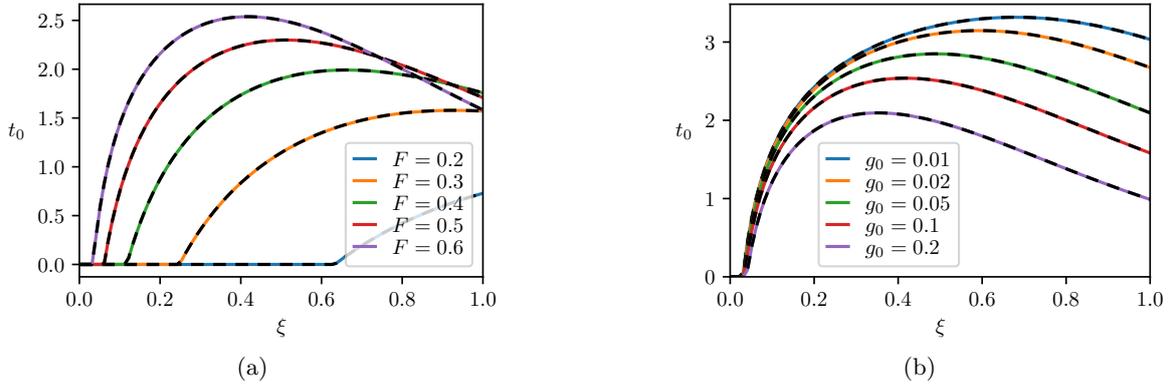(a)   (b)

Figure 4: $A < 2F$ and $T \leq T_{\mathrm{crit}}$. Switching times t_0 of the optimal control for various parameter values found using optimisation of the full ODE problem (coloured curves) and numerical solution of equation (46) $R(t_0) = 0$ (dashed curves). The horizontal axis shows the parameter ξ, by which we set $A = 2a + \xi(2F - 2a)$ to interpolate between $2a$ and $2F$. Parameters, when not specified, were $a = 0$, $F = 0.4$, $g_0 = 0.1$, and $T = 5$.
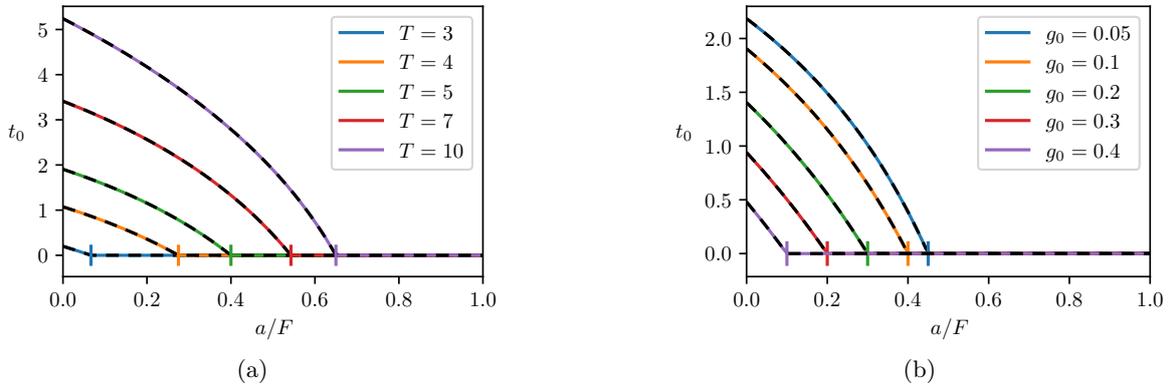

(a)   (b)

Figure 5: $A - a - F = 0$ and $T \leq T_{\mathrm{crit}}$. Switching times t_0 of the optimal control for various parameter values found using optimisation of the full ODE problem (coloured curves) and numerical solution of equation (55) $S(t_0) = 0$ (dashed curves). The horizontal axis shows the parameter $a \in (0, F)$. Parameters, when not specified, were $A = a + F$, $F = 0.4$, $g_0 = 0.1$, and $T = 5$. Vertical lines show the bifurcation point a given in (56) beyond which the optimal control is simply a $\mathcal{B}_+$ arc, that is, $t_0 = 0$.

where

$$h_1(t_0) := \lim_{c \to 0} g_1(t_0) = \left[1 + \left(\frac{1}{g_0} - 1\right) \mathrm{e}^{-Ft_0}\right]^{-1}. \tag{55b}$$

The function $S(t_0)$ is monotonically decreasing in t_0 and $S(T) = -1 < 0$, therefore, an admissible solution $t_0 \in (0, T)$ exists if and only if $S(0) > 0$, which may be written as $a < F(1 - g_0) - 1/T$. The condition (48) is necessary here too, and thus similarly $F > a$ must hold and T must be sufficiently large in order for an admissible solution to exist. We recap the case with no singular arc in the following result.

Proposition 14. *If $2a < A < 2F$ and no singular arc is formed, there is an extremal $\mathcal{B}_-\mathcal{B}_+$ control if and only if: $R(0) > 0$ when $A - a - F \neq 0$, $S(0) > 0$ when $A - a - F = 0$.*

In fig. 4 we illustrate switching times t_0 for the subcritical case of strong induction ($A < 2F$ and $T \leq T_{\mathrm{crit}}$). The analytical and numerical results remain indistinguishable. In this case, the absence of sufficient time to establish the singular arc requires a switching time t_0 that balances between an initial arc approaching (but not reaching) the optimal running pay-off state G_s and a final arc increasing the terminal yield. This trade-off gives rise to a fascinating non-monotonicity exhibited by the switching time t_0 as the growth rate A varies. For increasing induction strength F the switching time t_0 increases (for fixed A), reflecting the fact that a longer time can be spent approaching G_s with minimal induction before inducing to increase the terminal yield (fig. 5a). As in fig. 2b, a smaller initial grower population justifies a longer arc of minimal induction (fig. 5b).
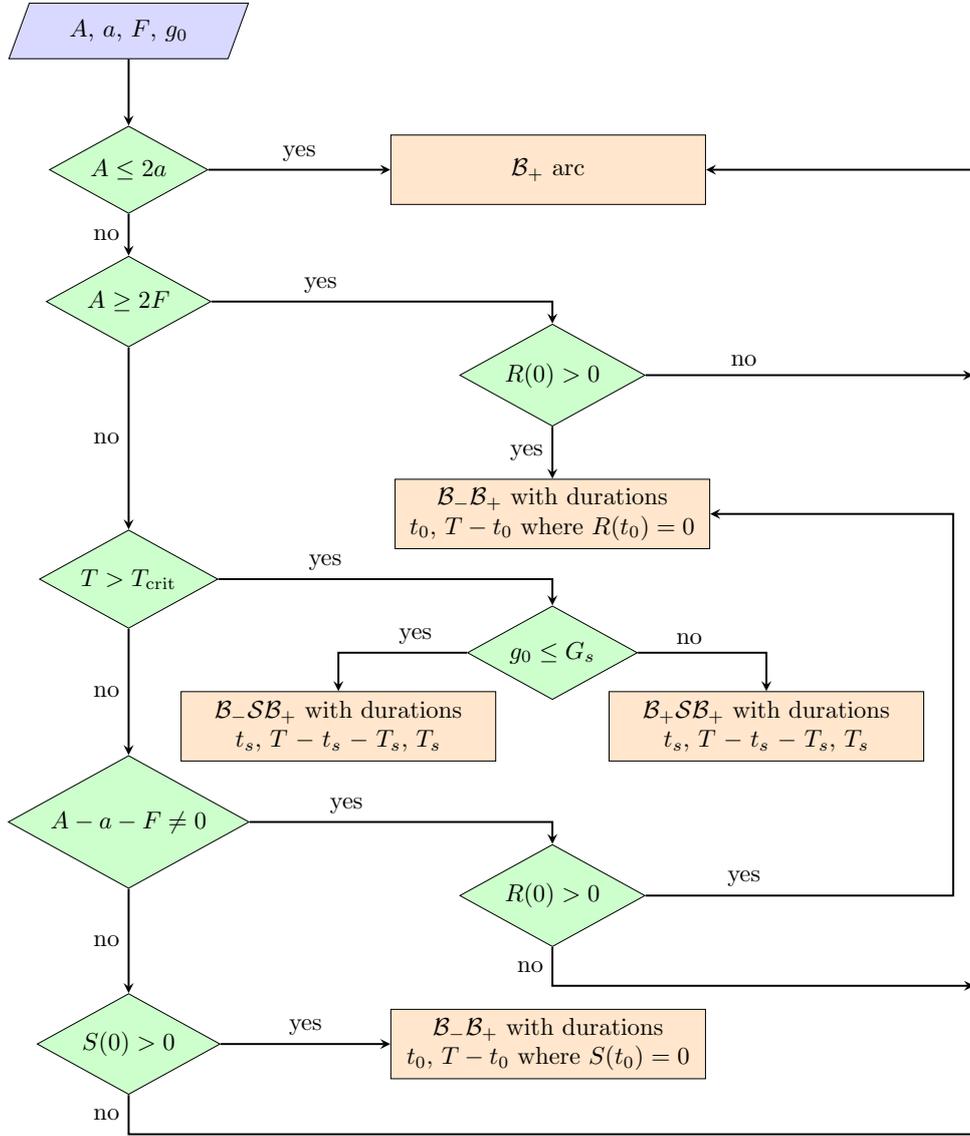
18

Figure 6: Flowchart summarising the optimal control synthesis for the ODE problem (29). Quantities G_s, $R(t_0)$, $S(t_0)$, t_s, and T_s are defined in (36), (46), (55), (52), (54), respectively, and are all functions of the input parameters A, a, F, and g_0, with $T_{\text{crit}} = t_s + T_s$.

In fig. 5 we plot the results for the removably singular subcritical case ($A - a - F = 0$ and $T < T_{\text{crit}}$). The analytical and numerical results are identical. Just as for the case of weak induction, we could consider the case of asymptotically large $T - t_0$ and $g_0 = \mathcal{O}(1)$. However, subcriticality imposes an upper bound on T, while the value t_s in (52) serves as an upper bound on t_0, therefore the assumption is severely limited. Nevertheless, we may determine bifurcation points, where the solution passes through $t_0 = 0$, satisfying (50). While the equation remains transcendental (except with respect to g_0), in the limiting case of $c \to 0$ it takes the simpler form $S(0) = 0$, which we write as

$$F(1 - g_0) - a = \frac{1}{T}. \tag{56}$$

Equation (56) admits explicit solutions with respect to a, F, g_0, and T. Solutions for a are illustrated as coloured vertical lines in fig. 5.

This concludes the analysis. Since we partitioned the parameter space into several regions and subregions, it is worthwhile to take stock of our results. Combining all the cases serves the dual purpose of summarising the results and concretely demonstrating that the entire parameter space is covered, as illustrated by the flowchart in fig. 6.

It is worth interpreting the analytical insight of the optimal control structure from the biological perspective. When growth arrest is weak $A \leq 2a$, that is, the producers' diminished growth rate a is

19

at least half the growers' growth rate A, the optimal control is full light, regardless of the optogenetic induction strength F or experimental time horizon T. This is a striking result: even if induction is strong, so that one can turn growers into producers rapidly, if the growth arrest is weak it is optimal to have as large a proportion of producers as possible irrespective of all else. When induction is weak $A \geq 2F$, the optimal solution is full light (unless the initial condition and time horizon are such that it is beneficial to have an initial dark transient $R(0) > 0$). When induction is strong $A < 2F$ and the time horizon suffices $T > T_{\text{crit}}$, the optimal control is of the turnpike form: first reach the singular arc where the running pay-off is optimised, and then finish with a transient of full light to account for the terminal pay-off. When the time horizon is insufficient, as in the case of weak induction, the optimal solution is full light (unless the initial condition and time horizon are such that it is beneficial to have an initial dark transient).

4   Conclusions

We study the modelling and optimal control of microbial consortia, in particular a two-species consortium generated by light-inducable differentiation. These two seemingly distinct aspects are merely two sides of the same coin.

By posing a fully stochastic model comprising individual cells, with stochastic chemical kinetics governing their internal state, alongside their stochastic division and removal, we begin by encapsulating a biologically faithful process. Successive reductions are required to obtain increasingly tractable surrogate models. The first reduction produces a model governing the expected population density. This reduction is exact, having made no approximations. Moreover, the model retains heterogeneity among the population as well as multiscale dynamics, driven by both the single-cell and population effects, while preserving the complexity of a single-cell (stochastic) model. Nevertheless, in capturing the heterogeneity, this model retains a large state space that makes analytical progress challenging. A further reduction, by taking low-order moments over the state space and appropriate assumptions, recovers the ODE models ubiquitous in the literature. The derivation highlights the precise conceptual nature of, and connections between, these three classes of models.

We then study the optimal control problem of maximising protein yield based on the reduced-order ODE model. This model finds important application in bioproduction, and its analysis allows for the fast computation of open-loop optimal controls that may be periodically updated via MPC to close the loop [21]. In applying Pontryagin's maximum principle, we partition the parameter space into several regions. One collection of regions admit explicit optimal controls, while in the remainder of parameter space the optimal control is of known structure where the control switching time is given as the solution of a transcendental equation. We prove that, in these latter cases, the solution of the equation exists on a known compact interval and is unique, and thus amenable to standard numerical root-finding methods. We verify our analytical results via direct numerical solution of the optimal control problem, finding excellent agreement between the two approaches.

It is our aim that this hierarchy of models of different fidelity constitute a family that may prove useful in practical bioproduction processes. The derivation technique is generalisable to a broad range of multiscale processes comprising stochastically evolving populations of stochastic agents. Our analysis demonstrates that the low-fidelity models provide extensive analytical insight into optimally controlling such processes. Optimising the coarse model provides a tractable path to parameter and control choice in the numerical simulation of higher-fidelity models. This reveals how the two distinct contributions of this work are in fact woven together in a mutual interplay. These models could also be fused in a hybrid manner to simultaneously exploit the strengths of each, providing a cohesive framework to model and control bioproduction processes.

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 (2021), pp. 1–10, https://doi.org/10.1038/s41467-021-26129-7.

[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 (2022), p. e2114438119, https://doi.org/10.1073/pnas.2114438119.

[3] 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 (2019), pp. 1–36, https://doi.org/10.5281/zenodo.1257968.

[4] P. Bittihn, A. Didovyk, L. S. Tsimring, and J. Hasty, *Genetically engineered control of phenotypic structure in microbial colonies*, Nat. Microbiol., 5 (2020), pp. 697–705, https://doi.org/10.1038/s41564-020-0686-0.

[5] K. Brenner, L. You, and F. H. Arnold, *Engineering microbial consortia: a new frontier in synthetic biology*, Trends Biotechnol., 26 (2008), pp. 483–489, https://doi.org/10.1016/j.tibtech.2008.05.004.

[6] D. A. Charlebois, K. Hauser, S. Marshall, and G. Balázsi, *Multiscale effects of heating and cooling on genes and gene networks*, P. Natl. Acad. Sci. USA, 115 (2018), pp. E10797–E10806, https://doi.org/10.1073/pnas.1810858115.

[7] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, *On the LambertW function*, Adv. Comput. Math., 5 (1996), pp. 329–359, https://doi.org/10.1007/BF02124750.

[8] J. B. Deris, M. Kim, Z. Zhang, H. Okano, R. Hermsen, A. Groisman, and T. Hwa, *The innate growth bistability and fitness landscapes of antibiotic-resistant bacteria*, Science, 342 (2013), p. 1237435, https://doi.org/10.1126/science.1237435.

[9] L. Duso and C. Zechner, *Stochastic reaction networks in dynamic compartment populations*, Proc. Natl. Acad. Sci. U.S.A., (2020), https://doi.org/10.1073/pnas.2003734117.

[10] D. T. Gillespie, *A general method for numerically simulating the stochastic time evolution of coupled chemical reactions*, J. Comput. Phys., 22 (1976), pp. 403–434, https://doi.org/10.1016/0021-9991(76)90041-3.

[11] D. T. Gillespie, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81 (1977), pp. 2340–2361, https://doi.org/10.1021/j100540a008.

[12] D. T. Gillespie, *A rigorous derivation of the chemical master equation*, Physica A, 188 (1992), pp. 404–425, https://doi.org/10.1016/0378-4371(92)90283-V.

[13] N. E. Grandel, K. Reyes Gamas, and M. R. Bennett, *Control of synthetic microbial consortia in time, space, and composition*, Trends Microbiol., 29 (2021), pp. 1095–1105, https://doi.org/10.1016/j.tim.2021.04.001.

[14] S. Klumpp, Z. Zhang, and T. Hwa, *Growth rate-dependent global effects on gene expression in bacteria*, Cell, 139 (2009), pp. 1366–1375, https://doi.org/10.1016/j.cell.2009.12.001.

[15] R. M. Lewis, *Definitions of order and junction conditions in singular optimal control problems*, SIAM J. Control Optim., 18 (1980), pp. 21–32, https://doi.org/10.1137/0318002.

[16] Z. Li, X. Wang, and H. Zhang, *Balancing the non-linear rosmarinic acid biosynthetic pathway by modular co-culture engineering*, Metab. Eng., 54 (2019), pp. 1–11, https://doi.org/10.1016/j.ymben.2019.03.002.

[17] D. Liberzon, *Calculus of variations and optimal control theory*, Princeton university press, 2011.

[18] D. Lunz, *On continuum approximations of discrete-state markov processes of large system size*, Multiscale Model. Sim., 19 (2021), pp. 294–319, https://doi.org/10.1137/20M1332293.

[19] 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 (2021), pp. 1–24, https://doi.org/10.1371/journal.pcbi.1009214.

[20] D. Lunz, J. F. Bonnans, and J. Ruess, *Revisiting moment-closure methods with heterogeneous multiscale population models*, Mathematical Biosciences, (Accepted).

[21] D. Lunz, J. F. Bonnans, and J. Ruess, *Optimal control of bioproduction in the presence of population heterogeneity*, (https://hal.inria.fr/hal-03445175).

[22] A. Miano, M. J. Liao, and J. Hasty, *Inducible cell-to-cell signaling for tunable dynamics in microbial communities*, Nat. Commun., 11 (2020), pp. 1–8, https://doi.org/10.1038/s41467-020-15056-8.

[23] D. Nevozhay, R. M. Adams, E. Van Itallie, M. R. Bennett, and G. Balázsi, *Mapping the environmental fitness landscape of a synthetic gene circuit*, PLOS Comput. Biol., 8 (2012), pp. 1–17, https://doi.org/10.1371/journal.pcbi.1002480.

[24] 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 (2016), pp. 4063–4068, https://doi.org/10.1073/pnas.1507291113.

[25] K. M. Rapp, J. P. Jenkins, and M. J. Betenbaugh, *Partners for life: building microbial consortia for the future*, Curr. Opin. Biotech., 66 (2020), pp. 292–300, https://doi.org/10.1016/j.copbio.2020.10.001.

[26] D. Salzano, D. Fiore, and M. di Bernardo, *Controlling reversible cell differentiation for labor division in microbial consortia*, bioRxiv, (2021), https://doi.org/10.1101/2021.08.03.454926.

[27] C. Tan, P. Marguet, and L. You, *Emergent bistability by a growth-modulating positive feedback circuit*, Nat. Chem. Biol., 5 (2009), pp. 842–848, https://doi.org/10.1038/nchembio.218.

[28] P. Thomas, *Intrinsic and extrinsic noise of gene expression in lineage trees*, Sci. Rep., 9 (2019), pp. 1–16, https://doi.org/10.1038/s41598-018-35927-x.

[29] E. Trélat and E. Zuazua, *The turnpike property in finite-dimensional nonlinear optimal control*, J. Differ. Equations, 258 (2015), pp. 81–114, https://doi.org/10.1016/j.jde.2014.09.005.

[30] 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 (2006), pp. 25–57, https://doi.org/10.1007/s10107-004-0559-y.

[31] 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), IEEE, 2019, pp. 2663–2668, https://doi.org/10.23919/ECC.2019.8795858.

[32] J. J. Winkle, B. R. Karamched, M. R. Bennett, W. Ott, and K. Josić, *Emergent spatiotemporal population dynamics with cell-length control of synthetic microbial consortia*, PLOS Comput. Biol., 17 (2021), pp. 1–23, https://doi.org/10.1371/journal.pcbi.1009381.

[33] 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 (2015), pp. 377–383, https://doi.org/10.1038/nbt.3095.