



HAL
open science

Stochastic chemical kinetics of cell fate decision systems: from single cells to populations and back

Jakob Ruess, Guillaume Ballif, Chetan Aditya

► **To cite this version:**

Jakob Ruess, Guillaume Ballif, Chetan Aditya. Stochastic chemical kinetics of cell fate decision systems: from single cells to populations and back. *Journal of Chemical Physics*, 2023, 159 (18), 10.1063/5.0160529 . hal-04275681

HAL Id: hal-04275681

<https://inria.hal.science/hal-04275681>

Submitted on 8 Nov 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License

Stochastic chemical kinetics of cell fate decision systems: from single cells to populations and back.

Jakob Ruess,^{1,2} Guillaume Ballif,¹ and Chetan Aditya^{2,3}

¹*Inria Saclay, 91120 Palaiseau, France.*

²*Institut Pasteur, Université Paris Cité, 75015 Paris, France.*

³*Department of Molecular Biology, Princeton University, Princeton NJ 08544, USA.*

(*Electronic mail: jakob.ruess@inria.fr)

(Dated: 12 October 2023)

Stochastic chemical kinetics is a widely used formalism for studying stochasticity of chemical reactions inside single cells. Experimental studies of reaction networks are generally performed with cells that are part of a growing population, yet the population context is rarely taken into account when models are developed. Models that neglect the population context lose their validity whenever the studied system influences traits of cells that can be selected in the population, a property that naturally arises in the complex interplay between single-cell and population dynamics of cell fate decision systems. Here, we represent such systems as absorbing continuous-time Markov chains. We show that conditioning on non-absorption allows one to derive a modified master equation that tracks the time evolution of the expected population composition within a growing population. This allows us to derive consistent population dynamics models from a specification of the single-cell process. We use this approach to classify cell fate decision systems into two types that lead to different characteristic phases in emerging population dynamics. Subsequently, we deploy the gained insights to experimentally study a recurrent problem in biology: how to link plasmid copy number fluctuations and plasmid loss events inside single cells to growth of cell populations in dynamically changing environments.

INTRODUCTION

Over the last decades, numerous studies have demonstrated that intracellular biochemical processes are inherently stochastic¹⁻³. As this realization became common knowledge, much effort has been focused on measuring such processes at the single-cell level, on developing stochastic reaction network models, governed by the chemical master equation, to describe their dynamics^{4,5}, and on studying the resulting cell-to-cell variability and its consequences⁶. For instance, many studies have focused on elucidating experimentally and theoretically how cell-to-cell variability can be exploited as a bet-hedging strategy for unknown environments⁷, how it can give rise to non-genetic resistance of cancer cells to treatments⁸ or of bacteria to antibiotics⁹, or how it leads to stochastic cell fate decisions in developmental and various other systems¹⁰. The underlying logic of such work is typically that the causal relation is such that emerging phenotypes in cell populations are a function of the biochemical process and not of selection processes that might be operating on those phenotypes¹¹.

In some sense the opposite perspective is usually taken in studies that focus on evolution^{12,13}. Here, an emerging phenotype (distribution) is typically a consequence of population dynamics and selection operating on the fitness of phenotypes¹⁴⁻¹⁶, whereas the stochastic biochemical processes that actually define or change phenotypes are only very rarely represented explicitly¹⁷. The main reason for these different viewpoints is that studies of intracellular biochemical processes have often focused on systems that are, or at least traditionally were, thought to not have an impact on fitness. While this is probably true in a number of cases, there also exist many examples where we need to question this assumption in light of today's knowledge¹⁸. If, for instance, synthetically constructed genetic circuits draw strongly upon cellular re-

sources or cause toxicity and affect the cells' growth rates^{19,20}, we might need to expect that cells with low gene expression levels will end up being overrepresented in the population²¹. Clearly, such an effect would manifest if the system is constructed with the purpose of modulating the cells' growth rates, as has already been demonstrated experimentally for a histidine production circuit in *Escherichia coli* growing in histidine depleted environments²². Apart from this, there exist many cases where selection does not act through differing growth rates but through selective removal of cells with particular phenotypes from the population, for instance in response to anti-cancer treatments²³ or antibiotics²⁴⁻²⁶, or due to selective recombination²⁷. Again, it is natural to anticipate that, upon treatment, there will be selective removal of cells with particular phenotypes and that the size of the remaining population, as well as its phenotype distribution, will depend both on the biochemical process that is generating phenotypes and the selection process that is removing some of them. More generally, any cell fate decision system in which the fate of a cell depends on the realization of a stochastic biochemical process will lead to selective removal of cells with particular cell-internal states from the sub-population of cells that have not triggered the cell fate decision yet and thereby change the sub-population phenotype distribution^{28,29}.

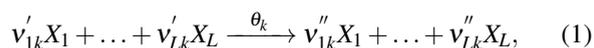
Here, representing cell fate decision systems as absorbing continuous-time Markov chains³⁰, we show mathematically how consistent (deterministic) population dynamics models can be derived from a specification of the (stochastic) single-cell process, how statistics of cells that have not triggered the cell fate decision yet can be tracked by a non-linear version of the system's master equation, and how the infinitesimal generator matrix of the single-cell model can be analyzed to provide insights into emergent behavior at the population scale. Taken together, these results allow us to study coupled single-cell and population scale dynamics of cell fate decision systems

without the need for any costly individual-based simulations.

We make use of these theoretical results to study population dynamics in response to environmental stimuli that trigger cell fate decisions and show that different characteristic phases in population dynamics emerge depending on whether stochasticity in cell fate decisions is due to pre-existing population variability or a consequence of stochasticity in chemical reactions that are triggered by the stimulus. We then study a biological problem that has continued to resurface already for decades: how to model fluctuations of plasmid copy numbers inside single cells, plasmid loss events, and emerging dynamics of microbial populations carrying plasmids. We equip yeast cells with 2-micron plasmids carrying a selection marker and show experimentally that characteristic population dynamics of a cell fate decision system with pre-existing variability emerge when cells are switched back and forth between selective and non-selective media. Deploying the theoretical results presented in this manuscript, we show that emerging population dynamics are quantitatively predictable from a single-cell model of plasmid copy number fluctuations.

BASICS OF STOCHASTIC CHEMICAL KINETICS

Stochastic biochemical processes inside cells are typically represented as reaction networks of a certain number of chemical species, X_1, \dots, X_L , that interact according to reactions



where the coefficients v'_{ik} and v''_{ik} determine how many molecules of the i -th species are consumed and produced in the k -th reaction, respectively, and θ_k is the reaction rate constant of the reaction. Taking into account the intrinsic stochasticity of biochemical reactions in small volumes, the dynamics of the reaction network can be described by a stochastic process $X(t)$ that takes states $n = [n_1 \dots n_L]^T \in \mathbb{N}^L$ and that counts the number of molecules of the L chemical species. The dynamics of $X(t)$ are determined by the reaction propensities

$$a_k(n) = \theta_k \prod_{i=1}^L \binom{n_i}{v'_{ik}}, \quad k = 1, \dots, K$$

and the stoichiometric transition vectors $v_k = [v''_{1k} \dots v''_{Lk}]^T - [v'_{1k} \dots v'_{Lk}]^T$, $k = 1, \dots, K$ that determine how molecule numbers in the system change upon the occurrence of reactions. If the reaction system is well-stirred, the probability for a reaction of type k to occur in an infinitesimally small time interval $[t, t + dt]$ given that $X(t) = n$ is $a_k(n) \cdot dt$ and $X(t)$ is a continuous time Markov chain (CTMC) on \mathbb{N}^L governed by a chemical master equation (CME)

$$\frac{d}{dt}p(n, t) = -p(n, t) \sum_{k=1}^K a_k(n) + \sum_{k=1}^K p(n - v_k, t) a_k(n - v_k), \quad (2)$$

where $p(n, t)$ is the probability for $X(t)$ to be in state n at time t (see³¹ for further details on the validity of the chemical master equation).

Since the number of states that $X(t)$ can attain grows exponentially in the number of reacting chemical species, it is often desirable to exploit specific properties of the reaction network such as conservation laws of chemical species or time-scale separation between different reactions to reduce the dimensionality of the process, which may lead to deviations from the standard formulation of chemical reactions in Eq.(1) and the chemical master equation in Eq.(2). A specific example for this that we will also use later is the well-know one-dimensional model of bursty gene expression that can be obtained from a two-dimensional standard reaction network by exploiting time-scale separation between fluctuations in messenger RNA and protein numbers³².

Irrespectively of whether or not the model is simplified, the number of states that $X(t)$ can, at least in principle, attain is infinite in most cases since there is no fixed boundary for molecule numbers that cannot be exceeded. However, reaction networks that are practically relevant are typically such that the probability for molecule numbers to increase beyond some boundary is very small since the cell volume is finite. In line with this, a popular approach to calculate with chemical master equations is to truncate the state space^{33,34}, i.e. to approximate the infinite state CTMC $X(t)$ by a finite state CTMC so as to obtain a finite number of equations of the form Eq.(2) that need to be solved to calculate the time evolution of the probability distribution of the chain. We will follow this approach here and, with a slight abuse of notation, refer to the approximate finite state CTMC with the same notation $X(t)$ as used for the infinite state CTMC in the mindset that the difference between the two processes is negligible as long as the truncation of the state space is chosen sufficiently large.

With a finite state space, we may choose some ordering of all states $n^{(1)}, \dots, n^{(m)}$, where m is the cardinality of the truncated state space and each $n^{(i)} \in \mathbb{N}^L$, $i = 1, \dots, m$ corresponds to some possible combination of molecule counts of the L reacting species. Given such an ordering, one can arrange the probabilities of all states in a vector $\tilde{\mathbf{p}}(t) := [p(n^{(1)}, t) \dots p(n^{(m)}, t)]^T$, write the chemical master equation in vector form as

$$\frac{d}{dt}\tilde{\mathbf{p}}(t) = \tilde{\mathbf{A}}\tilde{\mathbf{p}}(t), \quad (3)$$

and solve it, for instance, via calculation of the matrix exponential of $\tilde{\mathbf{A}}$.

CELL FATE DECISION SYSTEMS AS ABSORBING MARKOV CHAINS

Stochastic cell fate decision systems are processes where intracellular reactions eventually lead to events, such as differentiation or cell death, that irreversibly change the phenotype of the cell. Cell fate decisions can be incorporated into stochastic reaction network models by augmenting the CTMC with an absorbing state, denoted $n^{(d)}$ in the following. Transitions to this state then correspond to the cellular event that causes cell fate decisions and occur with a rate $f(n)$ that depends on the current state, n , of the system, i.e. the probability

for a cell fate decision to occur in an infinitesimally small time interval $[t, t + dt]$ is $f(n) \cdot dt$ given that $X(t) = n$ at time t . In all following parts of this paper, we will only consider the CTMC with the added absorbing state and to not further complicate the notation refer to this process with the same symbol $X(t)$ that was previously used for CTMCs.

Adding the probability of the absorbing state to the vector of all probabilities $\mathbf{p}(t) := [p(n^{(d)}, t) \ p(n^{(1)}, t) \ \cdots \ p(n^{(m)}, t)]^T$, we can write a new vector form master equation

$$\frac{d}{dt} \mathbf{p}(t) = \mathbf{A} \mathbf{p}(t) \quad (4)$$

in which \mathbf{A} has the specific structure

$$\mathbf{A} = \begin{bmatrix} 0 & \mathbf{c}_1 \\ \mathbf{0} & \mathbf{C} \end{bmatrix}$$

where $\mathbf{C} \in \mathbb{R}^{m \times m}$ and $\mathbf{c}_1 = [f(n^{(1)}) \ \cdots \ f(n^{(m)})]$. \mathbf{c}_1 comprises the rates of transition to the absorbing state for any possible system configuration $n^{(1)}, \dots, n^{(m)}$, which constitute inflow terms of probability to the absorbing state that must also appear as outflow terms on the diagonal of \mathbf{C} such that the column sums of \mathbf{A} are all zero and the total probability in Eq.(4) is preserved. Thus, it holds that $\mathbf{c}_1 = -\mathbf{1}^T \mathbf{C}$ is the negative of the column sums of \mathbf{C} .

It is sometimes useful to calculate conditional probabilities $p_c(n, t) := \mathbb{P}(X(t) = n \mid X(t) \neq n^{(d)})$ for the absorbing Markov chain to be in state n at time t given that it has not reached the absorbing state yet. We can derive an equation for the time evolution of these conditional probabilities from Eq.(4) (see Supporting Information, Section S.3 and²⁷) and obtain:

$$\frac{d}{dt} \mathbf{p}_c(t) = \mathbf{C} \mathbf{p}_c(t) + \mathbf{p}_c(t) \cdot (\mathbf{c}_1 \mathbf{p}_c(t)), \quad (5)$$

where $\mathbf{p}_c(t) := [p_c(n^{(1)}, t) \ p_c(n^{(2)}, t) \ \cdots \ p_c(n^{(m)}, t)]^T$.

In recent years, it has become increasingly popular to study probabilities or timing of cellular events, for instance by calculating the waiting time distribution for a certain protein in the cell to reach a particular threshold level given some initial amount of protein^{35–37}.

Based on Eq.(4), we can define the random time to the cellular event as

$$\tau := \min_{t \geq 0} \left\{ t : X(t) = n^{(d)} \right\}$$

and calculate its probability density function as

$$f_\tau(t) = \mathbf{p}_0^\top \exp(\mathbf{C}^\top t) \mathbf{c}_1^\top, \quad (6)$$

where \mathbf{p}_0 is the initial condition of the system, that is \mathbf{p}_0 specifies the probabilities to have any specific number of molecules present in a cell at time zero. Eq.(6) can be used to study important questions on the functioning of cell fate decision systems. For instance, one can investigate how the variance of the time to the cellular event, $\text{Var}(\tau)$, depends on the properties of the single-cell process that defines \mathbf{C} and \mathbf{c}_1 . However,

such questions are implicitly based on a single-cell perspective. Trying to lift the analysis to a population context, one encounters numerous questions that cannot be answered in any obvious way. How should we choose the initial condition \mathbf{p}_0 in Eq.(6)? Can we just interpret \mathbf{p}_0 as a population distribution and if we do so, then what population event is described by $f_\tau(t)$? What is the waiting time in remaining cells when a part of the population has already triggered the cell fate decision? What if cells are also growing and the pool of cells is continuously replenished? What is even an appropriate definition of a “time zero” from which on to calculate event timing in a context where growth and cell fate decisions are operating at the same time scale? And most importantly, how can we derive a model that tracks population dynamics of the different sub-populations from the single-cell perspective? To lay the basis for answering these questions, we first need to establish how classical stochastic chemical kinetics can be lifted to the context of a growing cell population.

STOCHASTIC REACTION NETWORKS IN GROWING CELL POPULATIONS

Originating from classical chemistry, the theory of chemical kinetics is based on assumptions that are adapted to chemical reactions in test tubes. The reaction system is assumed to be well-stirred while the reaction volume is assumed to remain constant. To be able to use a classical chemical master equation (CME) for reaction networks in cells, these assumptions need to be ported over to cells despite the fact that cells are growing and dividing. Since this clearly leads to a crude approximation of the real biological process, first efforts were made recently to develop models of chemical reactions inside cells coupled to explicit models of the cell cycle^{38–42}. Thomas and Shahrezaei pursued such an approach and introduced stochastic agent-based models that track all cells, their sizes, as well as the state of the cell-internal reaction network⁴³. To calculate with such models, they introduced a stochastic simulation algorithm based on the assumption that growth rates of cells are constant and do not depend on the state of the reaction network and that there is no death or removal of cells. Subsequently, they introduced a corresponding master equation, which they approximated using a linear noise approximation. The main problem of such models is that their analysis and simulation is extremely difficult, which often renders them impractical to use.

As a consequence, the CME remains a viable alternative for modeling stochastic reaction networks in growing cells despite the simplifying assumptions on which it is based. However, the CME is by design an equation for the time evolution of the probability distribution of the state of a single reaction network in a single cell. Contrary to that, in practice, the CME is often used to model the distribution of the state of the reaction network in a growing population and to match population snapshot data such as obtained, for instance, by flow cytometry. In the following, we will clarify under which assumptions the classical CME can be justified from the perspective of a full population model and show how the CME can be gener-

alized to incorporate population level dynamics such as cell growth or death/removal rates that depend on the state of the reaction network. Cell fate decision systems, which are the focus of this manuscript, then emerge as a special case of a growing population with state dependent cell removal rates.

To start, we need to define a population model in which the standard assumptions of classical chemical kinetics are lifted to the population scale. To this end, we assume that cell volume remains constant throughout the cell cycle. At cell division, a new cell is added to the population that inherits the internal state of the reaction network from its mother cell. The waiting time for division events of single cells is assumed to follow an exponential distribution $Exp(\lambda)$. Since this exponential distribution defines how quickly cells divide, its parameter λ needs to be the single-cell growth rate that may (or may not) depend on the state of the reaction network, that is $\lambda = \lambda(n)$, where $n \in \mathbb{N}^L$ and L is the number of species of the reaction network as in previous sections. We highlight that the assumption of exponentially distributed cell cycle times is essential here for tractability of the model: the state n of the reaction network may change in a cell during its cell cycle but this is not a problem because the exponential distribution is memoryless and allows one to simply recalculate waiting times for cell division events when internal states of cells change. Equivalently to cell divisions events, we assume that waiting times for cell death/removal events are exponentially distributed with a parameter $f(n)$ that may depend on the state of the reaction network.

With these assumptions, we can define the population composition as a continuous-time Markov chain $Y(t)$ following the update rules summarized in Table I and taking states $Y(t) = y$ that can be written as

$$y = \sum_{n \in \mathbb{N}^L} y(n) \delta_n$$

where $y(n)$ is the number of cells with internal state n in the population and δ_n is the Dirac mass at n (see Supporting Information, Section S.4 for a rigorous definition of the state space of $Y(t)$ and the δ_n functions).

At this point, it is useful to compare the population model to a standard single-cell model. Instead of tracking the state of a single reaction network, the population model tracks the number of cells with given internal states in the population. This number may change either due to reactions occurring inside single cells or due to population level events such as cell birth or death. The explicit incorporation of population level events in Table I is particularly useful: it defines what exactly is assumed at the population scale and allows us to modify these assumptions if necessary. For instance, changing the state y by a Dirac δ_n upon a birth event of a cell with internal state n corresponds to adding a new daughter cell to the population that inherits the exact state from its mother cell. To incorporate other assumptions on mother to daughter inheritance, the rules for state changes due to growth events could readily be modified. For instance, a state change $y \rightarrow y - \delta_n + \delta_{n'} + \delta_{n''}$ would correspond to a cell with internal state n splitting into two new cells with internal states n' and n'' .

Since the stochastic process $Y(t)$ that describes the dynam-

	Change of a cell in state n	Change of Y_t	Rate
Reaction R_k ($1 \leq k \leq K$)	$n \rightarrow n + v_k$	$y \rightarrow y + \delta_{n+v_k} - \delta_n$	$y(n)a_k(n)$
Cell birth event		$y \rightarrow y + \delta_n$	$y(n)\lambda(n)$
Cell death event	$n \rightarrow \emptyset$	$y \rightarrow y - \delta_n$	$y(n)f(n)$

TABLE I. Dynamics of the population composition. Rows: possible events modifying the state of the population composition. Columns: corresponding state changes of a single cell (first column), the population composition (second column) and rates of these events inside the population (third column). The rates are obtained as the as the single-cell rates of the event for the corresponding state n of the reaction network multiplied by the number of cells $y(n)$ that have internal state n for given state y of the population composition. In other words, each cell in the population is a Markov chain that has a rate of updating its state or firing a birth or death reaction. The rate of an event occurring inside the population is then the sum of the rates of all individual cell Markov chains, i.e. an event inside the population occurs when the first single cell fires a reaction.

ics of the population composition is a Markov process, it is conceptually very similar to a standard model of a single-cell reaction network and we can write down a master equation that tracks the time evolution of its probability distribution:

$$\begin{aligned} \frac{d}{dt} p(y, t) = & \sum_{k=1}^K \sum_{n \in \mathbb{N}^L} \left[\left(y(n - v_k) + 1 \right) a_k(n - v_k) p(y + \delta_{n - v_k} - \delta_n, t) \right. \\ & \left. - y(n) a_k(n) p(y, t) \right] \\ & + \sum_{n \in \mathbb{N}^L} \left[\left(y(n) - 1 \right) \lambda(n) p(y - \delta_n, t) - y(n) \lambda(n) p(y, t) \right] \\ & + \sum_{n \in \mathbb{N}^L} \left[\left(y(n) + 1 \right) f(n) p(y + \delta_n, t) - y(n) f(n) p(y, t) \right], \end{aligned} \quad (7)$$

where $p(y, t) := \mathbb{P}(Y(t) = y | Y(0) = y_0)$.

Eq.(7) captures the dynamics of the full probability distribution of the population composition in a finite size population. It is convenient to have such an equation at our disposal but trying to numerically solve Eq.(7) is rather hopeless since the state space of $Y(t)$ is much too complicated. What we can do, however, is to use Eq.(7) as the basis for deriving an equation for the expectation of $Y(t)$, that is the expected population composition. To this end, define

$$r(n, t) := \sum_y y(n) p(y, t),$$

where the sum goes over all possible states y that the population composition can attain and, as before, $y(n)$ is the number of cells with internal state n when $Y(t) = y$. Computing the time derivative of $r(n, t)$ from Eq.(7) then leads to (see Sup-

porting Information, Section S.4)

$$\begin{aligned} \frac{d}{dt}r(n,t) = & -r(n,t) \sum_{k=1}^K a_k(n) + \sum_{k=1}^K r(n-v_k,t) a_k(n-v_k) \\ & + r(n,t) (\lambda(n) - f(n)). \end{aligned} \quad (8)$$

We can notice that the sums over all chemical reactions on the right hand side are exactly the same as in the standard CME (Eq.(2)). The terms in the second row of the equation, however, stem from the addition and removal of cells at the population scale and are not present in the single-cell CME. Furthermore, these terms break the preservation of total probability in the classical CME, which is a consequence of the definition of $r(n,t)$ as the expected number of cells with internal state n that may grow or decay with the population size. The expected population size $z(t) = \sum_{n \in \mathbb{N}^L} r(n,t)$ follows

$$\frac{d}{dt}z(t) = \sum_{n \in \mathbb{N}^L} r(n,t) (\lambda(n) - f(n)), \quad (9)$$

which reduces to a simple exponential growth or decay model $\frac{d}{dt}z(t) = (\lambda - f_0)z(t)$ for the case when cell growth and death rates $\lambda(n) = \lambda$ and $f(n) = f_0$ do not depend on the state n of the reaction network.

To more directly compare Eq.(8) to a standard CME, one can derive a probability preserving version by defining a normalized expected population composition $\rho(n,t) := \frac{r(n,t)}{z(t)}$. Deriving the time derivative of $\rho(n,t)$, one obtains

$$\begin{aligned} \frac{d}{dt}\rho(n,t) = & -\rho(n,t) \sum_{k=1}^K a_k(n) + \sum_{k=1}^K \rho(n-v_k,t) a_k(n-v_k) \\ & + \rho(n,t) (\lambda(n) - f(n)) \\ & - \rho(n,t) \sum_{\tilde{n} \in \mathbb{N}^L} (\lambda(\tilde{n}) - f(\tilde{n})) \rho(\tilde{n},t). \end{aligned} \quad (10)$$

It is straightforward to see that if either $\lambda(n)$ or $f(n)$ do not depend on the state of the reaction network n (or \tilde{n} in the sum), the corresponding terms in the third row on the right hand side cancel with those in the second row. If both $\lambda(n)$ and $f(n)$ are independent of n (or \tilde{n}), the entire third row cancels with the second row. The first row is equivalent to the standard CME (Eq.(2)). It follows that the standard CME provides a model for the expected population composition under the assumption that neither growth nor death rates of cells depend on the state of the reaction network.

Since the focus of this paper is on cell fate decision systems, we will from now on assume that $\lambda(n) = \lambda$ is constant but maintain the generalization from the standard setting that allows for cell death/removal rates $f(n)$ to depend on the state of the reaction network.

LINKING EXPECTED POPULATION COMPOSITION AND ABSORBING MARKOV CHAINS.

Standard chemical master equations have been studied for a long time but little is known about the properties of the

equations for the expected population composition, Eq.(8) and Eq.(10), that we derived in the previous section. In this section, we will show that when the growth rate of cells does not depend on the state of the reaction network, $\lambda(n) = \lambda$, the vector form of Eq.(10) is equivalent to Eq.(5), that is it can be interpreted as the governing equation for the probability distribution of an absorbing Markov chain conditional on the chain not having reached the absorbing state.

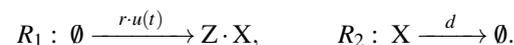
To this end, let us first return to the assumption/approximation that the state space of the model is finite and can be enumerated as $n^{(1)}, \dots, n^{(m)}$. To obtain a vector-based formulation as in the first sections of the paper, we can define the vectors $\mathbf{c}_1 := [f(n^{(1)}) \dots f(n^{(m)})]$ and $\rho(t) := [\rho(n^{(1)},t) \dots \rho(n^{(m)},t)]^T$, where $f(n^{(i)})$ and $\rho(n^{(i)},t)$, $i = 1, \dots, m$ refer to cell death/removal rate and normalized expected population composition for the corresponding state of reaction network, respectively. It follows from Eq.(10) that

$$\frac{d}{dt}\rho(t) = \tilde{\mathbf{A}}\rho(t) - \rho(t) \odot \mathbf{c}_1^T + \rho(t) \cdot (\mathbf{c}_1\rho(t)), \quad (11)$$

where \odot denotes component-wise multiplication. Here, $\tilde{\mathbf{A}}$ is the generator matrix of the CTMC describing the single-cell reaction network as introduced in Eq.(3). If we now absorb the subtraction $-\rho(t) \odot \mathbf{c}_1$ into the diagonal of $\tilde{\mathbf{A}}$ and define $\mathbf{C} := \tilde{\mathbf{A}} - \text{diag}(\mathbf{c}_1)$, we can see that Eq.(11) is equivalent to Eq.(5). It follows that under the assumption $\lambda(n) = \lambda$, Eq.(11) can be interpreted as an equation for conditional probabilities of absorbing CTMCs.

Intuitively, one can understand that if the transition to the absorbing state of an absorbing CTMC is interpreted as cell death/removal, then conditioning on the chain not having reached the absorbing state corresponds to taking only cells that have not yet been removed from the population into consideration, which is what a population model does by construction. The equivalence of Eq.(5) and Eq.(11) is useful since it allows us to deploy known theoretical results for absorbing Markov chains to study stochastic cell fate decision systems in growing populations. For the remainder of the paper, we will refer to conditional probabilities, Eq.(5), and the absorbing Markov chain interpretation, keeping in mind that Eq.(11) allows us to interpret the conditional probability distribution as an expected population composition.

To illustrate Eq.(5), we consider a simple example of inducible bursty production of a protein X, that is a reaction network consisting of a single species and the two chemical reactions



with propensity functions $a_1(n) = r \cdot u(t)$, $a_2(n) = d \cdot n$, and the slight generalization from the standard formulation in Eq.(1) that the model incorporates burstiness of protein production by allowing the number of protein molecules produced in the first reaction to be a random variable, Z, that follows a geometric distribution. This is a classical approach for representing stochasticity stemming from messenger RNA production and degradation in a model that tracks only protein copy numbers under the assumption that messenger RNA

fluctuations occur at a faster time scale³². Cell fate decisions are assumed to correspond to cell death events and to occur at a probability per unit time that is given by an increasing Hill function of the current cellular amount of protein $f(n) := a \frac{n^\eta}{\theta^\eta + n^\eta}$. One can then truncate the nominally infinite state space of the CTMC that counts the number of protein molecules by redirecting all burst events that would lead to molecule numbers exceeding a maximally allowed boundary to the boundary state, which leads to the **C**-matrix displayed on page 4 of the Supporting Information. Parameter values and simulation results for this model are provided in Figure 1.

Comparing the solution of Eq.(5) for this example to the solution of the standard master equation of the bursty protein production model shows that preferential death of cells with large protein numbers shifts the conditional probability distribution to lower levels. We then verified by stochastic simulation that the solution of Eq.(5) indeed agrees with population statistics calculated over all cells that have not died at any moment in time, contrary to the solution of the standard chemical master equation that does not lead to agreement with the simulation (Figure 1d). $\mathbf{p}_c(t)$ is thus the distribution that would need to be plugged into Eq.(6) to calculate waiting times on cell fate decisions in the population context but this only makes sense if $\mathbf{p}_c(t)$ is not time-varying and has converged to stationarity. We will study stationarity later and first focus more generally on how (deterministic) population dynamics models that are consistent with the single-cell process can be derived from Eq.(5).

POPULATION DYNAMICS OF STOCHASTIC CELL FATE DECISION SYSTEMS

We saw previously that Eq.(9) describes the expected population size in a general setting where both cell growth rates and death/removal rates may depend on the state of the reaction network. If we focus specifically on cell fate decision systems and assume that all cells grow at the same rate λ and note that $r(n,t) = z(t)\rho(n,t) = z(t)p_c(n,t)$, we obtain

$$\frac{d}{dt}z(t) = \lambda z(t) - z(t) \sum_{i=1}^m f(n^{(i)}) p_c(n^{(i)}, t)$$

as an equation for the expected population size. This implies that we can define a population death/removal rate as

$$k(t) := \mathbf{c}_1 \mathbf{p}_c(t) = \sum_{i=1}^m f(n^{(i)}) \cdot p_c(n^{(i)}, t), \quad (12)$$

and correspondingly a net population growth rate as $\lambda - k(t)$.

To test the accuracy of the resulting population model for a finite size population (i.e. when the actual population size may deviate from its expectation), we implemented an individual-based population simulation algorithm that tracks every chemical reaction in every cell of a growing population (see Supporting Information, Section S.2). Using this algorithm to simulate the cell fate decision system in Figure 1a in the context of a growing population, we find close agreement between

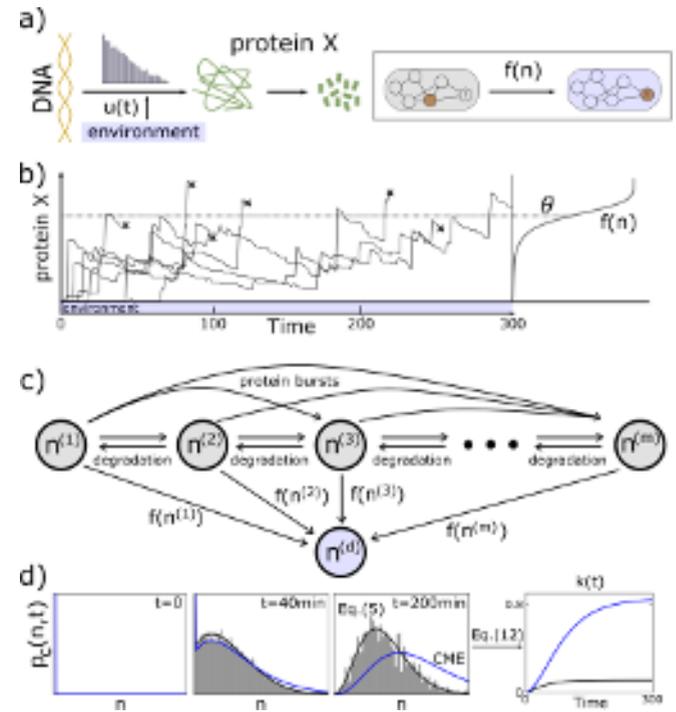


FIG. 1. Cell fate decision systems as absorbing Markov chains.

(a) Graphical visualization of our running example: proteins are degraded/diluted at a rate d and produced in geometrically distributed bursts of average size b with burst rate $r \cdot u(t)$ that depends on an environmental input $u(t)$. $u(t) = u$ is assumed to be constant in time for now. The probability of a cell fate decision taking place at any moment in time is determined by a Hill function $f(n) := a \frac{n^\eta}{\theta^\eta + n^\eta}$ of the number of protein molecules n . (b) Representative stochastic simulation of the system in panel a for a small number of cells. (c) Absorbing Markov chain representation of the system in panel a. For this 1-dimensional example, states of the CMTC $n^{(1)}, \dots, n^{(m)}$ correspond directly to molecule numbers of the single protein, i.e. to $0, \dots, m-1$ protein molecules being present in the system. Cell fate decisions take place when the chain moves to the absorbing state, $n^{(d)}$. (d) Eq.(5) can be solved to track the probability distribution, $p_c(n,t)$, of all cells that have not yet triggered a cell fate decision. $p_c(n,t)$ (black) agrees with statistics obtained by stochastic simulation (grey histograms, 10,000 cells initially) and is shifted towards lower levels compared to the solution of the CME of the reaction network (blue). The per cell rate of cell fate decisions, $k(t)$, (right panel, black) can be calculated from Eq.(12), increases initially as protein accumulates in cells, but eventually reaches a level that is lower than what would be obtained if the solution of the plain CME would have been used in Eq.(12) to calculate $k(t)$ (blue). For all panels, parameter values have been chosen as $d = 0.017/\text{min}$, $b = 10$, $r = 0.0765/\text{min}$, $u(t) = 1 \forall t$, $a = 1/\text{min}$, $\theta = 90$, $\eta = 5$ and we assumed that at time zero no protein is present in any cell.

the distribution of the simulated population and its expectation calculated from equation Eq.(5). As a final note, above we focused on cell fate decisions corresponding to death or removal of cells from the population that is tracked. If the size of the (sub-)population of cells that have triggered the cell fate decision system is also of interest, one can readily add an additional variable to the model to obtain a population dynamics

model of the form

$$\begin{aligned}\frac{d}{dt}z(t) &= \lambda z(t) - k(t)z(t), \\ \frac{d}{dt}z_c(t) &= \lambda_c z_c(t) + k(t)z(t),\end{aligned}\quad (13)$$

where $z(t)$ and $z_c(t)$ are the population sizes of cells that have not and have triggered the cell fate decision, respectively. If the two cell types are growing at different rates, λ and λ_c can be chosen differently and whether or not the two cell types can stably co-exist depends on the difference in growth rate and how exactly $k(t)$ emerges from the single-cell process according to Eq.(5) and Eq.(12).

We conclude this discussion by highlighting some important context and consequences. In practice, deterministic population dynamics models are often constructed without taking single-cell processes into account. Eq.(13) implies that such approaches are perfectly justifiable as models for the expected population size even in cases where the single-cell process is stochastic. However, if population models are assumed outright, the only reasonable way to set $k(t)$ is to determine all of its possible dynamics experimentally, which defies the purpose of making a model in the first place. In contrast to this, we have shown how $k(t)$ can be determined from first principles by deriving it from the stochastic chemical kinetics of intracellular reactions.

CONVERGENCE OF POPULATION STATISTICS AND FIRST-PASSAGE TIMES

We have seen in the previous section (Eq.(13)), that population dynamics models emerging from stochastic cell fate decision systems are in general not time-invariant. We may now ask ourselves when and why time-invariant models are justifiable from a single-cell perspective. Eq.(13) is time-invariant if $k(t)$ is constant in time. From Eq.(12), we can see that this is the case if $\mathbf{p}_c(t)$ does not change in time, that is if the normalized expected population composition has converged to some stationary distribution. Such convergence should be expected when the cell fate decision system is operating continuously and independently from any environmental input. For systems where cell fate decisions are triggered by an environmental input, we should expect that $\mathbf{p}_c(t)$ and $k(t)$ change in response to changes in the environment but converge whenever the environment stays constant for sufficiently long. For instance, for our simple example of inducible bursty production of a protein that triggers cell death (Figure 1a), application of the inducing input leads to transient changes in protein amounts and the population death rate. On the other hand, when the input remains constant for sufficiently long, the population protein distribution converges and a constant death rate of cells emerges at the population scale (Figure 1d). Notably, however, the emerging stationary population distribution is different from the stationary solution of the standard master equation that describes bursty protein production³² in the absence of cell death. Can we nevertheless predict this

stationary population distribution and corresponding population death rate from a specification of the single-cell process? Setting the time-derivative of $\mathbf{p}_c(t)$ in Eq.(5) to zero, we can recognize that the emerging stationary population distribution \mathbf{p}_c^s must be the (normalized right) eigenvector of the matrix \mathbf{C} corresponding to the largest eigenvalue γ_s (see³⁰ and Supporting Information, Section S.6). We can furthermore deduce from the relation $\mathbf{c}_1 = -\mathbf{1} \mathbf{C}$ that the population death rate, $k(t)$, approaches a constant $k_s = \mathbf{c}_1 \mathbf{p}_c^s = -\mathbf{1} \mathbf{C} \mathbf{p}_c^s = -\gamma_s$ that is equal to the negative of the largest eigenvalue of \mathbf{C} when the population distribution converges to stationarity. This implies that population rates of cell fate decisions in constant environments can be readily determined by a simple eigenvalue calculation on the generator matrix of the single-cell process.

To conclude, we return to the original single-cell analysis of event timing via first passage times. If we plug \mathbf{p}_c^s as initial condition into Eq.(6), the resulting distribution $f_\tau(t)$ describes the waiting time on cell fate decision of a randomly chosen cell in the population in stationary conditions. Importantly, \mathbf{p}_c^s is an eigenvector of \mathbf{C} and therefore of quite special nature for Eq.(6): it equilibrates the single-cell process “in all directions except the direction that leads towards cell fate decisions” (Supporting Information, Section S.6). As a consequence of these one-dimensional dynamics, $f_\tau(t)$ reduces to a simple exponential distribution with parameter $-\gamma_s$. This observation is what consolidates the single-cell and population perspective: exponential waiting times on events imply the Markov property, which ultimately is the reason why a time-invariant deterministic model with $-\gamma_s$ as rate of cell fate decisions is appropriate at the population scale in stationary conditions.

Finally, we point out that the convergence results provided in this section may also be useful to deduce population distributions at the initiation of experiments, i.e. $\mathbf{p}_c(0)$. In practice, cell populations will always be growing already before what the experimenter defines as “time zero” of the experiment. In cases where cell fate decisions are not (only) triggered by an environmental input that is applied at time zero but take place already before the start of the experiment, \mathbf{p}_c^s provides an equilibrated initial condition to which the population distribution will have converged in the growth period before time zero of the experiment. An example of this will be provided in our experimental study on plasmid copy number fluctuations provided at the end of this manuscript.

CLASSIFYING CELL FATE DECISION SYSTEMS

With the required theoretical background established, we can now demonstrate how our mathematical approach can be used to address concrete biological questions. To this end, we first investigate different types of cell fate decision systems that are commonly encountered in biology. As an illustrating example, we already discussed a system in which cell fate decisions are determined by some protein that is produced in response to an environmental input (Figures 1a and 2a, top). A key feature of this example is that the protein is absent in all cells before the input is applied, implying that all cells are

initially identical. The example can therefore be interpreted as the simplest motif of the class of systems in which stochasticity in cell fate decisions emerges exclusively as a consequence of intrinsic noise in chemical reactions that arises after the application of the environmental input. We will refer to such systems as type (i) cell fate decision systems in the following. In some sense the opposite type of cell fate decision system is obtained if protein production is always active in cells and does not depend on the environment, while instead the probability per unit time for cell fate decision events to take place depends directly on the environmental input in addition to the amount of protein in cells (Figure 2b, top). A typical example of this may be a protein that is activated by the environmental input instead of being produced in response to it. The key feature of such systems (referred to as type (ii) cell fate decision systems in the following) is that stochasticity in cellular processes will lead to cells having different amounts of relevant proteins already before the input that triggers cell fate decisions is applied. Differential responses of cells to the same environmental input will then largely be a consequence of pre-existing variability in the population.

To investigate the differences between the two types of systems, we considered a time-varying environmental input, $u = u(t)$, that is piecewise constant and switches between zero and one multiple times. We numerically solved Eqns.(5,12,13) for the two systems displayed in Figure 2 to calculate the full dynamics of $\mathbf{p}_c(t)$, $k(t)$, and $z(t)$ assuming, as before, that cell fate decisions correspond to cell death (i.e. without tracking $z_c(t)$ in Eq.(13)). Our results (Figure 2) show that the population dynamics of both the type (i) and the type (ii) system show three distinct phases: an adaptation phase after $u(t)$ switches from zero to one, a stationary phase in which convergence results apply, and a relaxation phase in which the population returns to its initial configuration. The characteristic features of these phases are, however, very different for the two systems. For the type (i) system, the adaptation phase is primarily a delay phase during which enough protein still needs to accumulate in cells for $k(t)$ to become large such that the effective population growth rate, $\lambda - k(t)$, starts to decrease only slowly. For the type (ii) system, $k(t)$ is largest directly after application of the environmental stimulus before it drops, together with the average amount of protein in cells, because cells with large protein content are now selectively removed. If the environmental input is maintained for sufficiently long, both systems eventually reach stationary conditions in which $\mathbf{p}_c(t)$ and $k(t)$ have converged and the population grows/decays exponentially at rate $\lambda + \gamma_s$ (which is negative in our examples). When $u(t)$ switches back to zero, $k(t)$ in the type (ii) system drops to zero immediately since $f(n, u(t)) = 0$ for all n . Therefore, the population resumes exponential growth at rate λ and the protein distribution relaxes back to its initial condition according to the dynamics of the standard CME because no more cells are selectively removed. For the type (i) system, $k(t)$ remains non-zero in the relaxation phase because the protein is not removed instantaneously. Therefore, the population returns only slowly to plain exponential growth at rate λ . We conclude by pointing out that both systems will show a phenotypic population memory

effect if the second environmental stimulus is applied before the system has fully returned to its initial condition during the relaxation phase: $k(t)$ will then reach its maximum faster for the type (i) system while it will spike less for the type (ii) system. The duration of this memory period is determined by the characteristic time scale(s) of the standard CME for the type (ii) system but needs to be calculated from Eq.(5) for the type (i) system because the CME does not govern the dynamics of the protein distribution during the relaxation phase in this case.

In practice, most cell-fate decisions will of course depend on more than the amount of a single protein. Real systems may then well contain elements of both type (i) and type (ii) systems. For instance, cell fate decisions may be caused by a protein that is stochastically produced in response to the environmental input but the rate of production of this protein may vary between cells due to pre-existing variability in the amount of ribosomes. For such mixed systems, it is a priori unclear what the main causes of variability in cell fate decisions are. Searching for characteristic features of type (i) and type (ii) systems in the dynamics of experimentally observable quantities may then help to understand dominant noise sources in the system.

In the following section, we will experimentally and theoretically study an example of a biological system that is not classically thought of as a cell fate decision system but that is, in fact, a perfect example of a pure type (ii) cell fate decision system according to our classification.

SINGLE-CELL AND POPULATION MODELS OF PLASMID LOSS

As an experimental example, we study a cell fate decision system that is of core importance for numerous natural and synthetically constructed biological processes: loss of plasmids that are needed for cellular growth or survival. Models of plasmid copy number fluctuations have been developed and applied already for some decades^{44,45}. At the population scale, models are typically deterministic and track the two sub-populations of cells that still contain at least one copy of the plasmid and cells that have lost all plasmids. Population dynamics are then determined by growth rates of cells with and without plasmids and a plasmid loss rate. Interpreting plasmid loss as a cell fate decision, and denoting the two sub-populations by $z(t)$ and $z_c(t)$, we can identify such models with Eq.(13) for constant $k(t) = k_s$. Such models assume a priori that the single-cell event of plasmid loss can be associated with a rate parameter at the population scale. While we have seen in the previous sections that this is valid whenever the plasmid copy number distribution of the population has converged to stationarity, it remains unclear how the plasmid loss rate at the population scale emerges from the single-cell process that governs plasmid copy number fluctuations.

At the single-cell scale, plasmid copy number fluctuations are sometimes modeled equivalently to standard biochemical reaction networks as CTMCs. Alternatively, discrete-time models that track cell generations can be used. The latter

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0160529

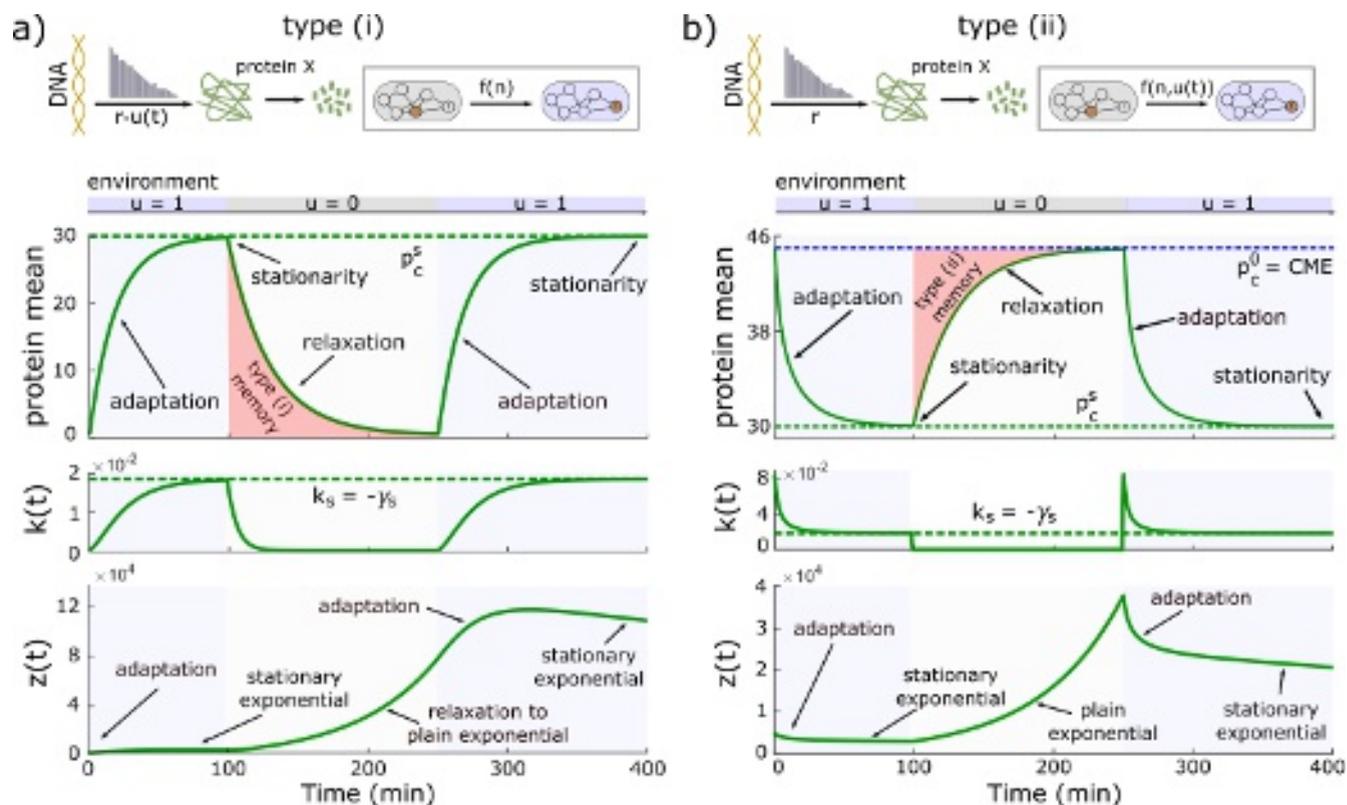
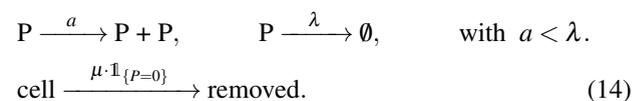


FIG. 2. Two types of cell fate decision systems. (a) Top: Basic motif of a type (i) system in which stochasticity in cell fate decisions is caused exclusively by intrinsic noise that arises after the environmental input, $u(t)$, is applied. The concrete model of this system and its parameter values are the same as already introduced for Figure 1, except that we now consider piecewise constant environmental inputs $u(t) \in \{0, 1\}$ as indicated by the blue and grey bars below the sketch of the system. We assumed that $u(t) = 0$ for sufficiently long before the start of the experiment such that no cell contains any protein at $t = 0$. Bottom: Typical type (i) system dynamics at the population scale ($z(t)$) and at the single-cell scale (protein mean, calculated from the solution of Eq.(5)) for a switching environment $u(t)$. Single-cell dynamics can be grouped into three phases: adaptation to the environmental input, stationarity in the presence of the input, and a phase where the population relaxes to its initial statistics after the input is removed. Early during the relaxation phase, the system still displays a memory of the input (red shaded area). These phases are transmitted to $z(t)$ via the population death rate $k(t)$ and create characteristic features in population growth curves. (b) Top: Basic motif of a type (ii) system in which variability in cell fate decisions is primarily caused by cell-to-cell variability that is already present before the environmental input is applied. For this example, we kept the same chemical reactions as in the type (i) example but removed the dependency on the environmental input $u(t)$ from the propensity functions ($a_1(n) = r$, $a_2(n) = d \cdot n$), whereas the rate of cell fate decisions is now assumed to depend explicitly on $u(t)$. Concretely, we used the function $f(n, u(t)) := u(t) \cdot a \frac{n^\eta}{\theta n + n^\eta}$ and the same parameter values as for the type (i) system. This implies that for $u(t) = 1$ the two models are identical and converge to the same stationary distribution (as can be seen from convergence to a protein mean of 30 for both models when $u(t)$ is maintained equal to one). To determine an appropriate initial condition for the type (ii) example, we need to take into account that protein production does not depend on $u(t)$ in this case, which implies that for $u(t) = 0$ for sufficiently long before the start of the experiment, the population protein distribution will have converged to the stationary distribution of the bursty protein production model in the absence of cell fate decisions. Bottom: type (ii) system dynamics are shown equivalently to panel (a).

makes it easier to mechanistically represent cellular processes such as plasmid replication and unequal distribution of plasmids among mother and daughter cell. However, even if a model is originally constructed in discrete time, it can still be interpreted in continuous time and associated with a CTMC as done, for instance, by Gnuegge et al.⁴⁵. While single-cell models have been extensively studied, it has remained somewhat unclear how they are to be used in a context where the population is growing in media that selects against cells that have lost the plasmid⁴⁶ or when the media is switched between selective and non-selective.

In²⁷, we introduced a simple birth-death process model

of an unregulated 2-micron plasmid carrying a selection marker and an engineered light-inducible recombination system. Concretely, plasmid copy number fluctuations were modeled as:



The first reaction models plasmid replication whereas the second reaction models effective plasmid dilution due to growth and λ is the growth rate of the cells. $a < \lambda$ then ensures that

plasmid copy numbers do not diverge and corresponds to assuming that every plasmid is on average replicated less than once per cell cycle. μ describes the rate at which cells without plasmids are removed in selective media and needs to be set to zero when cells are growing in non-selective media. Here, we focus on emerging population and plasmid copy number dynamics when cell populations are switched back and forth between selective and non-selective media.

Interpreting removal of cells without plasmids as cell fate decisions, we can use the theoretical results presented in the previous sections to determine that the population plasmid copy number distribution must follow Eq.(5) with matrix \mathbf{C} given in the Supporting Information Section S.10.2. whereas the total population size, $z(t)$, should evolve according to

$$\frac{d}{dt}z(t) = \lambda z(t) - k(t)z(t) = (\lambda - k(t))z(t), \quad (15)$$

where $k(t)$ is the population removal rate of cells that have lost the plasmid, follows Eq.(12), and can be calculated using Eq.(5) from the single-cell model.

For continuous growth in selective media, the plasmid copy number distribution should converge to stationarity and $k(t)$ should approach a constant $k_s = -\gamma_s$, where γ_s is the largest eigenvalue of \mathbf{C} as discussed earlier (see, for instance, Figure 2b). In non-selective media, the population should grow at maximal rate λ since $k(t) = 0$. Meanwhile, at the single-cell scale, the plasmid copy number distribution should converge to the stationary distribution of the single-cell model in the absence of selective removal of cells (which should be such that all cells have lost the plasmid with probability one). However, if cells are switched to non-selective media only for a comparably short time span such that stationarity is not reached and some cells still have plasmids, interesting dynamics of the effective population growth rate should emerge. In particular, we may wonder if growth dynamics in selective media depend on the time spent previously in non-selective media and if a type (ii) memory effect (Figure 2b) can be observed experimentally for this application.

To test this, we studied the dynamics of the population fluorescence distribution of a fluorescent reporter protein expressed from the plasmid when cells are switched between selective and non-selective media. To this end, we made use of a platform of parallelized bio-reactors⁴⁷ that operate in turbidostat mode (i.e. continuous dilution and feed of fresh media keep the cell density constant) and allow for external control of the media that is supplied to cells. We first grew cells in selective media for sufficiently long to ensure stationary conditions. Subsequently, we switched cells to growth in non-selective media for different durations in different reactors before switching back to selective media. To be able to observe the biological process at both population and single-cell scale, we monitored the dilution rate of the reactors (determines effective population growth rates) and simultaneously performed automated flow cytometry to quantify population distributions of the fluorescent reporter protein that is expressed from plasmids (provides an indirect readout of changes in plasmid copy numbers). We find that emerging population and single-cell dynamics clearly display all theoretically ex-

pected features of a type (ii) cell fate decision system (Figure 3). Upon switching cells from selective to non-selective media, effective population growth rates quickly increase to the maximal possible growth rate λ since no more cells are selectively removed (plain exponential growth phase). At the same time, relaxation dynamics are at play at the single-cell scale and plasmid copy numbers and cellular fluorescence converge exponentially towards the stationary condition in the absence of the selecting stimulus, which, in this case, is that plasmids are lost in all cells. When cells are switched back to selective media and the selecting stimulus is re-introduced, populations show the expected adaptation phase: directly upon media switch, $k(t)$ spikes high and the measured effective population growth rate drops very low (Figure 3a) since fewer cells have plasmids after the growth period in non-selective media (Figure 3b). Then, due to selective removal of cells without plasmids, the number of cells with plasmids in the population increases (Figure 3b), which leads to an effective population growth rate that increases until the initial stationary condition is re-established such that plasmid loss and removal of cells that have lost the plasmid are dynamically balanced and the population grows exponentially at an effective rate $\lambda - k_s$ (stationary exponential phase). Thereby, the magnitude of the decrease in effective population growth rates as well as the time that it takes for the stationary condition to be re-established depends on the time that populations previously spent in non-selective media (shades of green in Figure 3). These results establish agreement of experiments and theory on a qualitative level.

To test quantitative agreement, we fixed the parameter λ to the measured growth rate of cell populations in non-selective media, $\lambda \approx 0.39/h$, and assumed that cells that have lost the plasmid have a zero growth rate in selective media, which in our model is equivalent to assuming $\mu = \lambda$. The final free model parameter, a , can then be determined as $a = 0.3335/h = 0.85 \cdot \lambda$ by measuring the exponential decay of the average population fluorescence after switching the cells from selective to non-selective media (see Supporting Information, Section S.10.2 for an extended discussion on determining model parameters). A full comparison of the resulting model to experimental data (including replicates) is provided in the Supporting Information, Figures S.12-S.15. We find that, despite the rather simplistic representation of plasmid copy number fluctuations at the single-cell scale, the model leads to predictions on emerging growth rate dynamics and average plasmid copy numbers that are in close agreement with the data.

Finally, we tested one additional prediction that emerges from the presented theory: after a long period of growth in selective media, the measured stationary population fluorescence distribution should correspond to the single-cell process being fully equilibrated except for the direction that leads to removal of cells (see previous sections and Supporting Information, Section S.6). This implies that upon switching cells to non-selective media, we expect the plasmid distribution of the population to change but such that the shape of the distribution remains invariant with the probability of every non-zero plasmid copy number just being scaled by the increas-

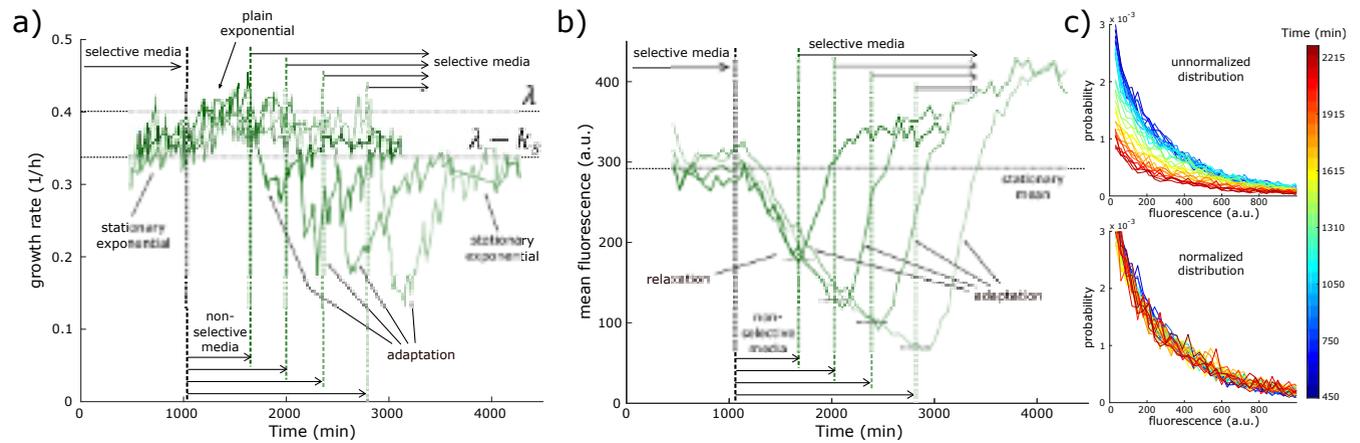


FIG. 3. Plasmid dynamics and population growth. (a) Growth rate dynamics in switching media. Upon switching from selective to non-selective media (black dashed line), effective cell population growth rates increase from stationary exponential growth (at rate $\lambda - k_s$) to plain exponential growth (at rate λ) since cells that lose the plasmid are not removed from the population anymore. When cells are subsequently switched back to selective media (green dashed lines), the effective population growth rates drops very low due to transiently large removal rates, $k(t)$, of cells without plasmids but then converges back towards the initial stationary condition (adaptation phase). The magnitude of the drop in growth rate and the time that it takes to re-establish the stationary condition depend on the length of the growth phase in non-selective media (shades of green). (b) Average levels of a fluorescent protein expressed from the plasmid decrease exponentially towards zero when cells are switched to non-selective media and provide a read-out on plasmid copy numbers in cells of the population. When populations are switched back to selective media, selective removal of cells without plasmids leads to increasing average plasmid copy numbers and fluorescence. Details of the experiment are provided in the Supporting Information, Section S.10. (c) Top: Dynamics of the full population fluorescence distribution after media switch for the experiment shown in the lightest shade of green in panels a,b. The peak of the distribution close to zero, presumably corresponding to cells without plasmids, has been omitted. Bottom: The same distributions as in the top panel are shown after normalization by the fraction of cells whose fluorescence is indistinguishable from autofluorescence.

ing fraction of cells that have lost the plasmid. To test this, we analyzed population fluorescence distributions for the experiment in Figure 3a in which cells are kept in non-selective media the longest (lightest shade of green, results for other experiments are provided in the Supporting Information, Section S.10). Upon media switch, the fraction of cells that have fluorescence levels that are distinguishable from autofluorescence decreases in time as expected (Figure 3c, top). Normalizing the relative frequencies of these fluorescent levels by the fraction of cells whose fluorescence is indistinguishable from autofluorescence at the corresponding time point (see Supporting Information, Section S.10), we find that the population distribution indeed remains invariant except for the increasing fraction of cells that have lost the plasmid (Figure 3c, bottom).

DISCUSSION

The study of stochastic cell fate decision systems is of fundamental importance for understanding emerging population behavior for a large class of important biological problems such as differential responses of cells to environmental perturbations with antibiotics²⁴, anti-cancer treatments²³, or other stimuli⁴⁸. We have shown here that representing stochastic cell fate decision systems as absorbing continuous-time Markov chains (CTMCs) allows one to predict emerging population behavior from theoretical analysis of the single-cell model. Concretely, we used our theoretical results to classify cell fate decision systems into systems of two types in which

randomness in cell fate decisions is either caused by stochasticity of chemical reactions that are triggered by the environmental stimulus (type (i)) or due to pre-existing variability (type (ii)). We then argued that plasmid copy number fluctuations, plasmid loss, and cellular growth in selective media provide an example of a type (ii) cell fate decision system and showed experimentally that theoretically expected dynamics of population growth rates and single-cell processes emerge when cells are switched back and forth between selective and non-selective media.

While the analysis of plasmid copy number fluctuations and plasmid loss constitutes an important example that continues to resurface in systems and synthetic biology, it is but one specific example of a stochastic cell fate decision system. More broadly, the presented theory can be used to study various other important problems. For instance, the required strength of a treatment stimulus that is needed to kill a heterogeneous population of cells can be identified with an eigenvalue of the generator matrix of the single-cell process that triggers cell death (Supporting Information, Figure S.6). Furthermore, our results can be used to classify and quantify different types of memory that populations exhibit in response to environmental stimuli and to study emerging population dynamics for repeated application of such stimuli (Supporting Information, Figures S.10 and S.11) or to study how population dynamics depend on the time scales of single-cell processes (Supporting Information, Section S.7). Apart from the study of natural biological systems, we expect that our results on predicting emerging population behavior from a specification of

the single-cell process will prove to be useful for the construction of synthetic circuits that aim to mimic natural phenotype switching⁴⁹ or differentiation systems⁵⁰ and that need to function in the presence of cell-to-cell variability. By construction, such systems create a coupling between single-cell and population dynamics and designing the synthetic circuit based only on its functionality in individual cells will lead to other than expected dynamics at the population scale, as we have already demonstrated previously²⁷.

The most important aspect of the work presented in this manuscript is that we have shown how one can derive models for the (expected) population composition that incorporate population level processes such as growth and selection within a standard CME-based modeling framework without increasing the dimensionality of the equations. This is in stark contrast to other approaches that have been proposed to model stochastic reaction network in growing cell populations such as the agent-based framework proposed by Thomas and Shahrezaei⁴³. Furthermore, we have shown that both cell growth rates and death/removal rates that depend on the state of the reaction network can readily be included in our framework. The price that to pay for this was that we needed to maintain the simplifying assumptions of constant cell volume and exponentially distributed cell cycle times that are implicit in the CME.

For stochastic cell fate decision systems, we have shown that a model for the expected population composition of a growing population, Eq.(5), can be obtained by no more than adding a single absorbing state to the reaction network and then conditioning the stochastic process on not having reached this state. This highlights that our approach is scalable to reaction networks that are more complex than the simple examples that were studied in this paper. In fact, Eq.(5) can even be more readily obtained by finite state projection³³ than approximations to the standard CME whenever transitions to the absorbing state confer a natural boundary on the number of molecules that the system is unlikely to exceed. That said, our modeling approach inherits the curse of dimensionality from the CME and will become intractable for reaction networks that contain more than just a few chemical species. It is therefore important to further investigate in the future how approaches such as approximations based on Langevin and Fokker-Planck equations^{51,52} or moment-based methods^{53–55}, which are routinely used for the study of stochastic single-cell processes, can be applied in the population context. Another important research direction for the future will be to focus on the case where growth rates of cells depend on the state of the reaction network. In this case, we can still use Eq.(10) to study dynamics of the expected population composition but we cannot anymore associate the equation with an absorbing Markov chain. It will therefore be necessary to establish new theoretical results on the existence of stationary distributions and uniqueness of largest eigenvalues for the case of state dependent growth.

SUPPLEMENTARY MATERIAL

The Supplementary Material contains additional results and discussions, a description of the experiments carried out for the paper, and all supplementary figures referenced in the main text.

ACKNOWLEDGMENTS

This work was partially supported by the ANR (Opt-MC project, ANR-22-CE45-0019) and the European Research Council (ERC-2022-STG, BridgingScales, grant agreement 101075989).

DATA AVAILABILITY STATEMENT

All data presented in the article is publicly available at <https://zenodo.org/record/7313869>.

- ¹H. McAdams and A. Arkin. Stochastic mechanisms in gene expression. *Proceedings of the National Academy of Sciences of the USA*, 94(3):814–819, 1997.
- ²M. Elowitz, A. Levine, E. Siggia, and P. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
- ³L. Potvin-Trottier, N.D. Lord, G. Vinnicombe, and J. Paulsson. Synchronous long-term oscillations in a synthetic gene circuit. *Nature*, 538(7626):514, 2016.
- ⁴C. Zechner, J. Ruess, P. Krenn, S. Pelet, M. Peter, J. Lygeros, and H. Koepl. Moment-based inference predicts bimodality in transient gene expression. *Proceedings of the National Academy of Sciences of the USA*, 109(21):8340–8345, 2012.
- ⁵G. Neuert, B. Munsky, RZ. Tan, L. Teytelman, M. Khammash, and A. van Oudenaarden. Systematic identification of signal-activated stochastic gene regulation. *Science*, 339:584–587, 2013.
- ⁶J. Raser and E. O’Shea. Noise in gene expression: Origins, consequences, and control. *Science*, 309(5743):2010–2013, 2005.
- ⁷M. Acar, J. Mettetal, and A. van Oudenaarden. Stochastic switching as a survival strategy in fluctuating environments. *Nature Genetics*, 40(4):471–475, 2008.
- ⁸S. Spencer and P. Sorger. Measuring and modeling apoptosis in single cells. *Cell*, 144(6):926–939, 2011.
- ⁹N. Rossi, I. El Meouche, and M. Dunlop. Forecasting cell fate during antibiotic exposure using stochastic gene expression. *Communications Biology*, 2:259, 2019.
- ¹⁰G. Balazsi and A. van Oudenaarden. Cellular decision making and biological noise: from microbes to mammals. *Cell*, 144:910–925, 2011.
- ¹¹A. Raj and A. van Oudenaarden. Nature, nurture, or chance: Stochastic gene expression and its consequences. *Cell*, 135(2):216–226, 2008.
- ¹²E. Kussell and S. Leibler. Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, 309(5743):2075–2078, 2005.
- ¹³S. Leibler and E. Kussell. Individual histories and selection in heterogeneous populations. *Proceedings of the National Academy of Sciences of the USA*, 107(29):13183–13188, 2010.
- ¹⁴A. Mayer, T. Mora, O. Rivoire, and A. Walczak. Diversity of immune strategies explained by adaptation to pathogen statistics. *Proceedings of the National Academy of Sciences of the USA*, 113(31):8630–8635, 2016.
- ¹⁵A. Mayer, T. Mora, O. Rivoire, and A. Walczak. Transitions in optimal adaptive strategies for populations in fluctuating environments. *Physical Review E*, 96:032412, 2017.
- ¹⁶T. Nozoe, E. Kussell, and Y. Wakamoto. Inferring fitness landscapes and selection on phenotypic states from single-cell genealogical data. *PLoS Genetics*, 13(3):e1006653, 2017.

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.

PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0160529

- ¹⁷T. Mora and A. Walczak. Effect of phenotype selection on stochastic gene expression. *The Journal of Physical Chemistry B*, 117:13194–13205, 2013.
- ¹⁸K. Mineta, T. Matsumoto, N. Osada, and H. Araki. Population genetics of non-genetic traits: Evolutionary roles of stochasticity in gene expression. *Gene*, 562(1):16–21, 2015.
- ¹⁹C. Tan, P. Marguet, and L. You. Emergent bistability by a growth-modulating positive feedback circuit. *Nature Chemical Biology*, 5:842–848, 2009.
- ²⁰C. González, J.C. Ray, M. Manhart, R.M. Adams, D. Nevozhay, A.V. Morozov, and G. Balázsi. Stress-response balance drives the evolution of a network module and its host genome. *Molecular Systems Biology*, 11(8): 827, 2015.
- ²¹D. Nevozhay, R.M. Adams, E. Van Itallie, M.R. Bennett, and G. Balázsi. Mapping the environmental fitness landscape of a synthetic gene circuit. *PLoS Computational Biology*, 8(4):e1002480, 2012.
- ²²S. Tsuru, N. Yasuda, Y. Murakami, J. Ushioda, A. Kashiwagi, S. Suzuki, K. Mori, B.-W. Ying, and T. Yomo. Adaptation by stochastic switching of a monostable genetic circuit in *Escherichia coli*. *Molecular Systems Biology*, 7:493, 2011.
- ²³F. Bertaux, S. Stoma, D. Drasdo, and G. Batt. Modeling dynamics of cell-to-cell variability in trail-induced apoptosis explains fractional killing and predicts reversible resistance. *PLoS Computational Biology*, 10(10): e1003893, 2014.
- ²⁴N. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler. Bacterial persistence as a phenotypic switch. *Science*, 305(5690):1622–1625, 2004.
- ²⁵Y. Wakamoto, N. Dhar, R. Chait, K. Schneider, F. Signorino-Gelo, S. Leibler, and J. McKinney. Dynamic persistence of antibiotic-stressed mycobacteria. *Science*, 339(6115):91–95, 2013.
- ²⁶M. Ciechonska, M. Sturrock, A. Grob, G. Larrouy-Maumus, V. Shahrezaei, and M. Isalan. Ohm's law for emergent gene expression under fitness pressure. *bioRxiv*, 2019.
- ²⁷C. Aditya, F. Bertaux, G. Batt, and J. Ruess. Using single-cell models to predict the functionality of synthetic circuits at the population scale. *Proceedings of the National Academy of Sciences of the USA*, 119(11): e2114438119, 2022.
- ²⁸G. Chalancon, C. Ravarani, S. Balaji, A. Martinez-Arias, L. Aravind, R. Jothi, and M. Babu. Interplay between gene expression noise and regulatory network architecture. *Trends in Genetics*, 28(5):221–232, 2012.
- ²⁹V. Shahrezaei and V. Marguerat. Connecting growth with gene expression: of noise and numbers. *Current Opinion in Microbiology*, 25:127–135, 2015.
- ³⁰J. Darroch and E. Seneta. On quasi-stationary distributions in absorbing continuous-time finite Markov chains. *Journal of Applied Probability*, 4(1):192–196, 1967.
- ³¹D. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188(1-3):404–425, 1992.
- ³²V. Shahrezaei and P. Swain. Analytical distributions for stochastic gene expression. *Proceedings of the National Academy of Sciences of the USA*, 105(45):17256–17261, 2008.
- ³³B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of Chemical Physics*, 124:044104, 2006.
- ³⁴A. Gupta, J. Mikelson, and M. Khammash. A finite state projection algorithm for the stationary solution of the chemical master equation. *The Journal of Chemical Physics*, 147:154101, 2017.
- ³⁵A. Singh and J.J. Dennehy. Stochastic holin expression can account for lysis time variation in the bacteriophage lambda. *Journal of the Royal Society Interface*, 11:20140140, 2014.
- ³⁶K. Ghusinga, J. Dennehy, and A. Singh. First-passage time approach to controlling noise in the timing of intracellular events. *Proceedings of the National Academy of Sciences of the USA*, 114(4):693–698, 2017.
- ³⁷M. Backenköhler, L. Bortolussi, and V. Wolf. Bounding Mean First Passage Times in Population Continuous-Time Markov Chains. *Lecture Notes in Computer Science*, 12289, 2019.
- ³⁸P. Thomas. Making sense of snapshot data: ergodic principle for clonal cell populations. *Journal of the Royal Society Interface*, 14(136):20170467, 2017.
- ³⁹P. Thomas, G. Terradot, V. Danos, and A. Weiße. Sources, propagation and consequences of stochasticity in cellular growth. *Nature Communications*, 9:4528, 2018.
- ⁴⁰C. Nieto-Acuna, C. Vargas-Garcia, A. Singh, and J. Pedraza. Efficient computation of stochastic cell-size transient dynamics. *BMC Bioinformatics*, 20:647, 2019.
- ⁴¹N. Totis, C. Nieto, A. Küper, C. Vargas-García, A. Singh, and S. Waldherr. Population-based approach to study the effects of growth and division rates on the dynamics of cell size statistics. *IEEE Control Systems Letters*, 5(2): 725–730, 2021.
- ⁴²C. Jia and R. Grima. Coupling gene expression dynamics to cell size dynamics and cell cycle events: Exact and approximate solutions of the extended telegraph model. *iScience*, 26(1):105746, 2023.
- ⁴³P. Thomas and V. Shahrezaei. Coordination of gene expression noise with cell size: analytical results for agent-based models of growing cell populations. *Journal of the Royal Society Interface*, 18(178):20210274, 2021.
- ⁴⁴J. Paulsson and M. Ehrenberg. Noise in a minimal regulatory network: plasmid copy number control. *Quarterly Reviews of Biophysics*, 34(1):1–59, 2001.
- ⁴⁵R. Nügge, T. Liphardt, and F. Rudolf. A shuttle vector series for precise genetic engineering of *Saccharomyces cerevisiae*. *Yeast*, 33:83–98, 2016.
- ⁴⁶B. Shao, J. Rammohan, D. A. Anderson, N. Alperovich, D. Ross, and C. Voigt. Single-cell measurement of plasmid copy number and promoter activity. *Nature Communications*, 12:1475, 2021.
- ⁴⁷F. Bertaux, S. Sosa-Carrillo, V. Gross, A. Fraisse, C. Aditya, M. Furstenheim, and G. Batt. Enhancing bioreactor arrays for automated measurements and reactive control with ReacSight. *Nature Communications*, 13: 3363, 2022.
- ⁴⁸S. Pelet, F. Rudolf, M. Nadal-Ribelles, E. de Nadal, F. Posas, and M. Peter. Transient activation of the HOG MAPK pathway regulates bimodal gene expression. *Science*, 332(6030):732–735, 2011.
- ⁴⁹J.-B. Lugagne, S. Sosa-Carrillo, M. Kirch, A. Köhler, G. Batt, and P. Hersen. Balancing a genetic toggle switch by real-time feedback control and periodic forcing. *Nature Communications*, 8:1671, 2017.
- ⁵⁰C. Aditya, F. Bertaux, G. Batt, and J. Ruess. A light tunable differentiation system for the creation and control of consortia in yeast. *Nature communications*, 12:5829, 2021.
- ⁵¹D. Gillespie. The chemical Langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.
- ⁵²D. Lunz, G. Batt, J. Ruess, and J.-F. Bonnans. Beyond the chemical master equation: stochastic chemical kinetics coupled with auxiliary processes. *PLoS Computational Biology*, 17(7):e1009214, 2021.
- ⁵³A. Singh and J. Hespanha. Approximate moment dynamics for chemically reacting systems. *IEEE Transactions on Automatic Control*, 56(2):414–418, 2011.
- ⁵⁴J. Ruess, A. Miliás-Argeitis, S. Summers, and J. Lygeros. Moment estimation for chemically reacting systems by extended Kalman filtering. *The Journal of Chemical Physics*, 135(16):10B621, 2011.
- ⁵⁵D. Lunz, J.-F. Bonnans, and J. Ruess. Revisiting moment-closure methods with heterogeneous multiscale population models. *Mathematical Biosciences*, 350:108866, 2022.