UNIVERSITY OF CALIFORNIA,
IRVINE

Computational Methods for Modeling and Mapping Cellular Decision Networks

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Chemical Engineering

by

Margaret Jing Yi Tse

Dissertation Committee:
Assistant Professor Elizabeth L. Read, Chair
Assistant Professor Jun Allard
Associate Professor Wendy F. Liu

2017

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank Elizabeth L. Read for going down the rare-event sampling rabbit hole with me for the last five years. It has been a real delight, and I don't have the vocabulary to describe how thankful I am. I am also grateful to the members of my committee, Jun Allard and Wendy F. Liu for their encouragement and exceedingly helpful insights.

Thanks to Tim D. Smith for a year of panicking productively to the point of writing most of Chapter 4 and to Wendy F. Liu for helping push the paper out the door!

I'd like to mention the Read Lab, current and former members, Brian Chu, Royce Sato, Cameron Gallivan, Chris Lee, Mahua Roy, and all our various undergrads for their help, coffee, and conversations.

I would also like to thank Elsevier and Biophysical Society for permission to reprint Chapter 2 and The Royal Society of Chemistry for permission to reprint Chapter 4.

CURRICULUM VITAE

Margaret Jing Yi Tse

EDUCATION

Doctor of Philosophy in Chemical Engineering **2017**
University of California, Irvine *Irvine, CA*

Master of Science in Chemical Engineering **2014**
University of California, Irvine *Irvine, CA*

Bachelors of Science in Chemical Engineering **2012**
California Institute of Technology *Pasadena, CA*

RESEARCH EXPERIENCE

Graduate Research Assistant **2013–2017**
University of California, Irvine *Irvine, CA*

TEACHING EXPERIENCE

Teaching Assistant **2009–2010**
University of California, Irvine *Irvine, CA*

CBEMS 45C **Spring 2013, Spring 2014**

CBEMS 149AB **Winter/Spring 2015, Winter/Spring 2016**

CBEMS 154 **Fall 2017**

REFEREED JOURNAL PUBLICATIONS

An RF-powered micro-reactor for the detection of astrobiological target molecules on planetary bodies **2012**
Review of Scientific Instruments

Scott, V. J., **Tse, M. J.**, Shearn, M. J., Siegel, P. H., and Amashukeli, X

DNA-Binding Kinetics Determines Mechanism of Noise-Induced Switching in Gene Networks **2015**
Biophysical Journal

Tse, M. J., Chu, B. K., Roy, M., and Read, E. L.

Regulation of macrophage polarization and plasticity by complex activation signals **2016**
Integrative Biology

Smith, T. D., **Tse, M. J.**, Read, E. L., and Liu, W. F.

Markov State Models of gene regulatory networks **2017**
BMC Systems Biology

Chu, B. K., **Tse, M. J.**, Sato, R. R., and Read, E. L.

Rare Event Sampling of Epigenetic Landscapes and Phenotype Transitions (*in review*) **2017**
PLOS: Computational Biology

Tse, M. J., Chu, B. K., and Read, E. L.

REFEREED CONFERENCE PUBLICATIONS

State and the single cell: understanding macrophage plasticity **Mar 2015**
CCBS Annual Retreat

Mapping Epigenetic Landscapes of Gene Regulatory Networks by Adaptive Weighted Ensemble Sampling **Feb 2016**
Biophysical Society 60th Annual Meeting

ABSTRACT OF THE DISSERTATION

Computational Methods for Modeling and Mapping Cellular Decision Networks

By

Margaret Jing Yi Tse

Doctor of Philosophy in Chemical Engineering

University of California, Irvine, 2017

Assistant Professor Elizabeth L. Read, Chair

Cell phenotypes are controlled by complex interactions between genes, proteins, and other molecules within a cell, along with extracellular signals. Gene regulatory networks (GRNs), which describe these interactions mathematically, are multi-stable dynamical systems, in which attractor states represent cell phenotypes. Transitions between these states are thought to underlie critical cellular processes, including cell fate-decisions, phenotypic plasticity, and carcinogenesis. In principle, a GRN model can produce a map of possible cell phenotypes and phenotype-transitions, potentially informing experimental strategies for controlling cell phenotypes. Such a map could have a profound impact on many medical fields, ranging from stem cell therapies to wound healing. As such, there is increasing interest in the development of theoretical and computational approaches that can shed light on the dynamics of these state-transitions in multi-stable gene networks. In this work, we approach the problem of understanding cellular decision-making on two fronts. First, we develop and extend rare-event stochastic simulation methods, toward efficient characterization of the global dynamics of multi-stable stochastic systems, such as GRNs. When applied to a mutual inhibition network motif and a model of pluripotency in stem cells, our sampling methods demonstrate that spontaneous cell phenotype transitions involve collective behavior of both regulatory proteins and DNA, and that transition dynamics are sensitive to parameters governing transcription factor-DNA binding kinetics. Our approach significantly expands the capability of stochastic

simulation to investigate gene regulatory network dynamics. In the second portion of this work, we model the dynamic response of macrophages to complex stimuli by inferring cell-decision networks from data, in the absence of detailed molecular information. Macrophage activation has been described as a continuum, and different stimuli lead to M1, M2, or mixed phenotypes. Flow cytometry experiments performed in the Liu lab at UCI found discovered that macrophages acquire a mixed activation state when exposed to a combination of LPS, IFN-γ, IL-4, and IL-13. Additionally, mathematical modeling of candidate regulatory networks indicates that a complex inter-dependence of M1- and M2-associated pathways underlies macrophage activation. Together these results corroborate a continuum model of macrophage activation and demonstrate that phenotypic markers evolve with time and with exposure to complex signals.

Chapter 1

Introduction

1.1 Gene Regulatory Networks

Many multicellular organisms are comprised of multiple cell types, most of which have the same genome. The vast difference in cellular function between, for example, a cell in a hair follicle and a cardiomyocyte is determined by differential expression across the genome. A gene regulatory network (GRN) is the complex network of biochemical interactions between genes, proteins, transcription factors, and other signaling molecules that determines the expression level of genes within a cell and, thereby, its phenotype. Nonlinear interactions in GRNs give rise to multiple high-probability states (metastability) which correspond to specific gene expression profiles that are unchanging in time or resistant to slight perturbations [124, 92], i.e. observable cell phenotypes [159]. Transitions between cell phenotypes underlie critical cellular processes such as differentiation [62], reprogramming [15], and carcinogenesis [136]. Understanding the molecular level interactions that control these transitions could potentially inform strategies for controlling cell phenotype. As such, there is considerable interest in generating a comprehensive map of phenotype states accessible to a model gene

regulatory network, the relative probabilities of these states, and the likely transition paths connecting them.

1.2 Stochasticity in Gene Regulatory Networks

Over the past twenty years, the inherent noisiness of gene expression has been amply demonstrated. Stochastic biomolecular fluctuations can significantly impact dynamics of expression due to small-number effects[19, 125, 104]. Positive feedback, which is common in regulatory networks, can amplify these fluctuations, giving rise to large-scale changes in gene expression programs. As such, intrinsic biomolecular noise is thought to underlie phenotypic variability in cell populations[13, 162, 71, 29]. Stochastic processes can also allow for spontaneous noise-induced transitions between metastable gene expression states. This stochastic state-switching can be advantageous by priming cells to diversify according to alternative developmental programs[27, 37], or by promoting survival of microorganisms or cancer cells in fluctuating environments[17, 5, 136]. These findings highlight the need to understand how cellular networks achieve—or remain resistant to—noise-induced switching.

1.3 Modeling of Gene Regulatory Networks

Due to the size and complexity of GRNs, there are many methodological frameworks for modeling and analyzing them at different levels of abstraction, trading biological accuracy for ease of computation and ability to incorporate a larger number of biochemical species (Figure 1.1). Among the simplest and most abstracted models are logical models, such as Boolean networks [74]. In Boolean networks, all chemical entities are abstracted as being in an activated or highly expressing state (1) or an inactive state (0). This abstraction greatly reduces the amount of biological data necessary to construct the network model while also simplifying

any methods to calculate dynamics of interest. More biological accuracy can be captured in continuous models, such as ordinary differential equations (ODEs) or stochastic differential equations (SDEs) [73]. However, continuous models often assume that gene expression levels depend directly and deterministically on concentrations of regulatory proteins (via "Hill"-type functions) [85], whereas increasingly it is understood that genes can stochastically transition between active and inactive states. These stochastic transitions have functional consequences, for exampling in priming cells to develop along alternative differentiation paths in developmental networks [3, 13] and in contributing to bistable expression patterns of cytokines in immune cells [126]. Likewise, SDEs suffer from ignoring molecular discreteness, especially due to the single-molecule nature of genes. The most complete models of GRNs are single-molecule models, such as the chemical master equation (CME) which describes the time evolution of the probability density vector of all possible cell state configurations. However, full analytical solutions of the CME are impossible due to the "curse-of-dimensionality" [118] where each additional chemical species exponentially increases the size of the system. To bypass this limit, single time trajectories of the CME can be simulated using the stochastic simulation algorithm [51]. A graphical outline of how the epigenetic landscape of all accessible phenotypes in a GRN can be mapped through the CME is presented in Fig. 1.2.

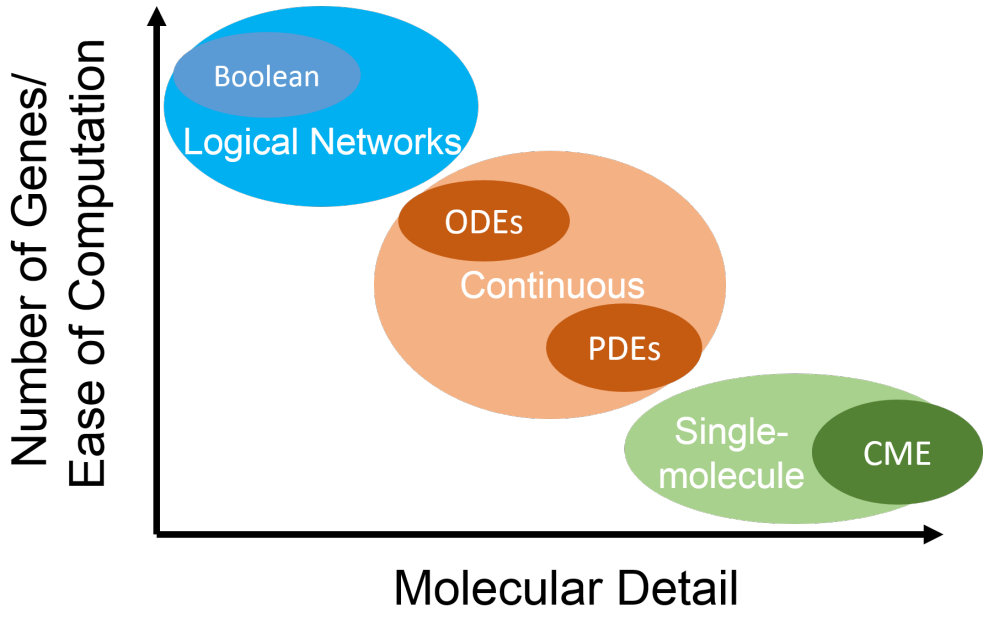

Figure 1.1: **Comparison of Gene Regulatory Network Models.** Models are listed along a scale of increasing molecular detail on the x-axis and increasing scalability and ease of simulation on the y-axis. The most abstracted and simplified models are logical networks, which can be used to analyze extremely large biochemical networks with a minimal amount of measured data. On the other side of the spectrum are single-molecule models of GRNs, where every change in gene state or copy number is simulated in a Monte Carlo step. While these networks have significantly more molecular detail, specialized computational methods are required to make their solution tractable for larger biochemical systems.

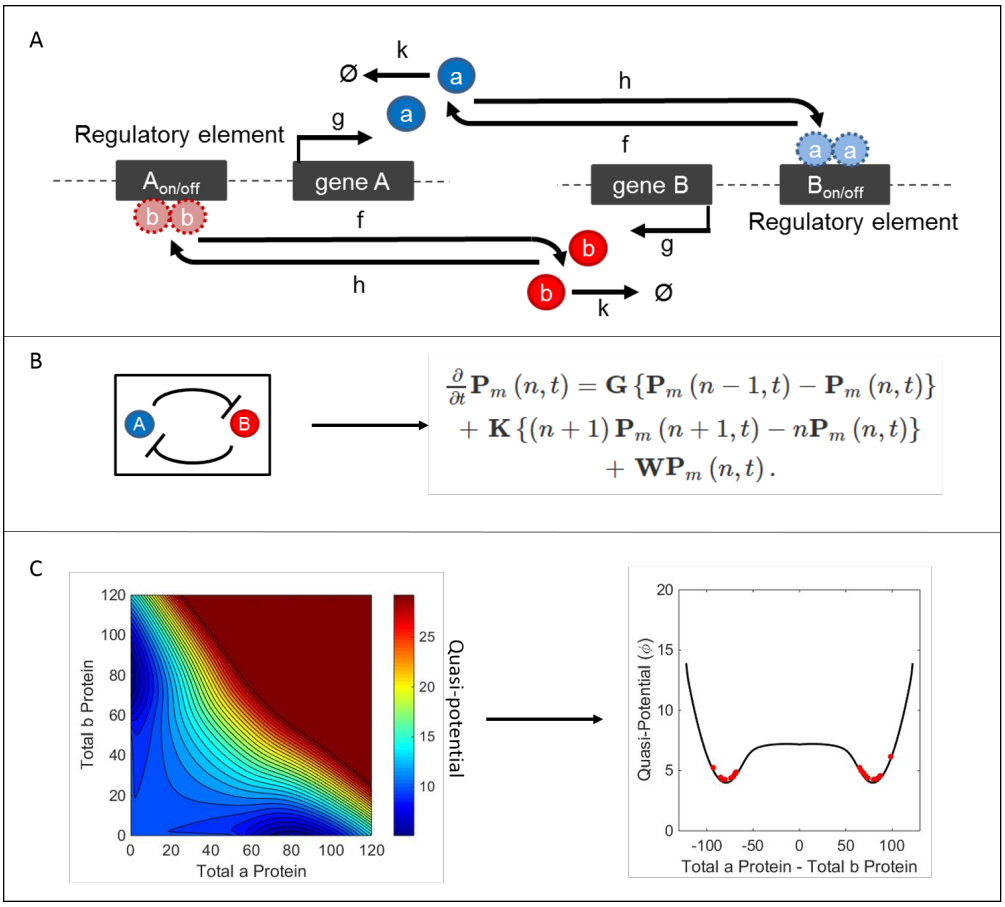

Figure 1.2: **Gene Regulatory Networks (A) Can be Expressed as a Chemical Reaction Network (B), Which Describes the State Space of All Possible Gene Expression Profiles (C).** (A) The epigenetic configurations of a two gene mutual repression genetic regulatory network. The binding and unbinding of transcription factors to the genes can then be described as a chemical reaction network of m genes (B), where G, K, and W describe the probabilities of protein production, protein degradation, and change in epigenetic configuration for each gene. Solving the reaction network gives the probability distribution of all possible gene expression profiles (all protein copy numbers and epigenetic configurations). The negative log of the probability distribution can be projected onto the levels of protein copy numbers alone to show the gradient of the system (C). Using an order parameter, a variable that describes progress along a transition, the system can be converted into a 1-dimensional potential, where trajectories (single solutions of the chemical reaction network) shown in red can spontaneously jump between wells.

1.4 Rare-event Sampling

While stochastic single molecule networks provide the most molecular detail and are most faithful to biological processes, mapping out the global dynamics of these networks is extremely computationally intensive, except for the smallest of systems. Furthermore, GRNs are metastable, and transitions between multiple metastable states can be infrequent on the time-scale of the long-lived states of the network. Conventional SSA of GRNs can spend long waiting times before capturing a single transition of interest. Specialized computationally methods called rare-event sampling methods help redistribute computational effort from high-probability regions of state space to low-probability regions without modifying the underlying dynamics. Rare-event sampling methods typically involve discretizing the state-space of the system in order to access long-timescale kinetics on the basis of short-timescale trajectories in each region [8]. The state-space is naturally defined in terms of copy numbers and conformational/binding configurations of all biomolecular species, but even the smallest 2-gene network (i.e., the toggle switch [48]) has >2 species (including DNA promoter or regulatory binding sites, mRNA transcripts, protein products with multiple configurations, and so on, depending on the level of detail of the model). This challenge can be addressed by sampling along a single, relevant order parameter, such as a progress coordinate for the transition of interest, as pioneered by Allen and Ten Wolde with the Forward Flux Sampling (FFS) [7] method. Alternatively weighted ensemble methods can use adaptive partitioning to re-frame high-dimensional state space into a set of quasi-1-dimensional sampling regions [70, 38, 33]. This approach enables sampling collective dynamics involving multiple system variables, without prior knowledge of an order parameter or the nature of dynamics in the transition region. In chapter 2 and 3 of this work, we present two methods to expand stochastic simulation methods to GRNs with rare-events.

1.5 Macrophage Polarization

In chapter 4 of this work, we explore methods for inferring GRNs from a limited set of biological data, specifically in macrophage polarization. Macrophages are key effector cells of the immune system. They have important roles in phagocytosing and clearing foreign material and dead cells, and they are responsible for many signals in the immune system. They are many (at least four) sub-phenotypes of macrophages, but there are two macrophage activation states that are most prominent in the immune response [109]. Classically activated, or M1 macrophages are typically induced by LPS (foreign antigens) or Th1 cytokines, and are the first responders to foreign material and pathogens in the body, engulfing them and attacking with an oxidative onslaught. Alternatively activated M2 macrophages, which are induced by Th2 cytokines or by parasitic infection, are responsible for anti-inflammatory signaling and enhanced matrix production [43]. Because of the vastly different duties of these macrophage phenotypes, controlling the dominant macrophage response can lead to drastically improved medical outcomes for a variety of diseases [121].

Previously, macrophages found to have both M1 and M2 characteristics were thought to be unstable states either in transit to either the M1 or M2 phenotype, or artificially held in place by continuous mixed polarizing signals [109]. However, there is experimental evidence that M1+M2 cells are stable phenotypes, and the macrophage effector response exists on a spectrum of M1 and M2 responses [138]. Tumor-associated macrophages (TAMs) in human cutaneous squamous cell carcinoma appear to represent a mixed population of M1, M2, and bi-activated M1+M2 cells [120]. Similarly, CD4+ T cells display a continuously tunable mixed state under mixed input conditions of polarizing signals [12]. A quantitative understanding of the prevalence of these mixed populations or their dynamics in GRNs is missing. By combining experimental data with stochastic simulations, we will characterize the response of macrophages to defined inputs and understand macrophage cross-activation at a population level.

Chapter 2

DNA-Binding Kinetics Determines Mechanism of Noise-Induced Switching in Gene Networks

2.1 Introduction

Multi-stable dynamics in gene regulatory networks has been proposed as the basis for the existence of diverse cell types[74, 68, 66]. In this view, the biochemical interactions encoding gene networks give rise to complex, nonlinear expression dynamics. Distinct gene expression states—cellular phenotypes—are self-stabilizing attractors of the dynamical system. Transitions between attractors correspond to critical cellular processes, including developmental fate decisions[167, 158], cellular reprogramming[66], phenotype switching[61], and carcinogenesis[69]. As such, there is interest in characterizing the global dynamics of complex multi-stable gene networks to gain insight into the relative stability of cell states and the processes by which transitions between states can occur.

Common approaches to modeling gene networks formulate system dynamics as a set of stochastic biochemical reactions encompassing, for example, transcription factor binding and unbinding to regulatory sites on DNA, gene transcription and translation, and degradation/dilution of mRNA and protein molecules. The associated Master Equations are amenable to Monte Carlo simulation (e.g., by the Gillespie algorithm[51]), which can exactly account for intrinsic stochastic fluctuations and molecular discreteness. Many studies of gene regulatory networks employ continuum approximations and dimensionality-reduction to aid in the analysis of complex dynamics. For example, occupancies of binding sites on DNA are often assumed to be in quasi-equilibrium with gene expression levels (i.e., concentrations of expressed proteins), which leads to nonlinear Hill function expressions for gene interactions[75, 56, 48]. This assumption stems from a separation of timescales, when rates of binding and unbinding of regulatory proteins to DNA sites are fast relative to the rates of protein synthesis and degradation. This kinetic regime has been termed the "adiabatic" limit, in analogy to electron transfer reactions[135, 131].

Recent studies have focused on gene network dynamics in cases where this separation of timescales does not apply[167, 130, 83, 45, 49]. In eukaryotic regulation, highly complex regulatory processes such as chromatin remodeling (in contrast to simple protein binding) drive fluctuations in gene expression. Such processes occur on relatively slow timescales; slow fluctuations in chromatin structure have been identified as a major source of gene expression noise in eukaryotes[125, 104]. While knowledge of the detailed biochemical reactions underlying such higher-order regulation is generally lacking, theoretical studies have explored how different kinetic regimes contribute to the stability of gene expression states. Notably, cell-states have been found to be most stable (most resistant to noise-induced switching) in the adiabatic limit, where fast kinetics at DNA regulatory sites allows rapid response of gene expression states to local concentrations of regulatory proteins[45, 157]. This has led to the hypothesis that a "weakly-adiabatic" kinetic regime allows for relatively frequent stochastic state-switching, and thus developmental plasticity, in pluripotent stem

cells[167, 130, 46].

In addition to exploring the principles governing the stability of gene expression states, recent theoretical studies have predicted transition paths (or reaction coordinates) of stochastic state-switching in regulatory networks. These paths describe the most probable step-wise changes over multiple species in the network (e.g., the time-dependent changes in gene expression patterns) that occur as the system moves from one metastable state to another. Approaches for calculating transition paths have largely adopted analytical or numerical methods based on large deviation theory from chemical physics [92, 127, 16, 170, 160], related path integral approaches[167, 158, 83, 168], or Monte Carlo simulations[46, 108, 145]. Valuable insights into the dynamics of noise-induced switching in gene networks have emerged from these studies, including the significance of nonequilibrium phenomena (e.g., the irreversibility of switching paths) in gene network dynamics.

Gaining quantitative insight to stochastic state-switching in gene networks is mathematically and computationally challenging. Stochastic fluctuations, overlapping temporal scales, and large numbers of interacting species can preclude the use of analytical approaches. Approximation methods may not accurately predict switching mechanisms driven by intrinsic biochemical fluctuations. For example, mean-field approximations have been found to distort multi-stable landscapes in some biochemical networks[133, 86, 95, 14]. On the other hand, brute force Monte Carlo simulations can be plagued by inefficiency because state-switching may be a rare event: long waiting times between transitions mean that simulations capture few, if any, switching events.

Numerical rare event sampling algorithms provide an alternative approach, by preferentially simulating events of interest without modifying the underlying system dynamics[7, 35, 38]. Moreover, these algorithms aid interpretation of large amounts of noisy simulation data by providing automated means of extracting essential dynamical properties[134]. These types of approaches have been adopted for the study of biochemical networks, including

circadian rhythms[36], enzymatic cycles[38], and genetic switches[108]. So-called "string"-based sampling algorithms[35, 152, 6] are especially well-suited to discovery of transition paths for rare events in complex dynamical systems.

In this study, we employed a recently developed weighted-ensemble string (WE-string) simulation method[6] to study noise-induced switching in a set of genetic toggle switch networks. By efficiently predicting transition paths involving collective changes over all species in the biochemical network, the approach uncovered new insights into how spontaneous switching depends on fluctuations of both protein number and DNA-binding occupancies. In particular, we found that the kinetics of protein binding and unbinding to DNA regulatory sites controls the mechanism of switching by determining whether the switching event is driven by these binding/unbinding events, or is instead driven by fluctuations in protein copy numbers. Our results potentially inform strategies to perturb gene network dynamics in order to stabilize or destabilize particular cellular states, or drive a desired cellular transition.

2.2 Modeling and Simulation Approach

2.2.1 Toggle switch models.

To explore the capability of the WE-string simulation method to predict transition paths in gene networks, we applied it to a set of related chemical network models for the genetic toggle switch. In this ubiquitous motif, two genes mutually repress one another through the action of their protein products, giving rise to a bistable gene expression pattern. Described in detail for the lysis/lysogeny decision of bacteriophage λ [123], the mutual repression motif is recognized as underlying a wide variety of cellular decisions, from embryonic development[113] to haematopoiesis[169] to specialization of T cell subsets[63]. The discovery of this motif paved the way for theory-based approaches to the design of synthetic gene circuits[48].

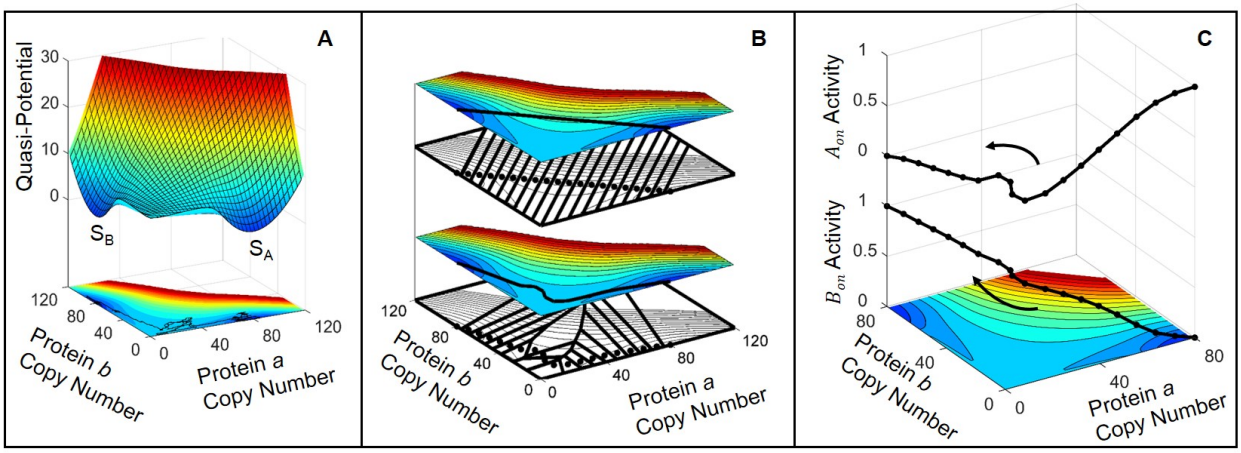

Figure 2.1: **The Weighted Ensemble String Method Identifies the Most Probable Transition Path in Many Sub-Spaces** (A) (Top) The quasi-potential surface $\phi = -\ln(P_S)$ for a bistable gene network. States with lower potential (bluer colors) are more stable (color scale is same for all panels). (Bottom) A single stochastic simulation trajectory that switches from the S_A attractor basin (high expression of protein a encoded by gene A) to the S_B basin (high expression of protein b by gene B). ϕ is projected onto sub-spaces representing a and b protein copy numbers.
(B) The simulation method predicts the most-probable transition path by utilizing a "string" connecting two states of interest to adaptively partition the state space and sample short stochastic trajectories initialized in each partition. (Top) The simulation is initialized with a string representing a guess path for the $S_A \rightarrow S_B$ transition (corresponding Voronoi partitions plotted underneath). (Bottom) After the simulation has converged, the string (and corresponding partitions) indicates the most probable transition path.
(C) The predicted transition path contains information on all network species: for the toggle switch network, it contains information on protein copy numbers and gene activities (DNA occupancies). (Top) Transition path for the $S_A \rightarrow S_B$ switch, projected onto a, b, and A_{on} sub-spaces, where A_{on} is the probability that the regulatory site for Gene A is unbound, rendering it "on" (active). (Bottom) Same transition path projected onto a,b, and B_{on} sub-spaces, where B_{on} is the probability that the regulatory site for Gene B is unbound.

2.2.2 Quasi-potential landscape.

The stability of cell states can be expressed by a potential landscape over the gene network state-space[158, 155, 67]. We use a definition for a quasi-potential function based on the stationary probability distribution: the quasi-potential ϕ is taken to be $\phi = -\ln(P_S)$, where P_S is the stationary (steady-state) probability over the state-space (Fig. 2.1A). Defined in this way, areas of state-space with low potential are most probable, and therefore most stable. However, this quasi-potential does not completely describe dynamics in non-equilibrium, open systems such as biochemical networks[158, 170, 35]. Additionally, characterizing the quasi-potential landscape of a GRN by through brute force Monte Carlo simulation is computationally intractable due to rare events in GRNs.

2.2.3 Transition path sampling algorithm.

Predicting the most probable switching path by brute-force simulations is inefficient, because switching events are infrequent, and individual trajectories may be highly variable (Fig. 2.1A, 6.1, and 6.2). To predict the most probable path between two states of interest, the WE-string algorithm preferentially samples dynamics in the transition region (a low-probability area of state-space), statistically merging information from many short trajectories. The algorithm achieves this by combining Weighted Ensemble rare event sampling [70] with a string-based adaptive discretization method for partitioning the state-space[41]. Adelman and Grabe recently introduced the method[6], and showed that it compares favorably to other path sampling methods in terms of efficiency, and could successfully sample dynamics in the space of many collective variables in an application to protein conformational change.

Rare-event sampling methods typically involve discretizing the state-space of the system in order to access long-timescale kinetics on the basis of short-timescale trajectories in each region[174]. The computational expense of discretizing a system evenly over N dimensions

scales exponentially with N, which quickly becomes intractable for gene networks. The state-space is naturally defined in terms of copy numbers and conformational/binding configurations of all biomolecular species, but even the smallest 2-gene network (i.e., the toggle switch) has >2 species (including DNA promoter or regulatory binding sites, mRNA transcripts, protein products with multiple configurations, and so on, depending on the level of detail of the model). This challenge can be addressed by sampling along a single, relevant order parameter, such as a progress coordinate for the transition of interest, as pioneered by Allen and Ten Wolde with the Forward Flux Sampling (FFS) method[8]. Alternatively, string-based methods adaptively partition space along a "string": a 1-dimensional path winding through a high-dimensional space, which connects two states of interest. This approach enables sampling collective dynamics involving multiple system variables, without prior knowledge of an order parameter or the nature of dynamics in the transition region[34].

The string is defined by a set of N_{str} evenly-spaced "nodes", or points in the full state-space; these nodes define the centers of N_{str} regions. All points in the space lie in a region which is defined by the nearest string node. Thus, the string nodes are the generating points of Voronoi polyhedra. The string is initialized as a "guess"-path connecting two states (Fig. 2.1B). For each iteration of the simulation, multiple weighted stochastic trajectories ("replicas") in each region are simulated for a time τ. After each iteration, the WE method statistically combines and duplicates weighted replicas as needed to ensure unbiased sampling of the partitioned space. After T_{move} iterations of τ, the string position is updated by moving the string nodes towards the average position of replicas in each region collected from the last T_{avg} iterations of τ, using a procedure of moving, smoothing, and reparameterization. Details of simulation parameters and algorithm implementation are in Methods. A more detailed description of the general algorithm is in reference [6].

Because the string is defined as a list of points in the system state-space, the converged string contains information about the probable step-wise changes along all network dimensions over

the course of the transition. For our model toggle switch network (Fig. 2.2), this means that the predicted transition path includes information both about gene expression levels (copy numbers of protein products, a and b) and gene activity states (i.e., likelihood of gene states A_{on} and B_{on}), which reflect the occupancy states of the DNA binding sites (Fig. 2.1C).

2.3 Methods

2.3.1 Toggle Switch Reaction Network

We considered two previously-studied variants of the toggle switch, which we term the "base network"[133] and the "network with explicit dimerization" (Fig. 2.2 and Equations 2.1,2.2,2.3)[108]. In both variants, protein dimers act as repressors of the other competitor gene. The switch is symmetric with respect to the rate parameters governing the behavior of each of the two genes, A and B. Translation and transcription are subsumed into a single biochemical reaction for synthesis of the two encoded proteins, a and b. The two genes exist in one of two occupancy states, giving rise to binary regulation ($A_{\text{on/off}}$, $B_{\text{on/off}}$), where only the "on" states contribute to protein synthesis. We furthermore studied the "general" and "exclusive" versions of the switch; the exclusive switch assumes that the a_2 and b_2 repressors competitively bind to a single regulatory site that controls both genes; therefore, both genes cannot simultaneously be in the "off" state (i.e., the doubly repressed state is disallowed (equations in SI 6.1)).

Protein synthesis and degradation reactions are given by:

$$\begin{aligned} \mathrm{A_{on}} &\xrightarrow{\mathrm{g}} \mathrm{A_{on}} + \mathrm{a} & \qquad \mathrm{B_{on}} &\xrightarrow{\mathrm{g}} \mathrm{B_{on}} + \mathrm{g} \\ \mathrm{a} &\xrightarrow{\mathrm{k}} \emptyset & \qquad \mathrm{b} &\xrightarrow{\mathrm{k}} \emptyset \end{aligned} \tag{2.1}$$

In the base network, dimerization is neglected, and gene repression is represented by:

$$\mathrm{A_{on}} + 2\,\mathrm{b} \underset{\mathrm{f}}{\overset{\mathrm{h}}{\rightleftharpoons}} \mathrm{A_{off}} \qquad \mathrm{B_{on}} + 2\,\mathrm{a} \underset{\mathrm{f}}{\overset{\mathrm{h}}{\rightleftharpoons}} \mathrm{B_{off}} \tag{2.2}$$

These reactions give rise to an overall rate of gene inactivation of $hx(x-1)/2$, where x is the repressor protein encoded by the competing gene. In the network with explicit dimerization, the above reactions are replaced by:

$$\begin{array}{ll} \mathrm{A_{on}} + \mathrm{b_2} \underset{\mathrm{f}}{\overset{\mathrm{h}}{\rightleftharpoons}} \mathrm{A_{off}} & \mathrm{B_{on}} + \mathrm{a_2} \underset{\mathrm{f}}{\overset{\mathrm{h}}{\rightleftharpoons}} \mathrm{B_{off}} \\ 2\,\mathrm{a} \underset{\mathrm{u}}{\overset{\mathrm{d}}{\rightleftharpoons}} \mathrm{a_2} & 2\,\mathrm{b} \underset{\mathrm{u}}{\overset{\mathrm{d}}{\rightleftharpoons}} \mathrm{b_2} \end{array} \tag{2.3}$$

Rate parameters for the reaction network are given in Table 2.1. The effect of DNA-binding kinetics was studied by varying the rates h and f (binding and unbinding of transcription factors to DNA, respectively). The binding equilibrium constant $X_{eq} = f/h$ was maintained at a constant value, such that the locations of the two stable states are preserved.

2.3.2 Calculation of quasi-potential

We calculated P_S and ϕ by expressing the Chemical Master Equation (CME) for the network in matrix form: $\frac{dP}{dt} = AP$, where A is the transition rate matrix over a truncated state-space ($0 \leq a,\ b \leq 120$), and solving for the eigenvector associated with the zero-eigenvalue[73] using MATLAB[1]. For parameter regimes giving relatively short waiting times between switching events (i.e., Parameter Sets II and III), we also obtained P_S from a single long Gillespie

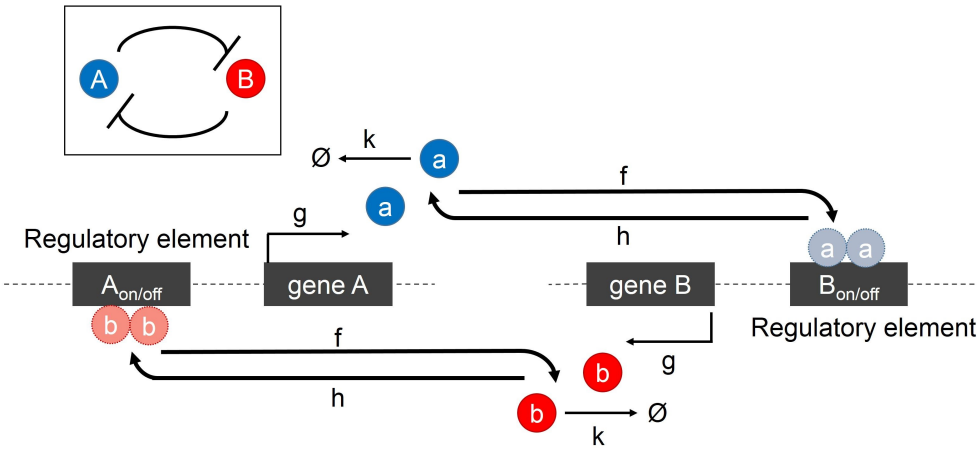

Figure 2.2: Schematic diagram of biochemical reactions in the genetic toggle switch network (inset: basic network motif of two mutually repressing genes). Gene A codes for repressor protein a (blue), and gene B codes for protein b (red). Each gene is controlled by a regulatory element (e.g., the promoter). When the regulatory element is unbound, the gene is on (active), or free to express protein. When it is bound by a homodimer of the opposing repressor, the gene is off (inactive).

simulation trajectory, using BioNetGen(ver. 2.2.2)[44]. The quasi-potential landscapes obtained from each method were indistinguishable.

2.3.3 Transition path simulation

The WE-string sampling method was adapted from Adelman *et al.*[6]. The procedure for string-smoothing, parameterization, and separate sampling of forward and backward paths (by separating replicas according to their most recently-visited basin) was adapted from Dickson, *et al.*[36]. String ends were fixed in the two stable attractor basins, S_A and S_B, corresponding to states with high Gene A expression and high Gene B expression, respectively. S_A and S_B were defined as hyper-spheres with unit radius in N-dimensional state-space (base network: N=4 a, b, $A_{on/off}$, $B_{on/off}$, explicit dimerization network: N=6 a, b, a_2, b_2, $A_{on/off}$, $B_{on/off}$). The basin centers were found by identifying the minima of the quasi-potential for each parameter set. The number of string nodes, N_{str}, was 20, and the number of replicas in each region (N_{rep}) was 150. Definitions and values of additional simulation parameters are given in Table 6.1. "Switch Progress" is defined as the normalized distance along the converged string. That is, for the forward path, the progress at the i$^{\mathrm{th}}$ string node is defined by D_i/D_{AB}, where D_i is the Euclidean distance to the i$^{\mathrm{th}}$ node from the center of S_A, and D_{AB} is the total distance along the string. The simulation was determined to have converged after the position of the string remained stable over > 1000 string-movements. Convergence was determined from the root mean-squared difference between the current node positions and the running average of the previous 100 positions (Figs. 6.3). Sampling code was written in MATLAB and Gillespie trajectories were simulated using BioNetGen[44].

2.3.4 Calculation of switching rates

Switching rate constants k_{AB} were estimated from brute force Gillespie simulations when possible (faster switching regime: parameter sets II and III) or WE sampling, in cases where capturing switching events from brute force simulations was intractable (rare-switching regime: parameter set I). The network parameters are symmetric, such that $k_{AB} = k_{BA}$. For the brute force simulations, k_{AB} was estimated by inverting the mean first passage time of $(S_A \rightarrow S_B)$ transitions from 1000 trajectories. The WE rate estimation followed [173], using a linear partitioning of state space along the progress coordinate λ with bins of unit length. λ is given by $\lambda = na - nb$, where $na = a + 2B_{off}$, $nb = b + 2A_{off}$ (general switch) or $na = a + 2a_2 + 2B_{off}$, $nb = b + 2b_2 + 2A_{off}$ (exclusive switch); this is the same progress coordinate as used previously[108, 8]. The simulation timestep was chosen to be the same as that of the WE-string simulation for the corresponding parameter set; 150 replicas were simulated in each bin.

Network Parameters	$h\ [k^{-1}]$	$f\ [k^{-1}]$	$g\ [k^{-1}]$	$k\ [k^{-1}]$	$k_{AB}\ [t^{-1}]$ (general)		$k_{AB}\ [t^{-1}]$ (exclusive)	
Baseline					*Mean*	*Std.*	*Mean*	*Std.*
I	10^2	10^4	80	1	1.1×10^{-6}	5×10^{-7}	5.9×10^{-10}	6×10^{-11}
II	10^{-1}	10	80	1	2×10^{-4}	9×10^{-5}	3×10^{-4}	2×10^{-4}
III	10^{-2}	1	80	1	6×10^{-3}	2×10^{-3}	2×10^{-3}	8×10^{-4}
Explicit-Dimerization							$k_{AB}\ [t^{-1}]$ (exclusive)	
	$h\ [k^{-1}]$	$f\ [k^{-1}]$	$g\ [k^{-1}]$	$k\ [k^{-1}]$	$d\ [k^{-1}]$	$u\ [k^{-1}]$	*Mean*	*Std.*
I	8.89×10^3	10^5	80	8.944	1	1	3.3×10^{-10}	7×10^{-11}

Table 2.1: Rate parameters and calculated $S_A \to S_B$ switching rates k_{AB}, for all studied genetic toggle switch variants (base and explicit-dimerization networks, general and exclusive). Rate parameters correspond to the schematic in Fig. 2 and Methods. h and f: binding and unbinding rates (respectively) of proteins to DNA. g: protein synthesis, k: protein degradation, d and u: binding and unbinding rates of monomer proteins to form homodimers. All network parameters are reported in units of $[k^{-1}]$ (inverse degradation rate) in the base network and $[u^{-1}]$ (inverse dimer unbinding rate) in the explicit-dimerization network. Parameter set I corresponds to the adiabatic regime, where the propensities of proteins binding and unbinding to DNA ($hx(x-1)/2$ and f, respectively) are large relative to g and k.

2.3.5 Validation of transition path simulation

Validation of switching mechanisms predicted by the transition path simulations was done by comparison to brute-force switching trajectories (where possible), and by committor analysis[94]. Successful switching trajectories harvested from brute force simulations were binned according to the progress coordinate λ. Additional simulations assessed committor probabilities along the converged strings: these reflect the probability that simulation replicas in a given region will next "commit" to the final target state, rather than returning to their most recently-visited basin.

2.4 Results

2.4.1 Stability of phenotypes depends on DNA-binding kinetics

We calculated the quasi-potential landscape for the genetic switch, and varied the parameters governing DNA-binding/unbinding kinetics between the adiabatic (fast kinetics, Figs. 2.3 and 2.4, parameter set I) and non-adiabatic (slow kinetics, Figs. 2.3 and 2.4, parameter set III) regimes. By varying rates of unbinding (f) and binding (h) while maintaining their constant ratio $X_{eq} = f/h$ the bistability and locations of the attractor basin centers are preserved, but the barrier height separating the two states changes (Figs. 6.4 and 6.5). Our results show that the barrier height decreases with decreasing adiabaticity, in agreement with previous results([133]). This difference is also reflected in the calculated rate constants k_{AB} for spontaneous switching from the attractor state S_A to S_B (2.1). The rate constant for spontaneous switching between metastable states has been identified as a measure of stability or robustness of gene expression states to noise. In agreement with previous studies[167, 130, 45, 157, 133], we find that the rate of switching slows with increasing adiabaticity, indicating that gene expression states are more stable with faster DNA-binding kinetics. This trend was preserved for all variants of the switch (2.1). For parameter set I, k_{AB} is much lower for the exclusive switch than for the general, while the switching rates for the two variants are of the same order of magnitude for sets II and III. Interestingly, the switching rates are nearly identical for parameter set II.

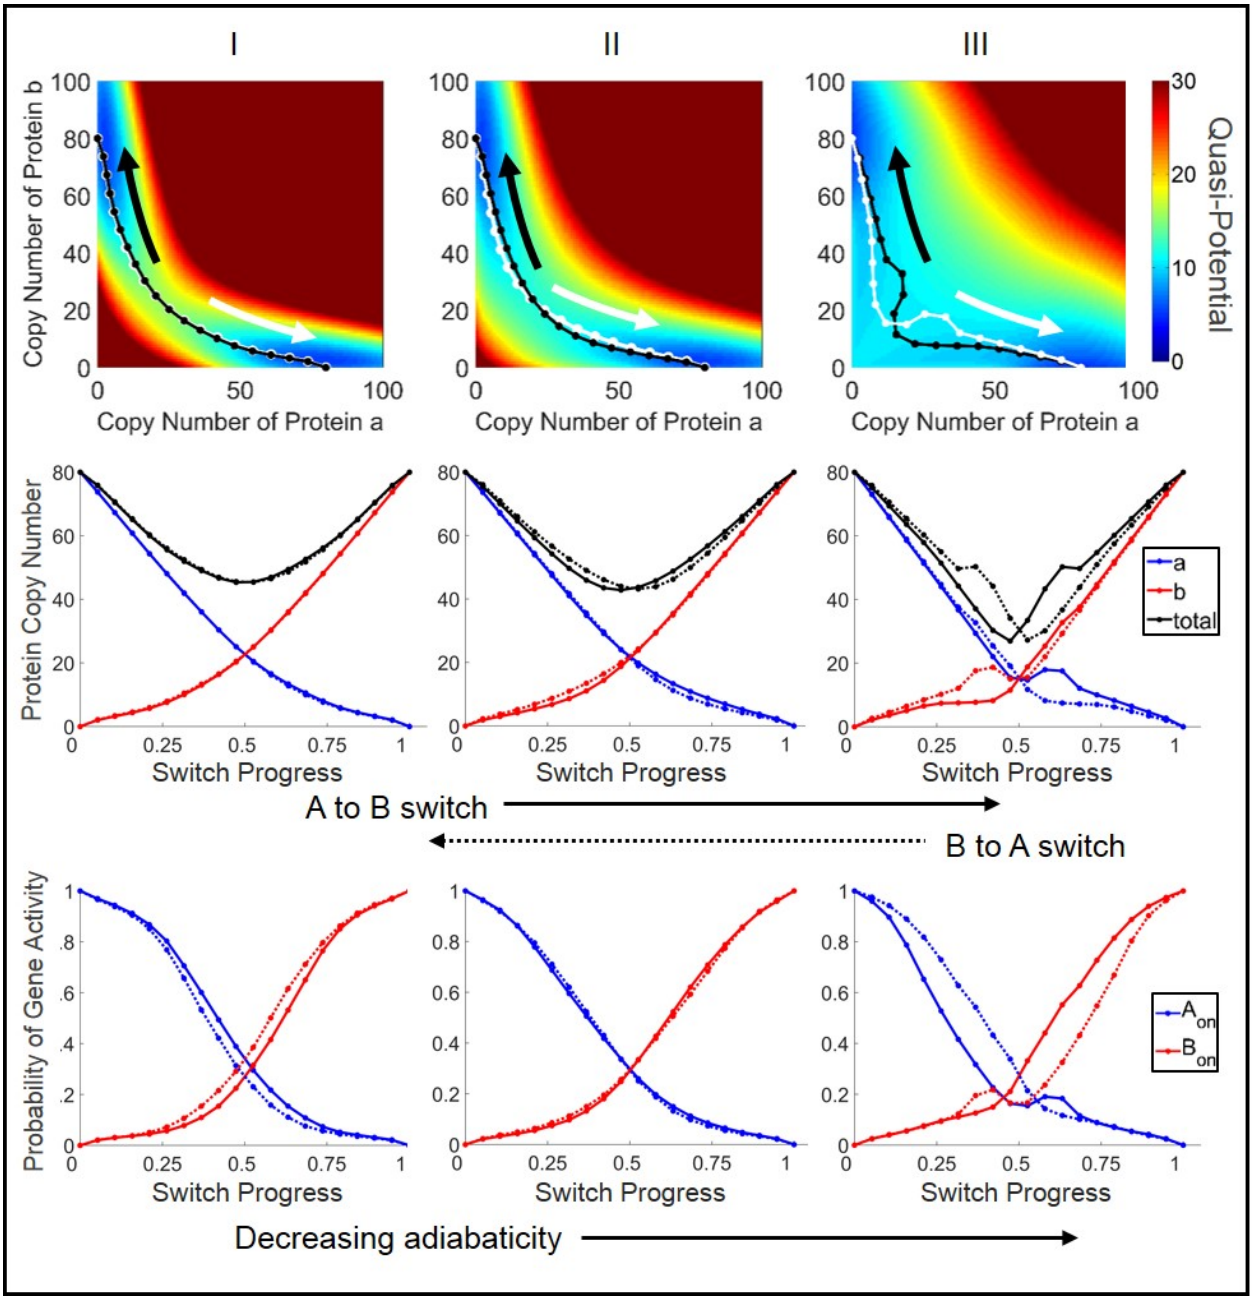

Figure 2.3: Quasi-potential surfaces and predicted transition paths for the base general toggle switch network. Left to right: network rate parameters with decreasing adiabaticity (slower DNA-binding kinetics) (see Table 2.1 for parameters) (Top). Transition paths superimposed on the 2D projection of the quasi-potential surface. The black and white paths illustrate the forward ($S_A \rightarrow S_B$) and backward ($S_B \rightarrow S_A$) transitions, respectively. (Middle) Transition paths plotted as protein copy numbers vs. switch progress (protein a (blue), protein b (red), and total protein number ($a+b$, black). Switch progress is defined as the normalized distance along the transition path (see Methods). Forward switching ($S_A \rightarrow S_B$) is displayed as solid lines, backwards ($S_B \rightarrow S_A$) as dotted lines. (Bottom) Transition paths plotted as gene activities vs. switch progress. The activity of gene A (A_{on}, blue curve), is the probability of the regulatory site of gene A being unbound, which renders it active. (Similar for B_{on}, red curve).

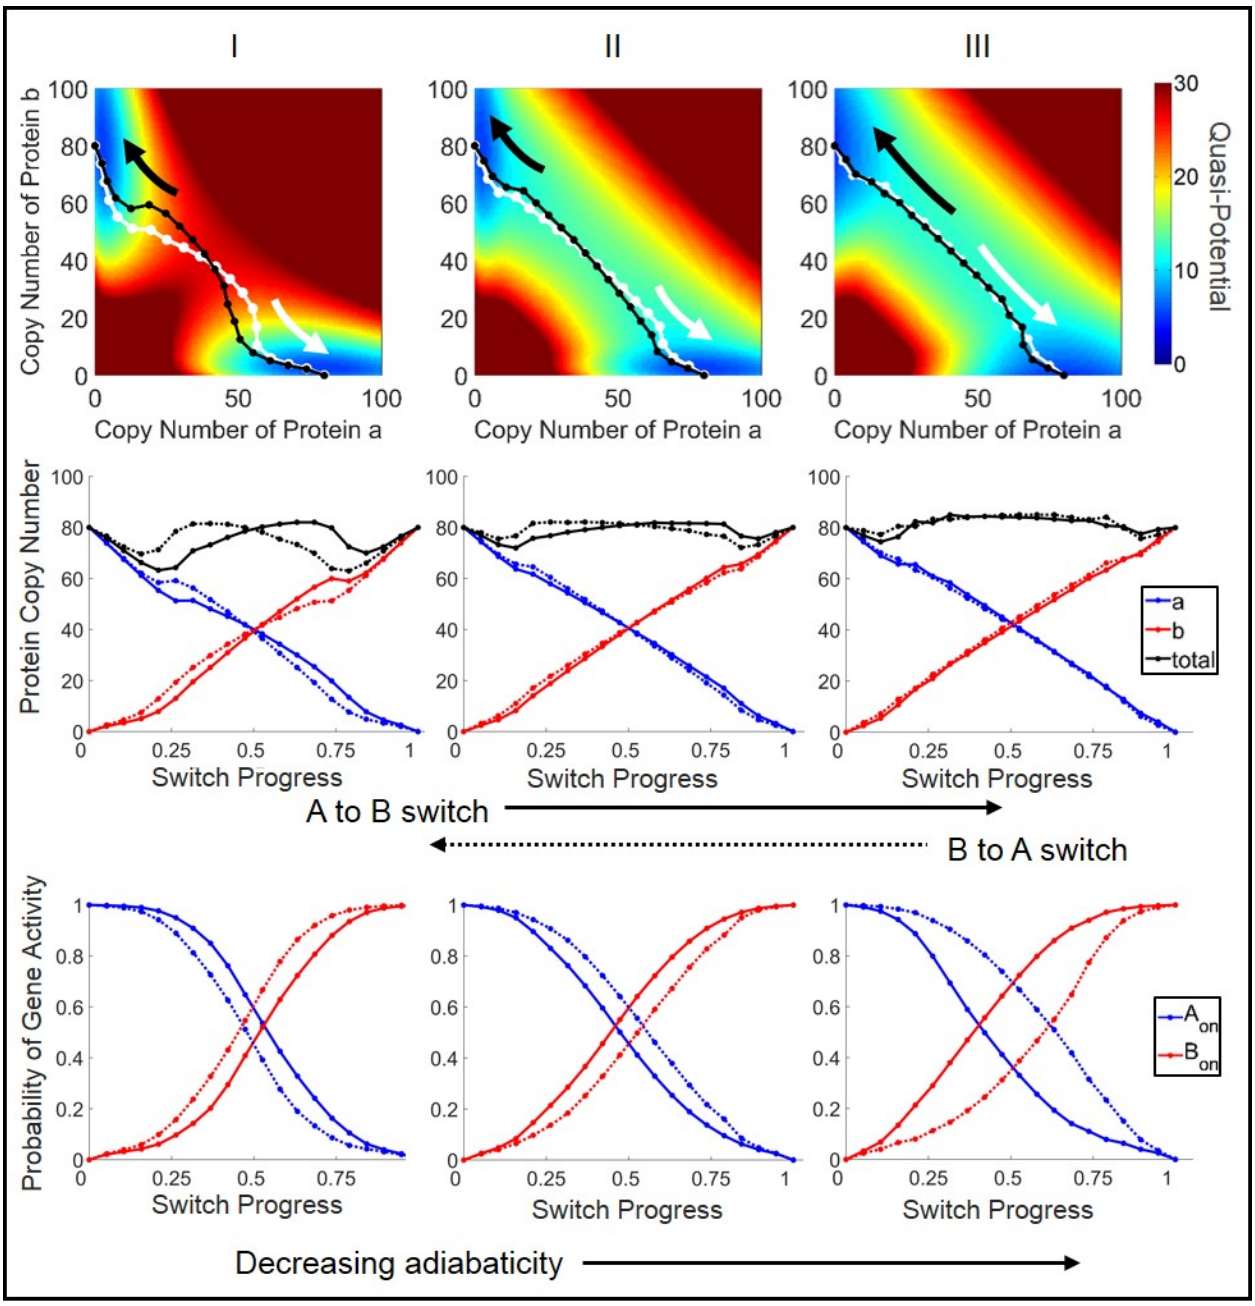

Figure 2.4: Quasi-potential surfaces and predicted transition paths for the exclusive toggle switch network, with competitive binding of repressors to regulatory sites. All definitions same as Fig. 2.3 Left to right: network rate parameters with decreasing adiabaticity (slower DNA-binding kinetics) (see 2.1 for parameters) (Top). Transition paths superimposed on the 2D projection of the quasi-potential surface. (Middle) Transition paths plotted as protein copy numbers vs. switch progress (Bottom) Transition paths plotted as gene activities vs. switch progress.

2.5 Sampling algorithm finds transition path for the toggle switch

Transition paths of noise-induced switching between gene expression states were predicted by performing WE-string sampling of stochastic simulations. Transition paths were predicted for both forward ($S_A \rightarrow S_B$) and backward ($S_B \rightarrow S_A$) switching events, and were compared to the numerically calculated quasi-potential landscapes for the three parameter sets. Although the WE-string simulations require no foreknowledge of the underlying potential in the transition region, the paths converge to the relatively low-potential (i.e., high probability density) transition tube connecting the two states (Figs. 2.3 and 2.4). In addition to mapping out the region of state-space traversed by successful transitions, committor analysis (Fig. 6.9 and 6.10, and Methods) indicates that, for most studied parameter sets, the string also accurately reflects the true reaction coordinate for the transition. Progress along the string corresponds to the most probable step-by-step changes undergone as the system progresses from one basin to another. However, slight discrepancies exist for the general switch in the non-adiabatic regime, likely due to a known limitation of string-based sampling methods in cases of wide transition tubes[41]. Predicted transition paths were robust to differences in simulation initial conditions (guess paths) (Fig. 6.6).

2.6 Transition paths reveal additional dynamic features and non-equilibrium phenomena

The predicted transition paths also reveal that the system dynamics is not governed solely by the topography of the potential surface. The paths reveal additional features, including pathway oscillations and non-overlapping forward and reverse paths, in agreement with

previous studies[158, 108]. These features can be seen by comparison to calculated Minimum Energy Paths (Figs. 6.7 and 6.8), which reflect only gradient dynamics over the projected quasi-potential surface. The degree to which the forward and reverse paths diverge depends both on the parameters and on the choice of sub-space on which the system dynamics is projected. For example, when projected onto the protein copy number sub-space, the paths show more divergence in the non-adiabatic regime than in the adiabatic regime for the general switch (Fig. 2.3) while the opposite is true for the exclusive switch (Fig. 2.4). For the exclusive switch in the non-adiabatic regime (Fig 2.4, parameter set III), there is clear divergence of the forward and reverse transition paths in the sub-space of DNA-occupancy states ("gene states"), yet the transition paths are nearly superimposed in the protein sub-space (Fig 2.4.).

Protein number oscillations are observed in the adiabatic regime for the exclusive switch (Fig. 2.4, parameter set I). This motion is not echoed in the DNA binding occupancies, suggesting that it results solely from the birth-death reactions for protein expression. In the non-adiabatic regime for the general switch (Fig. 2.3, parameter set III), oscillatory motion in the proteins is echoed in the DNA-binding occupancies ("gene activities"), suggesting that it could result from previously described "eddy currents" in gene networks, where slow transcription factor binding/unbinding events drive cyclic dynamics of protein expression in the non-adiabatic limit[130, 157, 168]. However, as discussed above, the detailed features of the string for the non-adiabatic general switch may not reflect dynamics along the true reaction coordinate.

2.6.1 Influence of DNA-binding kinetics on switching mechanism

The predicted transition paths for the toggle switch reveal that the switching mechanism is altered by the DNA-binding kinetic parameters. This is seen by considering the switch progress separately in terms of either protein copy numbers or gene activities: the switch progress

in these sub-spaces is not synchronized, and the degree to which they are asynchronous depends on the DNA-binding kinetics. In this symmetric system, when the system reaches the separatrix at ($a = b$), the transition can be considered to be half-complete. This crossing corresponds closely to the half-distance point along the transition path (Switch Progress = 0.5) (Figs. 2.3,2.4). Under a quasi-equilibrium assumption, the gene activities depend directly on the levels of protein expression, according to $A_{on} = 1/(1 + b^2/X_{eq})$ and $B_{on} = 1/(1 + a^2/X_{eq})$. As such, when protein expression from the two competing genes reaches equivalence, the genes would then also have equal probabilities of being in the active (unbound) states, giving ($A_{on} = B_{on}$). However, the transition paths reveal that this crossing is not necessarily reached simultaneously in both sub-spaces. Similar trends are seen in the parameter dependence for the different switch variants: in the adiabatic regime, the gene activities reach ($A_{on} = B_{on}$) after the protein numbers reach ($a = b$) (Figs. 2.3 and 2.4, left column), while in the non-adiabatic regime, the gene activities equalize before the proteins (Figs. 2.3 and 2.4, right column). In other words, when DNA-binding kinetics is slow, stochastic binding and unbinding events play a larger role in driving the switch towards completion, and protein expression follows. Conversely, when DNA-binding kinetics is fast, birth-death fluctuations in protein numbers drive the switch. This trend was preserved for additional studied parameter sets (Figs. 6.11 and 6.12), and also appeared in averaging of switching trajectories harvested from brute-force simulations (Figs. 6.1 and 6.2).

The simulations predict that the exclusive and general switch variants (see Methods) progress through different transition states. Our results are in agreement with the previous finding[108] that the general switch transitions through a state in which both genes are repressed, which is also reflected in the decrease in overall expression of proteins (Fig. 2.3, "total" expression curves). For the exclusive switch variant, the transition progresses through a region where each gene has a nearly equal probability of being "on" or "off". Additionally, the total protein number when the switch is half-complete (Switch Progress = 0.5) remains high after an initial decrease (Fig. 2.4, "total"). These transition states clearly reflect the difference in

the DNA occupancy states available to each toggle switch variant. In the general switch, all combinations of bound/unbound states at the regulatory sites of the two genes regulatory sites are allowed (i.e., $A_{on}/B_{on},\ A_{on}/B_{on},\ A_{off}/B_{on},\ A_{off}/B_{off}$), while in the exclusive switch, the A_{off}/B_{off} state is not accessible due to competitive binding in the case of overlapping regulatory sequences. This scenario is the limiting case of a reduced binding rate of a repressor if the competing repressor is already bound to the DNA.

2.6.2 Advantage of transition path simulation method for multi-dimensional networks with rare switching

We applied the sampling algorithm to the extended network with explicit dimerization (Methods). The parameters of the extended network were chosen such that the base and extended networks give identical total protein expression at steady-state, in terms of the variables na and nb (Methods and SI). The predicted transition paths showed qualitatively similar switching mechanisms for the base and extended networks. In particular, the trend of asynchronous switching in protein numbers and gene activities was preserved in the network with explicit dimerization, with gene activities lagging behind protein numbers in the adiabatic regime (Fig. 2.5). For this parameter set, it was not possible to access the full quasi-potential surface: because the size of the state-space increases exponentially with dimensionality (i.e., number of species), the addition of two more molecular species (the dimers a_2 and b_2) renders the transition rate matrix too large for straightforward numerical calculations (see Methods). This curse-of-dimensionality does not limit Gillespie simulations; however, brute-force simulation of 10^9 timesteps failed to capture a single switching event (Fig. 2.5A) due to the low probability of switching ($k_{AB} = 3.3 \times 10^{-10}$, Table 2.1). Our results thus demonstrate that the WE-string sampling method is particularly advantageous for studying transition paths for networks with several or more molecular species and long waiting times between events of interest.

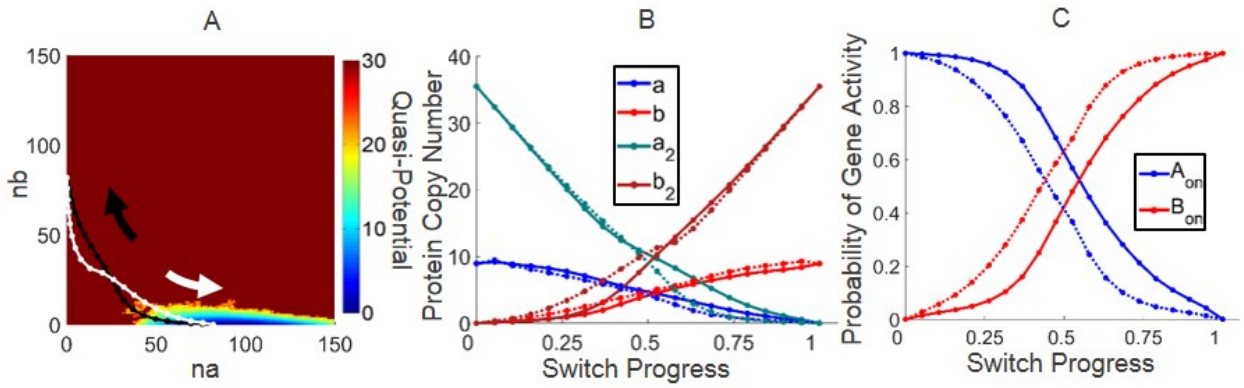

Figure 2.5: Transition paths predicted for the exclusive switch with explicit dimerization, parameter set I (see Table 2.1), and incomplete quasi-potential surface. A) Predicted transition paths for the $S_A \rightarrow S_B$ (black) and $S_B \rightarrow S_A$ (white) transitions. The 2D projection shows the quasi-potential surface obtained from a brute force Gillespie simulation of 10^9 timesteps initialized in SA. The long simulation fails to produce a single switching event, giving rise to a one-sided, incomplete quasi-potential surface. Despite the rarity of switching for this network ($k_{AB} = 3.3 \times 10^{-10}$, Table 2.1), the simulation method remains capable of predicting the most probable transition path. B) Transition paths plotted as protein copy numbers vs. switch progress C) Transition paths plotted as gene activities vs. switch progress.

2.7 Discussion

In this work, we applied recently-developed stochastic simulation methods to the study of gene-regulation dynamics, in order to predict switching mechanisms in a ubiquitous bistable gene network. Our simulations recapitulate several general aspects of gene network switching, including the increased stability of cell states to noise-induced switching with fast DNA-binding/unbinding kinetics, and evidence of non-equilibrium phenomena in switching dynamics, in agreement with previous studies[167, 158, 45, 157, 108, 133]. Our simulations also uncovered new insights into stochastic phenotype switching. In particular, we found that switching progresses asynchronously in the separate sub-spaces of protein copy numbers or DNA-occupancies (i.e., gene activities). We observed a trend whereby the gene activities lagged behind the protein numbers during the switching transition in the fast-DNA-binding (adiabatic) kinetic regime, and switched ahead of the proteins in the slow-DNA-binding (non-adiabatic) regime. Thus, our results reveal that the detailed mechanisms (and overall

rates, k_{AB}) of spontaneous switching depend on the network parameters. Our findings suggest that protein copy-number fluctuations play a more significant role than DNA-site-occupancy fluctuations in driving the switch toward completion in the adiabatic regime, whereas the converse is true in the non-adiabatic regime. Our study also demonstrates that rare-event sampling algorithms, in particular string-based methods[152], are well-suited to the discovery of transition paths that link metastable states in gene networks.

Stochastic state-switching in gene networks is thought to be a mechanism driving phenotypic heterogeneity in genetically-identical cell populations. This type of non-genetic heterogeneity induced by transient switching events appears to play a role in diverse biological contexts, including development and disease. In stem cell networks, probabilistic differentiation has been linked to stochastic switching of pluripotent cells between sub-states, distinguished by different levels of expression of key transcription factors[164, 141, 26, 72]. Stochastic state-switching is also thought to allow bacteria to switch into and out of antibiotic-resistant states[17], virus-infected cells to switch into and out of latency[162], and cancer cells to switch into and out of chemotherapy-resistant states[136]. Identification of the mechanisms by which these switching events occur can provide a route to discovery of novel strategies for cellular engineering (e.g., stem cell reprogramming) and drug treatment. However, mechanistic studies of switching can be experimentally challenging, particularly when only a fraction of cells in a population is poised to transition between states[23]. As such, quantitative models coupled with the computational approaches presented here could provide important insight into underlying mechanisms, e.g., by predicting patterns of gene expression associated with cells that are in transit between phenotypic states.

Chromatin remodeling contributes to gene expression noise by stochastically transitioning individual genes between on (active) and off (inactive) states[125, 104]. These transitions can occur on timescales that are of the same order or longer than mRNA and protein lifetimes[98]. Though the toggle switch model studied here is highly simplified, the slow-DNA-

binding (weakly- or non-adiabatic) kinetic regime broadly captures this complex epigenetic regulation. The relative instability of phenotypic states in this regime has been hypothesized to underlie developmental plasticity of pluripotent stem cells[167, 130]. Our results suggest that stochastic fluctuations at the level of the individual promoter sites, rather than changes in local regulatory protein concentration, play a critical role in driving the spontaneous, global switching events that poise pluripotent cells to choose probabilistically between alternative lineages.

The simulation method employed in this study simultaneously tackles two challenges encountered in modeling gene network dynamics: the rare-event problem (switching between gene states may occur rarely, giving rise to impossibly long simulation convergence times), and the curse of dimensionality (switching requires coordinated changes involving many species in the network). By applying the WE-string method to different toggle switch variants in multiple kinetic regimes, we demonstrate its potential flexibility for capturing stochastic dynamics in different gene network systems. We found that a major advantage of the method is its ability to simultaneously track dynamics of both protein products and DNA occupancy states (i.e. promoter states), which together control global gene network dynamics.

Previous studies of transition paths in gene networks treating DNA occupancies as being in quasi-equilibrium with protein numbers implicitly assumed that dynamics in the two network subspaces is identical[92, 127, 16, 170, 160].Our results show that, even in the adiabatic regime, where a separation-of-timescales assumption is justified, the switching dynamics of the gene activities and protein numbers have distinct features. Theoretical methods have been developed recently for studying noise-induced transitions in gene networks, which can account for non-adiabaticity (slow DNA-binding) by path integral approaches based on approximations to the CME[167, 168]. In this study, we explore an alternative approach in direct sampling, which circumvents the need for approximations to the CME, can be used in conjunction with available stochastic simulation software packages[44], and potentially scales

more easily to other types of biochemical networks with many species.

The application of rare-event sampling techniques to study switching dynamics in gene networks was pioneered previously with the FFS method[8]. Morelli, *et al.*[108] carried out a detailed study of the toggle switch by FFS; our network model is largely based on theirs, although we studied parameter regimes commensurate with those of Sasai, *et al.*[131], producing larger average protein numbers. The transition states we found for the general versus exclusive switch variants are in qualitative agreement with their study, although direct quantitative comparison of our results to theirs is not possible, because of the different parameter regimes, and because their study did not focus explicitly on the mechanistic dependence on adiabaticity.

In principle, FFS and string-based sampling methods can access the same information; the advantages and disadvantages of each method, making them potentially complementary, have been discussed previously[34]. FFS samples dynamics along a single order parameter, which must be defined *a priori*, and which tracks progress between initial and final states. Choice of this parameter is non-trivial, and a poor choice can lead to computational inefficiency, though statistical methods to optimize the order parameter have been developed[22]. Use of a single order parameter necessarily results in loss of information about dynamics along separate system coordinates; the progress coordinate λ is much more sensitive to protein copy numbers than to DNA-occupancy states. Adaptive string-based methods were developed to circumvent this problem and access collective dynamics in spaces of many order parameters[35]. Because progress through a transition is measured by successive string nodes, each of which represents a configuration in the full state-space, this approach offers a more intuitive description of dynamics in gene networks with many species. Moreover, the approach enabled us to directly isolate separate contributions of protein numbers and DNA-occupancy states to the switching mechanism, which in turn revealed the mechanistic dependence on adiabaticity.

However, string-methods also have potential limitations. Convergence of the string to the

most probable transition path may not occur in systems with highly complex quasi-potential landscapes with several or more local minima, so extension of the method to more complex gene networks remains to be explored. The transition path predicted by the converged string may not correspond exactly to the true chemical reaction coordinate for some networks (as defined by committor probabilities[94]), due to the difficulty of accurately partitioning wide transition tubes[41]. Our results show that, despite these potential limitations, the WE-string method was able to resolve switching dynamics over a wide range of parameters. Even in the non-adiabatic limit, the major mechanistic predictions (regarding the nature of the transition state and the asynchronous switch progression in the protein vs. gene-activity subspaces) were found to be robust by comparison to brute force trajectories. Our findings suggest that the method may be particularly useful for exploring dynamics of gene networks where little information (such as a suitable progress coordinate) is available *a priori*, and may prove particularly powerful in combination with other, complementary rare-event sampling approaches.

Chapter 3

Rare-Event Sampling of Epigenetic Landscapes and Phenotype Transitions

3.1 Introduction

In multicellular organisms, differentiation of pluripotent stem cells into tissue-specific cells was traditionally considered to be an irreversible process. The discovery of cell reprogramming revealed that a the identity of a cell is not irreversibly stable, but rather plastic and amenable to control by perturbation of gene regulatory interactions—for example, through over-expression of key transcription factors [148]. Cellular plasticity has also been observed in other contexts, where cells appear to spontaneously transition among phenotypically distinct states. For example, in embryonic stem cells, expression levels of key transcription factors show dynamic heterogeneity, which is thought to enable diversification of the population prior to lineage commitment [3, 37, 72, 141, 117]. This heterogeneity may result at least in part from stochastic

state-transitions between functionally distinct, metastable subpopulations [72, 140, 47, 64]. Stochastic state-transitions have also been proposed to play a role in cancer, by enabling cancer stem cells to arise *de novo* from non-stem subpopulations [54], or by enabling cells to reversibly transition to a drug-tolerant phenotype [136]. In microbial systems, stochastic phenotype switching has been identified as a survival mechanism for populations subjected to fluctuating environments [5, 17].

Mathematical modeling has provided a basis for understanding how gene regulatory mechanisms and network interactions control cellular identity, stability, and phenotype-transitions. These approaches yield a quantitative means of reinterpreting the long-standing conceptual framework known as Waddington's epigenetic landscape [155, 20, 158, 67]. In a mathematical framework, the "valleys" in the landscape that stabilize cell identities within distinct lineages correspond to attractor basins of a high-dimensional nonlinear dynamical system [68]. The nonlinearity results from positive feedback in transcriptional regulation and epigenetic barriers to chromatin remodeling, for example. These feedback mechanisms give rise to multiple, stable (or metastable) phenotype-states accessible to a given genome. Given the "bursty" nature of gene expression and ever-present molecular fluctuations in the cell [42, 78], an active area of research is in modeling the effects of so-called intrinsic noise on gene regulatory network (GRN) dynamics. These mathematical models support the idea that intrinsic noise can drive stochastic phenotype-transitions [16, 131, 46, 150, 49], which, though likely to be exceedingly rare in general cellular contexts, may explain the heterogeneity observed in embryonic stem cells where epigenetic barriers appear to be lowered [27].

Mathematical models of GRN dynamics that treat stochastic molecular processes are often formulated as probabilistic Master Equations, in which the system evolves probabilistically over a discrete state-space of molecular species and configurations according to a defined set of biochemical reaction rules. Another common framework is that of a coupled system of ODEs describing the expression levels of genes in the network, with the inclusion of

additive noise terms. The Master Equation framework is well-suited to studying how "local" stochastic molecular events (e.g., transcription factors interacting with DNA or chromatin state-transitions near promoters) impact "global" dynamics of phenotype stability and state-switching [130, 46, 167, 150, 49]. These molecular fluctuations affecting promoter activity have been shown to significantly impact the structure of epigenetic landscapes, motivating the use of Master Equation-based approaches. That is, the number and stability of phenotype-states accessible to a given GRN varies depending on the kinetic parameters governing these fluctuations [46, 150, 30]. Furthermore, ODE or "mean-field" models that average over these fluctuations can show qualitatively different landscape features [86, 133, 95].

Master Equation approaches face the well-known challenge of the "Curse-of-Dimensionality", as solving them requires enumeration of a state-space that grows exponentially with the number of molecular species in the network. For this reason, discrete stochastic models of GRNs are often studied by stochastic Monte Carlo simulation, via the Gillespie algorithm [51]. However, stochastic simulation can also be problematic: in systems with metastability, such as GRNs, stochastic simulation becomes highly inefficient. Transitions between metastable states are rare events (i.e., rare relative to the timescale of fluctuations within a metastable attractor basin), and thus difficult or impossible to observe. Often, these rare events are precisely the events of interest, such as in GRNs where infrequent state-transitions represent critical cell-fate transitions.

Rare-event sampling algorithms are designed to overcome these challenges, by redirecting computational resources towards events of interest, while maintaining statistical accuracy to global system dynamics [7, 172]. In this work, we present a rare-event simulation-based method for computing and analyzing epigenetic landscapes of stochastic GRN models. We combine rare-event methods with coarse-graining and analysis by Transition Path Theory–adopted from the field of Molecular Dynamics of protein folding [114]–and show that this unified framework provides an automated approach to map epigenetic landscapes and transition

dynamics in complex GRNs. The method quantifies the number of metastable phenotype-states accessible to a GRN, calculates the rates of transitioning among phenotypes, and computes the likely paths by which transitions among phenotypes occur. We apply the method to a model of pluripotency in mouse Embryonic Stem Cells. Our results reveal rare sub-populations and transitions in the network, demonstrate how global landscape structure depends on kinetic parameters, and reveal irreversibility in paths of differentiation and reprogramming. Our approach is not limited to gene regulatory networks; it is generalizable to other stochastic dynamics frameworks and is thus a potentially powerful tool for computing global dynamic landscapes in areas such as signal-transduction, population dynamics, and evolutionary dynamics.

3.2 Methods

A graphical overview of the computational pipeline presented in this paper can be found in Fig 3.1.

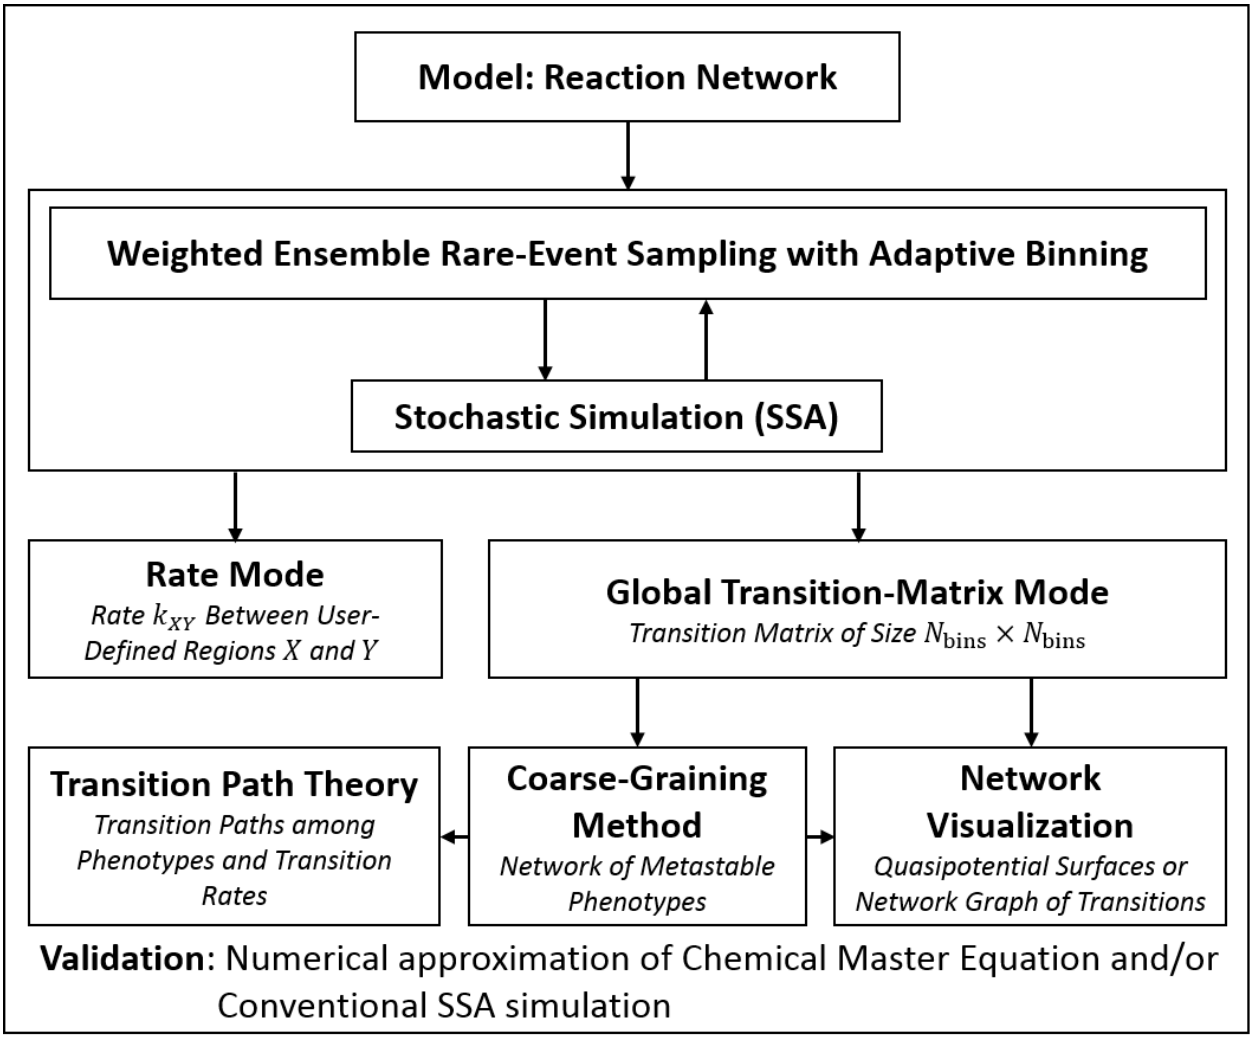

Figure 3.1: **Computational Pipeline for Rare-Event Sampling of Epigenetic Landscapes and Phenotype Transitions.** The input to the computational pipeline is a reaction network model of gene regulatory network dynamics. Stochastic simulations are performed using SSA [51] and Weighted Ensemble rare-event sampling [70]. The WE method can be run in two modes: *Rate Mode* computes the rate of transitioning between two user-defined regions of interest with high accuracy. *Transition-Matrix Mode* computes the pairwise transition probabilities among N_{bins} adaptively defined sampling bins that span the system state-space. Further visualization and analysis of the transition-matrix can be performed, including automatic designation of metastable phenotypes via the coarse-graining framework [122] and identification of likely transition paths [114].

3.2.1 Gene Regulatory Network Models

We demonstrate the rare-event sampling method for two representative GRN models. A small, two-gene network serves as a model system to validate the simulations. We then apply the method to a more complex model of pluripotency in mouse Embryonic Stem Cells (mESCs).

Exclusive Mutual Inhibition, Self-Activation Model

The Exclusive Mutual Inhibition, Self-Activation (ExMISA) model is a two-gene network representing an archetypal motif at cell-fate branch points [66, 53]. Each gene, denoted generically as A or B, encodes a transcription factor that activates its own transcription and represses transcription of the other gene. We adopt previous conventions [75, 131, 46] for stochastic GRN dynamic models. The full list of biochemical reactions and parameters can be found in the Supplement, Chapter 6.3 and Chapter 6.2. The model encompasses stochastic birth/death processes for transcription factor production and degradation, and stochastic binding and unbinding of transcription factors to DNA regulatory/promoter regions; the binding-states of these regions governs the production rate. Each transcription factor is assumed to bind to DNA as a homodimer, giving cooperative regulation. In the "exclusive" network variant, transcription factors compete for binding sites on DNA (only one transcription factor dimer can be bound to a gene's promoter at a time). The discrete state-vector, which completely describes the state of the system, is given by $\mathbf{x} = [A_{ij}, B_{ij}, n_a, n_b]$. A_{ij} and B_{ij} represent the three possible promoter binding-states for each gene (i.e., $A/B_{00}, A/B_{10}, A/B_{01}$ denote unbound, activator-bound, or repressor-bound states). The copy-numbers of expressed protein transcription factors are denoted by n_a and n_b for products of gene A and B, respectively, and may in principle take any nonnegative integer value. All processes related to transcription, translation, and assembly are subsumed into a single protein birth reaction.

For genes in state A/B_{ij}, this production occurs with rate constant g_{ij}. The production rate is high when the promoter is bound only by the activator (its own product). Otherwise, if unbound or repressor-bound, a low "basal" rate of expression is assumed, i.e. $g_{00} = g_{01} < g_{10}$. Degradation of protein products occurs with rate k, and stochastic binding/unbinding of transcription factors to DNA occur with h and f, respectively. The model is symmetric, with equivalent parameters for the two genes.

Pluripotency Network Model

The pluripotency network model of mESCs was developed by Zhang and Wolynes [167] on the basis of experimental literature and previous models. The 8-gene network shares the same stochastic reaction framework as the ExMISA model. The genes (NANOG, OCT4, SOX2, GCNF, KLF4, PBX1, GATA6, and CDX2) suppress and activate each other through homo- and heterodimers of their encoded transcription factors (OCT4 and SOX2 form a heterodimer; all other regulatory interactions occur via homodimers). Binding of transcription factors to promoters is not exclusive. The model has five kinetic parameters: g_{on}, g_{off}, h, f, and k, corresponding to the rate of gene expression in the activated state, the rate of gene expression in the un-activated state, binding of transcription factors to DNA, unbinding of transcription factors from DNA, and transcription factor degradation (or exit from the nucleus). Genes are expressed at the basal rate g_{off} except when bound by at least one activator and no repressor, in which case they are expressed with rate g_{on}. The exception to this logic rule is NANOG, which must be bound by the the KLF4 and PBX1 transcription factor homodimers and the heterodimer OCT4-SOX2 to be activated. Overall, these interactions lead to a total of 396 biochemical reactions, with a total of 88 "species" (counting 80 distinct gene promoter configurations and 8 protein species). The complete logic rules and list of reaction rate parameters can be found in the Supplement (Chapter 6.3, Tab. 6.3, and Tab. 6.4).

3.2.2 Theoretical Background: the Chemical Master Equation and Stochastic Transition-Matrix

The mathematical framework of the network models is the discrete Chemical Master Equation (CME) [51], which gives the time-evolution of the probability to observe the system in a given state. In vector-matrix form, the CME can be written

$$\frac{d\mathbf{p}(\mathbf{x},t)}{dt} = \mathbf{K}\mathbf{p}(\mathbf{x},t) \tag{3.1}$$

where $\mathbf{p}(\mathbf{x},t)$ is the probability over the system state-space ($\mathbf{x}$) at time t, and $\mathbf{K}$ is the reaction rate-matrix containing stochastic reaction propensities (diagonal elements $k_{jj} = -\sum_i k_{ij}$, i.e., columns sum to 0). Equation 3.1 assumes a well-mixed system of reacting species, and assumes that the technically infinite state-space described by $\mathbf{x}$ (containing molecular species numbers/configurations) may be limited to some finite number of "reachable" states, (i.e., with non-negligible probability) for an enumeration of N states of the system, $\mathbf{K} \in \mathbb{R}^{N\times N}$. The steady-state probability $\pi(\mathbf{x}) \equiv \mathbf{p}(\mathbf{x}, t \to \infty)$ over N states satisfies

$$\mathbf{K}\pi(\mathbf{x}) = \mathbf{0}. \tag{3.2}$$

Thus, $\pi(\mathbf{x})$ can be obtained from $\mathbf{K}$ as the normalized right-eigenvector corresponding to the zero-eigenvalue.

It is sometimes desirable to work with the time-dependent stochastic transition-matrix $\mathbf{T}(\tau)$ rather than the time-independent stochastic rate matrix $\mathbf{K}$ [122]. For example, $\mathbf{T}(\tau)$ may be more amenable to estimation by sampling (as we demonstrate in this work for the pluripotency

network, for which $\mathbf{K}$ is impractical to enumerate). For a CME with rate matrix $\mathbf{K}$, $\mathbf{T}(\tau)$ is given by

$$\mathbf{T}(\tau) = \exp(\tau \mathbf{K}^{\mathrm{T}}) \tag{3.3}$$

where exp denotes the matrix exponential. $\mathbf{T}(\tau) \in \mathbb{R}^{N \times N}_{0 \leq x \leq 1}$ then gives the conditional probability for the system to transition between each pair of states within a lagtime τ. That is, the elements t_{ij} give the probability that the system, if found in state i, will then be found in state j at a time τ later, and rows sum to 1. Using $\mathbf{T}(\tau)$, the evolution of probability over discrete intervals of the lagtime τ is given by the Chapman-Kolmogorov equation:

$$\mathbf{p}^T(\mathbf{x}, t + k\tau) = \mathbf{p}^{\mathrm{T}}(\mathbf{x}, t)\mathbf{T}^k(\tau). \tag{3.4}$$

Eigenvectors corresponding to dominant eigenvalues of the stochastic transition-matrix are associated with slow system processes. By Perron-Frobenius, for an irreducible stochastic matrix $\mathbf{T}(\tau)$ with eigenvalues λ_i, there exists $\lambda_1 = 1$, and all other eigenvalues satisfy $|\lambda_i| < 1$. Analogous to Equation (3.2) for $\mathbf{K}$, the steady-state probability can be obtained directly from $\mathbf{T}(\tau)$ according to $\pi^T(\mathbf{x}) = \pi^T(\mathbf{x})\mathbf{T}(\tau)$, i.e., as the normalized left-eigenvector corresponding to λ_1. Eigenvalues λ_i are related to global system timescales t_i by

$$t_i = -\frac{\tau}{\ln|\lambda_i(\tau)|}, \tag{3.5}$$

(with t_1 giving the infinite-time, stationary result) [122]. Additionally, the Mean First Passage Time ($\mathrm{MFPT}_{X,Y}$) where X and Y are individual states can be computed using the matrix

elements $T_{i,j}$ by [132, 60]:

$$\mathrm{MFPT}_{X,Y} = \tau \times \begin{cases} 0 & X = Y \\ 1 + \sum_z T_{X,Z} \mathrm{MFPT}_{Z,Y} & X \neq Y \end{cases}. \tag{3.6}$$

3.2.3 Weighted Ensemble Stochastic Simulation

Stochastic reaction kinetics can be simulated by the Stochastic Simulation Algorithm (SSA) [51], which produces numerically exact realizations of the CME (Eq 3.1). Simulation circumvents the need for enumerating the exceedingly large system state-spaces typical of gene network models, but suffers from inefficiency due to rare events. The Weighted Ensemble (WE) rare-event sampling algorithm [70] redistributes computational resources from high-probability regions of state-space to low-probability regions, which tend to be under-sampled in conventional simulation. The method thereby reduces computational effort in sampling rare transitions and improves accuracy of estimating probability density in, e.g., barrier-regions or tails of distributions. The method can be applied to any stochastic dynamics framework; in recent years, it has been widely applied to atom-scale Molecular Dynamics. Details of the methodology are discussed in a recent review [172] and references therein.

Briefly, the algorithm works as follows: state-space is divided up into bins that span transitions of interest. The number of bins, N_{bins}, is typically $\mathcal{O}(100)$, and a variety of binning procedures can be used (we use an adaptive procedure described below). Initially, a single simulation trajectory, or "replica", is assigned a weight of 1 and allowed to freely move within and between bins for a user-defined lagtime τ_{WE}. After each iteration of τ_{WE}, a splitting and culling procedure divides and/or combines replicas and their associated weights in such

a way as to reach and maintain an equal target number of weighted replicas, M_{targ}, in each bin. Over the course of the simulation, the combined weights of the replicas in a bin (averaged over successive iterations) will evolve toward the probability of the system to reside in that bin. By maintaining the same number of replicas in each bin (M_{targ}), with weights proportional to probability, the algorithm devotes comparable computational time to low- and high-probability regions. Effectively, the algorithm computes long-time processes on the basis of many short-time simulated trajectories.

Adaptive Binning Procedure

As with other enhanced sampling methods, the WE algorithm requires dividing of state-space into defined sampling regions or "bins". For high-dimensional systems, discretization poses a challenge because, for an N-dimensional, evenly spaced grid, the number of required sampling bins increases exponentially with the number of degrees of freedom. To address this challenge, a variety of Voronoi-polyhedra-based procedures have been developed [35, 33, 166]. These methods balance the need to focus simulation toward regions with non-negligible probability, while still enabling capture of rare transitions of interest. In addition to efficiently discretizing high-dimensional spaces, the methods have the benefit of requiring little to no *a priori* knowledge of system dynamics (e.g., of the locations of regions of interest, or of appropriate progress coordinates for transitions). We utilize an adaptive binning procedure from ref. [166]. Each bin (of user-defined number N_{bins}) is a Voronoi polyhedron with a generating node; the bin is defined as the region of state-space encompassing all points closer to the generating node than to nodes of any other region. After each lagtime τ_{WE}, new Voronoi regions are generated by successively selecting N_{bins} node-positions from the current replica positions in a way that maximizes the Euclidean distance between them. By this procedure, over the course of the simulation, bins spread to encompass all areas of state-space reached by any simulated trajectory. After sufficient iterations, the bin positions stop spreading to new areas

but continue to fluctuate. The procedure is shown by representative simulations in Fig. 6.13.

Computation of Transition Rates

One important output of WE sampling is the quantitative rate of transitions between regions of interest, which may be difficult or impossible to estimate from conventional simulation. WE sampling may be run in different modes, depending on whether the sought-after information concerns a specific transition of interest, or a more global picture of system dynamics, i.e., encompassing approximate rates of transitions among many system states. We term the two modes "rate" mode and "global transition-matrix" mode. The former can deliver a more accurate estimate for a particular state-transition, while the latter can yield a more comprehensive, but approximate, measure of global system dynamics.

In rate mode, the user specifies two regions of interest, X and Y, The flux of probability into/out of regions of interest can be estimated by recording the amount of weight transferred at the end of each simulation iteration. The mean first passage time of transitions from X to Y ($\mathrm{MFPT}_{X,Y}$) is given in general by the inverse of probability flux from X to Y. In practice, we apply a"labeling" scheme [146, 36], where each replica is labeled as belonging to either set $\mathcal{S}_X$ or $\mathcal{S}_Y$ according to its history, i.e., whether it most recently visited region X or Y, respectively. The summed weight of all replicas in $\mathcal{S}_X$ is given by $P_{\mathcal{S}_X}$, and $P_{\mathcal{S}_X} + P_{\mathcal{S}_Y} = 1$ satisfies probability conservation. Then,

$$\mathrm{MFPT}_{X,Y} = \frac{\overline{P}^{\mathrm{SS}}_{\mathcal{S}_X}}{\overline{\Phi}^{\mathrm{SS}}(Y|\mathcal{S}_X)} \tag{3.7}$$

where $\overline{\Phi}^{\mathrm{SS}}(Y|\mathcal{S}_X)$ is the average probability flux from $\mathcal{S}_X$ into Y at steady-state, which is measured by the weight of $\mathcal{S}_X$-labeled replicas entering Y during the simulation after

convergence to steady-state. The labeling scheme enables accurate estimates, including for non-Markovian transitions. For Markovian transitions well-described by a single rate-constant, $k_{X,Y} = 1/\mathrm{MFPT}_{X,Y}$.

Computation of Network Transition-Matrix

Running WE in transition-matrix mode enables visualization and analysis of global system dynamics on the basis of a single simulation, and requires no designation of regions of interest. In this mode, the previously-converged Voronoi bins are fixed, and simulations are used to estimate a coarse-grained stochastic transition-matrix $\widetilde{\mathbf{T}}(\tau)$ of size $N_{\mathrm{bins}} \times N_{\mathrm{bins}}$. The coarse-grained $\widetilde{\mathbf{T}}(\tau)$ approximates the true dynamics over the full state-space, as given by $\mathbf{T}(\tau)$. Thus, the procedure enables estimation of the global transition-matrix (and subsequent analysis) in systems where enumeration of states is not feasible. To estimate $\widetilde{\mathbf{T}}(\tau)$, the weight transferred between bins is recorded at each iteration, and the elements of the transition-matrix are estimated according to [146]:

$$\widetilde{T}_{i,j} = \frac{\langle w_{i,j} \rangle_2}{\langle w_i \rangle} \tag{3.8}$$

where $\langle w_{i,j} \rangle_2$ is the average weight transferred from bin i to bin j over the iteration time τ_{WE} (counting only after at least 2 transitions, and averaging over multiple iterations) and $\langle w_i \rangle$ is the average population (summed weight) in bin i. By construction, this is a row-stochastic transition-matrix with state-space "resolution" determined by N_{bins} (each state in the full state-space sampled by the simulation is assigned to its nearest neighboring Voronoi node). The lagtime τ of the transition-matrix corresponds to the sampled WE-time τ_{WE}. However, use of $\widetilde{\mathbf{T}}(\tau)$ to compute system dynamics imposes a Markovian approximation, by which equilibration of replicas within bins is assumed to be rapid on the timescale of τ, and hops

between states (i.e. bins) are memoryless. As such, while this mode of simulation has the advantage of acquiring a holistic view of global system dynamics, it has the disadvantage of introducing a Markovian approximation.

3.2.4 Coarse-Graining Procedure to Classify Phenotype-States

While the sampled $N_{\text{bins}} \times N_{\text{bins}}$ transition-matrix provides a global approximation of the epigenetic landscape and state-transitions, we apply a method to further coarse-grain dynamics, known as the Markov State Model framework [122, 114, 30]. This automated procedure produces a highly simplified representation of global dynamics in terms of a few (generally < 10) clustered sets and the transitions among them. Such highly-reduced models can be beneficial in terms of human intuition of system dynamics, comparison to experiments, and–in this application–automated designation of dynamic phenotype-states. The method utilizes the concept of metastability, i.e., system states that experience relatively fast transitions among them are clustered together into the same coarse-grained set. Collectively, the coarse sets experience relatively rare inter-cluster transitions and frequent intra-cluster transitions. We employ the metastability concept as a definition of cell phenotype, reasoning that a phenotype should be a relatively stable attribute of a cell, and stochastic inter-phenotype transitions should be relatively rare. In practice, we employ the Markov State Model framework to further reduce the sampled row-stochastic transition-matrix $\widetilde{\mathbf{T}}(\tau)$ from size $N_{\text{bins}} \times N_{\text{bins}}$ down to $C \times C$, where C is the number of coarse-grained clusters. As the Markov State Model (MSM) is itself a stochastic transition-matrix on a coarse-grained space, it implies a more severe Markovian approximation. It provides a way to describe global system dynamics in a highly simplified way while maintaining high accuracy to the slowest system dynamics as sampled by $\widetilde{\mathbf{T}}(\tau)$. In previous work, we demonstrated the application of this coarse-graining approach to automatically designate phenotypes in small gene networks [30]; here, we extend the applicability of the coarse-graining to large, complex networks by combining it with

rare-event sampling.

The coarse-graining procedure is a spectral clustering method based on the Perron Cluster Cluster Analysis (PCCA+) algorithm [128], which optimizes the (nearly)-block-diagonal structure of $\widetilde{\mathbf{T}}(\tau)$ for systems with metastability. The signature of such metastability is a separation-of-timescales for intra- and inter-basin dynamics, which may be seen as gaps in the eigenvalue spectrum [122]. As noted above, $\mathbf{T}(\tau)$ (or its sampled counterpart, $\widetilde{\mathbf{T}}(\tau)$) has $\lambda_1 = 1$, corresponding to the infinite time-limit. If a set of m dominant eigenvalues exists, such that for decreasing eigenvalues $\lambda_i \lessapprox 1$, $i \in \{2, ..., m\}$, and a gap is present, $\lambda_j << \lambda_m$ for $j > m$, this indicates the presence of m slow-timescale processes in the system, and further indicates that $\widetilde{\mathbf{T}}(\tau)$ may be re-ordered to give m nearly-uncoupled blocks. In practice, the algorithm attempts to find a coarse-graining onto C clusters, where C may be user-defined, or may be determined algorithmically, e.g., according to the spectral gap [128]. Here, we choose C clusters, where the last significant gap in the spectrum is seen between λ_C and λ_{C+1}. For the GRNs studied here, this corresponds to choosing C such that $\lambda_C/\lambda_{C+1} > 10$.

Transition Path Analysis

The coarse-grained model of system dynamics given by the MSM enables estimation of the ensemble of dominant transition paths among phenotypes, along with their relative probabilities. We adopt methods from Transition Path Theory according to Noe, *et al.* [114] (details therein). Briefly, $\widetilde{\mathbf{T}}(\tau)$ can be used to compute the effective flux of trajectories, along any edge in the coarse-grained network, contributing to transitions between states X and Y (where these designated states correspond to one or more coarse-grained phenotype-states produced by the MSM). A pathway decomposition algorithm on the matrix of effective fluxes for $X \to Y$ transitions then yields a set of dominant pathways and the relative contribution of each to the overall flux. Each state in the MSM is analogous to a cell phenotype, and transition path analysis is used to identify parallel phenotype transition paths and the relative

rates of transitioning between phenotypes.

3.2.5 Visualization of Epigenetic Landscapes

Both the sampled transition-matrix $\widetilde{\mathbf{T}}(\tau)$ and the coarse-grained MSM encode stationary and dynamic information about global dynamics–that is, they quantify the epigenetic landscape. For visualization, we use Gephi graph visualization software [18] using the Force Atlas algorithm. Every circle (or node) in the graph corresponds to a sampling bin or to a coarse-grained phenotype, and the area of a circle is proportional to its relative steady state probability according to $\ln(\gamma P_{SS})$, where P_{SS} is the steady state probability of the node and γ is a constant chosen to improve visibility of low probability regions of the landscape. Lines between circles (edges) correspond to transitions between sampling regions or coarse-grained phenotype. Their thickness and coloring correspond to their relative transition probability and source state, respectively.

3.2.6 Validation: Numerical Solution of the Chemical Master Equation

To validate the simulation method, we compare the simulated dynamics to the numerical solution to the CME. We choose the parameters of the ExMISA model in such a way as to restrict the effective state-space, so that a numerical solution of the CME is tractable. Building the reaction rate matrix $\mathbf{K} \in \mathbb{R}^{N \times N}$ requires enumeration of N system states. In general, if a system of S molecular species has a maximum copy number per species of n_{max}, then $N \approx n_{max}^{S}$. In the ExMISA model, the state-vector is given by $\mathbf{x} = [A_{ij}, B_{ij}, n_a, n_b]$. For enumeration, we neglect states with protein copy-numbers larger than a cutoff value which exceeds g_{10}/k (corresponding to the average number of transcription factors maintained

in the system from a gene while in its active state). For example, with model parameters $g_{10} = 18$ and $k = 1$, we truncate at $n_{a,max} = n_{b,max} = 41$ and assume that probability flux between states with n_a, $n_b \leq 41$ and states with n_a, $n_b > 41$ is assumed to be 0 (i.e., the boundaries of the state-space are reflective). Including the gene-binding states, this gives $N = 3 \times 3 \times 42 \times 42 = 15876$ states. This size is tractable for complete solution of the CME using matrix methods in MATLAB [1]. This truncation of the state-space introduces a small approximation error (see Fig. 6.14).

The pluripotency network has 8 genes with copy numbers of $\mathcal{O}(10^3)$ (determined by the parameters $g_{on}/k = 3900$). The number of distinct binding-promoter states for each gene are $16, 32, 8, 8, 2, 8, 4,$ and 2 for GATA6, NANOG, CDX2, OCT4, SOX2, KLF4, GCNF, PBX1, respectively (see Tab. 6.3). Together these combinations enumerate a state-space of $N > 10^{30} \approx 1000^8 \times 16 \times 32 \times 8 \times 8 \times 2 \times 8 \times 4 \times 2$. This size precludes solution of the CME, and we instead estimate the dynamics by WE sampling. Where possible, we validate the WE-sampling results by "conventional", i.e., by direct simulation using SSA.

Validation of Coarse-Grained Models

To check the validity of the coarse-grained MSM as a representation of the global dynamics, we use the Chapman-Kolmogorov test to compare the relaxation curves of the coarse-grained system to those found through direct SSA following Equation 3.4 [122]. If the coarse-graining is appropriate, the relaxation curves of the MSM probabilities will match the relaxation profile of long conventional (direct SSA) simulations initiated within each coarse-grained phenotype. Transition paths through the coarse-grained phenotype network are validated, where possible, against conventional SSA simulation.

3.2.7 Implementation and Software

Stochastic Gillespie (SSA) simulations were carried out using BioNetGen [44]. WE sampling was implemented with in-house software code written in MATLAB. Simulations were run on the high performance computing cluster (HPC) at the University of California, Irvine, and parallelization of BioNetGen SSA simulations was performed using the Sun Grid Engine scheduler. The coarse-graining procedure and transition path analysis was implemented in python scripts, adapted from MSMBuilder [55] and Pyemma [132], respectively. Transition-matrix and MSM visualization was carried out using Gephi software and the Force Atlas layout[18]. All simulation parameters can be found in the supplement Tab. 6.5. Pseudo-code for the adaptive binning procedure can be found in Chapter 6.5 and software can be found in `https://github.com/read-lab-uci/adaptive_we`.

3.3 Results

3.3.1 Rare States and Transitions in Gene Regulatory Networks are Accessible by Rare-Event Sampling

We first apply the computational pipeline to a small two-gene model (the exclusive Mutual Inhibition, Self-Activation model, ExMISA, see Methods), exhibiting an archetypal motif for cell fate-decisions[53, 66]. The model is tractable for computation of full, discrete stochastic dynamics to within a small approximation error using matrix methods. Thus, the model provides a numerical benchmark for assessing the accuracy of the simulation method, before extension to larger systems where solution of the Chemical Master Equation (CME) is intractable. For the chosen parameters, the ExMISA model shows four peaks in the steady-state probability distribution (projected onto protein copy numbers, n_a and n_b). Peaks

in probability correspond to basins in the so-called quasipotential landscape, defined by $U = -\ln(\pi(\mathbf{x}))$ (Fig 3.2). The four peaks/basins corresponds to four possible combinations of binarized A/B gene expression: hi/hi, hi/lo, lo/hi, and lo/lo. These four phenotype-states arise due to the combination of balanced repression and self-activation in the network, and the slow kinetic parameters (Supplementary Tab. 6.2) for transcription factor binding and unbinding to promoters that effect changes in individual gene-activity states between low and high expression rates [160, 30].

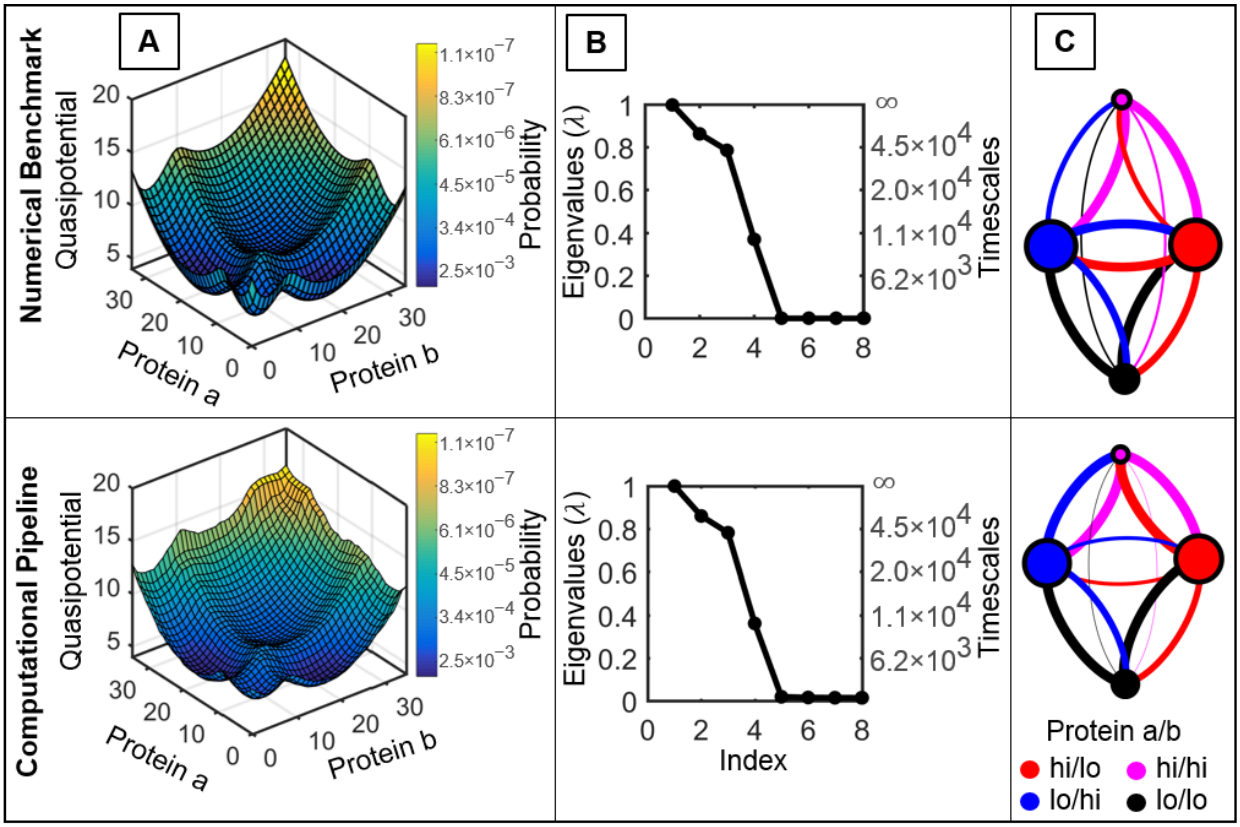

Figure 3.2: **Simulation Results Show Good Agreement with a Theoretical Benchmark for the 2-gene ExMISA (Mutual Inhibition, Self-Activation) Cell-Decision Circuit** The Chemical Master Equation for the 2-gene model, ExMISA, was solved numerically (see Methods) (top) and compared to simulation results from the computational pipeline presented in this paper (bottom). Shown for each are the Quasipotential Landscape (A), Eigenvalue Spectrum (B), and Markov State Model (C). (A) Quasipotential landscapes of the ExMISA network projected onto the two protein coordinates. Deep blue regions denote low potential (high probability) and yellow denote high potential (low probability). The four visible basins in both correspond to combinations of lo/hi expression for the two genes A and B. (For both rows, quasipotential surfaces estimated over discrete states/bins are smoothed for visualization). WE sampling captured both the basin structure and low probability edge and barrier regions. (B) Eigenvalue spectra and corresponding computed global transition timescales. Gaps in the eigenvalue spectrum indicate separation of timescales, i.e., the presence of metastability. C) Four-phenotype coarse-grained models automatically generated from the clustering algorithm (see Methods). Each colored circle represents a cell phenotype, sized proportionally to its probability. Edges are inter-phenotype transitions (colored by source-state, with width proportional to probability). The full CME and simulation pipeline identify similar metastable phenotype networks (see Fig. 6.23 for details).

The WE-based simulation method enabled estimation of global dynamics of the ExMISA

model. By redistributing computational resources from relatively high-probability to low-probability regions (see Methods), the WE method enabled uniform sampling of the quasipotential landscape, i.e., mapping basins (high-probability regions) along with high barriers (low probability regions) (Fig 3.2a). The simulation estimated individual steady-state bin-probabilities as low as 1.3×10^{-6} and showed good global agreement with the numerical CME benchmark (see Fig 3.2 and Supplement, Fig. 6.15).

In addition to sampling global dynamics, the WE method can be used to estimate rate constants for individual, rare transitions of interest. The Mean First Passage Time of the global network switch from the center of one polarized phenotype-state to another, i.e., $\mathrm{MFPT}_{X \to Y}$from protein a/b expression level hi/lo to lo/hi was estimated from WE to be 1.82×10^5 (see Tab. 6.7), in agreement with the CME result.

3.3.2 Phenotype Transitions can be Approximated by Markovian Jumps, Enabling Construction of Coarse-Grained Models

A network transition-matrix $\widetilde{\mathbf{T}}(\tau)$ over sampled bins ($N_{\mathrm{bins}} = 300$) was constructed from WE sampling for ExMISA and used for subsequent analysis of global system dynamics. By comparison, a full network transition-matrix $\mathbf{T}(\tau)$ over the enumerated system state-space was constructed from the CME ($N = 15876$, see Methods). The full, computed ($\mathbf{T}(\tau)$) and simulated ($\widetilde{\mathbf{T}}(\tau)$) transition-matrices showed qualitatively similar eigenvalue spectra with four dominant eigenvalues, indicating the presence of metastability (separation-of-timescales between intra-basin and inter-basin transitions) (Fig 3.2b). The slow system-timescales predicted by the full CME model corresponding to eigenvalues $\lambda_2, \lambda_3, \lambda_4$ were $t_2, t_3, t_4 = 6.8 \times 10^4,\ 4.2 \times 10^4,\ 1.0 \times 10^4$ respectively, in units of k^{-1} where k is the protein degradation rate (the Perron eigenvalue $\lambda_1 = 1$ is associated with the infinite-time (stationary) distribution). The corresponding values given by the WE-simulated $\widetilde{\mathbf{T}}(\tau)$ were

$6.1 \times 10^4,\ 3.5 \times 10^4,\ 9.4 \times 10^3$, respectively. These numbers demonstrate how the sampled $\widetilde{\mathbf{T}}(\tau)$ enables global approximation of slow system timescales to $< 20\%$ relative error. Quantitative error in these values depends on both "spectral" (lagtime) and discretization error [122] (see Fig. 5.2). In contrast, WE sampling in "rate mode" (see Methods) enabled highly accurate estimation of $\text{MFPT}_{X \to Y}$ to within 2% error (Tab. 6.7).

According to the Markov State Model framework, the presence of timescale separation indicates that a simplified model, retaining a few coarse-grained metastable states with Markovian transitions among them, can reasonably approximate the full system dynamics. Using this approach, we label the metastable sets as *phenotypes* accessible to the network, reasoning that a useful classification of cell phenotypes should be one that gives relatively stable, rather than transient, cell types. We apply the Markov State Model coarse-graining procedure to both the full $\mathbf{T}(\tau)$ and simulated $\widetilde{\mathbf{T}}(\tau)$, yielding similar results. The coarse sets (or metastable phenotype-states) in the reduced models for both cases are generated automatically, and map directly onto the four basins seen in the quasipotential landscape (i.e., the gene A/B expression hi/hi, hi/lo, lo/hi, and lo/lo cell phenotypes). The reduced models are visualized by network graphs, in which node sizes are proportional to steady-state probability, and the thicknesses and lengths of edges are proportional to the transition probability between them (on lagtime τ) (Fig 3.2c). Numerical values for the reduced models can be found in Tab. 6.6. The network graph can be considered to be an alternative representation of the global epigenetic landscape, which contains both stationary and dynamic information. (In contrast, the epigenetic landscape plotted as a quasipotential function does not explicitly contain dynamic information, due to non-gradient dynamics [158]).

Validation of the coarse-grained model can be carried out according to the Chapman-Kolmogorov test [122], which tests how well the relaxation dynamics initialized in the metastable phenotypes approximate the dynamics that are predicted either by the full model (CME) or simulated trajectories. According to this test, relaxation dynamics out of metastable

phenotypes from WE sampling was predicted with error values between 0.02 and 0.12 for all phenotypes (Fig. 6.16). Together, these results indicate (i) that a Markovian model of phenotype transitions is a good approximation of the full system dynamics for the ExMISA model, and (ii) that the WE-simulation based computational pipeline predicts a quantitatively similar coarse-grained phenotype-network to the full CME model.

3.3.3 The Method Maps the Epigenetic Landscape and Identifies Dominant Phenotypes in a Pluripotency Network Model

We apply the computational pipeline to a pluripotent fate-decision network from mouse Embryonic Stem Cells (mESCs) introduced by Zhang *et al.* [167] (Fig 3.3A). The network comprises eight interacting genes: NANOG, GATA6, CDX2, SOX2, OCT4, GCNF, and PBX1. Three of these genes, NANOG, SOX2, and OCT4 have been suggested to maintain pluripotency[26], and NANOG inhibits the expression of differentiation markers [139]. The GATA6 and CDX2 genes have been used in experiments as markers of differentiation, with the GATA6 transcription factor being a marker of the primitive endoderm cell lineage, and the CDX2 transcription factor being a marker of the trophectoderm lineage [57].

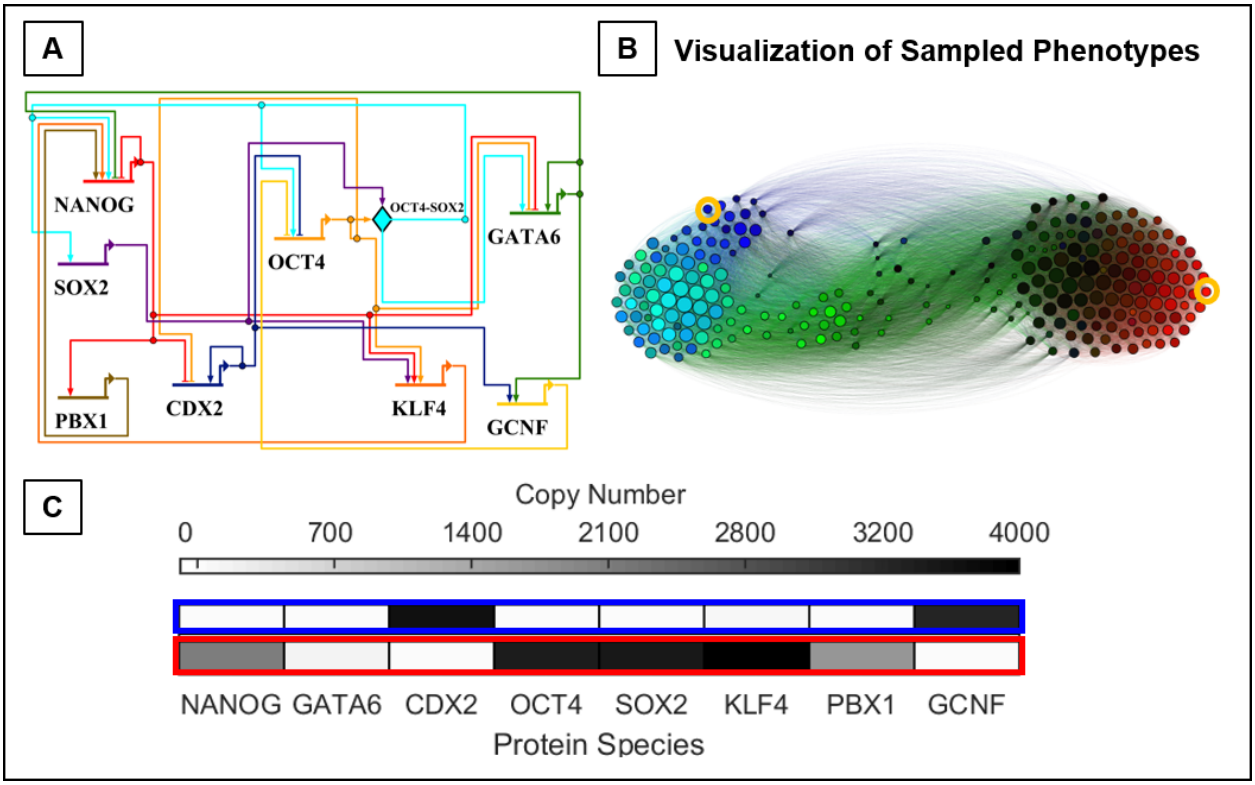

Figure 3.3: **Pluripotency Network Model and Simulation Results** A)Wiring diagram for the eight-gene pluripotency network model, adapted from [167]. Arrowheads represent positive interactions, while flat lines denote repression. B) Simulation results: state-transition graph of sampled network states. Circles represent aggregate gene-expression states sampled during the Weighted Ensemble simulation. Circle areas are proportional to the steady-state probability π_i in each state according to $\ln(\gamma\pi_i)$ with scaling factor $\gamma = 3.4$. States are colored according to the gene expression levels of three of the genes; red, green, and blue correspond to high NANOG, GATA6, and CDX2 expression respectively, while black corresponds to low or no gene expression. Edges connecting the states indicate possible state-transitions, colored according to the originating state. The graph is produced using Gephi [18] using a force-directed layout algorithm (Force Atlas), therefore short inter-state distances reflect higher probability of transitioning. C) Full protein compositions of two representative states, with either high CDX2 expression (blue) or high NANOG expression (red). States in (C) correspond to yellow circles in (B).

Using the WE-based computational pipeline, we estimate $\widetilde{\mathbf{T}}(\tau)$ with a resolution of $N_{\text{bins}} = 250$. To visualize the global landscape as a graph network at this resolution, we plot the converged $\widetilde{\mathbf{T}}(\tau)$ using a force-directed automated graph layout [18] (Fig 3.3B). The barbell shape of the network reflects the broad antagonism between pluripotency and differentiation genes, which is a general feature of the overall network topology. At the same time, each "pole" comprises

multiple distinct patterns of gene expression (seen in the graph as different colors with full compositions in Fig 3.3C), hinting at the existence of multiple phenotypes associated with both pluripotency and lineage-specification. Moreover, the network representation reveals numerous links between pluripotent and differentiated states, pointing to both direct and indirect transitions, through a network of relatively transient intermediate states.

To further analyze the global dynamics of the pluripotency network, we apply the Markov State Model coarse-graining framework. The simulated $\widetilde{\mathbf{T}}(\tau)$ shows gaps in the eigenvalue spectrum after four and after six eigenvalues (Fig 3.4a). The corresponding approximate timescales are given by $t_2,\ t_3,\ t_4,\ t_5,\ t_6 = 1.1 \times 10^5,\ 95,\ 51,\ 12,\ 12\ (k^{-1})$, respectively. These values, though only approximate, indicate the presence of a single long timescale process (t_2) corresponding to transfer between differentiated and pluripotent states, while transitions within those basins (t_3, etc.) occur at least four orders of magnitude more quickly. Applying the coarse-graining algorithm to achieve six clusters results in a reduced model (Fig 3.4b), with the clusters representing metastable phenotypes. The phenotypes can largely be distinguished in the subspace of NANOG, GATA6, and CDX2 expression levels; the differentiated phenotypes show expression of either GATA6 (primitive endoderm, PE), CDX2 (trophectoderm, TE), or both (denoted an intermediate cell type, IM). Phenotypes associated with pluripotency do not express high levels of GATA6 or CDX2, and may express high levels of NANOG (stem cell, SC). The coarse-grained model reveals two separate pluripotent phenotypes that are low in NANOG expression: one which expresses other pluripotent factors OCT4, SOX2, and KLF4 ("Low NANOG 1" LN1), and one which has low expression of all factors ("Low NANOG 2" LN2) (Fig 3.4c). Overall, these phenotypes broadly match experimentally-determined categories, coincide with steady-states of the stochastic model computed previously by a CME-approximation method[167], and coincide with phenotype-states identified in related pluripotency GRN models[83]. The steady-state probabilities associated with the phenotypes are highly nonuniform, with 95% of the population divided nearly evenly between the IM and LN1 phenotypes, which are associated with differentiation and pluripotency, respectively.

The LN2 state is rarest, comprising only $8 \times 10^{-4}\%$ of the population. Together, these results indicate that the clustering method identifies both common and exceedingly rare phenotypes in the *in silico* cell population modeled by simulation trajectories. Furthermore, the automated method identifies both expected and novel phenotypes.

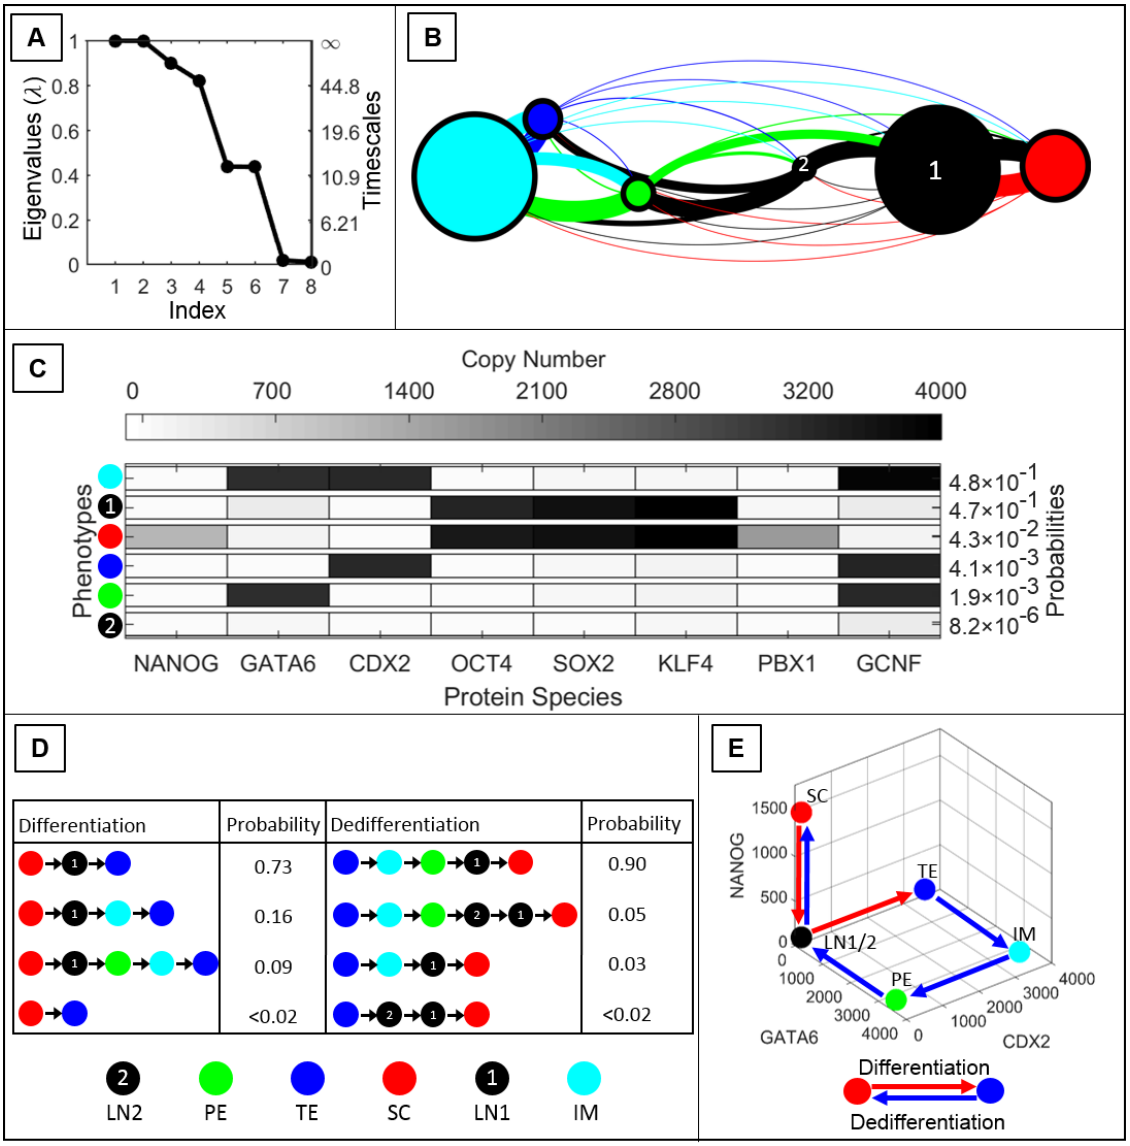

Figure 3.4: **Simulation Results for the Pluripotency Network (Parameter Set I). The Computational Pipeline Uncovers Six Metastable Phenotypes and Irreversible Phenotype Transitions.** A) Computed eigenvalue spectrum and global timescales indicating the presence of metastability in the network. The gap in the eigenvalue spectrum after the sixth eigenvalue suggests that a partitioning can be found into six metastable phenotypes. B) The coarse-grained network showing six algorithmically-identified phenotypes designated as Low NANOG 1 (LN1), Low NANOG 2 (LN2), Stem Cell (SC), Primitive Endoderm (PE), Trophectoderm (TE), and the Intermediate Cell (IM) state. C) The averaged gene expression levels (copy numbers) of each transcription factor for each phenotype and their respective steady-state probabilities. D) The four most probable transition pathways from the SC state to the TE state (differentiation) and from the TE state to the SC state (dedifferentiation). E) The highest probability transition paths projected onto three protein coordinates, NANOG, GATA6, and CDX2. Differentiation from SC to TE is visibly irreversible.

3.3.4 The Method Reveals Multiple, Irreversible Pathways for Phenotype Transitions in the Pluripotency Network

Previously, Markov State Models constructed on the basis of Molecular Dynamics simulations were used to analyze the ensemble of distinct pathways of protein-folding [114]. Here, we utilize the coarse-grained model of phenotype transitions in the pluripotency GRN in a similar manner, to analyze pathways of cell differentiation and dedifferentiation. Using Transition Path Theory, the method identifies the pathways that carry the greatest fraction of net probability flux, among sequences associated with successful SC→TE transitions (and reverse) (Fig 3.4d,e). Transition paths between the stem cell (SC) and PE phenotypes can be found in Fig. 6.17. For Parameter Set I, the method identifies three pathways encompassing $> 98\%$ of the probability flux for both forward and reverse transitions. While the SC→ TE transition is most likely to occur directly through the LN1 state (i.e., NANOG expression will shut off, followed by turning on CDX2), the reverse transition shows a different route through the IM and PE states (i.e., GATA6 expression turns on, then CDX2 turns off, then GATA6 turns off, and finally NANOG turns on).

Dynamic analysis of the coarse-grained model, including analysis of transition paths, relies on the Markovian approximation for inter-phenotype transitions. In the pluripotency network, stochastic transitions between pluripotency (SC, LN1, LN2) and differentiation (TE, IM, PE) basins are infrequent relative to transitions within those basins, justifying the Markovian assumption, since the system equilibrates within those basins much more rapidly than inter-basin transitions occur. However, the Markovian assumption may be less accurate for describing intra-basin transitions between phenotypes, which occur much more frequently. Despite the coarse-grained model encompassing transitions on highly disparate timescales, the qualitative results of transition path analysis were validated by collected conventional simulation trajectories (not subject to any Markovian assumption), which identified the same dominant transition paths (Fig. 6.18). Overall, these results indicate that a stochastic

excursion of a cell from the SC to TE phenotypes and back maps a cycle in gene-expression space, echoing previous studies indicating nonequilibrium dynamics in GRNs [158, 46]. The results further indicate that the Markov State Model, while a highly coarse-grained approximation, can provide an accurate estimation of inter-phenotype transition dynamics.

3.3.5 Cell Phenotype Landscape and Transition Dynamics are Sensitive to Kinetic Parameters

We applied the computational pipeline to the pluripotency network using two different rate parameters sets (see Chapter 6.3), which differ in rates of transcription factor binding and unbinding to DNA. In line with previous studies[46, 150, 30], we found that increasing the so-called adiabaticity (i.e., increasing h and f, or the rates of TF-binding relative to protein production and degradation, Parameter Set II) led generally to rarer inter-phenotype transitions (see Table 3.1). For example, in Parameter Set I, the Mean First Passage Time (MFPT) for transitions from SC $\rightarrow$ TE was calculated to be 1.36×10^5 in units of k^{-1}, as compared to 8.13×10^8 for Parameter Set II. The MFPTs of the reverse transition TE $\rightarrow$ SC for each set were 2.70×10^5 and 5.82×10^9, respectively (see Table 3.1 and S. Table 6.8). These differences in magnitude broadly reflect that moving toward the adiabatic regime leads to increased epigenetic barriers between phenotypes.

Transition	**SC → LN(1)**	**LN(1) → SC**	**SC → TE**	**TE → SC**
Parameter Set I ($f = 10$)	1.71×10^1	1.94×10^2	1.36×10^5	2.70×10^5
Parameter Set II ($f = 50$)	7.71×10^4	1.28×10^4	8.13×10^8	5.82×10^9

Table 3.1: **Computed Mean First Passage Times (MFPTs) of Phenotype Transitions in the Pluripotency Network.** MFPTs are shown for transitions between the pluripotency (high NANOG) state (SC) and low NANOG expression states (LN(1)) (left columns) and for transitioning between the pluripotency state (SC) and the trophectoderm state (TE) (right columns), in units of the inverse transcription factor decay rate, k^{-1}. Transitions for Parameter Set I were computed using the WE method in rate mode while transitions for Parameter Set II were estimated from the sampled transition matrix. The definitions of SC and LN(1) are analogous to the high NANOG production (N^{hi}) and low NANOG production (N^{lo}) transitions measured in experiments [47, 64]. Increasing the adiabaticity (i.e., the rates of DNA-(un)binding, h, f), leads to rarer inter-phenotype transitions. The simulations also show that, within the same gene network for a given parameter set, inter-phenotype transition times span four orders of magnitude.

In addition to generally slowing transitions, the increased adiabaticity of Parameter Set II gives rise to an epigenetic landscape structure that is distinct from that of Parameter Set I, with altered steady-state phenotype probabilities (Fig 3.5a). The eigenvalue spectrum shows qualitatively distinct features as well, with a gap after five values (Fig 3.6a). As such, the Markov State Model framework identifies five dominant phenotypes in the network, which correspond broadly to those of Parameter Set I, except that only a single Low-NANOG (LN) phenotype is identified (Fig 3.6b). Most of the steady-state probability is contained in the IM state (Fig 3.6c). In addition to altering the transition rates and relative phenotype probabilities, the kinetic parameters altered the dynamics of differentiation and dedifferentiation. The two likeliest pathways of forward (and reverse) SC → TE transitions follow the same route through LN and IM phenotypes (Fig 3.6d,e). Alternative differentiation pathways of forwards (and reverse) SC → PE transitions can be found in Fig. 6.20. These results indicate that, while the same GRN model with different kinetic parameters may give rise to qualitatively similar phenotypes, they differ in quantitative stationary and dynamic features, including relative steady-state probabilities, transition times, and likeliest transition

pathways.

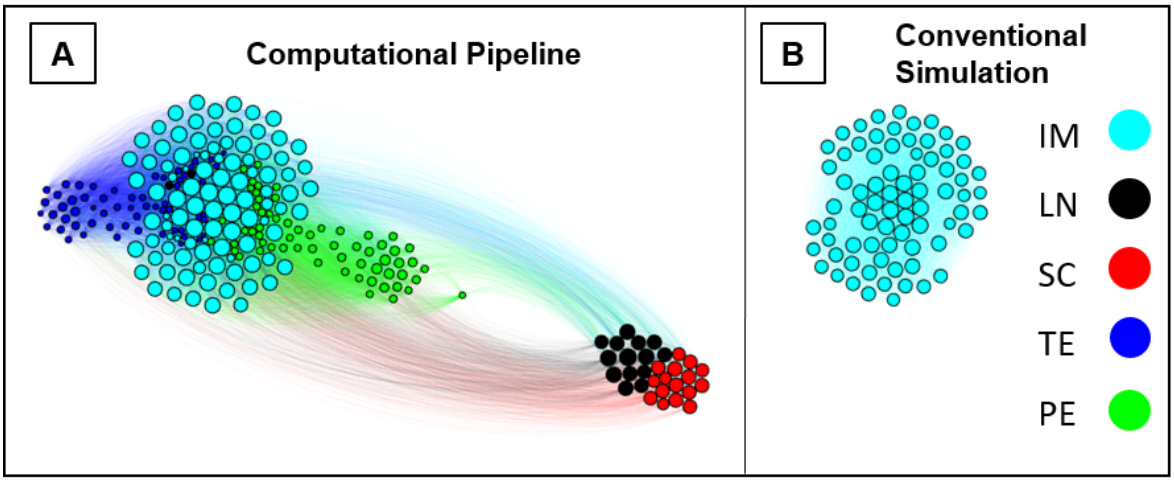

Figure 3.5: **The Rare-Event Sampling Pipeline Makes Rare States and Transitions Accessible to Simulation.** A) The global state-transition graph computed with the computational pipeline for the Pluripotency Network with rare transitions (Parameter Set II). The states are colored according to the coarse-grained (algorithmically-identified) phenotypes. In this parameter regime ($f = 50$) the differentiated (TE, PE, IM) and pluripotent phenotypes are cleanly separated, reflecting exceedingly rare transitions between the two phenotypes ($O(10^9)$, see Table 3.1). (B) States visited in conventional SSA simulation (using the same initialization, definitions, and placement as in (A)). In the conventional simulation, a transition out of the IM phenotype was never observed.

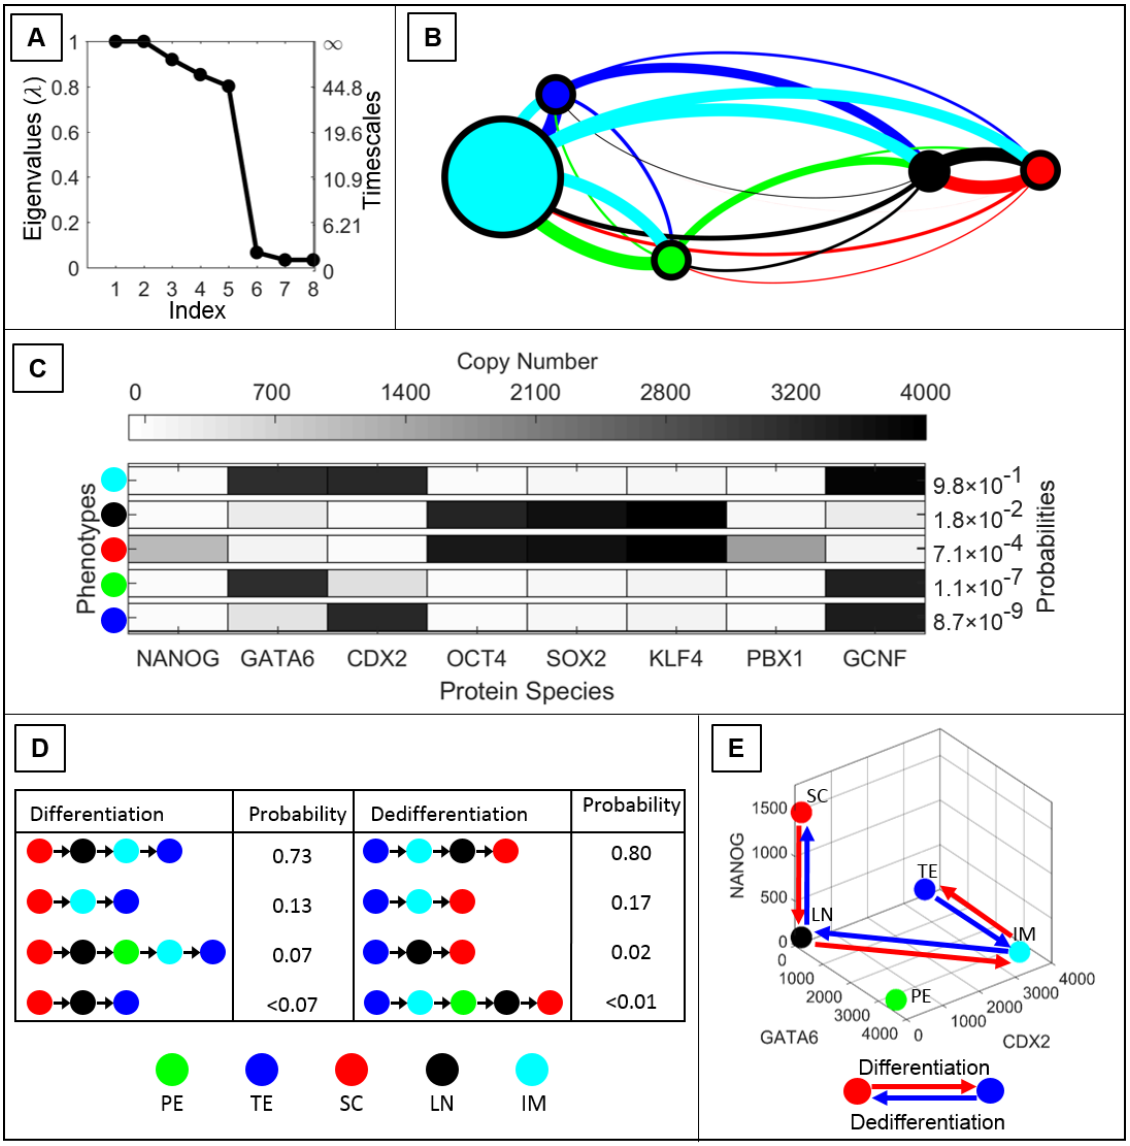

Figure 3.6: **Simulation Results for the Pluripotency Network (Parameter Set II). Changing DNA-Binding Kinetics Alters the Epigenetic Landscape.** A) Computed eigenvalue spectrum and global timescales. B) The coarse-grained Markov State Model showing five phenotypes corresponding to the LN1, SC, PPE, TE, and IM phenotypes of Parameter Set I. The majority of the steady state probability is in the IM phenotype (0.98). C) The gene expression levels for each phenotype and their respective steady-state probabilities. D) The four most probable differentiation pathways between SC and TE phenotypes. E)The dominant pathways of (de)differentiation projected onto the GATA6, CDX2, and NANOG coordinates. The change in DNA-binding kinetics shows different transition dynamics from Parameter Set I. Here, the forward and reverse paths are the same.

3.3.6 Efficiency of Rare-Event Sampling Compared to Conventional SSA

Phenotype transitions that are relatively rare can be difficult to observe with conventional SSA simulation. We compared simulated landscapes (based on estimated $\widetilde{\mathbf{T}}(\tau)$) from the computational pipeline to those obtained from an equivalent (large) number of SSA simulation steps (Fig 3.5a,b). Additional comparisons of synthetic cell populations using tSNE visualization reflect the rarity of phenotypes and phenotype transitions (Fig 3.7, Fig. 6.21). This comparison revealed that the WE-based method uncovers multiple phenotypes and associated transitions that are invisible to conventional simulation due to the rarity of exiting metastable basins. Quantitative estimates of efficiency gains for WE have often been based on comparing the number of simulation steps required to estimate a desired quantity (such as a rate constant) using WE versus conventional simulation [38]. Treating $\widetilde{\mathbf{T}}(\tau)$ as the desired output (as it contains holistic dynamic information for the system), we estimate the efficiency gain of our pipeline by computing:

$$E = \frac{\text{Sim. steps to estimate } \widetilde{\mathbf{T}}(\tau), \text{ Conv.}}{\text{Sim. steps to estimate } \widetilde{\mathbf{T}}(\tau), \text{ WE}}. \tag{3.9}$$

However, it is often difficult to acquire the required number of steps for conventional simulation, so an approximate lower bound for the denominator can be estimated according to:

$$\Big[\text{Sim. steps to estimate } \widetilde{\mathbf{T}}(\tau), \text{ Conv.}\Big] \gtrapprox \sum_{i,j} \text{MFPT}_{i,j}, \tag{3.10}$$

where simulation steps and transition times are measured using the same time-unit (here, k^{-1}). That is, the denominator is the sum over the MFPTs of transitions between each pair

of states (bins), i, j, where $i, j = 1...N_{\text{bins}}$. This approximation is based on the rationale that one requires simulation time $\mathcal{O}(\text{MFPT})$ to observe at least one transition between a given pair of states. From the WE-estimated transition-matrix $\widetilde{\mathbf{T}}(\tau)$, estimates of the MFPT for transitions between any pair of states (bins) can be obtained using Eq 3.6. According to Eq 3.9, we estimate that our pipeline provided efficiency gains of 3000 for ExMISA (Fig. 3.2), 200 for Pluripotency Parameter Set I (Fig. 3.3), and 4×10^7 for Parameter Set II (Fig. 3.6). These numbers show that the pipeline affords a significant speedup over conventional simulation in providing global dynamic information. The numbers further show that the efficiency gain is most pronounced for the Pluripotency network with exceedingly rare inter-phenotype transitions.

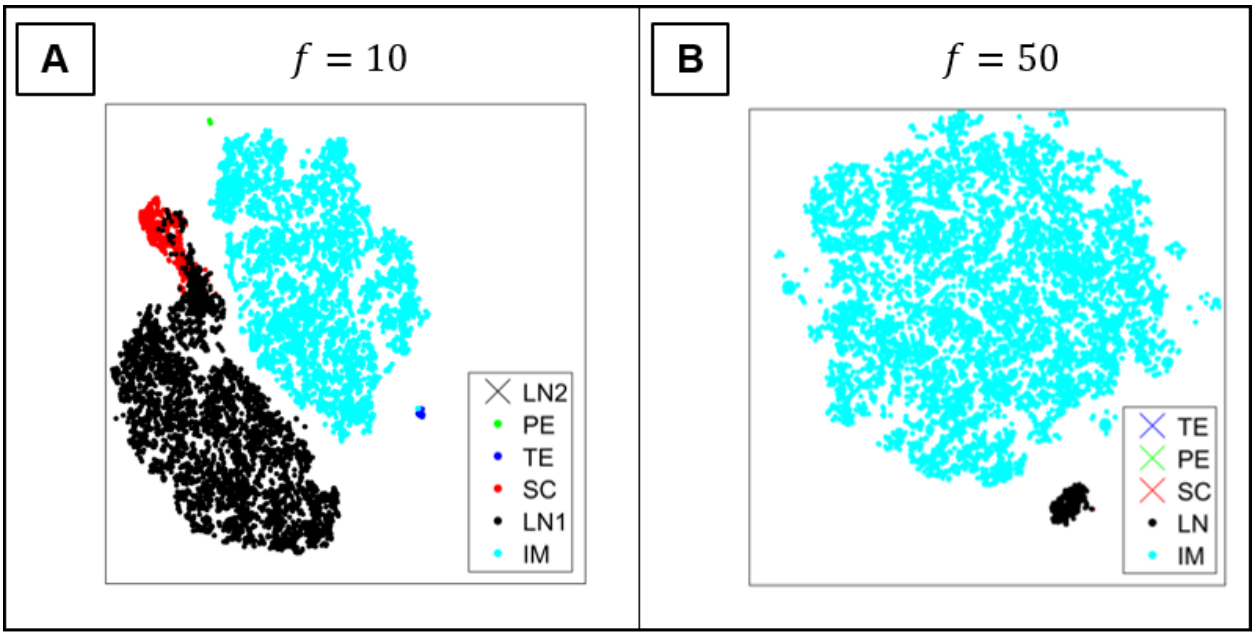

Figure 3.7: **"Synthetic" Cell Population Data Computed by the Rare-Event Sampling Pipeline, Visualized with a Single Cell Visualization Method (tSNE).** A) tSNE visualization of 15000 simulated 'cells' (replicas) drawn from WE sampling for Parameter Set I. Each cell is colored according to its phenotype after the coarse-graining. The population is heavily dominated by the IM and LN phenotypes, though all other phenotypes are sampled with the exception of the rare LN2 phenotype. b) tSNE visualization of Parameter Set II simulation data. Only the LN and IM phenotype-states are sampled in a synthetic population of size 15000.

3.4 Discussion

In this work, we present a method for efficient, automated computation of epigenetic landscapes, metastable phenotypes, and phenotype-transition dynamics of stochastic GRN models. Our computational pipeline was inspired by studies of metastability and barrier-crossing in Molecular Dynamics, and our application of the pipeline to cell-scale networks addresses a number of current challenges for stochastic GRN dynamics. First, it overcomes the curse-of-dimensionality of complex models, by leveraging available rule-based modeling tools for stochastic biochemical networks [44]. Second, it overcomes the challenge of efficiently simulating stochastic systems with rare events, by using enhanced Weighted Ensemble rare-event sampling [70]. Third, it addresses the challenge of extracting and interpreting essential dynamics of complex systems on the basis of simulated trajectories, by using the Markov State Model framework [114] to automatically generate a compact, approximate representation of global system dynamics. Combining these tools into a unified pipeline provides an automated means of computing and visualizing essential stationary and dynamic properties of stochastic GRNs, including the number and identities (i.e. state-space mapping) of metastable phenotypes, their steady-state probabilities, and most-likely pathways of inter-phenotype transitions and their transition rates. By advancing the capability to compute and interpret hypothesized or experimentally-derived stochastic GRN models, the method can yield insight into how "local" stochastic, molecular processes involved in epigenetic regulation affect "global" dynamics such as phenotypic stability and fate-transitions in cells. Moreover, it can help close the gap between dynamic, molecular-detailed models of gene regulation and cell-population level experimental data, to inform rational cell reprogramming strategies.

3.4.1 Insights from the Pluripotency Network Simulations

We used the pluripotency network as a model system to develop and demonstrate the simulation approach, but the results also yielded biological insights. For example, the simulations revealed a hierarchical structure of the epigenetic landscape. The network–exhibiting 5-6 metastable phenotypes–occupies a limited subspace from the vast possible gene combinations (e.g., $2^8 = 256$ possible distinct on/off combinations of gene expression states). The dominant feature of the global landscape is a high barrier/slow timescale between pluripotent and differentiated phenotypes. Within each of these categories, further sub-states were identified. The model revealed multi-timescale dynamics of phenotype transitions; the pluripotency network showed relatively rapid transitions between phenotype-states that differed in the expression-level (high vs. low) of a single gene, e.g. the high NANOG to low NANOG transition, whereas phenotype transitions involving a change in expression level of seven genes, e.g. the SC macrostate to the TE macrostate, occurred five orders of magnitude more slowly on average.

While the accessible phenotypes appear broadly similar across parameter sets, the relative stability and transition dynamics among phenotypes were sensitive to kinetic parameters governing transcription factor binding/unbinding. A global change in these parameters (affecting all individual transcription factor-DNA interactions equally) changed the shape of the landscape, altering the relative steady-state probabilities of different phenotypes and the likely transition pathways linking them. The DNA binding parameters capture the local epigenetic mechanisms that enable/disable transcription factors from accessing regulatory elements. A global rate change nevertheless has a varying influence on different genes because the number of regulators differs, as does the molecular logic by which activators and repressors exert combinatorial control on different genes. These results echo findings that global modification of chromatin regulators often have lineage-specific effects [31]. These results highlight both the need for, and the challenge, of informing cell reprogramming

strategies with quantitative network models, as they suggest that the dynamic response of cellular networks to perturbations is governed by the detailed kinetics of molecular regulatory mechanisms, which are generally difficult to parameterize.

3.4.2 Dynamic Definition of Cell Phenotype

The Markov State Model framework implicitly imposes a dynamic definition of cell phenotypes; the number of phenotypes was determined using spectral gap-analysis, and the coarse-graining algorithm automatically identified metastable aggregates (i.e., grouped sampled network states into larger clusters). This is different from the classifications of phenotypes that are generally used in analyzing experimental data, where gene expression or marker levels are often used to categorize cells. However, experiments have also revealed the potential need for a dynamic definition of cell phenotype, based not only on single-timepoint measurements of gene expression or phenotype-markers, but also on information from past or future timepoints[72, 47]. For example, Filipczyk *et al.* [47] identified distinct subpopulations within a compartment of NANOG-negative cells in mESCS, which differed in their propensity to re-express NANOG. At the same time, fluctuations between low- and high-NANOG expressing cells were not necessarily associated with any functional state change. The Markov State Model approach, based on kinetic/dynamic coarse-graining, thus provides a quantitative approach for classifying phenotype-states that is both completely generalizable rather than *ad hoc* (it requires no *a priori* knowledge or designation of markers/genes) and is in line with these recent experiments revealing the need for a dynamic definition of phenotype.

3.4.3 Timescales of Stochastic Phenotype Transitions

Markovian transitions (i.e., memoryless "hops") among cell phenotypes have been observed experimentally: examples include transitions among phenotypes in cancer cells, as measured

by flow cytometry [54], and among pluripotency-states in mESCs, as measured by time-lapse microscopy of fluctuating gene expression [140, 47, 64]. The compact nature of these data-inferred networks–showing hops among a limited set of broad phenotypes–suggests that the computed MSM framework advanced in this study provides an appropriate level of resolution at which to analyze GRN dynamics and may serve as a useful tool for comparing models to experimental data.

Experimental studies have quantified the timescales of Markovian transitions between NANOG-high and NANOG-low states in mESCs [47, 64]. From Hormoz *et al.*, the probability of transitioning from NANOG-high to NANOG-low in mESCs is 0.02 per cell cycle, while that of the reverse transition is 0.08. These values represent a relatively rapid transition rate, since NANOG expression is known to be particularly dynamic [139]. Similarly, plasticity has been observed in cancer cells where quantitative estimates of stochastic cell transitions between a stem cell cancer cell phenotype to a basal cancer cell phenotype were observed to be roughly on the order of 0.01 to 0.1 per cell cycle [54]. We can translate our model results to approximate biological timescales: the degradation rate, which sets the timeunit for model results (i.e., k is taken to be 1) was experimentally determined to be on the order of a few hours (in the E14 mouse embryonic stem cell line, the half-lives of NANOG, OCT4, and SOX2 are approximately 4.7, > 6, and 1.6 hours, respectively [2]). Assuming that degradation is unimolecular, $k = \ln(2)/t_{[NANOG]1/2}$, and the half-life of NANOG, $t_{[NANOG]1/2} = 5$ hours, the degradation rate is $k = 0.1$. Using a mESC cell cycle time of 12 hours [156], the simulations for Parameter Set I then predict NANOG-high to NANOG-low transitions occurring with a rate of 0.03 per cell cycle, and of 3×10^{-3} for the reverse. For Parameter Set II, the computed rates were 8×10^{-6} and 5×10^{-5}, respectively. Comparison of these computed and experimental rates of NANOG transitions indicates that Parameter Set I ($f = 10$) is more in line with experimental observations, while Parameter Set II ($f = 50$) gives transition rates that are three orders of magnitude too slow. These results are in agreement with previous findings from theoretical studies that GRNs in pluripotency networks operate in a

so-called "weakly-adiabatic" regime [130, 167, 150], in which the timescale of DNA-binding by transcription factors is on the order of transcription factor production and degradation.

3.4.4 Comparison to Other Models and Computational Approaches

A number of theoretical studies have elucidated dynamics of stochastic molecular-detailed GRN models (i.e., models that include molecular fluctuations and regulatory mechanisms, in contrast to Boolean models[28]). These studies have largely focused on small 1- or 2-gene motifs[[16, 75, 131, 46, 95, 150, 49]], but recent years have seen extension of stochastic methods to studies of more complex, experimentally derived GRN models encompassing $\mathcal{O}(10)$ genes. For example, determination of global dynamic properties of such networks has been achieved by combining information from long stochastic simulations of discrete models [130, 83], or of continuum SDE models, in combination with path integral approaches [160, 84]. The pluripotency network studied herein was developed by Zhang and Wolynes [167]; in their work, the authors developed a continuum approximation to the Chemical Master Equation that enabled quantitative construction of the epigenetic landscape. Here, we present an alternative approach that is unique in two major aspects: (1) the use of stochastic simulations (i.e., SSA [51]), which is enabled by use of the WE rare-event sampling algorithm, and (2) the automated Markov State Model framework for designating phenotypes and constructing a coarse-grained view of the epigenetic landscape. While we utilize a different framework (that of coarse-grained, discrete stochastic models) from Zhang and Wolynes to approximate and interpret dynamics, our results are broadly consistent with theirs. For example, the dominant identified phenotypes we found are the same as in their work (the only exception being the exceedingly rare LN2 phenotype identified by the coarse-graining algorithm for Parameter Set I).

3.4.5 Current Challenges

Several challenges and potential weaknesses with the pipeline exist, both with regard to sampling rare events, and in determining an appropriate coarse-grained model. Potential challenges with the WE algorithm have been described elsewhere[39, 172], and include the difficulty of determining a binning that captures slow degrees of freedom and the existence of time-correlations between sampled iterations of the simulation, which can impede unbiased sampling. The Voronoi-based binning procedure we employ here is related to a number of similar approaches [35, 166, 33, 150], and has the advantage of effectively tiling a high-dimensional space without the need for *a priori* knowledge. However, in practice, according to others and our own studies, the method is effective up to about 10 degrees of freedom. These challenges are the subject of continued study.

3.5 Supporting information

For more information, see Chapter 6.3 and Chapter 6.5.

3.6 Acknowledgments

We thank the administrators of the University of California, Irvine High-Performance Computing Cluster and we thank Jun Allard for helpful discussions.

Chapter 4

Regulation of Macrophage Polarization and Plasticity by Complex Activation Signals

4.1 Introduction

Macrophages are functionally complex cells of the immune system, and participate in both innate and acquired immunity. In addition to their roles as professional phagocytes and antigen-presenting cells, macrophages are sensitive integrators and transducers of biochemical signals with a wide repertoire of responses. A recent study of the macrophage transcriptomic response to a range of stimulatory molecules identified dozens of distinct mRNA coexpression modules [163]. Two particular macrophage response patterns which have been widely recognized are the M1 and M2 programs, named because they are respectively elicited by products of Th1 and Th2 cells [21, 105]. The M1 response, also described as classical activation, is typically evoked in vitro by treating cells with interferon-γ (IFN-γ) and lipopolysaccharide

(LPS), a bacterial cell wall component and TLR4 agonist. The M1 phenotype includes production of inflammatory cytokines including TNF-α and IL-1β[11]. M1 macrophages also undergo a metabolic transition towards glycolysis and secrete free radicals. The M2 response, also known as alternative activation, is evoked by IL-4 and IL-13 treatment and is characterized by increased expression of CD206, scavenger receptors, and—in mice—arginase 1. The M2 response has been further classified as the M2a response after other forms of alternative activation were identified. These stereotyped responses are likely only fully recognized *in vitro* but M1-like and M2-like phenotypes are readily identified in physiological contexts.

M1- and M2-activated macrophages exhibit characteristic transcriptional and secretory profiles. M1 activation is associated with STAT1 and IRF activation, and M2 activation is associated with STAT6 activity [99]. These pathways suppress each other; IL-4-induced STAT6 activation suppresses STAT1-dependent transcription in mouse macrophages [115], and STAT1 activation suppresses STAT6-dependent transcription [154]. Further, costimulation with IL-4 reduces the IFN-γ-dependent surface expression of FcγR on human monocytes[149]. Yet, despite the evidence of mutual repression, markers associated with both M1 and M2 phenotypes have been observed simultaneously on individual cells in vivo [144]. This co-expression may reflect simultaneous co-activation of M1 and M2 programs. A similar process is observed in T cells, where differentiation of CD4+ T cells to IFN-γ-secreting Th1 cells or IL-4-secreting Th2 cells in mixed culture conditions yields a tunable continuum of cell fates [12]. Modeling revealed that this outcome was consistent with gene regulation governed by a mutual inhibition, self-activation (MISA) network, a common motif thought to govern alternative fate-decisions in many cell types [171], including macrophages[138, 80].

M1 and M2 marker coexpression may also indicate that cells are shifting from one phenotype to another as the microenvironment changes. Indeed, macrophages have been shown to exhibit phenotypic plasticity *in vitro*. M1-activated macrophages induced by exposure to bacteria

[32] or IFN-γ, alone [121] or in combination with LPS [76], can be re-polarized to express markers associated with an anti-inflammatory phenotype upon subsequent treatment with IL-4, alone or in combination with IL-13. Similarly, macrophages treated with IL-4 will express inflammatory markers upon subsequent treatment with LPS or IFN-γ [32]. The plasticity of these influential cells has made them an attractive target for immunomodulation; scaffolds and materials that achieve controlled delivery of macrophage-activating agents is an active area of research for treatment of diseases involving macrophage dysregulation (reviewed in Ref. [10]). In atherosclerosis, sustained inflammation exacerbates oxidative stress in the plaque [107]. Plaque shoulders, which are prone to rupture, are typically dominated by macrophages expressing markers associated with M1 activation [144]. Reprogramming macrophages toward a M2 phenotype may prevent plaque rupture and promote plaque resolution by encouraging matrix deposition [102, 151]. In the case of cancer, tumor-associated macrophages (TAMs) are thought to induce anti-inflammatory signaling that helps protect the tumor from immune assault [142]. Recent evidence suggests presenting M1-activating factors to TAMs can help engage the immune system to attack the tumor [89]. An improved understanding of how macrophages respond to stimuli that redirect their phenotype should help develop better therapeutic approaches for these important pathologies.

In this study, we investigated how treatment with mixed M1 and M2 stimuli, either simultaneously or sequentially, regulates macrophage phenotype. We were motivated to understand whether a mixed phenotype represents a superposition of the M1 and M2 phenotypes, a transition between states, or, as some findings have suggested [119], a unique mixed program. To consider the expression state of individual cells, we used flow cytometry to assay surface expression of M1 marker CD86 and M2 marker CD206. Mouse bone marrow derived macrophages were stimulated with LPS/IFN-γ and/or IL-4/IL-13 at various concentrations over durations of 24-96 hours. Our findings suggest that macrophages adopt a mixed phenotype dependent on the relative strength of stimuli present, and that cells progress towards a M2 phenotype over time. These temporal changes in expression were found to be consistent with

a mathematical model comprising a modified MISA network. In addition, reprogramming of macrophages to the opposing phenotype is dependent on the extent of pre-polarization. More specifically, expression of CD206 in response to IL-4/IL-13 is enhanced by pre-polarization towards an M1 phenotype with LPS/IFN-γ. In contrast, expression of CD86 in response to LPS/IFN-γ, particularly at low concentrations, is inhibited by pre-polarization towards an M2 phenotype with IL-4/IL-13. Together, these data provide evidence of a macrophage phenotypic continuum by analysis of phenotypic markers at the single cell level, and suggest that macrophage reprogramming by combined activation signals is dependent on initial polarization state and dosage of stimulation.

4.2 Results

4.2.1 Co-stimulated macrophages express markers of both M1 and M2 activation

We first sought to establish a concentration range that would yield a submaximal response to facilitate detection of enhancement and repression effects. We exposed mouse bone marrow derived macrophages (BMDM) to varying concentrations of stereotypical M1 (LPS/IFN-γ) or M2 stimuli (IL-4/IL-13). To maximize the sensitivity of our assays, we selected concentrations from 0 to 0.3 ng/ml for LPS/IFN-γ and 0 to 1 ng/ml for IL-4/IL-13 since these ranges did not completely saturate expression of CD86 and CD206. Cells were exposed to stimulus for 48 hours and assayed for expression of M1 marker CD86, a T-cell costimulatory molecule, and M2 marker CD206, a mannose receptor, by flow cytometry. We found that median CD86 labeling intensity increased tenfold as the concentration of LPS/IFN-γ was increased from 0 to 0.3 ng/ml (Fig. S6.24a and b). Labeling intensity of CD206 increased threefold as IL-4/IL-13 was increased from 0 to 1 ng/ml (Fig. S6.24c and d). At these concentration

ranges, the expression of phenotypic markers was not saturating, so that the expression of markers generally increased with stimulation concentration.

To explore the effect of co-stimulation with M1 and M2 activation signals on macrophages, BMDM were exposed simultaneously to combinations of LPS/IFN-γ and IL-4/IL-13 at concentrations in the determined range for 48 hours. Expression of CD86 and CD206 was analyzed by flow cytometry (Fig. 4.1a and b). Notably, the population remained single-peaked in plots of CD86 expression vs CD206 expression, and did not show separation into distinct subpopulations. Cells generally did not individually commit to exclusive CD86 or CD206 expression. Indeed, CD86 and CD206 expression were only partially inhibited by exposure to their opposing polarization signal. Analysis of CD86 expression in LPS/IFN-γ-stimulated cells (0.3 ng/ml) showed that moderate amounts of co-added IL-4/IL-13 (0.1 ng/ml) in fact increased the median expression of CD86 40% ($\pm$ 22% SEM, $n = 3$; Fig. 4.1c). Further increasing the IL-4/IL-13 concentration to 1 ng/ml abrogated this enhancement. Meanwhile, the expression of CD206 in IL-4/IL-13 stimulated cells was not affected by the co-addition of LPS/IFN-γ stimuli at any concentration (Fig. 4.1d). In sum, these data demonstrate that macrophages exposed to combinations of the activation signals LPS/IFN-γ and IL-4/IL-13 express both CD86 and CD206 at 48 h of stimulation, and repression of the contrasting pathway was only partially observed with these phenotypic markers.

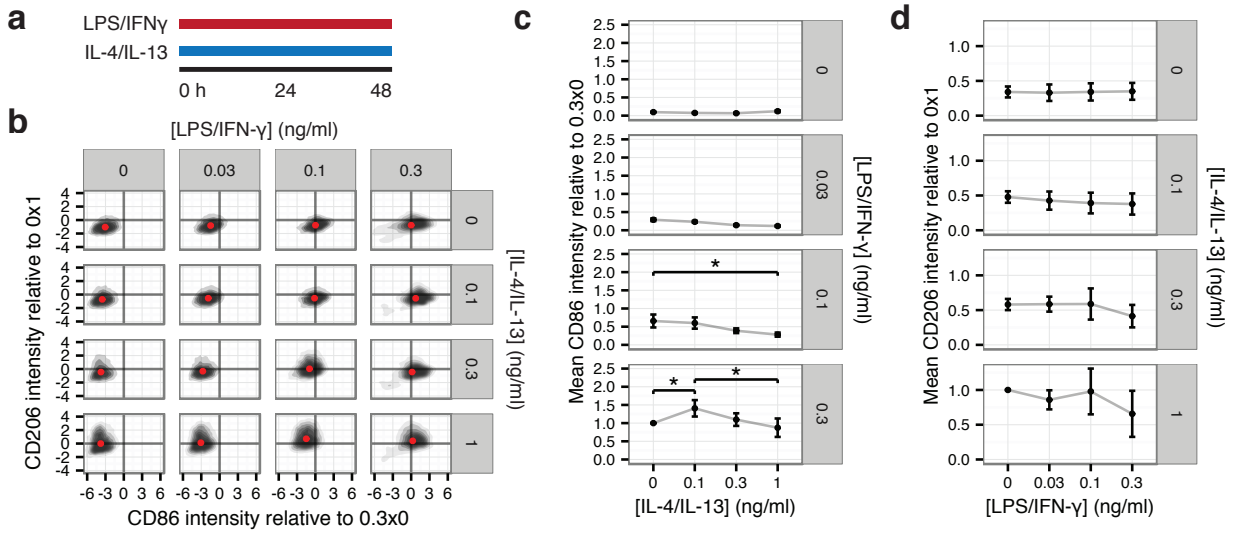

Figure 4.1: **Co-stimulation with LPS/IFN-γ and IL-4/IL-13 leads to expression of both CD86 and CD206.** (a) Schematic illustrating experimental conditions. Macrophages were exposed to LPS/IFN-γ and/or IL-4/IL-13 for 48 hours before analysis. (b) Density plots of normalized CD206 versus CD86 staining intensity of macrophages subjected to different concentrations of LPS/IFN-γ and/or IL-4/IL-13 (ng/ml) for 48 hours, assessed by flow cytometry. CD86 is normalized to the LPS/IFN-γ-only condition and CD206 is normalized to the IL-4/IL-13-only condition. Representative plots from a single experiment. Median position is indicated by a red dot. (c) Average median normalized CD86 intensity $\pm$ SEM (n = 3) of LPS/IFN-γ treated cells vs. co-added IL-4/IL-13 stimulus, grouped by LPS-IFNγ concentration, normalized per experiment as in B. (d) Average median CD206 intensity $\pm$ SEM (n = 3) of IL-4/IL-13 treated cells vs. co-added LPS/IFN-γ stimulus, grouped by IL-4/IL-13 concentration, normalized per experiment as in B. Asterisk indicates significant difference by two-sided t test, $p < 0.05$.

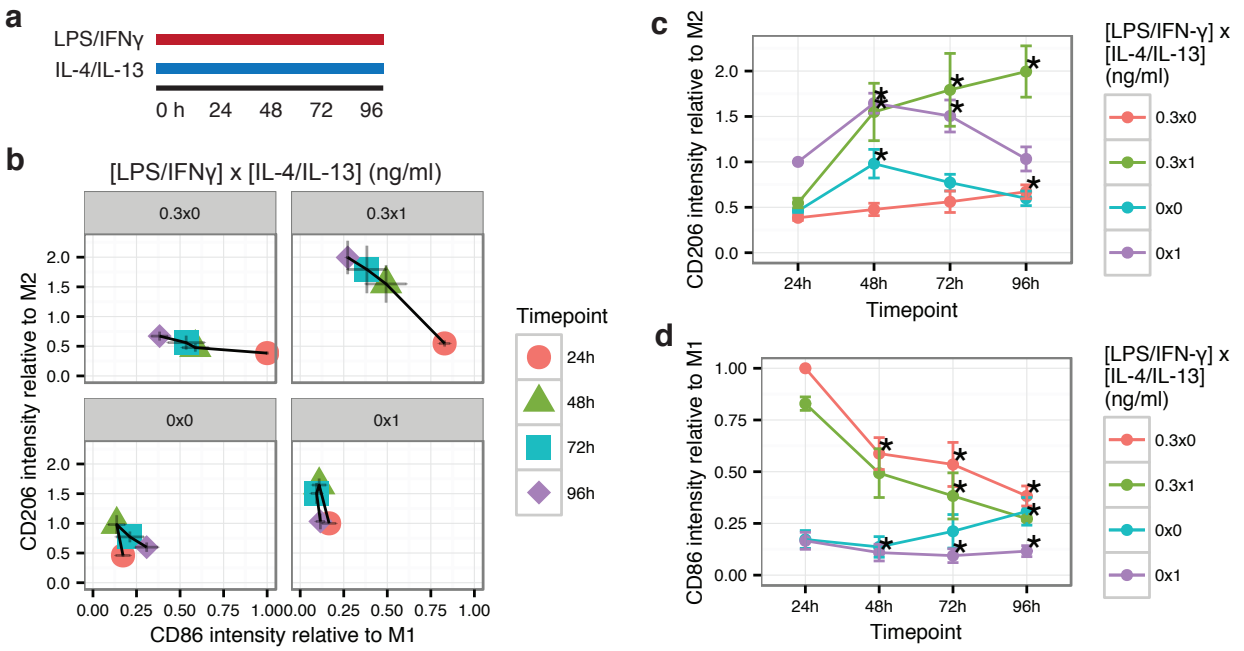

Figure 4.2: **Co-stimulated macrophages exhibit decreased CD86 expression and increased CD206 expression over time.** (a) Schematic illustrating experimental conditions. Macrophages were exposed simultaneously to LPS/IFN-γ and/or IL-4/IL-13 for 24, 48, 72, or 96 hours. Each experiment used BMDM isolated from a single mouse. (b) Expression of CD206 versus CD86 of different stimulation conditions over time. Average of median population location $\pm$ SEM ($n = 5$) is shown. (c) Expression of CD86 staining intensity over time for different stimulation conditions, normalized to the intensity of LPS/IFN-γ condition at 24 hours. Average of median population location $\pm$ SEM ($n = 5$) is shown. Asterisk indicates difference vs 24 hours by two-sided t test, $p < 0.05$. (d) Expression of CD206 staining intensity over time for different stimulation conditions, normalized to the intensity of IL-4/IL-13 condition at 24 hours. Average of median population location $\pm$ SEM ($n = 5$) is shown. Asterisk indicates difference vs 24 hours by two-sided t test, $p < 0.05$.

4.2.2 Modeling proposes a complex interdependence of M1- and M2-associated pathways

In order to gain further insight into the logic of macrophage activation, we performed mathematical modeling of CD86 and CD206 expression in response to the different costimulatory conditions. Our modeling strategy was designed to identify the key features of the regulatory logic linking CD86 and CD206 expression ("outputs") to stimulation by LPS/IFN-γ and/or

IL-4/IL-13 ("input"). To this end, we analyzed a suite of candidate models and performed model selection based on fitting to the experimental 96-hour timecourse data (Fig. 4.3). Mathematical descriptions of the models are provided as Supplementary Equations 6.6; the parameters are described in Supplementary Table S6.25 and the best-fit values are given in Supplementary Table S6.26. Rather than treating signaling and gene regulatory networks in detail (as quantitative parameters remain unknown), the network models comprise a small number of interacting nodes representing inputs, outputs, and M1- and M2-associated pathways. Models of T cell specialization [12, 97, 62] and fate-decisions in diverse cell types [171] commonly employ a core Mutual Inhibition, Self-Activation (MISA) network motif. We found that the basic MISA motif was insufficient to reproduce the observed temporal expression patterns, including the decrease of CD86 expression after 24 hours and the sharp increase of CD206 after 24 hours under costimulatory conditions.

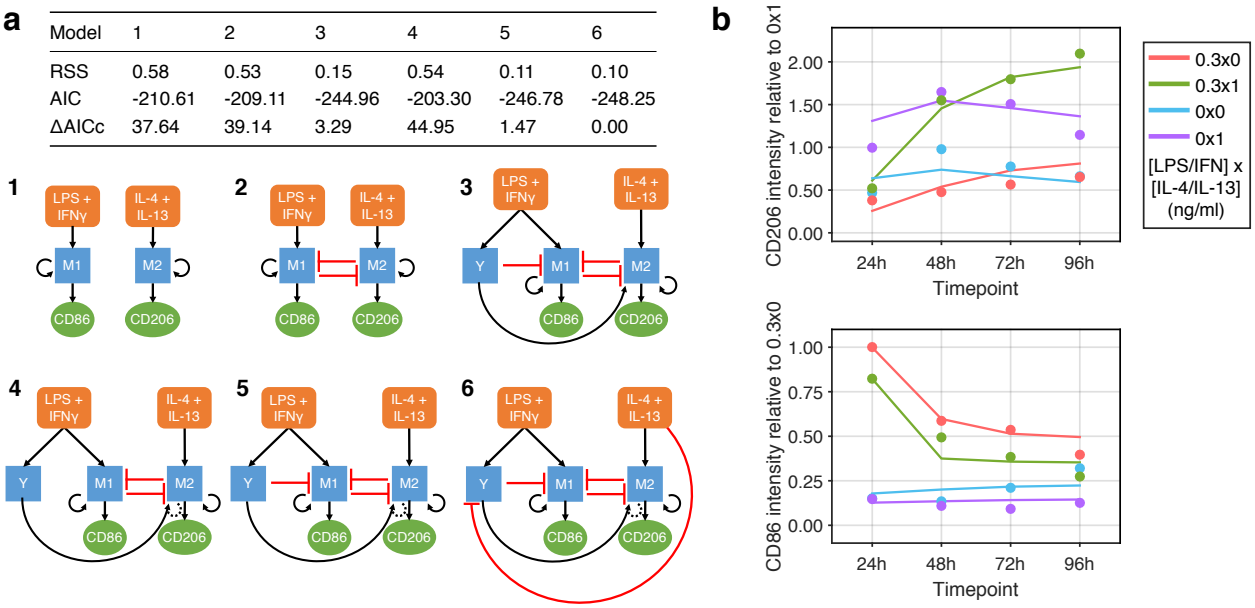

Figure 4.3: **Mathematical modeling of macrophage regulatory logic.** (a) A representative set of minimal models for activation of M1 and M2 pathways under costimulation, comprised of a modified MISA (Mutual Inhibition/Self-Activation) network. Y: an unspecified regulator. Dashed line indicates that positive regulation of M2 by Y occurs cooperatively with M2. Models are shown in order of increasing number of parameters. RSS score indicates goodness of fit, and AIC score measures model quality, while penalizing presence of additional parameters. Δ AICc is reported relative to best (lowest) value, corresponding to Model 6. In general, successful models (3,5, and 6) incorporated both an incoherent feed-forward loop on M1 and positive regulation of M2, mediated by Y. (b) Simulated expression of CD86 and CD206 stimulated with LPS/IFN-γ and/or IL-4/IL-13 for 24, 48, 72, or 96 hours. Simulated data are from the model with both the best AICc and RSS score, Model 6, fit to the normalized timecourse data.

We explored a number of additional small-network topologies, consistent with current knowledge of macrophage activation pathways. In the MISA paradigm, costimulus results in a mixed response, in which both markers are expressed simultaneously, albeit at a somewhat reduced level as compared to the strongly polarized case. While the temporal expression of CD86 shows this behavior (dampened, but qualitatively similar kinetics, with addition of IL-4/IL-13), the CD206 kinetics suggests a more complex response to costimulation. We found that successful models (as assessed by either the error or the AICc information criterion) required at least two features extending the core MISA: an incoherent feed-forward loop on the M1 pathway, and a mixed (both inhibiting and activating) character of M2 regulation by

the M1-axis (Fig. 4.3, Supplementary Fig. S6.26). In the network models, these interactions are mediated by a single additional intracellular regulator (labelled Y). The predictions of the candidate model are presented alongside experimental results in (Fig. 4.3b). Details of the mathematical modeling can be found in the Supplementary Information 6.6.

4.2.3 Multiparametric characterization of macrophage phenotype

To assess whether the expression of CD86 and CD206 are representative of macrophage function, we performed a multiplexed cytokine assay. We found that macrophages stimulated with LPS/IFN-γ alone for 24 hours exhibited the highest secretion of inflammatory cytokines including IL-6, IP-10, MIG, MIP-1a, MIP-2, MIP-1B, RANTES, and TNF-α, which were found at much lower levels in the supernatants of unstimulated cells or cells exposed only to IL-4/IL-13 (Fig. 4.4a). Cells that were stimulated with mixed LPS/IFN-γ and IL-4/IL-13 secreted somewhat lower levels of inflammatory cytokines, with the greatest proportional decrease observed in G-CSF, IL-6, IL-12, IL-15, and TNF-α. Cytokines that were expressed in greater quantity by M1-stimulated cells compared to naïve cells were also expressed by cells exposed to mixed stimuli. This was consistent with a moderate but significant (17% $\pm$ 3.5% SEM, $n = 6$) decrease in CD86 expression in cells treated with mixed cytokines compared to the LPS/IFN-γ only condition at 24 hours (Fig. 4.4C). Examining cytokines typically associated with M2 macrophages, we found that IL-10 was elevated, though not significantly, in the LPS/IFN-γ-only and mixed conditions. TGF-β 1 and 2 expression was similar among all conditions examined. VEGF was strongly suppressed in conditions containing IL-4/IL-13. Indeed, none of the analytes in the ELISA panel were preferentially produced by IL-4/IL-13 treated cells.

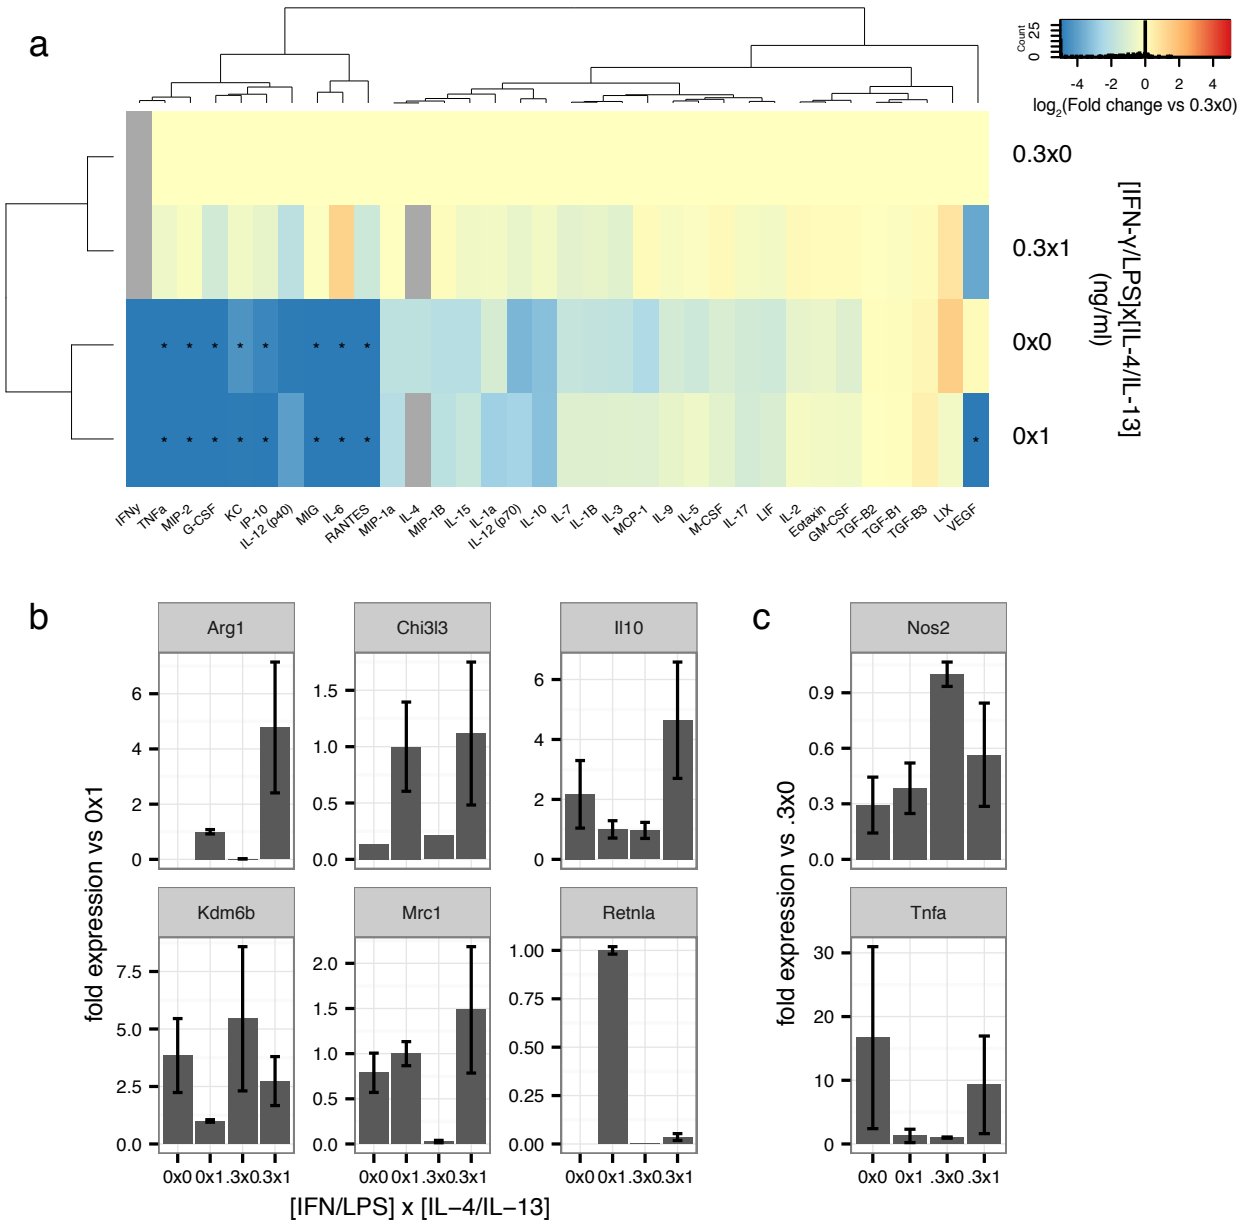

Figure 4.4: **Multiparametric characterization of macrophage phenotype.** (a) Dendrogram and heat map of cytokine release from macrophage cultures exposed to LPS/IFN-γ and/or IL-4/IL-13 for 24 hours. All conditions are normalized to inflammatory stimulus. $*$ indicates difference from LPS/IFN-γ-only condition by two-sided t test with Holm correction, p <0.05. Average of 3 experiments. (b) Mean $\log_2$ transformed fold difference mRNA expression vs. (B) IL-4/IL-13 only or (C) LPS/IFN-γ-only stimulation condition at 24 hours. Missing values indicate amplification below limit of detection. $n = 3$.

Over the course of 96 hours, we observed that most cytokines were either relatively stable

or exhibited decay (Supplementary Fig. S6.25). MCP-1 substantially increased in mixed conditions over this time period, which was not observed in any other condition (Supplementary Fig. S6.25). In addition, TGF-β increased in IL-4/IL-13-only and, to a lesser extent, mixed conditions, but not in LPS/IFN-γ-only or unstimulated conditions. Together, these data suggest that the time scale of M1 activation is shorter than that of M2 activation, and that cells stimulated with mixed cytokines progress towards a M2-like state. Moreover, expression of some M2 proteins and genes is enhanced under mixed conditions when compared to IL-4/IL-13-only stimulation conditions.

We also examined expression of genes associated with M2 (Fig. 4.4b) and M1 (Fig. 4.4c) activation by RT-qPCR at 24 h after stimulation. Transcript expression of *Cd206* as well as *Arg1* showed an increase in expression in mixed cytokine conditions compared to IL-4/IL-13-only conditions (Fig. 4.4b). However, expression of *Retnla* (Relmα/Fizz1) and *Chi3l3* (Ym1) was highest in the IL-4/IL-13-only condition, and co-addition of LPS/IFN-γ inhibited expression of these genes, suggesting that M2 markers may be heterogeneously expressed. Expression of Nos2 was highest in the LPS/IFN-γ only condition. *Tnfa* levels were lower in the LPS/IFN-γ condition compared to the unstimulated and mixed stimulus conditions at the examined timepoint (Fig. 4.4c), perhaps due to refractory downregulation after stimulation: *Tnfa* is coinduced with genes that degrade *Tnfa* transcripts, leading to a short transcript half-life [25], and activity at the *Tnfa* promoter stops by 18 hours after stimulation with LPS [88].

4.2.4 Macrophage state impacts reprogramming by a second activation signal

It is thought that macrophages in wound environments are plastic and can transition from M1-like to M2-like states as signals in their environment change [93] or are presented by

therapeutic materials [10]. To investigate this transition *in vitro*, we examined how pre-exposure of macrophages to an inflammatory stimulus influences their response to IL-4 and IL-13. Macrophages were treated with LPS/IFN-γ for 24 hours before IL-4/IL-13 were added for an additional 24 hours. Cells were assayed for CD86 and CD206 expression (Fig. 4.5a). We found that expression of CD86 in LPS/IFN-γ pre-treated macrophages showed a modest and non-statistically significant increase in response to subsequent addition of IL-4/IL-13 (Fig. 4.5b). In addition, pre-treatment of cells with LPS/IFN-γ did not block expression of CD206 upon IL-4/IL-13 stimulation, and thus cells pre-polarized to a M1 phenotype were still capable of acquiring characteristics of a M2 phenotype (Fig. 4.5c). Interestingly, at high concentrations of subsequent IL-4/IL-13 cytokine addition, the extent of CD206 expression was 30% ($\pm$ 9% SEM, $n = 5$) higher in LPS/IFN-γ pretreated cells when compared to unpretreated cells. These data demonstrate that pre-polarization towards an M1 phenotype with LPS/IFN-γ does not prevent subsequent M2 response to IL-4/IL-13, and in fact can enhance the expression of the M2 marker CD206.

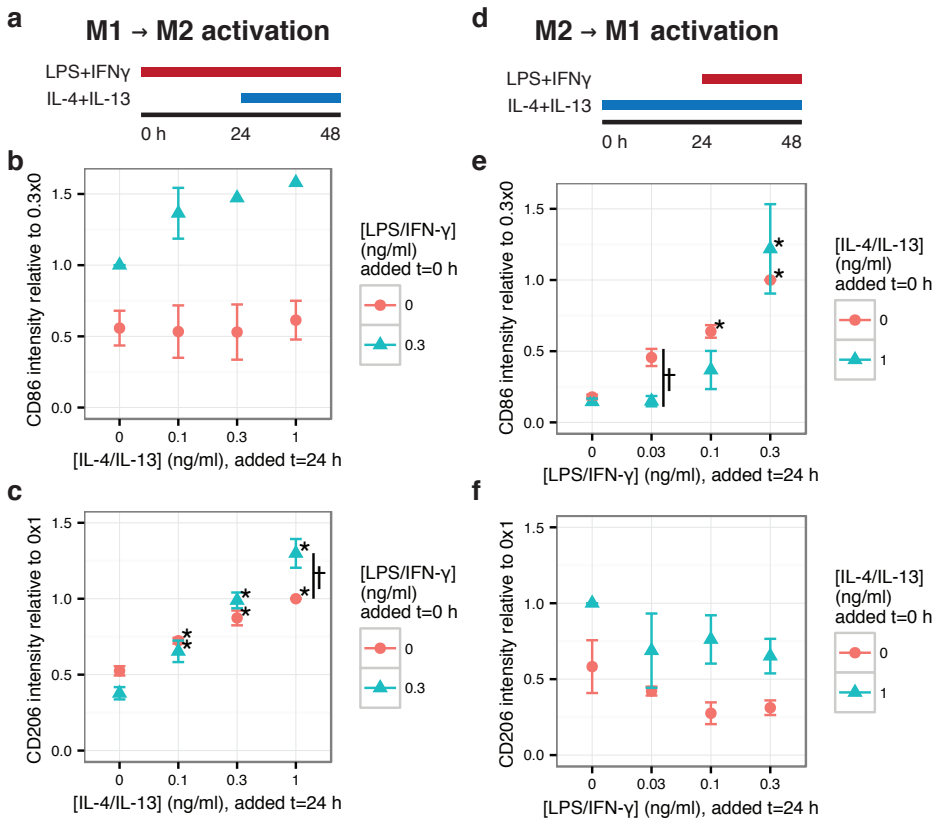

Figure 4.5: **Macrophage polarization state influences their plastic response to opposing activation signals**(a) Schematic illustrating experimental conditions. Macrophages were stimulated for 24 hours with inflammatory stimuli alone before anti-inflammatory stimuli were added for an additional 24 hours, and then assayed. (b) Expression of CD86 in cells in cells either pre-treated with 1 ng/ml LPS/IFN-γ or untreated for 24 hours and then subsequently treated with IL-4/IL-13 at the indicated dosages. Data are normalized to CD86 expression in 0.3 ng/ml LPS/IFN-γ-only condition, mean $\pm$ SEM ($n = 4$). (c) Expression of CD206 in the same conditions as B. Data are normalized to CD206 expression in 1 ng/ml IL-4/IL-13-only condition, mean $\pm$ SEM ($n = 4$). (d) Schematic illustrating experimental conditions. Macrophages were stimulated for 24 hours with anti-inflammatory stimuli alone before inflammatory stimuli were added for an additional 24 hours, and then assayed. (e) Expression of CD86 in cells either pre-treated with 1 ng/ml IL-4 and IL-13 or untreated for 24 hours and then subsequently treated with LPS and IFN-γ at the indicated dosages. Data are normalized to CD86 expression in 0.3 ng/ml LPS/IFN-γ-only condition, mean $\pm$ SEM ($n = 3$). (f) Expression of CD206 in the same conditions as E. Data are normalized to CD206 expression in 1 ng/ml IL-4/IL-13-only condition, mean $\pm$ SEM ($n = 3$). Asterisks indicate significant differences compared to unstimulated cells by two-sided t test, $p < 0.05$. Dagger indicates differences between groups with and without pre-treatment by two-sided t test, $p < 0.05$.

Conversely, perturbing M2-like macrophages towards a M1-like phenotype could be therapeuti-

cally useful in diseases including cancer, where reprogramming tumor-associated macrophages may help to inhibit tumor growth. To test M2-to-M1 plasticity *in vitro*, we pre-treated macrophages with IL-4/IL-13 for 24 hours, added LPS/IFN-γ for a subsequent 24 hours, and assayed cells for CD86 and CD206 expression (Fig. 4.5d). We found that at high concentrations of LPS/IFN-γ, pre-treatment with IL-4/IL-13 did not affect CD86 expression when compared to naïve cells, suggesting that M2 pre-polarization does not impact reprogramming towards an M1 phenotype. However, pre-treatment with IL-4/IL-13 inhibited the expression of CD86 upon exposure to low concentrations of LPS/IFN-γ (Fig. 4.5e). CD206 expression was modestly decreased in response to subsequent addition of LPS/IFN-γ in both the IL-4/IL-13 pre-treated and naïve conditions, although these differences were not statistically different (Fig. 4.5f). These data suggest that the ability of macrophages to acquire M1-like behavior after entering a M2-like state is dependent on the concentration of subsequent M1 stimulation, and that high concentrations may be necessary for reprogramming than for initial activation of naïve cells.

4.3 Discussion

We demonstrate that simultaneous exposure of macrophages to mixed cytokines leads to expression of both CD86 and CD206 in individual cells, which are established respective markers of M1 and M2 activation [80, 109, 100]. With 48 hours of stimulation, small quantities of co-added IL-4/IL-13 stimuli enhanced CD86 expression in LPS/IFN-γ-stimulated cells, whereas greater IL-4/IL-13 concentrations inhibited this enhancement. However, high concentrations of stimulus in the mixed condition induced expression of CD86 equivalent to expression in the LPS/IFN-γ-only condition, and expression of CD206 equivalent to expression in the IL-4/IL-13-only condition. These findings are likely dependent on the time point of observation and the concentrations of simuli, although an increase in M1 markers

in response to M2 stimuli has been previously described [96]. Individual cell expressions of CD86 and CD206 expression formed single-peaked, broad distributions, suggesting that individuals do not strongly polarize in mixed cytokine environments. Consistent with this, analysis of the secretome of these populations showed that the presence of IL-4/IL-13 along with LPS/IFN-γ only moderately dampens the level of inflammatory cytokine secretion when compared to LPS/IFN-γ only stimulated cells.

Although both M1 and M2 markers are present upon co-stimulation, their evolution over time is different. In cells stimulated in mixed conditions, the M1 marker CD86 decreases after the first 24 hours and returns almost to basal levels after 96 hours, whereas the M2 marker CD206 continues to increase peaking at 48-72 hours, and remained sustained relative to unstimulated macrophages even after 96 hours. This difference may be indicative of the natural progression of macrophages from inflammatory to anti-inflammatory phenotype during a host response to a wound or infection with pathogen [93]. Unexpectedly, macrophages exposed to mixed conditions had higher levels of CD206 when compared to cells treated with IL-4/IL-13 alone at the longer timepoints, suggesting that presence of an inflammatory stimulus may enhance the long term wound healing response. M2 marker *Arg1* transcripts measured by qRT-PCR were also elevated at each time point in the mixed stimulation condition compared to IL-4/IL-13 only (Supplementary Figure S6.27).

Our results suggest that macrophage reprogramming to a contrasting phenotype is dependent on initial polarization state and the strength of the second signal. For cells polarized towards an inflammatory phenotype with LPS/IFN-γ, CD206 expression with IL-4/IL-13 was enhanced compared to naïve cells. In contrast, cells polarized to an anti-inflammatory state with IL-4/IL-13 were more resistant to reprogramming with LPS/IFN-γ towards CD86 expression. This effect was observed specifically at the lower LPS/IFN-γ concentrations; expression of CD86 with high LPS-IFNγ concentrations was not significantly different between pre-polarized and naïve cells. These data suggest the anti-inflammatory phenotype is enhanced

by an initial inflammatory signal, and that macrophage progression from M1-like to M2-like phenotypes is favored.

While exploring signaling network topologies that could model macrophage responses, we discovered that the basic MISA motif was insufficient to account for the complex temporal expression patterns of CD86 and CD206, despite suggestions that elements of mutual inhibition and self activation may play a role in macrophage polarization [138, 80]. The mathematical models shed light on regulatory interactions which enable macrophages to achieve a spectrum of polarization states, depending quantitatively on microenvironmental cues. The models also suggest a regulatory logic by which individual cells co-stimulated by M1 and M2 signals can achieve transient M1 character followed by progression to a M2 phenotype. Although each node in the small-network models represents the combined action of many species, 'M1' and 'M2' likely reflect (at least in part) regulation by STAT1 and STAT6, respectively, consistent with their mutual antagonism induced by LPS/IFN-γ and IL-4/IL-13 [115, 154, 149, 116, 77, 147]. For construction of a parsimonious model in order to avoid overfitting, a single additional node ('Y') was introduced to mediate both the transient nature of CD86 expression, and the mixed inhibiting/activating effect of LPS-IFNγ on CD206. As such, 'Y' likely comprises feedback inhibition mechanisms, including those mediated by SOCS and STAT3 (reviewed in [65]). Furthermore, 'Y' may reflect regulation by NFκB, which is activated by LPS[91] and inhibited by IL-4 [171]. Several studies have suggested a cooperative interaction between NFκB and STAT6 to promote genes downstream of IL-4 [103, 4, 137, 52]. Incorporating these interactions into the mathematical model enabled us to construct small networks that captured the temporal response to both mixed and polarizing stimuli. A limitation of this approach is that the model is trained on two markers, which captures only some of the changes associated with macrophage activation. A more comprehensive dataset could lead to a more predictive model at the cost of increasing the complexity of the model.

These findings may have implications for therapeutic strategies involving macrophage repro-

gramming. For modulating the host response to biomaterial implants, delivery of IL-4 and IL-13 has been shown to increase expression of CD206 in surrounding macrophages [106]. Our results suggest that additional delivery of inflammatory cytokines, either concurrently or beforehand, may enhance anti-inflammatory activation and potentially improve the wound healing response. Along these lines, Spiller et. al. recently demonstrated that delivery of IFN-γ increases angiogenesis in response to a decellularized bone implant [143]. For cancer treatment, reprogramming tumor associated macrophages, which are thought to be anti-inflammatory, towards an inflammatory phenotype may require high concentrations of inflammatory stimuli, since low concentrations were not sufficient to induce this transition.

In summary, we find that macrophages exposed to both M1 and M2 activation signals express markers of both phenotypes, but the M1 markers decay over time while the M2 markers remain elevated. The distribution of markers suggest that macrophages do not exist in discrete polarized states. In addition, acquisition of the M2 phenotype appears to be enhanced by additional exposure to inflammatory stimulus, suggesting that inflammatory insult potentiates the wound healing response. Together, these results provide a better understanding of macrophage behavior in response to opposing activation signals, which is likely to be involved in the dynamic immune response to pathogens or injury. This improved understanding of macrophage activation will likely help design strategies for treatment of disease in which macrophages are involved.

4.4 Method

4.4.1 Cell Isolation and Culture

All protocols involving animals were approved by University of California Irvine's Institutional Animal Care and Use Committee, which is accredited by the Association for the Assessment

and Accreditation of Laboratory Animal Care International (AAALACi). Primary bone marrow derived macrophages were obtained by harvesting marrow from femurs of 6-12 week old female C57BL/6 mice, lysing red blood cells with ACK buffer, and then culturing cells for seven days on bacteriological polystyrene plates in DMEM supplemented with 10% FBS, 2% penicillin/streptomycin, 2 mM L-glutamine, and 10% conditioned media from CMG 12-14 cells expressing recombinant mouse M-CSF. Macrophages were treated with the indicated concentrations of LPS (Sigma), IFN-γ, IL-4, or IL-13 (all from Biolegend) for the indicated time. Macrophages were seeded at 3e5 cells/ml, allowed to adhere overnight, and treated with indicated concentrations of IL-4, IL-13, TNF-α, and IFN-γ, and then assayed for flow cytometry or cytokines as described below.

4.4.2 Flow Cytometry

Cells were fixed in 4% formaldehyde and stored at 4 °C until staining with anti-CD86 (clone GL-1, APC conjugate; Tonbo Biosciences) and anti-CD206 (clone C068C2, Alexa 488 conjugate; Biolegend) antibodies or isotype controls. Cells were analyzed on a BD LSR flow cytometer with post-processing in FlowJo (Tree Star). Cell populations were gated on forward and side scatter to select intact single cells. Events were acquired until 10,000 events were collected in a preliminary analysis gate or the sample was exhausted.

4.4.3 Cytokine Analysis

Supernatants were collected at 24, 48, 72, and 96 hours after stimulation and analyzed with a Luminex 31-plex mouse cytokine array (Eve Technologies). Hierarchical clustering was performed in R using a complete linkage method and presented with the gplots package [161].

4.4.4 RT-PCR

For gene expression analysis, cells were lifted from culture plates 24 hours after stimulation, pelleted by centrifugation, and frozen at −80 °C. Cells were lysed and RNA was extracted with the Qiagen RNeasy Mini kit. Reverse transcription was performed with the Qiagen Quantitect Reverse Transcription kit, which uses random priming and includes a DNase treatment. Resulting cDNA was observed to be free of contaminating gDNA by testing with the mVPA1 primer set [79]. qPCR was performed with BioRad SsoFast EvaGreen master mix on a BioRad CFX96 thermocycler using recommended cycling parameters (hot-start activation at 95 °C for 30 s, followed by 40 cycles of 5 s denaturation at 95 °C and 5 s annealing/extension at 55 °C, followed by melt curve collection from 65-95 °C in 0.5 °C increments at 5 s/step). Inhibition was cleared by diluting samples 1:100 in ddH_2O before analysis. Amplification was confirmed to be target-specific with Primer-BLAST [165] and by observing that melt curves had a single peak consistent with predicted amplicon melting temperature. Primer sequences and target and amplicon details are presented in Supplementary Table S6.9. Gene-of-interest expression was determined relative to an ensemble of Hprt, Gapdh, and Ldha expression using the GeNorm method [153] implemented by the package eleven (`https://github.com/tdsmith/eleven`).

4.4.5 Construction of Macrophage Models

We constructed mathematical models comprising minimal nonlinear Ordinary Differential Equation (ODE) networks. Our models were developed from previously described models of T cell subset specialization [97, 63, 12], adding additional connectivities and species to account for the complex kinetics of CD86 and CD206 after stimulation. We constructed 70 models with different topologies. A representative set of the studied topologies, models 1-6, are illustrated in Figure 4.2 and Fig. 4.6. A model described by the mutual-inhibition

self-activation (MISA) motif (Model 2) fits the data better than one in which the M1- and M2-associated pathways are activated independently (Model 1). However, the MISA is nevertheless insufficient to capture the complex kinetics of CD86 and CD206 expression after co-stimulation.

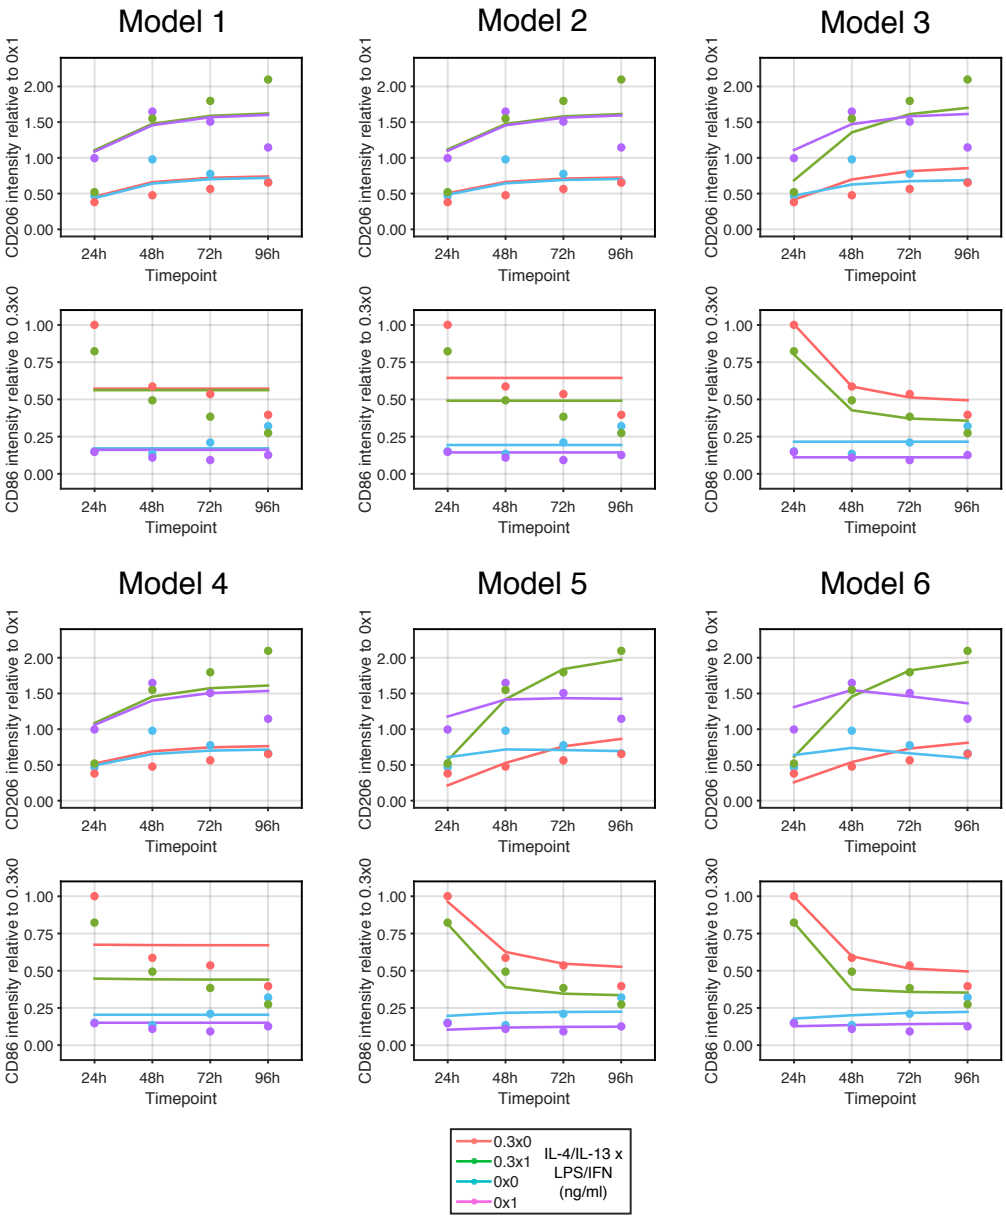

Figure 4.6: **Simulated timecourse experiment from each model shown in Fig. 4.4.** Points mark experimental data, as shown in Fig. 4.4B.

Models that extend the MISA topology by introducing a new species Y that interacts with

the M1- and M2-associated pathways, and representing unknown processes downstream of LPS/IFN and/or IL-4/IL-13 signaling events improve the model score (see Akaike Information Criterion (AICc) below) (Models 3, 5, 6). Specifically, we found that an incoherent feedforward loop on M1, mediated by an additional species Y, was necessary to capture the decay of CD86 expression in costimulated cells at 96 hours [9]. An activating link between Y and M2 was also consistent with the increased expression of CD206 of costimulated cells at later times. In particular, cooperative activation of M2 from Y and M2 improved the overall fit. To discover extended topologies, we were guided by the features of the temporal data and the literature on macrophage activation, as discussed in the main text.

Model quality was assessed based on optimization of parameters by fitting to the 96-hour time course data (Fig. 4.2) of four timepoints (24, 48, 72, 96 h) for four different stimulation conditions ({0.3,0}, {0.3,1}, {0,0}, {0,1} ng/ml) {LPS/IFN-γ, IL-4/IL-13}. The number of replicates was between three and five for each timepoint, giving 72 experimental data points. The error metric used was the sum of squared residuals (RSS) with normalized mean weighting. Parameter estimation was performed by minimizing the RSS of the model predicted CD86 and CD206 values to the normalized mean-weighted experimental values. The Matlab Optimization Toolbox and the trust-region-reflective algorithm were used to perform 1,000 individual fits. Parameters were initialized from a lognormal distribution with a mean and variance of 2, and were constrained to be positive. Parameters were optimized to the normalized timecourse data, and thus are expressed in arbitrary concentration and time units. Initial fits were performed using 400 trust-region-reflective iterations, or until convergence, using normalized unweighted experimental values. A second fit was then performed to the normalized mean-weighted experimental values. All models were assessed using the AICc, a scoring metric for model selection that includes penalties for increasing the number of fitted parameters [111]. Model parameters can be found in Tables S6.10 and S6.11.

To replicate the cell-to-cell variability in the flow cytometry data, individual cells were given

static parameters drawn from a distribution. Cell populations with between 3000 and 10000 cells were simulated, and model parameters for each cell were drawn from a lognormal distribution centered on the optimized parameters and with a variance of one percent of the mean. The resulting CD86 and CD206 expression levels for all models in Figure 4.2 and fitted parameters (Fig. 4.6 showed single-peaked distributions shifting with dosage, in qualitative agreement with the experimental density plots.

Chapter 5

Future Directions and Weighted Ensemble Sampling Methodology

5.1 Generalizability to Other Biochemical Frameworks

The WE-based computational pipeline presented in this work (Chapter 3) is uniquely suited to extracting global dynamics information for stochastic systems with metastability, using simulations. An advantage of this approach is that both the WE and coarse-graining algorithms are "dynamics-agnostic" [38], meaning that they can be applied to any type of stochastic dynamics framework. In the context of computational biology, our pipeline could be extended to other types of stochastic biochemical systems, such as systems with hybrid discrete-continuum dynamics [58], systems with spatial heterogeneity [39], or multi-level models [101]. In addition to this flexibility, simulation-based methods have the advantage of being able to leverage existing, widely-used open-source packages, which in turn facilitate model specification and model sharing. For example, BioNetGen [44] can interpret models specified in the Systems Biology Markup Language.

In ongoing work, we are studying models of cardiomyocyte differentiation [59], bacterial competence [110], DNA methylation maintenance [50], and receptor-ligand binding in a crowded environment [112] using WE-based sampling. While the WE-based computational pipeline and string method appear to work well in systems with well-defined attractor basins, even in systems with dynamical oscillations (pluripotency network, SC to LN(1/2) transitions), the coarse-graining method, PCCA+ [132], was ostensibly designed for systems that obey detailed balance, where, at equilibrium, each process is equilibrated with its reverse process. Bacterial competence and methylation maintenance involve cyclical oscillations, which can be observed and sampled using the WE framework, but these cycles explicitly do not satisfy detailed balance, potentially precluding the use of clustering based on separation of timescales. However, it is possible to construct a Markov model for a cyclic process [81], though further study is necessary to determine if such a model has biological relevance.

5.1.1 Alternatives to stochastic simulation algorithm

While the WE methods presented in this document use BioNetGen [44], other software to simulate trajectories of the SSA have been used, including StochKit [129] and the Stochastic Simulation Compiler (SSC) [87]. Development of methods to improve the computational efficiency and speed of SSA, e.g. tau-leaping [24], is still ongoing. Using a speed-up to the SSA could extend the WE sampling method to larger (> 10 genes) systems since, thanks to parallelization, the bottleneck of WE sampling is simulating the system dynamics. However, these methods are, by necessity, approximate, potentially leading to errors in small number regimes, such as transitions in gene state or changes in copy numbers close to zero.

5.2 Modification to Weighted Ensemble Binning Methods

Although the Weighted Ensemble method itself is generally efficient for studying systems with rare events [38], the method of adaptively binning is currently computationally inefficient. With every iteration, N_{bins} calculations of distance to M_{targ} replicas are necessary to fully repopulate the sampling region centers. Alternative adaptive binning methods, such as WExplore [33] can reducing the computational effort in exploring a large state space with a limited number of sampling regions by redistributing the distance calculations across multiple hierarchies of sampling regions. Alternative distance metrics, such as those covered in [90] might also provide alternative clustering methods for the identification of metastable states, and better sampling region coverage of phenotypes that are dynamically distinct, but close to each other in Euclidean distance, e.g. gene states.

5.3 Quantifying Parameter Dependence in Weighted Ensemble

To use the weighted ensemble method, we assume that the replicas in each sampling region are well-mixed and at equilibrium *within* that sampling region. The choice of sampling time τ_{WE} can have drastic effects on the efficiency and accuracy of WE sampling. We demonstrate the τ_{WE} dependence on an autocatalytic, trimolecular toy network, the Schlögl network with two attractor basins (Fig. 5.1):

$$2\,\mathrm{X} \underset{\mathrm{k_2}}{\overset{\mathrm{k_1}}{\rightleftharpoons}} 3\,\mathrm{X}$$
$$0 \underset{\mathrm{k_4}}{\overset{\mathrm{k_3}}{\rightleftharpoons}} \mathrm{X}$$

where $k_1 = 3 \times 10^{-2}$, $k_2 = 3.33 \times 10^{-5}$, $k_3 = 200$, and $k_4 = 3.5$. The mean first passage time calculated from a truncated enumeration of the CME was $MFPT_{AB} = 5.03 \times 10^4$, where the basins A and B are unit circles centered around $X = 82$ and $X = 583$, respectively. We find percent error of the flux calculated from the WE computational pipeline to that calculated through a long conventional simulation, using the same sampling time τ. Of the three simulations, $\tau_{WE} = 10,\ 50, \text{and} 100$, the simulation with $\tau_{WE} = 50$ has the smallest % error. However, the total amount of sampling time is different for all three simulations. Overall, it appears that the amount of fluctuation in the estimated flux decreases for $\tau_{WE} = 50$ and 100, and the simulations with longer τ appear to require fewer iterations to estimate a flux within 50% of the conventional simulation value. Unfortunately, the sensitivity of the simulation results to the WE parameters N_{bins} and τ_{WE} change with each network. A general 'rule-of-thumb' for choosing the sampling time τ_{WE} is that a replica should have an equal probability of leaving a sampling region and staying within it during the sampling time, meaning that the effect of N_{bins} and τ_{WE} on sampling accuracy and speed are highly correlated.

When we use weighted ensemble sampling for the transition matrix, we have another consideration with the sampling time τ. Estimation of timescales in the transition matrix is also dependent on sampling time $\tau_{trans.\ matrix}$ and the number of sampling regions, demonstrated on the ExMISA network (Fig. 5.2). In this case the weighted ensemble sampling time is the same as the lagtime used to estimate the transition matrix ($\tau_{WE} = \tau_{trans.\ matrix}$), and the Markovian assumption of the transition matrix is subject to the timescale of $\tau_{trans.\ matrix}$.

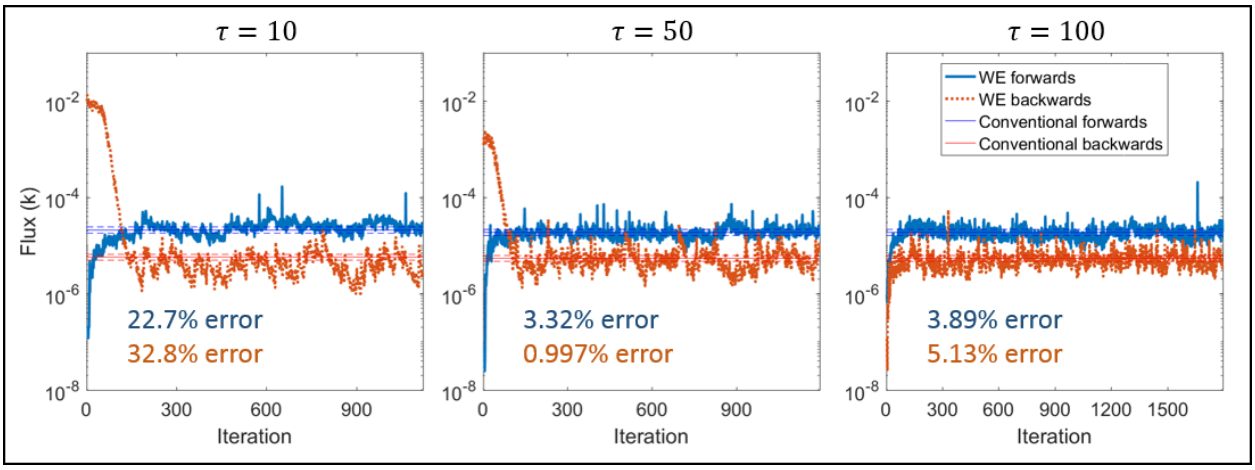

Figure 5.1: **Dependence of Weighted Ensemble Sampling Accuracy with Sampling Time τ, Demonstrated on the Schlögl Network** Forwards transition from basin A, defined as a unit circle centered around $X = 82$ to basin B, defined as a unit circle centered around $X = 583$. The A $\rightarrow$ B transition is shown in blue, with the mean of a long conventional simulation shown in light blue. The B $\rightarrow$ A transition is shown in orange. The percent error of the flux found through the computational pipeline using rate mode (averaged over the last 300 iterations) vs. flux found from a long conventional simulation (calculated with the same τ) is displayed for each sampling time.

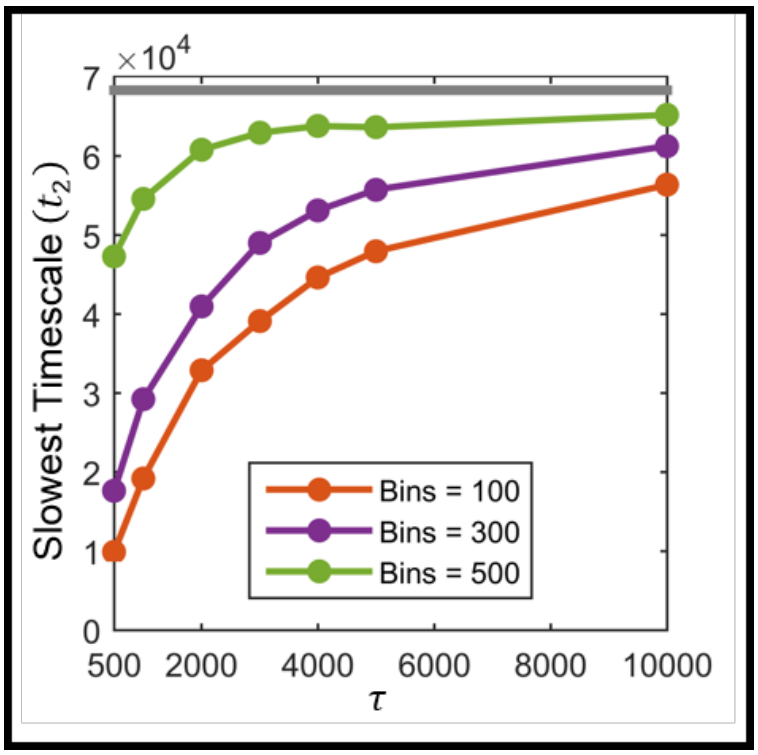

Figure 5.2: **Convergence of the slowest implied timescale t_2 with increasing number of sampling regions (bins) and increasing lagtime τ for the ExMISA network.** The lagtime calculated using the truncated CME is shown in gray. The accuracy of the WE approximation increases monotonically with increasing bin number and lagtime.

5.4 Model Inference in Macrophage Polarization

Previously, we used the normalized average expression level of CD86 and CD206 markers to select six candidate models for macrophage polarization. While our previous methods identified the necessary network topology to capture the mixed activation response (an incoherent feed-forward loop), we could not capture the memory effects of pre-polarization (Fig. 5.3), nor could we properly fit the distribution of the sixteen induction conditions with additive noise. We previously attempted to fit the distribution of the sixteen induction conditions with an SDE model with additive noise including cooperative mutual repression (Fig. 5.4, right). However, the model primarily fit the mean expression of CD86 and CD206

for each condition and had difficulty fitting the distribution for conditions with high M1 induction. We did not attempt to fit the larger models with incoherent feed-forward loops or with the additional species 'Y' due to the computational inefficiency of simulating 'synthetic' cell populations (approximately 10,000 cells) for all sixteen polarizing conditions. Attempts at fitting discrete models using Rule-Based simulation (BioNetGen) were largely unsuccessful for similar reasons.

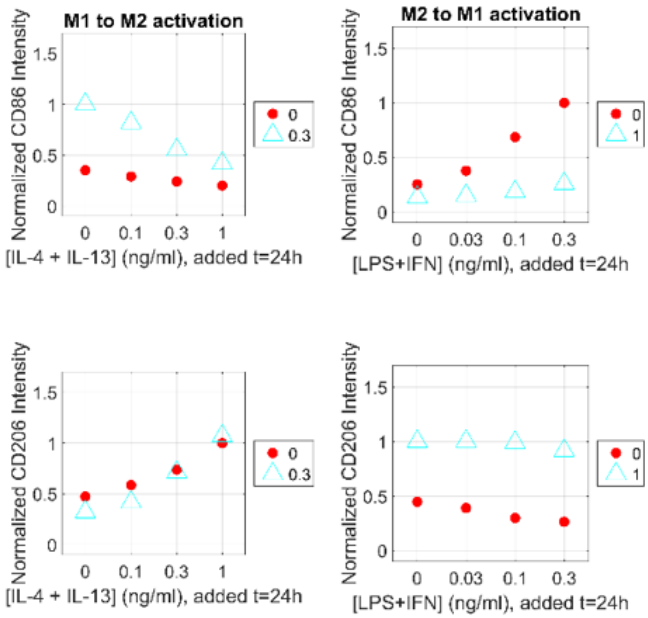

Figure 5.3: **Behavior of the model six to pre-polarization signals.** Although the model fits the time response data well according to the AICc goodness-of-fit measurement, it does not show the increased response in CD86 with M2 pre-polarization at higher doses of LPS/IFN-γ.

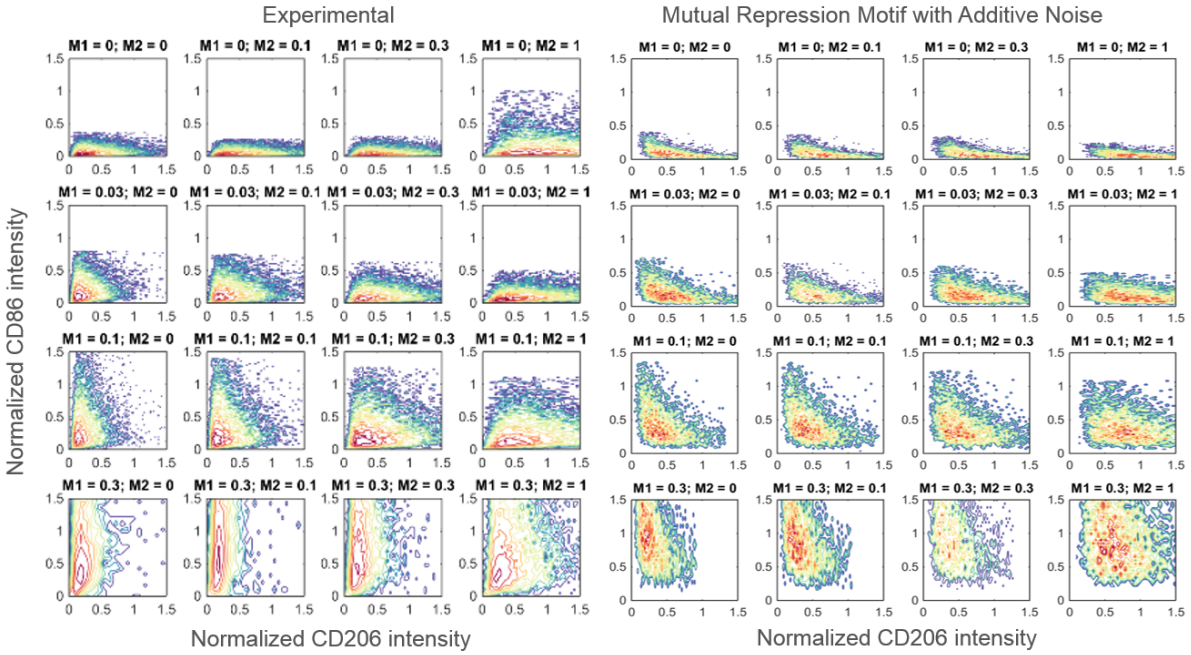

Figure 5.4: **Fit of Candidate Models to Sixteen Conditions of Macrophage Induction.** (Left) One representative replicate of the experimental flow cytometry results, normalized to the mean expression level of CD86 in the {0.3,0} condition and to the mean expression level of CD206 in the {0,1} condition. (Right) Mutual repression motif with cooperativity fit to the distribution of the sixteen experimental conditions. This model was fit to minimize the sum of the Kolmogorov-Smirnov test for all sixteen conditions, which fits the standard deviation of the distributions, but fails to capture the tails.

Instead of attempting to directly compare two noisy cell populations on largely varying scales (arbitrary units of fluorescence vs. discrete cell marker copy number), comparison of the probability distribution of CD86 and CD206 via Kullback–Leibler divergence or the Kolmogorov–Smirnov test could provide an easier value to minimize. Previously, to compress the two-dimensional response in CD86 and CD206 to a one-dimensional distribution, we converted the distribution of CD86 and CD206 expression to a single α distribution used to describe polarization in T-cells [12]. This metric measured how close the response of a cell population was to either polarized signal (e.g. where M1 induction would be associated with $\alpha = 0$ and M2 induction would be associaed with $\alpha = \pi/2$). Unfortunately, there was expression of CD86 under pure M2 induction, which precluded the use of the metric. However,

there are metrics that can measure distance from two-dimensional distributions, such as the Mahalanobis distance. A two-dimensional distribution of CD86 and CD206 expression can be estimated through the WE-based computational pipeline, which can be directly compared to the probability distribution of fluorescence data (for each triplicate). Introducing fully discrete chemical reaction network models to the candidate model set could lead to the inference of a network model that fully captures macrophage polarization dynamics, from the time-dependent dose-response, to the pre-polarization memory effects. Alternatively, to fit the general macrophage polarization response, iterative fitting procedures that first fit a model to the timecourse data, then to the sixteen induction conditions or the pre-polarization signals might improve the current model selection protocol.

Bibliography

[1] MATLAB and Parallel Computing Toolbox Release 2012b.

[2] E. Abranches, E. Bekman, and D. Henrique. Generation and Characterization of a Novel Mouse Embryonic Stem Cell Line with a Dynamic Reporter of Nanog Expression. *PLOS ONE*, 8(3):e59928, Mar. 2013.

[3] E. Abranches, A. M. V. Guedes, M. Moravec, H. Maamar, P. Svoboda, A. Raj, and D. Henrique. Stochastic NANOG fluctuations allow mouse embryonic stem cells to explore pluripotency. *Development*, 141(14):2770–2779, July 2014.

[4] Y. Abu-Amer. Il-4 abrogates osteoclastogenesis through stat6-dependent inhibition of nf-κb. *Journal of Clinical Investigation*, 107(11):1375, 2001.

[5] M. Acar, J. T. Mettetal, and A. van Oudenaarden. Stochastic switching as a survival strategy in fluctuating environments. *Nature Genetics*, 40(4):471–475, Apr. 2008.

[6] J. Adelman and M. Grabe. Simulating rare events using a weighted ensemble-based string method. *The Journal of chemical physics*, 138(23387566):044105, Jan. 2013.

[7] R. Allen, C. Valeriani, and P. ten Wolde. Forward flux sampling for rare event simulations. *Journal Of Physics-Condensed Matter*, 21(000271268400004):463102, Jan. 2009.

[8] R. J. Allen, P. B. Warren, and P. R. ten Wolde. Sampling Rare Switching Events in Biochemical Networks. *Physical Review Letters*, 94(1), Jan. 2005.

[9] U. Alon. *An introduction to systems biology: design principles of biological circuits*. CRC press, 2006.

[10] M. M. Alvarez, J. C. Liu, G. Trujillo-de Santiago, B.-H. Cha, A. Vishwakarma, A. M. Ghaemmaghami, and A. Khademhosseini. Delivery strategies to control inflammatory response: Modulating m1–m2 polarization in tissue engineering applications. *Journal of Controlled Release*, 240:349–363, 2016.

[11] J. M. Anderson, A. Rodriguez, and D. T. Chang. Foreign body reaction to biomaterials. In *Seminars in immunology*, volume 20, pages 86–100. Elsevier, 2008.

[12] Y. E. Antebi, S. Reich-Zeliger, Y. Hart, A. Mayo, I. Eizenberg, J. Rimer, P. Putheti, D. Pe'er, and N. Friedman. Mapping differentiation under mixed culture conditions reveals a tunable continuum of t cell fates. *PLoS biology*, 11(7):e1001616, 2013.

[13] A. Arkin, J. Ross, and H. McAdams. Stochastic kinetic analysis of developmental pathway bifurcation in phage lambda-infected Escherichia coli cells. *Genetics*, 149(000075232100002):1633–1648, Dec. 1997.

[14] M. N. Artyomov, M. Mathur, M. S. Samoilov, and A. K. Chakraborty. Stochastic bimodalities in deterministically monostable reversible chemical networks due to network topology reduction. *The Journal of Chemical Physics*, 131(19):195103, Nov. 2009.

[15] S. S. Ashwin and M. Sasai. Epigenetic Dynamics of Cell Reprogramming. *arXiv:1410.2337 [cond-mat, q-bio]*, Oct. 2014. arXiv: 1410.2337.

[16] E. Aurell and K. Sneppen. Epigenetics as a First Exit Problem. *Physical Review Letters*, 88(4), Jan. 2002.

[17] N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler. Bacterial Persistence as a Phenotypic Switch. *Science*, 305(5690):1622–1625, Sept. 2004.

[18] M. Bastian, S. Heymann, and M. Jacomy. Gephi: An Open Source Software for Exploring and Manipulating Networks. 2009.

[19] D. A. Beard and H. Qian. *Chemical Biophysics: Quantitative Analysis of Cellular Systems*. Cambridge University Press, May 2008. Google-Books-ID: kTmQAz6QXckC.

[20] S. Bhattacharya, Q. Zhang, and M. E. Andersen. A deterministic map of Waddington's epigenetic landscape for cell fate specification. *BMC Systems Biology*, 5:85, May 2011.

[21] S. K. Biswas and A. Mantovani. Macrophage plasticity and interaction with lymphocyte subsets: cancer as a paradigm. *Nature immunology*, 11(10):889–896, 2010.

[22] E. E. Borrero and F. A. Escobedo. Reaction coordinates and transition pathways of rare events via forward flux sampling. *The Journal of Chemical Physics*, 127(16):164101, Oct. 2007.

[23] Y. Buganim, D. A. Faddah, A. W. Cheng, E. Itskovich, S. Markoulaki, K. Ganz, S. L. Klemm, A. van Oudenaarden, and R. Jaenisch. Single-Cell Expression Analyses during Cellular Reprogramming Reveal an Early Stochastic and a Late Hierarchic Phase. *Cell*, 150(6):1209–1222, Sept. 2012.

[24] Y. Cao, D. T. Gillespie, and L. R. Petzold. Adaptive explicit-implicit tau-leaping method with automatic tau selection. *The Journal of chemical physics*, 126(22):224101, 2007.

[25] E. Carballo, W. S. Lai, and P. J. Blackshear. Feedback inhibition of macrophage tumor necrosis factor-α production by tristetraprolin. *Science*, 281(5379):1001–1005, 1998.

[26] I. Chambers, J. Silva, D. Colby, J. Nichols, B. Nijmeijer, M. Robertson, J. Vrana, K. Jones, L. Grotewold, and A. Smith. Nanog safeguards pluripotency and mediates germline development. *Nature*, 450(7173):1230–1234, Dec. 2007.

[27] H. H. Chang, M. Hemberg, M. Barahona, D. E. Ingber, and S. Huang. Transcriptome-wide noise controls lineage choice in mammalian progenitor cells. *Nature*, 453(7194):544–547, May 2008.

[28] R. Chang, R. Shoemaker, and W. Wang. Systematic Search for Recipes to Generate Induced Pluripotent Stem Cells. *PLoS Computational Biology*, 7(12), Dec. 2011.

[29] P. J. Choi, L. Cai, K. Frieda, and X. S. Xie. A Stochastic Single-Molecule Event Triggers Phenotype Switching of a Bacterial Cell. *Science*, 322(5900):442–446, Oct. 2008.

[30] B. K. Chu, M. J. Tse, R. R. Sato, and E. L. Read. Markov State Models of gene regulatory networks. *BMC Systems Biology*, 11:14, Feb. 2017.

[31] P. G. Constantinides, P. A. Jones, and W. Gevers. Functional striated muscle cells from non-myoblast precursors following 5-azacytidine treatment. *Nature*, 267(5609):364–6, May 1977.

[32] M. J. Davis, T. M. Tsang, Y. Qiu, J. K. Dayrit, J. B. Freij, G. B. Huffnagle, and M. A. Olszewski. Macrophage m1/m2 polarization dynamically adapts to changes in cytokine microenvironments in cryptococcus neoformans infection. *MBio*, 4(3):e00264–13, 2013.

[33] A. Dickson and C. L. Brooks. WExplore: Hierarchical Exploration of High-Dimensional Spaces Using the Weighted Ensemble Algorithm. *The Journal of Physical Chemistry B*, 118(13):3532–3542, Apr. 2014.

[34] A. Dickson and A. R. Dinner. Enhanced Sampling of Nonequilibrium Steady States. *Annual Review of Physical Chemistry*, 61(1):441–459, 2010.

[35] A. Dickson, A. Warmflash, and A. R. Dinner. Nonequilibrium umbrella sampling in spaces of many order parameters. *The Journal of Chemical Physics*, 130(7):074104, Feb. 2009.

[36] A. Dickson, A. Warmflash, and A. R. Dinner. Separating forward and backward pathways in nonequilibrium umbrella sampling. *The Journal of Chemical Physics*, 131(15):154104, Oct. 2009.

[37] J.-E. Dietrich and T. Hiiragi. Stochastic patterning in the mouse pre-implantation embryo. *Development*, 134(23):4219–4231, Dec. 2007.

[38] R. Donovan, A. Sedgewick, J. Faeder, and D. Zuckerman. Efficient stochastic simulation of chemical kinetics networks using a weighted ensemble of trajectories. *The Journal of chemical physics*, 139(24070313):115105, Sept. 2013.

[39] R. M. Donovan, J.-J. Tapia, D. P. Sullivan, J. R. Faeder, R. F. Murphy, M. Dittrich, and D. M. Zuckerman. Unbiased Rare Event Sampling in Spatial Stochastic Systems Biology Models Using a Weighted Ensemble of Trajectories. *PLoS Computational Biology*, 12(2), Feb. 2016.

[40] W. E, W. Ren, and E. Vanden-Eijnden. String method for the study of rare events. *Physical Review B*, 66(5):052301, Aug. 2002.

[41] W. E, W. Ren, and E. Vanden-Eijnden. Transition pathways in complex systems: Reaction coordinates, isocommittor surfaces, and transition tubes. *Chemical Physics Letters*, 413(1):242–247, Sept. 2005.

[42] M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic Gene Expression in a Single Cell. *Science*, 297(5584):1183–1186, Aug. 2002.

[43] L.-P. Erwig, D. C. Kluth, G. M. Walsh, and A. J. Rees. Initial cytokine exposure determines function of macrophages and renders them unresponsive to other cytokines. *The Journal of Immunology*, 161(4):1983–1988, 1998.

[44] J. R. Faeder, M. L. Blinov, and W. S. Hlavacek. Rule-Based Modeling of Biochemical Systems with BioNetGen. In *Systems Biology*, Methods in Molecular Biology, pages 113–167. Humana Press, 2009. DOI: 10.1007/978-1-59745-525-1_5.

[45] H. Feng, B. Han, and J. Wang. Landscape and Global Stability of Nonadiabatic and Adiabatic Oscillations in a Gene Network. *Biophysical Journal*, 102(5):1001–1010, Mar. 2012.

[46] H. Feng and J. Wang. A new mechanism of stem cell differentiation through slow binding/unbinding of regulators to genes. *Sci Rep*, 2:550, Jan 2012.

[47] A. Filipczyk, C. Marr, S. Hastreiter, J. Feigelman, M. Schwarzfischer, P. S. Hoppe, D. Loeffler, K. D. Kokkaliaris, M. Endele, B. Schauberger, O. Hilsenbeck, S. Skylaki, J. Hasenauer, K. Anastassiadis, F. J. Theis, and T. Schroeder. Network plasticity of pluripotency transcription factors in embryonic stem cells. *Nature Cell Biology*, 17(10):1235–1246, Oct. 2015.

[48] T. S. Gardner, C. R. Cantor, and J. J. Collins. Construction of a genetic toggle switch in *Escherichia coli*. *Nature*, 403(6767):339, Jan. 2000.

[49] H. Ge, H. Qian, and X. S. Xie. Stochastic Phenotype Transition of a Single Cell in an Intermediate Region of Gene State Switching. *Physical Review Letters*, 114(7):078101, Feb. 2015.

[50] P. Giehr, C. Kyriakopoulos, G. Ficz, V. Wolf, and J. Walter. The influence of hydroxylation on maintaining cpg methylation patterns: a hidden markov model approach. *PLoS computational biology*, 12(5):e1004905, 2016.

[51] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, Dec. 1977.

[52] S. Goenka and M. H. Kaplan. Transcriptional regulation by stat6. *Immunologic research*, 50(1):87, 2011.

[53] T. Graf and T. Enver. Forcing cells to change lineages. *Nature*, 462(7273):587–594, Dec. 2009.

[54] P. B. Gupta, C. M. Fillmore, G. Jiang, S. D. Shapira, K. Tao, C. Kuperwasser, and E. S. Lander. Stochastic State Transitions Give Rise to Phenotypic Equilibrium in Populations of Cancer Cells. *Cell*, 146(4):633–644, Aug. 2011.

[55] M. P. Harrigan, M. M. Sultan, C. X. Hernández, B. E. Husic, P. Eastman, C. R. Schwantes, K. A. Beauchamp, R. T. McGibbon, and V. S. Pande. MSMBuilder: Statistical Models for Biomolecular Dynamics. *Biophysical Journal*, 112(1):10–15, Jan. 2017.

[56] J. Hasty, J. Pradines, M. Dolnik, and J. J. Collins. Noise-based switches and amplifiers for gene expression. *Proceedings of the National Academy of Sciences*, 97(5):2075–2080, Feb. 2000.

[57] D. C. Hay, L. Sutherland, J. Clark, and T. Burdon. Oct-4 Knockdown Induces Similar Patterns of Endoderm and Trophoblast Differentiation Markers in Human and Mouse Embryonic Stem Cells. *STEM CELLS*, 22(2):225–235, Mar. 2004.

[58] B. Hepp, A. Gupta, and M. Khammash. Adaptive hybrid simulations for multiscale stochastic reaction networks. *The Journal of Chemical Physics*, 142(3):034118, Jan. 2015.

[59] F. Herrmann, A. Groß, D. Zhou, H. A. Kestler, and M. Kühl. A boolean model of the cardiac gene regulatory network determining first and second heart field identity. *PloS one*, 7(10):e46798, 2012.

[60] P. G. Hoel, S. C. Port, and C. J. Stone. *Introduction to Stochastic Processes*. Waveland Press, Dec. 1986. Google-Books-ID: vfUWAAAAQBAJ.

[61] T. Hong, J. Xing, L. Li, and J. J. Tyson. A Mathematical Model for the Reciprocal Differentiation of T Helper 17 Cells and Induced Regulatory T Cells. *PLOS Computational Biology*, 7(7):e1002122, July 2011.

[62] T. Hong, J. Xing, L. Li, and J. J. Tyson. A simple theoretical framework for understanding heterogeneous differentiation of cd4+ t cells. *BMC systems biology*, 6(1):66, 2012.

[63] T. Hong, J. Xing, L. Li, and J. J. Tyson. A simple theoretical framework for understanding heterogeneous differentiation of CD4+ T cells. *BMC Systems Biology*, 6:66, 2012.

[64] S. Hormoz, Z. S. Singer, J. M. Linton, Y. E. Antebi, B. I. Shraiman, and M. B. Elowitz. Inferring Cell-State Transition Dynamics from Lineage Trees and Endpoint Single-Cell Measurements. *Cell Systems*, 3(5):419–433.e8, Nov. 2016.

[65] X. Hu, S. D. Chakravarty, and L. B. Ivashkiv. Regulation of interferon and toll-like receptor signaling during macrophage activation by opposing feedforward and feedback inhibition mechanisms. *Immunological reviews*, 226(1):41–56, 2008.

[66] S. Huang. Reprogramming cell fates: reconciling rarity with robustness. *BioEssays: News and Reviews in Molecular, Cellular and Developmental Biology*, 31(5):546–560, May 2009.

[67] S. Huang. The molecular and mathematical basis of Waddington's epigenetic landscape: a framework for post-Darwinian biology? *BioEssays: News and Reviews in Molecular, Cellular and Developmental Biology*, 34(2):149–157, Feb. 2012.

[68] S. Huang, G. Eichler, Y. Bar-Yam, and D. E. Ingber. Cell fates as high-dimensional attractor states of a complex gene regulatory network. *Physical Review Letters*, 94(12):128701, Apr. 2005.

[69] S. Huang and S. Kauffman. How to escape the cancer attractor: rationale and limitations of multi-target drugs. *Seminars in Cancer Biology*, 23(4):270–278, Aug. 2013.

[70] G. Huber and S. Kim. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophysical Journal*, 70(8770190):97–110, Jan. 1996.

[71] D. A. Hume. Probability in transcriptional regulation and its implications for leukocyte differentiation and inducible gene expression. *Blood*, 96(7):2323–2328, Oct. 2000.

[72] T. Kalmar, C. Lim, P. Hayward, S. Muñoz-Descalzo, J. Nichols, J. Garcia-Ojalvo, and A. Martinez Arias. Regulated Fluctuations in Nanog Expression Mediate Cell Fate Decisions in Embryonic Stem Cells. *PLoS Biology*, 7(7):e1000149, July 2009.

[73] N. G. V. Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, Nov. 1992. Google-Books-ID: 3e7XbMoJzmoC.

[74] S. Kauffman. Control circuits for determination and transdetermination. *Science*, 181(4198229):310–8, July 1973.

[75] T. B. Kepler and T. C. Elston. Stochasticity in Transcriptional Regulation: Origins, Consequences, and Mathematical Representations. *Biophysical Journal*, 81(6):3116–3136, Dec. 2001.

[76] J. Khallou-Laschet, A. Varthaman, G. Fornasa, C. Compain, A.-T. Gaston, M. Clement, M. Dussiot, O. Levillain, S. Graff-Dubois, A. Nicoletti, et al. Macrophage plasticity in experimental atherosclerosis. *PloS one*, 5(1):e8852, 2010.

[77] A. S. Kristof, J. Marks-Konczalik, E. Billings, and J. Moss. Stimulation of signal transducer and activator of transcription-1 (stat1)-dependent gene transcription by lipopolysaccharide and interferon-γ is regulated by mammalian target of rapamycin. *Journal of Biological Chemistry*, 278(36):33637–33644, 2003.

[78] M. Kærn, T. C. Elston, W. J. Blake, and J. J. Collins. Stochasticity in gene expression: from theories to phenotypes. *Nature Reviews Genetics*, 6(6):451–464, June 2005.

[79] H. Laurell, J. S. Iacovoni, A. Abot, D. Svec, J.-J. Maoret, J.-F. Arnal, and M. Kubista. Correction of rt–qpcr data for genomic dna-derived signals with validprime. *Nucleic acids research*, 40(7):e51–e51, 2012.

[80] T. Lawrence and G. Natoli. Transcriptional regulation of macrophage polarization: enabling diversity with identity. *Nature reviews immunology*, 11(11):750–761, 2011.

[81] J. Lee. The derivation of markov processes without detailed balance. *arXiv preprint arXiv:1708.00184*, 2017.

[82] S. Legewie, H. Herzel, H. V. Westerhoff, and N. Blüthgen. Recurrent design patterns in the feedback regulation of the mammalian signalling network. *Molecular systems biology*, 4(1):190, 2008.

[83] C. Li and J. Wang. Quantifying Waddington landscapes and paths of non-adiabatic cell fate decisions for differentiation, reprogramming and transdifferentiation. *Journal of The Royal Society Interface*, 10(89):20130787, Dec. 2013.

[84] C. Li and J. Wang. Quantifying the Landscape for Development and Cancer from a Core Cancer Stem Cell Circuit. *Cancer Research*, 75(13):2607–2618, July 2015.

[85] V. Likhoshvai and A. Ratushny. Generalized hill function method for modeling molecular processes. *Journal of bioinformatics and computational biology*, 5(02b):521–531, 2007.

[86] A. Lipshtat, A. Loinger, N. Q. Balaban, and O. Biham. Genetic Toggle Switch without Cooperative Binding. *Physical Review Letters*, 96(18):188101, May 2006.

[87] M. Lis, M. N. Artyomov, S. Devadas, and A. K. Chakraborty. Efficient stochastic simulation of reaction–diffusion processes via direct compilation. *Bioinformatics*, 25(17):2289–2291, 2009.

[88] H. Liu, P. Sidiropoulos, G. Song, L. J. Pagliari, M. J. Birrer, B. Stein, J. Anrather, and R. M. Pope. Tnf-α gene expression in macrophages: regulation by nf-κb is independent of c-jun or c/ebpβ. *The Journal of Immunology*, 164(8):4277–4285, 2000.

[89] M. Liu, F. Luo, C. Ding, S. Albeituni, X. Hu, Y. Ma, Y. Cai, L. McNally, M. A. Sanders, D. Jain, et al. Dectin-1 activation by a natural product β-glucan converts immunosuppressive macrophages into an m1-like phenotype. *The Journal of Immunology*, 195(10):5055–5065, 2015.

[90] J. A. Long and T. A. Nelson. A review of quantitative methods for movement data. *International Journal of Geographical Information Science*, 27(2):292–318, 2013.

[91] C. J. Lowenstein and E. Padalko. inos (nos2) at a glance. *Journal of cell science*, 117(14):2865–2867, 2004.

[92] M. Lu, J. Onuchic, and E. Ben-Jacob. Construction of A Self-consistent Landscape for Multistable Gene Regulatory Circuits. Mar. 2014.

[93] T. Lucas, A. Waisman, R. Ranjan, J. Roes, T. Krieg, W. Müller, A. Roers, and S. A. Eming. Differential roles of macrophages in diverse phases of skin repair. *The Journal of Immunology*, 184(7):3964–3977, 2010.

[94] A. Ma and A. R. Dinner. Automatic Method for Identifying Reaction Coordinates in Complex Systems. *The Journal of Physical Chemistry B*, 109(14):6769–6779, Apr. 2005.

[95] R. Ma, J. Wang, Z. Hou, and H. Liu. Small-Number Effects: A Third Stable State in a Genetic Bistable Toggle Switch. *Physical Review Letters*, 109(24):248107, Dec. 2012.

[96] I. Malyshev and Y. Malyshev. Current concept and update of the macrophage plasticity concept: intracellular mechanisms of reprogramming and m3 macrophage "switch" phenotype. *BioMed research international*, 2015, 2015.

[97] L. Mariani, M. Löhning, A. Radbruch, and T. Höfer. Transcriptional control networks of cell differentiation: insights from helper t lymphocytes. *Progress in biophysics and molecular biology*, 86(1):45–76, 2004.

[98] L. Mariani, E. G. Schulz, M. H. Lexberg, C. Helmstetter, A. Radbruch, M. Löhning, and T. Höfer. Short-term memory in gene induction reveals the regulatory principle behind stochastic IL-4 expression. *Molecular Systems Biology*, 6:359, Apr. 2010.

[99] F. O. Martinez and S. Gordon. The m1 and m2 paradigm of macrophage activation: time for reassessment. *F1000prime reports*, 6, 2014.

[100] F. O. Martinez, L. Helming, R. Milde, A. Varin, B. N. Melgert, C. Draijer, B. Thomas, M. Fabbri, A. Crawshaw, L. P. Ho, et al. Genetic programs expressed in resting and il-4 alternatively activated mouse and human macrophages: similarities and differences. *Blood*, 121(9):e57–e69, 2013.

[101] C. Maus, S. Rybacki, and A. M. Uhrmacher. Rule-based multi-level modeling of cell biological systems. *BMC Systems Biology*, 5:166, Oct. 2011.

[102] H. J. Medbury, H. Williams, and J. P. Fletcher. Clinical significance of macrophage phenotypes in cardiovascular disease. *Clinical and translational medicine*, 3(1):63, 2014.

[103] B. Messner, A. Stütz, B. Albrecht, S. Peiritsch, and M. Woisetschläger. Cooperation of binding sites for stat6 and nf kappa b/rel in the il-4-induced up-regulation of the human ige germline promoter. *The Journal of Immunology*, 159(7):3330–3337, 1997.

[104] K. Miller-Jensen, S. S. Dey, D. V. Schaffer, and A. P. Arkin. Varying virulence: epigenetic control of expression noise and disease processes. *Trends in Biotechnology*, 29(10):517–525, Oct. 2011.

[105] C. D. Mills, K. Kincaid, J. M. Alt, M. J. Heilman, and A. M. Hill. M-1/m-2 macrophages and the th1/th2 paradigm. *The Journal of Immunology*, 164(12):6166–6173, 2000.

[106] N. Mokarram, A. Merchant, V. Mukhatyar, G. Patel, and R. V. Bellamkonda. Effect of modulating macrophage phenotype on peripheral nerve repair. *Biomaterials*, 33(34):8793–8801, 2012.

[107] K. J. Moore, F. J. Sheedy, and E. A. Fisher. Macrophages in atherosclerosis: a dynamic balance. *Nature Reviews Immunology*, 13(10):709–721, 2013.

[108] M. J. Morelli, S. Tănase-Nicola, R. J. Allen, and P. R. ten Wolde. Reaction Coordinates for the Flipping of Genetic Switches. *Biophysical Journal*, 94(9):3413–3423, May 2008.

[109] D. M. Mosser. The many faces of macrophage activation. *Journal of leukocyte biology*, 73(2):209–212, 2003.

[110] A. Mugler, M. Kittisopikul, L. Hayden, J. Liu, C. H. Wiggins, G. M. Süel, and A. M. Walczak. Noise expands the response range of the bacillus subtilis competence circuit. *PLoS computational biology*, 12(3):e1004793, 2016.

[111] J. A. Nelder and R. Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.

[112] J. Newby and J. Allard. First-passage time to clear the way for receptor-ligand binding in a crowded environment. *Physical review letters*, 116(12):128101, 2016.

[113] H. Niwa, Y. Toyooka, D. Shimosato, D. Strumpf, K. Takahashi, R. Yagi, and J. Rossant. Interaction between Oct3/4 and Cdx2 Determines Trophectoderm Differentiation. *Cell*, 123(5):917–929, Dec. 2005.

[114] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proceedings of the National Academy of Sciences*, 106(45):19011–19016, Nov. 2009.

[115] Y. Ohmori and T. A. Hamilton. Il-4-induced stat6 suppresses ifn-gamma-stimulated stat1-dependent transcription in mouse macrophages. *The Journal of immunology*, 159(11):5474–5482, 1997.

[116] Y. Ohmori and T. A. Hamilton. Interleukin-4/stat6 represses stat1 and nf-κb-dependent transcription through distinct mechanisms. *Journal of Biological Chemistry*, 275(48):38095–38103, 2000.

[117] Y. Ohnishi, W. Huber, A. Tsumura, M. Kang, P. Xenopoulos, K. Kurimoto, A. K. Oleś, M. J. Araúzo-Bravo, M. Saitou, A.-K. Hadjantonakis, and T. Hiiragi. Cell-to-cell expression variability followed by signal reinforcement progressively segregates early mouse lineages. *Nature cell biology*, 16(1):27–37, Jan. 2014.

[118] A. Pakes and P. McGuire. Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse'of dimensionality. *Econometrica*, 69(5):1261–1281, 2001.

[119] P. Pelegrin and A. Surprenant. Dynamics of macrophage polarization reveal new mechanism to inhibit il-1β release through pyrophosphates. *The EMBO journal*, 28(14):2114–2127, 2009.

[120] J. S. Pettersen, J. Fuentes-Duculan, M. Suárez-Fariñas, K. C. Pierson, A. Pitts-Kiefer, L. Fan, D. A. Belkin, C. Q. Wang, S. Bhuvanendran, L. M. Johnson-Huang, et al. Tumor-associated macrophages in the cutaneous scc microenvironment are heterogeneously activated. *Journal of Investigative Dermatology*, 131(6):1322–1330, 2011.

[121] F. Porcheray, S. Viaud, A.-C. Rimaniol, C. Leone, B. Samah, N. Dereuddre-Bosquet, D. Dormont, and G. Gras. Macrophage activation switching: an asset for the resolution of inflammation. *Clinical & Experimental Immunology*, 142(3):481–489, 2005.

[122] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé. Markov models of molecular kinetics: Generation and validation. *The Journal of Chemical Physics*, 134(17):174105, May 2011.

[123] M. Ptashne. *A genetic switch: phage lambda revisited*. CSHL press, 2004.

[124] H. Qian and L. M. Bishop. The chemical master equation approach to nonequilibrium steady-state of open biochemical systems: linear single-molecule enzyme kinetics and nonlinear biochemical reaction networks. *International journal of molecular sciences*, 11(9):3472–3500, 2010.

[125] A. Raj and A. van Oudenaarden. Nature, Nurture, or Chance: Stochastic Gene Expression and Its Consequences. *Cell*, 135(2):216–226, Oct. 2008.

[126] M. A. Rieger, P. S. Hoppe, B. M. Smejkal, A. C. Eitelhuber, and T. Schroeder. Hematopoietic cytokines can instruct lineage choice. *Science*, 325(5937):217–218, 2009.

[127] D. M. Roma, R. A. O'Flanagan, A. E. Ruckenstein, A. M. Sengupta, and R. Mukhopadhyay. Optimal path to epigenetic switching. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics*, 71(1 Pt 1):011902, Jan. 2005.

[128] S. Röblitz and M. Weber. Fuzzy spectral clustering by PCCA+: application to Markov state models and data classification. *Advances in Data Analysis and Classification*, 7(2):147–179, June 2013.

[129] K. R. Sanft, S. Wu, M. Roh, J. Fu, R. K. Lim, and L. R. Petzold. Stochkit2: software for discrete stochastic simulation of biochemical systems with events. *Bioinformatics*, 27(17):2457–2458, 2011.

[130] M. Sasai, Y. Kawabata, K. Makishi, K. Itoh, and T. P. Terada. Time Scales in Epigenetic Dynamics and Phenotypic Heterogeneity of Embryonic Stem Cells. *PLOS Computational Biology*, 9(12):e1003380, Dec. 2013.

[131] M. Sasai and P. G. Wolynes. Stochastic gene expression as a many-body problem. *Proceedings of the National Academy of Sciences*, 100(5):2374–2379, Mar. 2003.

[132] M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Pérez-Hernández, M. Hoffmann, N. Plattner, C. Wehmeyer, J.-H. Prinz, and F. Noé. PyEMMA 2: A Software Package for Estimation, Validation, and Analysis of Markov Models. *Journal of Chemical Theory and Computation*, 11(11):5525–5542, Nov. 2015.

[133] D. Schultz, A. M. Walczak, J. N. Onuchic, and P. G. Wolynes. Extinction and resurrection in gene networks. *Proceedings of the National Academy of Sciences of the United States of America*, 105(49):19165–19170, Dec. 2008.

[134] C. Schütte, F. Noé, J. Lu, M. Sarich, and E. Vanden-Eijnden. Markov state models based on milestoning. *The Journal of Chemical Physics*, 134(20):204105, May 2011.

[135] E. Segal and J. Widom. From DNA sequence to transcriptional behaviour: a quantitative approach. *Nature Reviews. Genetics*, 10(7):443–456, July 2009.

[136] S. V. Sharma, D. Y. Lee, B. Li, M. P. Quinlan, F. Takahashi, S. Maheswaran, U. McDermott, N. Azizian, L. Zou, M. A. Fischbach, K.-K. Wong, K. Brandstetter, B. Wittner, S. Ramaswamy, M. Classon, and J. Settleman. A Chromatin-Mediated Reversible Drug-Tolerant State in Cancer Cell Subpopulations. *Cell*, 141(1):69–80, Apr. 2010.

[137] C.-H. Shen and J. Stavnezer. Interaction of stat6 and nf-κb: direct association and synergistic activation of interleukin-4-induced transcription. *Molecular and cellular biology*, 18(6):3395–3404, 1998.

[138] A. Sica and A. Mantovani. Macrophage plasticity and polarization: in vivo veritas. *The Journal of clinical investigation*, 122(3):787, 2012.

[139] J. Silva, J. Nichols, T. W. Theunissen, G. Guo, A. L. van Oosten, O. Barrandon, J. Wray, S. Yamanaka, I. Chambers, and A. Smith. Nanog Is the Gateway to the Pluripotent Ground State. *Cell*, 138(4):722–737, Aug. 2009.

[140] Z. S. Singer, J. Yong, J. Tischler, J. A. Hackett, A. Altinok, M. A. Surani, L. Cai, and M. B. Elowitz. Dynamic Heterogeneity and DNA Methylation in Embryonic Stem Cells. *Molecular Cell*, 55(2):319–331, July 2014.

[141] A. M. Singh, T. Hamazaki, K. E. Hankowski, and N. Terada. A heterogeneous expression pattern for Nanog in embryonic stem cells. *Stem Cells (Dayton, Ohio)*, 25(10):2534–2542, Oct. 2007.

[142] G. Solinas, G. Germano, A. Mantovani, and P. Allavena. Tumor-associated macrophages (tam) as major players of the cancer-related inflammation. *Journal of leukocyte biology*, 86(5):1065–1073, 2009.

[143] K. L. Spiller, S. Nassiri, C. E. Witherel, R. R. Anfang, J. Ng, K. R. Nakazawa, T. Yu, and G. Vunjak-Novakovic. Sequential delivery of immunomodulatory cytokines to facilitate the m1-to-m2 transition of macrophages and enhance vascularization of bone scaffolds. *Biomaterials*, 37:194–207, 2015.

[144] J. L. Stöger, M. J. Gijbels, S. van der Velden, M. Manca, C. M. van der Loos, E. A. Biessen, M. J. Daemen, E. Lutgens, and M. P. de Winther. Distribution of macrophage polarization markers in human atherosclerosis. *Atherosclerosis*, 225(2):461–468, 2012.

[145] M. Strasser, F. J. Theis, and C. Marr. Stability and multiattractor dynamics of a toggle switch based on a two-stage model of stochastic gene expression. *Biophysical Journal*, 102(1):19–29, Jan. 2012.

[146] E. Suárez, S. Lettieri, M. C. Zwier, C. A. Stringer, S. R. Subramanian, L. T. Chong, and D. M. Zuckerman. Simultaneous Computation of Dynamical and Equilibrium Information Using a Weighted Ensemble of Trajectories. *Journal of Chemical Theory and Computation*, 10(7):2658–2667, July 2014.

[147] A. Szanto, B. L. Balint, Z. S. Nagy, E. Barta, B. Dezso, A. Pap, L. Szeles, S. Poliska, M. Oros, R. M. Evans, et al. Stat6 transcription factor is a facilitator of the nuclear receptor pparγ-regulated gene expression in macrophages and dendritic cells. *Immunity*, 33(5):699–712, 2010.

[148] K. Takahashi, K. Tanabe, M. Ohnuki, M. Narita, T. Ichisaka, K. Tomoda, and S. Yamanaka. Induction of Pluripotent Stem Cells from Adult Human Fibroblasts by Defined Factors. *Cell*, 131(5):861–872, Nov. 2007.

[149] A. A. te Velde, R. d. Waal Malefijt, R. J. Huijbens, J. De Vries, C. G. Figdor, et al. Il-10 stimulates monocyte fc-gamma-r surface expression and cytotoxic activity: distinct regulation of antibody-dependent cellular cytotoxicity by ifn-gamma, il-4 and il-10. 1992.

[150] M. J. Tse, B. K. Chu, M. Roy, and E. L. Read. DNA-Binding Kinetics Determines the Mechanism of Noise-Induced Switching in Gene Networks. *Biophysical Journal*, 109(8):1746–1757, Oct. 2015.

[151] F. M. van der Valk, D. F. van Wijk, and E. S. Stroes. Novel anti-inflammatory strategies in atherosclerosis. *Current opinion in lipidology*, 23(6):532–539, 2012.

[152] E. Vanden-Eijnden and M. Venturoli. Revisiting the finite temperature string method for the calculation of reaction tubes and free energies. *The Journal of Chemical Physics*, 130(19):194103, May 2009.

[153] J. Vandesompele, K. De Preter, F. Pattyn, B. Poppe, N. Van Roy, A. De Paepe, and F. Speleman. Accurate normalization of real-time quantitative rt-pcr data by geometric averaging of multiple internal control genes. *Genome biology*, 3(7):research0034–1, 2002.

[154] C. Venkataraman, S. Leung, A. Salvekar, H. Mano, and U. Schindler. Repression of il-4-induced gene expression by ifn-γ requires stat1 activation. *The Journal of Immunology*, 162(7):4053–4061, 1999.

[155] C. Waddington and H. Kacser. *The Strategy of the Genes*. Routledge, 1957.

[156] T. Wakayama, I. Rodriguez, A. C. F. Perry, R. Yanagimachi, and P. Mombaerts. Mice cloned from embryonic stem cells. *Proceedings of the National Academy of Sciences of the United States of America*, 96(26):14984–14989, Dec. 1999.

[157] A. M. Walczak, J. N. Onuchic, and P. G. Wolynes. Absolute rate theories of epigenetic stability. *Proceedings of the National Academy of Sciences of the United States of America*, 102(52):18926–18931, Dec. 2005.

[158] J. Wang, K. Zhang, L. Xu, and E. Wang. Quantifying the Waddington landscape and biological paths for development and differentiation. *Proceedings Of The National Academy Of Sciences Of The United States Of America*, 108(21536909):8257–62, May 2011.

[159] P. Wang, C. Song, H. Zhang, Z. Wu, X.-J. Tian, and J. Xing. Epigenetic state network approach for describing cell phenotypic transitions. *Interface focus*, 4(3):20130068, 2014.

[160] P. Wang, C. Song, H. Zhang, Z. Wu, X.-J. Tian, and J. Xing. Epigenetic state network approach for describing cell phenotypic transitions. *Interface Focus*, 4(3), June 2014.

[161] G. R. Warnes, B. Bolker, L. Bonebakker, R. Gentleman, W. Huber, A. Liaw, T. Lumley, M. Maechler, A. Magnusson, S. Moeller, et al. gplots: Various r programming tools for plotting data. *R package version*, 2(4):1, 2009.

[162] L. S. Weinberger, J. C. Burnett, J. E. Toettcher, A. P. Arkin, and D. V. Schaffer. Stochastic Gene Expression in a Lentiviral Positive-Feedback Loop: HIV-1 Tat Fluctuations Drive Phenotypic Diversity. *Cell*, 122(2):169–182, July 2005.

[163] J. Xue, S. V. Schmidt, J. Sander, A. Draffehn, W. Krebs, I. Quester, D. De Nardo, T. D. Gohel, M. Emde, L. Schmidleithner, et al. Transcriptome-based network analysis reveals a spectrum model of human macrophage activation. *Immunity*, 40(2):274–288, 2014.

[164] Y. Yamanaka, F. Lanner, and J. Rossant. FGF signal-dependent segregation of primitive endoderm and epiblast in the mouse blastocyst. *Development (Cambridge, England)*, 137(5):715–724, Mar. 2010.

[165] J. Ye, G. Coulouris, I. Zaretskaya, I. Cutcutache, S. Rozen, and T. L. Madden. Primer-blast: a tool to design target-specific primers for polymerase chain reaction. *BMC bioinformatics*, 13(1):134, 2012.

[166] B. Zhang, D. Jasnow, and D. Zuckerman. The "weighted ensemble" path sampling method is statistically exact for a broad class of stochastic processes and binning procedures. *The Journal of chemical physics*, 132(20136305):054107, Feb. 2010.

[167] B. Zhang and P. G. Wolynes. Stem cell differentiation as a many-body problem. *Proceedings of the National Academy of Sciences*, 111(28):10185–10190, July 2014.

[168] K. Zhang, M. Sasai, and J. Wang. Eddy current and coupled landscapes for nonadiabatic and nonequilibrium complex system dynamics. *Proceedings of the National Academy of Sciences*, 110(37):14930–14935, Sept. 2013.

[169] P. Zhang, G. Behre, J. Pan, A. Iwama, N. Wara-aswapati, H. S. Radomska, P. E. Auron, D. G. Tenen, and Z. Sun. Negative cross-talk between hematopoietic regulators: GATA proteins repress PU.1. *Proceedings of the National Academy of Sciences of the United States of America*, 96(15):8705–8710, July 1999.

[170] J. X. Zhou, M. D. S. Aliyu, E. Aurell, and S. Huang. Quasi-potential landscape in complex multi-stable systems. *Journal of the Royal Society Interface*, 9(77):3539–3553, Dec. 2012.

[171] J. X. Zhou and S. Huang. Understanding gene circuits at cell-fate branch points for rational cell reprogramming. *Trends in Genetics*, 27(2):55–62, Feb. 2011.

[172] D. M. Zuckerman and L. T. Chong. Weighted Ensemble Simulation: Review of Methodology, Applications, and Software. *Annual Review of Biophysics*, 46(1):43–57, 2017.

[173] M. C. Zwier, J. L. Adelman, J. W. Kaus, A. J. Pratt, K. F. Wong, N. B. Rego, E. Suárez, S. Lettieri, D. W. Wang, M. Grabe, D. M. Zuckerman, and L. T. Chong. WESTPA: An Interoperable, Highly Scalable Software Package for Weighted Ensemble Simulation and Analysis. *Journal of Chemical Theory and Computation*, 11(2):800–809, Feb. 2015.

[174] M. C. Zwier and L. T. Chong. Reaching biological timescales with all-atom molecular dynamics simulations. *Current Opinion in Pharmacology*, 10(6):745–752, Dec. 2010.

Chapter 6

Supporting Info

6.1 Stable states of the base network and the network with dimerization

Parameters for the base network were taken to be consistent with Schultz *et al.*[133]. Each parameter set has the same stable state in the protein sub-spaces ($a = 0$, $b = g/k$ and $a = g/k$, $b = 0$) and the same dissociation constant of the binding proteins to DNA ($X_{eq} = f/h$). We use an order parameter in terms of DNA-operator occupancy (gene state) and protein copy number, $\lambda = na - nb$, where $na = a + 2a_2 + 2B_{off}$ and $nb = b + 2b_2 + 2A_{off}$. For the network with dimerization, we use a parameter $X_{ad} = g_d/k_d$, where g_d and k_d are the monomer production and degradation rates for the network, respectively. The stable states in the protein sub-spaces of the network with dimerization are located where $a = 0$, $a_2 = 0$, $b = X_{ad}$, $b_2 = d(2u)^{-1}(X_{ad}^2 - X_{ad})$ and $a = X_{ad}$, $a_2 = d(2u)^{-1}(X_{ad}^2 - X_{ad})$, $b = 0$, $b_2 = 0$. The rate parameters for the network with dimerization were chosen to keep the stable states in term of the order parameter λ of both networks equivalent. To maintain the equivalency, the DNA unbinding rate and the monomer production rate are kept the same between both

networks ($g = g_d$; $f = f_d$) while the parameter X_{ad} is scaled in the dimerization network to $X_{ad} = (1 - u/d)/2 + (1/2)[(1 - u/d)^2 + 4gu/dk]^{1/2}$ and the DNA binding rate is scaled to $h_d = hu\frac{(g/k)^2 - g/k}{d(X_{ad}^2 - X_{ad})}$.

6.2 Exclusive switch chemical reaction network

In the general switch, four distinct gene states are present due to the binary regulation of each gene (i.e., $A_{on}/B_{on},\ A_{on}/B_{on},\ A_{off}/B_{on},\ A_{off}/B_{off}$). In the exclusive switch, competitive binding means that the A_{off}/B_{off} gene state is disallowed. This is encoded in the reaction network by transforming the reactions to the operator notation used by Morelli *et al.*[108]:

$$O(\text{unbound operator}) = A_{on}/B_{on}$$
$$OA_2 = A_{on}/B_{off}$$
$$OB_2 = A_{off}/B_{on}$$
$$OA_2B_2 = A_{off}/B_{off} \text{ (not present in exclusive switch reaction)}$$

The reactions for the exclusive switch are thus transformed to:

$$\mathrm{O} \xrightarrow{\mathrm{g}} \mathrm{a} \qquad \mathrm{O} \xrightarrow{\mathrm{g}} \mathrm{b}$$
$$\mathrm{OA_2} \xrightarrow{\mathrm{g}} \mathrm{a} \qquad \mathrm{OB_2} \xrightarrow{\mathrm{g}} \mathrm{b}$$
$$\mathrm{a} \xrightarrow{\mathrm{k}} 0 \qquad \mathrm{b} \xrightarrow{\mathrm{k}} 0$$
$$\mathrm{O} + 2\,\mathrm{b} \underset{\mathrm{f}}{\overset{\mathrm{h}}{\rightleftharpoons}} \mathrm{OB_2} \qquad \mathrm{O} + 2\,\mathrm{a} \underset{\mathrm{f}}{\overset{\mathrm{h}}{\rightleftharpoons}} \mathrm{OA_2}$$

The parameters used in these equations are the same as those for the general toggle switch. In this formulation, a single operator (O) has a single binding site for either two a proteins or two b proteins. Binding by one repressor prevents the other from binding. The unbound

operator can produce either a protein or b protein, but a bound operator can only express one protein. The probabilities that the operator is in any of its three states ($P(O)$, $P(OA_2)$, and $P(OB_2)$) can be converted into the probabilities that gene A is on ($P(A_{on})$) or the probability that gene B is on ($P(B_{on})$) to be consistent with the notation of the general toggle switch is shown below:

$$P(A_{on}) = P(O) + P(OA_2)$$
$$P(B_{on}) = P(O) + P(OB_2)$$
$$P(A_{off}) = 1 - P(A_{on})$$
$$P(B_{off}) = 1 - P(B_{on})$$

6.3 ExMISA Network

Two-gene network with Mutual Inhibition, Self-Activation, and exclusive transcription factor binding.

$$\mathrm{A_{00}} + 2\,\mathrm{a} \underset{\mathrm{f}_a}{\overset{\mathrm{h}_a}{\rightleftharpoons}} \mathrm{A_{10}} \qquad \mathrm{B_{00}} + 2\,\mathrm{b} \underset{\mathrm{f}_a}{\overset{\mathrm{h}_a}{\rightleftharpoons}} \mathrm{B_{10}}$$

$$\mathrm{A_{00}} + 2\,\mathrm{b} \underset{\mathrm{f}_r}{\overset{\mathrm{h}_r}{\rightleftharpoons}} \mathrm{A_{01}} \qquad \mathrm{B_{00}} + 2\,\mathrm{a} \underset{\mathrm{f}_r}{\overset{\mathrm{h}_r}{\rightleftharpoons}} \mathrm{B_{01}}$$

$$\mathrm{A_{00}} \xrightarrow{\mathrm{g_0}} \mathrm{A_{00}} + \mathrm{a} \qquad \mathrm{B_{00}} \xrightarrow{\mathrm{g_0}} \mathrm{B_{00}} + \mathrm{b}$$

$$\mathrm{A_{01}} \xrightarrow{\mathrm{g_0}} \mathrm{A_{01}} + \mathrm{a} \qquad \mathrm{B_{01}} \xrightarrow{\mathrm{g_0}} \mathrm{B_{01}} + \mathrm{b}$$

$$\mathrm{A_{10}} \xrightarrow{\mathrm{g_1}} \mathrm{A_{10}} + \mathrm{a} \qquad \mathrm{B_{10}} \xrightarrow{\mathrm{g_1}} \mathrm{B_{10}} + \mathrm{b}$$

$$\mathrm{a} \xrightarrow{\mathrm{k}} 0 \qquad \mathrm{b} \xrightarrow{\mathrm{k}} 0$$

6.4 Pluripotency network

There are eight genes (encoding transcription factors) in the pluripotency network. Transcription factors bind as homodimers with the exception of the OCT4-SOX2 heterodimer. Only three transcription factors interact with their own gene, CDX2, NANOG, and GATA6. Transcription factors bind as dimers with the rate h and unbind with the rate f. When a gene is bound by any activator and no repressors, it expresses at a rate g_{on}, otherwise, it expresses at a rate g_{off}. The only exception is NANOG, which must be bound by all three of its activators and no repressors to be activated.

6.5 Pseudo-code for the Weighted-Ensemble Based Computational Pipeline

6.5.1 Weighted Ensemble Exploration Mode

1. Format the reaction network into a BioNetGen file.
2. Choose M_{targ}, the target number of replicas per sampling region, and N_{bins}, the target number of sampling regions or bins.
3. Initialize M_{targ} replicas in a single starting location.
4. Simulate replicas for a simulation time τ_{WE}. Replicas are simulated in parallel.
5. Chose N_{bins} new bin positions.
   (a) Chose one random replica as the first new bin position.
   (b) Chose the replica furthest from the set of bin positions to be the next new bin position.

(c) Repeat (b) until N_{bins} new positions have been chosen.

6. Perform the WE step.

   (a) For a given bin, if the number of replicas in the bin is less than M_{targ}, split the replica with the largest weight into n equally weighted replicas until there are M_{targ} replicas.

   (b) For a given bin, if the number of replicas in the bin is greater than M_{targ}, combine the weight of n replicas and randomly chose one to receive the combined weight such that there are M_{targ} replicas in the bin.

   (c) Repeat (a) or (b) for each sampling bin

7. Repeat steps 4-6 for a chosen number of simulation steps.

6.5.2 Transition-Matrix Mode

1. Start from the end of exploration mode.

2. Simulate replicas for a time τ_{WE}.

3. Collect weights transferred from bin i to bin j over the simulation period τ into a transition matrix.

4. Perform the WE step.

5. Repeat steps 2-4 for a chosen number of simulation steps.

6.5.3 Rate-Estimation Mode

1. Start from the end of exploration mode.

2. Label replicas as having most recently visited region of interest X or visited region of interest Y.

3. Simulate replicas for a time τ_{WE}.

4. Collect weights transferred from X to Y over the simulation period τ and change the replica label as necessary.

5. (Optional) Chose N_{bins} new sampling regions.

6. Perform the WE step.

7. Repeat steps 3-6 for a chosen number of simulation steps.

6.5.4 Coarse-Graining Procedure

1. Find the left-eigenvalues and eigenvectors of the row-stochastic transition matrix calculated from transition-matrix estimation mode.

   (a) The probability distribution of the system is estimated by the left-eigenvector associated with the eigenvalue $\lambda = 1$

2. Perform the PCCA+ algorithm using MSMBuilder software to cluster the N_{bins} sampling regions into macrostates. MSMBuilder software outputs a Markov State Model of the reaction network.

3. Use transition path analysis (using PyEMMA software) on the resulting MSM to obtain parallel transition paths or estimate the rate of transitioning between any two states.

4. Gephi 0.7 is used to visualize the row-stochastic transition matrix and the MSM.

6.6 Models of Macrophage Polarization

The parameters of these equations are described in Supplementary Table S2 and fitted numeric values are given in Supplementary Table S3.

Equation S1. Model 1: Self-activation

$$\frac{d[M1]}{dt} = k_1 \frac{S1}{S1 + K_{ind1}} + k_3 \frac{[M1]}{[M1] + K_{\text{act}}} + k_5 - d_1[M1]$$

$$\frac{d[M2]}{dt} = k_2 \frac{S2}{S2 + K_{ind2}} + k_4 \frac{[M2]}{[M2] + K_{\text{act}}} + k_6 - d_2[M2]$$

$$\frac{d[CD86]}{dt} = g_1[M1] - d_4[CD86]$$

$$\frac{d[CD206]}{dt} = g_2[M2] - d_5[CD206]$$

Equation S2. Model 2: MISA

$$\frac{d[M1]}{\text{dt}} = \frac{k_1 \frac{S1}{S1+K_{ind1}} + k_3 \frac{[M1]}{[M1]+K_{\text{act}}}}{1 + \left([M2]/K_{rep2}\right)^n} + k_5 - d_1[M1]$$

$$\frac{d[M2]}{dt} = \frac{k_2 \frac{S2}{S2+K_{ind2}} + k_4 \frac{[M2]}{[M2]+K_{\text{act}}}}{1 + \left([M1]/K_{rep1}\right)^n} + k_6 - d_2[M2]$$

$$\frac{d[CD86]}{dt} = g_1[M1] - d_4[CD86]$$

$$\frac{d[CD206]}{dt} = g_2[M2] - d_5[CD206]$$

Equation S3. Model 3: MISA with IFFL and inhibition between Y and M1

$$\frac{d[M1]}{\mathrm{dt}} = \frac{k_1 \frac{S1}{S1+K_{ind1}} + k_3 \frac{[M1]}{[M1]+K_{\mathrm{act}}}}{\left(1+\left([M2]/K_{rep2}\right)^n\right)\left(1+\left([Y]/K_Y\right)^n\right)} + k_5 - d_1[M1]$$

$$\frac{d[M2]}{dt} = \frac{k_2 \ \frac{S2}{S2+K_{ind2}} + k_4 \ \frac{[M2]}{[M2]+K_{\mathrm{act}}}}{(1+([M1]/K_{rep1})^n)} + k_6 + k_7 \frac{[Y]}{[Y]+K_{\mathrm{CY}}} - d_2[M2]$$

$$\frac{d[Y]}{dt} = k_8 \frac{S1}{S1+K_{ind1}} - d_3[Y1]$$

$$\frac{d[CD86]}{dt} = g_1[M1] - d_4[CD86]$$

$$\frac{d[CD206]}{dt} = g_2[M2] - d_5[CD206]$$

Equation S4. Model 4: MISA with cooperative IFFL

$$\frac{d[M1]}{\mathrm{dt}} = \frac{k_1 \frac{S1}{S1+K_{ind1}} + k_3 \frac{[M1]}{[M1]+K_{\mathrm{act}}}}{\left(1+\left([M2]/K_{rep2}\right)^n\right)} + k_5 - d_1[M1]$$

$$\frac{d[M2]}{\mathrm{dt}} = \frac{k_2 \ \frac{S2}{S2+K_{ind2}} + k_4 \ \frac{[M2]}{[M2]+K_{\mathrm{act}}} + k_7 \frac{[M2][Y]}{(K_{CM2}+[M2])(K_{\mathrm{CY}}+[Y])}}{(1+([M1]/K_{rep1})^n)} + k_6 - d_2[M2]$$

$$\frac{d[Y]}{dt} = k_8 \frac{S1}{S1 + K_{ind1}} - d_3[Y1]$$

$$\frac{d[CD86]}{dt} = g_1[M1] - d_4[CD86]$$

$$\frac{d[CD206]}{dt} = g_2[M2] - d_5[CD206]$$

Equation S5. Model 5: MISA with cooperative IFFL and inhibition between Y and M1

$$\frac{d[M1]}{\mathrm{dt}} = \frac{k_1 \frac{S1}{S1+K_{ind1}} + k_3 \frac{[M1]}{[M1]+K_{\mathrm{act}}}}{\left(1 + \left([M2]/K_{rep2}\right)^n\right)\left(1 + \left([Y]/K_Y\right)^n\right)} + k_5 - d_1[M1]$$

$$\frac{d[M2]}{\mathrm{dt}} = \frac{k_2 \ \frac{S2}{S2+K_{ind2}} + k_4 \ \frac{[M2]}{[M2]+K_{\mathrm{act}}} + k_7 \frac{[M2][Y]}{(K_{CM2}+[M2])(K_{\mathrm{CY}}+[Y])}}{\left(1 + \left([M1]/K_{rep1}\right)^n\right)} + k_6 - d_2[M2]$$

$$\frac{d[Y]}{dt} = k_8 \frac{S1}{S1 + K_{ind1}} - d_3[Y1]$$

$$\frac{d[CD86]}{dt} = g_1[M1] - d_4[CD86]$$

$$\frac{d[CD206]}{dt} = g_2[M2] - d_5[CD206]$$

Equation S6. Model 6: MISA with cooperative IFFL, inhibition between Y and

M1, and inhibition on Y

$$\frac{d[M1]}{\mathrm{dt}} = \frac{k_1 \frac{S1}{S1+K_{ind1}} + k_3 \frac{[M1]}{[M1]+K_{\mathrm{act}}}}{\left(1+\left([M2]/K_{rep2}\right)^n\right)\left(1+\left([Y]/K_Y\right)^n\right)} + k_5 - d_1[M1]$$

$$\frac{d[M2]}{\mathrm{dt}} = \frac{k_2 \ \frac{S2}{S2+K_{ind2}} + k_4 \ \frac{[M2]}{[M2]+K_{\mathrm{act}}} + k_7 \frac{[M2][Y]}{(K_{CM2}+[M2])(K_{\mathrm{CY}}+[Y])}}{\left(1+\left([M1]/K_{rep1}\right)^n\right)} + k_6 - d_2[M2]$$

$$\frac{d[Y]}{dt} = \frac{k_8 \frac{S1}{S1+K_{ind1}}}{1+\left(S2/K_{ind2}\right)^n} - d_3[Y1]$$

$$\frac{d[CD86]}{dt} = g_1[M1] - d_4[CD86]$$

$$\frac{d[CD206]}{dt} = g_2[M2] - d_5[CD206]$$

6.7 Supporting Figures

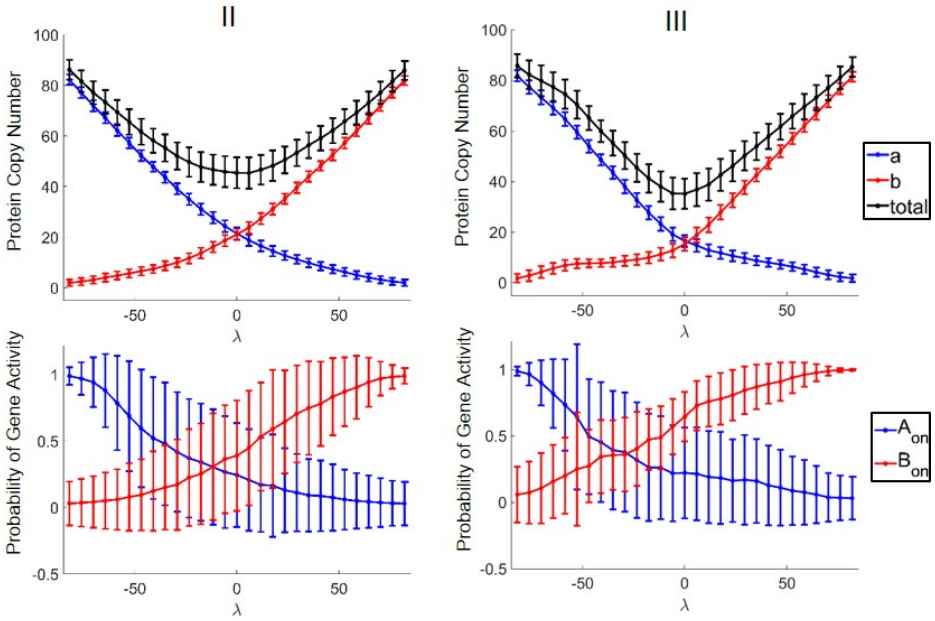

Figure 6.1: **General switch: Transition paths predicted from brute force simulations.** The mean and standard deviation of the $S_A \rightarrow S_B$ transition is shown for the general toggle switch as averaged from 300 successful brute force switching-trajectories for parameters II and III. Stochastic switching trajectories are averaged by binning according to the progress coordinate λ (see Methods) (Top) The transition path projected onto the protein numbers vs. λ. (Bottom) The transition path projected onto the probabilities of gene activity vs. λ. With decreasing adiabaticity, the switch occurs earlier in the gene activity subspace relative to the proteins. In the protein numbers, $a = b$ near $\lambda = 0$ in all cases, whereas $A_{on} = B_{on}$ at increasingly negative values of lambda with decreasing adiabaticity. (It was not possible to obtain transition paths from brute force simulation for Parameter Set I: Simulating for 10^6 timesteps failed to produce any successful switching trajectories.)

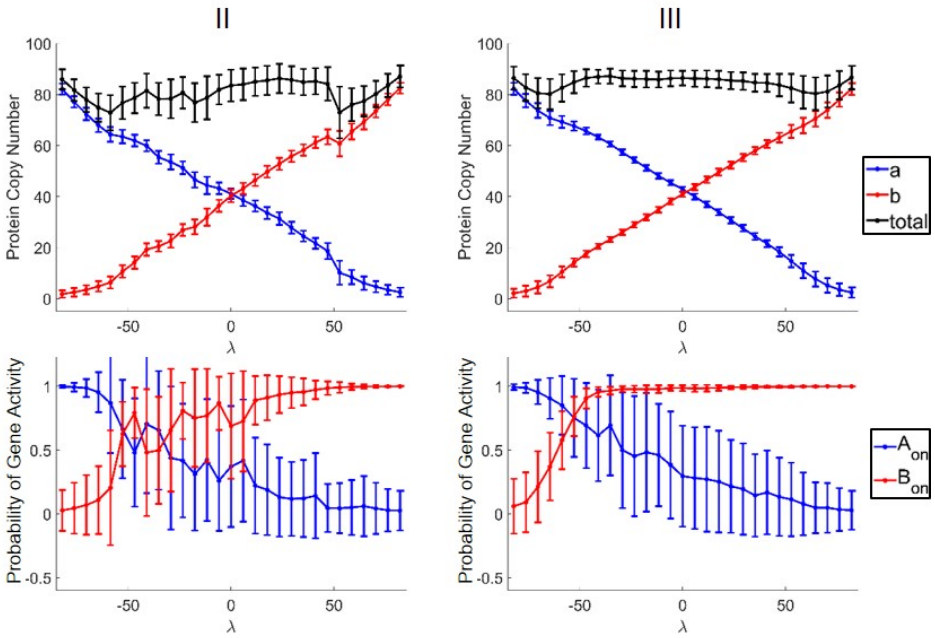

Figure 6.2: **Exclusive switch: Transition paths predicted from brute force simulations.** The mean and standard deviation of the $S_A \rightarrow S_B$ transition is shown for the exclusive toggle switch as averaged from 300 successful brute force switching-trajectories for parameters II and III. All definitions are the same as Fig. 6.1.

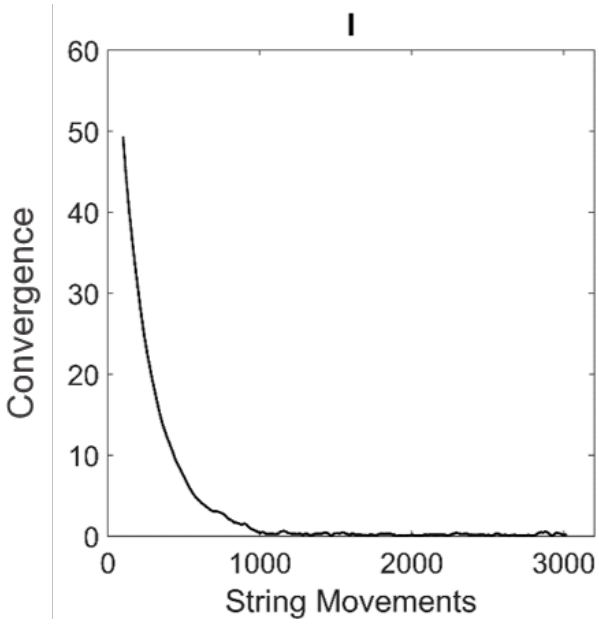

Figure 6.3: **Convergence is defined as the root mean squared distance (RMSD) between a moving average of the last 100 string movements and the current string position.** The convergence of the baseline general toggle switch for parameter set I is illustrative of the trend in convergence for the other parameter sets and chemical network variants.

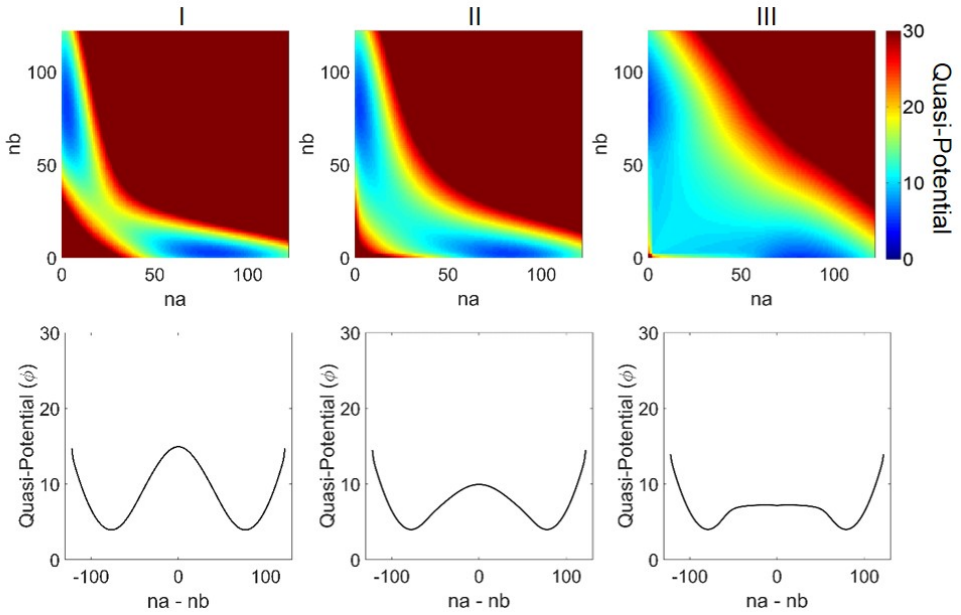

Figure 6.4: **General switch: The barrier height in the quasi-potential between the two attractor states decreases with decreasing adiabaticity.** (Top) Quasi-potential projected onto the total number of proteins na in the x-axis and nb in the y-axis. (Bottom) 1-D quasi-potential of the order parameter $na - nb$.

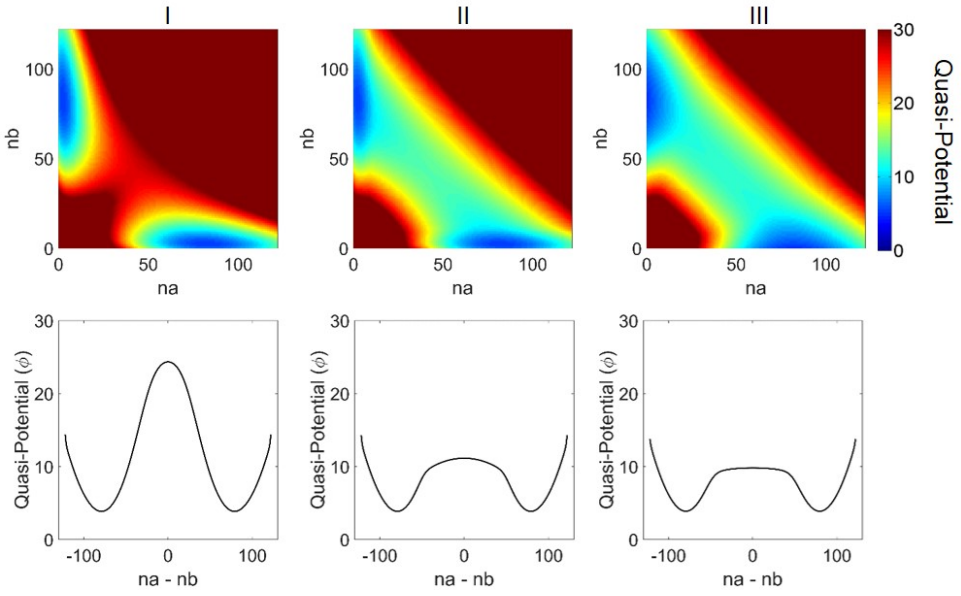

Figure 6.5: **Exclusive switch: The barrier height in the quasi-potential between the two attractor states decreases with decreasing adiabaticity.** All definitions are the same as Fig. 6.4

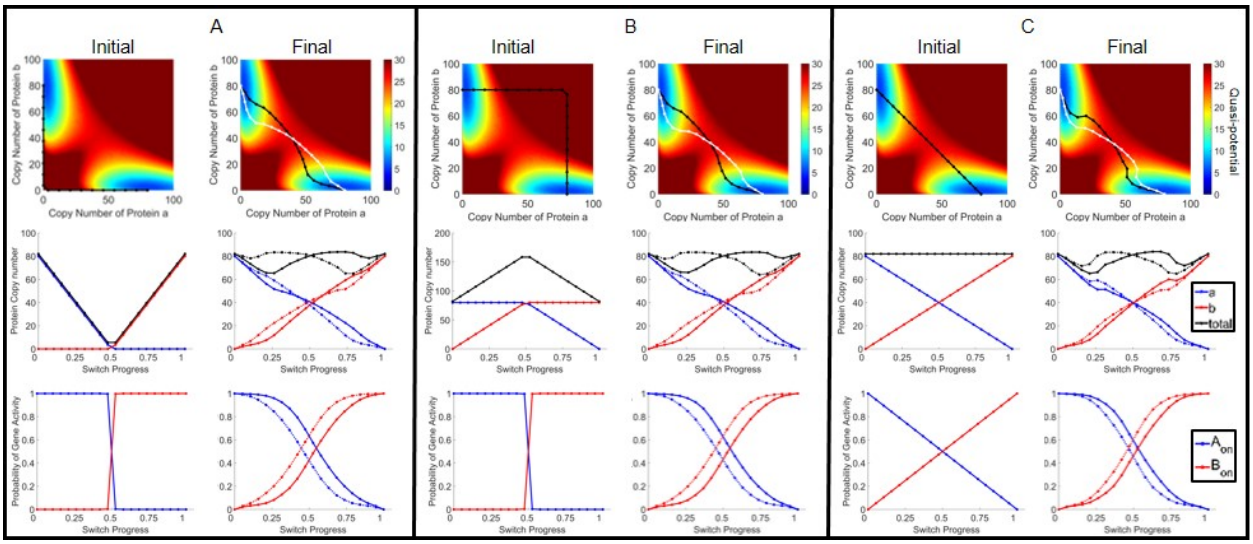

Figure 6.6: **Results of transition path prediction are robust to large differences in initial guess path.** Shown here are the predicted transition paths and quasi-potential surfaces of the exclusive switch initialized in three different regions of state space for the most adiabatic parameter set (I). Despite the difference in initial guess paths, all three calculations return transitions with similar features and separation of forwards and backwards transitions. Panels A) and B) contain the initial transition guess path and final predicted transition for a string initialized in low probability areas of state space. Panel C) displays the initial transition guess path and final predicted transition for a typical string initialization.

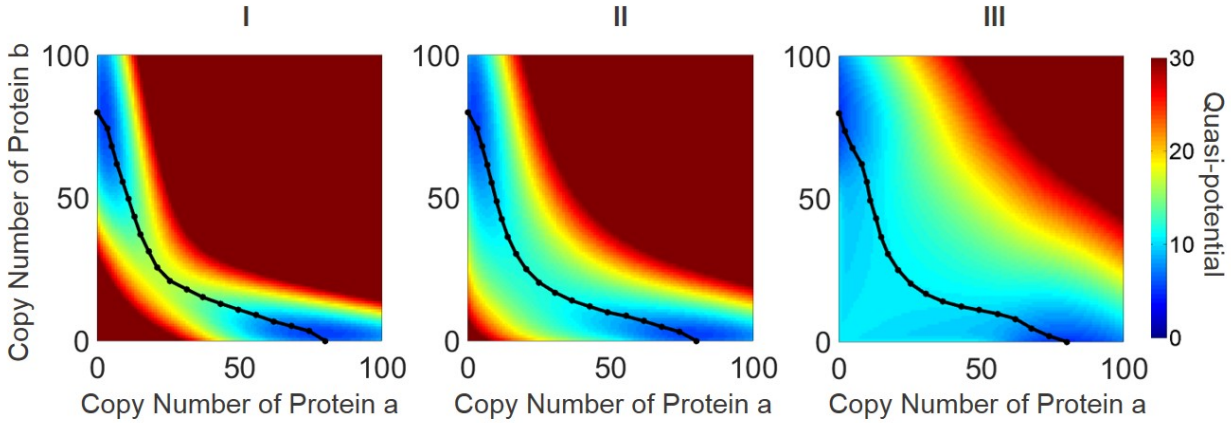

Figure 6.7: **Predicted minimum energy paths for the baseline general toggle switch network determined only from the gradient dynamics of the quasi-potential surface (ϕ).** Since the quasi-potential surface is symmetric, the resulting strings are also symmetric along the ascent out of a basin of attraction and the descent into a basin. The string method (following E. *et al.* ([40])) was used to calculate the minimum energy path by iteratively moving 20 string nodes according to the gradient dynamics of the quasi-potential and reparameterizing the nodes.

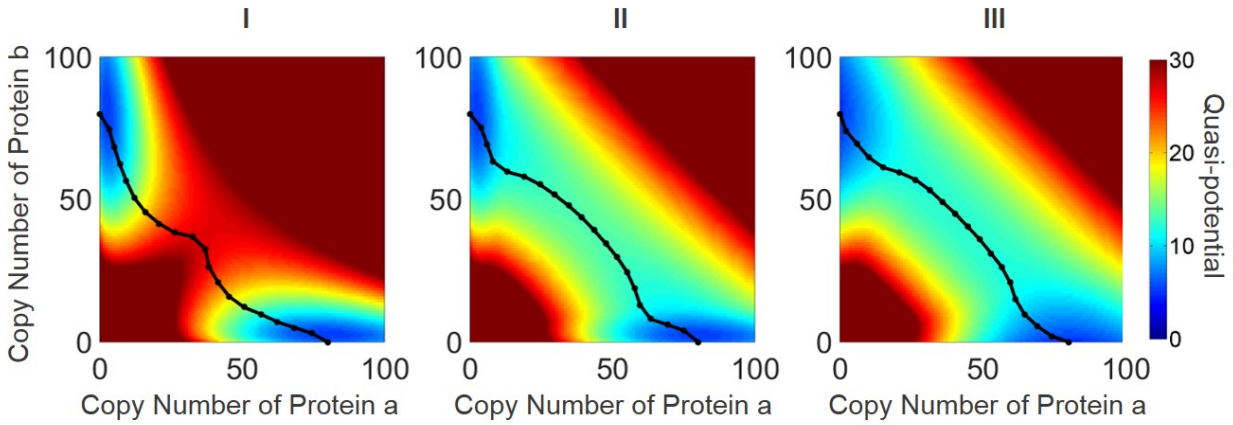

Figure 6.8: **Predicted minimum energy paths for the baseline exclusive toggle switch network determined only from the gradient dynamics of the quasi-potential surface (ϕ).** All definitions are the same as Fig. 6.7.

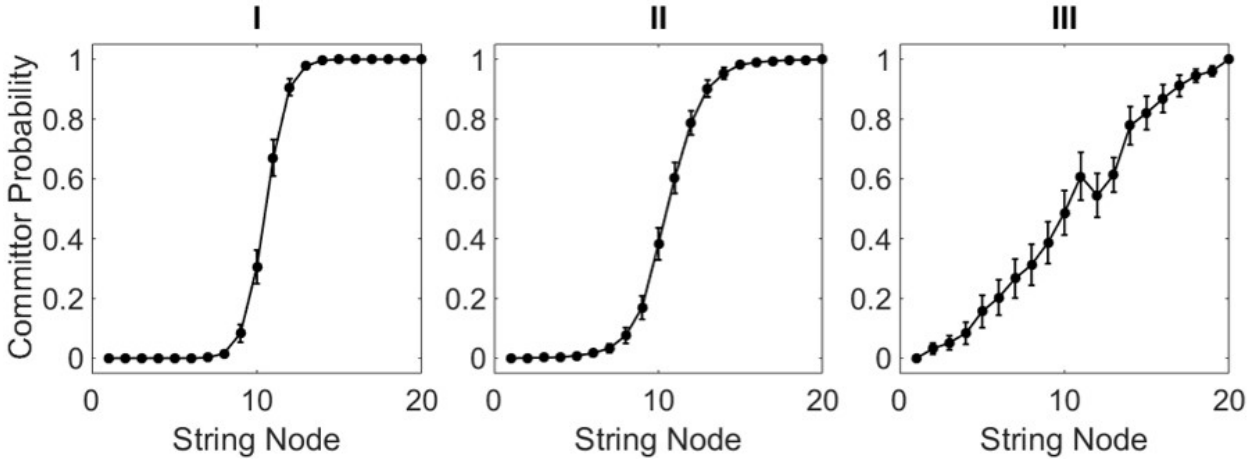

Figure 6.9: **Committor probabilities for the transition from $S_A \rightarrow S_B$ in the baseline general toggle switch network for each region of state space defined by the string nodes.** For each Voronoi region, one hundred trajectories are initialized for each replica in the region. The committor for one such replica is the number of trajectories that successfully reach S_B before they reach S_A. The committors in each region are averaged to find the committor probability of the corresponding string node. The standard deviation is plotted for each string node.

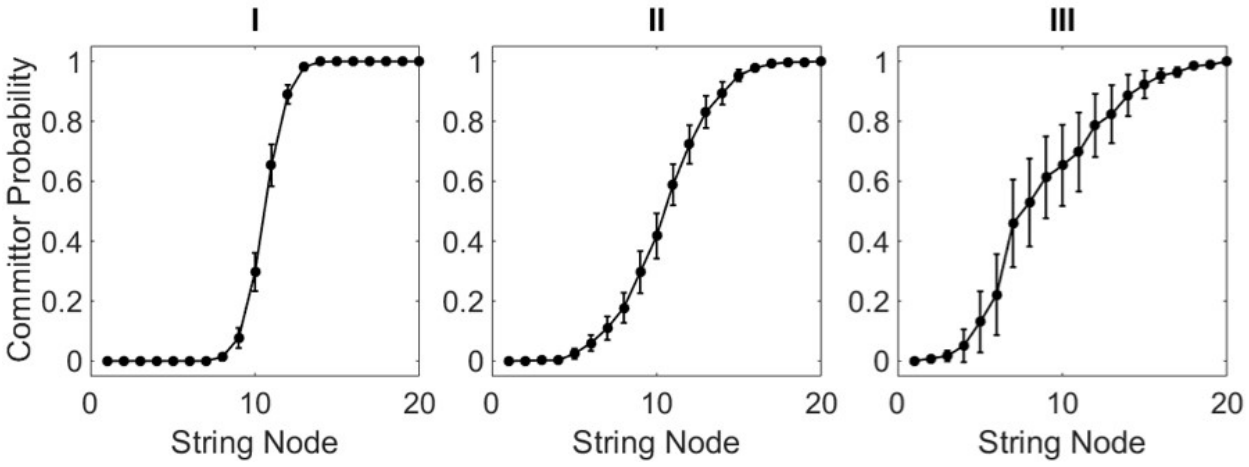

Figure 6.10: **Committor probabilities for the transition from $S_A \rightarrow S_B$ in the baseline exclusive toggle switch network for each region of state space defined by the string nodes.** All definitions are the same as Fig. 6.9

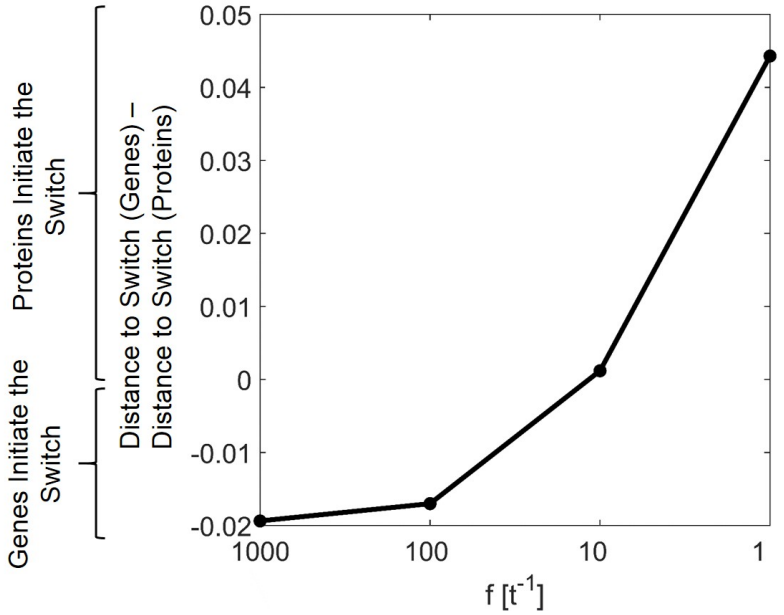

Figure 6.11: **General switch: when adiabaticity decreases, the transition changes from being initiated by fluctuations in DNA-binding occupancy (changes in gene activity) to fluctuations in protein copy number.** The distance to the switch in the gene states minus the distance to the switch in the protein sub-spaces is plotted vs. decreasing adiabaticity. This difference in distance is defined in terms of the switch progress (Methods): $(D_{A_{on}-B_{on}} - D_{a=b})/D_{AB}$. The switch occurs in the gene sub-spaces when the activity of the two genes is equivalent ($A_{on} = B_{on}$), and the switch occurs in the protein sub-spaces when the two proteins have equal copy number ($a = b$).

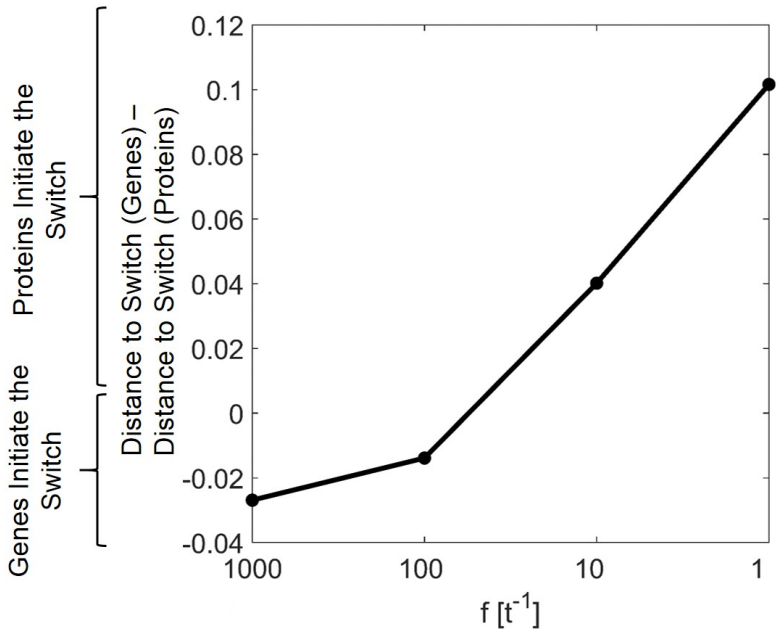

Figure 6.12: **Exclusive switch: when adiabaticity decreases, the transition changes from being initiated by fluctuations in DNA-binding occupancy (changes in gene activity) to fluctuations in protein copy number.** All definitions are the same as Fig. 6.11

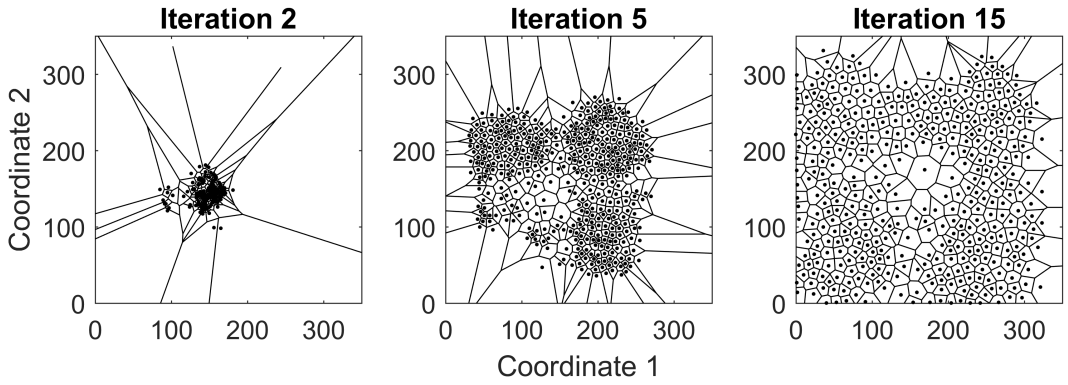

Figure 6.13: **Movement of Voronoi Centers during weighted ensemble sampling.** Starting from the left are shown three successive iterations of the adaptive WE simulation for a representative network.

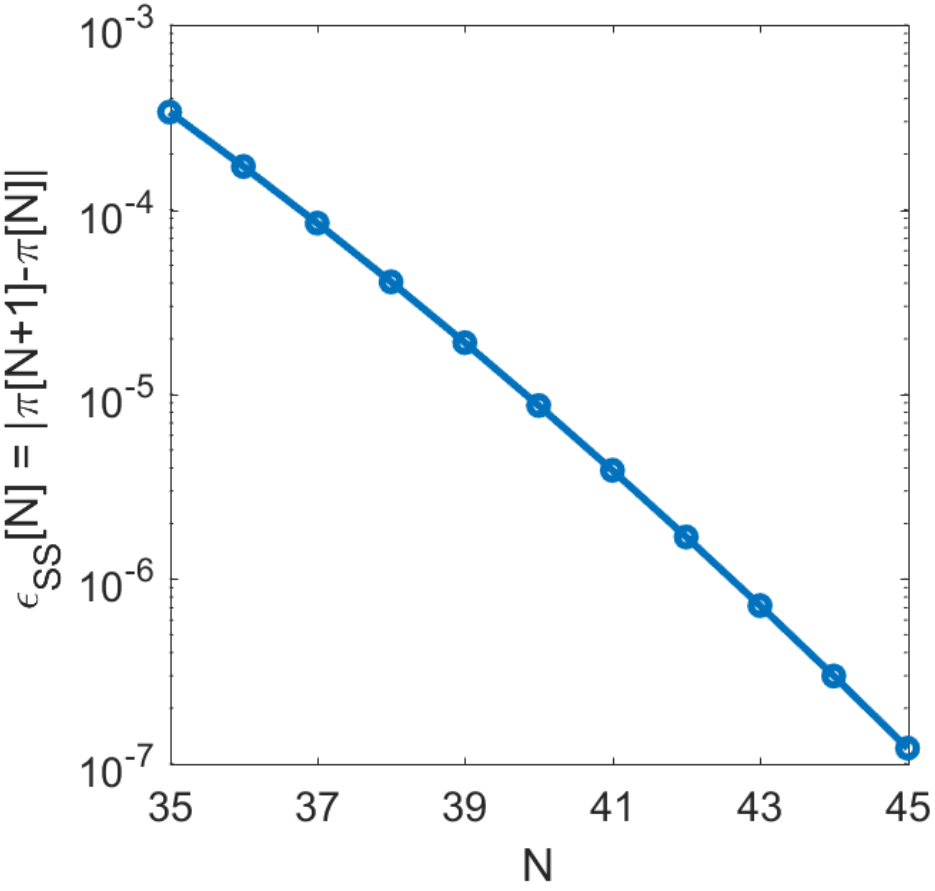

Figure 6.14: **Error in computed steady-state probability as a function of N, the number of protein states retained in the state-space truncation.** N corresponds to the maximum allowed copy-number of transcription factors a and b in the ExMISA network. For a truncation to N, probability flux between states with n_a, $n_b \leq N$ and states with n_a, $n_b > N$ is assumed to be 0 (i.e., the boundaries of the state-space are reflective). The error $\epsilon_S S[N]$ is defined by $\sum_i |\pi[N+1] - \pi[N]|$, where i runs over all enumerated states of the state-space with truncation to $N+1$ (all states outside the boundary have probability 0). That is, the error is computed as the sum of the absolute difference between steady-state probabilities for each state, comparing $\pi[N]$ (steady-state probability computed with truncation to N) to $\pi[N+1]$ (truncated to $N+1$).

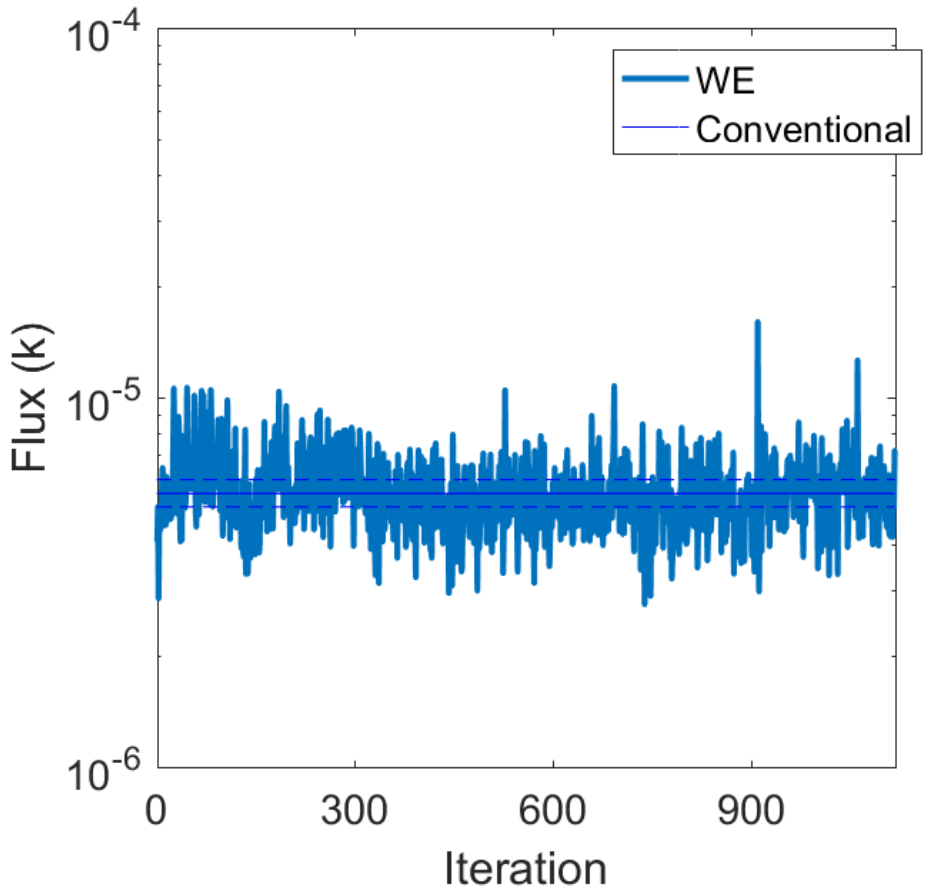

Figure 6.15: **Convergence of the flux of the transition between the polarized phenotype-states in the ExMISA network.** The 5% and 95% confidence intervals for the long conventional simulation are shown in dotted blue lines. The flux between the a/b hi/lo and lo/hi phenotypes was calculated using WE sampling with parameters: $\tau = 200$, 300 bins, and 50 replicas per bin. The system was sampled for 1100 iterations of τ.

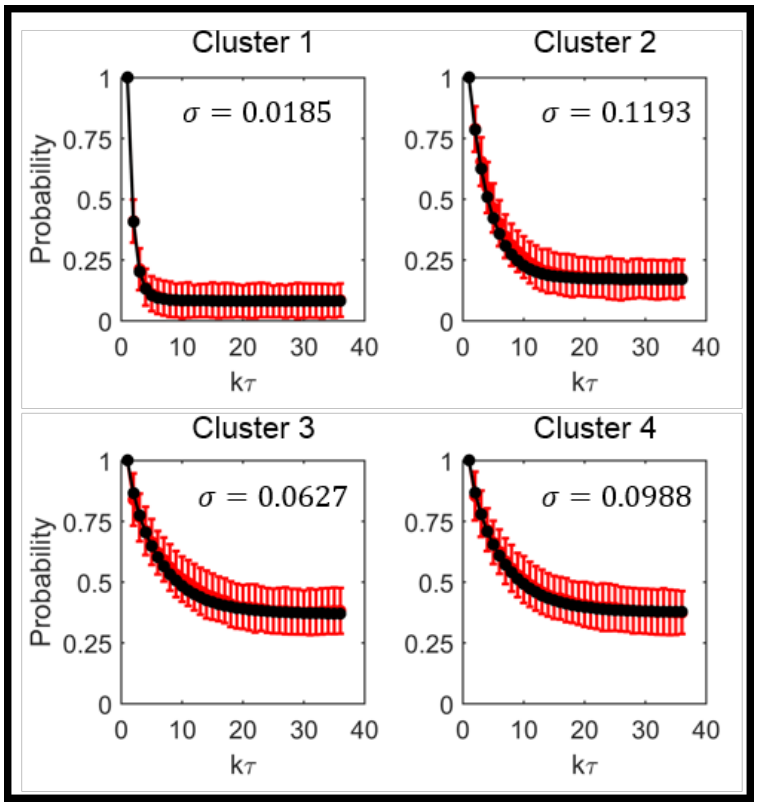

Figure 6.16: **The Chapman-Kolmogorov test on the four Markov State Model phenotypes of the sampled ExMISA network.** The relaxation curves and variance of a 1000τ trajectory are shown in red. The relaxation curve predictions from the MSM transition matrix is shown in black. The total error σ is measured as the 2-norm containing the differences between the two estimates of the relaxation curve.

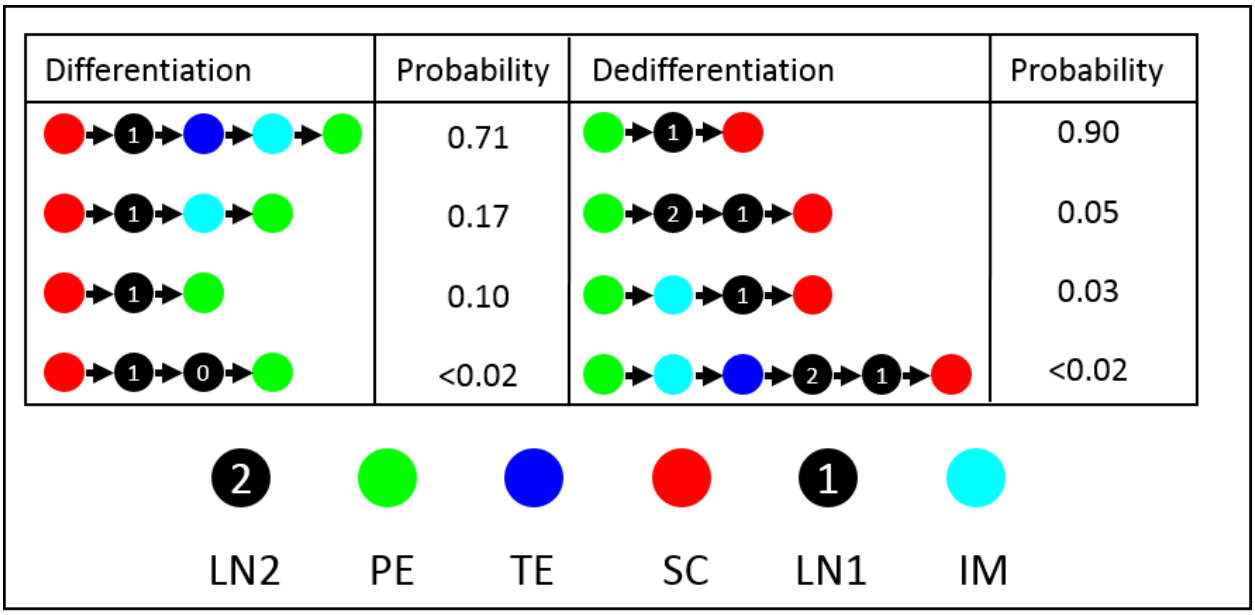

Figure 6.17: **Pathway decomposition for the SC → PE transition for $f = 10$**

Method	Differentiation	Probability	Dedifferentiation	Probability
Conventional Simulation		0.73		0.90
		0.25		0.06
		<0.02		<0.03
		<0.01		<0.02
Simulated Markov State Model		0.73		0.90
		0.16		0.05
		0.09		0.03
		<0.02		<0.02

Figure 6.18: **Validation of the SC → TE transition pathway calculated through weighted ensemble sampling.** The parallel transition pathways are compared against those calculated from a single long conventional simulation.

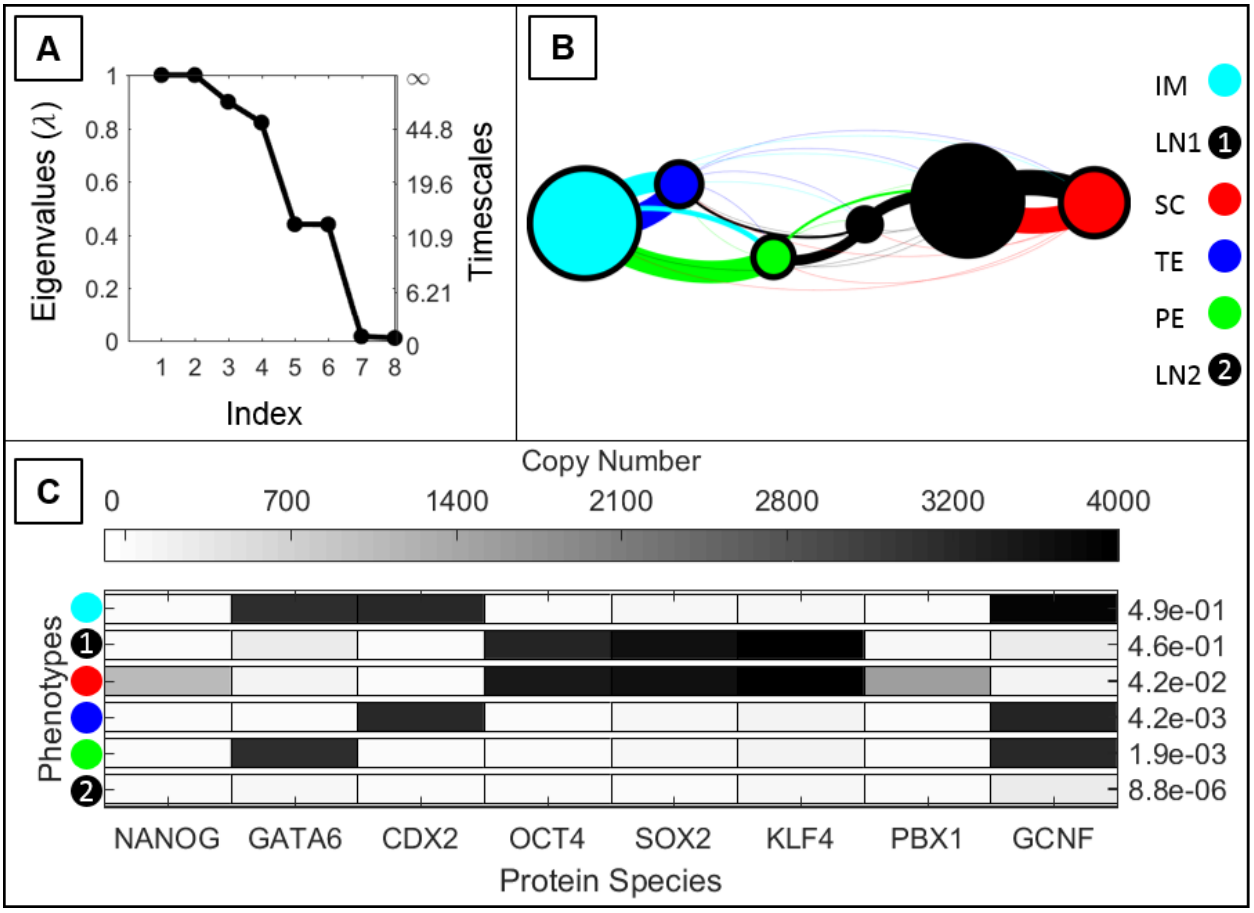

Figure 6.19: **Reproducibility of the weighted ensemble sampling of the pluripotency network.** The second WE sampling of $f = 10$ parameter set was initialized in the same manner as the first. A) Eigenvalues and timescales. B) MSM C) Macrostate compositions

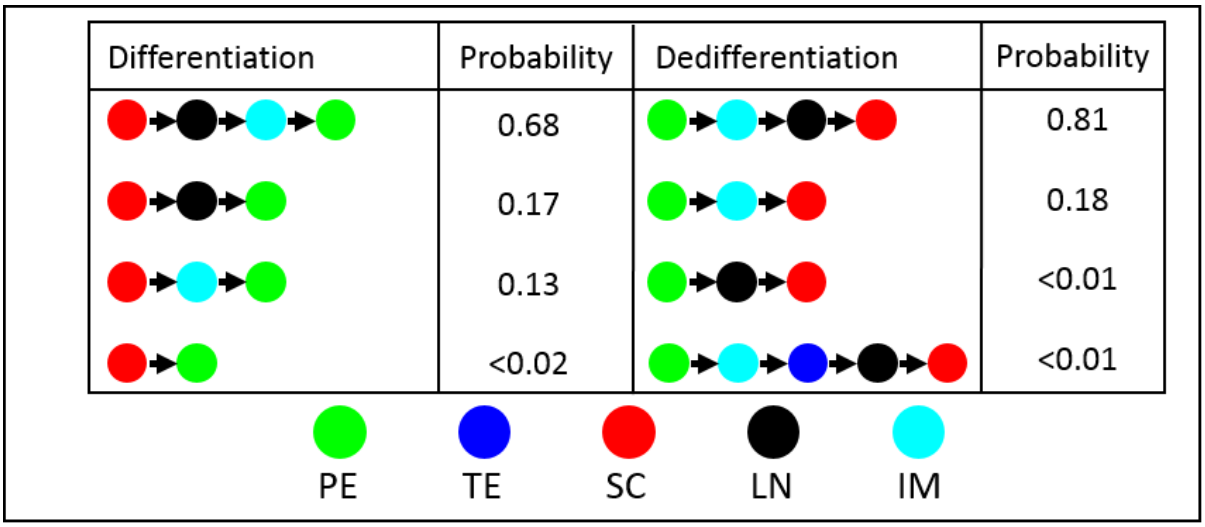

Figure 6.20: **Pathway decomposition for the SC → PE transition for** $f = 50$

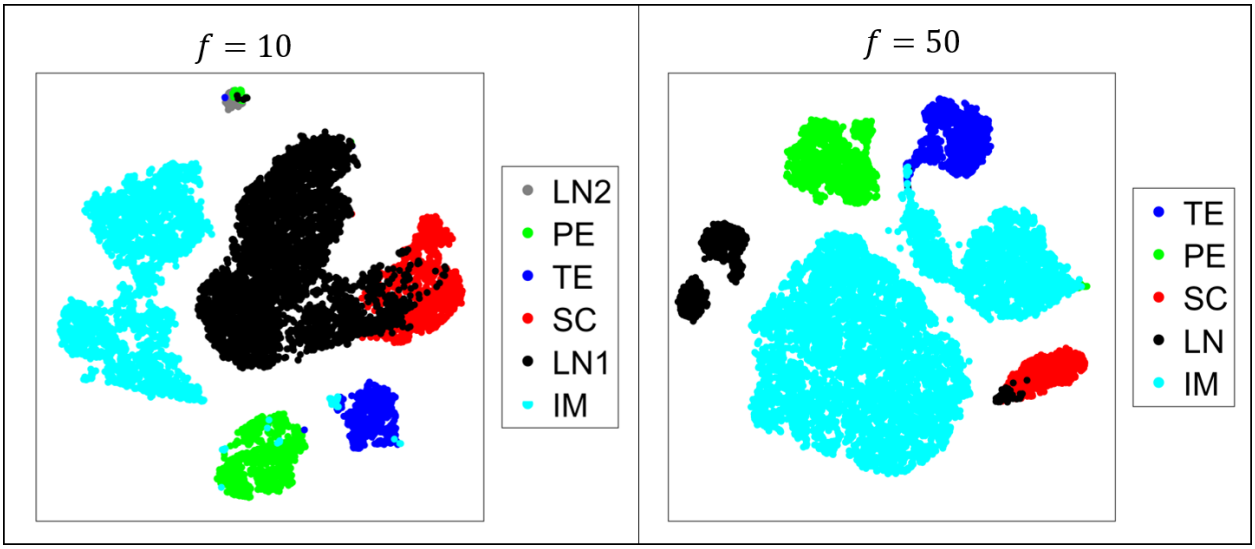

Figure 6.21: **tSNE plot of replica populations with rescaled weight** $(\ln(P_s) + 30)$

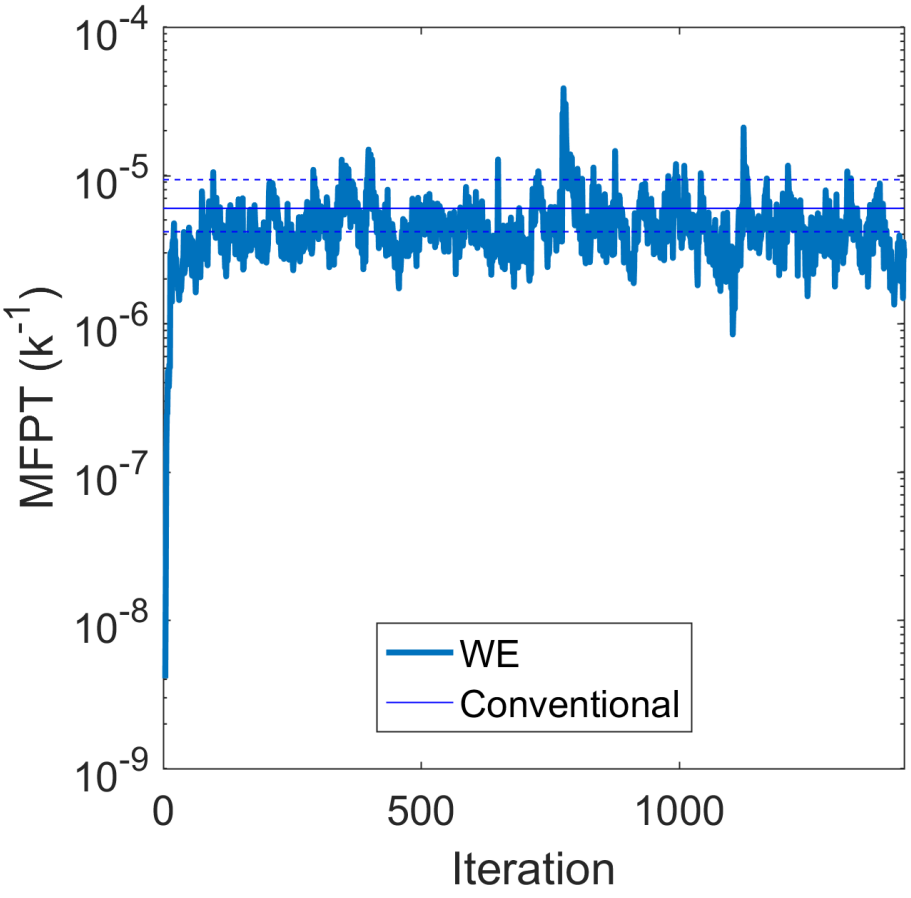

Figure 6.22: **Convergence of the flux of the TE $\rightarrow$ SC transition in the pluripotency network with** $f = 10$. The 5% and 95% confidence intervals for the long conventional simulation are shown in dotted blue lines. The flux was calculated using WE sampling with parameters: $\tau = 50$, 250 bins, and 100 replicas per bin. The system was sampled for 1400 iterations of τ.

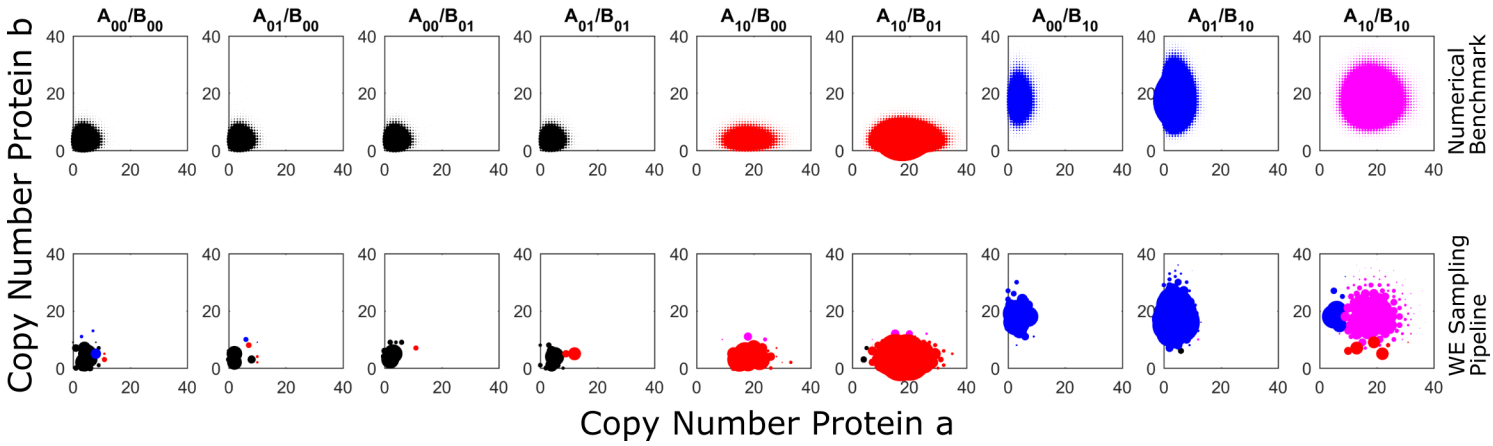

Figure 6.23: **Difference in Coarse-Grained clustering for the 2-gene ExMISA cell decision network studied through the numerical benchmark (top) and the WE sampling pipeline (bottom).** The color of each coarse-grained phenotype cluster corresponds to the expression level of protein a/b: lo/lo (black), hi/lo (red), lo/hi (blue), hi/hi (magenta). All enumerated state phenotypes are sized proportionally to their probability for each of the nine gene configurations for the numerical benchmark, while only the centers of each sampling region are shown for the WE sampling computational pipeline. While the centers of the sampling regions are mostly well separated according to their gene configuration, the phenotype states assigned to the sampling region extend across multiple gene configurations due to the choice of euclidean distance metric in assigning phenotype states to sampling regions.

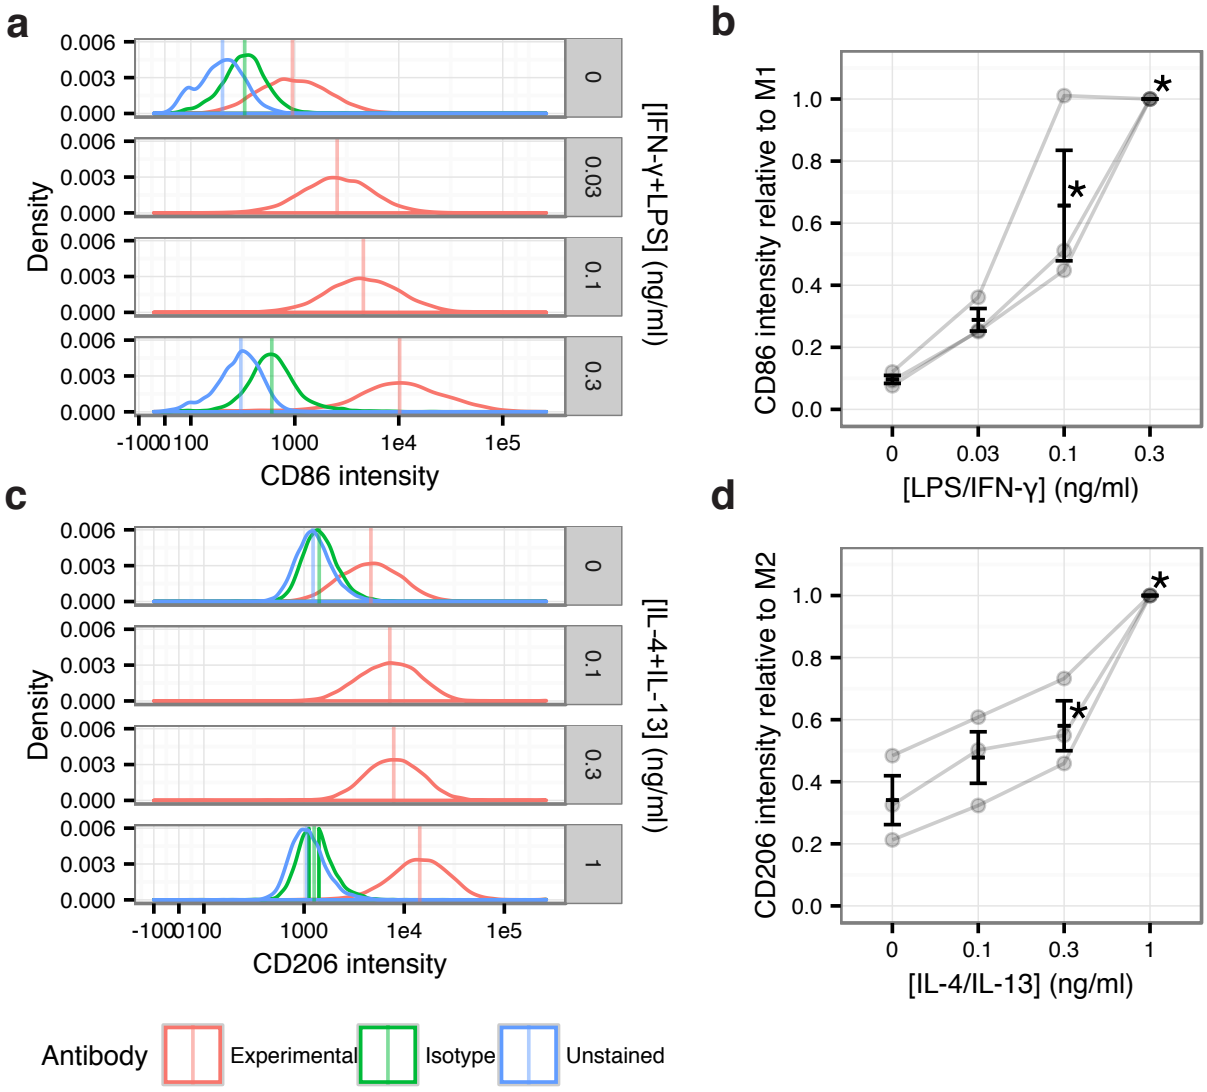

Figure 6.24: **Expression of phenotypic markers CD86 and CD206 in response to LPS/IFN-γ and IL-4/IL-13 stimuli is dose-dependent.** (A) Representative flow cytometry histograms of CD86 intensity after 48 hours of treatment with indicated dose of LPS/IFN-γ. (B) Average median normalized CD86 expression $\pm$ SEM as a function of increasing LPS/IFN-γ dose. Data are normalized to 0.3 ng/ml treatment condition. Asterisk indicates difference vs. untreated, $p < 0.05$; $n = 3$. (C) Representative flow cytometry histograms of CD206 intensity after 48 hours of treatment with indicated dose of IL-4/IL-13. (D) Average median normalized CD206 intensity $\pm$ SEM as a function of increasing IL-4/IL-13 dose. Data are normalized to 1 ng/ml treatment condition. Asterisk indicates difference vs. untreated by two-sided t test, $p < 0.05$; $n = 3$.

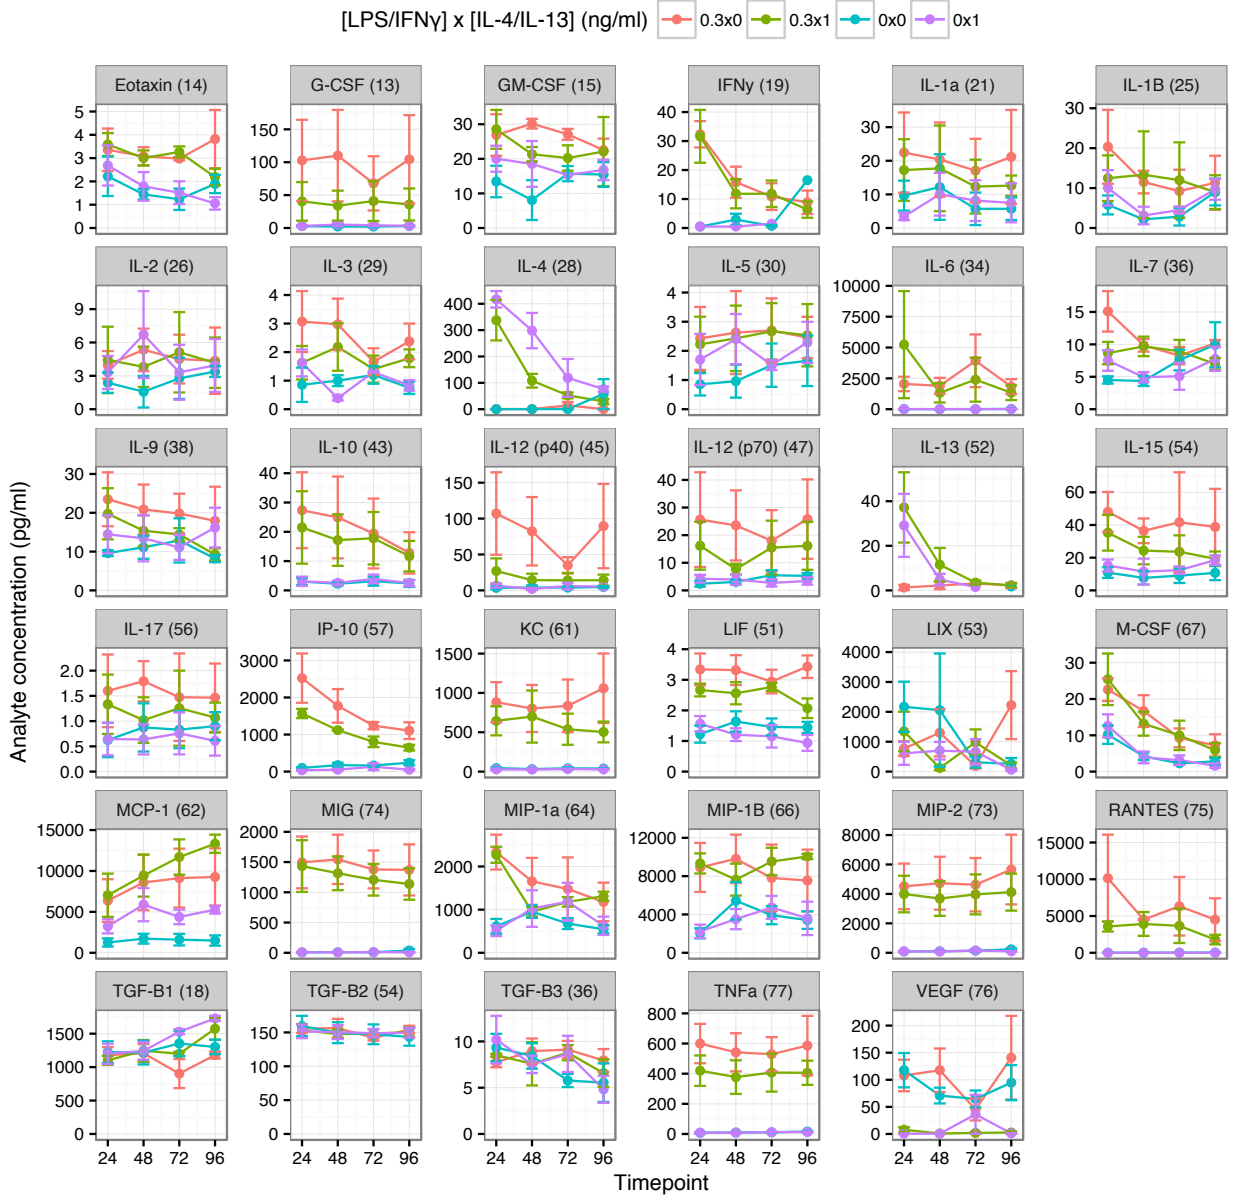

Figure 6.25: **Cytokines present in culture media of macrophages exposed to LPS/IFN-γ $\pm$ IL-4/IL-13 for the indicated time in hours, assessed by multiplex ELISA.** Data presented as mean $\pm$ SEM, n=3.

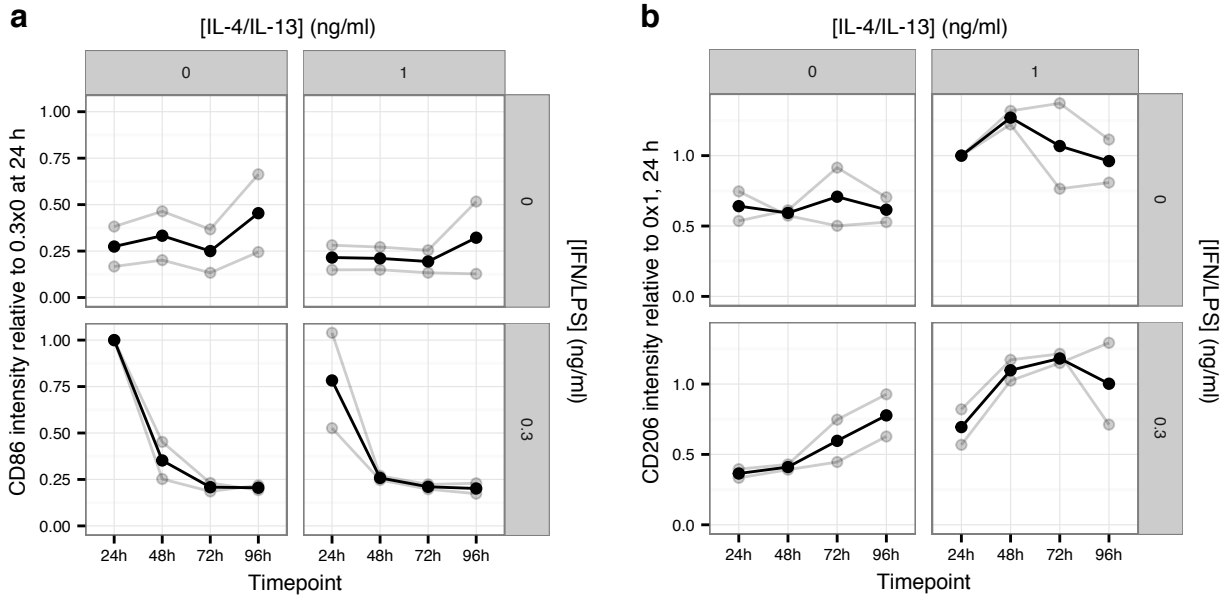

Figure 6.26: **Pulse-chase experiment**. Cells were treated with LPS/IFN-γ or IL-4/IL-13 at t=0. Media was replaced with untreated media at t=24 hours. Cells were collected at t=24, 48, 72, and 96 hours for analysis of CD86 and CD206 expression by flow cytometry to observe how marker expression evolved over time in the absence of continued stimulus. Median normalized fluorescence intensity from each experiment (n=2) is plotted in gray and the mean is plotted in black.

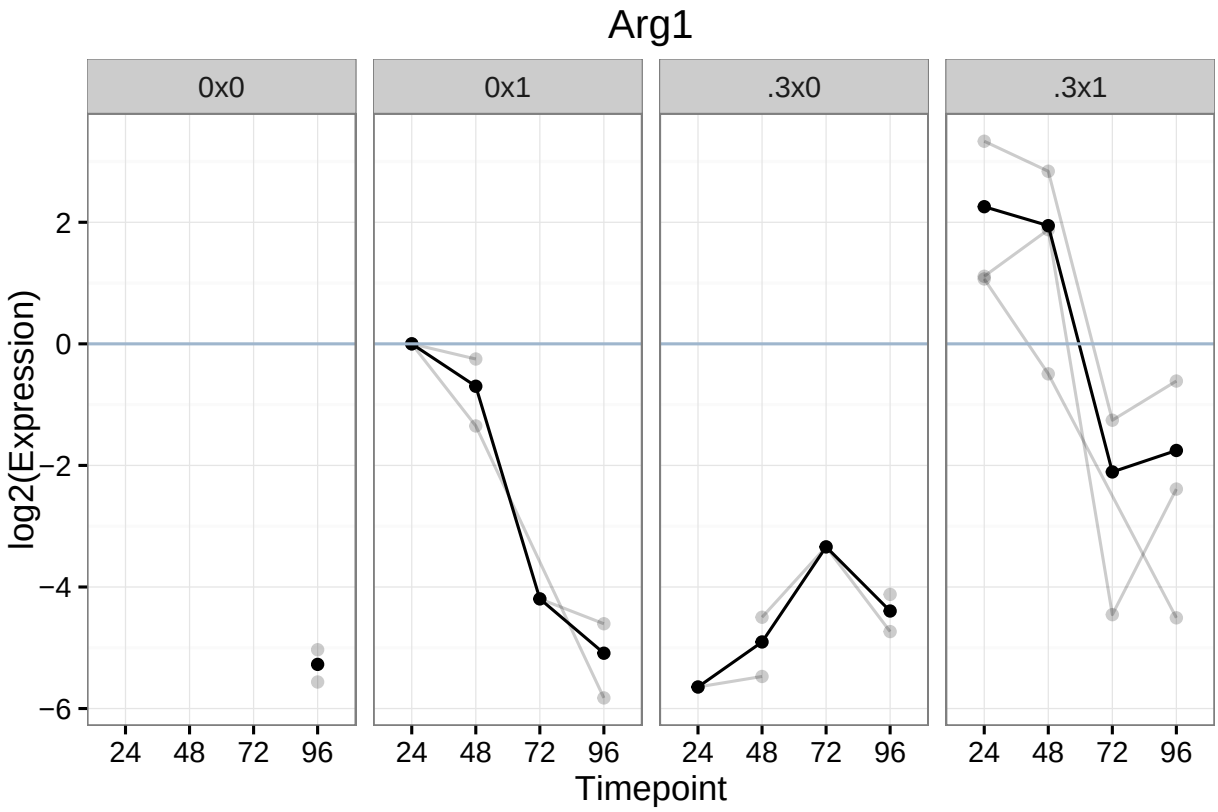

Figure 6.27: **qRT-PCR timecourse data for *Arg1* expression, relative to the IL-4/IL-13-only condition at 24 hours.** Gray points are individual observations and gray lines connect points from the same experiment; the black points and line represent the average. Missing points indicate missing data due to signals below limit of quantitation. Headings describe stimulation condition as ng/ml concentration of LPS/IFN-γ x IL-4/IL-13. Timepoint is in hours.

6.8 Supporting Tables

Table 6.1: **String parameters for each WE string simulation**

String Parameters	**String Nodes** ($N_{\mathbf{str}}$)	ξ	κ	N_{rep}	T_{move}	T_{avg}	τ $[t^{-1}]$
Baseline							
I	20	0.0025	0.005	150	10	5	0.003
II	20	0.0025	0.005	150	10	5	0.05
Explicit Dimerization	20	0.0025	0.01	150	10	5	0.5

The number of string nodes (N_{str}), the string movement parameter (ξ), the number of replicas (N_{rep}), the number of simulations steps between string movements (T_{move}), and the simulations steps used to move the string (T_{avg}), are kept the same for all simulations. Definitions of the parameters are the same as in Adelman and Grabe ([6]).The string smoothing parameter κ is modified between the base and the explicit dimerization networks. The simulation length τ is maintained for the base parameters II-III, but reduced for parameter I.

Table 6.2: **ExMISA Network Parameters**

ExMISA Parameters	**Value in** $[1/k]$	**Description**
g_0	4.0	basal/ repressed expression rate
g_1	18.0	activated expression rate
h_a	1×10^{-5}	binding rate of activator
h_r	1×10^{-1}	binding rate of repressor
f_a	1×10^{-5}	unbinding rate of activator
f_r	1	unbinding rate of repressor
k	1	transcription factor degradation rate

Parameters of the ExMISA network

Table 6.3: **Pluripotency Network**

Gene	**Activators**	**Repressors**
$PBX1$	NANOG	—
$CDX2$	CDX2	NANOG, OCT4
$NANOG$	PBX1, OCT4-SOX2,KLF4	NANOG, GATA6
$GATA6$	GATA6, OCT4-SOX2	NANOG, OCT4
$GCNF$	CDX2, GATA6	—
$KLF4$	NANOG, OCT4, SOX2	—
$OCT4$	OCT4-SOX2	GCNF, CDX2
$SOX2$	OCT4-SOX2	—

Interaction rules for genes in the pluripotency network.

Table 6.4: **Pluripotency Network Parameters**

Parameter	Set I	Set II	**Description**
g_{off}	100	100	basal/ repressed expression rate
g_{on}	3900	3900	activated expression rate
h	1×10^{-5}	5×10^{-5}	binding rate of transcription factor
f	10	50	unbinding rate of transcription factor
k	1	1	transcription factor degradation rate

Parameters of the pluripotency network in units of k^{-1}.

Table 6.5: **Weighted Ensemble Simulation Parameters**

WE Parameters	**ExMISA (Voronoi Movement)**	**ExMISA (Transition Matrix Mode)**	**Pluripotency f = 10 and f = 50 (Voronoi Movement)**	**Pluripotency f = 10 and f = 50 (Transition Matrix Mode)**
τ	10000	10000	10	10
Simulation regions	300	300	250	250
Replicas per region	100	100	500	500
Iterations	60	600	60	600

WE parameters for all networks.

Table 6.6: **Transition Matrices of Metastable Phenotype Clusters (MSMs)**

ExMISA Network				
	State 1 (lo/lo)	**State 2 (hi/hi)**	**State3 (lo/hi)**	**State 4 (hi/lo)**
State 1 (lo/lo)	8.96×10^{-1}	5.63×10^{-4}	5.29×10^{-2}	5.05×10^{-2}
State 2 (hi/hi)	1.24×10^{-4}	9.53×10^{-1}	2.22×10^{-2}	2.45×10^{-2}
State 3 (lo/hi)	1.16×10^{-2}	1.99×10^{-2}	9.68×10^{-1}	5.52×10^{-4}
State 4 (hi/lo)	1.05×10^{-2}	2.15×10^{-2}	6.78×10^{-4}	9.67×10^{-1}

Pluripotency Network Parameter Set I						
	State 1 (LN2)	**State 2 (PE)**	**State3 (TE)**	**State 4 (SC)**	**State 5 (LN1)**	**State 6 (IM)**
State 1 (LN2)	2.97×10^{-1}	2.50×10^{-1}	6.04×10^{-2}	3.47×10^{-3}	2.68×10^{-1}	2.07×10^{-2}
State 2 (PE)	1.42×10^{-3}	8.90×10^{-1}	3.06×10^{-4}	1.83×10^{-4}	1.11×10^{-2}	9.66×10^{-2}
State 3 (TE)	1.91×10^{-4}	2.11×10^{-4}	8.03×10^{-1}	3.03×10^{-7}	1.00×10^{-4}	1.96×10^{-1}
State 4 (SC)	5.09×10^{-12}	4.30×10^{-6}	9.34×10^{-6}	4.30×10^{-1}	5.70×10^{-1}	1.39×10^{-7}
State 5 (LN1)	2.06×10^{-6}	5.16×10^{-6}	3.53×10^{-5}	5.16×10^{-2}	9.48×10^{-1}	8.20×10^{-6}
State 6 (IM)	2.72×10^{-7}	4.14×10^{-4}	1.64×10^{-3}	1.01×10^{-9}	1.36×10^{-6}	9.98×10^{-1}

Pluripotency Network Parameter Set II					
	State 1 (TE)	**State 2 (PE)**	**State3 (SC)**	**State 4 (LN)**	**State 5 (IM)**
State 1 (TE)	8.05×10^{-1}	2.66×10^{-6}	1.31×10^{-6}	2.92×10^{-3}	1.92×10^{-1}
State 2 (PE)	3.09×10^{-7}	9.20×10^{-1}	3.13×10^{-7}	1.65×10^{-4}	7.98×10^{-2}
State 3 (SC)	4.70×10^{-8}	2.21×10^{-7}	8.23×10^{-1}	1.77×10^{-7}	2.38×10^{-6}
State 4 (LN)	5.25×10^{-9}	4.92×10^{-8}	6.86×10^{-3}	9.93×10^{-1}	5.02×10^{-7}
State 5 (IM)	1.60×10^{-9}	8.22×10^{-9}	2.22×10^{-9}	1.01×10^{-8}	9.99×10^{-1}

Markov State Models of metastable phenotype-cluster transitions found through the computational pipeline for all three simulated networks. There are four different combinations of a/b protein expression levels in the coarse-grained phenotype network of the ExMISA network: lo/lo, hi/hi, lo/hi, and hi/lo. The steady-state probabilities of the lo/lo, hi/hi, lo/hi, and hi/lo cell phenotypes are predicted by the computational pipeline to be 1.71×10^{-2}, 7.67×10^{-2}, 3.71×10^{-1}, 3.80×10^{-1}, respectively. The steady state probabilities of the six and five coarse-grained phenotype networks in the pluripotency network Parameter Set I and Parameter Set II can be found in figures 3.4 and 3.6, respectively.

Table 6.7: **Computed Mean First Passage Times in the ExMISA Network– Comparison of Different Methods**

Method	**Start/End State**	**MFPT**	**5% Confidence**	**95% Confidence**
CME - numeric benchmark	Basin centers	1.84×10^5	—	—
Conventional SSA simulation	Basin centers	1.82×10^5	1.67×10^5	1.98×10^5
WE - rate mode	Basin centers	1.82×10^5	1.78×10^5	1.85×10^5
WE - transition matrix mode	Coarse-grained polarized phenotype	2.34×10^5	—	—
Coarse-grained phenotype network	Coarse-grained polarized phenotype	1.70×10^5	—	—

Computed Mean First Passage Times (MFPTs, time-units k^{-1}) of the ExMISA network, using different computation methods. For each row, $\mathrm{MFPT}_{XY} = \mathrm{MFPT}_{YX}$ due to symmetry in the network, and the start- and end-state (X and Y) for the transition are defined either with respect to distance from the centers of the polarized phenotype basins, or in terms of aggregated states in the coarse-grained phenotype definition. For basin centers, State X is defined as a hypersphere of radius 1 centered around the state vector [4,16,0,0,1,0,1,0], corresponding to the species: [a,b,A_{00},A_{10},A_{01},B_{00},B_{10},B_{01}]. State Y is a hypersphere centered around [16,4,0,1,0,0,0,1]. For the coarse-grained phenotype definition, states correspond to the polarized a/b hi/lo and lo/hi phenotypes.

Table 6.8: **Computed Mean First Passage Times of Inter-Phenotype Transitions in the Pluripotency Network (Parameter Set I)**

Method	**Transition**			
	SC → LN1	**LN1 → SC**	**SC → TE**	**TE → SC**
CME	—	—	—	—
ConventionalSSA	$1.85(1.75, 2.33) \times 10^1$	$2.69(2.30, 3.43) \times 10^2$	$1.65(1.07, 2.34) \times 10^5$	$3.59(2.50, 5.18) \times 10^5$
Weighted Ensemble Rate Mode	$1.71(1.64, 1.78) \times 10^1$	$1.94(1.85, 2.05) \times 10^2$	$1.36(1.02, 1.77) \times 10^5$	$2.70(2.48, 2.91) \times 10^5$
Weighted Ensemble Transition Matrix Mode	2.09×10^1	4.13×10^2	2.19×10^5	2.21×10^5
MSM	2.23×10^1	3.89×10^2	2.06×10^5	2.20×10^5

The Mean First Passage Times of NANOG Fluctuation and (De)differentiation in the pluripotency network calculated using $\tau = 10$ and $f = 10$. The MFPT reported for the WE is a block average over the last 500 iterations. The MFPT and standard deviation are found for transitioning between the stem cell phenotype (SC) and the pluripotent phenotype with low NANOG expression (LN1) (analogous to high NANOG production (N^{hi}) and low NANOG production (N^{lo}) transitions measured in experiments) and for transitioning between the stem cell phenotype (SC) and the trophectoderm phenotype (TE), calculated on the timescale of the protein degradation rate k. The SC, TE, and LN1 regions of interest (ROI) are defined as the SC, TE, and LN1 phenotypes derived from the MSM reduction of the sampled transition matrix.

Table 6.9: **qPCR primers**

Gene	Direction	Sequency	Amplicon length (bp)
Arg1	F	CTCTGTCTTTTAGGGTTACGG	152
	R	CTCGAGGCTGTCCTTTTGAG	
Chi3l3	F	AGTGCTGATCTCAATGTGGATTC	142
	R	TAGGGGCACCAATTCCAGTC	
Gapdh	F	GTCAAGCTCATTTCCTGGTATGAC	131
	R	TCTCTTGCTCAGTGTCCTTGC	
Hprt	F	TGGACAGGACTGAAAGACTTGCTCG	81
	R	CCTTGAGCACACAGAGGGCCAC	
Il10	F	CCCACTTCCCAGTCGGCCAG	300
	R	GGAGAAATCGATGACAGCGCCTC	
Kdm6b	F	GGTTCACTTCGGCTCAACTTAG	75
	R	CTCCACCGTATGTTCACCGC	
Ldha[A]	F	TGTCTCCAGCAAAGACTACTGT	155
	R	GACTGTACTTGACAATGTTGGGA	
Mrc1	F	TGTTTTGGTTGGGACTGACC	269
	R	TGCAGTAACTGGTGGATTGTC	
mVPA1 [79]	F	GGAGCCCAGTGTAGAAGAGCA	87
	R	AGCCAGCGAACCATATCCTGA	
Nos2	F	TTGGGTCTTGTTCACTCCAC	211
	R	TGTATTGTTGGGCTGAGAACAG	
Retnla	F	GCCAATCCAGCTAACTATCCC	187
	R	AGTCAACGAGTAAGCACAGG	
Sdha[B]	F	CTTGAATGAGGCTGACTGTG	87
	R	ATCACATAAGCTGGTCCTGT	
Tnfa	F	CCCACGTCGTAGCAAACCACCA	172
	R	TCGGGGCAGCCTTGTCCCTT	

[A] RTPrimerDB[7] 3720; [B] RTPrimerDB 3875

Table 6.10: **Macrophage Model Parameter Definitions**

Parameter	Meaning
k_1	Maximum stimulation rate of M1 cascade under induction with S1
k_2	Maximum stimulation rate of M2 cascade under induction with S2
k_3	Maximum stimulation rate of M1 cascade under self-activation
k_4	Maximum stimulation rate of M2 cascade under self-activation
k_5	Basal rate of M1 activation
k_6	Basal rate of M2 activation
k_7	Maximum rate of M2 stimulation from Y and M2 cooperative stimulation
k_8	Maximum rate of Y production under S1 induction
K_Y	Level of Y to reach half-maximum inhibition of M1
K_{CY}	Level of Y to reach half-maximum cooperative activation of M2
K_{CM2}	Level of M2 to reach half-maximum cooperative activation of M2
K_{rep1}	Level of M1 to reach half-maximum inhibition of M2
K_{rep2}	Level of M2 to reach half-maximum inhibition of M1
K_{ind1}	Level of S1 to reach half-maximum induction of M1
K_{ind2}	Level of S2 to reach half-maximum induction of M2
K_{act}	Level of M1 or M2 to reach half-maximum self-activation
d_1	M1 decay rate
d_2	M2 decay rate
d_3	Y decay rate
d_4	CD86 decay rate
d_5	CD206 decay rate
g_1	CD86 production rate
g_2	CD206 production rate
n	Hill coefficient

Table 6.11: **Macrophage Model Parameter Values**

Model	1	2	3	4	5	6
k_1	0.4341	0.5871	10.421	0.7595	2.2698	2.8511
k_2	1.1124	1.6691	1.4815	1.4998	0.8939	1.0146
k_3	0.8322	0.791	1.0512	0.6856	1.1102	0.7833
k_4	1.6865	0.311	0.9216	0.4332	1.5412	1.1121
k_5	0.0456	0.0679	0.0421	0.0968	0.0332	0.1067
k_6	0.0172	0.5813	0.3239	0.541	0.0831	0.1589
k_7	N/A	N/A	0.1977	4.3814	2.8281	2.2411
k_8	N/A	N/A	0.1096	0.1594	0.7288	0.1206
K_{M2}	N/A	N/A	N/A	9.9178	1.3292	0.0012
K_{CY}	N/A	N/A	0.3438	8.4118	5.4604	5.4965
K_Y	N/A	N/A	N/A	N/A	0.0209	0.9182
K_{rep1}	N/A	2.8882	2.2441	1.9004	2.4051	1.0306
K_{rep2}	N/A	14.228	1.4216	2.0858	1.0162	1.195
K_{ind1}*	1	1	1	1	1	1
K_{ind2}*	0.3	0.3	0.3	0.3	0.3	0.3
K_{act}*	1	1	1	1	1	1
d_1*	1	1	1	1	1	1
d_2*	1	1	1	1	1	1
d_3*	N/A	N/A	0.05	0.05	0.05	0.05
d_4*	0.05	0.05	0.05	0.05	0.05	0.05
d_5*	0.05	0.05	0.05	0.05	0.05	0.05
g_1*	1	1	1	1	1	1
g_2*	1	1	1	1	1	1
n	N/A	2	N/A	2	2	2
Free parameters	6	8	11	12	13	13

Representative best fit parameter values for each model from optimization. Parameters are in arbitrary units of concentration and time, relative to the rate of degradation of the M1 species (d_1), which is approximated to be 1 [/hr] according to the half-life of STAT1 [82]. The parameters were optimized to normalized CD86 and CD206 expression levels. Parameters with an asterisk were fixed to constrain parameter space during optimization. Fixed values were chosen based on initial parameter searches. Alternative constraints yielded different quantitative values, but the same ordering of model scores according to the AICc. The threshold parameters for induction, K_{ind1} and K_{ind2} are based on the dose-response of CD86 and CD206 under the single-stimulus conditions. K_{act} is approximated from the experimental data condition with no induction stimulus at 24 hours. A Hill coefficient of 2 was used for all parameter sets. Parameter sets estimated using a Hill coefficient of 1 produced AICc scores equivalent or worse than the AICc scores using a Hill coefficient of 2.