2015

Physics of epigenetic landscapes and statistical inference by cells

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

PHYSICS OF EPIGENETIC LANDSCAPES AND STATISTICAL INFERENCE BY CELLS

by

ALEX H. LANG

B.S., University of Wisconsin-Madison, Madison, USA, 2010
M.S., Boston University, Boston, USA, 2012

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
2015

Approved by

First Reader

Pankaj Mehta, Ph.D.
Assistant Professor of Physics

Second Reader

Kirill Korolev, Ph.D.
Assistant Professor of Physics

Acknowledgments

I am extremely grateful for many people who helped me successfully complete my PhD over these past five years.

First I would like to thank my research advisor Pankaj Mehta. He has provided a great environment for me to achieve my full research potential. Pankaj provided the right mix of support and independence to allow me to develop into a complete scientist.

Besides Pankaj, I would like to thank the rest of the research group. I had many great conversations and learned countless number of facts from Javad Noorbakhsh, Charles Fisher, Steven Ahn, Joseph Samaniego-Evans, and Ben Dickens.

I would also like to thank the my PhD committee Kirill Korolev, Claudio Chamon, Martin Schmaltz, Ahmed Khalil, Laertis Ikonomou for their time and dedication to helping me graduate.

I would also like to thank my collaborators on various projects that I have yet to mention. These include Darrell Kotton, Katie Benson, Keri Dame, Andrew Wilson and the rest of the BU Center for Regenerative Medicine (CReM), Daniel Segre and the rest of the COMETs collaboration, David Schwab, Jim Collins, Hu Li, Horacio Castillo, and Sai Teja Pusuluri.

During my time at Boston University, I have made many great friends that made my time in Boston extremely fun. I would like to especially thank Gustavo Marques Tavares, Maira Constantino, Colin Howard, Maria Gonzalez, Alex Long, Jeff Gustafson, Joe Hardcastle, Michael Chernicoff, and Nick Lubbers.

I would like to thank my parents, Hunter and Laura, and my siblings, Amanda, Aaron, and Amy, who have been very supportive of me throughout my life.

Finally and most importantly, I would like to thank my lovely wife Tina. Without her love and support I would never have finished my PhD. I also should thank my dog Buddy for secretly doing all of my research.

PHYSICS OF EPIGENETIC LANDSCAPES AND STATISTICAL INFERENCE BY CELLS

(Order No.)

ALEX H. LANG

Boston University Graduate School of Arts and Sciences, 2015

Major Professor: Pankaj Mehta, Assistant Professor of Physics

ABSTRACT

Biology is currently in the midst of a revolution. Great technological advances have led to unprecedented quantitative data at the whole genome level. However, new techniques are needed to deal with this deluge of high-dimensional data. Therefore, statistical physics has the potential to help develop systems biology level models that can incorporate complex data. Additionally, physicists have made great strides in understanding non-equilibrium thermodynamics. However, the consequences of these advances have yet to be fully incorporated into biology.

There are three specific problems that I address in my dissertation. First, a common metaphor for describing development is a rugged "epigenetic landscape" where cell fates are represented as attracting valleys resulting from a complex regulatory network. I introduce a framework for explicitly constructing epigenetic landscapes that combines genomic data with techniques from spin-glass physics. The model reproduces known reprogramming protocols and identifies candidate transcription factors for reprogramming to novel cell fates, suggesting epigenetic landscapes are a powerful paradigm for understanding cellular identity.

Second, I examine the dynamics of cellular reprogramming. By reanalyzing all available time-series data, I show that gene expression dynamics during reprogramming follow a simple one-dimensional reaction coordinate that is independent of both the time and details

of experimental protocol used. I show that such a reaction coordinate emerges naturally from epigenetic landscape models of cell identity where cellular reprogramming is viewed as a "barrier-crossing" between the starting and ending cell fates. Overall, the analysis and model suggest that gene expression dynamics during reprogramming follow a canonical trajectory consistent with the idea of an "optimal path" in gene expression space for reprogramming.

Third, an important task of cells is to perform complex computations in response to external signals. Intricate networks are required to sense and process signals, and since cells are inherently non-equilibrium systems, these networks naturally consume energy. Since there is a deep connection between thermodynamics, computation, and information, a natural question is what constraints does thermodynamics place on statistical estimation and learning. I modeled a single chemical receptor and established the first fundamental relationship between the energy consumption and statistical accuracy of a receptor in a cell.

Contents

List of Figures

List of Abbreviations

ATP	Adenosine triphosphate
C+OSKM	CEBPα plus OSKM
CLP	Common Lymphoid Progenitor
CMP	common myeloid progenitor
DNA	Deoxyribonucleic acid
EpiSC	epiblast stem cell
ESC	embryonic stem cell
GFP	green flourescent protein
GMP	granulo-monocytic progenitors
GPCR	G-protein coupled receptors
HM	histone modification
iCM	induced cardiomyocytes
iHep	induced hepatocytes
iN	induced neuron
iNPC	induced NPC
iPSC	induced pluripotent stem cells
K4	Histone 3 tri-methylation at lysine 4
K27	histone 3 tri-methylation at lysine 27
MEF	mouse embryonic fibroblasts
MEP	megakaryocyte-erythroid progenitors
MLE	maximum likelihood estimation
mRNA	messenger RNA
MSC	Mesenchymal stem cells
NESS	nonequilibrium steady states

NK	Natural Killer cells
NPC	neural progenitor cells
NSC	neural stem cells
OSKM	Yamanaka reprogramming factors: Oct4, Sox2, Klf4, and Myc
PC	principal components
PCA	principal components analysis
PRC	Partially reprogrammed cells
RNA	Ribonucleic acid
RNA-Seq	RNA sequencing
SAS	sensory adaptive system
STD	Standard Deviation
TF	transcription factors

1 Introduction

The section Epigenetic Landscapes and Cellular Identity has been adapted from Pankaj Mehta's 2012 NIH New Innovator Grant. The section Statistical Inference by Cells has been adapted from Landauer in the age of synthetic biology: energy consumption and information processing in biochemical networks by Pankaj Mehta, Alex H. Lang, and David J. Schwab. arXiv 1505.02474.

1.1 Overview

Biology is currently in the midst of a revolution. The human genome project successfully sequenced the entire genome and heralded in an era of unprecedented quantitative data at the whole genome level [1]. While the human genome project took a giant collaboration over a decade to complete, now microarrays and RNA-Seq have ushered in a new genomic phase in biology in which any lab can quickly and cheaply produce genomic data. Labs routinely measure and share whole-genome expression data, leading to terrabytes of data publicy available to any researcher. This explosion of genomic data has been deemed the biological Big Data revolution [2].

Parallel theoretical advances have led to an emerging area of biology: systems biology. Since the discovery of DNA, biology has focused on single genes and their proteins. However, systems biology shifts the emphasis to the cellular level and is most interested in networks of interacting genes and proteins and their emergent phenotypic behaviors [3]. This systems level thinking fits in naturally with statistical physics, the subfield of physics devoted to thinking about the universal properties of large systems. Therefore, statistical physics has and will continue to provide novel tools for systems biology and the Big Data revolution.

Physicists have also recently made great strides in understanding non-equilibrium thermodynamics. While physicists have mainly focused on simple physical systems such as cold atoms, biology is fundamentally a non-equilibrium system driven at the cellular level by ATP. Therefore, the advances made in non-equilibrium thermodynamics need to be fully

incorporated into biology.

My PhD research focused on two distinct areas. First, I used statistical physics models to address an outstanding question in systems biology: what is the nature of cellular identity. This model incorporates gene expression data to explicitly model cellular identity. Second, I created a non-equilibrium model of a cellular chemical sensor and analyzed its properties, focusing on the key differences between equilibrium and non-equilibrium sensors.

My thesis is organized as follows. In the Introduction, I will first provide the background information on epigenetic landscapes and cellular identity in section 1.2, and then I will give an overview of the intersection of non-equilibrium thermodynamics and biological implementations of statistical inference in section 1.3. Then in chapter 2 I will introduce the model for cellular identity landscapes and analyze the static properties. In chapter 3 I will use the previously introduced model to examines the dynamics of cellular reprogramming. Finally, in chapter 4 I will introduce a model for non-equilibrium biological sensors and analyze their properties.

1.2 Epigenetic Landscapes and Cellular Identity

1.2.1 Cellular Reprogramming

Cellular reprogramming, the direct reprogramming of one cell type into another, holds enormous therapeutic potential for studying and treating human disease [4]. Epigenetic reprogramming of differentiated cells into embryonic stem cells (ESC), as well as direct conversion or transdifferentiation into other differentiated cell types, has been achieved by overexpression of select transcription factors and RNAs [5]. Despite these enormous advances, cellular reprogramming of differentiated cells to cells resembling ESCs, dubbed, induced pluripotent stem cells (iPSCs), produces low yields ($< 1\%$) and debate remains about the fidelity or accuracy of iPSCs vs ESCs [4]. Furthermore, so far direct conversion protocols only exist for a handful of cell types [6–11]. For this reason, there is a critical need for new theoretical frameworks and computational tools to understand and manipulate the genetic networks underlying cellular differentiation. A major goal of the first part of

this disserataion was to develop new theoretical and computational tools to help facilitate the development of rational reprogramming strategies between arbitrary cell types. To do so, we utilized an interdisciplinary approach that combines large-scale genomic and epigenetic data sets with methods from statistical physics, machine learning, and computational neuroscience.

The molecular basis of cellular identity and differentiation remains one most the fundamental questions in biology. The work of Takahashi and Yamanaka [12] demonstrating that the overexpression of just four transcription factors is sufficient to convert somatic fibroblasts to IPSCs has revolutionized the field. Great progress has been made in identifying the molecules and genes that shape cellular identity using modern genomic techniques such as whole-genome expression studies and epigenetic profiling [13, 14]. Yet, how these molecular components interact to give rise to the 'epigenetic landscape' underlying cellular identity remains elusive due to the large number of genes and molecules involved, the plethora of interaction between genes, and the importance of stochasticity in cellular reprogramming [15, 16]. My dissertation research focused on building new probabilistic, mathematical models for the biomolecular networks underlying cellular identity using existing genomic and epigenetic data sets. The models build upon insights from statistical physics and attractor neural networks and are intrinsically stochastic and large scale. Despite having no free parameters, our model reproduces known reprogramming protocols to embryonic stem cells, heart (cardiomyocytes), and liver (hepatocytes) and identifies candidate transcription factors for reprogramming to novel cell types, suggesting it represents a powerful paradigm for understanding cellular identity.

The molecular circuitry underlying cellular differentiation is intricate and involves transcriptional regulation, chromatin regulators, non-coding and micro RNAs, and signal transduction pathways [15, 17]. Transcription factors (TFs) play a crucial role in cell type maintenance, cellular differentiation and reprogramming. Transcription factors often act cooperatively to regulate genes by forming large protein complexes and recruiting chromatin regulators [15]. Depending on their chromatin marks, genes can be active, silenced, or

in a ‘poised state [18]. Active genes are commonly associated with trimethylation of histone H3 lysine 4 (H3K4me3), whereas silenced genes are trimethylated at histone H3 lysine 27 (H3K27me3). Poised genes exhibit ‘bivalent markings with both active and repressive marks and are thought to be poised for activation in response to appropriate developmental cues. Importantly, epigenetic modifications are reversible, highly dynamic, and often change depending on cell type [18].

A given cell type is associated with a typical genome-wide expression and epigenetic profile [19–21]. Consequently, it is common in the literature to think of cell type as a systems-level property of the genetic regulatory networks underlying cellular identity [15–17]. In particular, cell type is now thought to arise from the dynamic interaction of components within a regulatory circuit [17, 22]. This perspective raises new questions. Cell type is an extremely stable phenotypic trait. Even pluripotent cells such as stem cells can self-renew and propagate in the absence of external signals [23]. Yet it is possible to reprogram cells by overexpressing just a few carefully chosen transcription factors. In particular, cellular reprogramming protocols exist for reprogramming between differentiated cells and embryonic stem cells [12], closely related somatic cells [24], and more recently, between the distantly related cell types [6–11]. How can one reconcile these two seemingly conflicting observations with the dynamics of the underlying gene regulatory network? A further complication is that cellular reprogramming is stochastic with only a small number of cells switching cellular states [25–27]. Cellular reprogramming can also be incomplete, giving rise to “partially reprogrammed” cells [28, 29]. How do we unite these diverse characteristics and experimental facts?

1.2.2 Existing Models

In light of these observations, there has been a renewed interest in the idea of an ‘epigenetic landscape. The idea of an epigenetic landscape, originally introduced by Waddington in the 1930s to understand development, has become the dominant paradigm for thinking about cellular identity and reprogramming [30–34]. A cells state is represented by a time-

dependent vector in a high-dimensional state space that reflects its time-dependent genetic or epigenetic expression profile [20–22]. In analogy with simple genetic switches [35–37], a cell type is argued to correspond to a stable attractor for the stochastic dynamics of the underlying gene regulatory network. In terms of landscapes, these attractors are visualized as valleys separated by "high-energy" barriers. The barriers ensure robustness of the cell type phenotype against small changes in gene expression but can be traversed by rare stochastic fluctuations as in cellular reprogramming [26].

1.2.3 Model Overview

The landscape analogy offers a attractive framework for thinking about cellular identity and reprogramming. Mounting experimental evidence shows that cell type can be viewed as a dynamical attractor in a high-dimensional state space [20–23, 25, 26]. Consistent with this view are experiments showing that cellular reprogramming can occur by any method that biases gene expression toward a particular attractor including overexpression of transcription factors [12], transfer of the RNA transcriptome [21], or even by expressing cell-type-specific non-coding RNAs [38]. Yet, a quantitative model of the landscape remains elusive. Current mathematical models are confined to small networks consisting of a few genes [39]. This contrasts sharply with experiments that suggest the dynamics of gene regulatory networks governing cellular identity are intrinsically high-dimensional [22]. While dimensional reduction techniques such as Principal Component Analysis (PCA) can help reduce the dimension of the state space, cellular reprogramming between distant cell types usually requires multiple transcription factors suggesting that any quantitative model must be intrinsically high-dimensional. Current mathematical models also fail to account for the complicated epigenetic regulation of genes and the dense web of interactions between regulatory proteins suggested by experiment [13, 18]. Furthermore, current models cannot identify transcription factors that bias cells toward a desired cell-type and have no way of mapping the epigenetic landscape governing cellular identity. These shortcomings highlight the crucial need for new mathematical and theoretical frameworks for understanding

cellular identity and cellular reprogramming on a quantitative level.

The mathematical models proposed here overcome the shortcomings outlined above by combining techniques from statistical physics and computational neuroscience [40–43] with publicly available genomic and epigenetic data sets. The data presented in chapters 2 and 3 suggests our model is feasible. It can quantitatively capture most experimental observations, and most importantly, reproduces current cellular reprogramming protocols. Using this model, we identify transcription factors for direct cellular programming. This model utilizes a correlation-based approach, which have been extremely successful in computational neuroscience and systems biology [40–43]. Underlying a correlation-based approach is the idea that genes whose expression is strongly correlated across cell types are likely to be co-regulated. Correlation-based approaches offer three advantages over more traditional models based on differential equations. First, no explicit knowledge of how genes and proteins interact is required. This is especially important given our incomplete knowledge of the complex molecular interactions that regulate cellular identity [15]. Second, unlike differential equation based approaches, correlation-based approaches are scalable, allowing for high-dimensional models with many interacting components. Third, these correlation models can easily incorporate stochastic dynamics. Taken together, these results suggest that epigenetic landscapes represent a powerful framework for understanding the molecular circuitry and dynamics that gives rise to cell fate.

1.3 Statistical Inference by Cells

1.3.1 Introduction

Cells live in complex and dynamic environments. They sense and respond to both external environmental cues and to each other through cell-to-cell communication. Adapting to changing environments often requires cells to perform complex information processing, and cells have developed elaborate signaling networks to accomplish this feature. These biochemical networks are ubiquitous in biology, ranging from the quorum-sensing [44] and chemotaxis networks [45] in single-celled organisms to developmental networks in higher

organisms [46]. Inspired by both these natural circuits and physical computing devices, synthetic biologists are designing sophisticated synthetic circuits that can perform complicated "computing-like" behaviors. Synthetic biologists have designed gene circuits executing a wide range of functionalities including switches [35], oscillators [47], counters [48], and even cell-communicators [49].

Despite these successes, many challenges to harnessing the full potential of synthetic biology persist [50–55]. While there are guiding principles to synthetic biology [56], actual construction of synthetic circuits often proceeds in an ad-hoc manner through a mixture of biological intuition and trial-and-error. Furthermore, the functionality and applicability is limited by a dearth of biological components [57]. For this reason, it would be helpful to identify general principles that can improve the design of synthetic circuits and help guide the search for new biological parts. One promising direction along these lines is recent work examining the relationship between the information processing capabilities of these biochemical networks and their energetic costs (technically this is usually a cost in free energy, but for the sake of simplicity we will refer to this as energy). Energetic costs place important constraints on the design of physical computing devices [58] as well as on neural computing architectures in the brain and retina [59], suggesting that thermodynamics may also influence the design of cellular information processing networks. As the field of synthetic biology seeks to assemble increasingly complex biochemical networks that exhibit robust, predictable behaviors, natural questions emerge: What are the physical limitations (thermodynamic and kinetic) on the behavior and design of these biological networks? How can one use energy consumption improve the design of synthetic circuits?

In a classic paper written at the advent of modern computing [58], Landauer asked analogous questions about physical computing devices. He argued that a central component of any general purpose computing device is a memory module that can be "reset" to a predefined state. He pointed out that such a device must obey certain thermodynamic and kinetic constraints. In particular, he convincingly argued that resetting memory necessarily leads to power dissipation, implying that heat generation and energy consumption are

unavoidable consequences of the computing process itself. The paper also outlined three general sources of error resulting from kinetic and thermodynamic considerations: incomplete switching between memory states due to long switching times, the decay of stored information due to spontaneous switching, and what he called a "Boltzmann" error due to limited energy supplies. Furthermore, the paper showed that there exist fundamental trade-offs between these types of errors and energetic costs in these memory devices. These considerations suggested general strategies for designing new devices and parts for physical memory modules.

The goal of this review is to synthesize recent theoretical work on thermodynamics and energy consumption in biochemical networks and discuss the implications of this work for synthetic biology. Theoretical papers in this field are often highly technical and draw on new results in non-equilibrium statistical mechanics. For this reason, our goal is to organize the insights contained in these papers [60–83] into a few simple, broadly applicable principles. We find that energy consumption in cellular circuits tends to serve four basic purposes: (1) increasing specificity, (2) reducing variability and increasing precision, (3) signal amplification and (4) erasing information. Furthermore, for each of these categories, there exist implicit tradeoffs between power consumption and dynamics.

Beyond synthetic biology, biochemical networks offer a unique setting to explore fundamental physics questions in non-equilibrium statistical mechanics. Recently there has been a surge of interest among physicists in the relationship between information and thermodynamics [84, 85]. For example, using sophisticated optical traps groups have recently experimentally tested Landauers principle [86, 87], and there is an active debate on how to extend Landauers principle to quantum regimes. A flurry of recent works has focused on extending concepts like entropy and free-energy to non-equilibrium regimes, often using information theoretic concepts [88–94]. Living systems are perhaps the most interesting example of non-equilibrium systems, and thinking about information and thermodynamics in the context of cells is likely to yield new general insights into non-equilibrium physics.

1.3.2 Increasing Specificity

One common role of energy consumption in biochemical circuits is to increase the specificity of an enzyme or signaling pathway. The most famous example of this is kinetic proofreading. In a landmark paper [95], John Hopfield showed how it is possible to increase the specificity of an enzyme beyond what would be expected from equilibrium thermodynamics by consuming energy and driving the system out of equilibrium. Kinetic proofreading-type mechanisms are also thought to underlie the exquisite specificity of eukaryotic pathways such as the TCR signaling network [96], in which a few-fold difference in the affinities between molecules can lead to several orders of magnitude difference in response. A full review of kinetic proofreading and all its applications is beyond the scope of this review, but we highlight some important lessons for synthetic biology.

The first general principle that emerges from kinetic proofreading is that greater specificity requires greater energy consumption. In particular, the error rate in kinetic proofreading depends exponentially on the amount of energy consumed in each step of the proofreading cascade. This increased specificity comes at the expense of a more sluggish dynamic response (see [78, 97] for an interesting exploration of this tradeoff). This highlights a second general theme about energy consumption: there generally exist trade-offs between greater specificity and other desirable properties such as a fast dynamical response or sensitivity to small signals.

The latter trade-off is clearest in the context of non-specific activation of an output in a synthetic circuit. For example, in a transcriptional synthetic circuit an output protein may be produced at low levels even in the absence of an input signal. A common strategy for dealing with such background levels of activation is to place a strong degradation tag on the protein that increases its degradation rate. This ensures that in the absence of an activating signal, proteins are quickly degraded. However, increasing the degradation rate clearly comes at a steep energetic cost as more proteins have to be produced to reach the same steady-state. At the same time, the gene circuit loses sensitivity to small input signals

due to their fast degradation.

One of the most interesting recent examples of how energy consumption can be used to increase specificity is the recent work of retroactivity [98–100]. The central problem addressed in these papers is the observation that biochemical signal transduction circuits often have their dynamical behavior altered upon coupling to external outputs due to sequestration of proteins, a property dubbed "retroactivity". Such coupling is particularly undesired when there are a number of downstream outputs. These works demonstrate, both theoretically and experimentally, that it is possible to introduce insulating elements that reduce the magnitude of this retroactivity and thereby restore the modular dynamical behavior of synthetic circuits. A key property of these insulating elements is that they utilize enzymatic futile cycles and hence actively consume energy. Moreover, a detailed theoretical analysis shows that the effectiveness of an insulating element is directly related to its energy consumption [99].

To demonstrate these concepts, we will consider the simple example of a protein Z that is produced at a time-dependent rate $k(t)$ and is degraded at a rate δ (see Figure 1.1). In addition, Z regulates a family of promoters, with concentration $p_{\rm tot}$, by binding/unbinding to the promoter to form a complex C at rates $k_{\rm on/off}$. The kinetics of this simple network is described by the set of ordinary differential equations

$$\begin{aligned} \frac{dZ}{dt} &= k(t) - \delta Z - \tau^{-1}[k_{on} Z(p_{\rm tot} - C) + k_{\rm off} C], \\ \frac{dC}{dt} &= \tau^{-1}[k_{\rm on}\tau Z(p_{\rm tot} - C) + k_{\rm off}\tau C], \end{aligned} \tag{1.1}$$

where we have introduced an overall dimensionless timescale τ for the binding/unbinding dynamics. Notice that if $\tau^{-1} \gg 1$, then the timescale separation between the Z and C dynamics means that the Z dynamics are well approximated by setting $\frac{dC}{dt} = 0$ so that

$$\frac{dZ}{dt} \approx k(t) - \delta Z. \tag{1.2}$$

Thus, when Z is coupled to a system with extremely fast dynamics, the retroactivity term, $\tau^{-1}[k_{on} Z(p_{\rm tot} - C) + k_{\rm off} C]$, is negligible.

This basic observation motivates the idea behind kinetic insulators. Instead of coupling Z directly to the complex C, one couples Z to C indirectly through an intermediary insulating element with fast kinetics. Similar analysis of this more complex network shows that this dramatically decreases the amount of retroactivity. In practice, the insulating element is a phosphorylation/dephosphorylation cycle with fast kinetics (see Figure 1.1). The faster the intermediary kinetics, and hence the more energy consumed by the futile cycle, the better the quasi-static approximation and the more effective the insulator (see [99, 100] for details).

1.3.3 Reducing Variability

Biochemical circuits can also consume energy to reduce variability and increase reproducibility. One of the best studied examples of this is the incredibly reproducible response of mammalian rod cells in response to light stimulation (see [101] and references therein). This reproducibility of the rod cell response is especially surprising given that the response originates from the activation of a single rhodopsin molecule. A simple biophysically plausible model for an active rhodopsin is that its lifetime is exponentially distributed (i.e. the deactivation of rhodopsin is a Poisson process). In this case, the trial-to-trial variability, measured by the squared coefficient of variation, $CV^2 = \sigma^2/\mu^2$, would be equal to 1, Surprisingly, the actual variability is much smaller than this naive expectation.

Experiments indicate that discrepancy is at least partially explained by the fact that the shut-off of active rhodopsin molecules proceeds through a multi-step cascade (i.e the active rhodopsin molecule starts in state 1, then transitions to state 2, etc. until it reaches state L). If each of these steps were identical and independent, then from the central limit theorem the coefficient of variation of the L step cascade would be L times smaller than that of a single step, i.e. $\sigma^2/\mu^2 = 1/L$.

Notice that in order for such a multi-step cascade to reduce variability it is necessary that each of the transitions between the L states be irreversible. If they were not, then one could not treat the L-steps as independent and the progression of the rhodopsin molecule through

the various states would resemble a random walk, greatly increasing the variability [101]. For this reason, reducing variability necessarily consumes energy. Consistent with this idea is the observation that the variability of rhodopsin seems to depend on the number of phosphorylation sites present on a rhodopsin molecule.

In fact, it is possible to directly compute the coefficient of variation [102, 103] as a function of the ratio of the forward and backward rates at each step, θ. The logarithm of this ratio is simply the free-energy consumed at each step, $\Delta G = \log \theta$. Figure 1.2 shows that the coefficient of variation is a monotonically decreasing function of ΔG and hence the energy consumed by the cascade. Note that this decrease in the variability comes at the expense of a slower dynamic response, since the mean completion time scales linearly in the cascade length.

Recent calculations have applied these ideas to the problem of a non-equilbrium receptor that estimates the concentration of an external ligand [104]. It was shown that by forcing the receptor to cycle through a series of L states, one can increase the signal-to-noise ratio and construct a biochemical network that performs Maximum Likelihood Estimation (MLE) in the limit of large L. Since MLE is the statistically optimal estimator, these works suggest that it should be possible to improve the performance of synthetic biology based biodetectors by actively consuming energy.

Moreover, this trade-off between variability and energy consumption is likely to be quite general. Analytical arguments and numerical evidence suggest there may exist a general thermodynamic uncertainty relation relating the the variance, of certain quantities in biochemical networks and the energy consumption [105]. In particular, achieving an uncertainty, σ^2, in a quantity such as the number of consumed/produced molecules in a genetic circuit or the number of steps in a molecular motor, requires an energetic cost of $2k_BT/\sigma^2$. This suggest that any strategy for reducing noise and variability in synthetic circuits will require these circuits to actively consume energy.

1.3.4 Signal Amplification

Biochemical networks can also consume energy to amplify upstream input signals. Signal amplification is extremely important in many eukaryotic pathways designed to detect small changes in input such as the phototransduction pathway in the retina [106] or the T cell receptor signaling pathway in immunology. In these pathways, a small change in the steady-state number of input messenger molecules, dI, leads to a large change in the steady-state number of output molecules, dO. The ratio of these changes is the number gain, often just called the gain,

$$g_0 = \frac{dO}{dI} \tag{1.3}$$

with $g_0 > 1$ implying the ratio of output to input molecules is necessarily greater than 1.

Before proceeding further, it is worth making the distinction between the number gain, which clearly measures changes in absolute number, with another commonly employed quantity used to describe biochemical pathways called logarithmic sensitivity [106]. The logarithmic sensitivity, $\frac{d\log[O]}{d\log[I]}$, measures the logarithmic change in the concentration of an output signal as a function of the logarithmic change in the input concentration and is a measure of the fractional or relative gain. Though logarithmic sensitivity and gain are often used interchangeably in the systems biology literature, the two measures are very different [106]. To see this, consider a simple signaling element where a ligand, L binds to a protein X and changes its conformation to X^*. The input in this case is L and the output is X^*. To have $g_0 > 1$, a small change in the number of ligands, dL must produce a large change in the number of activated X^*. Notice that by definition, in equilibrium, $\frac{dX^*}{dL} < 1$ since each ligand can bind only one receptor. If instead n ligands bind cooperatively to each X, then one would have $\frac{dX^*}{dL} < 1/n$. Thus, cooperatively in fact reduces the number gain. In contrast, the logarithmic sensitivity increases dramatically, $\frac{d\log[X]}{d\log[L]} = n$. But which gain is related to energy consumption? The energy is related to the number gain, not the logarithmic gain. In physics language, since energy is an extensive quantity, it depends on the actual numbers, not relative gain. In biology language, the fundamental unit of energy

is ATP and each time an energy consuming reaction occurs, ATP must be used. Therefore, ATP consumption directly scales with the number and not relative gain. In conclusion, an important consequence of this is that amplification of input signals necessarily requires a non-equilibrium mechanism that consumes energy.

In biochemical networks, this signal amplification is accomplished through enzymatic cascades, where the input signal couples to an enzyme that can catalytically modify (e.g. phosphorylate) a substrate. Such basic enzymatic "push-pull" amplifiers are the basic building block of many eukaryotic biochemical pathways, and are a canonical example of how energy consumption can be used to amplify input signals (see Figure 1.3). A push-pull amplifier consists of an activating enzyme E_a and a deactivating enzyme E_d that interconvert a substrate between two forms, X and X^*. Importantly, the post-translational modification of X is coupled to a futile cycle such as ATP hydrolysis. The basic equations governing a push-pull amplifier are

$$\frac{dX^*}{dt} = \Gamma_a(E_a)X - \Gamma_d(E_d)X^*, \tag{1.4}$$

where $\Gamma_a(E_a)$ is the rate at which enzyme E_a converts X to X^* and $\Gamma_d(E_d)$ is the rate at which enzyme E_d converts X^* back to X. This rate equation must be supplemented by the conservation equation on the total number of X molecules,

$$X + X^* = X_{\text{tot}}. \tag{1.5}$$

In the linear-response regime where the enzymes work far from saturation, one can approximate the rates in (1.4) as $\Gamma_a(E_a) \approx k_a[E_a]$ and $\Gamma_d(E_d) \approx k_d[E_d]$, with $k_a = k_a^{\text{cat}}/K_a$ and $k_d = k_d^{\text{cat}}/K_d$ the ratios of the catalytic activity, k_cat, to the Michelis-Mentin constant, K_M, for the two enzymes. It is straightforward to show that the steady-state concentration of activated proteins is

$$\bar{X}^* = \frac{X_{\text{tot}} k_a[E_a]}{k_a[E_a] + k_d[E_d]} \tag{1.6}$$

Furthermore, one can define a "response time", τ, for the enzymatic amplifier to be the rate at which a small perturbation from steady-state $\delta X^* = X^* - \bar{X}^*$ decays. This yields

(see [106] for details)

$$\tau = (k_a[E_a] + k_d[E_d])^{-1}. \tag{1.7}$$

As discussed above, a key element of this enzymatic amplifier is that it works out of equilibrium. Each activation/deactivation event where the substrate cycles between the states $X \mapsto X^* \mapsto X$ is coupled to a futile cycle (e.g. ATP hydrolysis) and hence dissipates an energy ΔG_{cycle}. At steady-state, the power consumption of the enzymatic amplifier is

$$P = k_a[E_a]\bar{X}\Delta G_{cycle} = k_d[E_d]\bar{X^*}\Delta G_{cycle}. \tag{1.8}$$

The input of the enzymatic amplifier is the number of activating enzymes E_a and the output of the amplifier is the steady-state number of active substrate molecules X^*. This is natural in many eukaryotic signaling pathways where E_a is often a receptor that becomes enzymatically active upon binding an external ligand. Using (1.8), one can calculate the static gain and find

$$g_0 = (P/[E_a])\tau(\Delta G_{cycle})^{-1}. \tag{1.9}$$

This expression shows that the gain of an enzymatic cascade is directly proportional to the power consumed per enzyme measured in the natural units of power that characterize the amplifier: $\Delta G_{cycle}/\tau$. This is shown in Figure 1.3 where we plot the gain as a function of power consumption for different response times.

Notice that the gain can be increased in two ways, by either increasing the power consumption or increasing the response time. Thus, at a fixed power consumption, increasing gain comes at the cost of a slower response. This is an example of a general engineering principle that is likely to be important for many applications in synthetic biology: the gain-bandwidth tradeoff [106]. In general, a gain in signal comes at the expense of a reduced range of response frequencies (bandwidth). If one assumes that there is a maximum response frequency (ie a minimal time required for a response, a natural assumption in any practical engineering system), the gain-bandwidth tradeoff is equivalent to tradeoff between gain and response time. For this reason, energy consumption is likely to be an important consideration for synthetic circuits such as biosensors that must respond quickly to small

changes in an external input. More generally, the gain-bandwidth tradeoff highlights the general tension between signal amplification, energy consumption, and signaling dynamics.

1.3.5 Erasing Information

Memory is a central component of all computing devices. In a seminal 1961 paper [58], Landauer outlined the fundamental thermodynamic and kinetic constraints that must be satisfied by memory modules in physical systems. Landauer emphasized the physical nature of information and used this to establish a connection between energy dissipation and erasing/resetting memory modules. This was codified in what is now known as Landauers principle: any irreversible computing device must consume energy.

The best understood example of a cellular computation from the perspective of statistical physics is the estimation of a steady-state concentration of chemical ligand in the surrounding environment by a biochemical network. This problem was first considered in the seminal paper [107] by Berg and Purcell who showed that the information a cell learns about its environment is limited by stochastic fluctuations in the occupancy of the receptors that detect the ligand. In particular, they considered the case of a cellular receptor that binds ligands at a concentration-dependent rate and unbinds particles at a fixed rate. They argued that cells could estimate chemical concentrations by calculating the average time a receptor is bound during a measurement time.

In these studies, the biochemical networks downstream of the receptors that perform the desired computations were largely ignored because the authors were interested in calculating fundamental limits on how well cells can estimate external concentrations. However, calculating energetic costs requires an explicit model of the downstream biochemical networks that implement these computations. As Feynman emphasized in his book on computation [108], "Information is physical."

Recently, we considered a simple two-component biochemical network that directly computes the Berg-Purcell estimator [109]. Information about external ligand concentration is stored in the levels of a downstream protein (shown in Figure 1.4). Such two-component

networks are common in bacteria and are often used to sense external signals with receptors phosphorylating a downstream response regulator. Receptors convert a downstream protein from an inactive form to an active form at a state-dependent rate. The proteins are then inactivated at a state-independent rate. Interestingly, one can explicitly map components and functional operations in the network onto traditional computational tasks (see Figure 1.4). Furthermore, it was shown within the context of this network, that computing the Berg-Purcell statistic necessarily required energy consumption. The underlying reason for this is that erasing/resetting memory requires energy (we note that while Landauer emphasized that erasing and not writing requires energy [58], a recent paper argues that writing also requries energy [81]). These results seem to be quite general and similar conclusions have been reached by a variety of authors examining other biochemical networks.

These ideas have important implications for synthetic biology. Much as memory is central to the function of modern computers, biological memory modules are a crucial component of many synthetic gene circuits [110]. Any reusable synthetic circuit must possess a memory module that it can write and erase. Currently, synthetic circuits use two general classes of memory modules: protein-based bistable genetic switches [35] and recombinase-based DNA memory [110]. In both cases, resetting the memory involves consuming energy by expressing and degrading proteins (proteins involved in bistability and recombinases, respectively). Although this energy consumption is fundamental to any reusable memory module, it is desirable to find less energetically costly memories that can still be stable over many generations. As synthetic circuits become increasingly complex, these energetic costs are likely to be ever more important.

1.3.6 Using energy consumption to improve synthetic circuits

Energy consumption is a defining feature of most information processing networks found in living systems. The theoretical work reviewed here provides new insights into biochemical networks. The greatest difference between equilibrium and non-equilibrium systems is that in equilibrium, the energy differences between states fundamentally determines the dynam-

ics of the system, while in a non-equilibrium system the energy differences and dynamics become decoupled. This can be used in a variety of ways in biochemical network and we broadly divided up the useful cases into relatively independent roles: increasing specificity, reducing variability, amplifying signal, and erasing information. We believe that focusing on examples of each role will allow theorists and experimentalists to establish a common language and further both non-equilibrium physics and synthetic biology. One beautiful outcome of the interplay of theory and experiment is the recent work showing that a kinetic insulator that actively consumes energy can restore modularity and eliminate retroactivity in a simple synthetic circuit [100].

We believe that the theoretical results can be summarized into several broad lessons on energy consumption that may prove useful synthetic biology as well as providing theorists with future connections to experiments.

- **Fundamental Trade-Offs.** The ultimate limits of response speed, sensitivity, and energy consumption are in direct competition.

- **Saturation of Trade-Offs.** Current works suggest that saturation effects are ubiquitous [78, 91, 97] in energy consumption of biochemical networks and therefore only a few ATP may be enough [104] to nearly achieve the fundamental limits.

- **Futile Cycles are NOT Futile.** Futile cycles appear to be useless when only considering energy costs, but can provide benefits in terms of the fundamental trade-offs.

- **Reusable Logic Must Consume Energy.** This is just the biological realization of Landauer's principle. Memory is especially important for circuits that function in stochastic environments where it is necessary to time-average over stochastic input signals.

- **Chains are Useful.** While it may seem redundant to have long chains of identical parts, if the chain consumes energy this can improve specificity and reduce variation.

- **Time Reversal Symmetry.** While equilibrium systems respect time reversal symmetry (forward and backwards flows are equivalent), energy consumption and non-equilibrium systems necessarily break this symmetry.

- **Manipulate Time Scales.** Consuming energy can be useful to change the time scale of dynamics, as illustrated by the example of retroactivity and the introduction of energy consuming insulators.

- **Information is Physical.** Theorists should heed Feynman's advice and always attempt to translate theoretical advances into actual physical devices.

We will end by focusing on one specific example that we believe is especially timely for synthetic biology. In naturally occurring biochemical networks, the primary source of energy for biochemical networks are futile cycles associated with post-translational modifications such as phosphorylation and methylation of residues. In contrast, energy dissipation in most synthetic circuits takes the form of the production and degradation of proteins. From the viewpoint of both energy and dynamics, protein degradation is an extremely inefficient solution to the problem. Proteins are metabolically expensive to synthesize, especially when compared to post-translational modifications. This may be one reason that most of the information processing and computation in eukaryotic signaling pathways is done through enzymatic cascades. Designing synthetic circuits that can reap the full benefits of energy consumption requires developing new biological parts built based around post-translational modification. Such a "post-transcriptional" synthetic biology would allow to harness the manifold gains in performance that come from actively consuming energy without the extraordinary metabolic costs associated with protein synthesis. Currently, the power of this approach is limited by the dearth of circuit components that act at the level of post-translational modifications. However, recent advances in this area, suggest that it should be possible to overcome these limitations using recently developed techniques.

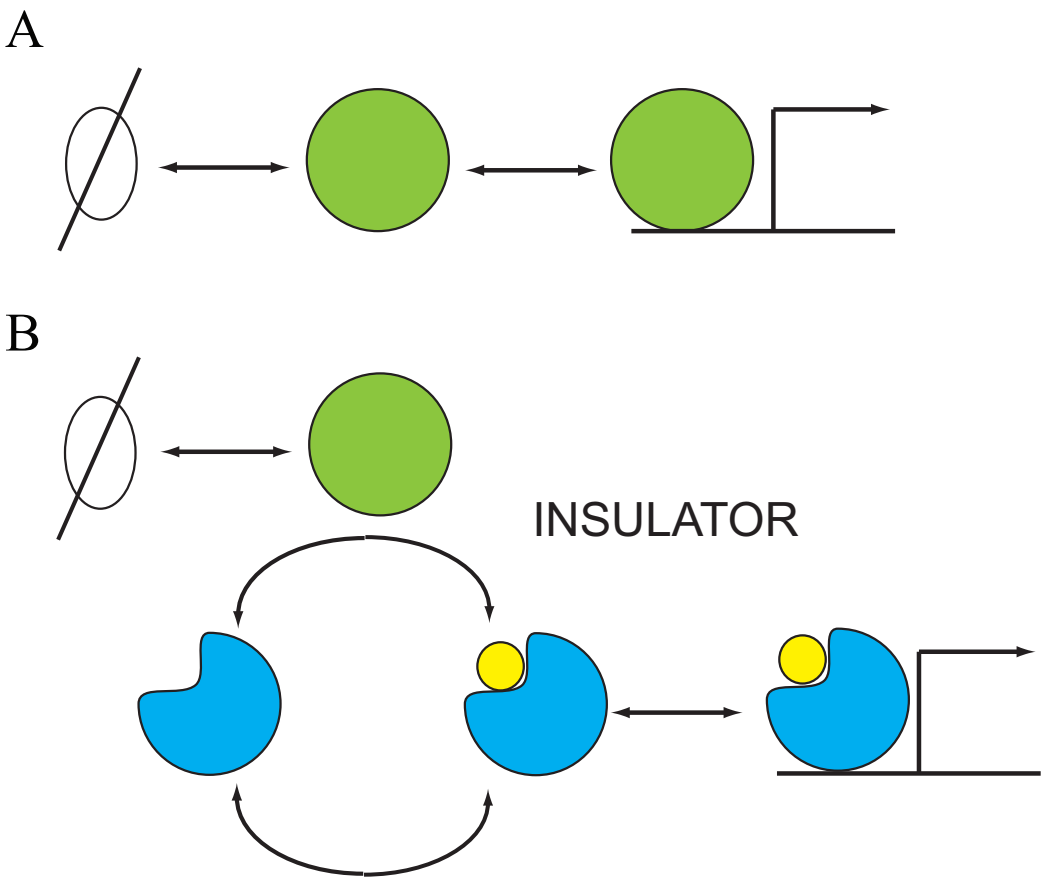

Figure 1.1: **Consuming energy to increase modularity.** (A) A transcription factor regulates downstream promoters. Sequestration of the transcription factor upon binding to promoters can lead to "retroactivity", i.e. a change in the dynamics of the transcription factor levels as a result of coupling to outputs. (B) Coupling the transcription factor through an insulating element consisting of a phosphorylation/dephosphorlyation cycle with fast dynamics reduces the effect of retroactivity.

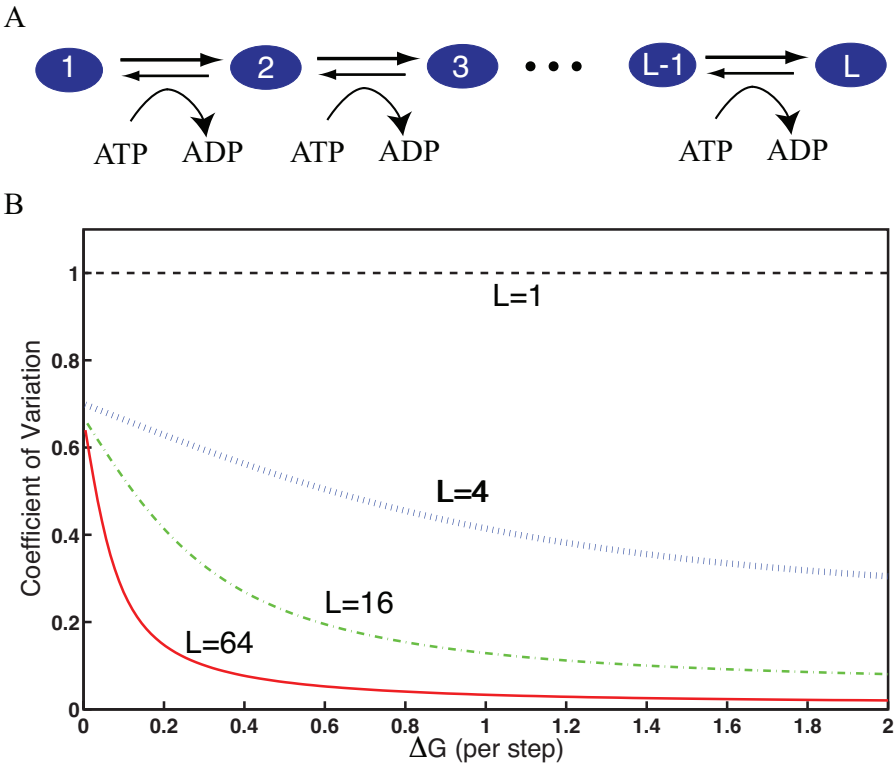

Figure 1.2: **Reducing variability in a multi-step cascade through energy consumption.** (A) A protein (blue ovals) is repeatedly phosphorylated L times. (B) The coefficient of variation, defined as the variance over the mean squared of the time it takes to complete L phosphorylations, as a function of the free-energy consumed during each step in the cascade, ΔG, for $L = 1, 4, 16, 64$.

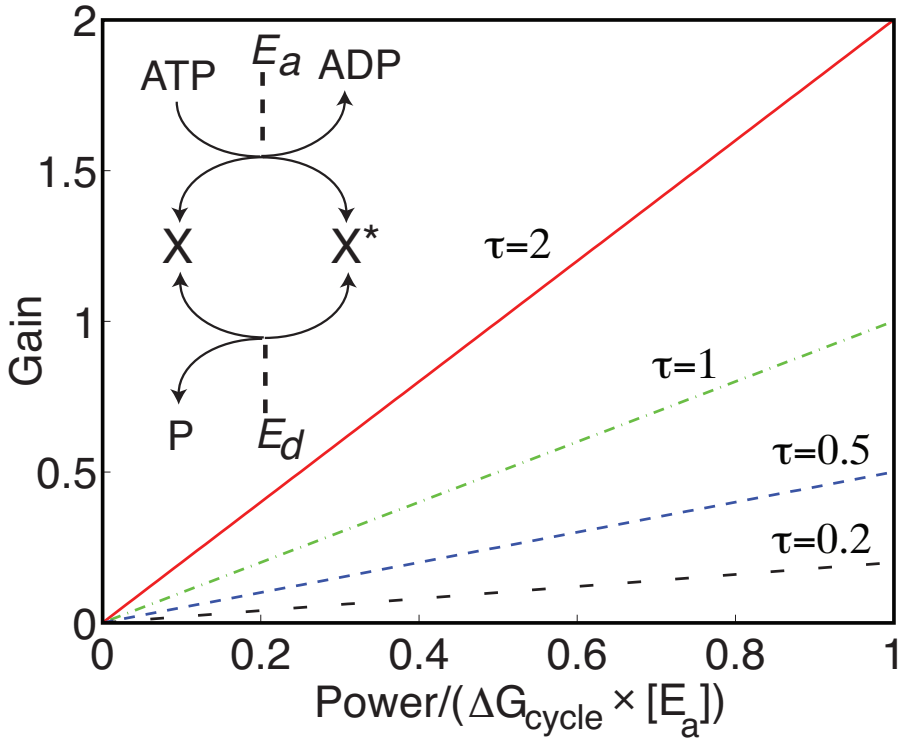

Figure 1.3: **Amplifying signals in a push-pull amplifier by consuming energy.** Schematic illustrates a simple push-pull amplifier where a kinase, E_a, modifies a protein from X to X^* and a phosphatase, E_d, catalyzing the reverse reaction. The plot illustrates that larger gain can be accomplished at the expense of a slower response time τ.

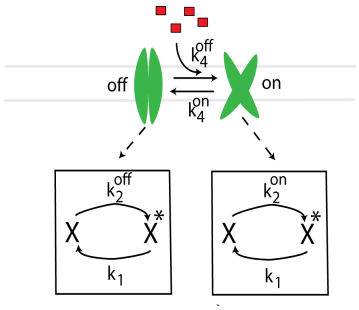

Computational Element	**Biological Realization**	**Biochemical Analogue**
Memory	Number of phosphorylated proteins	X^*
Writing to memory	Phosphorylate protein	k_2^{on}
Erasure of memory	Dephosphorylate protein	k_1
Energy dissipation	Entropy production due to a lack of detailed balance	

Figure 1.4: **A two-component network as a computational module.** (A) Cellular network that calculates the Berg-Purcell statistic for estimating the concentration of an external-ligand. (B) Table summarizing the relationship between the network and standard computational elements and techniques.

2 Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes

The following chapter and the corresponding Appendix and figures have been adapted from: Alex H. Lang, Hu Li, James J. Collins, Pankaj Mehta. Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes. PLoS Comput. Biol. (2014) 10(8): e1003734.

2.1 Introduction

Understanding the molecular basis of cellular identity and differentiation is a major goal of modern biology. This is especially true in light of the work of Takahashi and Yamanaka demonstrating that the overexpression of just four transcription factors (TFs) is sufficient to convert somatic fibroblasts into cells resembling embryonic stem cells (ESCs), dubbed induced pluripotent stem cells (iPSCs) [12]. The idea of using a small set of TFs to reprogram cell fate has proven to be extremely versatile and reprogramming protocols now exist for generating neurons [6], cardiomyocytes [7], liver cells [8, 9], neural progenitor cells (NPC) [10], and thyroid [11] (see reviews [5, 111] for more details). Despite these revolutionary experimental advances, cell fate is still poorly understood mechanistically and theoretically. Recent experiments suggest cell fates can be viewed as high-dimensional attractor states of the gene regulatory networks underlying cellular identity [22]. In particular, cell fates are characterized by a robust gene expression and epigenetic state resulting from the complex interplay of transcriptional regulation, chromatin regulators, non-coding and microRNAs, and signal transduction pathways.

These experiments have renewed interest in the idea of an 'epigenetic landscape' that underlies cellular identity [30–34]. The landscape picture requires several key features to be consistent with experimental observations (see Figure 2.1). All cell fates must be robust attractors, yet allow cells to change fate through rare stochastic transitions [5, 112] as in cellular reprogramming experiments (Figure 2.1A). A common result of reprogramming is

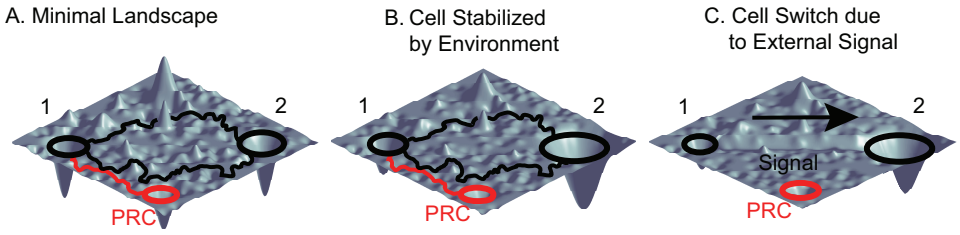

Figure 2.1: **Phenotypic Landscape.** These are illustrative cartoons of the cell fate attractor landscape. (A) The minimal cellular identity landscape. Each cell fate is a basin of attraction (black circles). Reprogramming between different cell fates (1 and 2) can occur probabilistically via different trajectories (black paths). Partially reprogrammed cells (PRC) exist as smaller, spurious, basins of attraction (red circle) that can be experimentally observed by reprogramming experiments (example trajectory in red). (B) Same cellular identity landscape in the presence of a stabilizing environment (ex. favorable culturing medium) for cell fate 2. The environment increases the radius and depth of the cell fate 2 basin of attraction. (C) Landscape in the presence of an external signal that gives rise to differentiation from cell fate 1 to cell fate 2 (ex. growth factors associated with differentiation). Notice the low energy path between the cell fates that drives switching from cell fate 1 to cell fate 2.

not the desired cell fate, but partially reprogrammed cells [28, 29]. These results suggest that the landscape is rugged and may contain additional spurious attractors corresponding to cell fates that do not naturally occur *in vivo*. In addition, environmental and external signals can control cell fates. Some environments stabilize particular cell fates (Figure 2.1B). A dramatic example of this is a protocol for reprogramming to neural progenitor cells (NPCs) that is identical to Yamanaka's protocol for reprogramming to ESC except for the culturing media [113]. Other external signals deterministically switch cell fates, as occurs in normal development (Figure 2.1C) [114]. Together, these imply the landscape is a dynamic entity that depends on environmental signals.

The recent experimental progress has inspired several different theoretical approaches to understand the epigenetic landscape and the underlying gene regulatory networks governing cell fates. One focus has been on explicit construction of landscapes for specific cell fate decisions such as the erythroid vs myeloid choice in hemopoietic development [115], pancreatic cell fates [116], or *C. elegans* vulva development [117]. Other network based approaches use experimental data to constrain the possible networks [118, 119]. A second area of work is based on understanding the underlying gene regulatory network [120, 121]. A recent paper [122] attempts to combine the network and landscape picture by using the network entropy to define a landscape. On a more abstract level, there has been a renewed interest in understanding Waddington's landscape mathematically using ideas from dynamical systems and nonequilibrium statistical mechanics [34, 123]. Most of these models focus on *in vivo* developmental decisions and hence consider the dynamics of a few genes or proteins.

Here, we present a new modeling framework to construct a global (i.e. all cell fates and all TFs) epigenetic landscape that combines techniques from spin glass physics with whole genome expression profiles. We were inspired by the successful application of spin glasses to model neural networks [40–43] and protein folding landscapes [124]. Here, we construct an epigenetic landscape model for cellular identity with 63 stable cell fates and 1337 TFs using cell-fate specific, mouse microarray gene expression data. Each cell fate is a robust

attractor, yet cells can deterministically switch fates in response to external signals. Our model provides a unified framework to discuss differentiation and reprogramming. It also naturally explains the existence of partially reprogrammed cell fates as 'spurious' attractors resulting from the high dimensionality of the landscape. Our model predicts, and we verify, that partially reprogrammed cells are hybrids that co-express TFs of multiple naturally occurring cell fates. Finally, our model reproduces known reprogramming protocols to iPSCs, heart, liver, NPC, and thyroid, and has the potential for designing reprogramming protocols to novel cell fates. Taken together, these results suggest that epigenetic landscapes represent a powerful framework for understanding the molecular circuitry and dynamics that gives rise to cell fate.

The organization of the paper is as follows. First, we explain the motivation for using an attractor neural network to model the epigenetic landscape. Second, we define the state space for the model and the actual biological data used to construct the state space. Third, we give an overview of our landscape model (with details given in Tables A.1 and A.2 and section A.2. Next, we show that our mathematical model captures the essential experimental features of cellular identity. We then show that our model naturally explains the existence of partially reprogrammed cells and makes predictions about their gene expression profiles. We verify this by reanalyzing experimental data. Finally, we show that our model can identify key reprogramming genes in existing reprogramming protocols, suggesting it can be used to identify candidate TF for reprogramming to novel cell fates. We conclude by discussing the implications of our mathematical model for understanding cellular identity and reprogramming.

2.2 Results

2.2.1 Motivation from attractor neural networks

The Takahashi and Yamanaka reprogramming experiments [12] are reminiscent of content-addressable memory and attractor neural networks. First, let us introduce a content-addressable memory with a paraphrasing of the original Hopfield paper. A content-addressable

memory allows one to retrieve a full memory based on sufficient partial information. For example, suppose the complete stored memory is "John J. Hopfield, Neural networks and physical systems with emergent collective computational abilities (1982)." A content-addressable memory is capable of retrieving the full memory based on partial, incomplete input. Therefore, the details "Hopfield," "Neural networks," and "1982" could be enough to recall the full memory.

In the Yamanaka reprogramming protocol, overexpressing only four TFs is enough for a fibroblast to "recall" the global TF expression of an ESC. A content-addressable memory is naturally represented as a basin of attraction in a dynamical system, with partial recall corresponding to entering the basin of attraction and full recall corresponding to reaching the minimum of the basin. Hopfield attractor neural networks [40, 41, 43] are a general method to take an input set of vectors ("memories") and explicitly construct a unique, global, landscape such that each input vector is a global minimum and has a basin of attraction. In what follows, we will exploit the analogy between associative memory in attractor neural networks and cellular reprogramming to explicitly construct the epigenetic landscape underlying cellular identity.

2.2.2 The epigenetic landscape

Our goal is to model the global epigenetic landscape involving all cell fates by using genome wide data. Currently, microarrays are the only technology with genome wide data for a multitude of cell fates (although RNA-seq and other technologies will likely be useful in the future). Specifically, we compiled a dataset of 601 mouse whole genome microarrays (details in section A.1) resulting in the gene expression for $N = 1337$ transcription factors for $p = 63$ cell fates. We restricted our considerations to TFs due to their importance in cellular reprogramming and differentiation. However, our model can be easily generalized to include other important genes. To robustly compare microarrays from multiple platforms, we converted the raw expression data into a rank ordered list. We assumed that gene expression is log-normal distributed (the minimal-assumption model for positive-definite

random numbers such as gene expression) and assigned a z-score to each TF. The final output of this procedure is that it assigns each TF in every cell fate a z-score gene expression.

This continuous gene expression could be used to construct our epigenetic landscapes. However, for mathematical convenience, we discretize the continuous gene expression data into high expression (+1 for z-score >= 0) and low expression (−1 for z-score < 0). See section A.5 for an extended discussion on continuous vs discrete TF expression in attractor neural networks.

This discretization process is biologically plausible. Cellular identity and differentiation are largely controlled by epigenetics, especially histone modifications (HMs) [127] (Figure 2.2A). Epigenetics primarily controls the accessibility of DNA and depending on the HM, the DNA can be stabilized in an open or closed configuration. Using global HM data [125, 126] and comparing it to microarray data, we created a conditional probability distribution of having a HM given a TF expression level (Figure 2.2B). We find that between a z-score of −0.5 to 0.5 there is a sharp threshold which distinguishes genes with the activating modification of histone 3 tri-methylation at lysine 4 (K4) from genes with the inactivating modification of histone 3 tri-methylation at lysine 27 (K27) and poised/bivalent genes (both K4 and K27). This provides a potential biological justification to our discretization. In summary, we take the continuous gene expression and binarize (Figure 2.2C). These binary (i.e. on/off) TF data are the only biological input into our model.

In order to precisely describe the landscape results, we need to define the correct way to measure distances. One possible measure is the overlap (aka dot product or magnetization), defined for cell fate μ as:

$$m^{\mu} = \frac{1}{N} \sum_{i=1}^{N} \xi_i^{\mu} S_i \tag{2.1}$$

where S_i is an arbitrary expression state and ξ_i^{μ} is the gene expression in the natural cell fate μ. The overlap between cell fate μ and state S_i for exactly correlated, anti-correlated, or uncorrelated states is 1, −1, or 0 respectively.

Cell fates from similar lineages (ex. blood) often have similar gene expression patterns. For example, B cells and T cells have a 77% overlap in their gene expression profiles. Such

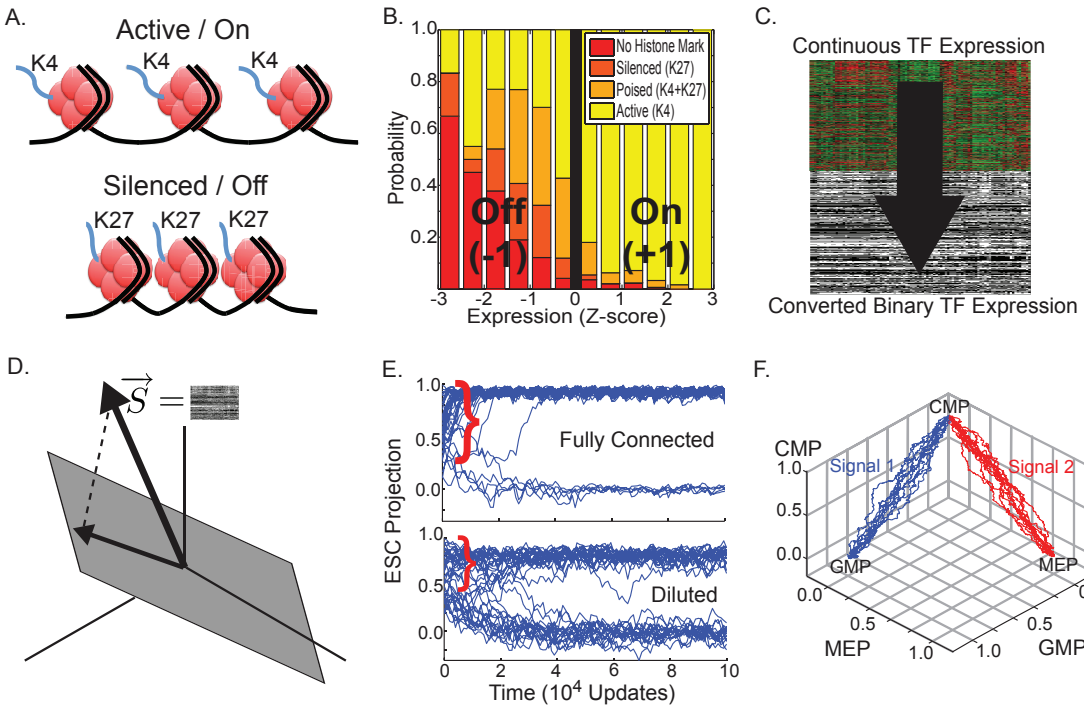

Figure 2.2: **Overview of model.** (A) Histone K4 is associated with active genes, while histone K27 is associated with repressed genes. (B) Conditional probability distribution of HM given TF expression levels derived by comparing microarray data with HM data from [125, 126]. (C) For mathematical convenience, we take the continuous TF expression levels and convert it to binary states. (D) An arbitrary state is represented by a vector $\vec{S}$ of ± 1, with each dimension in the vector space representing the state of a TF. The natural cell fates form a subspace (gray plane). The landscape model is based on the orthogonal projection of the TF state onto this subspace. (E) The dynamics of the landscape model for different initial conditions for a fully connected interaction matrix J_{ij} and a diluted (non-equilibrium) interaction matrix. Plot shows the projection of S on ESC as function of time. Notice the large basins of attraction (red bracket). Parameters used were $\beta = 2.2$ and burst errors of 2% every 5000 spin updates. (F) Simulations showing how a CMP can differentiate into either GMP or MEP in response to two distinct external signals. All trajectories used $\beta = 2.2$. For signal 1, we set $G^{GMP,CMP} = 0.5$ and all other $G^{\mu\nu} = 0$. For signal 2, we set $G^{MEP,CMP} = 0.5$ and all other $G^{\mu\nu} = 0$.

large correlations between cell fates makes the overlap, $\mathbf{m}$, a poor distance measure. In order to measure distances between highly correlated vectors, it is helpful to define the "projection" a^μ of a gene expression state S_i on a cell fate μ by

$$a^\mu = \sum_{\nu=1}^{p} (A^{-1})^{\mu\nu} m^\nu \tag{2.2}$$

where A^{-1} is the inverse correlation matrix and m^ν is the overlap on cell fate ν and is given by

$$A^{\mu\nu} = \frac{1}{N} \sum_{i=1}^{N} \xi_i^\mu \xi_i^\nu \tag{2.3}$$

The projection a^μ measures the orthogonal projection of a state S_i onto the subspace spanned by naturally occurring cell fates, ξ (see Figure 2.2D and section A.5), and a perfect projection onto state μ is given by $a^\mu = 1$. In contrast with the overlap, B cells have zero projection on T cells, and vice versa.

Our landscape assigns an "energy" to every global expression state. We emphasize that this energy does not correspond to physical energy consumption of ATP; instead it is an abstract energy that corresponds to stability and developmental potential of cell fates. The complete landscape H can be thought of as arising from four terms with a simple interpretation (see Figure 2.1):

$$H \quad = \quad H_{basin} + H_{bias} + H_{culture} + H_{switch} \tag{2.4}$$

The first term, H_{basin}, ensures that observed cell fates are valleys in our landscape (Figure 2.1A). The second term, H_{bias}, describes biasing of specific TFs by experimentalists (not shown in Figure 2.1). The third term, $H_{culture}$, increases the radius and depth of cell fates that are favored by the environment or culturing conditions (Figure 2.1B). Finally, in the presence of an external signal that gives rise to differentiation (ex. growth factors associated with differentiation), the fourth term, H_{switch}, opens a low energy path between the initial and final cell fates (Figure 2.1C). We give a complete mathematical description of the model in the section A.2 and a summary in Tables and .

Table 2.1: **Mathematical model of cell identity landscape. Part 1.**

Landscape Term: Index Notation	Landscape Term: Matrix Notation (dim.)	Biological Interpretation
$H = H_{basin}+H_{bias}+H_{culture}+H_{switch}$		Total landscape.
$H_{basin} = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j\neq i}^{N} S_i J_{ij} S_j$	$H_{basin} = -\frac{1}{2}\mathbf{S}^\mathsf{T} J \mathbf{S}$	Produces cell basins of attraction.
$H_{bias} = -\sum_{i=1}^{N} B_i S_i$	$H_{bias} = -\mathbf{B}^\mathsf{T}\mathbf{S}$	External control of individual genes, i.e. inducible expression.
$H_{culture} = -N\sum_{\mu=1}^{p} b^\mu a^\mu$	$H_{culture} = -N\mathbf{b}^\mathsf{T}\mathbf{a}$	External control of specific cell basins, i.e. culturing conditions.
$H_{switch} = -\frac{N}{2}\sum_{\mu=1}^{p}\sum_{\nu=1}^{p} m^\mu G^{\mu\nu} a^\nu$	$H_{switch} = -\frac{N}{2}\mathbf{m}^\mathsf{T} G \mathbf{a}$	Cell switching by signals, i.e. *in vivo* development.
N		Number of TFs, labeled by i, j. In this paper $N = 1337$.
p		Number of cell fates, labeled by μ, ν. In this paper $p = 63$.
S_i	$\mathbf{S}$ $(N$ x $1)$	State (± 1) of i^{th} TF.
ξ_i^μ	ξ $(p$ x $N)$	State (± 1) of i^{th} TF in cell fate μ.

This table provides a summary of the landscape model and the biological interpretation of each term. The first column is written in index notation, while the second column is the same term in matrix notation with the dimension of the term given in parenthesis. If no dimension is listed, the term is a single number.

2.2.3 Cell fates are dynamic attractors that are responsive to signals

We performed self-consistency checks for our model using two *in silico* experiments (see details in Materials and Methods: Simulations). To verify that naturally occurring cell

Table 2.2: **Mathematical model of cell identity landscape. Part 2.**

Landscape Term: Index Notation	Landscape Term: Matrix Notation (dim.)	Biological Interpretation
$A^{\mu\nu} = \frac{1}{N}\sum_{i=1}^{N} \xi_i^\mu \xi_i^\nu$	$A = \frac{1}{N}\xi\xi^\mathsf{T}$ $(p \text{ x } p)$	Correlation between cell fate μ and ν.
$J_{ij} = \frac{1}{N}\sum_{\mu=1}^{p}\sum_{\nu=1}^{p} \xi_i^\mu (A^{-1})^{\mu\nu} \xi_j^\nu$	$J = \frac{1}{N}\xi^\mathsf{T} A^{-1}\xi$ $(N \text{ x } N)$	Interaction strength between i and j.
B_i	$\mathbf{B}$ $(N \text{ x } 1)$	External control of i^{th} TF.
b^μ	$\mathbf{b}$ $(p \text{ x } 1)$	External control of μ^{th} cell fate.
$m^\mu = \frac{1}{N}\sum_{i=1}^{N} \xi_i^\mu S_i$	$\mathbf{m} = \frac{1}{N}\xi\mathbf{S}$ $(p \text{ x } 1)$	Overlap of $\mathbf{S}$ on cell fate μ.
$a^\mu = \sum_{\nu=1}^{p}(A^{-1})^{\mu\nu} m^\nu = \sum_{i=1}^{N} \eta_i^\mu S_i$	$\mathbf{a} = A^{-1}\mathbf{m} = \eta\mathbf{S}$ $(p \text{ x } 1)$	Projection of $\mathbf{S}$ on cell fate μ.
$\eta_i^\mu = \frac{1}{N}\sum_{\nu=1}^{p}(A^{-1})^{\mu\nu}\xi_i^\nu$	$\eta = \frac{1}{N}A^{-1}\xi$ $(p \text{ x } N)$	Predictivity of i^{th} TF in cell fate μ.
$G^{\mu\nu}$	G $(p \text{ x } p)$	Signal dependent coupling that drives cell fate ν to cell fate μ

This table provides a summary of the landscape model and the biological interpretation of each term. The first column is written in index notation, while the second column is the same term in matrix notation with the dimension of the term given in parenthesis. If no dimension is listed, the term is a single number.

fates are dynamic attractors, we randomly perturbed the gene expression profile of cells from the ESC state and then tracked the gene expression over time. Real biology has many potential sources of noise, and the asynchronous dynamics introduced above will likely underestimate the noise. To show that our model is still robust to other large sources of noise, in our simulations we also add in periodic bursts of noise by flipping a fixed percentage

of TF states (2%) to mimic the observation that cellular divisions produce HM errors [128]. Figure 2.2E shows the projection of the TF state on the ESC state as a function of time. For a large number of starting conditions, after an initial transient, the system relaxes back to the ESC state (red bracket), explicitly demonstrating the existence of a large basin of attraction [22]. This is true even when we break detailed balance by making the interaction matrix asymmetric by randomly deleting 20% of interactions (Figure 2.2E Diluted).

Our model can also deterministically switch between cell fates in response to differentiation signals. For example, the common myeloid progenitor (CMP) is a blood cell fate that *in vivo* can differentiate into either granulo-monocytic progenitors (GMP) or megakaryocyte-erythroid progenitors (MEP). In Figure 2.2F, we show *in silico* validation where we start the system in the CMP state and show the trajectories after applying either the GMP (signal 1, blue) or MEP (signal 2, red) differentiation signal, resulting in branching to two distinct cell fates.

2.2.4 Partially reprogrammed cells as "spurious" attractors

When performing a reprogramming experiment, besides the initial cell fate and the end goal cell fate, experimentalists often produce "novel cell fates", dubbed partially reprogrammed cells [28, 29]. These partially reprogrammed cells have the characteristics of a stable cell fate (i.e. they can be passaged indefinitely in culture), but may express a mix of key markers for multiple cell fates and have a global gene expression that does not match any *in vivo* cell fate [29].

While the existence of partially reprogrammed cells was surprising to experimentalists, they have a natural interpretation in our model. One of the most generic properties of all attractor neural networks is that in addition to the desired attractors, ξ_i^μ, the non-linearity of the dynamical process and topology of high-dimensional (in our case $N = 1337$) vector spaces induces additional attractors, which are termed spurious attractors [43]. In our model, since the natural cell fates are the input vectors, these spurious attractors can be interpreted as potential cell fates that do not occur *in vivo*. These spurious attractors are

predicted to be low-dimensional combinations, or hybrids (see section A.4 and section A.5 for details) that should also be stable attractors but with smaller basins of attraction.

A priori, there are several valid hypotheses for the relationship between partially reprogrammed cells and natural cell fates. In the original experiments [28, 29], it was expected that partially reprogrammed cells should be a hybrid of the starting and goal cell fate only (i.e. have a significant projection only on the starting or ending cell fate). Another hypothesis was that in a high-dimensional landscape, randomly chosen vectors should be orthogonal (Figure A.1) (i.e. have a projection of $a \approx 0$ with all cell fates). However, our model predicts that partially reprogrammed cells should be low-dimensional hybrids of existing cell fates, but that they do not necessarily have to be a combination of the starting and goal cell fate. Mathematically, we predict that partially reprogrammed cells should only have a projection $|a| > 0.106$ (2 std above 0, see Figure A.1) for a small number of natural cell fates. Reanalyzing existing genome-wide datasets on partially reprogrammed cells (Table A.3) validates the prediction of our model that partially reprogrammed cells are low-dimensional hybrids of existing cell fates. This qualitative agreement between the predicted spurious attractors and the partially reprogrammed states is independent of details of our landscape function. Importantly, such hybrid states are a generic property of *all* attractor-based landscape models and hence represents an important criteria for judging whether attractor-based models are suitable for describing epigenetic landscapes.

2.2.5 Identifying transcription factors for cellular reprogramming

Our landscape model provides a quantitative method to identify "predictive" TFs for a given cell fate. These predictive TFs can be used as markers of a cell fate and are potential candidates for reprogramming protocols. We expect reprogramming TFs to be a subset of all predictive TFs but not all predictive TFs will lead to successful reprogramming. For example, cell-specific downstream targets of reprogramming TFs are likely to also be highly predictive for a cell type but may not lead to successful reprogramming.

Most reprogramming experiments follow an experimental protocol similar to the one

Table 2.3: **Partially reprogrammed cells as spurious attractors. Part 1.**

Cell line	Start	Goal	Highest projecting states (projection)
1A2 [28]	MEF	ESC	**ESC (0.178)**, MSC (0.158), myoblast (0.142), MEP (0.129), blood vessel (0.113), keratinocyte (0.112), medullary thymic epithelial (-0.111), adipose - brown (-0.117), NK (-0.130), CMP (-0.138)
1B3 [28]	MEF	ESC	**ESC (0.222)**, **MSC (0.161)**, blood vessel (0.139), myoblast (0.138), GMP (0.127), kidney (0.111), MEP (0.107), cornea (0.107), NK (-0.129)
BIV1+ [29]	B Cell	ESC	**myoblast (0.181)**, **prostate (0.164)**, MSC (0.154), MEP (0.138), keratinocyte (0.136), cornea (0.125), ESC (0.111), intestine - Paneth cell (-0.111), CMP (-0.122)

Partially reprogrammed cell lines (first column) and their significant projections (2 std above noise or $|a| > 0.106$) onto "natural" cell fates based on microarray data. Bold indicates 3 std above noise or $|a| > 0.159$. Abbreviations: CLP, Common Lymphoid Progenitor; CMP, Common Myeloid Progenitor; EpiSC, epiblast stem cell; ESC, embryonic stem cell; GMP, Granulocyte-Monocyte Progenitor; MEF, mouse embryonic fibroblast; MEP, Megakaryocyte-Erythroid Progenitor; MSC, Mesenchymal stem cells; NK, Natural Killer cells; NSC, neural stem cells.

outlined by Takahashi and Yamanaka in their seminal paper [5, 12]. Initially the starting cells (usually mouse embryonic fibroblasts, MEFs) are infected with viruses containing all the TFs of interest. The original Yamanaka experiment over-expressed 24 TFs [12], while more recent experiments usually start with about 10 TFs [6–10]. Several days after infection, the cells are switched to culturing conditions that support the desired final cell fate. If an experiment is successful, cells resembling the desired cell fate will appear after a few weeks. This original list is then pruned to identify a "minimal" (essential) set of TFs that still

Table 2.4: **Partially reprogrammed cells as spurious attractors. Part 2.**

Cell line	Start	Goal	Highest projecting states (projection)
BIV1- [29]	B Cell	ESC	**ESC (0.382)**, **EpiSC (0.184)**, **MEP (0.160)**, myoblast (0.145), NSC (-0.108), T Cell (-0.115), skeletal muscle (-0.117), CMP (-0.154)
MCV6 [29]	MEF	ESC	MEP (0.155), myoblast (0.150), ESC (0.149), keratinocyte (0.145), CLP (0.107), GMP (0.107), cornea (0.107), CMP (-0.130)
MCV8 [29]	MEF	ESC	**ESC (0.203)**, **MEP (0.191)**, **myoblast (0.160)**, cornea (0.119), prostate (0.113), skeletal muscle (-0.141), CMP (-0.142)

Partially reprogrammed cell lines (first column) and their significant projections (2 std above noise or $|a| > 0.106$) onto "natural" cell fates based on microarray data. Bold indicates 3 std above noise or $|a| > 0.159$. Abbreviations: CLP, Common Lymphoid Progenitor; CMP, Common Myeloid Progenitor; EpiSC, epiblast stem cell; ESC, embryonic stem cell; GMP, Granulocyte-Monocyte Progenitor; MEF, mouse embryonic fibroblast; MEP, Megakaryocyte-Erythroid Progenitor; MSC, Mesenchymal stem cells; NK, Natural Killer cells; NSC, neural stem cells.

allows for successful reprogramming. In many cases, the viruses are excised [129] to confirm that the the reprogramming does not depend on viral expression. Furthermore, recent experiments indicate that the same TFs can be used to reprogram to a desired cell fate from multiple initial cell fates [112]. These experiments suggest that reprogramming TFs should be based on final, not initial, cell fate.

Intuitively, reprogramming candidates should be both highly expressed and highly "predictive" of the desired cell fate. The TF z-score naturally defines high and low TF expression levels. Within our landscape, the "predictivity" η_i^μ of the i^{th} TF for a given cell fate μ, is measured by its contribution to the potential energy of that cell fate, and is mathematically

defined as:

$$\eta_i^\mu = \sum_{\nu=1}^{p} \left(A^{-1}\right)^{\mu\nu} \xi_i^\nu \tag{2.5}$$

where A^{-1} is the cell fate correlation matrix and ξ_i^ν is the expression of TF i in cell fate ν. We note that the projection and predictivity are directly related as can be seen by

$$a^\mu = \sum_{i=1}^{N} \eta_i^\mu S_i \tag{2.6}$$

where η_i^μ is the predictivity of TF i in cell fate μ and S_i is an arbitrary gene expression state.

For a desired target cell fate, TFs that are high (low) in both predictivity and expression in that cell fate are candidates for over expression (knock out) in reprogramming (see Figure 2.3A). For a simple, single measure of reprogramming efficacy of a TF, the predictivity and expression can be multiplied together to give a "reprogramming score", where the top (bottom) rank order TFs are the best candidates for over expression (knock out). Figure 2.3 shows the expression and predictivity for TFs in a variety of cell fates. In Figure 2.3B, we have explicitly labeled the TFs used in the original Yamanaka protocol for reprogramming to ESC. Consistent with our model, these TFs are both predictive and highly expressed. Figure 2.3C shows TFs that have been successfully used in any reprogramming protocol to ESCs [5] as well as the pluripotency genes (involved in maintaining stem cell fate) *Zfp42* (*Rex1*) [130] and *Nr0b1* (*Dax1*) [131]. Once again these genes are highly predictive for ESCs. As a further check on the biological validity of our predictions, we analyzed the GO Annotation of our top 50 candidates for ESC reprogramming (Tables and). Within these top TFs, 12 have successfully been used in reprogramming, 7 are known pluripotency TFs, 16 are involving in cell differentiation, while 15 have no known function and are intriguing reprogramming candidates. Taken together this suggest that we are capturing the essential biology despite minimal biological data for input.

While ESC have been studied in the most detail, recent experiments have reprogrammed (aka direct conversion) to other cell fates such as cardiomyocytes [7] (Figure 2.3D), liver [8, 9] (Figure 2.3E), and thyroid [11] (Figure 2.3F). Once again we have explicitly labeled the TFs

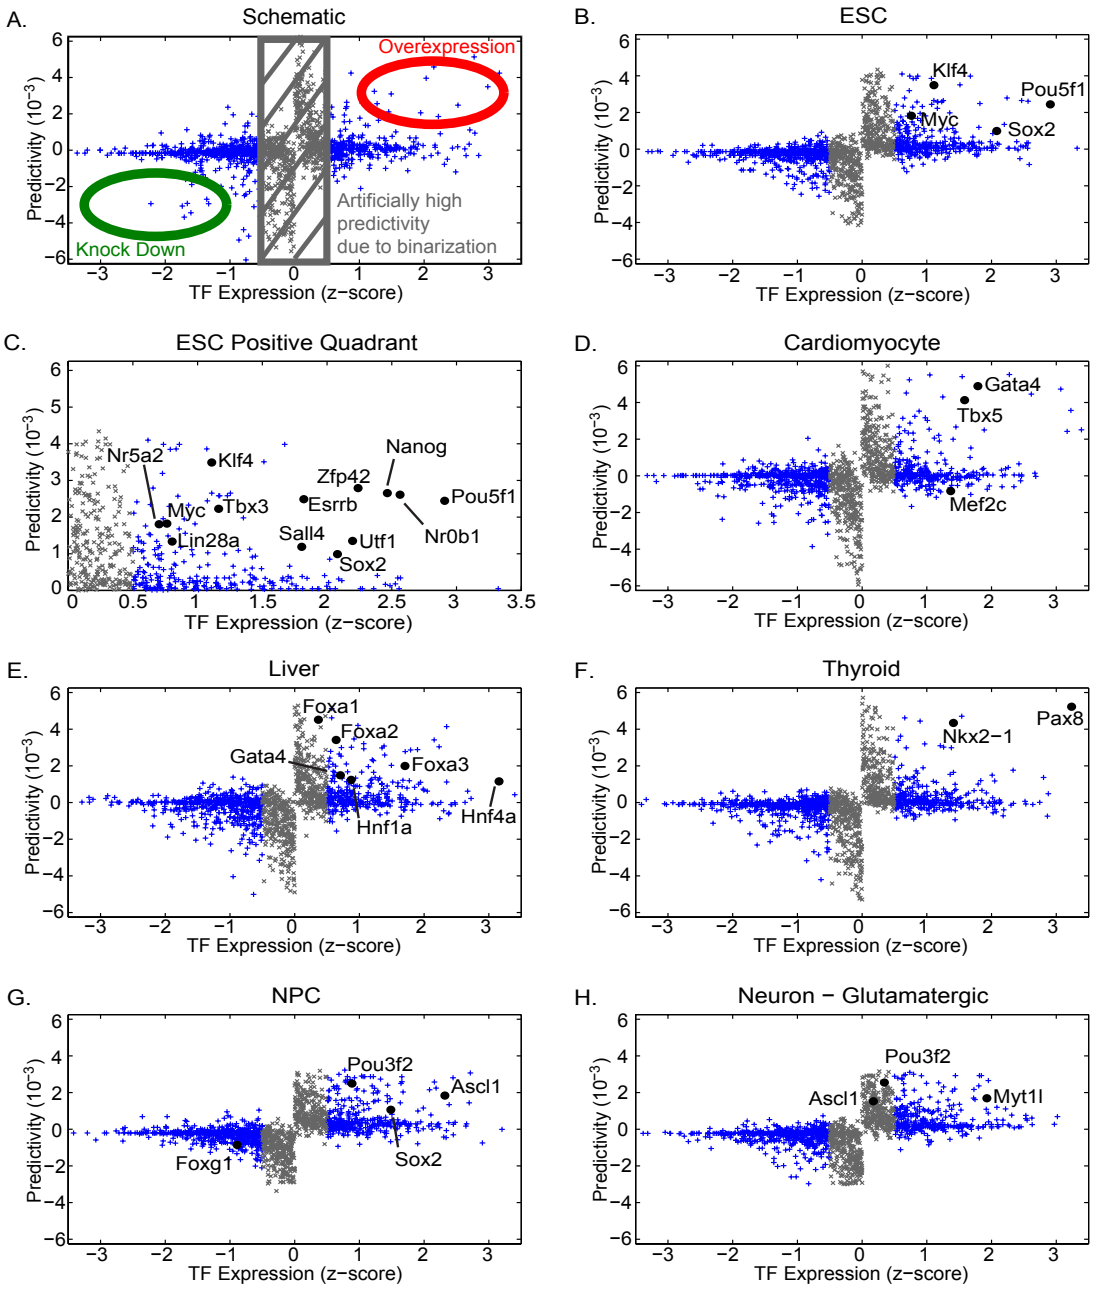

Figure 2.3: **Identifying reprogramming candidates.** See section 2.4 for extended figure details.

that have been successfully used for direct conversion. Notice that all of these TFs (except *Mef2c*) are highly predictive and highly expressed. Note that *p19Arf* [9] used in the direct conversion to liver was not differentially expressed in our microarrays and therefore was not included in our model.

We also examined TFs used in direct conversion to neural lineages. As discussed in [6], these TFs were chosen because they were known to be important in either neurons or neural progenitor cells (NPC). Figure 2.3F and 2.3G show the expression and predictivity of TFs for neural progenitor cells (NPC) [10] (Figure 2.3G), and neurons [6] respectively. Induced NPC were made using a four TF cocktail consisting of *Pou3f2* (*Brn2*), *Sox2*, and *Foxg1* [10]. Our analysis shows that the first two of these TFs are predictive for NPCs while *Foxg1* is predictive for neural stem cells (NSC) (see Figure A.3). Induced neurons (iN) can be made using the TFs *Myt1l*, *Pou3f2*, and *Ascl1* [6]. Consistent with their experimental design, we find that *Myt1l* is highly predictive for mature neurons, while the remaining TFs (*Pou3f2*, *Ascl1*) are predictive for NPCs.

While it is not possible to perform statistical tests to test our examples due to the scarcity of reprogramming protocols, we performed a simple numerical exercise to gauge the predictive power of our model. The four Yamanaka factors are all in the top 50 when ranked by their reprogramming score for ESCs (where the reprogramming score of a TF is defined as the product of the expression and predictivity scores of a TF). We randomly permuted TF labels and asked how often all four Yamanaka factors remained in the top 50. For a million independent permutations, this occurred only once, confirming that our model is capturing many essential aspects of cellular reprogramming.

2.3 Discussion

A common biological metaphor used to describe development and cellular reprogramming is a rugged "epigenetic landscape" which emerges from a complex gene regulatory network, with cell fates corresponding to attracting valleys in the landscape. Despite decades of biological innovation, the large number of genes and their complex interactions has pre-

vented the quantitative modeling of a global epigenetic landscape. To meet this challenge, we have developed a new quantitative framework of cellular identity to directly model the global, high-dimensional epigenetic landscape. Using whole genome expression data, we constructed an epigenetic landscape based on techniques from spin glass physics and neural networks. Our landscape only depends on the experimentally determined gene expression of natural cell fates. Yet, it explains the existence of spurious cell fates (known as partially-reprogrammed cells) and can reproduce known reprogramming protocols to embryonic stem cells, heart, liver, thyroid, neural progenitor cells, and neurons. More importantly, our model can be used to identify candidate transcription factors for reprogramming to novel cell fates.

An interesting question is if spurious attractors are ubiquitous throughout the landscape, why does standard development not produce partially reprogrammed cells? The key is the difference in the dynamics. In cellular reprogramming, the starting cell fate is forced to express a small number of TF and this leads to a stochastic conversion to the desired cell fate (Figure 2.1A). During this stochastic exploration of the landscape, there is only a weak bias towards the final state, so it is easy for the cells to get trapped in a metastable state. However, during standard development, the external signals actively reshape the landscape and open up low energy valleys between cell fates (Figure 2.1C). This strong bias towards the final cell state results in a deterministic switch during which the spurious attractors are only a small road bump on the path to the final cell state. Therefore, it is not a surprise that partially reprogrammed cells are only found during cellular reprogramming and not during standard development.

Epigenetic landscapes can also be used to identify important, or predictive, TFs for cell fates. The predictivity of a TF for a cell fate generalizes the idea of specificity. A TF is specific to a cell fate if it is expressed only on in a small subset of cell fates. In contrast with specificity, predictivity weighs the global correlations amongst cell fates when assessing the importance of a TF for a cell fate. Thus, the predictivity not only picks out important specific TFs, but also TFs that are lineage markers. For example, *Brachyury* (T) [132] is

a general marker of mesodermal lineages. Since it is highly expressed in large a number of cell fates, it is not specific to any given cell fate. However, it is predictive because its expression is a strong indicator that a given cell fate is a mesodermal lineage.

The concept of predictivity also yields new insights into the Yamanaka protocol. When the Yamanaka factors were first published, two of the four TFs, *Pou5f1* (*Oct4*) and *Sox2* were known to be important for ESCs. In contrast, the role of the other two TFs, *Klf4* and *Myc*, was not well understood [133]. It was quickly shown that *Myc* was was not essential to reprogramming (*Oct4*, *Sox2*, and *Klf4* can reprogram alone), but nonetheless enhanced the efficacy of reprogramming [134]. The importance of *Klf4* was surprising given that it is neither highly expressed nor specific for ESC. However, *Klf4* is highly predictive of ESC (Table A.3). For this reason, our model actually explains why *Klf4* is a prime candidate for reprogramming to ESCs.

We make several experimentally verifiable predictions. First, our model predicts the partially reprogrammed cells should be hybrids of existing natural cell fates. As more partially reprogrammed cells are studied, if they are found to either have high projection on only one cell fate ($a^\mu \approx 1$ for one μ) or no projections on any cell fates ($a^\mu \approx 0$ for all μ), this would call into question whether partially reprogrammed cells are truly the spurious attractors of an attractor neural network. Second, our model can be used to identify important, or predictive, TFs for cell fates. TFs with large positive (negative) predictivity should be positive (negative) markers for a cell fate. Additionally, for cellular reprogramming we predict that TFs with large positive (negative) predictivity and expression could be over expressed (knocked out) to reprogram to a desired cell fate. Therefore, our model has several predictions that can be tested against future experimental progress in the field.

Our model has several limitations. First, a generic limitation for any method relying on microarrays to define gene expression is that one cannot distinguish between direct, causal, interactions and indirect, correlative, interactions. Therefore, predictivity can establish the importance of a gene, but further experiments are needed to determine if the predictive gene is the controller of the cell type or just a passive indicator of a cell type. Second, it

fails to accurately capture the dynamics of reprogramming. Simulations of reprogramming with known protocols, such as the Yamanaka protocol, lead to rates of reprogramming that are comparable to the rates from a reprogramming simulation with a randomly selected protocol. This is likely due to the fact that cell fates are extremely stable and hence reprogramming is extremely rare. Third, our model does not directly explain the importance of the non-specific transcription factor *Myc*. Many protocols use *Myc* [5], but it can be replaced (with no deleterious effect) by short hairpin RNAs (shRNAs) [135], or dropped completely from protocols at the expense of speed and less efficient reprogramming [134]. This suggests that *Myc* may have an alternative role and instead of being a biasing field, B_i, it may instead raise the effective noise of the system (i.e. decrease β). Another limitation is that based on the currently available experimental data, our landscape construction cannot definitively be distinguished from alternative constructions. For example, the interaction network could be constructed by such that it does not weigh each cell fate equally (as is currently done). This would have the effect of changing the relative stability of cell fates. Therefore, in the absence of more experimental data, our landscape and a weighted landscape cannot be distinguished.

A popular approach to inferring landscapes from biology data are "Maximum Entropy" models. This method has been used to model firing neurons [136], protein configurations [137, 138], and antibody diversity [139]. The Maximum Entropy approach takes as input large samples of biological data and a set of constraints and outputs a landscape that maximizes the entropy. While Maximum Entropy models can be used to infer landscapes with basins of attraction [140], it can quickly become a computationally challenging problem. Our approach differs from Maximum Entropy models in the following way. Since our goal is to model a landscape with basins of attractions, we make the ansatz that the landscape can be described by a Hopfield neural network. Then we insert real biological data, ξ, to construct the landscape exactly. Our method requires no computational inference of parameters.

There are several natural extensions of the model discussed in this paper. The landscape

could be constructed with additional biological input such as other genes, microRNAs, or histone modification data. This opens up possibilities of improving upon the high reprogramming rates achieved by overexpressing microRNAs [141] or synthetic mRNAs [142]. Another attractive element of the framework presented here is that it allows for a quantitative analysis of whole genome-wide expression states (see Tables 2.3 and 2.4). This is likely to yield a more accurate classification of reprogrammed cells. Finally, directed differentiation protocols [143] attempt to mimic standard development in vitro and have proven to have high efficiency and fidelity. Future work will try to use our landscape to predict the necessary signaling factors for rationally designing more efficient directed differentiation protocols. Overall, epigenetic landscapes provide a unifying framework for cell identity, reprogramming, and directed differentiation, and our results suggest these landscapes can provide crucial insight into the molecular circuitry and dynamics that gives rise to cell fate.

2.4 Extended Figure 2.3 Caption

For a given cell fate, we plot every differentially expressed transcription factor's (TF) predictivity (aka energy projection-contribution, η_i^μ) vs TF expression level (z-score normalized). Unless otherwise stated all existing reprogramming protocols to a given cell fate are labeled. (A) Schematic illustrating predictivity vs expression level plots. The large positive (negative) predictivity and large positive (negative) gene expression TFs are candidates for over expression (knock out) in a reprogramming protocol. The TFs with z-score between -0.5 and 0.5 are highlighted in gray because Figure 2.2B suggests these TFs predictivity may be prone to extra noise induced by the data discretization. (B) Embryonic stem cell, ESC (induced pluripotent stem cells, iPSC). Original Takahashi and Yamanaka factors *Pou5f1* (*Oct 4*), *Sox2*, *Klf4*, and *Myc* [12]. (C) Inset of ESC positive predictivity and gene expression. *Zfp42* (*Rex1*) [130] and *Nr0b1* (*Dax1*) [131] are pluripotency markers that are not necessary to overexpress for reprogramming, while combinations of the remaining labeled TFs have been successfully used in reprogramming protocols [5]. (D) Heart (induced cardiomyocytes, iCM) [7]. (E) Liver (induced hepatocytes, iHep). There are two published protocols. One

protocol used *Hnf4a* plus any of *Foxa1*, *Foxa2*, or *Foxa3* [8] while another used *Gata4*, *Foxa3*, *Hnf1a*, and deletion of *p19Arf* [9]. *p19Arf* was not differentially expressed in our microarrays and is not shown. (F) Thyroid [11]. (G) Neural Progenitor Cells, NPC (induced NPC, iNPC) used *Pou3f2* (*Brn2*), *Sox2*, and *Foxg1* [10]. With our microarrays we find that *Foxg1* is not predictive for NPC but is predictive of neural stem cells (NSC) (see Figure A.3). (H) Neurons (induced neuron, iN) [6]. The reprogramming protocol used a combination of factors that were known to be important to ether mature neurons (*Myt1l*) or NPCs (*Pou3f2*, *Ascl1*). (G) shows that *Pou3f2* and *Ascl1* are predictive of NPCs.

3 Cellular reprogramming dynamics follow a simple one-dimensional reaction coordinate

The following chapter and the corresponding Appendix and figures have been adapted from: Alex H. Lang, Sai Teja Pusuluri, Pankaj Mehta, and Horacio E. Castillo. Cellular reprogramming dynamics follow a simple one-dimensional reaction coordinate. arXiv (2015), 1505.03889.

3.1 Introduction

Biology is in the midst of the revolution spearheaded by the pioneering work of Takahashi and Yamanaka on cellular reprogramming showing that it is possible to reprogram mouse embryonic fibroblasts (MEFs) to cells resembling embryonic stem cells (ESCs), commonly called induced pluripotent stem cells (iPSCs), by manipulating the expression of just four transcription factors (TFs). The idea of manipulating small sets of TFs to alter cell fates has proven extremely versatile and it is now possible to create iPSCs from a variety of cell types [5], as well as perform direct conversions between two differentiated cell types such as MEFs and neurons [6]. Most reprogramming experiments have a similar design [144] (Fig 3.1A). The starting cell type (e.g. MEF) is engineered with a construct containing the desired reprogramming genes. These genes are induced at the start of the experiment. After several days, the cell culturing conditions are switched to a medium favorable to the desired cell type (e.g. stem cell media). At a later time, typically a few weeks, the exogenous genes are turned off. If all goes well, a small percentage ($\approx 0.01 - 1\%$) of cells successfully reprogram to the desired cell type.

Significant progress has been made towards understanding the mechanisms underlying cellular reprogramming [4, 145] (which from now on we will use to include both reprogramming to iPSC as well as direct conversion), yet many questions remain. Cellular reprogramming requires global changes in gene expression involving hundreds of transcription factors and thousands of genes, but how cells dynamically alter their gene expression

profile during reprogramming is still not well understood. Reprogramming rates seem to depend on the exact protocol used and can be changed by several orders of magnitude through careful genetic manipulations [26, 146]. Experiments have also measured whole genome time courses during reprogramming but the high-dimensional nature of the measured trajectories make them difficult to interpret [147]. Other experiments have examined gene-level events during reprogramming. Buganim et al [112] analyzed reprogramming dynamics at the single-cell level and concluded that reprogramming initially is probabilistic but ends with a hierarchichal (i.e. ordered), deterministic stage. In contrast, Polo et al [147] analyzed reprogramming dynamics with both population level and single-cell level measurements and concluded that reprogramming follows an early deterministic phase with many gene changes, followed by an intermediate phase with fewer changes, and ending with a deterministic phase with many gene changes. Recently, Chung et al [148] measured single cell reprogramming dynamics and proposed that the intermediate phase of reprogramming is a loosely ordered probabilistic phase in which the timing between events is probabilistic, but the order of events is relatively deterministic. This highlights the need for a better understanding gene expression dynamics during reprogramming.

Reprogramming involves global changes in gene expression and hence is intrinsically high dimensional. For this reason, it is common to use dimensional reduction techniques such as Principal Component Analysis (PCA) to project the dynamics onto a low-dimensional sub-space. However, dimensional reductions techniques such as PCA have several key limitations. The principal component vectors have no clear biological interpretation, making it difficult to extract biological meaning from the resulting low-dimensional dynamics. PCA also depends on the type and quality of the data included in the dataset, making it cumbersome to compare dynamical data across experiments and systems.

To overcome these challenges, we introduce a new technique for visualizing high-dimensional reprogramming dynamics inspired by "epigenetic landscape" models for cellular identity. In Waddinton's original landscape idea [30], cell types correspond to basins of attraction in an abstract cell identity landscape. This idea has been refined by a variety of researchers,

and has yielded a number of insights into the genetic basis of cellular identity [16, 31–33, 39, 116, 122, 123, 149–154]. Two of us recently proposed a landscape model [155] that takes global gene expression profiles (microarrays or RNA-Seq) and uses techniques inspired by spin physics and the Hopfield model to explicitly construct a cell identity landscape. This model provided a natural explanation for the existence of partially-reprogrammed cell types and can identify TFs that have been used to successfully reprogram to multiple cell types. In this paper, we extend our previous work to analyze reprogramming dynamics. Using a new linear-algebra based analysis method inspired by our landscape model, we show that the experimentally observed gene expression dynamics during reprogramming follow a simple, one-dimensional reaction coordinate. This reaction coordinate emerges naturally in numerical simulations of our landscape model, suggesting that reprogramming can be understood as a "barrier crossing" between landscape minima.

3.2 Results

3.2.1 Mathematical model and Data Analysis Method

Here, we briefly summarize the relevant features of the landscape model (see section B and Lang et al. [155] for details). Cell types are stable basins of attraction (minima of the landscapes) and reprogramming between basins proceeds through stochastic fluctuations resulting from gene expression noise (Fig 3.1B). The landscape is constructed directly from the genome wide expression profiles of natural cell types using a curated dataset of microarrays for $p = 63$ cell types and approximately $N \sim 1400$ TFs (see section B). This data is summarized in a cell type matrix, ξ_i^μ, whose entries contain the expression level of TF i in cell type μ (e.g. MEF, ESC). This construction can easily be extended to include genes beyond TFs.

The global gene expression level of TFs can be summarized using a N-dimensional expression state vector S_i whose entries encode the expression level of TF i with $i = 1 \ldots N$. Expression levels are treated as continuous variables when analyzing experimental data and as binary variables which can be either on or off ($S_i = + \pm 1$) when performing numerical

simulations (see section B). To analyze experimental data, it is useful to define a "distance" measure between an arbitrary expression vector S_i and the expression vector, ξ_i^μ, for cell type μ. One natural distance measure in gene expression space is the overlap or dot product, $m^\mu = 1/N \sum_i S_i \xi_i^\mu$, which measures the correlations between S_i and ξ_i^μ. The overlap between cell type μ and state S_i is 1, -1, or 0 for a perfectly correlated, anti-correlated, or uncorrelated states, respectively. In practice, the dot product is a poor measure of distance because cell types are highly correlated with each other. For example, blood cell types share a common core set of gene expression and thus B cells and T cells have a 87% overlap in their gene expression profiles.

For this reason, it is useful to introduce an alternative measure of distance we call projections, with a^μ denoting the projection of S_i on the expression profile of cell type μ. The projection has a simple geometric interpretation depicted in Figure 3.1C and is calculated by first projecting (ie casting a shadow) of S onto the hyperplane defined by the p cell types in the matrix ξ (represented as the gray plane) and then measuring the distance to the cell type μ within this cellular subspace. The benefit of this construction is that it naturally accounts for the correlations between cell types: the projection of a B cell with itself is one, while a B cell's projection on T cells is zero, and vice versa. This is in stark contrast with correlation based measure of distance in gene expression space.

Projections arise naturally when constructing landscape models for cellular identity. In Lang et. al [155], it was shown that it is possible to define a Lyapunov function (commonly called an energy), H, that characterizes the landscape. In terms of the projections a^μ and overlaps m^μ, the energy or Lyapunov function takes the form (see section B and [155]):

$$\begin{aligned} H &= H_{\text{basin}} + H_{\text{culture}} \\ &= -\frac{N}{2} \sum_{\mu=1}^{p} m^\mu a^\mu - \sum_{\mu=1}^{p} b^\mu a^\mu. \end{aligned} \tag{3.1}$$

We emphasize that this Lyapunov function represents an abstract "cellular identity energy surface" characterizing the stability of cell states and cannot be directly related to metabolism or ATP consumption. In this expression, the first term H_{basin} arises from the

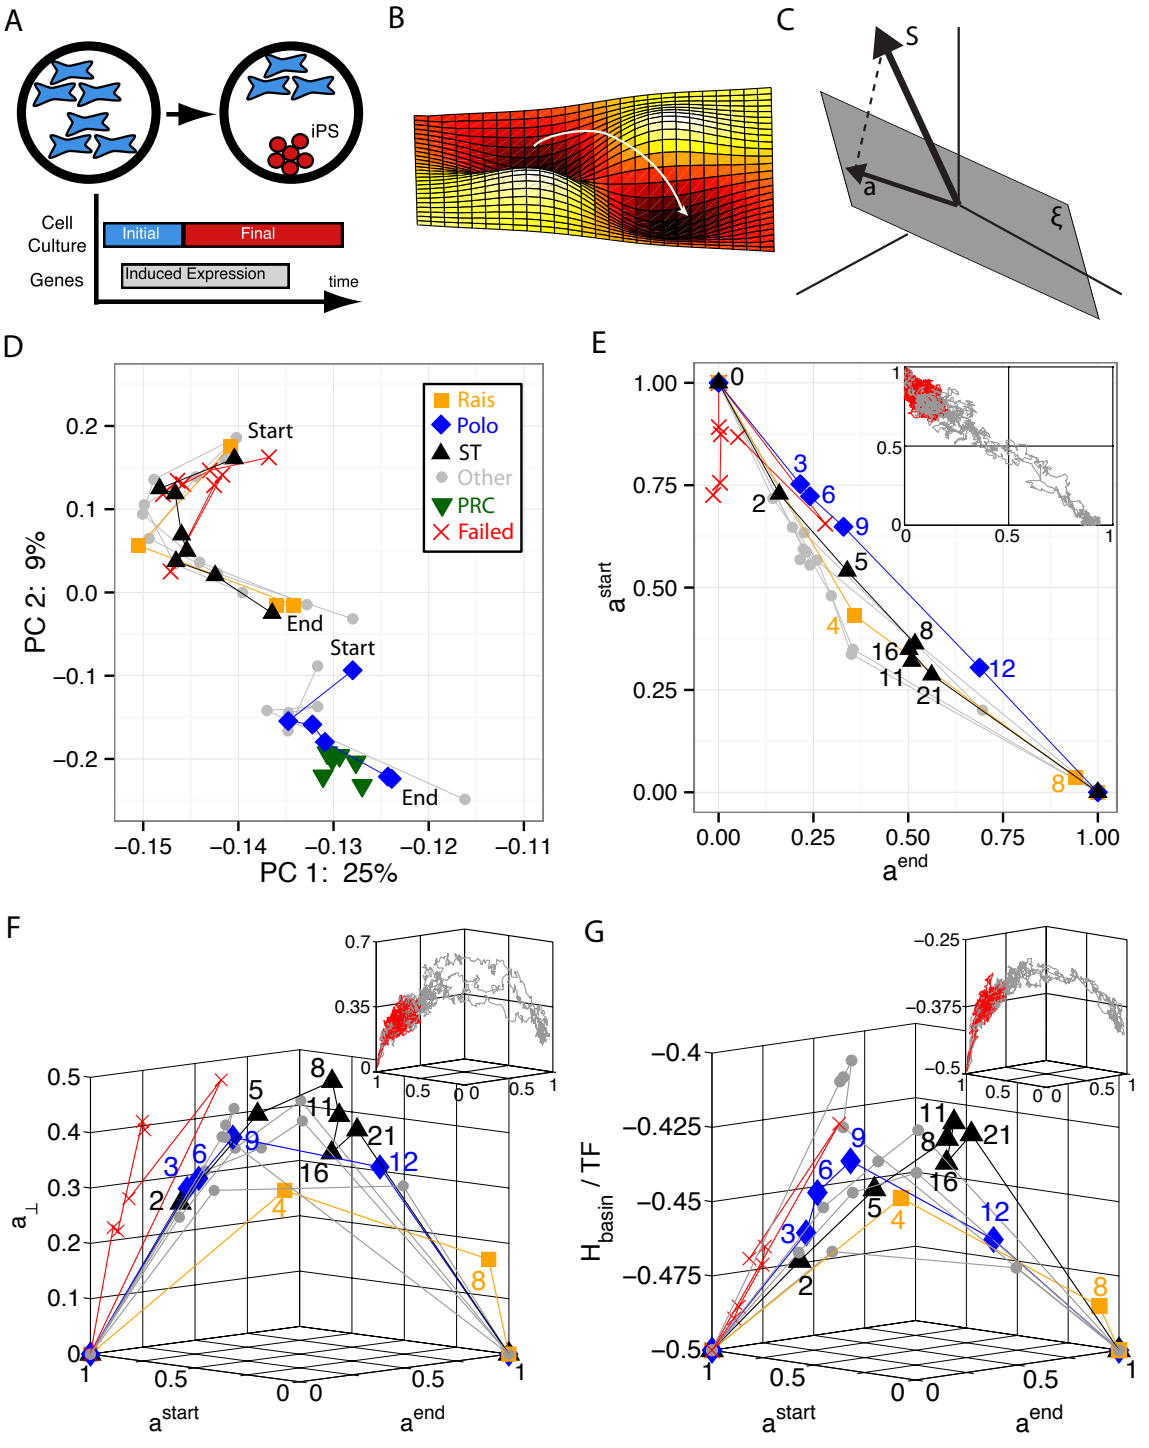

Figure 3.1: **Cellular Reprogramming Reaction Coordinate.** See extended figure caption in section 3.4.

"effective" interaction between genes and ensures that all cell types $\mu = 1 \dots p$ are attractors of the dynamics that have large basins of attraction. This can be seen by noting that in a given cell type (say $\mu = 1$), $S_i = \xi_i^1$ and $a^1 = m^1 = 1$, while the projection on all other cell types is zero, $a^\nu = 0$ ($\nu = 2 \dots p$). Plugging these results into Eq. 3.1 shows that each cell type is a global minimum with energy $H_{min} = -\frac{N}{2}$. The second term H_{culture} represents the stabilizing effect of the culturing conditions on a particular cell type. For example, when cells are grown in MEF culture, then only $b^{MEF} \neq 0$, while in ESC culture, only $b^{ESC} \neq 0$. Finally, to incorporate the fact that some transcription factors are overexpressed in the experiments (see section B) the dynamics of the variables S_i corresponding to over expressed TF are locked in the "on" state.

3.2.2 MEF reprogramming dynamics

We begin by reanalyzing the experimentally available time series data on reprogramming. Fig 3.1D, shows the first two principal components (PC) for 10 different reprogramming trajectories from MEF to iPSC from multiple labs. In the analysis, we have included partially reprogrammed cells (PRC), which are novel cell states only found during unsuccessful reprogramming experiments. The plot shows dynamics projected onto the first two PCs, but in reality this system is high-dimensional and it takes 21 PCs to explain 80% of the variation in the data (see Figure B.3 for details). The PCA plot illustrates several important findings. First, reprogramming trajectories seem to group into two distinct clusters, and within each cluster, the starting points (day 0) and ending points (final iPSC) are near each other. Therefore, even for different experimental protocols, reprogramming seems to follow only a few paths. Second, these paths are distinct from partially reprogrammed cells (PRC). While several reprogramming data points seem to be near PRCs, this is an artifact of keeping only two PCs in our visualization. In fact, the PRCs only have a Spearman correlation of 90% with the closest reprogramming data point and approximately 80% correlation with the two closest trajectories. Third, the final state of failed trajectories (trajectories that did not successfully reprogram to iPSCs) is closer to their starting point rather than to iPSCs,

suggesting that failed trajectories do not leave the basin of attraction of the initial cell type. While PCA allows easy visualization of the data, the Principal Components have no clear biological meaning making it difficult to interpret the lower dimensional PCA dynamics.

In Figures 3.1E-3.1G, we have replotted the same time-series data as in the PCA plots using projections on the starting and ending cell type. As in the PCA plot, the various symbols represent the actual data, while the lines connecting data show the time order of experimental points. In these plots, the starting (ending) states for each trajectory are defined as the initial (final) time point for the corresponding experiment. When calculating projections, the start (end) states replace MEF (ESC) in our cell type matrix ξ. This allows us to plot each experiment against its own start and end points. This additional step is necessary because different experiment define MEFs and iPSs differently.

The result of this analysis is shown in Figure 3.1E. In contrast to the PCA plot which contained two clusters (Figure 3.1D), the reprogramming trajectories in the projected basis all follow a similar path. This suggest cells follow a simple one-dimensional reaction coordinate during reprogramming: a straight line joining the starting cell type with the ending cell type in projection space. This data collapse is more remarkable when considering the extreme heterogeneity in reprogramming rates across the plotted experiments. The Polo et al experiment [147] represents a typical time course with reprogramming taking approximately two weeks, while Rais et al [146] is the fastest trajectory (8 days) and ST (Samavarchi-Tehrani et al) [156] is the slowest trajectory in our dataset (30 days).

In order to better understand how trajectories with such different reprogramming rates can still follow the same coordinate, it is useful to extend the analysis to account for how reprogramming trajectories project on other cell types besides the starting and ending cell types. To do so, we introduce a new quantity,

$$a_{\perp} = \sqrt{\sum_{\substack{1 \leq \nu \leq p \\ \nu \neq (\text{start,end})}} \left(a^{\nu}\right)^2}, \tag{3.2}$$

that measures the magnitude of the projections perpendicular to the plane spanned by the starting and ending cell type. This is shown in Figure 3.1F. Notice that faster trajectories

have a smaller perpendicular projection on the remaining cell types than slower trajectories. Furthermore, the difference in speed between experiments arises largely from the fact that slower trajectories also appear to get stuck at particular points along the reaction coordinate for as many as two weeks.

To compare these experimental trajectories to our mathematical model, it is useful to visualize this data in yet another way. In Figure 3.1G, we have replotted the same data taking the z-axis as the energy per TF, which can be calculated directly from gene expression profiles using our landscape construction (H_{basin}/N). In making these plots we have ignored the contributions of the culture terms in Eq. 3.1 to the energy in our model (see Figures B.3 through B.3 and section). Notice that the faster trajectories follow a lower energy path while the slowest trajectory (ST) follows a high energy path and appears to spend time stuck in two different barriers between days 8 and 21. These observation suggest that the experimentally observed reprogramming dynamics are consistent with the idea of a "barrier crossing" between the starting and ending cell types in a rough landscape (see Figure 3.1B).

Further evidence for this barrier-crossing picture comes from numerical simulation using our landscape model (see section B). The insets in Fig. 3.1E-3.1G show failed and successful reprogramming trajectories from Monte-Carlo simulations. There is a striking similarity between the model trajectories and experiment. Like in experiment, successful reprogramming trajectories in our model follow a simple one-dimensional reaction coordinate in the projection space and reprogramming requires crossing a significant energy barrier. Supplementary Figures B.3 through B.3 contain more examples of successful and unsuccessful simulation trajectories.

Finally, we note that the reaction coordinate can also be visualized using more traditional measures of distances such as the overlap (dot product) of the gene expression profile with the starting and ending states (see Figure B.3A). However, when using overlaps, each experiment has its own starting and ending point, making it hard to compare across experiments. Furthermore, overlaps are unable to discern the "barrier crossing" picture that emerges naturally from using projections (see Figure B.3B).

3.2.3 B cell reprogramming dynamics

The previous section considered reprogramming from MEF to iPSC. Here, we extend this analysis to consider two additional reprogramming experiments from B cells to IPSs [157, 158]. In the first experiment, the standard Yamanaka reprogramming protocol (OSKM) [12] was used to reprogram B cell to iPSC. Unlike in MEFs, in B cells the OSKM protocol resulted in extremely low reprogramming yields. To increase the reprogramming yield, the protocol was then modified so that OSKM expression was preceded by pulsed expression of CEBPα (abbreviated C+OSKM). This modified protocol significantly increased the reprogramming yield. Figure 3.2A shows that for both experiments, reprogramming trajectories once again follow a simple reaction coordinate in projection space. Figure 3.2B extends these plots to the energy vs reaction coordinate plane. Notice, that in both experiments, the energy of the trajectories first increase and then decrease. The higher yield trajectory (C+OSKM) makes steady progress over the energy barrier, while the low yield trajectory (OSKM) appears to meander through inefficient directions. Thus the reprogramming dynamics of B cells are similar to the reprogramming dynamics of MEF: in all cases reprogramming follows a simple one-dimensional reaction coordinate and can be understood as a barrier crossing between minima.

The insets in these figures show results from numerical simulations using the landscape model. The simulations reveal a simple reaction coordinate. However unlike in experiment, the simulated trajectories for the two protocols exhibit nearly identical dynamics. This likely reflects the limitations of the coarse-graining approximation used to construct the landscape model. In the model, TFs are treated as binary variables and all TFs are treated on equal footing – no distinction is made between global chromatin remodelers like CEBPα and more specific downstream factors. Despite these limitations, the phenomenological model still captures the qualitative phenomena seen in the experiments.

The similarity of the reprogramming trajectories from MEFs and B cells suggest a universal reaction coordinate for reprogramming: a straight line connecting the starting

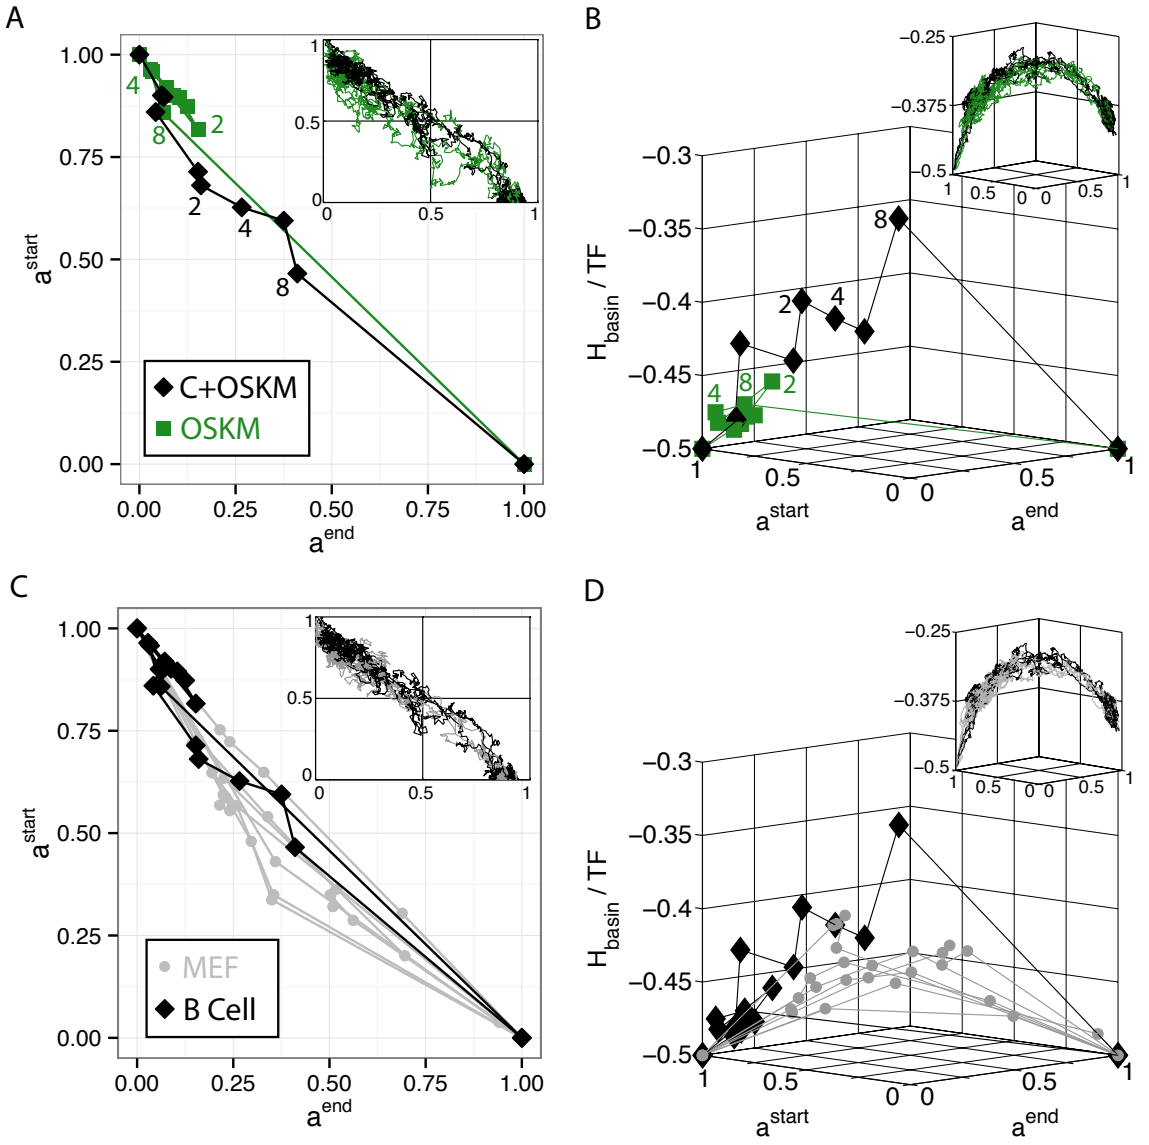

Figure 3.2: **Universal Reaction Coordinate..** A. Direct conversion from a^{start} (B Cells) to a^{end} (iPSC) by Di Stefano et al [157]. OSKM is the standard Yamanaka protocol, while C+OSKM is a pulse of C/EBPα followed by OSKM which led to higher reprogramming yield. All insets are simulation data of same data shown in main figure. See Figure B.3 for larger version of simulations. B. Energy landscape of basins of attraction, H_{basin}, per transcription factor (TF) vs reaction coordinate. See Figure B.3 for larger version of simulations. C. Data collapse of trajectories to a^{start} vs a^{end} for both MEF to iPSC (gray) and B Cell to iPSC (black). See Figure B.3 for larger version of simulations. D. Data collapse of trajectories when viewed as energy vs reaction coordinate. See Figure B.3 for larger version of simulations.

and ending cell type in projection space. This can be seen best in Figures 3.2C and 3.2D where we have plotted reprogramming dynamics from both MEFs and B Cells on the same plots. These experimental data are consistent with numerical simulations using our landscape model which show that reprogramming trajectories always follow a straight line in projection space for both choices of starting cell type.

3.2.4 Insight into dynamics from our mathematical model

Given the strong agreement between experiment and the landscape model, it is interesting to ask if the model can provide further insights into reprogramming dynamics beyond those that can be directly gleaned from analyzing experimental time series. As discussed in the introduction, there is an ongoing debate in the reprogramming literature about the order and organization of gene-level events during reprogramming [112, 147, 148]. To address this, we performed detailed simulations that allowed us to probe gene-level events during reprogramming from MEF to iPSC (see section B). Experimentally, reprogramming times (as measured by reporters for pluripotency markers) are well described as a Poisson process, implying the existence of a single rate limiting step [26]. Our simulation results support the idea of a single rate limiting step to the turning on of pluripotency markers (see Fig 3.3A). In our simulations, the time to turn-on pluripotency markers is calculated by measuring the time it takes a trajectory to have a significant projection on an iPSC state ($a^{end} = 0.3$) Additionally, our simulations show that the later phase of reprogramming (defined as the period of time when trajectories go from having a projection $a^{end} = 0.3$ to $a^{end} = 0.8$) follows a narrowly peaked distribution. Once reprogramming has started, it is very fast: the median time for the later phase is approximately 40 times shorter than the median time for the early phase. Consistent with experiment [26, 146], we find that almost all trajectories eventually reprogram. These results are inconsistent with an "elite" model of reprogramming [159] in which only a special subset of cells are amenable to reprogramming.

To ask about the order of gene level events, we probed the gene level dynamics of 10 genes known to be specific for either MEFs or ESCs (*Snai1*, *Snai2*, *Prrx1*, *Twist2*, *Twist1*

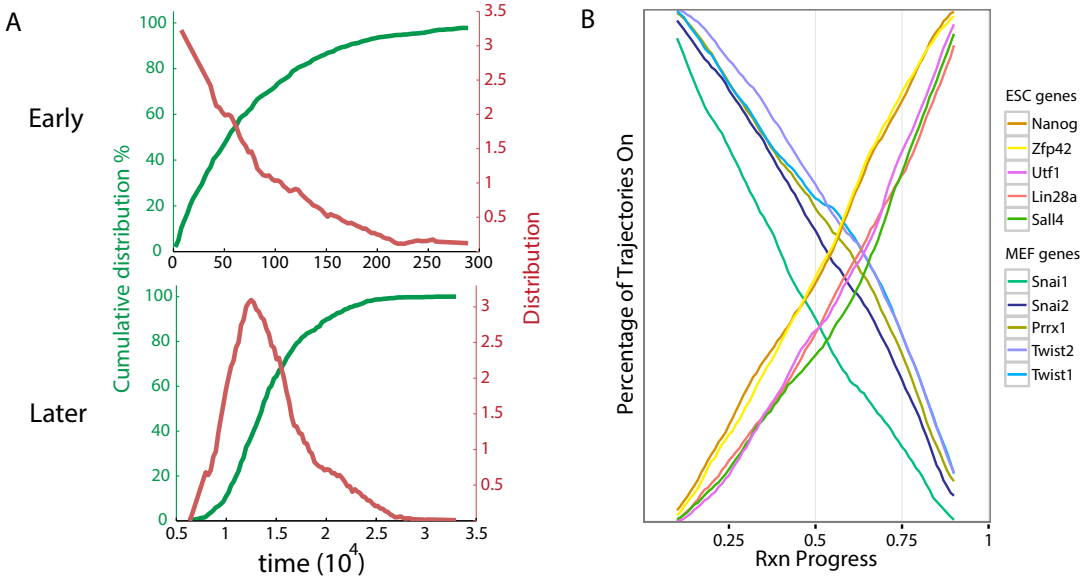

Figure 3.3: **Nature of Reprogramming Dynamics in the Landscape Model..** A. Cumulative distributions of timing show that the early ($a^{end} = 0$ to $a^{end} = 0.3$) and later ($a^{end} = 0.3$ to $a^{end} = 0.8$) stages of reprogramming are respectively a Poisson and a narrowly peaked distribution. See SI Figure 11 for early ($a^{end} = 0$ to $a^{end} = 0.3$), middle ($a^{end} = 0.3$ to $a^{end} = 0.7$) and late ($a^{end} = 0.7$ to $a^{end} = 0.8$) phases of reprogramming as Poisson, narrowly peaked and narrowly peaked distributions respectively. In order to study the complete timing distribution, the data shown here and in Figure B.3 were obtained in a simulation of duration $t = 3 \times 10^6$ MC steps, which is 30 times longer than the simulations reported on in all other figures. B. Percentage of trajectories in which a gene is on vs reaction coordinate. Data shown is a moving average of MEF (ESC) genes turning off (on) over time. See Figure B.3 for example of non-averaged data.

and *Zfp42*, *Nanog*, *Utf1*, *Lin28a*, *Sall4*, respectively) for 224 successful reprogramming trajectories out of a total of 3000 attempts. Recall, that in our model, each gene is represented by a binary variable and can either be 'on' or 'off'. Since the dynamics of our landscape model are stochastic, these genes turn on and off at different values of reaction coordinate in each of these 224 trajectories. To understand if there is any structure in the gene level dynamics, we counted the percentage of trajectories for which a gene was on on at a given reaction coordinate using a moving average (see section B). The results are shown in Fig 3.3B (see Figure B.3 for an example of non-averaged data). The MEF (ESC) genes gradually turn off (on) over time as expected. Furthermore, the order in which genes turn on and off is relatively stable, at least when averaged over trajectories. In contrast, individual simulation trajectories show much more variability in the order which genes turn on. However, if we consider individual pairs of TFs, we find that their ordering tends to be consistent with what one would expect from Fig 3.3B. For example, *Nanog* turns on before *Sall4* in 58% of trajectories, and *Snai1* turns off before *Twist1* in 71% of trajectories, but for *Twist1* and *Twist2*, there is no clear trend for one or the other to turn off first.

All the qualitative features of our simulations are consistent with the idea that reprogramming trajectories correspond to successful "barrier crossing" between two minima in a landscape. An important qualitative prediction of all barrier crossing is that reprogramming trajectories should be dominated by a small number of optimal paths, with some amount of fluctuations around those paths [36, 160]. In particular, the facts that the early phase of reprogramming is well described by a Poisson process, the later phase is described by a narrow distribution of times, and that the median time for the early phase is much longer than the median time for the later phase are all features that would be expected of a simple barrier-crossing process. Furthermore, these simulations show that in a high-dimensional barrier crossing, genes can turn on in a temporally ordered manner (at least when averages over many reprogramming attempts) even though the process is driven entirely by stochasticity.

3.3 Discussion

A common metaphor used to describe cellular identity is Waddington's landscape or the idea of a rugged "epigenetic landscape" in which cell types are basins of attraction. In this picture, cellular reprogramming is envisioned as a process in which one cell type is externally driven out of its basin of attraction, across a barrier, and eventually ends up in the basin of attraction of the desired cell type. Previously, we used ideas from spin physics to introduce a model of cellular identity that can be built from genome expression data. In this paper, we reanalyzed experimental data on reprogramming dynamics in terms of our model and found good agreement between the experiments and simulations of our model.

Our model provides several interesting insights into reprogramming dynamics. We find that reprogramming dynamics proceed along a simple one-dimensional reaction coordinate and must cross a significant energy barrier. Somewhat surprisingly, this reaction coordinate is independent of reprogramming dynamics. In terms of projections, we can simply describe the reaction coordinate as a straight line from ($a^{start} = 1$, $a^{end} = 0$) to ($a^{start} = 0$, $a^{end} = 1$). What makes this simple picture especially interesting is that we demonstrated its validity for two different types of reprogramming experiments (MEF or B Cell to iPSC). Based on simulations with our model, we believe that any cellular interconversion (reprogramming or direct conversion), will proceed along a similar, universal, reaction coordinate when described in terms of a^{start}, a^{end}, and energy.

Our model also gives insight into the ongoing debate about the phases of reprogramming dynamics. A priori, reprogramming dynamics may be either probabilistic or deterministic with respect to both the timing and order of gene level events. Our simulations show the the initial phase of reprogramming follows a Poisson distribution – initiating reprogramming is a rare event. However, once initiated, reprogramming proceeds quickly and efficiently. This is reflected in our simulations by the observation that the dynamics of the reprogramming process at later stages are well described by a narrowly peaked distribution. Furthermore, we find that when averaged over many successful reprogramming trajectories, the order of

gene level events are relatively reproducible. Our simulations strongly support Chung et al [148] description of reprogramming as a "loosely ordered probabilistic process".

Why have different dynamics experiments led to such drastically different conclusions? So far, each experiment has used different techniques, each of which have their own limitations. GFP reporters (for example [26]) provide precise timing data but are limited to small numbers of genes. Whole genome expression data (for example [147]) provides data on all genes, but both microarrays and RNA-Seq require populations of cells. Finally, single cell gene expression data (for example [112]) provides accurate details of gene expression, but only for a subset of genes (currently 48 with standard Fluidigm chips [112]). Therefore, depending on which technique is utilized, each experimentalist rightfully sees a different picture of reprogramming dynamics. However, viewing reprogramming as a loosely ordered probabilistic process unifies all of these different experimental pictures.

Besides examining the gene level reprogramming dynamics, our model provides a clearer picture of the global mechanism behind reprogramming. One of the most surprising aspects of reprogramming is that the over expression of just a few TFs (out of thousands) can lead to such drastic changes in the global gene expression profile. Our simulations suggest the underlying reason for this is the important role played by culturing conditions. In our model, inducing the OSKM TFs in MEFs only changes the energy by 0.5% , which at the noise levels considered here, do not lead to any successful reprogramming event. However, by including the effect of cell culture in our simulations, we achieve 7.43% reprogramming rates. This suggests that culturing conditions likely play an important role in dictating reprogramming efficiencies. For example, it is possible to use the OSKM factors, normally used to reprogram to iPSC, to instead reprogram to blood cells just by changing culture conditions [113]. This highlights an important issue of experimental design for direct conversions to a given cell type. Before one searches for TFs to manipulate, it is essential to understand the correct culturing conditions for the desired cell type. Without the correct medium, direct conversion may prove exceedingly difficult. In our simulations, we have found that the culture term for a given cell type decreases the size of the basin of attraction of all the other cell types.

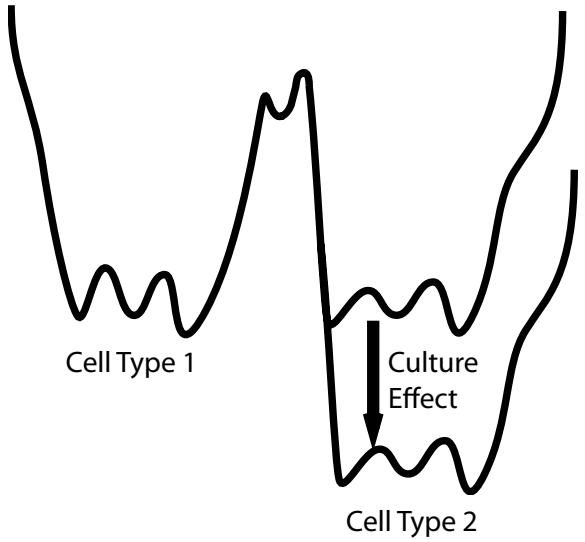

Figure 3.4: **Culture Schematic..** The correct culture conditions plays an essential role in reprogramming by stabilizing the final cell type.

We even find some reprogramming events when we bias the system just by introducing the culture term, without forcing expression of the OSKM TFs (this likely reflects the limitations of the model). However, when we compare simulations of MEF to ESC reprogramming at a certain noise level and for a certain duration, the ones where expression of the OSKM TFs is forced and the ESC culture term is present have a success rate 5 times higher than the ones where the ESC culture term is present but OSKM expression is not forced. In the future, it will be interesting to further explore this tradeoff between stability and plasticity of cell types.

The experimental analysis and simulations presented here suggest that reprogramming can be viewed as a "barrier crossing" in rugged landscape (see Figure 3.4). In all barrier crossings, the dynamics are dominated by a few "optimal paths", suggesting that reprogramming dynamics are likely to be low-dimensional and fairly reproducible at the gene level. A natural consequence of this picture is the existence of a simple reaction coordinate that describes the progress along the optimal path. If the landscape picture is correct, the

existence of a reaction coordinate is likely to be a generic feature of all reprogramming and direction conversion protocols. Directed differentiation is a closely related experimental technique that instead of using TFs to convert between cell types focuses on recapitulating embryonic development through sequences of signaling molecules [161]. It will be interesting to see if projections are also a useful reaction coordinate for directed differentiation experiments.

The results presented here are also likely to be applicable to other systems. Recently, it has been suggested the evolutionary dynamics of viruses such as HIV can also be understood using a Hopfield-inspired landscape model [162]. In evolutionary landscapes, crossing fitness valleys in rugged landscapes can naturally be understood in terms of barrier crossings. For this reason, it is likely that the techniques developed here in the context of cellular reprogramming can be adapted to visualize evolutionary data on fitness crossing dynamics. More generally, landscapes have proven to be an important tool for furthering our understanding a variety of biological problems such a protein folding [124, 163, 164]. The intuitions developed in the context of these other problems are also likely to be applicable to cellular reprogramming and in the future, it will be interesting to explore these connections further.

3.4 Extended Figure 3.1 Caption

A. Transient expression of reprogramming genes plus switching culturing conditions probabilistically leads to the desired cell type. B. Reprogramming is commonly described as the crossing of a barrier in a high-dimensional landscape. C. Our proposed cellular identity landscape is based on the projection, a, of an arbitrary gene expression, S, onto the subspace (gray plane) spanned by the natural cell types, ξ. D. Principal component analysis (PCA) of reprogramming from mouse embryonic fibroblasts (MEF) to induced pluripotent stem cells (iPSC) with start marking day 0 and end marking iPSC. Rais [146], Polo [147], and ST (Samavarchi-Tehrani) [156] are three successful trajectories in which the explicit time in days is labeled on plots E, F, and G. Other represents additional successful trajectories, PRC are partially reprogrammed cells, and failed trajectories do not reprogram.

E. Projection onto a^{start} (MEF) and a^{end} (iPSC) only. All successful trajectories follow a simple reaction coordinate in projection space, a straight line from ($a^{start} = 1$, $a^{end} = 0$) to ($a^{start} = 0$, $a^{end} = 1$). Insets in E, F, and G are simulation data with failed trajectories in red and successful trajectories in gray. See Figure B.3 for larger version of simulations. F. Measure of projection on all other cell types, $a_{\perp}$ vs the reaction coordinate. See Figure B.3 for larger version of simulations. G. Energy landscape of basins of attraction, H_{basin}, per transcription factor (TF) vs reaction coordinate. See Figure B.3 for larger version of simulations.

4 Thermodynamics of statistical inference by cells

The following chapter and the corresponding Appendix and figures have been adapted from: Alex H. Lang, Charles K. Fisher, Thierry Mora, Pankaj Mehta. Thermodynamics of statistical inference by cells. Phys. Rev. Lett. (2014) 113 (14), 148103.

4.1 Introduction

Cells often perform complex computations in response to external signals. These computations are implemented using elaborate biochemical networks that may operate out of equilibrium and consume energy [45, 75, 107, 165–168]. Given that energetic costs place important constraints on the design of physical computing devices [58] and neural computing architectures [59], one may conjecture that thermodynamic constraints also influence the design of cellular information processing networks. This raises interesting questions about the relationship between the information processing capabilities of biochemical networks and energy consumption [65, 66, 69, 70, 109]. Indeed, we will show that thermodynamics places fundamental constraints on the ability of biochemical networks to perform statistical inference. More generally, statistical inference is intimately tied to the manipulation of information and hence offers a rich setting to study the relationship between information and thermodynamics [67, 68, 86, 88, 90].

In order for a cell to formulate an appropriate response to an environmental signal, it must first estimate the concentration of an external signaling molecule using membrane bound receptors [45, 107, 165–169]. The biophysics and biochemistry of cellular receptors is highly variable. Whereas some simple receptor proteins behave like two-state systems (i.e. unbound and ligand bound) with dynamics obeying detailed balance [170], other receptors, such as G-protein coupled receptors (GPCRs), can actively consume energy as they cycle through multiple states. This naturally raises questions about how energy consumption by cellular receptors affects their ability to perform statistical inference. Here, we address these

questions by analyzing the accuracy of statistical inference (i.e. learning) as a function of energy consumption in a simple but biophysically realistic model. We show that learning more accurately always requires expending more energy, suggesting that the accuracy of a statistical estimator is fundamentally constrained by thermodynamics.

4.2 Results

Cells estimate the concentration of an external ligand using ligand-specific receptors expressed on the cell surface. A ligand (usually a small molecule), at a concentration c in the environment, binds the receptor at a concentration-dependent rate, k_+c, and unbinds at a concentration-independent rate, k_- [107] (see Fig. 4.1A). Upon ligand binding, the receptor protein undergoes conformational changes or chemical modifications that alter its activity, sending a signal that the ligand is bound to downstream portions of the biochemical network. During a time interval T, the receptor can undergo multiple stochastic transitions between the unbound nonsignaling state and the bound signaling states. This information is contained in the time series of signaling and nonsignaling intervals (see Fig. 4.1B). After a time T, the cell converts this time series into an estimate for the external concentration. A longer time series T always gives a better estimate for the concentration; however the cell needs to make a decision in a finite time, so we consider T to be fixed to a large but finite value. In principle, the estimate for the concentration could be computed using one of many different statistics that can be obtained from this time series (e.g. average bound time, average unbound time, etc.). Each of the resulting estimators for the external ligand concentration has a different accuracy. Following Berg and Purcell (BP) [107], we measure the accuracy of an estimator for the concentration using its "uncertainty," defined as:

$$\text{uncertainty} := \frac{\langle (\delta c)^2 \rangle}{\bar{c}^2} \tag{4.1}$$

where $\bar{c}$ is the mean and $\langle (\delta c)^2 \rangle$ is the variance of the estimated concentration.

Several methods have been proposed for how a cell may estimate the concentration of the external signaling molecule. In their pioneering paper, Berg and Purcell suggested

estimating the concentration using the average time the receptor was bound during the time T [107]. They showed that the minimal uncertainty a receptor could achieve with this estimator was

$$\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} = \frac{2}{\overline{N}} \tag{4.2}$$

where $\overline{N}$ is the expected number of binding events during the time interval T. For 30 years, many thought that the BP estimator placed a fundamental limit on the accuracy of a cellular receptor. However, in 2009, Endres and Wingreen [166] showed that a cell using maximum likelihood estimation (MLE) based on the average nonsignaling time could reduce its uncertainty by half to

$$\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} = \frac{1}{\overline{N}}. \tag{4.3}$$

However, the increased accuracy of MLE comes at an energetic cost. Previous work [168] established that BP sets a limit for the best possible estimator in equilibrium, implying that any receptor that performs MLE must operate out of equilibrium and consume energy.

In order to study the relationship between thermodynamics and the accuracy of statistical estimators, we introduce a new family of biophysically inspired cellular receptors that interpolate between BP and MLE. In our model, receptors can actively consume energy by operating out of equilibrium (for example by hydrolyzing adenosine triphosphate or ATP). Using this family of models, we show that there is a direct connection between the energy consumed by a receptor and the uncertainty of the resulting estimator. We find that in order to learn more information (decrease its uncertainty), the receptor must always expend more energy (increase entropy production). Note that, in this paper, we restrict ourselves to modeling the receptor and ignore the downstream signaling network that converts the signal from the receptor into a cellular response [70, 109]. Thus, the energies computed here represent lower bounds on the total energy consumed by the statistical estimation network.

Fig. 4.1C shows the simple two-state receptor considered by BP. The binding of an external ligand to the receptor induces a change in the receptor from a nonsignaling state to a signaling state (see Fig 4.1B). The dynamics of this simple two-state receptor always obey detailed balance. Thus, in order to model nonequilibrium receptors, we must consider

receptors with more than two states. Fig 4.1D shows a receptor with three states: one nonsignaling state to which ligands can bind, and two signaling states to which ligands cannot bind. With this extra state, the dynamics of the receptor can break detailed balance by coupling the conformational change in the receptor to another reaction such as the hydrolosis of ATP. In particular, by consuming energy it is possible to drive the system preferentially through a series of state changes [171], for example clockwise in Fig 4.1F and Fig 4.1G. This results in a nonzero probability flux through the state space and positive entropy production.

In order to relate the thermodynamic properties of these receptors to their ability to perform statistical inferences, it is useful to represent receptors as Markov chains. For example, the two-state receptor shown in Fig. 4.1C can be represented as a two-state Markov chain with a state 0 corresponding to the unbound nonsignaling state and state 1 corresponding to the signaling state (see Fig. 4.1E). We choose the transition rates between states in the Markov chain to be identical to the transition rates between conformations of the receptor. The three-state receptor can also be modeled as a three-state Markov chain with a ring structure, with state 0 once again corresponding to the unbound, nonsignaling state (Fig. 4.1F). In this more abstract notation, it is easy to generalize the three-state receptor considered above to a receptor with $L+1$ states (see Fig. 4.1G): L of these states are signaling states that cannot bind the ligand, while the remaining state, 0, corresponds to the nonsignaling state that can bind ligands. For ease of analysis, in this paper, we consider receptors arranged in a ring only. However, our model is a good approximation for more complicated receptors with multiple pathways, so long as the receptor has a single path (for example, of length L^*) that dominates the probability flux, see [97] for details. In that case, the complicated receptor reduces to a single ring of length L^*.

A straightforward calculation shows that for the architectures in Fig 4.1 (see section C for details), the uncertainty of an estimate for the concentration is given by [166]:

$$\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} = \frac{1}{\overline{\overline{N}}}\left[1 + \frac{\langle(\delta\tau_S)^2\rangle}{\bar{\tau}_S^2}\right] \equiv \frac{\mathcal{E}}{\overline{\overline{N}}} \tag{4.4}$$

where N is the number of binding events, $\overline{\tau}_S$ is the mean time spent in the signaling state after binding a ligand, and $\langle(\delta\tau_S)^2\rangle$ is the variance of the time spent in the signaling states. In the second equality, we have defined the coefficient $\mathcal{E}$ which measures the accuracy of an estimator; e.g. $\mathcal{E} = 2$ for the Berg-Purcell limit and $\mathcal{E} = 1$ for MLE. For a given estimator (i.e. a specific architecture and a set of rates $\vec{k}$), we can calculate the mean and the variance of the signaling time by a first passage calculation similar to that in [103] and section C.

Here we provide some intuition for Eq. (4.4). Notice that all the information about the ligand concentration is contained in the event of a ligand binding to the receptor, and the unbinding of the ligand, or the exiting of the signaling state, is independent of concentration. Thus, any variation in the duration of the signaling state adds additional noise to the estimate but does not contain any more information about the concentration. Therefore, the optimal estimator is one where the signaling intervals are completely deterministic and $\langle(\delta\tau_S)^2\rangle = 0$. Comparing Eqs. (4.4) and (4.3), we see that this corresponds to MLE. This is consistent with the well-known fact that MLE is the optimal unbiased estimator for large sample sizes. When the durations of the signaling times are exponentially distributed, like for a two-state receptor, $\langle(\delta\tau_S)^2\rangle = \overline{\tau}_S^2$, then Eq. (4.4) reduces to the BP result given in Eq. (4.2). Finally, in all cases, the uncertainty scales inversely with the average number of binding events $\overline{N}$ during the time interval T. This scaling law follows from the central limit theorem by treating each binding event as an independent sample of the concentration.

The Markov representation allows us to calculate the energy consumption using ideas from nonequilibrium statistical physics. We focus on long time intervals, $T \gg 1$, with many binding events, where the receptor dynamics can be modeled by nonequilibrium steady states (NESS). The entropy production of the Markov process is the energy per unit time (power) required to maintain this NESS, and therefore calculating the entropy production is equivalent to calculating the energy consumed by the biochemical network [109, 171]. The entropy production is given by [172]

$$e_p = \sum_{i=0}^{L}\sum_{j\neq i}^{L} p_i^{ss} k_{ij} \ln\frac{k_{ij}}{k_{ji}}, \tag{4.5}$$

with p_i^{ss} is the steady state probability of state i, k_{ij} is the transition rate from state i to state j, and we have set $k_BT = 1$, see section C for details. For the architectures where the Markov process forms a ring, the entropy production simplifies to

$$e_p = (p_0^{ss}k_{01} - p_1^{ss}k_{10}) \ln \frac{k_{01}k_{12}...k_{L0}}{k_{0L}k_{10}...k_{L,L-1}} = J \ln \gamma \tag{4.6}$$

where J is the net flux around the ring and $\ln \gamma$ is the free energy per cycle [171] (see section C for details). For later reference, the total energy released in ATP hydrolysis is approximately 20 k_BT at room temperature [173]. We note that previous work investigating trade-offs between accuracy and energy in Markov chains used a non-thermodynamically feasible energy [174].

Our goal is to find the best performing estimator for a given receptor architecture and entropy production (energy consumption) rate. However, there are several biological constraints that need to be considered when optimizing over choices of kinetic parameters. First, the rate at which a chemical ligand binds to a receptor is set by diffusion limited binding [107] and hence k_{01} is not controlled by the cell. Therefore we set $k_{01} = 1$ and do not optimize over this rate. Second, a receptor needs to be specific. In principle, both "correct" ligands (i.e. the ligands the receptor has evolved to detect) and "wrong" ligands (any other chemical) can bind the receptor. However, nonspecific ligands quickly unbind and cause the receptor to switch back to the nonsignaling state. Thus, the specificity of a receptor is set by the mean duration of the signaling state in the presence of the correct ligand, $\overline{\tau}_S$. This is incorporated by requiring a small nonspecific binding rate ($k_{0L} = \epsilon \ll 1 = k_{01}$) and we do not optimize over k_{0L}. Lastly, since any statistical estimator is always improved with more samples, to fairly compare different families of estimates, we will fix the sampling rate, $\overline{n} = \overline{N}/T$, where $\overline{N}$ is the expected number of samples and T is the signal integration time. By fixing the nonspecific binding rate (k_{0L}) to be small (see section C for details), this implies $\overline{\tau}_S \approx \overline{n}^{-1} - 1$. But since we are also fixing the sampling rate, $\overline{n}$, this fixes $\overline{\tau}_S$. In summary, our goal is to find the global minima for uncertainty, given the above constraints.

We begin by analyzing the three-state receptor (Fig. 4.1F). Figure 4.2 shows the un-

certainty as a function of entropy production for the optimal three state receptor for four different choices of the ligand binding rate, $\bar{n} = \overline{N}/T$. To generate these plots, we have used an analytic ansatz (see section C for details) for the optimal parameters which we have checked using simulated annealing (with agreement within 1.25%). Notice that learning more accurately (reducing uncertainty) always increased energy consumption (entropy production). At low energy consumption, the receptor approaches the equilibrium BP limit ($\mathcal{E} = 2$), while at high energy consumption (corresponding approximately to the energy of ATP hydrolysis) the optimal performance asymptotically approaches the infinite entropy production analytic limit of

$$\frac{\langle (\delta c)^2 \rangle}{\bar{c}^2} \sim \frac{3}{2\overline{N}} \tag{4.7}$$

One striking observation is that these curves exhibit a data collapse when plotted as a function of the energy consumption per ligand binding rate, $e_p/\overline{n}$. The inset of Fig. 4.2 shows the same curves as the main graph as a function of $e_p/\overline{n}$. Since each ligand binding event can be viewed as an independent sample of the external concentration, this data collapse suggests that the natural variable linking thermodynamics and inference is the energy per independent sample consumed in constructing an estimator.

The three-state receptor is not able to reach the MLE limit of $\mathcal{E} = 1$ for any level of entropy production. To reach the MLE limit, we consider a receptor with $L + 1$ states, L of which are signaling states (see Fig. 4.1G). This Markov chain has $2L$ independent parameters, which makes it hard to find the global optimum. For this reason, we analyzed a simplified, but still biophysically realistic, rate structure (without performing any optimization over parameters) where k_{01}, k_{0L}, k_{10}, k_{L0} can independently vary but all other forward rates are fixed to be identical, $k_{i,i+1} = f$ and all other backward rates chosen so that $k_{i+1,i} = b$, where $i = 1 \ldots L-1$, see section C for details. Once again, for all choices of L, the optimal uncertainty exhibits a data collapse as a function of the energy consumption per ligand binding rate, $e_p/\overline{n}$ (see Fig. 4.3). At low energy consumption, the uncertainty approaches the BP limit ($\mathcal{E} = 2$), while at high energy consumption (corresponding approximately to the energy of ATP hydrolysis) asymptotically approaches the infinite entropy

production analytic limit of

$$\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} \sim \left(1+\frac{1}{L}\right)\frac{1}{\overline{N}} \tag{4.8}$$

Thus, receptors with large energy consumption and many signaling states ($L \gg 1$) approach the MLE limit. In order to perfectly achieve the MLE limit, all backward rates b would need to be 0, leading to infinite entropy production. An interesting feature of these curves is that beyond some scale (which can be achieved by hydrolysis of only a few ATP), the marginal gain in improvement that results from consuming more energy becomes negligible. This is reminiscent of the recently found transition in kinetic proofreading where adding additional energy only marginally improves the error threshold [97, 102]. It will be interesting to see if this is a generic feature of many biochemical information processing circuits.

4.3 Discussion

In conclusion, by analyzing the ability of cells to estimate the concentration of an external chemical signal using nonequilibrium receptors we have established an unexpected link between statistical inference and thermodynamics. Specifically, we found that the efficacy of an estimator for the concentration of a ligand depends on the energy consumed per independent sample by the receptor. Extrapolating this result suggests that there may be fundamental thermodynamic bounds on statistical inference. The trade-off between accuracy and energy is general and may be relevant for other signal transduction systems, such as gene regulation [175], light-activated proteins [101] or ligand-gated ion channels [176]. We note that following the tradition of Berg and Purcell, in this paper we only considered estimating a concentration after a long time T. However, in many related cases, such as transcription [177], the speed is an important trade-off in addition to accuracy and energy consumption. In the context of phosphorelays, it is likely that the circuits can respond quickly even for multistep cascades. For example, the four-stage phosphorelay utilized for phototransduction in the retina can still respond to stimuli in about half a second [106]. Nonetheless, understanding these trade-offs represents an important future research direction.

We conjecture that our observed scaling, $(e_p/\overline{n})$, reflects a general principle: the efficiency of a statistical estimator is limited by the energy consumed per sample during its construction. Of course, much more investigation is needed to see if this conjecture holds in general. In particular, it will be interesting to see if these results change for receptors modeled by heterogeneous Markov networks that are not strictly ringlike in nature. Recent work indicates that at large entropy production the dynamics of such networks may be independent of details of the underlying topology, suggesting that our basic picture should hold even for more complicated nonequilibrium receptors [178]. An additional extension to our model would be to consider externally varying concentrations by implementing a sensory adaptive system (SAS) as was done in recent papers [79, 81]. These papers found that the accuracy and energy consumption of the SAS depends on the time scale of external concentration fluctuations. Finally, it is well known that many receptors, such as GPCRs, actively consume energy in order to operate. Our model presents one possible explanation for this observation. The energy consumption may help reduce noise in the downstream signal, allowing cells to more accurately determine external concentrations. Our model also shows that hydrolysis of only one or two ATP nearly achieves the theoretical minima of uncertainty. This may explain why cell sensors often require only a few phosphorylation sites.

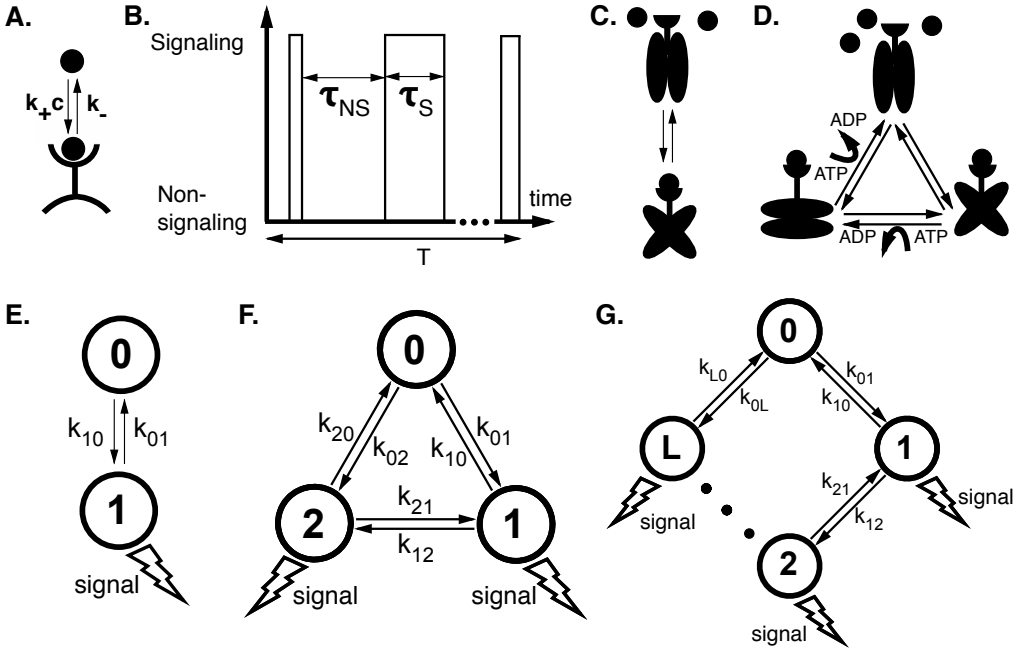

Figure 4.1: **Schematic of a cell receptor and our model of a receptor.** (A) A chemical ligand at concentration c binds to the receptor at rate k_+c and unbinds at rate k_-. (B) Example time series of a receptor binding. While unbound, the receptor is in nonsignaling state, but upon ligand binding it transitions to a signaling state. After a long time T, the receptor has a series of nonsignaling times τ_{NS} and signaling times τ_S from which to estimate the concentration. (C) Two-state and (D) three-state biochemical models of a receptor. Upon ligand binding the receptor undergoes a physical change (represented as a conformational change) that transmits signals to the downstream biochemical network. (E) Two-state, (F) three-state, and (G) L-state Markov models of a receptor, where the chain of states 3, 4, ... $L-1$ has been suppressed.

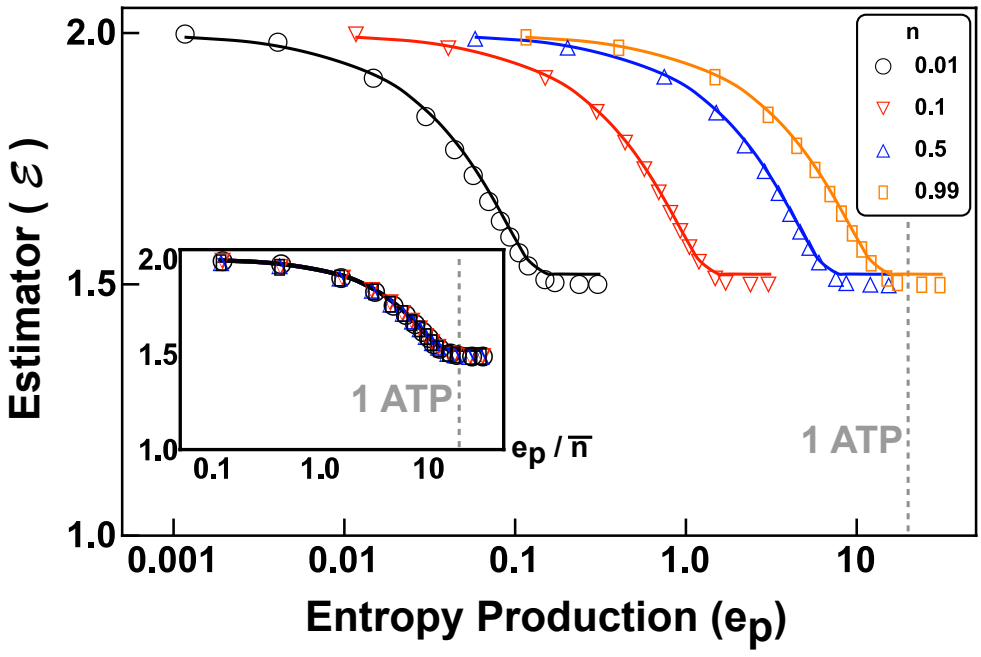

Figure 4.2: **Two signaling state estimator performance.** For varying sampling rate $\overline{n} = \overline{N}/T$, the plot shows estimator performance ($\mathcal{E}$) versus entropy production (e_p with units of $k_BT = 1$). The symbols represent results from simulated annealing, where $k_{01} = 1$ and $k_{02} = \epsilon = 10^{-3}$ while the other four rates are optimized. The continuous lines represent our ansatz (see section C for details) for the global minima. At high entropy production the estimators asymptotically approach 1.5. The inset shows the data collapse when the estimator performance ($\mathcal{E}$) is plotted versus entropy production per sampling rate ($e_p/\overline{n}$). The vertical dashed line corresponds to the approximate energy released by hydrolysis of a single ATP.

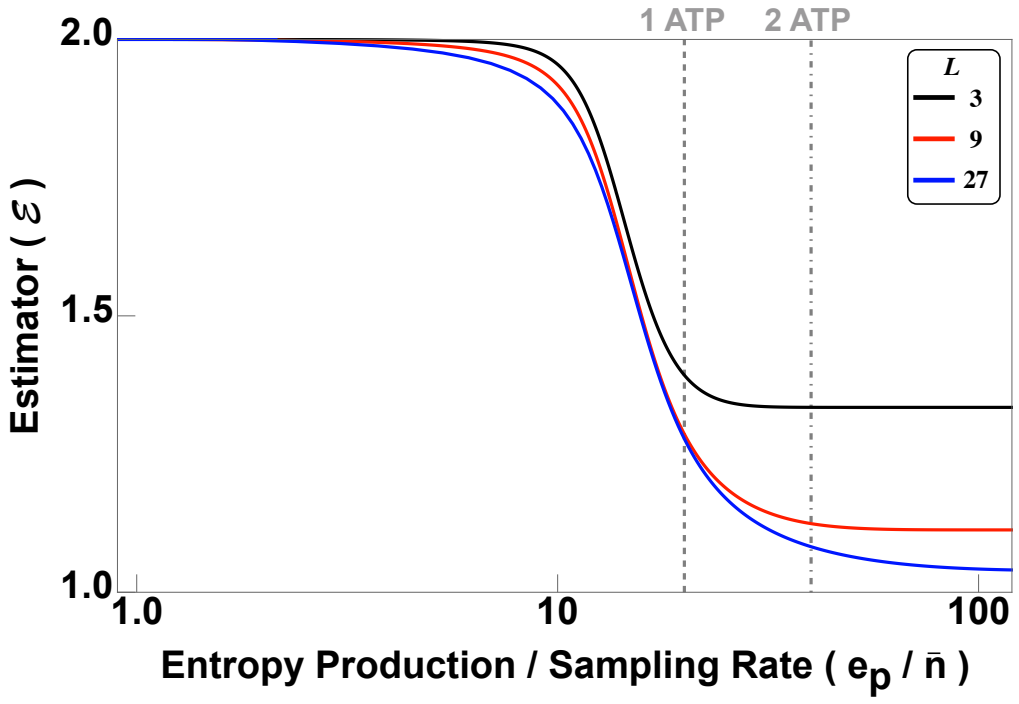

Figure 4.3: **Illustrative example of L signaling state estimator performance.** For a varying number of signaling states L, the plot shows estimator performance ($\mathcal{E}$) versus energy consumption ($e_p/\overline{n}$). For an increasing L, at high energy consumption the estimator approaches the maximum likelihood limit of 1. The following parameters are fixed at $\overline{n} = 0.99$, $k_{01} = 1$, $k_{0L} = 10^{-3}$, $\alpha = k_{10}/b = 10^{-3}$, and $\omega = k_{L0}/f = 1$, while b was varied to keep $\overline{n}$ fixed, and $\theta = f/b$ was varied to change the estimator and the energy consumption. These parameters were chosen for convenience and are not global optima. The vertical lines correspond to the approximate energy released by hydrolysis of a single ATP (dashes) or two ATPs (dot dashes).

Appendices

A Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes

A.1 Data Analysis

An older version of this manuscript, Arxiv v3 [179], has additional microarrays available that are unused in this version of the text. All microarrays used in this paper were taken from the public databases ArrayExpress (www.ebi.ac.uk/arrayexpress) or GEO (www.ncbi.nlm.nih.gov/geo).

There are two datasets, the natural cell fates and the partially reprogrammed cells. For the natural cell fates, we only used the Affymetrix GeneChip Mouse Gene 1.0 ST platform due to the large number of available microarrays on ArrayExpress (www.ebi.ac.uk/arrayexpress) and the better technical design of the platform (1.0 ST has probe matches throughout a gene in contrast to just the 3' UTR in Affymetrix GeneChip Mouse Genome 430 2.0). There is limited data on partially reprogrammed cells so we used microarrays from Affymetrix GeneChip Mouse Genome 430 2.0.

The raw microarray data was converted to an expression level as follows. Microarray probe-to-gene map was created with Bioconductor 2.10. All raw microarray files were initially processed by robust mean averaging (RMA) in MATLAB, and genes with multiple microarray probes were averaged. We did additional processing of this output for two reasons. First, we need to compare microarrays from multiple platforms, but the standard RMA output can vary significantly from platform to platform. Second, since gene expression is a set of positive definite numbers, the minimal assumption model of gene expression is a log-normal distribution. Therefore, to make robust comparisons across platforms, we used order statistics [180]. The RMA output was converted to a rank order. Next, we want to convert this rank order to the z-score of a log-normal distribution. We converting the rank to a percentile (for N genes, divide by $N+1$), and then this percentile into a normal

z-score. For later mathematical convenience, we used a biased estimator (normalize by N not $N-1$) since then the Euclidean norm of each microarray gene expression is N.

At this point, the natural dataset consisted of 601 microarrays with 20719 genes. Since we were interested in cellular identity, only transcription factors, transcription factor co-factors, or chromatin remodeling genes were kept (for short hand, referred to as transcription factors (TF) throughout the text) [181], leaving 1715 TFs.

As explained in the main text, since continuous (sigmoidal input) attractor neural networks and discrete attractor neural networks are known to have the same stable fixed points [182], we used the binarized gene expression. We binarized the gene expression by setting a positive z-score to $+1$ and a negative z-score to -1. While this was mainly done for mathematical convenience, this is potentially biologically justified. Histone modifications (HM) either leave chromatin in an open, accessible configuration or a closed, inaccessible state [127]. We found global HM data for embryonic stem cells (ESC), mouse embryonic fibroblasts (MEF), and neural progenitor cells (NPC) [125, 126]. Consequently, we used the global HM data for these three cell fates and compared them to microarray TF expression levels. This allowed us to create a conditional probability distribution of each HM for a given TF expression level (Figure 2B). We found a sharp cutoff (that coincides with a z-score of 0) which distinguished TFs with the activating modification of histone 3 tri-methylation at lysine 4 (K4) from TFs with the inactivating modification of histone 3 tri-methylation at lysine 27 (K27), poised/bivalent TFs (both K4 and K27), and no HM (most likely DNA methylation). This shows that our mathematical assumption is justified by the HM data.

After the binarization of TF expression, all TFs that were not differentially expressed across cell fates (i.e. TFs that are always on / always off in every cell fate) were dropped, leaving 1337 TFs. The binarized TF expression for the 63 cell fates was found by first binarizing all 601 microarrays and then taking the majority vote for each cell state (with ties broken by averaging the continuous data). The final result was the binary expression state for 63 cell fates.

Microarrays for partially reprogrammed cells were on the Affymetrix GeneChip Mouse

Genome 430 2.0 Array. The same procedure was used to convert raw microarray data to z-score expression. However, since different microarrays do not have the same genome coverage, the analysis comparing partially reprogrammed cells and natural cell fates used the $N = 1329$ TFs common to both platforms.

Several self-consistency checks were performed on the data. First, the correlation matrix $A^{\mu\nu}$ (explained in main text and below) was calculated for the original continuous data and for the binarized data (Figure S1). Both correlation matrices are consistent with each other showing binarization does not change the global correlations. Note that in the correlation matrix, cell fates have been grouped by tissue type, leading to a block diagonal form. Second, the expression state of all cell fates was constructed from multiple microarray experiments. These different experiments were compared with each other and were within 2 standard deviations (std equal to $1/\sqrt{N} \approx 0.027$) for all cell fates. This demonstrates that microarrays from multiple laboratories can be directly compared.

A.2 Landscape Model

Here we give an overview of our epigenetic landscape model. The model is summarized in Tables and , and section A.5 provides a supplementary overview of attractor neural networks.

A.2.1 State Space

Each TF (labeled by i, j) can be in a state $S_i = \pm 1$ where $+1$ indicates the TF is active while -1 indicates it is inactive. A general cell state is given by $\mathbf{S}$, an $N = 1337$ dimensional vector. There are $p = 63$ cell fates (labeled by μ, ν). In cell type μ, the state of TF i is given by ξ_i^μ. The complete cell type data ξ is a p by N matrix determined using our microarray data described above and these ξ are the only biological input into the landscape.

A.2.2 Full Landscape

The complete landscape H can be written as the following terms:

$$H = H_{basin} + H_{bias} + H_{culture} + H_{switch} \tag{A.1}$$

Our landscape assigns an "energy" to every global expression state. We emphasize that this energy does not correspond to physical energy consumption of ATP; instead it is an abstract energy that corresponds to stability and developmental potential of cell fates. Each of the four terms has a simple interpretation (see Figure 2.1). The first term, H_{basin}, ensures that observed cell fates are valleys in our landscape (Figure 2.1A). The second term, H_{bias}, describes biasing of specific TFs by experimentalists (not shown in Figure 2.1). The third term, $H_{culture}$, increases the radius and depth of cell fates that are favored by the environment or culturing conditions (Figure 2.1B). Finally, in the presence of an external signal that gives rise to differentiation (ex. growth factors associated with differentiation), the fourth term, H_{switch}, opens a low energy path between the initial and final cell fates (Figure 2.1C).

A.2.3 Landscape Details: H_{basin}

The gene expression profiles of naturally occurring cell fates must be minima of our landscape. This is ensured by the landscape term

$$H_{basin} = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j\neq i}^{N} S_i J_{ij} S_j \tag{A.2}$$

In order to guarantee that cell fates are basins of attraction, we need to choose the "effective interaction" matrix, J_{ij}, which encodes how the jth TF influences the ith TF. Since we have highly correlated cell fates, we use the projection-method [42] (see section A.5.5 for extended discussion on this choice), which defines the interaction matrix as:

$$J_{ij} = \frac{1}{N}\sum_{\mu=1}^{p}\sum_{\nu=1}^{p} \xi_i^{\mu}(A^{-1})^{\mu\nu}\xi_j^{\nu} \tag{A.3}$$

where ξ_i^{μ} are the natural cell fates and A^{-1} is the inverse of the correlation matrix between cell fates. Since our construction is based on correlations between gene expression

profiles, J_{ij} includes the effect of "indirect" interactions between TFs i and j that are mediated through other TFs (see section A.5 for additional mathematical explanation of this construction). While the current definition implies J_{ij} is symmetric, this can easily be generalized to an asymmetric J_{ij} (see later section A.2.8 for details).

A.2.4 Landscape Details: H_{bias}

The term H_{basin} ensures that all cell fates are global minima of the landscape. However, additional terms in the landscape are needed in order to incorporate key experimental features.

First, biologists can directly manipulate gene expression. For example, during the Yamanaka experiment, the TFs *Pou5f1 (Oct4)*, *Sox2*, *Klf4*, and *Myc* are overexpressed in fibroblasts. Mathematically, we represent the overexpression of TF i by a local biasing field B_i that ensures that $S_i = 1$. Therefore the Yamanaka reprogramming protocol enters the landscape as:

$$H_{bias} = -\sum_{i=1}^{N} B_i S_i \tag{A.4}$$

where for the Yamanaka protocol, $B_{Pou5f1} = B_{Sox2} = B_{Klf4} = B_{Myc} \to \infty$ and for any other TF i, the field $B_i = 0$.

A.2.5 Landscape Details: $H_{culture}$

Currently, the basins of attraction H_{basin} are all set to the same minima value. However, environmental signals (such as cell culture conditions) can stabilize and destabilize specific cell fates (see Figure 2.1B). This can be incorporated into our landscape by terms such as:

$$H_{culture} = -N\sum_{\mu=1}^{p} b^{\mu} a^{\mu} \tag{A.5}$$

$$= -\sum_{i=1}^{N} C_i S_i \tag{A.6}$$

where b^{μ} represents the culture biasing, and a^{μ} is the projection onto cell fate μ. This bias can be equivalently expressed at the level of TFs by defining a culture bias, C_i, for the ith

TF given by:

$$C_i = \sum_{\mu=1}^{p} \sum_{\nu=1}^{p} b^\mu \left(A^{-1}\right)^{\mu\nu} \xi_i^\mu \tag{A.7}$$

For example during the Yamanaka protocol, cells are cultured in conditions favorable to ESC, which is mathematically represented by $b^{ESC} > 0$, while for all other cell fates μ, $b^\mu = 0$.

A.2.6 Landscape Details: H_{switch}

During standard development, cells switch fates deterministically in response to external signals. We mathematically represent this using the term:

$$H_{switch} = -\frac{N}{2} \sum_{\mu=1}^{p} \sum_{\nu=1}^{p} m^\mu G^{\mu\nu} a^\nu \tag{A.8}$$

$$= -\frac{1}{2} \sum_{i=1}^{N} \sum_{j \neq i}^{N} S_i K_{ij} S_j \tag{A.9}$$

where m^μ is the overlap on cell fate μ, a^ν is the projection onto cell fate ν, and the matrix $G^{\mu\nu}$ is the developmental signal matrix that is a dynamic entity and a function of developmental time and external signals. We can equivalently write this in terms of transcription factors using the gene-interaction matrix, K_{ij}, defined as:

$$K_{ij} = -\frac{1}{N} \sum_{\mu=1}^{p} \sum_{\nu=1}^{p} \sum_{\rho=1}^{p} \xi_i^\mu G^{\mu\nu} \left(A^{-1}\right)^{\nu\rho} \xi_j^\rho \tag{A.10}$$

where ξ is the natural cell fate states, $G^{\mu\nu}$ is the developmental signal matrix, and A^{-1} is the inverse correlation matrix. Since $G^{\mu\nu}$ is asymmetric, K_{ij} is also asymmetric and explicitly breaks detailed balance (see later section A.2.8 for details).

We now explain the development signal matrix in more details. If $G^{\mu\nu} > 0$, this opens up a low energy path between cell fate ν and cell fate μ. For example, during blood development, the common myeloid progenitor (CMP) can differentiate into either granulomonocytic progenitors (GMP) or megakaryocyte-erythroid progenitors (MEP). The complicated external signals that induce switching from a CMP to GMP leads to $G^{GMP,CMP} > 0$ and all other $G^{\mu\nu} = 0$, while the signals that induce switching from a CMP to MEP leads to

$G^{MEP,CMP} > 0$ and all other $G^{\mu\nu} = 0$. We emphasize that this term is purely phenomenological and further research will be needed to directly connect the developmental biology signals (such as $TGF\beta$, WNT, etc) to the matrix elements $G^{\mu\nu}$.

A.2.7 Dynamics

We have uniquely defined the landscape H. However, there are multiple ways to implement dynamics on this landscape. In this paper, we are primarily interested in the behavior of the stable fixed points and not dynamical trajectories. Therefore, we follow the standard convention in the attractor neural network literature and update the network by random, asynchronous updates (Glauber dynamics) [43]. Therefore, at each update, a random TF, i, is selected and updated according to the probability

$$P[S_i(t+1)] = \frac{e^{\beta h_i(t) S_i(t+1)}}{e^{\beta h_i(t)} + e^{-\beta h_i(t)}} \tag{A.11}$$

where S_i is the expression state of the ith TF, β is an effective noise parameter, h_i is the local field, and t is the time index. The local field h_i is the gradient of the landscape (covariant derivative) defined for the full landscape H as:

$$h_i \;=\; \sum_{j \neq i}^{N} J_{ij} S_j + B_i + C_i + \sum_{j \neq i}^{N} K_{ij} S_j \tag{A.12}$$

where J_{ij} is the basin-inducing interaction matrix, B_i is the experimentally induced bias on the ith TF, C_i is the culturing-condition specific bias on the ith TF, and K_{ij} is the developmental interaction matrix.

We have introduced the effective noise parameter $\beta = 1/T$ (i.e. inverse temperature) that controls the level of stochasticity resulting from biochemical noise. When $\beta \to \infty$, the update approaches a deterministic step function, while when $\beta \to 0$ each state is equally likely. Based on the currently available static genomic data, this update time cannot be directly related to biological time. Finally, we emphasize that since in this paper we are primarily concerned with the structure of the landscape, our results are independent of our choice of dynamics (see section A.5 for detailed discussion on dynamics).

A.2.8 Landscape vs Pseudo-Landscape

Currently, the interactions between TFs, J_{ij}, are symmetric. In real biology, this is unlikely to be true. We can introduce asymmetry into the interactions by randomly deleting interactions (for example Figure 2.2E Diluted). This asymmetry means that influence of TF i on TF j no longer equals the influence of TF j on TF i. This asymmetry breaks detailed balance and implies a non-Lyapunov pseudo-potential [39, 43, 123] and has been shown to be an additional source of noise on the basins of attraction [43].

We also note that the landscape term H_{switch} is explicitly non-equilibrium and breaks detailed balance. Under one set of environmental conditions, $G^{\mu\nu} > 0$ while $G^{\nu\mu} = 0$ driving switching from $\nu \to \mu$, while under a different set of environmental conditions, $G^{\nu\mu} > 0$ while $G^{\mu\nu} = 0$ driving switching from $\mu \to \nu$. Therefore, by including H_{switch} we are actually making our landscape a pseudo-landscape (i.e. non-Lypanouv).

A.3 Simulations

Here we include details of the simulations in this paper. For all simulations, we set $\beta = 1/0.45 \approx 2.2$ and evolved the system for $100,000$ TF updates.

In Figure 2.2E, we demonstrate that we have basins of attraction. The initial conditions were created by taking the ESC expression vector and randomly flipping 15% of the TFs. After every 5000 updates of asynchronous dynamics, burst errors were introduced by randomly flipping 2% of TFs. For the asymmetric dilution, the standard interaction matrix J_{ij} was created. Then 20% of matrix entries were randomly set to 0.

In Figure 2.2F, we demonstrate that the landscape can deterministically switch between basins. The initial conditions were always the CMP expression vector. For signal 1, we set $G^{GMP,CMP} = 0.5$ and all other $G^{\mu\nu} = 0$. For signal 2, we set $G^{MEP,CMP} = 0.5$ and all other $G^{\mu\nu} = 0$.

A.4 Spurious Attractors

Here we provide more details on spurious attractors and hybrid cell fates. As explained in more detail in section A.5, for the traditional Hopfield model, these spurious attractors take the form of odd-majority vote mixtures [43] (i.e. majority vote at each TF of $3, 5, 7, \ldots$ of the ξ_i^μ). The projection method also has the additional spurious attractors of any linear combination of ξ_i^μ that spans the discrete state space (see geometric interpretation given in section A.5) [42]. For convenience, we use the word hybrid as the collective term for either majority vote mixtures or linear combinations of existing cell fates.

As discussed in the main text, the prediction of spurious attractors in the projection method inspired us to reexamine data on existing partially reprogrammed cells. Surprisingly, we found that partially reprogrammed cells could be thought of as hybrids of existing cell fates. However, we are currently only able to obtain qualitative agreement between partially reprogrammed cells and the predicted nature of the spurious attractors. While it is known that the projection method retains these odd-majority vote mixtures spurious attractors, the correlations between states implies these spurious attractors may no longer be symmetric mixtures. However, the exact nature of these spurious attractors is not known and will be explored in future work.

A.5 Attractor Neural Networks: Additional Details

This supplementary text gives a brief introduction to Hopfield neural networks [40, 41] and how they can be adapted to study epigenetic landscapes. We begin by reviewing the basic principles underlying the original Hopfield neural network. We then show how to generalize this to continuous spins [182] as well as discrete spins with correlated cell fates [42] (projection method). For an in-depth introduction to neural networks, please see the beautiful book by Amit [43].

A.5.1 Discrete, Standard Hopfield

There are N genes and each gene i is either on or off, with the output denoted by $S_i = \pm 1$. Alternatively, we could use the variables $\widetilde{S} = \frac{1}{2}(S+1) = 1, 0$ with the corresponding substitutions in all equations below.

The input to a given gene i is denoted by the local field

$$h_i = \sum_{j \neq i}^{N} J_{ij} S_j + B_i \tag{A.13}$$

where J_{ij} is the interaction between gene i and gene j and B_i is the external (i.e interaction independent) bias of gene i. Both J_{ij} and B_i are assumed to be independent of S_i.

The landscape H is given by

$$H = -\frac{1}{2} \sum_{i=1}^{N} \sum_{j \neq i}^{N} S_i J_{ij} S_j - \sum_{i=1}^{N} B_i S_i \tag{A.14}$$

$$= -\frac{N}{2} \sum_{\mu=1}^{p} (m^\mu)^2 - N \sum_{\mu=1}^{p} b^\mu m^\mu \tag{A.15}$$

where in equation A.15 we have introduce the order parameter for the overlap (dot product or "magnetization") of a spin configuration with a given cell fate μ as m^μ and also introduced the cell fate bias b^μ. The overlap is defined in terms of the cell fate vectors ξ_i^μ as:

$$m^\mu = \frac{1}{N} \sum_{i=1}^{N} \xi_i^\mu S_i \tag{A.16}$$

To prove that H is a Lypanov function (i.e. has stable equilibrium states and follows the standard definition of an "energy"), it is necessary to show that H is a decreasing function and bounded below. To do so, consider flipping a single S_i. The resulting change in H is

$$\Delta H = -\frac{1}{2} \left[\sum_{j \neq i}^{N} J_{ij} S_j + \sum_{j \neq i}^{N} S_j J_{ji} + B_i \right] \Delta S_i \tag{A.17}$$

When we have symmetric interactions, $J_{ij} = J_{ji}$, this simplifies to

$$\Delta H = -\left[\sum_{j \neq i}^{N} J_{ij} S_j + B_i \right] \Delta S_i = -h_i \Delta S_i \tag{A.18}$$

To determine the sign of ΔH we need the relation between h_i and ΔS_i. For deterministic (stochastic) dynamics, as long as ΔS_i and h_i are always (usually) the same sign, we always (usually) have $\Delta H < 0$. Therefore, any set of dynamics that stochastically matches the sign of ΔS_i and h_i will lead to H being a Lypanov function. This implies that any choice of dynamics leads to the same stable fixed points, but may give rise to different trajectories, limit cycles, and sizes of basins of attraction for fixed points, see Amit [43] section 2.2 and 3.5 for a detailed analysis. Therefore, in this paper we focus on predictions that are independent of the exact dynamics. This is equivalent to thinking about the stationary properties of the model.

We will follow the standard convention for neural networks and physics and implement Glauber dynamics which is an asynchronous, stochastic update rule. In this update scheme, at each time step, one gene is selected at random and probabilistically updated according to its local field

$$P[S_i(t+1)] = \frac{e^{\beta h_i(t) S_i(t+1)}}{e^{\beta h_i(t)} + e^{-\beta h_i(t)}} \tag{A.19}$$

with h_i defined above (or equivalently $h_i = -\frac{\partial H}{\partial S_i}$) and t time measured in discrete updates. Also, $\beta = 1/T$ is the inverse temperature and characterizes the slope of the sigmoid function. When $\beta \to \infty$, the sigmoid approaches a deterministic step function, while when $\beta \to 0$ each state is equally likely.

Now we need to specify the gene interaction J_{ij} and establish the global minima of the system. There are p cell fates and the state of gene i in cell fate μ is given by ξ_i^μ. The gene interaction is a correlation based interaction and in the standard Hopfield neural network it is defined as

$$J_{ij} = \frac{1}{N} \sum_{\mu=1}^{p} \xi_i^\mu \xi_j^\mu \tag{A.20}$$

In the standard Hopfield network, the cell fates have two assumptions. First, each cell fate is assumed to on average be unbiased (i.e. equal number of positive and negative spins)

$$\frac{1}{N}\sum_{i=1}^{N}\xi_i^\mu \approx 0 \tag{A.21}$$

and second every pair of cell fates is approximately orthogonal

$$\frac{1}{N}\sum_{i=1}^{N}\xi_i^\mu\xi_i^\nu \approx \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \tag{A.22}$$

These two assumptions can be relaxed in extensions of the standard Hopfield neural network, see later sections for one example (the projection method) that can incorporate correlated cell fates.

Now we can prove that each cell fate is a global minima of the landscape. For no external fields, the landscape can be written as:

$$H = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j\neq i}^{N} S_i J_{ij} S_j = -\frac{N}{2}\sum_{\mu=1}^{p}\left(\frac{1}{N}\sum_{i=1}^{N}\xi_i^\mu S_i\right)^2 + \frac{1}{2N}\sum_{i=1}^{N}\sum_{\mu=1}^{p} S_i\xi_i^\mu\xi_i^\mu S_i \tag{A.23}$$

This can be rewritten in terms of the overlap as:

$$H = -\frac{N}{2}\mathbf{m}^2 + \frac{1}{2}p \tag{A.24}$$

Then as long as N is large compared to p, whenever we are in a given cell fate the energy is $H = -N/2$ and this is the lowest bound since $\mathbf{m}^2 \leq 1$. We have shown that for $p \ll N$, H is a decreasing, bounded function and hence is a Lypanov function. When p and N are both large, a full replica calculation shows that H remains a Lypanov function [183].

While we have established that the landscape is a Lypanov function, we also need to examine the dynamical stability of the cell fates and the existence of spurious attractors. In the absence of stochastic update noise ($\beta \to \infty$), we can examine the signal-to-noise ratio of the cell fates. If a state is dynamically stable, one needs $S_i h_i > 0$. When the state is in a given cell fate (without loss of generality assume cell fate 1), we have that

$$\xi_1^1 h_1 = \frac{1}{N}\sum_{j\neq i}^{N}\sum_{\mu}^{p}\xi_1^1\xi_1^\mu\xi_j^\mu\xi_j^1 \tag{A.25}$$

which can be broken into a signal term (first term) and noise term (second term) as follows:

$$\xi_1^1 h_1 = \frac{N-1}{N} + \frac{1}{N}\sum_{j\neq i}^{N}\sum_{\mu\neq 1}^{N}\xi_1^1\xi_1^\mu\xi_j^\mu\xi_j^1 \tag{A.26}$$

For large N, the signal term approaches 1. We can evaluate the noise term by recognizing that it is an unbiased sum of $(N-1)(p-1) \approx Np$ random steps, and therefore has mean 0 and standard deviation $\sqrt{pN}$, giving us

$$\xi_1^1 h_1 = 1 + \mathcal{O}\left(\sqrt{\frac{p}{N}}\right) \tag{A.27}$$

Therefore as long as N is much larger than p, every cell fate is a fixed point. This rough signal-to-noise argument can be made more rigorous by a spin-glass replica calculation [183] which finds that cell fates are stable (in the case $\beta \to \infty$) as long as the ratio of p/N is less than 0.138.

Here is an intuitive argument of why the landscape must be rugged, which implies the scaling of stable states with N. From looking at small systems, a naive guess would be that the number of stable states should scale with the size of the state space 2^N. This scaling could be achieved if each minima occurred when a single TF state is turned on while all the other TFs are off. However, this implies that each minima is only marginally stable; any spin flip will move the state out of the minima. In order to have a basin of attraction, more TFs are needed to determine the minima. A simple error correction or redundancy could be implemented by using r redundant TFs, but this would require exponentially more states r^N. Instead, stable states could be determined by overlapping sets of TFs, as in the Hopfield neural network. This form of error-correction leads to frustration and Gaussian noise between the stable states, hence the scaling of stable states with N and not 2^N.

An unavoidable consequence of the non-linearity (ruggedness) of the Hopfield network is that in addition to the desired attractors (the input cell fates), there are additional spurious, metastable, attractors. There are a variety of spurious attractors, but the most common

are symmetric mixtures of odd states [41], for example without loss of generality we can make a spurious state with the first three cell types, $S_{spur} = \text{major}\,(\xi_1 + \xi_2 + \xi_3)$, where major stands for majority vote (equivalently the sign function) at each spin. The most common spurious attractor are symmetric mixtures of 3 states (as in the example above). A signal-to-noise analysis can also be done to establish that these spurious attractors are stable attractors, but with a smaller basin of attraction than the input cell fates (see Amit 4.3 for details [43]).

A.5.2 Continuous, Standard Hopfield

The previous section describes the basic ideas of Hopfield neural networks. Here, we show how discrete Hopfield neural networks can be considered a limiting case of continuous differential equations of gene expression. We start by defining continuous spins, Σ_i, that can take on real number between -1 and 1. For continuous dynamics, we must modify the dynamics of the corresponding local field. In particular, if the local field decays in time with a time constant τ_i we have

$$\frac{dh_i}{dt} = \sum_{j \neq i}^{N} J_{ij}\Sigma_j + B_i - \tau_i^{-1} h_i \tag{A.28}$$

where the J_{ij} are the same as in the discrete case and the spin Σ_i is related to the local field by some monotonic function $\Sigma_i = g_i\,[h_i]$.

Now the landscape is given by

$$H = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j \neq i}^{N} \Sigma_i J_{ij} \Sigma_j - \sum_{i=1}^{N} B_i \Sigma_i + \sum_{i=1}^{N} \tau_i^{-1} \int_{-1}^{\Sigma_i} g_i^{-1}\,[\Sigma]\, d\Sigma \tag{A.29}$$

where the first two terms are the same as in the discrete case while the third is the new term for continuous only. Taking derivatives with respect to time gives us

$$\frac{dH}{dt} = -\sum_{i=1}^{N} \frac{d\Sigma_i}{dt}\left(\sum_{j \neq i}^{N} J_{ij}\Sigma_j + B_i - \tau_i^{-1} h_i\right) = -\sum_{i=1}^{N} \frac{d\Sigma_i}{dt}\frac{dh_i}{dt} \tag{A.30}$$

Then since $h_i = g_i^{-1}\,[\Sigma_i]$, we can relate the derivative of h_i to the derivative Σ_i. Then using the fact that g_i is monotonically increasing we can show that the change in H is

always negative:

$$\frac{dH}{dt} = -\sum_{i}^{N} g_i^{-1}\left[\Sigma_i\right]\left(\frac{d\Sigma_i}{dt}\right)^2 \leq 0 \tag{A.31}$$

The decrease in H along with the fact that H is bounded below, shows that we have a Lypanuv function. It is easy to see that every discrete stable point is also a stable point in the continuous model; however, the continuous Hopfield neural networks can have additional stable points.

A.5.3 Continuous Gene Expression

A popular approach to model gene interactions is based on the genetic toggle switch [35] and represents gene interactions by a Hill function. For now, we will use the general variable $\widetilde{\sigma} \in [\sigma_{min}, \sigma_{max}]$.

In the most general case, we have that

$$\widetilde{\sigma}_i = \text{sign}(h_i)\frac{a_i|h_i|^{n_i}}{k_i^{n_i} + |h_i|^{n_i}} + b_i \tag{A.32}$$

where the input h_i is in the range $[-\infty, \infty]$ and the output σ_i is in the range $[-a_i + b_i, a_i + b_i]$.

If we rescale every gene by its dynamic range and center the Hill function at zero, we get that $\widetilde{\sigma} = \Sigma \in [-1, 1]$ and

$$\Sigma_i = \text{sign}(h_i)\frac{|h_i|^{n_i}}{k_i^{n_i} + |h_i|^{n_i}} \tag{A.33}$$

Using the function above for $\Sigma_i = g_i\left[h_i\right]$ allows one to relate continuous Hopfield neural networks to gene expression using Hill coefficients.

A.5.4 Discrete as Limit of Continuous

How can we relate the continuous model of gene expression to the previous discrete model? There are two limits. First, if we take the discrete time limit with the update time much

greater than the input memory, we get

$$h_i(t+1) = \sum_{j \neq i}^{N} J_{ij} S_j(t) + B_i \tag{A.34}$$

Second, in the genetic toggle switch language, when the cooperativity is large $n \gg 1$, then $S_i \to \pm 1$. This gives us a deterministic, discrete model of gene expression. If we introduce stochasticity through Glauber dynamics, we completely recover the discrete Ising model of gene expression.

A.5.5 Discrete, Projection Method

The standard Hopfield attractor neural network assumes that the "memories" (cell fates) have nearly no correlations amongst themselves. However, cell fates are highly correlated (see Figure). Therefore, instead of the standard Hopfield attractor neural networks, we will implement the projection method neural networks [42].

The correlation between cell fate μ and ν is given by

$$A^{\mu\nu} = \frac{1}{N} \sum_{i=1}^{N} \xi_i^\mu \xi_i^\nu \tag{A.35}$$

Now the inferred correlation-based, TF interaction matrix is

$$J_{ij} = \frac{1}{N} \sum_{\mu=1}^{p} \sum_{\nu=1}^{p} \xi_i^\mu (A^{-1})^{\mu\nu} \xi_j^\nu \tag{A.36}$$

Then the landscape can be rewritten as

$$H = -\frac{1}{2} \sum_{i=1}^{N} \sum_{j \neq i}^{N} S_i J_{ij} S_j = -\frac{1}{2N} \sum_{i=1}^{N} \sum_{j \neq i}^{N} \sum_{\mu=1}^{p} \sum_{\nu=1}^{p} S_i \xi_i^\mu (A^{-1})^{\mu\nu} \xi_j^\nu S_j \tag{A.37}$$

$$= -\frac{N}{2} \sum_{\mu=1}^{p} m^\mu a^\mu \tag{A.38}$$

where in equation A.38 we have introduced the projection order parameter a^μ which is the orthogonal projection of a spin vector onto the subspace spanned by the stable cell fates

$$a^\mu = \sum_{\nu=1}^{p} (A^{-1})^{\mu\nu} m^\nu = \sum_{\nu=1}^{p} \sum_{i=1}^{N} (A^{-1})^{\mu\nu} \xi_i^\nu S_i \tag{A.39}$$

A simple geometric picture illustrates that H makes each cell fate a global minimum of the landscape. An arbitrary vector can be rewritten in terms of its projection in the cell fate subspace and its orthogonal component δS_i,

$$S_i = \sum_{\mu=1}^{p} a^\mu \xi_i^\mu + \delta S_i \tag{A.40}$$

Then, the distance of an arbitrary vector $\mathbf{S}$ to the cell fate subspace is given by Δ,

$$\Delta = \left(\sum_{i=1}^{N} (\delta S_i)^2 \right)^{1/2} \tag{A.41}$$

which can be rewritten as

$$\frac{\Delta^2}{N} = 1 - \sum_{\mu=1}^{p} a^\mu m^\mu \tag{A.42}$$

This allows us to rewrite the stabilizing term of the landscape as

$$H = -\frac{N}{2} + \frac{1}{2}\Delta^2 \tag{A.43}$$

This provides a very clear interpretation of the landscape as the global distance of an arbitrary vector $\mathbf{S}$ to the natural cell fate subspace [42].

Again, let's examine the signal-to-noise of cell fates in the absence of stochastic update noise. If a state is dynamically stable, one needs $S_i h_i > 0$. When the state is a given cell fate (without loss of generality assume cell fate 1), we have that

$$\xi_1^1 h_1 = \frac{1}{N} \sum_{j \neq i}^{N} \sum_{\mu=1}^{p} \xi_1^1 \xi_1^\mu \left(A^{-1}\right)^{\mu\nu} \xi_j^\nu \xi_j^1 \tag{A.44}$$

$$= \sum_{\mu=1}^{p} \xi_1^1 \xi_1^\mu \left(A^{-1}\right)^{\mu\nu} A^{\nu 1} = 1 \tag{A.45}$$

Therefore, the stability of cell fate 1 has no noise interference from the other cell fates, and we have that cell fates are stable up to $p/N = 1$.

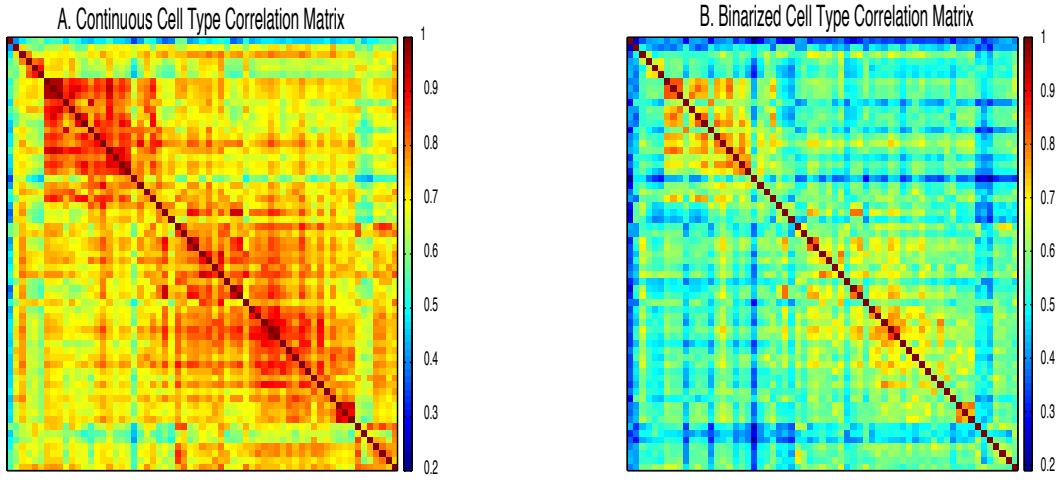

Figure A.1: **Cell fate correlation matrices.** (A) Correlation matrix between cell fates for continuous data. (B) Correlation matrix for binarized data.

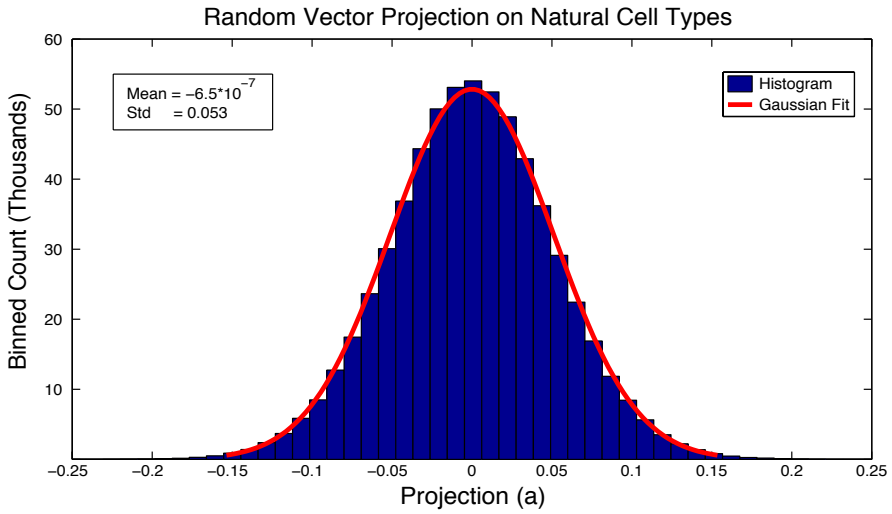

Figure A.2: **Projection of a random vector on a given cell fate.** Ten thousand binarized random vectors were created in MATLAB and projected onto the cellular subspace. The histogram shows the distribution of the projections. The red line is a Gaussian fit to the histogram. The mean is practically zero while the standard deviation is 0.053.

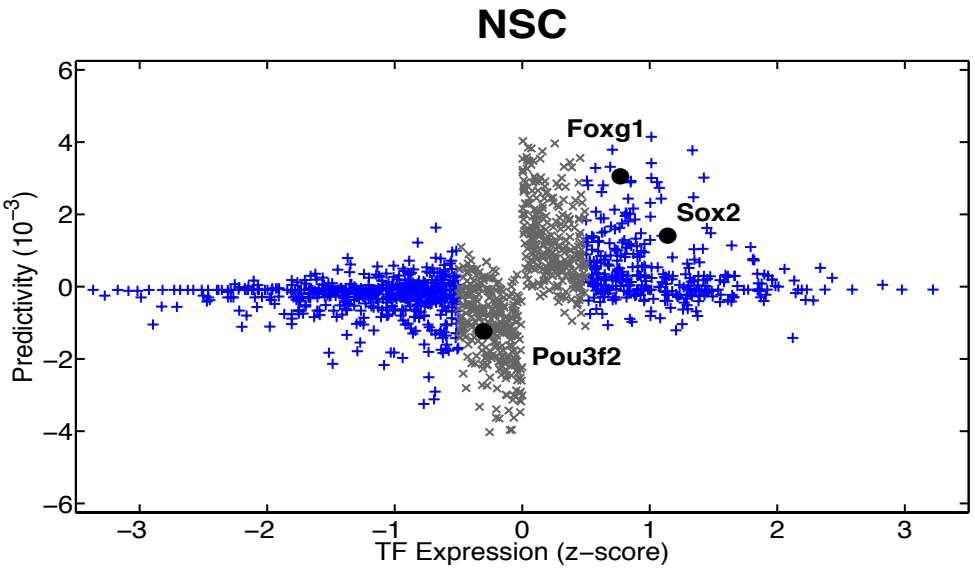

Figure A.3: **Predictivity vs Expression for NSC.** Same type of plot as Figure 2.3. Labeled TFs are part of reprogramming protocol to NPC [10]. This illustrates that *Foxg1* is predictive for NSC, even though it is not for NPC.

A.6 Additional Tables

A.6.1 Classifying Top ESC Reprogramming Candidates

Tables A.1 and A.2 have top 50 embryonic stem cell (ESC) reprogramming candidates (as ranked by z-score times predictivity, η_i^{μ}). Classification of each TF is either justified by paper citation or GO Process term. Reprogramming TFs are in a pre-existing reprogramming protocol, pluripotency TFs help maintain the ESC state but are non-essential for reprogramming, differentiation TFs are expressed in ESC but help induce cell fate change *in vivo*, and unknown TFs have no known function. Abbreviations: biological process (BP), bone morphogenesis or skeletal system morphogenesis (bone), skeletal muscle cell differentiation or regulation of skeletal muscle cell differentiation (muscle), erythrocyte differentiation (red blood), regulation of transcription or transcription, DNA-templated or regulation of transcription, DNA-templated(RT), ectoderm development (ect dev), negative regulation of transcription (NRT), lung development (lung), neural tube development (neural), embryo implantation (embryo), lateral mesoderm development (mesoderm), genetic imprinting (GI), chromatin modification (CM), telomere lengthening (telomere), brown fat cell differentiation (brown), and cell differentiation (cell diff).

Table A.1: Top ESC Candidates. Part 1.

TF	Z Score	η (10^{-3})	Rank Z$*\eta$	Classification	Term
Pou5f1 (*Oct4*)	2.77	2.45	1	Reprogramming	[5]
Gm13242	1.59	3.98	2	Unknown	BP
Nr0b1	2.44	2.59	3	Pluripotency	[131]
Nanog	2.30	2.65	4	Reprogramming	[5]
Zfp42	2.04	2.74	5	Pluripotency	[12]
Hsf2bp	1.42	3.49	6	Unknown	BP
Esrrb	1.74	2.49	7	Reprogramming	[5]
Zscan4f	1.01	3.86	8	Reprogramming	[184]
Klf4	1.04	3.25	9	Reprogramming	[5]
Zfp459	0.83	3.98	10	Unknown	BP
Zscan4c	0.82	3.86	11	Pluripotency	telomere
Zic3	1.17	2.65	12	Pluripotency	[185]
Zfp936	1.15	2.66	13	Unknown	BP
Zfp229	0.76	3.84	14	Unknown	BP
Zfp600	0.71	3.98	15	Unknown	BP
Zfp640	1.10	2.55	16	Differentiation	bone
Gm10324	1.09	2.55	17	Unknown	BP
Zscan10	1.04	2.65	18	Pluripotency	[186]
Utf1	2.03	1.30	19	Reprogramming	[5]
2610305D13Rik	1.02	2.45	20	Unknown	BP
Tfcp2l1	1.26	1.90	21	Pluripotency	[187]
Klf8	0.58	4.12	22	Differentiation	[188]
Epas1	0.70	3.18	23	Differentiation	red blood
Tbx3	1.09	2.03	24	Reprogramming	[5]
Tcf15	0.89	2.37	25	Differentiation	[189]

Table A.2: Top ESC Candidates. Part 2.

TF	Z Score	η (10^{-3})	Rank Z$*\eta$	Classification	Term
Tcfl5	0.82	2.56	26	Unknown	RT
Sall4	1.72	1.17	27	Reprogramming	[5]
Zfp553	0.87	2.22	28	Unknown	RT
Sox2	1.96	0.97	29	Reprogramming	[5]
Grhl3	0.61	2.75	30	Differentiation	ect dev
Zbtb10	0.75	2.22	31	Unknown	NRT
Mycn	1.90	0.85	32	Differentiation	lung
Sap30	0.93	1.72	33	Differentiation	muscle
Zbtb8a	0.83	1.88	34	Unknown	RT
Klf5	1.23	1.25	35	Differentiation	muscle
Sall1	1.30	1.18	36	Differentiation	neural
AA987161	0.60	2.36	37	Unknown	BP
Klf9	0.70	1.96	38	Differentiation	embryo
Myc	0.73	1.86	39	Reprogramming	[5]
Rarg	0.87	1.54	40	Differentiation	bone
Tead2	1.03	1.15	41	Differentiation	mesoderm
Dnmt3b	1.33	0.88	42	Pluripotency	GI
Nr5a2	0.67	1.75	43	Reprogramming	[5]
Nr1d2	0.74	1.53	44	Differentiation	muscle
Cbx7	1.14	0.99	45	Differentiation	CM
Bnip3	1.40	0.77	46	Differentiation	brown
Rbpms	1.63	0.64	47	Unknown	RT
Zfp7	0.91	1.15	48	Unknown	RT
Lin28a	0.78	1.31	49	Reprogramming	[5]
Zfp423	0.55	1.79	50	Differentiation	cell diff

Table A.3: Examining Yamanaka Factors in Detail

TF	(A) Exp.	(B) Diff. Exp.	(C) Specificity	(D) Predictivity	(E) Exp*Pred
Oct4	2	1	1.6%	70	1
Sox2	22	11	0.0%	201	29
Klf4	122	124	22.2%	28	9
Myc	213	1183	66.7%	106	39

A.6.2 Examining Yamanaka Factors in Detail

In Table A.3 we reexamine the Yamanaka transcription factors (TFs) in light of our model. When the Yamanaka results were first published, *Klf4* and *Myc* were counterintuitive factors [133]. *Myc* was quickly shown to enhance the efficiency of reprogramming but was dispensable [134]. *Klf4* remained a surprise, but this table demonstrates the power of predictivity by establishing the importance of *Klf4*. The columns (A),(B), and (C) are data about TFs available to Yamanaka, while (D) and (E) are data from our model. Unless otherwise stated, the numbers represent rank order (out of 1337) relative to the other TFs. To understand the importance of rank order, the original Yamanaka experiment used 24 TFs while most later studies test around 10 TFs at once. (A) Exp. is TF expression rank in embryonic stem cells (ESC). (B) Diff. Exp. is the differential expression rank between ESC and mouse embryonic fibroblasts (MEF), the starting cell fates in the Yamanaka protocol. (C) Specificity is the percentage of cell fates (out of our 63) which have expression at the same or higher level as the ESC. (D) Predictivity is the novel measure of TF importance generated by our model. (E) Exp*Pred is the rank of the product of expression and predictivity of highly expressed TFs and is an attempt to find a single quantity signifying reprogramming potential. The data available to Yamanaka illustrates that *Pou5f1* (*Oct4*) and *Sox2* were natural choices. *Myc* is an oncogene that enhances proliferation but was found to be non-essential for reprogramming [134], so we will ignore it. The power of predictivity is illustrated by examining *Klf4* which is not highly expressed (A), differentially expressed (B), or specific (C). However, it is very predictive of ESC (D) and is a top choice

when examining Exp*Pred (E). Note that *Klf4* illustrates that predictivity is not exactly the same as specificity. While *Klf4* is expressed in many cell fates, since predictivity takes into account correlations between cell fate expression patterns, predictivity can filter out the uncorrelated expression pattern and highlight the importance of *Klf4* for ESC.

B Cellular reprogramming dynamics follow a simple one-dimensional reaction coordinate

B.1 Data Analysis

Here we present details of the data analysis. Microarrays were taken from public datasets and come from a variety of different microarray platforms. In order to compare the different platforms, the following analysis was done. The raw microarray data was converted to an expression level as follows. Microarray probe-to-gene map was created with Bioconductor 3.0. All raw microarray files were initially processed by robust mean averaging (RMA) and genes with multiple microarray probes were averaged. Since we were interested in cellular identity, only transcription factors, transcription factor co-factors, or chromatin remodeling genes were kept (for short hand, referred to as transcription factors (TF) throughout the text) [190].

While the above analysis was done for both experimental data and simulations, from this point on the analysis differed between the two cases. For the experimental data analysis, we only used TFs that were common to all of the different microarray platforms, leaving $N = 994$ TFs. In order to make robust comparisons across platforms the RMA output was converted to a rank order. Next, we wanted to convert this rank order to the z-score of a log-normal distribution. We converted the rank to a percentile (for N genes, by dividing by N+1), and then this percentile into a normal z-score. For later mathematical convenience, we used a biased estimator (i.e. we normalized by N and not N-1) since then the Euclidean norm of each microarray gene expression was N. Therefore, for the data analysis each sample is described by a Gaussian distribution with a Euclidean norm of $N = 994$.

For the simulations, we followed similar steps to produce continuous TF expression levels for the cell type basis vector. However, in order to reduce the computational cost, we binarized the gene expression so that each TF is either on $(+1)$ or off (-1). We then

dropped all TFs that were always on or always off in every cell type, leaving $N = 1436$ TFs for the simulations.

B.2 Cellular Identity Landscape

Here we summarize our model for the cellular identity landscape [155]. The N transcription factors (TF) are labeled by Latin indices i and the p cell types are labeled by Greek indices μ. When analyzing experiments, we keep the $N = 994$ TFs common to all of the experimental datasets. Each sample is a Gaussian distribution with mean equal to 0 and Euclidean norm equal to N. This implies a standard deviation of $\frac{N}{N-1} \approx 1$. When performing simulations, we use the complete set of $N = 1436$ TFs and each TF is either on $(+1)$ or off (-1). A general network state is represented by a vector S_i of length N. A cell type μ is represented by the vector ξ_i^μ. The correlation (dot-product, overlap, or magnetization) between an arbitrary state and the cell type μ is given by

$$m^\mu = \frac{1}{N} \sum_{i=1}^{N} \xi_i^\mu S_i. \tag{B.1}$$

The correlation matrix between cell types is given by

$$A^{\mu\nu} = \frac{1}{N} \sum_{i=1}^{N} \xi_i^\mu \xi_i^\nu. \tag{B.2}$$

The projection onto each cell type is

$$a^\mu = \sum_{\nu=1}^{p} (A^{-1})^{\mu\nu} m^\nu. \tag{B.3}$$

We require all ξ_i^μ to be attractors in the landscape. This is ensured by constructing a correlation-based interaction network given by

$$J_{ij} = \frac{1}{N} \sum_{\mu=1}^{p} \sum_{\nu=1}^{p} \xi_i^\mu (A^{-1})^{\mu\nu} \xi_j^\nu \tag{B.4}$$

with $A^{\mu\nu}$ the correlation matrix between cell types. This interaction network produces stable basins of attraction and is written in terms of the cellular identity landscape as

$$H_{\text{basin}} = -\frac{1}{2} \sum_{ij=1}^{N} S_i J_{ij} S_j = -\frac{N}{2} \sum_{\mu=1}^{p} m^\mu a^\mu, \tag{B.5}$$

which is a Lyapunov function. There is also a culture term

$$H_{\text{culture}} = -\sum_{\mu=1}^{p} b^{\mu} a^{\mu}, \tag{B.6}$$

which stabilizes specific cell types. Normally only one cell type μ is stabilized, by choosing $b^{\mu} > 0$, and all the other coefficients $b^{\mu'}$ are zero.

The complete landscape is defined by an abstract energy H, which is composed of the two terms just discussed:

$$\begin{aligned} H &= H_{\text{basin}} + H_{\text{culture}} \\ &= -\frac{N}{2}\sum_{\mu=1}^{p} m^{\mu} a^{\mu} - \sum_{\mu=1}^{p} b^{\mu} a^{\mu} \end{aligned} \tag{B.7}$$

In addition, the model allows for certain transcription factors i to be locked, typically by choosing them to be "on", i.e. $S_i = +1$.

B.3 Dynamics and Simulations

To numerically study the dynamics of the model, we assume that the TFs can probabilistically switch states. To save computational effort, we also assume that the expression values are binarized, i.e. the only possible expression values are $+1$ and -1. Each TF is biased towards a state by its interactions with the network through its local field $h_i = -\frac{\partial H}{\partial S_i}$. The evolution is probabilistic and controlled by a global noise parameter β (i.e. inverse temperature $\beta = 1/T$). At each simulation update, u, one randomly chosen TF i is updated. The probability of the value S_i at update $u+1$ is related to the local field $h_i(u)$ at update u by

$$P[S_i, u+1] = \frac{e^{\beta h_i(u) S_i}}{e^{\beta h_i(u)} + e^{-\beta h_i(u)}}. \tag{B.8}$$

Additionally, as indicated above, in some of the simulations a subset of the transcription factors is locked at a certain value. Concretely, in many of the simulations we discuss, the OSKM TFs are fixed to have the value $+1$ ("on").

We performed Monte Carlo (MC) simulations of a system containing $N = 1436$ TFs using the update rule given by Eq. (B.8), with noise parameter $\beta = 1.62$ (i.e. $T \approx 0.617$).

When a culture term was introduced, it was to bias the system towards the ESC cell type, with $b^{\mu} = 0.03$ for $\mu = \text{ESC}$ and $b^{\mu'} = 0$ for all other cell types.

Most of the results reported in this paper correspond to simulations where the total number of steps was $t = 10^5$. For the case of simulations of MEF to ESC reprogramming, the OSKM transcription factors were locked "on" for the whole simulation, and the culture term was present from step $t = 5000$ until the end. In this case, 3000 trajectories were simulated, out of which 224 successfully reprogrammed, i.e. the reprogramming rate was 7.43%. For the simulations of B-cell to ESC reprogramming, the protocol was similar, and in this case 205 trajectories reprogrammed successfully out of a total of 3000, corresponding to a reprogramming rate of 6.83%.

In order to obtain additional details about the probability distributions of times associated with the reprogramming, which we show in Figure 3.3A and Figure B.3, we performed an additional set of simulations of MEF to ESC reprogramming, with the only change being that the total number of steps was 30 times larger, i.e. $t = 3 \times 10^6$ instead of $t = 10^5$. In this set of much longer simulations, 2937 trajectories out of 3000 successfully reprogrammed from MEF to ESC, which corresponds to a reprogramming rate of 97.90%.

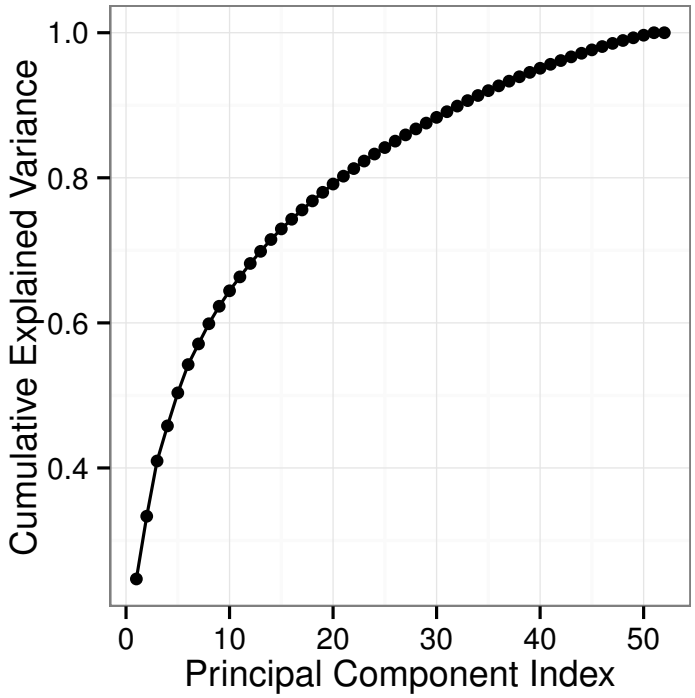

Figure B.1: **Principal components and explained variance.** This plot provides extended details of the principal component analysis (PCA) in Figure 3.1D. The cumulative fraction of explained variance vs principal component shows that in terms of PCA, the reprogramming dataset is high dimensional.

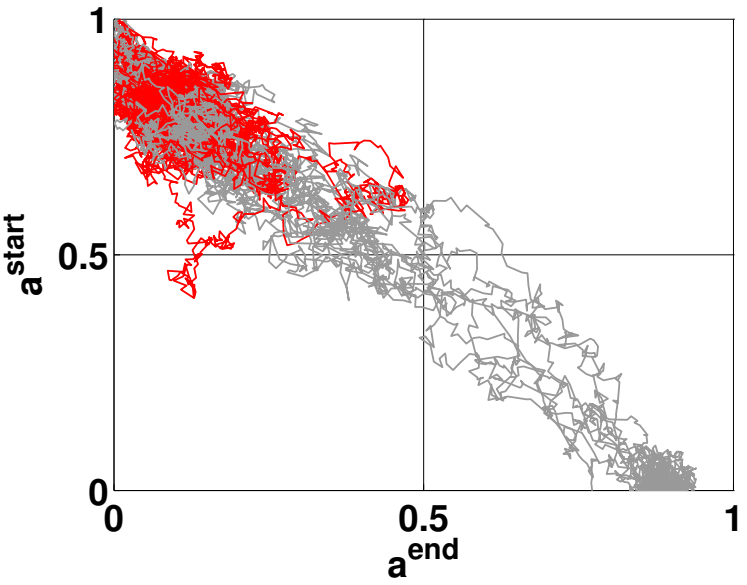

Figure B.2: **Fig 3.1E simulations.** Figure 3.1E simulation inset enlarged and with more trajectories.

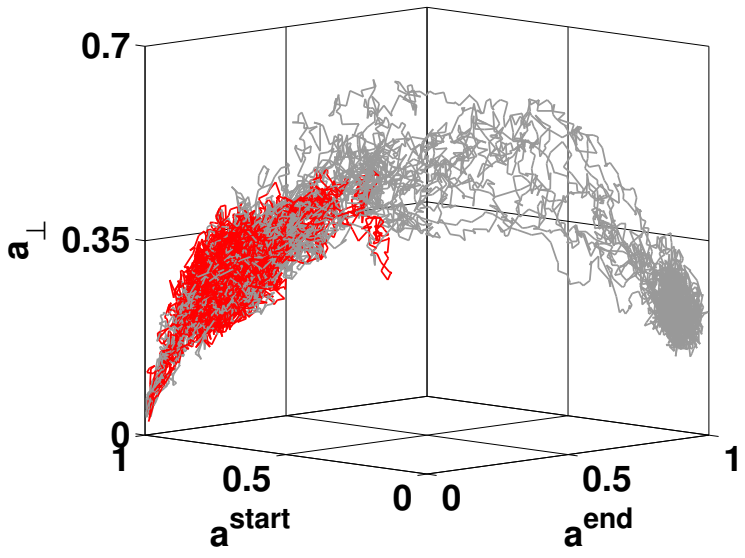

Figure B.3: **Fig 3.1F simulations.** Figure 3.1F simulation inset enlarged and with more trajectories.

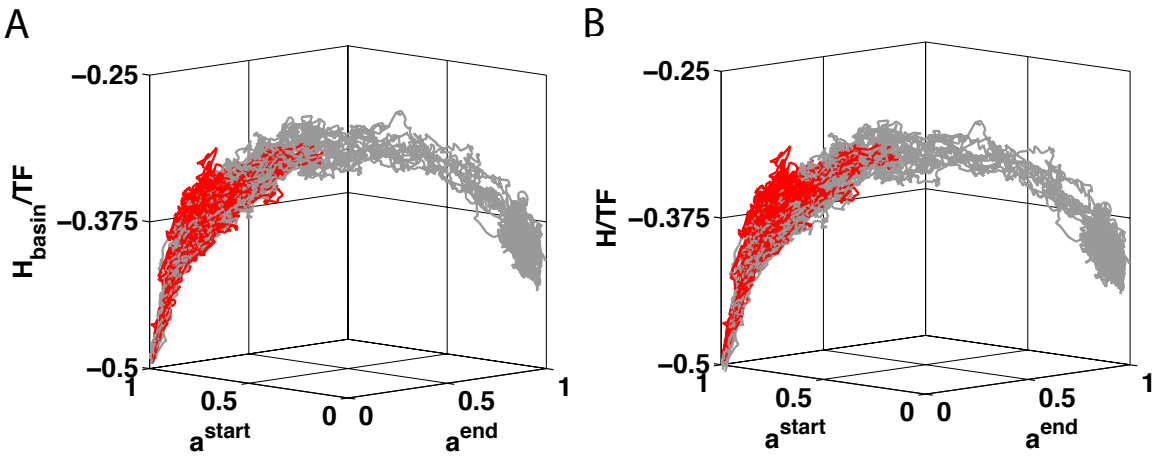

Figure B.4: **Fig 3.1G simulations.** A. Figure 3.1G simulation inset enlarged and with more trajectories. B. Figure 3.1G simulations including the small correction due to the culture term H_{culture}.

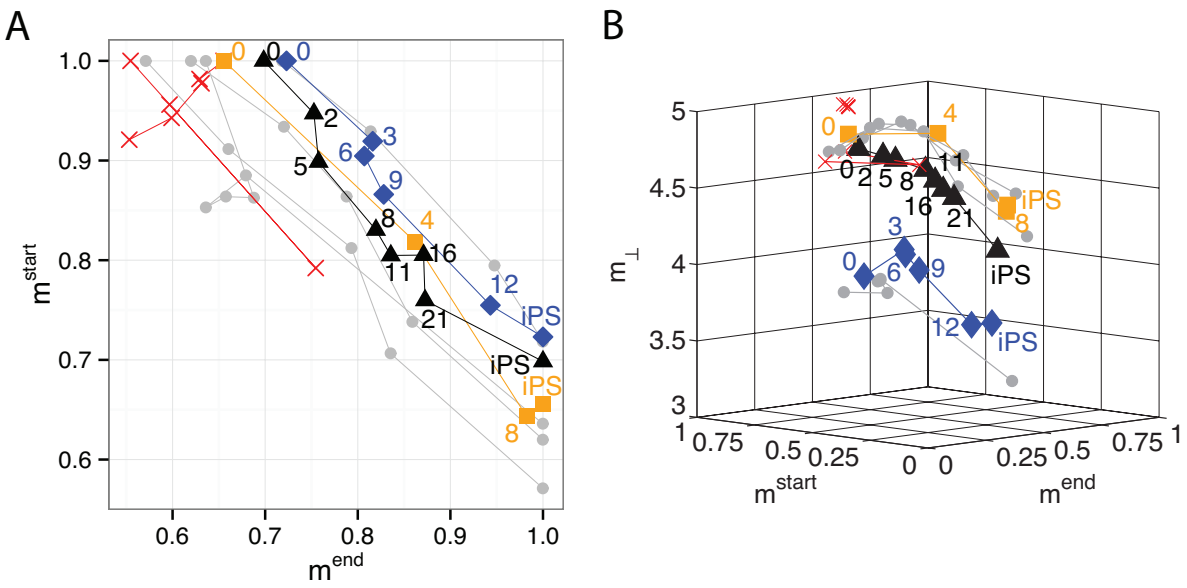

Figure B.5: **Alternative reaction coordinate and barrier.** A. This figure shows the same data presented in Fig 3.1E but instead of using projections (a), we have plotted dot products (m). B. This figure shows the same data presented in Fig 3.1F but instead of using projections (a), we have plotted dot products (m).

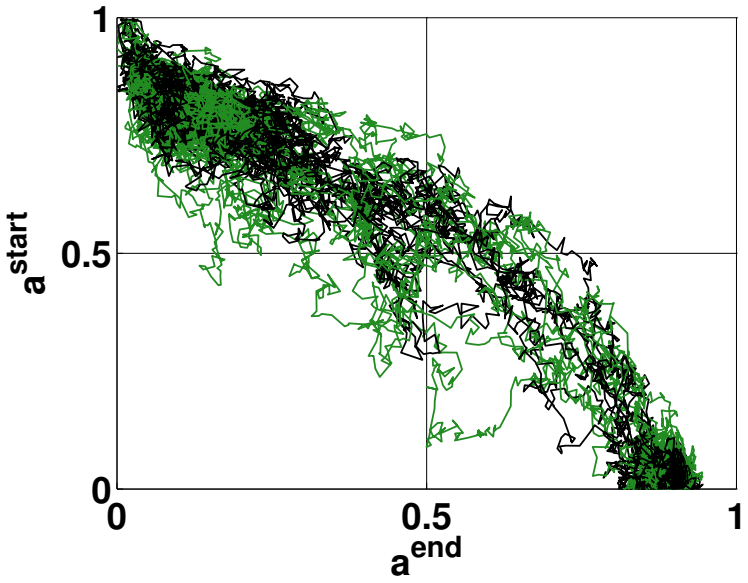

Figure B.6: **Fig 3.2A simulations.** Figure 3.2A simulation inset enlarged and with more trajectories.

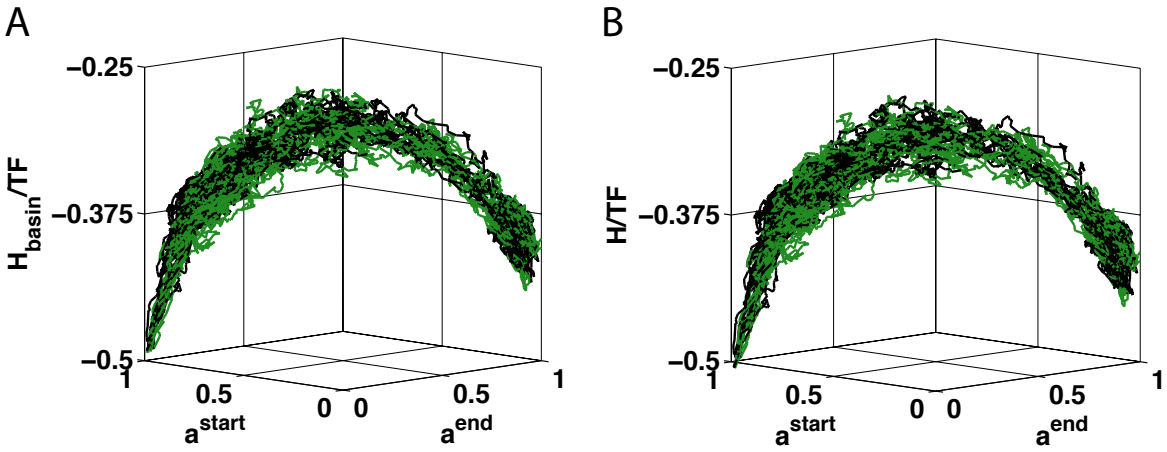

Figure B.7: **Fig 3.2B simulations.** A. Figure 3.2B simulation inset enlarged and with more trajectories. B. Figure 3.2B simulations including the small correction due to the culture term H_{culture}.

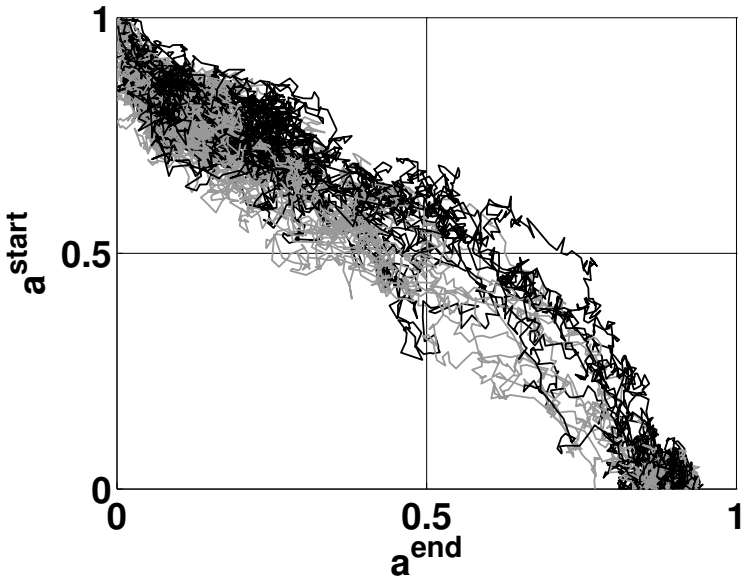

Figure B.8: **Fig 3.2C simulations.** Figure 3.2C simulation inset enlarged and with more trajectories.

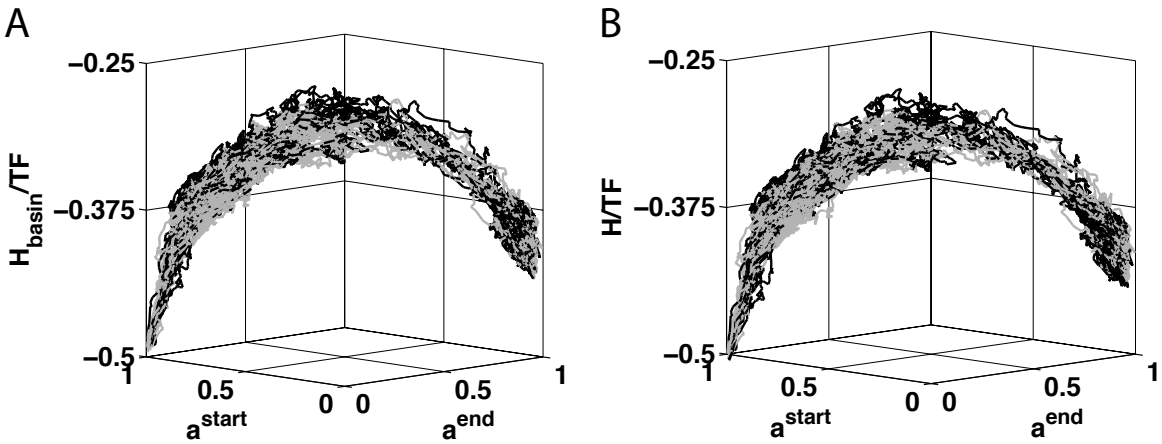

Figure B.9: **Fig 3.2D simulations.** A. Figure 3.2D simulation inset enlarged and with more trajectories. B. Figure 3.2D simulations including the small correction due to the culture term H_{culture}.

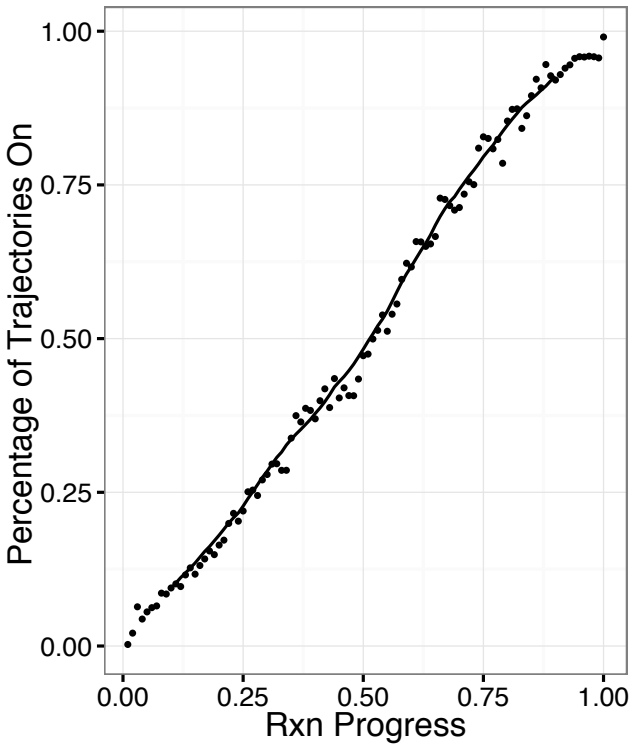

Figure B.10: **Percentage of trajectories in which a gene is on vs reaction coordinate.** Nanog is shown as an example where dots represent actual binned data, while the line is a 20 bin moving time average. This is just an example of the moving averages shown in Fig 3.3B.

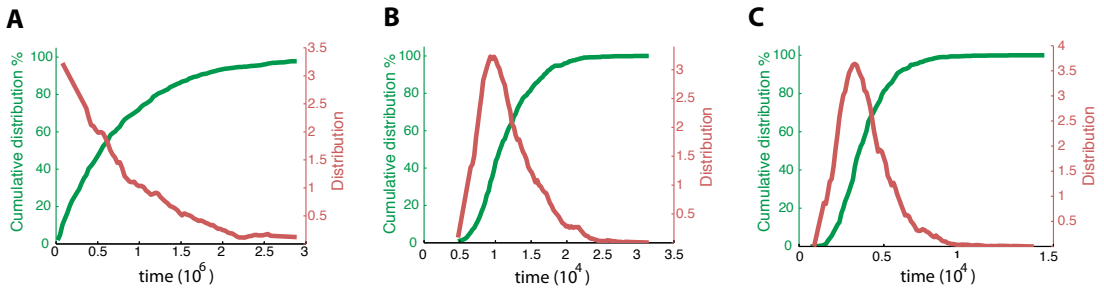

Figure B.11: **Additional Timing Details.** A. Cumulative distributions of timing show that the early ($a^{end} = 0$ to $a^{end} = 0.3$), B. middle ($a^{end} = 0.3$ to $a^{end} = 0.7$), and C. late ($a^{end} = 0.7$ to $a^{end} = 0.8$) stages of reprogramming are respectively a Poisson, a narrowly peaked, and a narrowly peaked distribution.

C Thermodynamics of statistical inference by cells

C.1 Notation

Here we provide more details of the results in the section 4. First, we outline our notation. The time dependent probability of state i is $p_i = p_i(t)$, while the steady state probability of state i is p_i^{ss}. The Laplace transformed probability of state i is $P_i(s)$. The rate to go from state i to state j is k_{ij}. The probability to transition from state i to state j is q_{ij}. The time it takes to transition from state i to j is τ_{ij}. The first passage time is given by $f(t)$ while the Laplace transformed first passage time is $F(s)$. The lifetime of state i is ρ_i.

C.2 Detailed Derivation of General Uncertainty

Here we derive formulas for the accuracy of statistical inference when the activated signaling states continuously produce signals. Following Berg and Purcell [107], we will measure the accuracy of a receptor by the "uncertainty" of the concentration estimate:

$$\text{uncertainty} := \frac{\langle (\delta c)^2 \rangle}{\bar{c}^2} \tag{C.1}$$

where $\bar{c}$ is the mean and $\langle (\delta c)^2 \rangle$ is the variance of the estimated concentration.

Let us consider the case where activated signaling states produce downstream signaling molecules at a rate α. We will define $\overline{\tau}_S$ as the mean lifetime of the signaling states and $\overline{\tau}_{NS}$ as the mean non-signaling time. Then, we know that the mean number of signaling molecules $\overline{u}$ produced after a time T is given by

$$\overline{u} = \alpha T \left(\frac{\overline{\tau}_S}{\overline{\tau}_S + \overline{\tau}_{NS}} \right) \equiv \alpha T \bar{\phi} \tag{C.2}$$

This follows by noting that $\bar{\phi}$ is just the fraction of time the receptor is in the signaling states. Notice that by definition, α and T are independent of the concentration c. The signaling time $\overline{\tau}_S$, can in principle depend on concentration, and for L signaling states is

given by

$$\overline{\tau}_S = \sum_{i=1}^{L} p_{0i}\tau_{i0} \tag{C.3}$$

where p_{0i} is the probability to transition from state 0 to state i, and τ_{i0} is the mean time to return from state i to state 0. Since we assume the receiving state is strongly biased (i.e. k_{01} is much larger than any other rate k_{0i} from non-signaling, 0, to signaling state i), then the derivative of the signaling time with respect to concentration is:

$$\frac{d\overline{\tau}_S}{dc} = -\sum_{i=2}^{L} \frac{k_{0i}}{k_{01}} \left(\tau_{10} + \tau_{i0}\right) \tag{C.4}$$

Since this is by assumption small, we will approximate $\overline{\tau}_S$ as independent of concentration, and thus all the concentration dependence comes from $\overline{\tau}_{NS}$. Thus, using the usual error-propagation formulas one has

$$\frac{\delta u}{\overline{u}} = -\frac{d\overline{\tau}_{NS}}{dc}\frac{1}{\overline{\tau}_S + \overline{\tau}_{NS}}\delta c \tag{C.5}$$

which gives the uncertainty for the concentration:

$$\frac{\langle(\delta c)^2\rangle}{\overline{c}^2} = \left(\overline{c}\frac{d\overline{\tau}_{NS}}{dc}\right)^{-2} \left(\overline{\tau}_{NS} + \overline{\tau}_S\right)^2 \frac{\langle(\delta u)^2\rangle}{\overline{u}^2} \tag{C.6}$$

The formula above reduces the problem to calculating the uncertainty in the number of signaling molecules produced in a time T. To calculate this, notice that $\overline{u}$ comes from on average $\overline{N} = T/(\overline{\tau}_S + \overline{\tau}_{NS})$ independent binding cycles (state 0 to state 1 transition). Thus, the variance in the fraction of time bound during a time T will just be $\overline{N}^{-1}$ times the variance in a single binding cycle. In particular, the coefficient of variation in a single cycle is given by

$$\frac{\delta\phi}{\overline{\phi}} = \frac{\overline{\tau}_{NS}}{\overline{\tau}_S + \overline{\tau}_{NS}}\left[\left(\frac{\delta\tau_S}{\overline{\tau}_S}\right) - \left(\frac{\delta\tau_{NS}}{\overline{\tau}_{NS}}\right)\right] \tag{C.7}$$

Noting that the signaling and non-signaling events are independent, we get

$$\frac{\langle(\delta u)^2\rangle}{\overline{u}^2} = \frac{1}{\overline{N}}\left(\frac{\overline{\tau}_{NS}}{\overline{\tau}_S + \overline{\tau}_{NS}}\right)^2 \left[\frac{\langle(\delta\tau_{NS})^2\rangle}{\overline{\tau}_{NS}^2} + \frac{\langle(\delta\tau_S)^2\rangle}{\overline{\tau}_S^2}\right] \tag{C.8}$$

Plugging this expressions into (C.6) gives

$$\frac{\langle(\delta c)^2\rangle}{\overline{c}^2} = \frac{1}{\overline{N}}\left(\overline{c}\frac{d\log\left(\overline{\tau}_{NS}\right)}{dc}\right)^{-2} \left[\frac{\langle(\delta\tau_{NS})^2\rangle}{\overline{\tau}_{NS}^2} + \frac{\langle(\delta\tau_S)^2\rangle}{\overline{\tau}_S^2}\right] \tag{C.9}$$

Therefore the complicated response of a receptor is reduced to its mean and variance of the time in both the signaling and non-signaling states. In this paper, we will examine the case where there is a single non-signaling state (0) and there are L signaling states arranged in a ring. In this case, the above expression simplifies to (leading order k_{0L}/k_{01}):

$$\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} = \frac{1}{\overline{\overline{N}}}\left[1 + \frac{\langle(\delta\tau_S)^2\rangle}{\overline{\tau}_S^2}\right] \tag{C.10}$$

For a two state process as considered by Mora and Wingreen [168], there is only the receiving state and one signaling state. These are just Poisson processes which each have an uncertainty of 1 and we recover the Berg and Purcell [107] limit

$$\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} = \frac{2}{\overline{\overline{N}}} \tag{C.11}$$

C.3 General First Passage Time

We need to calculate the first passage properties of the Markov chain, specifically the mean and variance of the first passage time. This can be calculated as follows [103, 191]. The master equation that we want to solve is $\frac{dp}{dt} = Kp(t)$.

First apply the Laplace transform

$$P_i(s) = \int_0^\infty p_i(t)e^{-st}dt \tag{C.12}$$

which leads to the master equation

$$(s - K)\,P(s) = p(t = 0) \tag{C.13}$$

with K the matrix of transitions for the full system but with the transition rates leaving the absorbing states set to zero.

The first passage time to return to state 0 is

$$f(t) = \frac{dp_0(t)}{dt} \tag{C.14}$$

$$F(s) = sP_0(s) \tag{C.15}$$

For our purposes, we only need the mean and variance of the first passage time. This is easily obtained by the uncentered moments

$$M^{(m)} = \int_0^\infty t^m f(t) = (-1)^m \left. \frac{d^m F(s)}{ds^m} \right|_{s=0} \tag{C.16}$$

where $m = 1$ is the mean and $m = 2$ is the uncentered second moment.

In general we know that τ_x, the spent in state x, is drawn from a mixture where it can switch to states $j = 1, 2, ...$ The variance of mixtures is $X = \sum_i w_i X_i$, where w_i are arbitrary weights and X_i are random variables drawn from distributions with mean μ_i and variance σ_i. Combining equations we get:

$$\text{Var(X)} = \sum_i w_i \left[(\mu_i - \mu)^2 + \sigma_i^2 \right] \tag{C.17}$$

with $\mu = \sum_i w_i \mu_i$.

We can get the time spent in state x, τ_x, by using the variance mixture formula combined with $\overline{\tau}_{ix}$ and $\text{Var}(\tau_{ix})$, respectively the mean and variance first passage time of starting in state i and ending in state x. This gives us

$$\overline{\tau}_x = \sum_i q_{xi} \overline{\tau}_{ix} \tag{C.18}$$

$$q_{xi} = \frac{k_{xi}}{\sum_j k_{xj}} = k_{xi} \rho_x \tag{C.19}$$

$$\rho_x = \left(\sum_j k_{xj} \right)^{-1} \tag{C.20}$$

$$\text{Var}(\tau_x) = \sum_i q_{xi} Var(\tau_{ix}) + \sum_i q_{xi} \left(\overline{\tau}_{ix} - \sum_k q_{xk} \overline{\tau}_{kx} \right)^2 \tag{C.21}$$

where q_{xi} is the probability of transitioning from state x to state i, k_{xi} is the rate to go from state x to state i, and ρ_x is the lifetime of state x.

In this paper, we have one non-signaling state and the other L states are signaling. Therefore, we will let state 0 be the absorbing state, and it can initially transition to state

1 and state L. The above equations then simplify to

$$\overline{\tau_0} = q_{01}\overline{\tau}_{10} + q_{0L}\overline{\tau}_{L0} \tag{C.22}$$

$$\mathrm{Var}(\tau_0) = q_{01}Var(\tau_{10}) + q_{0L}Var(\tau_{L0}) + 2q_{01}q_{0L}\left(\overline{\tau}_{10} - \overline{\tau}_{L0}\right)^2 \tag{C.23}$$

$$q_{0L} = 1 - q_{01} \tag{C.24}$$

C.4 First Passage Time: 2 Signaling States

Here we calculate the mean and variance of the first passage time to return to state 0 from either state 1 or 2. The master equation that we need to solve is $\frac{dp}{dt} = Kp(t)$. The matrix rates are:

$$K_{ij} = \begin{cases} k_{10} & \text{for } i = 0 \text{ and } j = 1 \\ k_{12} & \text{for } i = 2 \text{ and } j = 1 \\ k_{20} & \text{for } i = 0 \text{ and } j = 2 \\ k_{21} & \text{for } i = 1 \text{ and } j = 2 \\ -(k_{10} + k_{12}) & \text{for } i = 1 \text{ and } j = 1 \\ -(k_{20} + k_{21}) & \text{for } i = 2 \text{ and } j = 2 \\ 0 & \text{everywhere else} \end{cases} \tag{C.25}$$

While the initial conditions are set by the rates k_{01} and k_{02}, for the purposes of the first passage time calculation, the rates from 0 to 1 (k_{01}) and from 0 to 2 (k_{02}) are both set to zero, $k_{01} = k_{02} = 0$.

The Laplace transform for the initial condition of starting in state 1 is:

$$F(s) = sP_0(s) = k_{10}P_1 + k_{20}P_2 \tag{C.26}$$

with

$$P_1 = \left[\Gamma_1 - \frac{k_{12}k_{21}}{\Gamma_2}\right]^{-1} \tag{C.27}$$

$$P_2 = \frac{k_{12}}{\Gamma_2}P_1 \tag{C.28}$$

$$\Gamma_i = s + \rho_i^{-1} \tag{C.29}$$

We can obtain mean and variance from

$$\overline{\tau} = -\left.\frac{dF}{ds}\right|_{s=0} \tag{C.30}$$

$$\mathrm{Var}(\tau) = \left.\frac{d^2F}{ds^2}\right|_{s=0} - \overline{\tau}^2 \tag{C.31}$$

The mean and variance of the first passage time from starting in either state 1 or state 2 is:

$$\overline{\tau}_{10} = \rho_1 \frac{1 + k_{12}\rho_2}{1 - k_{12}k_{21}\rho_1\rho_2} = \frac{k_{12} + k_{20} + k_{21}}{\xi} \tag{C.32}$$

$$\overline{\tau}_{20} = \rho_2 \frac{1 + k_{21}\rho_1}{1 - k_{12}k_{21}\rho_1\rho_2} = \frac{k_{10} + k_{12} + k_{21}}{\xi} \tag{C.33}$$

$$\mathrm{Var}(\tau_{10}) = \tau_{10}^2 \left[1 + 2\rho_2^2 \frac{k_{12}\left(k_{10} - k_{20}\right)}{(1 + k_{12}\rho_2)^2}\right] = \tau_{10}^2 + 2\frac{k_{12}(k_{10} - k_{20})}{\xi^2} \tag{C.34}$$

$$\mathrm{Var}(\tau_{20}) = \tau_{20}^2 \left[1 + 2\rho_1^2 \frac{k_{21}\left(k_{20} - k_{10}\right)}{(1 + k_{21}\rho_1)^2}\right] = \tau_{20}^2 - 2\frac{k_{21}(k_{10} - k_{20})}{\xi^2} \tag{C.35}$$

$$\xi = k_{10}k_{20} + k_{10}k_{21} + k_{12}k_{20} \tag{C.36}$$

where the second equality holds as long as $\xi \neq 0$.

C.5 First Passage Time: L Signaling States

C.5.1 Derivation

Here we calculate the mean and variance of the first passage time in a $L+1$ state chain. The master equation that we need to solve is $\frac{dp}{dt} = Kp(t)$. The matrix is indexed from 0 to L and the rates are:

$$K_{ij} = \begin{cases} k_{10} & \text{for } i = 0 \text{ and } j = 1 \\ k_{L0} & \text{for } i = 0 \text{ and } j = L \\ f & \text{for } i = j + 1 \text{ and } 1 < j < L \\ b & \text{for } i = j - 1 \text{ and } 1 < j < L \\ -(f + k_{10}) & \text{for } i = 1 \text{ and } j = 1 \\ -(f + b) & \text{for } i = j \text{ and } 1 < j < L \\ -(k_{L0} + b) & \text{for } i = L \text{ and } j = L \\ 0 & \text{everywhere else} \end{cases} \tag{C.37}$$

While the initial conditions are set by the rates k_{01} and k_{0L}, for the purposes of the first passage time calculation, the rates from 0 to 1 (k_{01}) and from 0 to L (k_{0L}) are both set to zero, $k_{01} = k_{0L} = 0$.

For later convenience we define the following ratio of rates:

$$\theta = \frac{f}{b} \tag{C.38}$$

$$\alpha = \frac{k_{10}}{b} \tag{C.39}$$

$$\omega = \frac{k_{L0}}{f} \tag{C.40}$$

We can use a transfer matrix to find a general solution (for non-degenerate eigenvalues, i.e. $\theta \neq 1$) to the state probability as

$$P_i(s) = C_+\lambda_+^{i-1} + C_-\lambda_-^{i-1} \tag{C.41}$$

Solving for the the expressions $1 < i < L$ leads to

$$\lambda_\pm = \frac{1}{2b}\left(s + f + b \pm \sqrt{(s+f+b)^2 - 4fb}\right) \tag{C.42}$$

$$= \frac{1}{2}\left(\sigma \pm \sqrt{\sigma^2 - 4\theta}\right) = \frac{1}{2}\left(\sigma \pm \psi\right) \tag{C.43}$$

$$\sigma = \frac{s}{b} + \theta + 1 \tag{C.44}$$

$$\psi = \sqrt{\sigma^2 - 4\theta} \tag{C.45}$$

With the initial condition of starting in P_1, the boundary equations for P_1 and P_L are:

$$(\sigma + \alpha - 1)\left(C_+ + C_-\right) = 1/b + \left(C_+\lambda_+ + C_-\lambda_-\right) \tag{C.46}$$

$$(\sigma + (\omega - 1)\theta)\left(C_+\lambda_+^{L-1} + C_-\lambda_-^{L-1}\right) = \theta\left(C_+\lambda_+^{L-2} + C_-\lambda_-^{L-2}\right) \tag{C.47}$$

Solving these equations gives

$$C_- = -C_+\Lambda^L M \tag{C.48}$$

$$C_+ = \frac{1}{b\left[\lambda_- + \alpha - 1 - (\lambda_+ + \alpha - 1)\Lambda^L M\right]} \tag{C.49}$$

$$\Lambda = \frac{\lambda_+}{\lambda_-} \tag{C.50}$$

$$M = \frac{1 + (\omega - 1)\lambda_-}{1 + (\omega - 1)\lambda_+} \tag{C.51}$$

And then the probabilities are

$$P_1(s) = C_+\left(1-\Lambda^L M\right) \tag{C.52}$$

$$P_L(s) = C_+\lambda_+^{L-1}\left(1-\Lambda M\right) \tag{C.53}$$

The full Laplace transform F is:

$$F(s) = \frac{\alpha(1-\Lambda^L M)+\omega\theta\lambda_+^{L-1}(1-\Lambda M)}{\lambda_- +\alpha-1-(\lambda_+ +\alpha-1)\Lambda^L M} \tag{C.54}$$

C.5.2 Results

To get the mean and variance of the first passage time, we need

$$\overline{\tau}_{10} = -\left.\frac{dF}{ds}\right|_{s=0} \tag{C.55}$$

$$\mathrm{Var}(\tau_{10}) = \left.\frac{d^2F}{ds^2}\right|_{s=0} - \overline{\tau}^2 \tag{C.56}$$

The mean return time to state 0 when starting in state 1 is:

$$\overline{\tau}_{10} = \frac{\overline{\tau}_{10,num}}{\overline{\tau}_{10,den}} \tag{C.57}$$

$$\overline{\tau}_{10,num} = (\omega L-\omega+1)\theta^{L+1}-(\omega L+1)\theta^L+(\omega-1)\theta+1 \tag{C.58}$$

$$\overline{\tau}_{10,den} = b\left[\theta-1\right]\left[\omega\theta^{L+1}+\omega(\alpha-1)\theta^L+\alpha(1-\omega)\theta-\alpha\right] \tag{C.59}$$

The variance of the return time to state 0 when starting in state 1 is:

$$\begin{aligned}
\mathrm{Var}(\tau_{10}) &= \frac{\mathrm{Var}(\tau_{10})_{num}}{\mathrm{Var}(\tau_{10})_{den}} && \text{(C.60)}\\
\mathrm{Var}(\tau_{10})_{num} &= \theta^{2L+3}\left[\omega^2(L-1)+1\right] && \text{(C.61)}\\
&+ \theta^{2L+2}\left[\omega^2\left(L^2\alpha - L(3\alpha+1)+2\alpha-3\right)+2\omega((L-2)\alpha+1)+2\alpha-3\right]\\
&- \theta^{2L+1}\left[\omega^2\left(2L^2\alpha-4L\alpha+L+4\alpha-4\right)+\omega((4L-6)\alpha+4)+4\alpha-3\right]\\
&+ \theta^{2L}\left[\omega(\omega L+2)(L\alpha-\alpha+1)+2\alpha-1\right]\\
&+ \theta^{L+3}(\omega-1)\left[2(\omega-1)\alpha+3\omega L^2\alpha+L(\omega(4-5\alpha)+4\alpha)+2\right]\\
&+ \theta^{L+2}\left[-2\omega^2\left(3L^2\alpha+L(4-6\alpha)+\alpha-2\right)\right]\\
&+ \theta^{L+2}\left[\omega\left(9L^2\alpha+L(12-23\alpha)+8\alpha-6\right)+6(2L-1)\alpha+6\right]\\
&+ \theta^{L+1}\left[\omega\left(-9L^2\alpha+L(19\alpha-12)-6\alpha+6\right)\right]\\
&+ \theta^{L+1}\left[\omega^2(L-1)((3L-4)\alpha+4)+6(-2L\alpha+\alpha-1)\right]\\
&+ \theta^{L}\left[\alpha\left(3L^2\omega-5L\omega+4L+2\omega-2\right)+(4L-2)\omega+2\right]\\
&- \theta^3(\omega-1)^2(2\alpha-1)\\
&- \theta^2(\omega-1)(\omega+4\alpha-3)\\
&+ \theta(-2\omega-2\alpha+3)\\
&- 1\\
\mathrm{Var}(\tau_{10})_{den} &= b^2\left[\theta-1\right]^3\left[\omega\theta^{L+1}+\omega(\alpha-1)\theta^L+\alpha(1-\omega)\theta-\alpha\right]^2 && \text{(C.62)}
\end{aligned}$$

While the results here are for initial condition of being in state 1, one can easily find the results for the initial condition of state L if one makes the following substitutions $\theta \Leftrightarrow 1/\theta$, $b \Leftrightarrow f$, and $\alpha \Leftrightarrow \omega$.

C.6 Steady State Probabilities

In general, we are considering a Markov chain with $L+1$ nodes (labeled 0 to L). We have the master equation

$$\frac{dP(t)}{dt} = KP(t) \qquad \text{(C.63)}$$

with K the matrix of transition rates. The rates are labeled as k_{ij} where i is the initial state and j is the final state. For later convenience, define the lifetime of a state as

$$\rho_i = \left(\sum_{j \neq i} k_{ij} \right)^{-1} \tag{C.64}$$

The steady state distributions are easily obtained by solving $Kp^{ss} = 0$. The solution can be written in a compact form [192] as

$$P_i^{ss} = \frac{z_i}{Z} \tag{C.65}$$

$$Z = \sum_i z_i \tag{C.66}$$

and z_i is the matrix minor of K at (i,i) i.e. the determinant of K with the ith row and column removed.

For the two signaling state system we have that

$$p_0^{ss} = \frac{\rho_1^{-1}\rho_2^{-1} - k_{12}k_{21}}{Z} = \frac{k_{10}k_{20} + k_{10}k_{21} + k_{12}k_{20}}{Z} \tag{C.67}$$

$$p_1^{ss} = \frac{\rho_0^{-1}\rho_2^{-1} - k_{02}k_{20}}{Z} = \frac{k_{01}k_{20} + k_{01}k_{21} + k_{02}k_{21}}{Z} \tag{C.68}$$

$$p_2^{ss} = \frac{\rho_0^{-1}\rho_1^{-1} - k_{01}k_{10}}{Z} = \frac{k_{01}k_{12} + k_{02}k_{10} + k_{02}k_{12}}{Z} \tag{C.69}$$

$$Z = \sum_{i \neq j} \left(\rho_i^{-1}\rho_j^{-1} - k_{ij}k_{ji} \right) \tag{C.70}$$

For the L signaling state with the simplified rates, we will just present the result for state 0:

$$p_0^{ss} = \frac{p_{0,num}^{ss}}{p_{0,den}^{ss}} \tag{C.71}$$

$$p_{0,num}^{ss} = b(\theta - 1)\left(\omega\theta^{L+1} + \omega(\alpha - 1)\theta^L + \alpha(1 - \omega)\theta - \alpha \right) \tag{C.72}$$

$$p_{0,den}^{ss} = -\alpha\epsilon + \alpha b + \alpha L\epsilon + \epsilon + 1 \tag{C.73}$$

$$+ \quad \theta\left(\alpha b\omega - 2\alpha b - \alpha L\epsilon + \omega - \epsilon - 1 \right) \tag{C.74}$$

$$+ \quad \alpha b\theta^2 (1 - \omega) \tag{C.75}$$

$$+ \quad \theta^L (b\omega + \alpha\epsilon - L\omega - 1 - \epsilon - \alpha b\omega) \tag{C.76}$$

$$+ \quad \theta^{1+L} (\alpha b\omega - 2b\omega + L\omega - \omega + 1 + \epsilon) \tag{C.77}$$

$$+ \quad b\omega\theta^{L+2} \tag{C.78}$$

The rates from 0 to 1 is $k_{01} = 1$, from 1 to 0 is k_{10} (with $\alpha = k_{10}/b$), from 0 to L is $k_{0L} = \epsilon \ll 1$, and from L to 0 is k_{L0} (with $\omega = k_{L0}/f$). All other forward rates are f and backward rates are b and the ratio of rates is $\theta = f/b$.

C.7 Average Sampling Rate: $\overline{n}$

The average sampling rate is

$$\overline{n} = \frac{N}{T} = k_{01} p_0^{ss} \tag{C.79}$$

where N is the number of samples (i.e. number of binding events), T is the total integration time, k_{01} is the rate from state 0 to state 1, and p_0^{ss} is the steady state probability of being in state 0.

Since we are assuming that $k_{01} = 1$ and $k_{L0} = \epsilon \ll 1$, we have the mean signaling time becomes $\overline{\tau}_S \approx \tau_{10}$. With these rates we have

$$\overline{n} \approx (1 + \overline{\tau}_S)^{-1} \tag{C.80}$$

C.8 Entropy Production: e_p

For a general Markov process with states labeled by i, steady state probabilities p_i^{ss}, and transition rate k_{ij} from state i to state j, the non-equilibrium steady state (NESS) entropy production [109, 172] is given by

$$e_p = \sum_{i=0}^{L} \sum_{j \neq i}^{L} p_i^{ss} k_{ij} \ln \frac{k_{ij}}{k_{ji}} \tag{C.81}$$

where the summation is over both i and j. Alternatively, the entropy production can be written as a sum over the flux between each connected node as

$$e_p = \sum_{i=0}^{L} \sum_{j>i}^{L} (p_i^{ss} k_{ij} - p_i^{ss} k_{ij}) \ln \frac{k_{ij}}{k_{ji}} \tag{C.82}$$

where now we have an unrestricted sum over i but a restricted sum over j.

Since we are modeling our receptor as a ring, the entropy production simplifies to

$$e_p = (p_0^{ss} k_{01} - p_1^{ss} k_{10}) \ln \frac{k_{01} k_{12} \ldots k_{L0}}{k_{0L} k_{10} \ldots k_{L,L-1}} = J \ln \gamma \tag{C.83}$$

where the flux $J = p_0^{ss} k_{01} - p_1^{ss} k_{10}$ between each neighboring state is equal and the $\ln \gamma$ is the free energy difference of a cycle.

For 2 signaling states, the entropy production per sampling rate is given by:

$$\frac{e_p}{\overline{n}} = \left[1 + \frac{k_{10}}{k_{12}} + \frac{k_{10}k_{21}}{k_{12}k_{20}}\right]^{-1} \frac{\gamma - 1}{\gamma} \ln \gamma \tag{C.84}$$

$$\gamma = \frac{k_{01}k_{12}k_{20}}{k_{10}k_{21}k_{02}} \tag{C.85}$$

For the L signaling states arranged in a ring, the entropy production per sampling rate is given by:

$$\frac{e_p}{\overline{n}} = \left[1 + \frac{\alpha}{\omega}\theta^{-L} + \alpha\theta^{-1}\frac{1-\theta^{1-L}}{1-\theta^{-1}}\right]^{-1} \frac{\gamma - 1}{\gamma} \ln \gamma \tag{C.86}$$

$$\gamma = \frac{k_{01}\omega}{k_{0L}\alpha}\theta^{L} \tag{C.87}$$

where $\omega = k_{L0}/f$, $\alpha = k_{10}/b$, $\theta = f/b$, f is all the forward rates (except k_{01} and k_{L0}), and b is all the backward rates (except k_{10} and k_{0L}).

C.9 Ansatz for 2 Signaling State Receptor

Here are the details of the ansatz for the minimum uncertainty for the 2 signaling state system.

The rates are as follows:

- $k_{01} = 1$
- $k_{10} = \frac{k}{2}(1 - x)$
- $k_{12} = kx$
- $k_{21} = k\delta$
- $k_{20} = k$
- $k_{02} = \epsilon$

where $\epsilon \ll 1$ (and in this paper $\epsilon = 10^{-3}$), $0 < x < 1$, $\delta \ll 1$ (and in this paper $\delta = 0.04$), and k is varied to fix the mean sampling rate $\overline{n}$.

For the ansatz, the mean, coefficient of variation, and entropy production simplifies to

$$\overline{\tau}_S \approx \frac{2}{k} \tag{C.88}$$

$$\frac{\langle(\delta\tau_S)^2\rangle}{\overline{\tau}_S^2} \approx 1 - \frac{x}{1+x} \tag{C.89}$$

$$\frac{e_p}{\overline{n}} \approx \left(1 + \frac{1-x}{2x}\right)^{-1} \frac{\gamma - 1}{\gamma} \ln\gamma \tag{C.90}$$

$$\gamma = \frac{2x}{\epsilon\delta(1-x)} \tag{C.91}$$

C.10 Simulated Annealing

Simulated annealing is a meta-heuristic algorithm for global optimization in which one uses the Metropolis algorithm to perform a random walk in parameter space while periodically lowering the temperature. We used a simulated annealing algorithm to search for the parameters of a model describing a receptor with 2 signaling states that minimizes a cost function given by

$$\text{cost} = \frac{\langle(\delta c)^2\rangle}{\bar{c}^2} + \lambda_{e_p}(\ln e_p - \ln \hat{e}_p)^2 - (\lambda_n \hat{n} - 1)\ln\overline{n} - (\lambda_n(1-\hat{n}) - 1)\ln(1-\overline{n}) \tag{C.92}$$

That is, we minimize the uncertainty of the resulting estimator ($\langle(\delta c)^2\rangle/\bar{c}^2$) subject to soft constraints on the energy production (e_p) and sampling rate ($\overline{n}$), which are constrained to $\hat{e}_p$ and $\hat{n}$, respectively. Here, λ_{e_p} and λ_n implement the constraints. We chose $\lambda_{e_p} = 20$ and $\lambda_n = 20/\max\{\hat{n}, 1-\hat{n}\}$.

Let Ω_1 denote a set of parameters describing a receptor with 2 signaling states (i.e. all of the various rate constants). A new set of trial parameters Ω_2 was generated in the following way: for each $k \in \Omega_1$ set the corresponding $k' \in \Omega_2$ to $\ln k' = \ln k + \eta$ where η is a random variable with from a Normal distribution centered at zero. The width of the Normal distribution was chosen adaptively so that approximately 25% of the steps were accepted. Making the random perturbations to the logarithm of the rate constants ensures that they are always positive. The trial move was accepted according to the Metropolis

criterion with probability $\min[1, \exp((\text{cost}(\Omega_1) - \text{cost}(\Omega_2))/T)]$. The temperature T was initialized to $T = 10$ and adjusted by $T \leftarrow 0.95T$ every 2000 steps. The best solution obtained during the chain was stored in Ω_B, and the chain was re-initialized from $\Omega_1 = \Omega_B$ every 2000 steps to prevent the chain from getting stuck in a poor local minimum. This simulated annealing algorithm was run until convergence of $\langle(\delta c)^2\rangle/\bar{c}^2$, e_p and n.

C.11 Scaling with Temperature

In the section 4, we worked in the units of $k_B T = 1$. However, here we examine the general temperature dependence. Experimentally, it is known that rates of biochemical reactions doubles for every 10 °C [193, 194]. Therefore, a general rate k at a temperature T (measured in degrees Celsius) is related to initial rate k_0 and initial temperature T_0 by:

$$k = k_0 2^{\frac{T-T_0}{10}} \tag{C.93}$$

Now we need to determine the general scaling of various entities in this paper, which is summarized below in terms of a general rate k:

- Mean signaling time, $\overline{\tau}_S \sim k^{-1}$
- Variance in signaling time, $\langle(\delta\tau_S)^2\rangle \sim k^{-2}$
- Coefficient of variation of signaling time, $\frac{\langle(\delta\tau_S)^2\rangle}{\overline{\tau}_S^2} \sim 1$
- Sampling rate, $\overline{n} \sim k$
- Uncertainity, $\frac{\langle(\delta c)^2\rangle}{\bar{c}^2} \sim k^{-1}$
- Entropy production, $e_p \sim k$

While increasing temperature increases both the mean and variance of the signaling time, since the estimator ($\mathcal{E} = 1 + \frac{\langle(\delta\tau_S)^2\rangle}{\overline{\tau}_S^2}$) only depends on the coefficient of variation of signaling time, the estimator is independent of temperature. The sampling rate $\overline{n}$ does increase with increasing temperature, and therefore increasing temperature decreases the uncertainty. However, this decrease in uncertainty costs energy. While the free energy per

cycle ($\ln \gamma$) remains constant, the probability flux (J) is proportional to a rate, and since the entropy production is given by $e_p = J \ln \gamma$, we see that that decrease in uncertainty is directly related to the increase in entropy production.

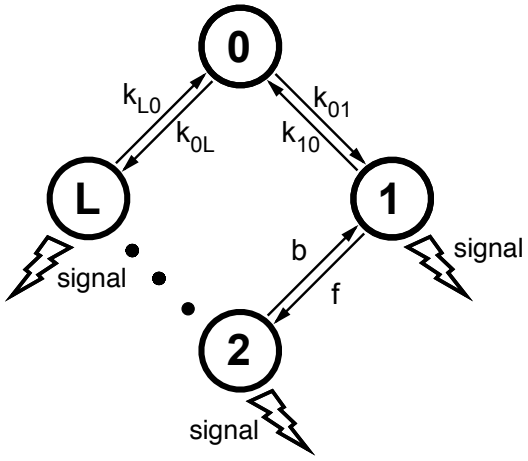

Figure C.1: **Simplified rate structure considered for L signaling states first passage time calculation.** The rates k_{01}, k_{10}, k_{L0}, k_{0L} are unconstrained, while the remaining forward rates are equal, $f = k_{12} = k_{23} = \ldots = k_{L-1,L}$ and the remaining backward rates are equal, $b = k_{21} = k_{32} = \ldots = k_{L,L-1}$.

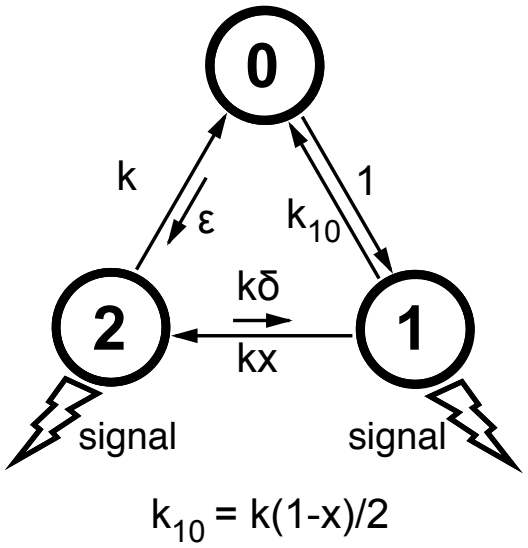

Figure C.2: **Rate structure for ansatz of minimum uncertainty for the $L = 2$ signaling state system.** The rates are as follows: $k_{01} = 1$, $k_{10} = \frac{k}{2}(1-x)$, $k_{12} = kx$, $k_{21} = k\delta$, $k_{20} = k$, and $k_{02} = \epsilon$. The mean signaling time is set by k. The other rates are $\epsilon, \delta \ll 1$ and $0 < x < 1$.

Bibliography

[1] Francis S Collins, Michael Morgan, and Aristides Patrinos. The human genome project: lessons from large-scale biology. *Science*, 300(5617):286–290

[2] Vivien Marx. Biology: The big challenges of big data. *Nature*, 498(7453):255–260

[3] Hiroaki Kitano. Systems biology: a brief overview. *Science*, 295(5560):1662–1664

[4] Shinya Yamanaka. Induced pluripotent stem cells: Past, present, and future. *Cell stem cell*, 10(6):678–684, 06 2012.

[5] Federico González, Stéphanie Boué, and Juan Carlos Izpisúa Belmonte. Methods for making induced pluripotent stem cells: reprogramming àla carte. *Nature Reviews Genetics*, 12(4):231–242, 04 2011.

[6] Thomas Vierbuchen, Austin Ostermeier, Zhiping P. Pang, Yuko Kokubu, Thomas C. Südhof, and Marius Wernig. Direct conversion of fibroblasts to functional neurons by defined factors. *Nature*, 463(7284):1035–1041, 02 2010.

[7] Masaki Ieda, Ji-Dong Fu, Paul Delgado-Olguin, Vasanth Vedantham, Yohei Hayashi, Benoit G. Bruneau, and Deepak Srivastava. Direct reprogramming of fibroblasts into functional cardiomyocytes by defined factors. *Cell*, 142(3):375–386, 08 2010.

[8] Sayaka Sekiya and Atsushi Suzuki. Direct conversion of mouse fibroblasts to hepatocyte-like cells by defined factors. *Nature*, 475(7356):390–393, 07 2011.

[9] Pengyu Huang, Zhiying He, Shuyi Ji, Huawang Sun, Dao Xiang, Changcheng Liu, Yiping Hu, Xin Wang, and Lijian Hui. Induction of functional hepatocyte-like cells from mouse fibroblasts by defined factors. *Nature*, 475(7356):386–389, 07 2011.

[10] Ernesto Lujan, Soham Chanda, Henrik Ahlenius, Thomas C. Südhof, and Marius Wernig. Direct conversion of mouse fibroblasts to self-renewing, tripotent neural

precursor cells. *Proceedings of the National Academy of Sciences*, 109(7):2527–2532, 2012.

[11] Francesco Antonica, Dominika Figini Kasprzyk, Robert Opitz, Michelina Iacovino, Xiao-Hui Liao, Alexandra Mihaela Dumitrescu, Samuel Refetoff, Kathelijne Peremans, Mario Manto, Michael Kyba, and Sabine Costagliola. Generation of functional thyroid from embryonic stem cells. *Nature*, 491(7422):66–71, 11 2012.

[12] Kazutoshi Takahashi and Shinya Yamanaka. Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors. *Cell*, 126(4):663–676, 8 2006.

[13] Laurie A. Boyer, Tong Ihn Lee, Megan F. Cole, Sarah E. Johnstone, Stuart S. Levine, Jacob P. Zucker, Matthew G. Guenther, Roshan M. Kumar, Heather L. Murray, Richard G. Jenner, David K. Gifford, Douglas A. Melton, Rudolf Jaenisch, and Richard A. Young. Core transcriptional regulatory circuitry in human embryonic stem cells. *Cell*, 122(6):947–956, 9 2005.

[14] K. Kim, A. Doi, B. Wen, K. Ng, R. Zhao, P. Cahan, J. Kim, M. J. Aryee, H. Ji, L. I. R. Ehrlich, A. Yabuuchi, A. Takeuchi, K. C. Cunniff, H. Hongguang, S. Mckinney-Freeman, O. Naveiras, T. J. Yoon, R. A. Irizarry, N. Jung, J. Seita, J. Hanna, P. Murakami, R. Jaenisch, R. Weissleder, S. H. Orkin, I. L. Weissman, A. P. Feinberg, and G. Q. Daley. Epigenetic memory in induced pluripotent stem cells. *Nature*, 467(7313):285–290, 09 2010.

[15] Richard A. Young. Control of the embryonic stem cell state. *Cell*, 144(6):940–954, 03 2011.

[16] Sui Huang. Reprogramming cell fates: reconciling rarity with robustness. *BioEssays*, 31(5):546–560, 2009.

[17] Jacob H. Hanna, Krishanu Saha, and Rudolf Jaenisch. Pluripotency and cellular reprogramming: Facts, hypotheses, unresolved issues. *Cell*, 143(4):508–525, 11 2010.

[18] Yusuke Hirabayashi and Yukiko Gotoh. Epigenetic control of neural precursor cell fate during development. *Nature Reviews Neuroscience*, 11(6):377–388, 06 2010.

[19] Fuchou Tang, Catalin Barbacioru, Siqin Bao, Caroline Lee, Ellen Nordman, Xiaohui Wang, Kaiqin Lao, and M. Azim Surani. Tracing the derivation of embryonic stem cells from the inner cell mass by single-cell rna-seq analysis. *Cell Stem Cell*, 6(5):468–478, 5 2010.

[20] Junhyong Kim and James Eberwine. Rna: state memory and mediator of cellular phenotype. *Trends in cell biology*, 20(6):311–318, 06 2010.

[21] Jai-Yoon Sul, Chia-wen K. Wu, Fanyi Zeng, Jeanine Jochems, Miler T. Lee, Tae Kyung Kim, Tiina Peritz, Peter Buckley, David J. Cappelleri, Margaret Maronski, Minsun Kim, Vijay Kumar, David Meaney, Junhyong Kim, and James Eberwine. Transcriptome transfer produces a predictable cellular phenotype. *Proceedings of the National Academy of Sciences*, 106(18):7624–7629, 2009.

[22] Sui Huang, Gabriel Eichler, Yaneer Bar-Yam, and Donald E. Ingber. Cell fates as high-dimensional attractor states of a complex gene regulatory network. *Physical Review Letters*, 94(12), 04 2005.

[23] Qi-Long Ying, Jason Wray, Jennifer Nichols, Laura Batlle-Morera, Bradley Doble, James Woodgett, Philip Cohen, and Austin Smith. The ground state of embryonic stem cell self-renewal. *Nature*, 453(7194):519–523, 05 2008.

[24] H Weintraub, S J Tapscott, R L Davis, M J Thayer, M A Adam, A B Lassar, and A D Miller. Activation of muscle-specific genes in pigment, nerve, fat, liver, and fibroblast cell lines by forced expression of myod. *Proceedings of the National Academy of Sciences*, 86(14):5434–5438, 1989.

[25] Balazsi, Alexander van Oudenaarden, and James J. Collins. Cellular decision making and biological noise: From microbes to mammals. *Cell*, 144(6):910–925, 03 2011.

[26] Jacob Hanna, Krishanu Saha, Bernardo Pando, Jeroen van Zon, Christopher J. Lengner, Menno P. Creyghton, Alexander van Oudenaarden, and Rudolf Jaenisch. Direct cell reprogramming is a stochastic process amenable to acceleration. *Nature*, 462(7273):595–601, 12 2009.

[27] Hannah H. Chang, Martin Hemberg, Mauricio Barahona, Donald E. Ingber, and Sui Huang. Transcriptome-wide noise controls lineage choice in mammalian progenitor cells. *Nature*, 453(7194):544–547, 05 2008.

[28] Rupa Sridharan, Jason Tchieu, Mike J. Mason, Robin Yachechko, Edward Kuoy, Steve Horvath, Qing Zhou, and Kathrin Plath. Role of the murine reprogramming factors in the induction of pluripotency. *Cell*, 136(2):364–377, 01 2009.

[29] Tarjei S. Mikkelsen, Jacob Hanna, Xiaolan Zhang, Manching Ku, Marius Wernig, Patrick Schorderet, Bradley E. Bernstein, Rudolf Jaenisch, Eric S. Lander, and Alexander Meissner. Dissecting direct reprogramming through integrative genomic analysis. *Nature*, 454(7200):49–55, 07 2008.

[30] C. H. Waddington. *The Strategy of the Genes. A Discussion of Some Aspects of Theoretical Biology.* Allen and Unwin, London, 1957.

[31] S. A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution.* Oxford University Press, New York City, 1993.

[32] Tariq Enver, Martin Pera, Carsten Peterson, and Peter W. Andrews. Stem cell states, fates, and the rules of attraction. *Cell Stem Cell*, 4(5):387–397, 05 2009.

[33] Joseph X. Zhou and Sui Huang. Understanding gene circuits at cell-fate branch points for rational cell reprogramming. *Trends in Genetics*, 27(2):55–62, 02 2011.

[34] James E. Ferrell. Bistability, bifurcations, and waddington's epigenetic landscape. *Current Biology*, 22(11):R458–R466, 6 2012.

[35] Timothy S. Gardner, Charles R. Cantor, and James J. Collins. Construction of a genetic toggle switch in escherichia coli. *Nature*, 403(6767):339–342, 01 2000.

[36] Pankaj Mehta. Exponential sensitivity of noise-driven switching in genetic networks. *Physical Biology*, 5(2), 2008.

[37] J. Wang, K. Zhang, and E. Wang. Kinetic paths, time scale, and underlying landscapes: A path integral framework to study global natures of nonequilibrium systems and networks. *The Journal of chemical physics*, 133:125103, 2010.

[38] Sabine Loewer, Moran N Cabili, Mitchell Guttman, Yuin-Han Loh, Kelly Thomas, In Hyun Park, Manuel Garber, Matthew Curran, Tamer Onder, and Suneet Agarwal. Large intergenic non-coding rna-ror modulates reprogramming of human induced pluripotent stem cells. *Nature genetics*, 42(12):1113–1117

[39] Jin Wang, Li Xu, Erkang Wang, and Sui Huang. The potential landscape of genetic circuits imposes the arrow of time in stem cell differentiation. *Biophysical Journal*, 99(1):29–39, 7 2010.

[40] J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.

[41] Daniel J. Amit, Hanoch Gutfreund, and H. Sompolinsky. Spin-glass models of neural networks. *Physical Review A*, 32(2):1007–1018, 08 1985.

[42] I. Kanter and H. Sompolinsky. Associative recall of memory without errors. *Physical Review A*, 35(1):380–392, 01 1987.

[43] D.J. Amit. *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge Univ. Press, Cambridge, 1992.

[44] Melissa B. Miller and Bonnie L. Bassler. Quorum sensing in bacteria. *Annual Review of Microbiology*, 55(1):165–199, 2015/04/23 2001.

[45] Victor Sourjik and Ned S Wingreen. Responding to chemical gradients: bacterial chemotaxis. *Current Opinion in Cell Biology*, 24(2):262–268, 4 2012.

[46] Catriona Y. Logan and Roel Nusse. The wnt signaling pathway in development and disease. *Annual Review of Cell and Developmental Biology*, 20(1):781–810, 2015/04/23 2004.

[47] Michael B Elowitz and Stanislas Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403(6767):335–338

[48] Ari E. Friedland, Timothy K. Lu, Xiao Wang, David Shi, George Church, and James J. Collins. Synthetic gene networks that count. *Science*, 324(5931):1199–1202, 2009.

[49] Tal Danino, Octavio Mondragon-Palomino, Lev Tsimring, and Jeff Hasty. A synchronized quorum of genetic clocks. *Nature*, 463(7279):326–330, 01 2010.

[50] Shankar Mukherji and Alexander van Oudenaarden. Synthetic biology: understanding biological design from synthetic circuits. *Nature Reviews Genetics*, 10(12):859–871, 12 2009.

[51] Priscilla E. M. Purnick and Ron Weiss. The second wave of synthetic biology: from modules to systems. *Nature Reviews Molecular Cell Biology*, 10(6):410–422, 06 2009.

[52] Ahmad S. Khalil and James J. Collins. Synthetic biology: applications come of age. *Nature Reviews Genetics*, 11(5):367–379, 05 2010.

[53] Nagarajan Nandagopal and Michael B. Elowitz. Synthetic biology: Integrated gene circuits. *Science*, 333(6047):1244–1248, 2011.

[54] Warren C. Ruder, Ting Lu, and James J. Collins. Synthetic biology moving into the clinic. *Science*, 333(6047):1248–1252, 2011.

[55] Albert J. Keung, J. Keith Joung, Ahmad S. Khalil, and James J. Collins. Chromatin regulation at the frontier of synthetic biology. *Nature Reviews Genetics*, 16(3):159–171, 03 2015.

[56] Jeffrey C Way, James J Collins, Jay D Keasling, and Pamela A Silver. Integrating biological redesign: where synthetic biology came from and where it needs to go. *Cell*, 157(1):151–161

[57] Roberta Kwok. Five hard truths for synthetic biology. *Nature News*, 463(7279):288–290

[58] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191 0018–8646, 1961.

[59] Simon B Laughlin. Energy as a constraint on the coding and processing of sensory information. *Current opinion in neurobiology*, 11(4):475–480 2001.

[60] Hong Qian and Daniel A. Beard. Thermodynamics of stoichiometric biochemical networks in living systems far from equilibrium. *Biophysical Chemistry*, 114(2–3):213–220, 4 2005.

[61] Hong Qian. Phosphorylation energy hypothesis: Open chemical systems and their biological functions. *Annual Review of Physical Chemistry*, 58(1):113–142, 2007.

[62] Xue-Juan Zhang, Hong Qian, and Min Qian. Stochastic theory of nonequilibrium steady states and its applications. part i. *Physics Reports*, 510(1–2):1–86, 1 2012.

[63] Hao Ge, Min Qian, and Hong Qian. Stochastic theory of nonequilibrium steady states. part ii: Applications in chemical biophysics. *Physics Reports*, 510(3):87–118, 1 2012.

[64] AC Barato, D Hartich, and U Seifert. Information-theoretic versus thermodynamic entropy production in autonomous sensory networks. *Physical Review E*, 87(4):042104, 2013.

[65] Ganhui Lan, Pablo Sartori, Silke Neumann, Victor Sourjik, and Yuhai Tu. The energy-speed-accuracy trade-off in sensory adaptation. *Nature Physics*, 8(5):422–428, 05 2012.

[66] Christopher C. Govern and Pieter Rein ten Wolde. Fundamental limits on sensing chemical concentrations with linear biochemical networks. *Physical Review Letters*, 109(21):218103–, 11 2012.

[67] Takahiro Sagawa and Masahito Ueda. Nonequilibrium thermodynamics of feedback control. *Physical Review E*, 85(2):021104–, 02 2012.

[68] Susanne Still, David A. Sivak, Anthony J. Bell, and Gavin E. Crooks. Thermodynamics of prediction. *Physical Review Letters*, 109(12):120604–, 09 2012.

[69] A. C. Barato, D. Hartich, and U. Seifert. Information-theoretic versus thermodynamic entropy production in autonomous sensory networks. *Physical Review E*, 87(4):042104–, 04 2013.

[70] Christopher C Govern and Pieter Rein ten Wolde. How biochemical resources determine fundamental limits in cellular sensing. *arXiv preprint arXiv:1308.1449*, 2013.

[71] Nils B Becker, Andrew Mugler, and Pieter Rein ten Wolde. Prediction and dissipation in biochemical sensing. *arXiv preprint arXiv:1312.5625*, 2013.

[72] Stefano Bo, Marco Del Giudice, and Antonio Celani. Thermodynamic limits to information harvesting by sensory systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(1):P01014, 2015.

[73] Christopher C. Govern and Pieter Rein ten Wolde. Energy dissipation and noise correlations in biochemical sensing. *Physical Review Letters*, 113(25):258102–, 12 2014.

[74] Christopher C. Govern and Pieter Rein ten Wolde. Optimal resource allocation in cellular sensing systems. *Proceedings of the National Academy of Sciences*, 2014.

[75] Kazunari Kaizu, Wiet de Ronde, Joris Paijmans, Koichi Takahashi, Filipe Tostevin, and Pieter Rein ten Wolde. The berg-purcell limit revisited. *Biophysical Journal*, 106(4):976–985, 2 2014.

[76] Garud Iyengar and Madan Rao. A cellular solution to an information-processing problem. *Proceedings of the National Academy of Sciences*, 111(34):12402–12407, 2014.

[77] Jangir Selimkhanov, Brooks Taylor, Jason Yao, Anna Pilko, John Albeck, Alexander Hoffmann, Lev Tsimring, and Roy Wollman. Accurate information transmission through dynamic biochemical signaling networks. *Science*, 346(6215):1370–1373, 2014.

[78] Arvind Murugan, David A Huse, and Stanislas Leibler. Discriminatory proofreading regimes in nonequilibrium systems. *Physical Review X*, 4(2):021016, 2014.

[79] Andre C Barato, David Hartich, and Udo Seifert. Efficiency of cellular information processing. *New Journal of Physics*, 16(10):103024

[80] Sosuke Ito and Takahiro Sagawa. Maxwell's demon in biochemical signal transduction. *arXiv preprint arXiv:1406.5810*, 2014.

[81] Pablo Sartori, Léo Granger, Chiu Fan Lee, and Jordan M. Horowitz. Thermodynamic costs of information processing in sensory adaptation. *PLoS Computional Biology*, 10(12):e1003974 EP –, 12 2014.

[82] Thomas E Ouldridge, Christopher C Govern, and Pieter Rein ten Wolde. On the connection between computational and biochemical measurement. *arXiv preprint arXiv:1503.00909*, 2015.

[83] David Hartich, Andre C Barato, and Udo Seifert. Nonequilibrium sensing and its analogy to kinetic proofreading. *arXiv preprint arXiv:1502.02594*, 2015.

[84] Charles H Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940

[85] C.H. Bennett. Notes on landauer's principle, reversible computation, and maxwell's demon. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 34(3):501–510, 2003.

[86] Antoine Berut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz. Experimental verification of landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 03 2012.

[87] Yonggun Jun, Momčilo Gavrilov, and John Bechhoefer. High-precision test of landauer's principle in a feedback trap. *Physical Review Letters*, 113(19):190601–, 11 2014.

[88] Suriyanarayanan Vaikuntanathan and Christopher Jarzynski. Modeling maxwell's demon with a microcanonical szilard engine. *Physical Review E*, 83(6):061120–, 06 2011.

[89] Christopher Jarzynski. Equalities and inequalities: Irreversibility and the second law of thermodynamics at the nanoscale. *Annual Review of Condensed Matter Physics*, 2(1):329–351, 2015/04/22 2011.

[90] Dibyendu Mandal and Christopher Jarzynski. Work and information processing in a solvable model of maxwell's demon. *Proceedings of the National Academy of Sciences*, 109(29):11641–11645, 2012.

[91] Suriyanarayanan Vaikuntanathan, Todd R. Gingrich, and Phillip L. Geissler. Dynamic phase transitions in simple driven kinetic networks. *Physical Review E*, 89(6):062108–, 06 2014.

[92] M. Cristina Diamantini and Carlo A. Trugenberger. Generalized landauer bound as a universal thermodynamic entropy in continuous phase transitions. *Physical Review E*, 89(5):052138–, 05 2014.

[93] Moupriya Das. Capturing the landauer bound through the application of a detailed jarzynski equality for entropic memory erasure. *Physical Review E*, 90(6):062120–, 12 2014.

[94] Juan M. R. Parrondo, Jordan M. Horowitz, and Takahiro Sagawa. Thermodynamics of information. *Nature Physics*, 11(2):131–139, 02 2015.

[95] J. J. Hopfield. Kinetic proofreading: A new mechanism for reducing errors in biosynthetic processes requiring high specificity. *Proceedings of the National Academy of Sciences*, 71(10):4135–4139, 10 1974.

[96] Timothy W Mckeithan. Kinetic proofreading in t-cell receptor signal transduction. *Proceedings of the national academy of sciences*, 92(11):5042–5046, 1995.

[97] Arvind Murugan, David A. Huse, and Stanislas Leibler. Speed, dissipation, and error in kinetic proofreading. *Proceedings of the National Academy of Sciences*, 109(30):12034–12039, 2012.

[98] Domitilla Del Vecchio, Alexander J Ninfa, and Eduardo D Sontag. Modular cell biology: retroactivity and insulation. *Molecular Systems Biology*, 4(1), 2008.

[99] John P Barton and Eduardo D Sontag. The energy costs of insulators in biochemical networks. *Biophysical journal*, 104(6):1380–1390

[100] Deepak Mishra, Phillip M Rivera, Allen Lin, Domitilla Del Vecchio, and Ron Weiss. A load driver device for engineering modularity in biological networks. *Nature biotechnology*, 32(12):1268–1275

[101] William Bialek. *Biophysics: Searching for Principles*. Princeton University Press, 2012.

[102] Brian Munsky, Ilya Nemenman, and Golan Bel. Specificity and completion time distributions of biochemical processes. *The Journal of Chemical Physics*, 131(23):–, 2009.

[103] Golan Bel, Brian Munsky, and Ilya Nemenman. The simplicity of completion time distributions for common complex biochemical processes. *Physical Biology*, 7(1):016003, 2010.

[104] Alex H. Lang, Charles K. Fisher, Thierry Mora, and Pankaj Mehta. Thermodynamics of statistical inference by cells. *Physical Review Letters*, 113(14):148103–, 10 2014.

[105] Andre C. Barato and Udo Seifert. Thermodynamic uncertainty relation for biomolecular processes. *Physical Review Letters*, 114(15):158101–, 04 2015.

[106] Peter B. Detwiler, Sharad Ramanathan, Anirvan Sengupta, and Boris I. Shraiman. Engineering aspects of enzymatic signal transduction: Photoreceptors in the retina. *Biophysical Journal*, 79(6):2801–2817, 12 2000.

[107] H. C. Berg and E. M. Purcell. Physics of chemoreception. *Biophysical Journal*, 20(2):193–219, 11 1977.

[108] Richard Phillips Feynman, JG Hey, and Robin W *Feynman lectures on computation*. Addison-Wesley Longman Publishing Co., Inc., 1998.

[109] Pankaj Mehta and David J. Schwab. Energetic costs of cellular computation. *Proceedings of the National Academy of Sciences*, 109(44):17978–17982, 2012.

[110] Piro Siuti, John Yazbek, and Timothy K Lu. Synthetic circuits integrating logic and memory in living cells. *Nature biotechnology*, 31(5):448–452

[111] Thomas Vierbuchen and Marius Wernig. Molecular roadblocks for cellular reprogramming. *Molecular Cell*, 47(6):827–838, 9 2012.

[112] Yosef Buganim, Dina A. Faddah, Albert W. Cheng, Elena Itskovich, Styliani Markoulaki, Kibibi Ganz, Sandy L. Klemm, Alexander van Oudenaarden, and Rudolf Jaenisch. Single-cell expression analyses during cellular reprogramming reveal an early stochastic and a late hierarchic phase. *Cell*, 150(6):1209–1222, 9 2012.

[113] Janghwan Kim, Jem A. Efe, Saiyong Zhu, Maria Talantova, Xu Yuan, Shufen Wang, Stuart A. Lipton, Kang Zhang, and Sheng Ding. Direct reprogramming of mouse fibroblasts to neural progenitors. *Proceedings of the National Academy of Sciences*, 108(19):7838–7843, 2011.

[114] E.H. Davidson. *The Regulatory Genome: Gene Regulatory Networks in Development and Evolution.* Academic Press, London, 2006.

[115] Sui Huang, Yan-Ping Guo, Gillian May, and Tariq Enver. Bifurcation dynamics in lineage-commitment in bipotent progenitor cells. *Developmental Biology*, 305(2):695–713, 5 2007.

[116] Joseph Xu Zhou, Lutz Brusch, and Sui Huang. Predicting pancreas cell fate decisions and reprogramming with a hierarchical multi-attractor model. *PLoS ONE*, 6(3):e14752–, 03 2011.

[117] Francis Corson and Eric Dean Siggia. Geometry, epistasis, and developmental patterning. *Proceedings of the National Academy of Sciences*, 109(15):5568–5575, 2012.

[118] Adrien Henry, Françoise Monéger, Areejit Samal, and Olivier C Martin. Network function shapes network structure: the case of the arabidopsis flower organ specification genetic network. *Molecular BioSystems*, 9(7):1726–1735, 2013.

[119] M. Zagorski, A. Krzywicki, and O. C. Martin. Edge usage, motifs, and regulatory logic for cell cycling genetic networks. *Physical Review E*, 87(1):012727–, 01 2013.

[120] Ben D. MacArthur, Avi Ma'ayan, and Ihor R. Lemischka. Systems biology of stem cell fate and cellular reprogramming. *Nature Reviews Molecular Cell Biology*, 10(10):672–681, 10 2009.

[121] Rong Lu, Florian Markowetz, Richard D. Unwin, Jeffrey T. Leek, Edoardo M. Airoldi, Ben D. MacArthur, Alexander Lachmann, Roye Rozov, Avi Ma'ayan, Laurie A. Boyer, Olga G. Troyanskaya, Anthony D. Whetton, and Ihor R. Lemischka. Systems-level dynamic analyses of fate change in murine embryonic stem cells. *Nature*, 462(7271):358–362, 11 2009.

[122] Christopher R. S. Banerji, Diego Miranda-Saavedra, Simone Severini, Martin Widschwendter, Tariq Enver, Joseph X. Zhou, and Andrew E. Teschendorff. Cellular

network entropy as the energy potential in waddington's differentiation landscape. *Scientific Reports.*, 3, 10 2013.

[123] Jin Wang, Kun Zhang, Li Xu, and Erkang Wang. Quantifying the waddington landscape and biological paths for development and differentiation. *Proceedings of the National Academy of Sciences*, 108(20):8257–8262, 2011.

[124] Joseph D. Bryngelson, José N Onuchic, Nicholas D. Socci, and Peter G. Wolynes. Funnels, pathways, and the energy landscape of protein folding: A synthesis. *Proteins: Structure, Function, and Bioinformatics*, 21(3):167–195, 1995.

[125] Tarjei S. Mikkelsen, Manching Ku, David B. Jaffe, Biju Issac, Erez Lieberman, Georgia Giannoukos, Pablo Alvarez, William Brockman, Tae-Kyung Kim, Richard P. Koche, William Lee, Eric Mendenhall, Aisling O'Donovan, Aviva Presser, Carsten Russ, Xiaohui Xie, Alexander Meissner, Marius Wernig, Rudolf Jaenisch, Chad Nusbaum, Eric S. Lander, and Bradley E. Bernstein. Genome-wide maps of chromatin state in pluripotent and lineage-committed cells. *Nature*, 448(7153):553–560, 08 2007.

[126] Alexander Meissner, Tarjei S. Mikkelsen, Hongcang Gu, Marius Wernig, Jacob Hanna, Andrey Sivachenko, Xiaolan Zhang, Bradley E. Bernstein, Chad Nusbaum, David B. Jaffe, Andreas Gnirke, Rudolf Jaenisch, and Eric S. Lander. Genome-scale dna methylation maps of pluripotent and differentiated cells. *Nature*, 454(7205):766–770, 08 2008.

[127] Thomas Jenuwein and C. David Allis. Translating the histone code. *Science*, 293(5532):1074–1080, 2001.

[128] Uri Ben-David, Yoav Mayshar, and Nissim Benvenisty. Large-scale analysis reveals acquisition of lineage-specific chromosomal aberrations in human adult stem cells. *Cell Stem Cell*, 9(2):97–102, 08 2011.

[129] Cesar A. Sommer, Andreia Gianotti Sommer, Tyler A. Longmire, Constantina Christodoulou, Dolly D. Thomas, Monica Gostissa, Fred W. Alt, George J. Murphy,

Darrell N. Kotton, and Gustavo Mostoslavsky. Excision of reprogramming transgenes improves the differentiation potential of ips cells generated with a single excisable vector. *STEM CELLS*, 28(1):64–74, 2010.

[130] Shinji Masui, Satoshi Ohtsuka, Rika Yagi, Kadue Takahashi, Minoru S H Ko, and Hitoshi Niwa. Rex1/zfp42 is dispensable for pluripotency in mouse es cells. *BMC Developmental Biology*, 8:45, 2008.

[131] Olfa Khalfallah, Matthieu Rouleau, Pascal Barbry, Barbara Bardoni, and Enzo Lalli. Dax-1 knockdown in mouse embryonic stem cells induces loss of pluripotency and multilineage differentiation. *STEM CELLS*, 27(7):1529–1537, 2009.

[132] Terry P. Yamaguchi, Shinji Takada, Yoshiaki Yoshikawa, Nongying Wu, and Andrew P. McMahon. T (brachyury) is a direct target of wnt3a during paraxial mesoderm specification. *Genes and Development*, 13(24):3185–3190, 1999.

[133] Rudolf Jaenisch. Nuclear cloning and direct reprogramming: The long and the short path to stockholm. *Cell Stem Cell*, 11(6):744–747, 12 2012.

[134] Marius Wernig, Alexander Meissner, John P. Cassady, and Rudolf Jaenisch. c-myc is dispensable for direct reprogramming of mouse fibroblasts. *Cell Stem Cell*, 2(1):10–12, 1 2008.

[135] Tamer T. Onder, Nergis Kara, Anne Cherry, Amit U. Sinha, Nan Zhu, Kathrin M. Bernt, Patrick Cahan, Ogan. B. Mancarci, Juli Unternaehrer, Piyush B. Gupta, Eric S. Lander, Scott A. Armstrong, and George Q. Daley. Chromatin-modifying enzymes as modulators of reprogramming. *Nature*, 483:598–602, 03 2012.

[136] Elad Schneidman, Michael J. Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 04 2006.

[137] W. Bialek and R. Ranganathan. Rediscovering the power of pairwise interactions. *Arxiv preprint*, arXiv:0712.4397, 2007.

[138] Simona Cocco, Remi Monasson, and Martin Weigt. From principal component to direct coupling analysis of coevolution in proteins: Low-eigenvalue modes are needed for structure prediction. *PLoS computational biology*, 9(8):e1003176

[139] Thierry Mora, Aleksandra M. Walczak, William Bialek, and Curtis G. Callan. Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences*, 107(12):5405–5410, 2010.

[140] Gašper Tkačik, Olivier Marre, Dario Amodei, Elad Schneidman, William Bialek, and Michael J Berry II. Searching for collective behavior in a large network of sensory neurons. *PLoS computational biology*, 10(1):e1003408

[141] Andrew S. Yoo, Alfred X. Sun, Li Li, Aleksandr Shcheglovitov, Thomas Portmann, Yulong Li, Chris Lee-Messer, Ricardo E. Dolmetsch, Richard W. Tsien, and Gerald R. Crabtree. Microrna-mediated conversion of human fibroblasts to neurons. *Nature*, 476(7359):228–231, 08 2011.

[142] Luigi Warren, Philip D. Manos, Tim Ahfeldt, Yuin-Han Loh, Hu Li, Frank Lau, Wataru Ebina, Pankaj K. Mandal, Zachary D. Smith, Alexander Meissner, George Q. Daley, Andrew S. Brack, James J. Collins, Chad Cowan, Thorsten M. Schlaeger, and Derrick J. Rossi. Highly efficient reprogramming to pluripotency and directed differentiation of human cells with synthetic modified mrna. *Cell Stem Cell*, 7(5):618–630, 11 2010.

[143] Tyler A. Longmire, Laertis Ikonomou, Finn Hawkins, Constantina Christodoulou, Yuxia Cao, J. C. Jean, Letty W. Kwok, Hongmei Mou, Jayaraj Rajagopal, Steven S. Shen, Anne A. Dowton, Maria Serra, Daniel J. Weiss, Michael D. Green, Hans-Willem Snoeck, Maria I. Ramirez, and Darrell N. Kotton. Efficient derivation of purified lung

and thyroid progenitors from embryonic stem cells. *Cell Stem Cell*, 10(4):398–411, 4 2012.

[144] Kazutoshi Takahashi, Keisuke Okita, Masato Nakagawa, and Shinya Yamanaka. Induction of pluripotent stem cells from fibroblast cultures. *Nature Protocols*, 2(12):3081–3089, 11 2007.

[145] Jun Xu, Yuanyuan Du, and Hongkui Deng. Direct lineage reprogramming: Strategies, mechanisms, and applications. *Cell Stem Cell*, 16(2):119–134, 2 2015.

[146] Yoach Rais, Asaf Zviran, Shay Geula, Ohad Gafni, Elad Chomsky, Sergey Viukov, Abed AlFatah Mansour, Inbal Caspi, Vladislav Krupalnik, Mirie Zerbib, Itay Maza, Nofar Mor, Dror Baran, Leehee Weinberger, Diego A. Jaitin, David Lara-Astiaso, Ronnie Blecher-Gonen, Zohar Shipony, Zohar Mukamel, Tzachi Hagai, Shlomit Gilad, Daniela Amann-Zalcenstein, Amos Tanay, Ido Amit, Noa Novershtern, and Jacob H. Hanna. Deterministic direct reprogramming of somatic cells to pluripotency. *Nature*, 502(7469):65–70, 10 2013.

[147] Jose M. Polo, Endre Anderssen, Ryan M. Walsh, Benjamin A. Schwarz, Christian M. Nefzger, Sue Mei Lim, Marti Borkent, Effie Apostolou, Sara Alaei, Jennifer Cloutier, Ori Bar-Nur, Sihem Cheloufi, Matthias Stadtfeld, Maria Eugenia Figueroa, Daisy Robinton, Sridaran Natesan, Ari Melnick, Jinfang Zhu, Sridhar Ramaswamy, and Konrad Hochedlinger. A molecular roadmap of reprogramming somatic cells into ips cells. *Cell*, 151(7):1617–1632, 12 2012.

[148] Kyung-Min Chung, Frederick W. Kolling IV, Matthew D. Gajdosik, Steven Burger, Alexander C. Russell, and Craig E. Nelson. Single cell analysis reveals the stochastic phase of reprogramming to pluripotency is an ordered probabilistic process. *PLoS ONE*, 9(4):e95304 EP –, 04 2014.

[149] Sui Huang. The molecular and mathematical basis of waddington's epigenetic landscape: A framework for post-darwinian biology? *BioEssays*, 34(2):149–157, 2012.

[150] Chunhe Li and Jin Wang. Quantifying cell fate decisions for differentiation and reprogramming of a human stem cell network: Landscape and biological paths. *PLoS Computional Biology*, 9(8):e1003165 EP –, 08 2013.

[151] Merja Heinaniemi, Matti Nykter, Roger Kramer, Anke Wienecke-Baldacchino, Lasse Sinkkonen, Joseph Xu Zhou, Richard Kreisberg, Stuart A Kauffman, Sui Huang, and Ilya Shmulevich. Gene-pair expression signatures reveal lineage control. *Nature Methods*, 10(6):577–583, 06 2013.

[152] Li Xu, Kun Zhang, and Jin Wang. Exploring the mechanisms of differentiation, dedifferentiation, reprogramming and transdifferentiation. *PLoS ONE*, 9(8):e105216 EP –, 08 2014.

[153] Chunhe Li and Jin Wang. Landscape and flux reveal a new global view and physical quantification of mammalian cell cycle. *Proceedings of the National Academy of Sciences*, 111(39):14130–14135, 2014.

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

[155] Alex H. Lang, Hu Li, James J. Collins, and Pankaj Mehta. Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes. *PLoS Computional Biology*, 10(8):e1003734 EP –, 08 2014.

[156] Payman Samavarchi-Tehrani, Azadeh Golipour, Laurent David, Hoon-ki Sung, Tobias A. Beyer, Alessandro Datti, Knut Woltjen, Andras Nagy, and Jeffrey L. Wrana. Functional genomics reveals a bmp-driven mesenchymal-to-epithelial transition in the initiation of somatic cell reprogramming. *Cell Stem Cell*, 7(1):64–77, 7 2010.

[157] Bruno Di Stefano, Jose Luis Sardina, Chris van Oevelen, Samuel Collombet, Eric M. Kallin, Guillermo P. Vicent, Jun Lu, Denis Thieffry, Miguel Beato, and Thomas Graf. C/ebpalpha poises b cells for rapid reprogramming into induced pluripotent stem cells. *Nature*, 506(7487):235–239, 02 2014.

[158] Bruno Di Stefano, Samuel Collombet, and Thomas Graf. Time-resolved gene expression profiling during reprogramming of c/ebp-pulsed b cells into ips cells. *Scientific Data*, 1:EP –, 05 2014.

[159] Shinya Yamanaka. Elite and stochastic models for induced pluripotent stem cell generation. *Nature*, 460(7251):49–52, 07 2009.

[160] David Marin Roma, Ruadhan A. O'Flanagan, Andrei E. Ruckenstein, Anirvan M. Sengupta, and Ranjan Mukhopadhyay. Optimal path to epigenetic switching. *Physical Review E*, 71(1):011902–, 01 2005.

[161] Andrew A. Wilson, Lei Ying, Marc Liesa, Charis-Patricia Segeritz, Jason A. Mills, Steven S. Shen, Jyhchang Jean, Geordie C. Lonza, Derek C. Liberti, Alex H. Lang, Jean Nazaire, Adam C. Gower, Franz-Josef Müeller, Pankaj Mehta, Adriana Ordóñez, David A. Lomas, Ludovic Vallier, George J. Murphy, Gustavo Mostoslavsky, Avrum Spira, Orian S. Shirihai, Maria I. Ramirez, Paul Gadue, and Darrell N. Kotton. Emergence of a stage-dependent human liver disease signature with directed differentiation of alpha-1 antitrypsin-deficient ips cells. *Stem Cell Reports*, (0):–, 2015.

[162] Andrew L. Ferguson, Jaclyn K. Mann, Saleha Omarjee, Thumbi Ndung'u, Bruce D. Walker, and Arup K. Chakraborty. Translating hiv sequences into quantitative fitness landscapes predicts viral vulnerabilities for rational immunogen design. *Immunity*, 38(3):606–617, 3 2013.

[163] JoséNelson Onuchic, Zaida Luthey-Schulten, and Peter G. Wolynes. Theory of protein folding: The energy landscape perspective. *Annual Review of Physical Chemistry*, 48(1):545–600, 2015/04/28 1997.

[164] JoséNelson Onuchic and Peter G Wolynes. Theory of protein folding. *Current Opinion in Structural Biology*, 14(1):70–75, 2 2004.

[165] William Bialek and Sima Setayeshgar. Physical limits to biochemical signaling. *Proceedings of the National Academy of Sciences of the United States of America*, 102(29):10040–10045, 2005.

[166] Robert G. Endres and Ned S. Wingreen. Maximum likelihood and the single receptor. *Physical Review Letters*, 103(15):158101–, 10 2009.

[167] Bo Hu, Wen Chen, Wouter-Jan Rappel, and Herbert Levine. Physical limits on cellular sensing of spatial gradients. *Physical Review Letters*, 105(4):048104–, 07 2010.

[168] Thierry Mora and Ned S. Wingreen. Limits of sensing temporal concentration changes by single cells. *Physical Review Letters*, 104(24):248101–, 06 2010.

[169] Xiang Cheng, Lina Merchan, Martin Tchernookov, and Ilya Nemenman. A large number of receptors may reduce cellular response time variation. *Physical Biology*, 10(3):035008, 2013.

[170] Juan E. Keymer, Robert G. Endres, Monica Skoge, Yigal Meir, and Ned S. Wingreen. Chemosensing in escherichia coli: Two regimes of two-state receptors. *Proceedings of the National Academy of Sciences of the United States of America*, 103(6):1786–1791, 2006.

[171] Terrell L Hill. Free energy transduction and biochemical cycle kinetics. 1989.

[172] JoelL. Lebowitz and Herbert Spohn. A gallavotti–cohen-type symmetry in the large deviation functional for stochastic dynamics. *Journal of Statistical Physics*, 95(1-2):333–365, 1999.

[173] Donald Voet and JG Voet. *Biochemistry*. John Wiley and Sons New York, 3rd edition edition, 2004.

[174] Sean Escola, Michael Eisele, Kenneth Miller, and Liam Paninski. Maximally reliable markov chains under energy constraints. *Neural computation*, 21(7):1863–1912

[175] David M. Suter, Nacho Molina, David Gatfield, Kim Schneider, Ueli Schibler, and Felix Naef. Mammalian genes are transcribed with widely different bursting kinetics. *Science*, 332(6028):472–474, 2011.

[176] László Csanády, Paola Vergani, and David C. Gadsby. Strict coupling between cftr's catalytic cycle and gating of its cl ion pore revealed by distributions of open channel burst durations. *Proceedings of the National Academy of Sciences*, 107(3):1241–1246, 2010.

[177] Martin Depken, Juan M.R. Parrondo, and Stephan W. Grill. Intermittent transcription dynamics for the rapid production of long transcripts of high fidelity. *Cell Reports*, 5(2):521 – 530, 2013.

[178] Suriyanarayanan Vaikuntanathan, Todd R Gingrich, and Phillip L Geissler. Dynamic phase transitions in simple driven kinetic networks. *arXiv preprint arXiv:1307.0801*, 2013.

[179] Alex H Lang, Hu Li, James J Collins, and Pankaj Mehta. Epigenetic landscapes explain partially reprogrammed cells and identify key reprogramming genes. *arXiv*, 2012.

[180] Herbert A David and Haikady N Nagaraja. *Order Statistics,*, volume 7. John Wiley & Sons, Hoboken, 2003.

[181] Hong-Mei Zhang, Hu Chen, Wei Liu, Hui Liu, Jing Gong, Huili Wang, and An-Yuan Guo. Animaltfdb: a comprehensive animal transcription factor database. *Nucleic Acids Research*, 40(D1):D144–D149, 2012.

[182] J J Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81(10):3088–3092, 1984.

[183] Daniel J. Amit, Hanoch Gutfreund, and H. Sompolinsky. Storing infinite numbers of patterns in a spin-glass model of neural networks. *Physical Review Letters*, 55(14), 09 1985.

[184] Jing Jiang, Wenjian Lv, Xiaoying Ye, Lingbo Wang, Man Zhang, Hui Yang, Maja Okuka, Chikai Zhou, Xuan Zhang, Lin Liu, and Jinsong Li. Zscan4 promotes genomic stability during reprogramming and dramatically improves the quality of ips cells as demonstrated by tetraploid complementation. *Cell Research*, 23(1):92–106, 01 2013.

[185] Linda Shushan Lim, Yuin-Han Loh, Weiwei Zhang, Yixun Li, Xi Chen, Yinan Wang, Manjiri Bakre, Huck-Hui Ng, and Lawrence W. Stanton. Zic3 is required for maintenance of pluripotency in embryonic stem cells. *Molecular Biology of the Cell*, 18(4):1348–1358, 2007.

[186] Hong-bing Yu, Galih Kunarso, Felicia Huimei Hong, and Lawrence W. Stanton. Zfp206, oct4, and sox2 are integrated components of a transcriptional regulatory network in embryonic stem cells. *Journal of Biological Chemistry*, 284(45):31327–31335, 2009.

[187] Shoudong Ye, Ping Li, Chang Tong, and Qi-Long Ying. Embryonic stem cell self-renewal pathways converge on the transcription factor tfcp2l1. *EMBO Journal*, 32(19):2548–2560, 10 2013.

[188] Haemi Lee, Hyo Jung Kim, Yoo Jeong Lee, Min-Young Lee, Hyeonjin Choi, Hyemin Lee, and Jae-woo Kim. Krüppel-like factor klf8 plays a critical role in adipocyte differentiation. *PLoS ONE*, 7(12):e52474 EP –, 12 2012.

[189] Owen R. Davies, Chia-Yi Lin, Aliaksandra Radzisheuskaya, Xinzhi Zhou, Jessica Taube, Guillaume Blin, Anna Waterhouse, Andrew J. H. Smith, and Sally Lowell. Tcf15 primes pluripotent cells for differentiation. *Cell Reports*, 3(2):472–484, 2 2013.

[190] Hong-Mei Zhang, Teng Liu, Chun-Jie Liu, Shuangyang Song, Xiantong Zhang, Wei Liu, Haibo Jia, Yu Xue, and An-Yuan Guo. Animaltfdb 2.0: a resource for expres-

sion, prediction and functional study of animal transcription factors. *Nucleic Acids Research*, 2014.

[191] Sidney Redner. *A Guide to First-Passage Processes*. Cambridge, 2001.

[192] Richard A. Blythe. Nonequilibrium phase transitions and dynamical scaling regimes. Master's thesis, Edinburgh, 2001.

[193] Faiza Hussain, Chinmaya Gupta, Andrew J. Hirning, William Ott, Kathleen S. Matthews, Krešimir Josić, and Matthew R. Bennett. Engineered temperature compensation in a synthetic genetic clock. *Proceedings of the National Academy of Sciences*, 111(3):972–977, 2014.

[194] Irwin H Segel. *Enzyme kinetics*, volume 360. Wiley, New York, 1975.

5 Curriculum Vitae