

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL: http://wrap.warwick.ac.uk/88919

Copyright and reuse:

This thesis is made available online and is protected by original copyright. Please scroll down to view the document itself. Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



Control and feedback in stochastic thermodynamics

by

Michael Maitland

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Complexity Science

August 2016



Contents

List of Figures v			\mathbf{v}
Acknowledgments vii			vii
Declara	Declarations viii		
Abstra	Abstract ix		
Table o	of nota	tion	x
Chapte	er 1 I	ntroduction	1
1.1	The se	econd law	1
1.2	Maxw	ell's dæmon	2
1.3	Inform	nation thermodynamics	5
1.4	The cu	urrent work	7
Chapte	er 2 K	Key concepts	10
2.1	Stocha	astic thermodynamics	11
2.2	Notati	on and framework	12
	2.2.1	Markov chains	13
	2.2.2	Time reversal	18
2.3	Trajec	tory functionals	20
	2.3.1	Definition of a functional	20
	2.3.2	Reversal of protocol	22
	2.3.3	Adjoint dynamics	23
	2.3.4	Adjoint dynamics with reversal of protocol	24
	2.3.5	Relationship between functionals	25
2.4	Relatio	on to entropy and heat	25
	2.4.1	Entropy	26
	2.4.2	Heat	28

2.5	Funda	amental fluctuation relationships
	2.5.1	Fundamental relation for reversal of protocol $\ldots \ldots \ldots$
	2.5.2	Fundamental relation for reversal of dynamics
	2.5.3	Fundamental relation for reversal of protocol and dynamics .
2.6	Large	deviations
	2.6.1	Definition of the large deviation principle
	2.6.2	Obtaining the large deviation rate function
	2.6.3	Rate function from modified Markov matrix
2.7	Fluctu	uation relations
	2.7.1	Integral fluctuation relations
	2.7.2	Detailed fluctuation relations
	2.7.3	Crooks-type fluctuation relations
	2.7.4	Fluctuation relations in the long-time limit
2.8	Concl	usion
Chapt	er 3 (Control and feedback
3.1	Contr	ol theory
	3.1.1	Control protocols
	3.1.2	Open-loop control
	3.1.3	Closed loop control
3.2	Inform	nation theory
	3.2.1	Shannon entropy
	3.2.2	Mutual information
	3.2.3	Change in uncertainty
3.3	Rever	sal of feedback protocol
	3.3.1	General form of fluctuation relationship $\ldots \ldots \ldots \ldots \ldots$
	3.3.2	Trajectory through state space
	3.3.3	Reverse trajectory
	3.3.4	Modified-dynamics
	3.3.5	Open-loop control
3.4	Gener	alised second law
	3.4.1	Interpretation of the generalised law
	3.4.2	Szilárd engine
3.5	Concl	usion
Chapt	er 4 I	Fluctuations in an Information engine
4.1	Inform	nation engines
	4.1.1	Types of information engine

	4.1.2	Model motivation	71
4.2	Model	description	74
	4.2.1	Calculating $P(\mathbf{y}_t)$	77
4.3	Entro	py production and integral fluctuation relations	78
4.4	Fluctu	ations of information	79
	4.4.1	Large Deviation Analysis	80
	4.4.2	Fluctuations of ΔI_s	81
	4.4.3	Trajectory Analysis	82
4.5	2-site	model	86
	4.5.1	System matrices	86
	4.5.2	Distribution of ΔI_s	87
	4.5.3	Large deviations of information	90
4.6	L-site	models	90
	4.6.1	General behaviour	90
	4.6.2	Maximum and Minimum of ΔI_s	95
	4.6.3	Markov chain model	96
	4.6.4	Beyond Markovian analysis	98
4.7	Discus	ssion	99
	4.7.1	Summary	99
	4.7.2	Possible extensions	99
	4.7.3	Conclusion	100
Chapte	er 5 I	Modelling of the measurement system	102
5.1	Work	and the 2^{nd} law \ldots	103
	5.1.1	Extraction of work through feedback	103
5.2	Model	lling the measurement process	105
	5.2.1	Langevin model	105
	5.2.2	Discretised model	107
	5.2.3	Functional for the process	109
	5.2.4	Protocol for Measurement and Exploitation	110
	5.2.5	Optimal protocols	111
5.3	Simula	ation	112
	5.3.1	Data	113
	5.3.2	Limiting behaviour	115
5.4	Doubl	e feedback	116
	5.4.1	Exploitation of the measurement device $\ldots \ldots \ldots \ldots$	117
	5.4.2	Double dæmon simulation	117

	5.4.3	Second-law violating?	118
5.5	Conclu	usion	121
	5.5.1	Related frameworks	122
	5.5.2	Experimental considerations	123
Chapte	er 6 C	Outlook and Discussion	125
6.1	Summ	ary and outlook	126
6.2	The se	cond law II	127
	6.2.1	Exorcising the dæmon	127
	6.2.2	The nature of information $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$	128
	6.2.3	Information as perspective	129

List of Figures

2.1	A cyclic chemical reaction	11
3.1	Bayesian causal diagram of a Markov chain	45
3.2	Bayesian causal diagram of a single repeated feedback process	48
3.3	Graph of the Shannon entropy $H(p)$ for a coin-flip	50
3.4	Venn diagram illustrating the relationship between Shannon entropy,	
	conditional Shannon entropy, joint Shannon entropy and Mutual in-	
	formation	53
3.5	State-space plot showing three different interpretations of a trajectory	
	through the system.	59
3.6	Causal diagram for generating the time-reversed trajectory when $\mathbb{Q}=$	
	P	59
3.7	State-space plot showing two different trajectories	60
3.8	Causal diagram for generating the reversed trajectory through path-	
	space	61
3.9	Bayesian causal diagram for a trajectory subject to open-loop feed-	
	back control.	64
3.10	Control cycle of the Szilárd engine	67
4.1	Schematic of an information engine model	73
4.2	Numerical tests of (2.89) , (3.34) and (4.16)	79
4.3	Baseline values of ΔI for unbiased and biased systems	84
4.4	Trajectories of \mathbf{x}_t , and \mathbf{y}_t plotted with the information gain ΔI_s and	
	$\ln(P(\mathbf{y}_{s-1})/P(\mathbf{y}_s))$	85
4.5	Cumulative density function for ΔI_s for different system sizes $L.~$	88
4.6	Large deviation rate function (4.22) for $L = 2$, $p = 0.2$, and $r = 0.9$.	91
4.7	Scatter plot of ΔI_{s+1} against ΔI_s	93
4.8	Sample average autocorrelation function of ΔI_s (4.42) for biased and	
	unbiased systems	94

4.9	Numerical results for $Max(\Delta I_s)$ and $-Min(\Delta I_s)$ with and predicted	
	values from (4.43) and (4.44)	95
4.10	The large deviation rate function (4.22) for $L = 10, p = 0.5, r = 0.9$.	96
4.11	The large deviation rate function (4.22) for $L = 10, p = 0.2, r = 0.9$.	97
5.1	Cartoon diagram depicting the dynamics for the X particle described	
	by (5.8) and (5.10)	108
5.2	The average work done over the course of a trajectory, with several	
	different time-scales for the evolution of the control parameter. $\ . \ .$	113
5.3	The average work done by each oscillator	114
5.4	Convergence of $\langle \mathcal{W}_m \rangle - \langle \Delta \mathcal{I} \rangle$ and $\langle \mathcal{W}^X \rangle + \langle \Delta \mathcal{I} \rangle$ to zero with increasing	
	timescale. \ldots	115
5.5	Control flow diagrams illustrating various feedback scenarios	116
5.6	The average work done over the course of a trajectory for the double	
	dæmon	118
5.7	The average work done by each oscillator when both oscillators are	
	exploited	119
5.8	Control flow diagrams illustrating various feedback scenarios	120

Acknowledgments

This thesis would not exist had it not been for the wisdom and guidance of my supervisors Stefan Grosskinsky and Rosemary J. Harris. For their patience and kindness throughout my time studying, I owe my deepest thanks. I would also like to thank everybody with whom I've communicated and collaborated over the course of my study, particularly Ian Ford and Sosuke Ito.

I would also like to thank all staff, students and alumni of the Complexity Science DTC for providing a fantastic environment to work in and being like a second family. I would especially like to thank my comrades Arran Tamsett, Yuri Lifanov, Dario Papavassiliou, Ellen Webborn, Pete Dawson, Tom Machon, Stefano Ugliano, Ömer Yüksel and Kārlis Goba; as well as my colleagues Mike Irvine, Ben Collyer, Adam Miller, Daniel Sprague, Yu-Xi Chau, Janis Klaiše, Joel Nicholls, Iris Ren, Tom Rafferty, Joe Hilton, Andreas Pizzoferrato and Paul Chleboun for day-to-day japes, frolics and merry-making.

To my parents, Chris and Julie Maitland, and my sisters, Alice and Hannah, who have both supported and encouraged me throughout, thank you for your understanding, patience and belief in me. Special thanks also to Sonoko Ashida, Matthew Barnard, Adam Thompson, Leo Jofeh, Chai Natan, Minsu Kim, Seong-Young Her, Daniel Halverson and everybody involved in *The Philosopher's Meme* for their friendship and support that filled me with determination.

Finally, I also thank the department coffee machine for its years of service. Gone but not forgotten.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

The work presented in Ch. 4 has been previously published in Maitland et al. [2015]. The model and results presented in Ch. 5 are inspired by a collaboration between myself and Ian Ford and as such are similar to those found in Ford and Maitland [2016].

This thesis was typeset with LATEX $2\varepsilon^1$ by the author.

¹LAT_EX 2ε is an extension of LAT_EX. LAT_EX is a collection of macros for T_EX. T_EX is a trademark of the American Mathematical Society. The style package *warwickthesis* was used.

Abstract

Stochastic thermodynamics provides a clear definition of entropy production and the framework of fluctuation theorems has allowed many general results to be obtained for nonequilibrium systems. Notably, relationships for work and free energy in nonequilibrium systems as well as second law-like inequalities have been obtained. A recent extension to this framework has been the addition of feedback and the generalisation of the previously obtained inequalities to incorporate information quantities related to control and feedback.

In this thesis we contribute to this framework by first providing a close analysis of the nature of 'time-reversal' for feedback systems. Time reversal is a key ingredient in the formulation of the fluctuation relationships from which one obtains the nonequilibrium work functions and so we consider how to meaningfully construct a time-reverse conjugate process for a system with feedback and provide a justification for the Sagawa-Ueda fluctuation relation. We then introduce a simple model of a feedback engine and use it to analyse the fluctuation properties of the information flow between the controlled system and the feedback controller.

Finally, we focus on the possibilities for feedback and consider a model whereby feedback is enacted symmetrically between two coupled systems and find that such a system has entropy-reducing dynamics. Since the dynamics appear to violate the second law of thermodynamics, we comment on their validity and argue that mutual feedback may be unphysical.

Table of notation

Symbol	Notes
χ	State space, arbitrary finite set.
$t \in \mathbb{N}$	Time horizon, $s \in \{0, \ldots, t\}$.
χ^t	Path space (time-evolution trajectories).
X_s	Configuration/State of system at time $s, X_s \in \chi$
\mathbf{X}_t	Stochastic trajectory $\mathbf{X}_t = \{X_0, \dots, X_t\} \in \chi^t$.
\mathbf{x}_t	Trajectory $\{x_0, x_1, \ldots, x_t\}$
$ ilde{\mathbf{x}}_t$	Reversed trajectory. $\{x_t, x_{t-1}, \ldots, x_0\}$, i.e. $\tilde{x}_s = x_{t-s}$.
$\pi_s(x)$	Probability distribution on χ to be in state x at time s .
	$\pi_s(x) = \mathbb{P}_{\pi_0} [X_s = x_s]$ where π_0 is the initial distribution.
\mathbb{P}_{π_0}	Path-space measure on χ^t , also defined by initial distribution
	π_0 and described by transition matrices Ω_s
$\omega_s(x \to x')$	Transition probability from x to x' at time s. $\sum_{x' \in \chi} \omega_s(x \to x')$
	x') = 1
Ω_s	Transition matrices, row sum is unity. $\Omega_s = \{\omega_s(x \to x') :$
	$x, x' \in \chi\}$
$\mu(x)$	Stationary probability distribution on χ of the process \mathbb{P} , if
	it exists.

"The only laws of matter are those which our minds must fabricate, and the only laws of mind are fabricated for it by matter."

> James Clerk Maxwell, Analogies in Nature

Introduction

1.1 The second law

In its simplest form, the second law of thermodynamics dictates that the entropy of an isolated system will increase. Entropy is an easily misunderstood quantity, but can be thought of as a measure of disorder of a system. The more disordered a system becomes, the harder it is to extract useful work from that system. This law is thought to be one of the most important empirical laws in all of physics; not only does it place fundamental limits on the abilities of thermodynamic processes, it is also intimately connected to the notion of directionality in time Halliwell et al. [1996]; Mackey [2003]; Carroll [2010] and allows speculation about the ultimate fate of the universe Kelvin [1862a,b]; Lebowitz [1993].

The origins of the second law are found in the work of the 19th century French scientist Sadi Carnot Carnot et al. [1890]; Prigogine and Kondepudi [1998]. While studying the principles of heat flow in the 1820s¹ Carnot developed the notion of a reversible process. Concluding that reversible processes did not exist in nature, Carnot laid the groundwork for the second law of thermodynamics. Perhaps the best concrete statement was given by Rudolf Clausius Clausius [1850] later that century: "Heat can never pass from a colder to a warmer body without some other

¹Which was relevant at the time for the understanding of steam engines.

change, connected therewith, occurring at the same time". The second law can be very loosely interpreted as the statement that many common phenomena are in fact irreversible. The consequence of this is that one cannot extract work from a system without dissipating energy into the environment, and once energy dissipates into the environment it is lost, which is exactly Kelvin's statement of the second law Kelvin et al. [1851]. Similar statements to the same effect have been made by Planck [1921] and Uhlenbeck et al. [1963].

Since its discovery, the second law has been one of the most empirically verified postulates in all of physics; every event that occurs at the macroscopic level appears to be in agreement with the second law. Eddington even went as far as to say the he believed the second law held a 'supreme position among the laws of nature', going on to say:

"If someone points out to you that your pet theory of the universe is in disagreement with Maxwell's equations, then so much the worse for Maxwell's equations. If it is found to be contradicted by observation, well, these experimentalists do bungle things sometimes. But if your theory is found to be against the second law of thermodynamics I can give you no hope; there is nothing for it but to collapse in deepest humiliation." Eddington [2012]

1.2 Maxwell's dæmon

One might be forgiven for drawing a line underneath the second law and accepting it *prima facie*. However, there have been many challenges to its validity and debates over its interpretation. The most significant of these was put forward in 1867 by the Scottish physicist James Clerk Maxwell. In a letter to Kelvin, Maxwell proposed a thought experiment that considered the microscopic events occurring in a thermodynamic system Kelvin [1879]. The experiment considers a box with adiabatic walls² containing a gas. A barrier is inserted into the box that divides the gas in two. Initially the two sides of the partition are at the same temperature. Located near the barrier is a sentient agent or 'dæmon'³ who is able to open and close a small hole in the barrier. If fast (slow) particles approach from the left (right) box, he lets them through into the right (left) which will cause a increase (decrease) in the temperature of the right (left) box. By monitoring the motion of the gas particles

²No heat enters or leaves the system through the walls.

³Throughout, we will use the spelling $d \ amon$ as it refers to an agent who works in the background. However, it is more common to see the spelling *demon*. Indeed, Kelvin himself used it,

and reacting appropriately, the dæmon is able to transfer heat from one side to the other. This results in a situation whereby heat is being transferred from the colder right side into the hotter left side. This result appears to be in direct contradiction of Clausius' and Kelvin's statements.

The most succinct way of expressing the second law would be to write

$$\langle \Delta \mathcal{S}_t \rangle \ge 0, \tag{1.1}$$

where the angle brackets denote the ensemble average over many realisations. Maxwell's thought experiment points out that on the microscopic scale, events do occur that decrease the entropy. This is reflected by (1.1) which can be read as stating that while the ensemble average of ΔS_t should be greater than zero, individual realisations of ΔS_t can be negative. The thought experiment was not so much a contradiction as a clarification; pointing out the existence of microscopic fluctuations and their potential relevance in special cases. Whilst decreases in entropy are thus theoretically possible, actually achieving a decrease in entropy on average⁴ would require the intervention of a sentient agent capable of witnessing and reacting to microscopic events. At the time of the thought experiment in the late 19th century such a 'neat-fingered being' was an implausibility, and the thought experiment (now dubbed 'Maxwell's dæmon') was not believed to be a contradiction of the 2nd law. However, while first conceived to elucidate the statistical subtleties of the second law of thermodynamics, the experiment has since sparked many debates Szilard [1929]; Brillouin [1951]; Landauer [1961]; Bennett [1982]; Leff and Rex [2010] on how the second law could be systematically violated. That is, violations that manifest in the ensemble average rather than the temporary violations occurring on the level of individual trajectories or events Gallavotti and Cohen [1995a]; Evans et al. [1993]; Evans and Searles [1994].

Much thought has been put into the realisability of real 'Maxwell's dæmons' and their compatibility with the second law that had been the theoretical foundation of much innovation during the industrialisation of Europe. To further clarify the implications of the dæmon, in 1929, the Hungarian physicist Léo Szilárd developed a new concrete model of a 'Maxwell's dæmon' that described exactly the ingredient that permitted the violation Szilard [1929]. In Szilárd's model, we again consider a box with adiabatic walls, however this time the box contains only a single molecule of gas. The system is allowed to reach thermal equilibrium and then a barrier is inserted. An observer then determines which side of the box the molecule is found

despite referring to a hidden agent rather than a malevolent entity.

⁴As opposed to on an individual trajectory.

on, and allows that side of the box to expand adiabatically until it is the same volume as the original box, whereupon the system is once again allowed to equilibrate. By the expansion of the gas, useful work is extracted from the system. However, the system has begun and ended in the same state. Since the change was adiabatic, the free energy change is zero. Through this procedure, the 'Szilárd engine' provides an even more direct confrontation to the second law. Unlike the original Maxwell dæmon, the ability to witness and react to microscopic thermodynamic fluctuations *as they occur* is not required.

Although Szilárd's model at first seems more shocking, it was in fact proposed to make clear a fact that Szilárd claimed had been ignored. In this model, it is far clearer to see the point at which the action of the dæmon influences the system and thus attempt to understand its effect. Noticing that the dæmon's measurement causes the process to branch (depending on which side of the box the particle is found on), Szilárd reasoned that the measurement process plays an important role. Szilárd first noted that the dæmon must have some method of measuring the system state in order to determine how to act. This act of measurement must produce entropy such that the second law was still valid for the whole system. Szilárd essentially argued that one should consider the gas *and* the dæmon as a whole system since they are interacting with one another, and that the second law would hold for the system as a whole.

In 1960, German American physicist Rolf Landauer further refined this understanding by pointing out that the process of information erasure was logically irreversible Landauer [1961], and that it was thus the *erasure* of information that produces entropy. Landauer suggested that having made a measurement of the system state, the dæmon must now *reset* his memory device in order to be prepared for the next measurement, and to achieve that, the dæmon must erase information and thus produce entropy. Since then, it has been further shown Bennett [1982] that even if the dæmon never erases any information from his memory, he will eventually run out of storage space and must either discontinue his activities or erase information he had previously gathered. While both Szilárd and Landauer concluded that the dæmon was not a violation of the second law for different reasons, both acknowledged the importance of information and the processing of information in thermodynamics. Building on the understandings of Maxwell and Boltzmann, the importance of the concept of 'information' to thermodynamics was made clear.

1.3 Information thermodynamics

In the words of Landauer [1996]; "information is physical". Information plays an important role in physics, and this insight has led to the development of novel thermodynamic models such as the 'information heat engine' and 'information refrigerator'. 'Information heat engines' are a class of thermodynamic systems that use information processing to do thermodynamic work without the need for a change in free energy Toyabe et al. [2010]; Park et al. [2013]. Methods such as feedback control allow the engine to use information gained about a physical system to decrease the system entropy and hence extract useful work from the system Lloyd [1997]. This does not constitute a violation of the 2nd law as it is understood that the operation of the feedback device entails an amount of entropy production at least equal and opposite to that change in the system Piechocinska [2000]; Landauer [1961]. The Szilárd engine is a very good example of an information heat engine. By using the information gained from the measurement, the dæmon is able to control the system and extract work. The information resource can be treated thought of as being similar to a 'heat reservoir' from traditional thermodynamics.

Operations involving information and its processing, generation, erasure, encoding, transmission and interpretation, are all processes that play a part in the thermodynamic behaviour of a system⁵. This understanding has allowed information to be included into thermodynamics and led to the development of a new framework of 'information thermodynamics'. The information theoretic description of Maxwell's dæmon thus includes the amount of information obtained by the dæmon. By considering the daemon as operating some control protocol based on the outcomes of his measurement, the damon is studied as an instance of closed-loop feedback control. In keeping with Maxwell's original conception, it is not necessary for the dæmon to expend any energy in operating his feedback control; it is only necessary that he possesses information about the state of the system. There is much research discussing the place of Maxwell's dæmon in physics, both in a modelling context and in experimental systems Jayannavar [1996]; Llovd [1997]; Mandal and Jarzynski [2012]; Mandal et al. [2013]; Barato and Seifert [2013]; Strasberg et al. [2013]; Deffner [2013]; Ford [2016]. Despite his age (and despite numerous attempted 'exorcisms'), Maxwell's dæmon still plays an important role in understanding thermodynamics.

For Maxwell and his contemporaries, it was impossible to observe or control a system on the level of thermodynamic fluctuations. However, due to recent tech-

⁵That is, these processes can influence how heat flows in a system as well as its equilibrium state(s).

nological developments, the thermodynamics of feedback control are increasingly relevant. Modern experiments are capable of probing regimes heavily influenced by thermal fluctuations. In the pursuit of ever faster and ever smaller devices, the understanding of thermodynamic fluctuations is very important. Even more important though, is that information thermodynamics allows us the exciting possibility of exploiting those very fluctuations to our own ends through the construction of real 'Maxwell dæmons'.

In addition to simulating the effect of feedback control on thermal systems; recent experiments have managed to realise the effect of reducing the entropy of a physical system. Using a high-speed camera and phase shifting electric fields controlled by a computer, the Maxwell dæmon effect has been observed operating on a Brownian particle Toyabe et al. [2010]. The role of information and information processing in control of thermodynamic systems has been confirmed in several other experiments for example Lopez et al. [2008]; Koski et al. [2014b,a]; Camati et al. [2016]. Recent theoretical studies have also shown that systems driven by an information engine can outperform chemically driven systems in terms of efficiency Horowitz et al. [2013]. In all of these models, the information thermodynamical framework has given insights into how to construct the control protocols Abreu and Seifert [2011]; Granger and Kantz [2011]; Sagawa and Ueda [2012] and how best to approach the maximum efficiency of these systems. As the realm of technology tends towards very small length and time scales, the possibility of utilising the natural thermodynamic fluctuations in the operation of these devices allows for a wide range of applications in many cutting edge technological and scientific scenarios.

A quantitative relationship between entropy and information is provided by the information thermodynamic framework Barato and Seifert [2014a], which gives the universal upper bound on the mean negative entropy production that can be obtained by feedback control Sagawa and Ueda [2008]. To be precise, the '2nd law of information thermodynamics' states that the entropy production ΔS_t , of a system up to time t is related to the information, ΔI_t , gained by the dæmon via the inequality Touchette and Lloyd [2000]

$$\langle \Delta \mathcal{S}_t \rangle \ge -\langle \Delta \mathcal{I}_t \rangle,$$
 (1.2)

In fact, this turns out to be a simple corollary of the generalised integral fluctuation theorem Sagawa and Ueda [2010] (itself a generalisation of the Jarzynski equality Jarzynski [1997b]),

$$\langle e^{-\Delta \mathcal{S}_t - \Delta \mathcal{I}_t} \rangle = 1. \tag{1.3}$$

This implies that for $\Delta \mathcal{I}_t \neq 0$,

$$\langle e^{-\Delta \mathcal{S}_t} \rangle \neq 1,$$
 (1.4)

that is, the standard integral fluctuation theorem does not hold in the presence of feedback Sagawa and Ueda [2012]. These results have been experimentally verified in small systems, where thermal fluctuations have a strong influence Toyabe et al. [2010]; Bérut et al. [2012]; Koski et al. [2014b]. In previous theoretical studies, the quantity $\Delta \mathcal{I}_t$ and its relation to ΔS_t has been discussed in the context of Langevin equations and continuous-time Markov chains Hartich et al. [2014]; Horowitz and Esposito [2014], and the mutual information between the feedback controller and the stochastic system has been considered for systems with discrete events Horowitz et al. [2013].

There is still much interest in probing and potentially overcoming the second law. Over the last century and a half the 'law' has been repeatedly challenged and revised, to the point where the once immutable and non-negotiable law and its implications are now "more what you'd call 'guidelines' than actual rules" Verbinski [2003]. It is not yet known whether the law can be finally and definitively broken or if further generalisations will continue to rescue and ultimately refine it. Whatever the case, the words of James Gleick seem appropriate;

"'It sometimes seems as if curbing entropy is our quixotic purpose in the universe." Gleick [2011]

1.4 The current work

In this thesis we will discuss recent developments in the physics of feedback. Specifically, we discuss the modifications necessary to the existing stochastic thermodynamics framework in order to accommodate measurement-based feedback. In particular we analyse the notion of a 'reverse process' in the presence of feedback and present a justification for the choice of 'reverse process' used to obtain the Sagawa-Ueda fluctuation relation (1.3). We also test the information thermodynamic framework in two specific models that include feedback. In particular, we use these models to study aspects of feedback that have previously not been considered. Firstly, we study fluctuations of information gained by a measurement device in a feedback system by presenting an analytically solvable model of such a system. We then explore the possible consequences when feedback between system components is cleverly organised.

The thesis is structured as follows. Chapter Ch. 2 is largely an introduction

to notation and framework as well as a review of important concepts and results in the field of stochastic thermodynamics. In this chapter we describe the framework of Markov chains and stochastic thermodynamics. We then detail the notion of 'time-reversal' and what this means for stochastic processes. Having detailed time-reversal we provide an introduction to trajectory functionals and how they are connected to standard thermodynamic quantities such as entropy and work. The most important of these quantities is the 'entropy production' which is a nonequilibrium analogue of entropy and is the primary quantity of interest in nonequilibrium stochastic thermodynamics. Finally we re-derive some standard results from stochastic thermodynamics using our framework and notation. In this chapter we also discuss large deviations and how the theory of large deviations can be used to characterise the fluctuation properties of a system.

Ch. 3 extends the framework from Ch. 2 to include the concepts of 'control' and 'feedback'. We first discuss how systems can be controlled, and how this is included in our earlier notation, before describing closed-loop or 'feedback' control. To do this we use some terminology and ideas from control theory. Since feedback control involves information about a system state being used to manipulate the system evolution, we provide a brief primer on information theory. We describe some information theoretic quantities relevant for discussing measurement-based feedback. We then go on to discuss several possible time-reverse processes that could be used to obtain meaningful trajectory functionals related to the functionals from the previous chapter.

In Ch. 4 we present and study our concrete model of an 'information engine', a system that exploits an information resource to do work using feedback control. This work in this chapter is based on Maitland et al. [2015]. We study the fluctuation properties of the flow of information between the system and the feedback controller by obtaining the large deviation rate function. We obtain an exact expression for the rate function for a two site system. For larger systems we provide an approximate analysis based on numerical data.

In Ch. 5 we study another model of a feedback device that employs feedback control symmetrically between two sub-systems. Systems that use feedback control in this way are speculated to be able to violate the second law and we confirm this with numerical simulation. We comment on the nature of this contradiction and provide a rationalisation of the result. A rigorous treatment of a Langevin model of a 'double dæmon' with similar results to those found here can be found in Ford and Maitland [2016].

Finally in Ch. 6 we summarise and provide an outlook on future research.

We conclude by discussing the relevance of Maxwell's dæmon to thermodynamics and speculating on the future of the second law within physics. "In this house, we obey the laws of thermodynamics!!"

Homer Simpson, The Simpsons, episode 124

2 Key concepts

In this chapter, we will detail the main concepts necessary for the formulation of the stochastic thermodynamic framework before providing an overview of previous results obtained as a form of literature review. We begin by describing the basic framework of Markov chains and introducing our notation in Sec. 2.2. In Sec. 2.3, we discuss the functionals over trajectories that will lead us to obtaining the key fluctuation relations that later chapters will be based on. We then briefly make an aside in Sec. 2.4 to discuss the relationship of these trajectory functionals to classical physical concepts such as entropy and dissipated heat. Sec. 2.5 introduces the fundamental fluctuation relationships by examining the generating functions of the previously defined functionals. Sec. 2.6 details the basics of the theory of large deviations which will be relevant in later sections discussing fluctuation symmetries and the fluctuation theorems in detail and show how they are related to the fundamental fluctuation relationship. Throughout this chapter, in discussing previously obtained results we draw on reviews by Harris and Schütz [2007] and Seifert [2012].

2.1 Stochastic thermodynamics

The primary goal of statistical mechanics is to derive macroscopic and empirically observable laws from microscopic principles. In this thesis, we do not quite start at the microscopic level, but instead satisfy ourself with a mesoscopic description. Specifically, we use a Markovian description of the systems under consideration that abstracts the large number of deterministic microscopic equations into probabilities. This probabilistic description renders all of the processes as *stochastic*, yet still allows us to define, obtain and use familiar thermodynamic quantities and relations. Before we detail this Markovian approach and the abstract stochastic dynamics, it is important to understand what exactly is being abstracted.



Figure 2.1: A cyclic chemical reaction. State E represents a single unbound Enzyme. States ES and EP represent bound states to substrate and product. In completing a cycle clockwise (blue arrows), a substrate is taken from the environment, binds with the enzyme and is converted into a product and then released. The same process can occur backwards, with a product being taken from the environment, binding to the enzyme and converted back to substrate before being released.

As a concrete example, consider the cyclic Michaelis-Menten Michaelis and Menten [1913] chemical reaction with a reversible product generation, shown in Fig. 2.1

$$E + S \rightleftharpoons ES \rightleftharpoons EP \rightleftharpoons E + P,$$
 (2.1)

where E is an enzyme, S is a substrate molecule and P is a product molecule.

ES and EP are complexes. We could consider a single Enzyme in solution with many substrate and product molecules. It is then typical to partition this situation into a 'system' and 'environment'. The 'environment' is the solution containing the substrate and product molecules and has a temperature T. The 'system' is the Enzyme itself as it forms the ES and EP complexes. The 'system' and environment are in thermal contact. In this example, the system and environment can exchange particles since a substrate (product) can be taken from the environment, converted into a product (substrate) and then released.

In general, the generation of a product or substrate involves numerous events that each involve many physical processes. For a single product to be generated, the substrate must come into contact with the enzyme, the two must bind together, changes in the shape of the substrate occur whilst they are bound and finally the changed substrate is released as a product. Rather than considering each of these microscopic processes, we can instead use the following description: The process completes a forward cycle – and converts a single substrate into a product molecule – with rate k_+ and completes a backward cycle with rate k_- . These rates fulfil a generalised detailed balance condition

$$\frac{k_+}{k_-} = e^{-[\mu_P - \mu_S]/k_{\rm B}T},\tag{2.2}$$

where μ_i is the chemical potential of *i* and k_B is Boltzmann's constant. If we assume that the chemical concentrations of *S* and *P* are large, such that the chemical potentials are essentially unchanging on the timescale of observing the process, then the cycle will be driven by the chemical potential difference between the two species. If $\mu_S > \mu_P$, then the process will on average convert substrate molecules to product molecules. This is to say that these rates provide a dimensionless measure of how likely the reaction is to proceed forward as it is backward. The transition rates k_+ and k_- abstract the underlying process yet are still related to physical quantities such as the temperature of the environment and the chemical potential driving the process.

2.2 Notation and framework

Here we outline the basic mathematical framework and notation we shall be using throughout. We consider discrete-time Markov chains on a discrete state space, and use a version of the quantum hamiltonian formalism. While the name alludes to a framework from physics, in reality, this is a somewhat fancy name for computing the evolution of the distribution of a stochastic process via a series of linear difference equations (differential equations in continuous-time systems) and writing these as a matrix operation on vectors in a cleverly chosen basis. We further use the bra/ket notation for notational simplicity, and will attempt to clarify in cases where the notation may be overly condensed. For further reading on the uses of this framework and its connections to field theory, see Mattis and Glasser [1998]; Schütz [2001].

2.2.1 Markov chains

Firstly, we detail the notation for Markov chains and the stochastic dynamics and the types of Markov chains that we are interested in. A Markov chain is a stochastic process that transitions between states in its state space in a 'memoryless' fashion, meaning that the probability to transition to any state depends on on the current state and not any previous states.

State space

We are concerned with dicrete-time Markov processes occurring in some state space χ . In general, the state of the system could refer to the position of a single particle; the configuration of a set of particles or more abstractly, a set of values taken by a set of variables. The system state at time s is $x_s \in \chi$. The probability to find the system in a state x at time s is written $\pi_s(x)$. A particular trajectory of length t is given by $\mathbf{x}_t = \{x_s\}_{s=0}^t$, a chronologically ordered list of the system states. This trajectory \mathbf{x}_t can be thought of as a realisation of the random variable \mathbf{X}_t . The random variable \mathbf{X}_t can be referred to as the stochastic process as it is the family of indexed random variables with a shared state space. The total number of states/configurations visited on a given trajectory is t + 1. The probability of observing a trajectory \mathbf{x}_t is written $\mathbb{P}[\mathbf{X}_t = \mathbf{x}_t]$ where \mathbb{P} is the path space measure. We use a subscript notation to denote the initial condition of the chain e.g. if the starting state of the chain is chosen from some distribution π we use \mathbb{P}_{π} .

A concrete example of this would be a 1D discrete-time random walk (1DDTRW) on a lattice with L sites. The state space is the position of the walker on the lattice, i.e. $\chi = \{1, \ldots, L\}$, and a trajectory would be a list of these sites in the order they are visited.

Stochastic dynamics

The probability of a transition between states x and x' is $\omega(x \to x')$, which forms a transition matrix

$$\Omega_{x',x\in\chi} = \omega(x \to x'). \tag{2.3}$$

For a process with time-dependent transition probabilities, we write

$$\Omega_{s\in\mathbb{N},x',x\in\chi} = \omega_s(x\to x'). \tag{2.4}$$

At each time step, the probability of transitioning to a new state (or staying in the current state) sums to one owing to conservation of probability, expressed as

$$\sum_{x'\in\chi}\omega_s(x\to x') = 1. \tag{2.5}$$

In general, the values of these probabilities are determined by some external parameter we call the 'control' parameter. The evolution of the control parameter itself is then called the 'protocol'. For now it is sufficient to consider that the transition probabilities depend only on time and refer to this time-dependence of the transition probabilities as the protocol. We will revisit the notion of 'control' parameters later in Sec. 3.1.

In all of the following unless explicitly stated, we consider processes that satisfy a weak reversibility condition known as microscopic reversibility (originally proposed by G.N. Lewis as *The law of entire equilibrium* Lewis [1925]; Tolman [1938]). In our notation this condition is written,

if
$$\omega_s(x \to x') > 0$$
 then $\omega_s(x' \to x) > 0 \quad \forall x, x' \in \chi, \ \forall s' \neq s.$ (2.6)

The microscopic reversibility condition states that if a transition can occur from x to x' under the dynamics at time s, then same transition must also be possible backwards – it must be possible to transition from x' to x at time s. This is a 'weak' reversibility as it only implies that the reverse transitions and trajectories are *possible*, i.e. they have non-zero probability of occurring¹. Note that in (2.6) we also consider that if a transition is allowed at some time, then it is allowed at all times. If transitions could become impossible at some specific time, then it would be possible to create processes that do not satisfy (2.17) which is required for a meaningful 'time-reversal' which is covered below in Sec. 2.2.2. Later in Ch. 3, we

¹In fact, many physicists are happy to refer to (2.6) as 'reversibility'. Mathematicians use the term to refer to the far stronger condition of (2.14).

will relax (2.6) but still take care to preserve the possibility of some meaningful 'reversibility'. The requirements we must fulfil are described below in Sec. 2.2.1 and given by (2.16) and (2.17).

Master equation

The master equation for a Markov process describes the time-evolution of the system by quantifying the probability $\pi_s(x)$. The master equation for a discrete-time Markov chain is written as

$$\pi_{s+1}(x) = \sum_{x' \in \chi} \pi_s(x')\omega_s(x' \to x).$$
(2.7)

Note that this equation captures the memoryless nature of the stochastic dynamics of a Markov chain. The state at time s + 1 is determined only by the state at time s and the stochastic dynamics in the form of the transition probabilities.

We use the quantum Hamiltonian formalism to write out the master equation. The quantum Hamiltonian formalism as an approach to stochastic dynamics has a history at least as far back as Glauber [1963]. A comprehensive review of this approach is found in Schütz [2001]. Essentially, this allows us to write the master equation in matrix form by using the bra/ket notation and a suitable choice of basis. Each state of the system x is assigned a vector $|x\rangle$ together with a transposed vector $\langle x|$. The probability distribution can then also be written as a vector

$$|\pi_s\rangle = \sum_{x \in \chi} \pi_s(x) |x\rangle \,. \tag{2.8}$$

Here it should be noted that $|x\rangle$ is a column vector². With a scalar product $\langle x | x' \rangle = \delta_{x,x'}$, the probability can be obtained by the scalar product $\pi_s(x) = \langle x | \pi_s \rangle$.

The master equation is linear, and can be written as

$$\pi_t \rangle = \Omega_{t-1} |\pi_{t-1}\rangle$$
$$= \left(\prod_{s=0}^{t-1} \Omega_s\right) |\pi_0\rangle.$$
(2.9)

In the case of time-independent transition probabilities, Eq. (2.9) reduces to

$$|\pi_t\rangle = \Omega^t |\pi_0\rangle. \tag{2.10}$$

 $^{^{2}}$ Once again, a difference between Physics and Mathematicians appears. The typical approach in Physics is to use column vectors for the state vector and transition matrices with normalised

To compute expectation values, we introduce a summation vector

$$\langle 1| = \sum_{x \in \chi} \langle x|, \qquad (2.11)$$

and so by construction we have $\langle 1 | \pi_s \rangle = 1$, which sums probability over every state, and $\langle 1 | \Omega = \langle 1 |$, which expresses conservation of probability in a discrete-time Markov system as stated above.

Stationarity

For a time-independent transition matrix Ω , the right eigenvector associated with the principal eigenvalue is a probability vector $|\mu\rangle$. This probability vector is known as the stationary distribution and satisfies

$$\left|\mu\right\rangle = \Omega \left|\mu\right\rangle,\tag{2.12}$$

which can be written out explicitly for each state x as

$$\mu(x) = \sum_{x' \in \chi} \omega(x \to x') \mu(x'), \qquad (2.13)$$

A Markov chain on a finite state space has at least one stationary distribution. This is proven by noting that since $\langle 1 | \Omega = \langle 1 |$, then $\langle 1 |$ is an eigenvector with eigenvalue 1. Since this is a left eigenvector there will be a corresponding right eigenvector with non-negative entries also associated to this eigenvalue, which is a result of the Perron-Frobenius theorem since the matrix Ω has only positive entries. This right eigenvector is exactly the stationary distribution of the Markov chain.

For time-dependent transition rates, we can think of the subscript s as an index over different transition matrices Ω_s . As each of these different transition matrices still refer to finite-state Markov chains, there will exist a distribution $\mu_s(x)$ that satisfies $\Omega_s |\mu_s\rangle = |\mu_s\rangle$. In the same way that the s subscript can be thought of as indexing different matrices, in μ_s the subscript s indexes the different stationary distributions associated to the transition matrices Ω_s rather than describing a single distribution that changes in time.

An irreducible³, aperiodic Markov chain with finite state space is ergodic. That is, it has a unique stationary distribution μ and $\pi_t \to \mu$ as $t \to \infty$ and so the

columns. Mathematicians use row vectors and matrices with normalised rows respectively.

³Stated in words, the condition of irreducibility is that every state $x' \in \chi$ should be 'accessible' from any state $x \in \chi$. For a state to be 'accessible', it is simply required that there is a non-zero probability that the system will transition from one state to the other within a finite time or finite

stationary distribution is the limiting distribution of the Markov chain. Irreducibility also tells us that this stationary distribution is unique and satisfies $\mu(x) > 0$ for all x. It should be noted that in general, a system only ever *approaches* stationarity and reaches stationarity only in the infinite-time limit. For time-dependent transition probabilities, the system distribution could be evolving towards a different stationary distribution as each time step and so there is no guarantee that the system will ever relax into any one of those stationary distributions.

Reversibility

Aside from the weak reversibility condition (2.6), we also consider the strong reversibility condition known as detailed balance Maxwell [1867a]; Van Kampen [1992] which is a condition on the distribution μ and the dynamics

$$\mu(x)\omega(x \to x') = \mu(x')\omega(x' \to x). \tag{2.14}$$

The detailed balance condition is stronger than the microscopic reversibility condition (2.6) and implies the latter for finite-state irreducible Markov chains because $\mu(x) > 0$ for all x. Detailed balance implies the existence of a stationary distribution, whereas the existence of a stationary distribution does not necessarily imply detailed-balance.

In the case of time-dependent transition probabilities, we use the term 'timedependent detailed balance' to refer to systems that satisfy

$$\mu_s(x)\omega_s(x \to x') = \mu_s(x')\omega_s(x' \to x) \quad \forall s \in \mathbb{N}_0$$
(2.15)

It is important to note here that even if a time dependent system satisfies (2.15) the distribution of the system will not *necessarily* be the stationary distribution i.e. in general $\pi_s \neq \mu_s$. All of that is to say that the term 'detailed balance' and the related conditions, (2.14) and (2.15), refer to conditions that the transition probabilities must satisfy, not the trajectory or the system.

Ergodic consistency

In order for the fluctuation relations we will obtain in Sec. 2.7 to be well defined, we require that the systems we consider satisfy a necessary⁴ condition known as 'ergodic consistency' Ehrenfest and Ehrenfest [2002]. The condition of ergodic consistency

number of time steps. That is, there must exist a path between any two pairs of states. Absorbing states or partitions of the state space that cannot be transitioned across clearly violate this condition

⁴Necessary but not sufficient by itself.

is written

$$\pi_t(x) > 0 \implies \pi_0(x) > 0 \quad \forall x \in \chi \; \forall t > 0, \tag{2.16}$$

that is, all time-evolved final states at time t are present in the initial distribution with non-zero probability. We also require that

$$\pi_0(x) > 0 \implies \pi_t(x) > 0 \quad \forall x \in \chi \; \forall t > 0. \tag{2.17}$$

For continuous time random walks, (2.17) is satisfied automatically (there is always a probability to not make any jump within any amount of time and hence remain in the starting state). However, for the discrete-time finite-state Markov chains that we will consider, (2.17) is an additional condition on the dynamics we can consider. Consider the 1DDTRW with zero probability to remain in the same state, i.e. $\omega(x \to x) = 0$. In this case, only on even numbered time steps is it possible to find the particle in the same state it started in. The condition is satisfied in the case of a 'lazy' random walk, where there is a probability for the system to remain in the same state at each time step. Conditions (2.17) and (2.16) are not just conditions on the dynamics, it is also a condition on the initial distribution of the system.

As another example, consider a lazy (i.e. $\omega(x \to x) > 0 \quad \forall x$) 1DDTRW random walk with a fixed initial condition $\pi_0(x_0) = 1$. In this case, $\pi_t(x_0) > 0$ as the particle may remain in the same state of any length of time t, but the system may be able to transition to a state x such that $\pi_t(x) > 0$ and $\pi_0(x) = 0$ contrary to (2.16). Alternatively, if $\omega(x \to x_0) = 0 \quad \forall x$, then (2.17) could be violated since the system has zero probability to end in the state it started in. To avoid these cases, we always consider systems for which $\pi_0(x) > 0$ for all $x \in \chi$ and choose our dynamics such that $\pi_s(x) > 0$ for all $x \in \chi$ and $s \in \mathbb{N}_0$.

2.2.2 Time reversal

Fluctuation relations are obtained by comparing a trajectory of some system to some conjugate trajectory that may nor may not be generated by a corresponding conjugate process. While in general one can consider any conjugate process for a fluctuation relation, meaningful physical interpretations usually follow by choosing a process or trajectory that can be considered a time reversal of the process or trajectory under consideration. Here we describe what this notion of time reversal refers to in several concrete senses before continuing.

It is necessary to distinguish between time-reversed trajectories, reversal of the protocol that determines the transition probabilities and time-reversal of the stochastic dynamics (the adjoint dynamics).

Reverse protocol

In the case of time-dependent transition probabilities, we call the specific evolution of the rates with time the 'protocol'. It is possible that we can reverse this protocol to generate a 'backward' protocol (so-called in order to distinguish it from the adjoint process), so that if the transition probabilities in the forward protocol at time sare $\omega_s^F(x \to x')$, then the transition probabilities at time s in the corresponding backward protocol are given by $\omega_s^B(x \to x') = \omega_{t-s-1}^F(x \to x')$ which defines a Markov process with path space distribution \mathbb{P}^B . The -1 in the subscript of the backward protocol transition rate is to ensure that the 'first' jump in the backward process (counting backward from the finishing time t) occurs at the same time as the 'last' jump (counting forward from initial time 0) in the forward protocol. In the forward process there are t transitions occurring at $\{t - 1, \ldots, 0\}$.

Time reversal of a trajectory

For time reversal of trajectories, we simply mean a trajectory with the start and endpoint switched, i.e. if the trajectory is $\mathbf{x}_t = \{x_s\}_{s=0}^t$, then the reversed trajectory is $\tilde{\mathbf{x}}_t = \{\tilde{x}_s\}_{s=0}^t = \{x_{t-s}\}_{s=0}^t$, so that if the state at time s is x_s in the forward trajectory, the state at time s in the reverse trajectory is $\tilde{x}_s = x_{t-s}$.

A corollary of (2.6) and (2.16) is the statement

$$\mathbb{P}_{\pi_0}\left[\mathbf{X}_t = \mathbf{x}_t\right] > 0 \implies \mathbb{P}_{\pi_0}\left[\mathbf{X}_t = \tilde{\mathbf{x}}_t\right] > 0, \tag{2.18}$$

which tells us that if a trajectory is realisable, the time-reversed trajectory is also realisable under the same dynamics. We can use the distribution of the time-reversed trajectory to define a new distribution $\tilde{\mathbb{P}}_{\pi_0} [\mathbf{X}_t = \mathbf{x}_t] = \mathbb{P}_{\pi_0} [\mathbf{X}_t = \tilde{\mathbf{x}}_t]$. If the chain is stationary (begins in the stationary distribution μ), then \mathbb{P}_{μ} describes a Markov chain. Further, if the stationary distribution is reversible (satisfies detailed balance) then $\mathbb{P}_{\mu} = \tilde{\mathbb{P}}_{\mu}$.

Adjoint process

For time-independent transition probabilities, the adjoint dynamics are given by

$$\Omega^{\dagger} = \bar{\mu} \Omega^{\mathrm{tr}} \bar{\mu}^{-1}, \qquad (2.19)$$

where $\bar{\mu}$ is a diagonal matrix with $\mu(x)$ as its diagonal entries. The transition probabilities of the adjoint dynamics are

$$\left[\Omega^{\dagger}\right]_{x',x} = \omega^{\dagger}(x \to x') = \frac{\mu(x')}{\mu(x)}\omega(x' \to x).$$
(2.20)

The adjoint transition rates are normalised which can be checked by summing (2.20) over x'. The adjoint dynamics have the same allowed transitions as the original process and also the same stationary distribution. The adjoint transition probabilities can only be computed if the stationary distribution is known. The adjoint dynamics can be generalised for time-dependent transition probabilities, defining time-dependent adjoint dynamics with time-dependent transition transition probabilities.

In general, one could define any adjoint process by modifying the rates with distributions on the state space, but only using the stationary measure as in (2.20) produces a normalised set of transition probabilities that describes time-reversal of a process (that process being the stationary process). The adjoint dynamics defined in this way describe a time-reversal of the stationary microscopic dynamics of the system, with a distribution \mathbb{P}^{\dagger} . The process described by the adjoint dynamics is a Markov chain, and is stationary with $\mathbb{P}^{\dagger}_{\mu} = \tilde{\mathbb{P}}_{\mu}$. In the case where the system satisfies detailed balance, $\mathbb{P}^{\dagger} = \mathbb{P}$.

2.3 Trajectory functionals

The fluctuation relations we will obtain later in Secs. 2.7 are concerned with probability distributions of trajectory-dependent functionals. In this section, we introduce the type of functionals we are interested in. We then proceed to detail how to obtain several key results in the field of nonequilibrium thermodynamics from these functionals. Our description of these functionals and their connection to fluctuation theorems follows the framework in Harris and Schütz [2007] and uses several key results from Seifert [2012], however we work only in discrete time and convert the expressions found there to our own notation.

2.3.1 Definition of a functional

A functional of a trajectory depends on the complete history of the process, i.e. on each specific state visited. Throughout, whenever considering a trajectory dependent quantity, we will use 'script-style' characters⁵. We consider functionals which

⁵i.e. rather than A, B and C, we use \mathcal{A}, \mathcal{B} and \mathcal{C} .

have explicit dependence only on the initial and final states of the system $(x_0 \text{ and } x_t)$ and on terms depending on pairs of states transitioned between along the trajectory. Specifically, we are interested in functionals of the form

$$\mathcal{R}_F(\mathbf{x}_t, f, g) = \sum_{s=0}^{t-1} r_s(x_s \to x_{s+1}) + \ln\left[f(x_0)\right] - \ln\left[g(x_t)\right].$$
(2.21)

where the functions f and g are arbitrary positive functions acting on χ . The r_s are a collection of functions that map pairs of states on to \mathbb{R} . The whole functional can be thought of as a kind of counting process, where an amount $r_s(x \to x')$ is counted if the system transitions from state x to x' at time s. This could be any quantity (mass, energy, charge etc) that changes when the system transitions between the two states. The boundary terms $\ln f$ and $-\ln g$ may also be related to observable physical quantities. If f and g refer to the same physical quantity then they can be regarded as weights on the initial and final states of the counting process. The above functional is known as the 'forward' functional. A corresponding 'backward' functional is then given as

$$\mathcal{R}_B(\mathbf{x}_t, g, f) = \sum_{s=0}^{t-1} r_{t-s-1}(x_{t-s} \to x_{t-s-1}) + \ln\left[g(x_0)\right] - \ln\left[f(x_t)\right],$$

$$= \sum_{s=0}^{t-1} r_s(\tilde{x}_s \to \tilde{x}_{s+1}) + \ln\left[g(\tilde{x}_t)\right] - \ln\left[f(\tilde{x}_0)\right], \qquad (2.22)$$

where $\tilde{x}_s = x_{t-s}$. This backward functional essentially evaluates the same counting process along the trajectory in the opposite direction in time. We may often suppress the argument notation of these functionals for brevity.

These functionals can be split into two components (taking the forward functional as an example)

$$\mathcal{R}_F(\mathbf{x}_t, f, g) = B(f, g) + \mathcal{J}_t(\mathbf{x}_t), \qquad (2.23)$$

where B and \mathcal{J}_t are referred to as the 'boundary' and 'current' terms. The boundary term is a log ratio of weights attached to the initial and final states of the trajectory, whereas the current term is associated with the dynamics/evolution of the process. In the case of antisymmetric increments, $r_s(x \to x') = -r_s(x' \to x)$, then the current term literally counts some integrated current across whatever spatio-temporal bond that links state x to x'. The subscript t on the current term is included as a reminder that this quantity is extensive in time. In the case where r is not only antisymmetric but also time independent, then the current term can be written

$$\mathcal{J}_t(\mathbf{x_t}) = \sum_{x,x'} r(x \to x') J_t(x \to x'), \qquad (2.24)$$

where $J_t(x \to x')$ is the integrated net number of transitions from state x to x', i.e.

$$J_t(x \to x') = N_t(x \to x') - N_t(x' \to x),$$
 (2.25)

with $N_t(x \to x')$ being the total number of times a transition occurs from x to x' within time t.

For suitably chosen f, g and r, functionals of this form can yield several relationships of interest. We will see later that choosing r to be the logarithm of ratios of transition probabilities is especially useful as it yields functionals that are related to probability measures on trajectory space. If r is the logarithm of the ratio of transition probabilities from the same process (Sec. 2.3.2) or a related process (Sec. 2.3.3 and Sec. 2.3.4) then the functional can take on meaning as a measure of reversibility. We will see later in Sec. 2.4 that using these types of functionals as a measure of reversibility allows us to make a connection to the concept of entropy and entropy production from physics.

2.3.2 Reversal of protocol

Consider that we time-reverse the protocol such that if the rates under the forward protocol are written as $\omega_s^F(x \to x')$, then the rates in the backward protocol are given by $\omega_s^B(x \to x') = \omega_{t-s-1}^F(x \to x')$. We count using the antisymmetric increments

$$r_s(x \to x') = \ln\left[\frac{\omega_s(x \to x')}{\omega_s(x' \to x)}\right],\tag{2.26}$$

for a transition from x to x' at time s. The forward functional then becomes

$$\mathcal{R}_F = \ln\left[\frac{f(x_0)}{g(x_t)}\right] + \sum_{s=0}^{t-1} \ln\left[\frac{\omega_s^F(x_s \to x_{s+1})}{\omega_s^F(x_{s+1} \to x_s)}\right],\tag{2.27}$$

where the first and second terms are the boundary and current terms respectively. The backward functional is written

$$\mathcal{R}_B = \ln\left[\frac{g(x_0)}{f(x_t)}\right] + \sum_{s=0}^{t-1} \ln\left[\frac{\omega_s^B(x_s \to x_{s+1})}{\omega_s^B(x_{s+1} \to x_s)}\right].$$
(2.28)

We note that the forward functional can be written

$$\mathcal{R}_F = \ln\left[\frac{f(x_0)}{g(x_t)}\right] + \sum_{s=0}^{t-1} \ln\left[\frac{\omega_s^F(x_s \to x_{s+1})}{\omega_{t-s-1}^B(x_{s+1} \to x_s)}\right].$$
 (2.29)

If we choose f and g such that they are the initial distribution π_0 then we obtain

$$e^{\mathcal{R}_F} = \frac{\pi_0(x_0)}{\pi_0(x_t)} \prod_{s=0}^{t-1} \frac{\omega_s^F(x_s \to x_{s+1})}{\omega_{t-s-1}^B(x_{s+1} \to x_s)} = \frac{\mathbb{P}_{\pi_0}(\mathbf{x}_t)}{\mathbb{P}_{\pi_0}^B(\tilde{\mathbf{x}}_t)},$$
(2.30)

which will only be well-defined for all trajectories if we have satisfied ergodic consistency (2.16). We can interpret $e^{\mathcal{R}_F}$ as the ratio between the probability of the forward trajectory \mathbf{x}_t under the forward protocol, and the probability of the reversed trajectory $\mathbf{\tilde{x}}_t$ under the backward protocol starting from the same initial condition. If the rates are time-independent and obey detailed balance then the current term of (2.21) reduces to a product of ratios of stationary probabilities. Further, if detailed balance is satisfied and the f and g are chosen to be the stationary distribution, then all the boundary terms and stationary probabilities cancel and $\mathcal{R}_F = \mathcal{R}_B = 0$ for all trajectories. This means that any trajectory is exactly as likely as its time-reversed analogue, which exactly the expectation for a system that is in equilibrium.

2.3.3 Adjoint dynamics

Instead, consider the increments

$$r_s(x \to x') = \ln \left[\frac{\omega_s(x \to x')}{\omega_s^{\dagger}(x \to x')} \right],$$

= $\ln \left[\frac{\omega_s(x \to x')\mu_s(x)}{\omega_s(x' \to x)\mu_s(x')} \right],$ (2.31)

where $\omega_s^{\dagger}(x \to x') = (\mu_s(x')/\mu_s(x)) \omega_s(x' \to x)$. This adjoint process with initial condition π_0 is described by the path space measure $\mathbb{P}_{\pi_0}^{\dagger}$ defined previously in Sec. 2.2.2. We note that for time-independent transition probabilities where $\pi_0 = \mu$ then $\mathbb{P}_{\mu}(\mathbf{x}_t) \equiv \mathbb{P}_{\mu}^{\dagger}(\tilde{\mathbf{x}}_t)$.

The functional counting these increments is

$$\mathcal{K} = \sum_{s=0}^{t-1} \ln\left[\frac{\omega_s(x_s \to x_{s+1})}{\omega_s^{\dagger}(x_s \to x_{s+1})}\right] = \sum_{s=0}^{t-1} \ln\left[\frac{\omega_s(x_s \to x_{s+1})\mu_s(x_s)}{\omega_s(x_{s+1} \to x_s)\mu_s(x_{s+1})}\right],$$
(2.32)

which is the log probability ratio of a trajectory \mathbf{x}_t to the same trajectory under the

adjoint dynamics. That is, we can write

$$\frac{\mathbb{P}_{\pi_0}\left[\mathbf{x}_t\right]}{\mathbb{P}_{\pi_0}^{\dagger}\left[\mathbf{x}_t\right]} = e^{\mathcal{K}}.$$
(2.33)

Note that neither (2.32) nor (2.33) have any terms relating to initial conditions. Since the trajectories are the same and we assume that both systems would be initialised in the same manner (as there would be no reason to weight the initial state differently in one process over the other), the terms for initial distributions cancel in (2.32). We write π_0 as a subscript on both \mathbb{P} and \mathbb{P}^{\dagger} in (2.33) just to remind that the initial condition should be the same for both sets of dynamics. For a system that satisfies detailed balance (i.e. satisfies Eq. (2.14)), we obtain $\mathcal{K} = 0$.

As an aside, we also note here that the transition probabilities in both (2.32) and later in (2.35) are those of the forward process from (2.21), but could also be defined using the transition probabilities from the backward process used in (2.22). We will assume that the forward process and hence forward protocol is the default choice of transition probabilities and that unless otherwise stated, $\omega_s(x \to x') = \omega_s^F(x \to x')$.

2.3.4 Adjoint dynamics with reversal of protocol

We can also combine the notion of time-reversal of the protocol and dynamics by considering the antisymmetric increments

$$r_s(x \to x') = \ln \left[\frac{\omega_s(x \to x')}{\omega_s^{\dagger}(x' \to x)} \right],$$

= $\ln \left[\frac{\mu_s(x')}{\mu_s(x)} \right],$ (2.34)

where again where $\omega_s^{\dagger}(x \to x') = (\mu_s(x')/\mu_s(x)) \omega_s(x' \to x)$. This leads to the functional

$$\mathcal{T} = \sum_{s=0}^{t-1} \ln\left[\frac{\mu_s(x_{s+1})}{\mu_s(x_s)}\right] + \ln\frac{f(x_0)}{g(x_t)},\tag{2.35}$$

which is now the ratio between the probability of a trajectory \mathbf{x}_t in the forward process with a forward protocol and the probability of the same trajectory in a process with time-reversed dynamics and with a reversed protocol.

It is important to stress, that although both (2.32) and (2.35) feature the stationary probability, the distribution of the system itself will not necessarily be stationary, i.e. $\pi_s \neq \mu_s$.
2.3.5 Relationship between functionals

It is interesting to note that the functionals given by (2.27), (2.32) and (2.35) are all related via

$$\mathcal{R}_F = \mathcal{K} + \mathcal{T}.\tag{2.36}$$

This is only mentioned as an aside for now but will become useful in Sec. 2.4.2 when we apply a physical interpretation to \mathcal{R}_F .

2.4 Relation to entropy and heat

In this section we discuss how thermodynamic concepts such as entropy and heat can be defined for athermal stochastic systems. This ultimately follows from the suggestion by Sekimoto [1997, 1998] to use thermodynamic concepts such as the first and second law in order to study stochastic dynamics. For an incremental change in a physical system, the first law reads

$$\mathcal{W} = \Delta U + \Delta \mathcal{Q} \tag{2.37}$$

where ΔW is the work done on the system, ΔU is the change in the internal energy and ΔQ is the heat dissipated to the environment⁶. It is easily seen that the first law is essentially a statement of the conservation of energy. If the internal energy of the system decreases, then the system can either be doing work (i.e. energy is being extracted from the system), dissipating heat into the environment, or both. If the internal energy of the system does not change, then it is possible to simply equate the work done on the system to the heat dissipated

$$\mathcal{W} = \Delta \mathcal{Q}.\tag{2.38}$$

That is to say, for the system to be able to do work, heat must be taken up by the environment. When heat is dissipated into the environment, there is an associated change in the entropy of the environment ΔS_{env} . For a system at temperature T, the dissipated heat is given by

$$\Delta Q = T \Delta S_{\text{env}}.$$
 (2.39)

 $^{^{6}}$ Having adopted here the convention that work applied *to* the system and the heat dissipated to the environment are both positive.

The second law (1.1) was the statement that on average, the total entropy $\Delta S_{\text{tot}} = \Delta S_{\text{sys}} + \Delta S_{\text{env}}^{7}$ will tend to increase. We can see then that to extract work from the system requires ΔQ be negative. Together, (1.1) and (2.39) suggest that to be able to extract work on average requires that we increase the system entropy ΔS_{sys} whilst decreasing the environmental entropy ΔS_{env} . We can extract work from a system as long as it is in a low entropy state. Through the operation of this work-extracting process, we expect to eventually reach an equilibrium whereby the system entropy can no longer be increased on average and thus can no longer withdraw heat from the environment in order to facilitate the performance of work.

In the following subsections, we describe the standard definitions of entropy, heat and work on the level of individual trajectories in the framework of stochastic thermodynamics. We then show later in Sec. 2.7 how these definitions allow us to obtain the second law (and other nonequilibrium quantities) for stochastic systems.

2.4.1 Entropy

Stochastic entropy production is essentially a statement of the 'reversibility' of a system trajectory and hence can be related to the trajectory functionals discussed above. Since it is a trajectory dependent quantity, it can be defined and computed both in equilibrium and nonequilibrium settings, making it a powerful concept. Following the standard definition in stochastic thermodynamics Maes and Netočnỳ [2003]; Seifert [2005, 2012], the total entropy production of a given trajectory is defined (in our notation) as

$$\Delta S_{\text{tot}}(\mathbf{x}_t) \coloneqq \Delta S_{\text{env}}(\mathbf{x}_t) + \Delta S_{\text{sys}}(x_0, x_t), \qquad (2.40)$$

where $\Delta S_{\text{sys}}(x_0, x_t)$ is the change in system entropy and only depends on the initial and final states of the system evolution. $\Delta S_{\text{env}}(\mathbf{x}_t)$ is the entropy change in the environment as the system evolves and is a functional of the system trajectory. The environmental entropy change is made up of contributions from each transition along the trajectory of the system, specifically it is the sum of the entropy production of each transition

$$\Delta \mathcal{S}_{\text{env}}(\mathbf{x}_t) \coloneqq \sum_{s=0}^{t-1} \Delta S_s(x_s, x_{s+1}), \qquad (2.41)$$

with the instantaneous changes in entropy being defined as

$$\Delta S_s = \ln \frac{\omega_s(x_s \to x_{s+1})}{\omega_s(x_{s+1} \to x_s)}.$$
(2.42)

⁷Note that since ΔS_{sys} is not trajectory dependent, we do not use a script-style character.

As an aside, we note that this expression of entropy production also applies to Markov chains with control (Crooks [1999]) by substituting the transition probabilities for their control-dependent counterparts which we will encounter later in Ch. 3. If we choose $r(x \to x')$ equal to (2.42), we associate the current term \mathcal{J}_t of (2.23) with the environmental entropy change.

For the boundary term, we begin by assigning an 'entropy' to the system of the form of the usual Gibbs entropy

$$\langle S_{\rm sys}(s) \rangle = -\sum_{x \in \chi} \pi_s(x) \ln \pi_s(x), \qquad (2.43)$$

where the angle brackets denote the ensemble average. We then make the conceptual leap that at a specific time s, an individual realisation of the system has an entropy contribution quantified by the summand in (2.43)Seifert [2005], i.e.

$$S_{\rm sys}(s) = -\ln \pi_s(x).$$
 (2.44)

If the we choose the weights f and g in the boundary term to be the distributions π_0 and π_t , then we can write

$$B = \ln\left[\frac{\pi_0(x_0)}{\pi_t(x_t)}\right] = -\left(\ln\pi_t(x_t) - \ln f(\pi_0)\right) = \Delta S_{\text{sys}},\tag{2.45}$$

and hence associate the boundary term with the change in system entropy.

The splitting of entropy into system and environment contributions in (2.40) follows from the typical approach used in physics. In physical systems, one often wishes to differentiate between the change in the entropy of the system studied, which is simply the difference between the entropy of its final and initial distribution, and the change in the entropy of the surrounding environment or heat-bath that supplies the energy for the transitions, which will be trajectory dependent⁸.

Combining these choices (setting $f = \pi_0$ and $g = \pi_t$ and choosing transition weights $r_s(x \to x') = \ln [\omega_s(x \to x')/\omega_s(x' \to x)]$) sets (2.21) to be exactly the total change in entropy. In this case, we can identify the boundary term of (2.23) to be the change in system entropy and the current term to be the change in environment entropy.

⁸For example, $\Delta S_{\rm sys}$ could depend on the number of transitions that occur.

2.4.2 Heat

Given (2.39) and that we have identified the current term in (2.21) with ΔS_{env} , we can think of that current-term as the dissipated heat divided by a system temperature T. That is, in direct analogy with thermodynamics, some authors like to write

$$\Delta S_{\rm env} = \frac{\Delta Q}{T}.$$
 (2.46)

For an abstract mathematical system, direct measurement of ΔQ or T is not possible, not do the terms have any clear meaning by themselves. However, their ratio still represents a meaningful quantity.

For functionals of the form (2.32) and (2.35) that include the notion of some conjugate process, we argue along the lines of Oono and Paniconi [1998]; Hatano and Sasa [2001] that the functionals can be identified with components of the heat dissipated (expressed here as entropies)

$$\Delta S_{env} = \Delta S_{hk} + \Delta S_{ex}, \qquad (2.47)$$

where ΔS_{hk} is the 'housekeeping' entropy associated with maintaining a nonequilibrium steady-state and ΔS_{ex} is the 'excess' entropy which results from the timedependent variation of the transition probabilities. By noting that $\mathcal{R} = \mathcal{K} + \mathcal{T}$, we can associate each of the functionals with components of the dissipated heat from (2.47),

$$\mathcal{K} = \Delta \mathcal{S}_{hk}, \tag{2.48}$$

$$\mathcal{T} = \Delta S_{\text{ex}} + \ln \frac{f(x_0)}{g(x_t)}.$$
(2.49)

 \mathcal{K} is associated to the housekeeping heat, since for an equilibrium system, the housekeeping heat is zero, and $\mathcal{K} = 0$ for systems that satisfy detailed balance. \mathcal{T} is then the excess heat that measures the heat dissipated by changing transition probabilities and corresponds to a generalised entropy within the framework of steady state thermodynamics Sasa and Tasaki [2006].

From (2.38) it follows that as long as the internal energy of the system is not changing, we can consider $\Delta S_{env}/T$ to also quantify the work done on the system. That is, if ΔS_{env} , then we could say that the system has had work applied to it, whereas a negative ΔS_{env} is analogous to the system having performed some work. Again, for abstract mathematical systems, the notion of work is not necessarily meaningful, but where a mathematical model is related to a physical system it may be possible to connect this work quantity to physical work done on or by the system being modelled.

2.5 Fundamental fluctuation relationships

To obtain the basic form of a fluctuation relation, we write generating functions for the functionals defined above and note their symmetries, from which the full fluctuation relations in Sec. 2.7 can be obtained. All of the fundamental fluctuation relationships shown here first appeared in Harris and Schütz [2007].

2.5.1 Fundamental relation for reversal of protocol

To obtain the generating function of the functional \mathcal{R}_F , we first must define a the weighted transition matrix $\hat{\Omega}_s(k)$, obtained by multiplying each element of Ω_s , by the corresponding increments (2.26) of the transition that element relates to. That is,

$$\left[\hat{\Omega}_s(k)\right]_{x',x} = \omega_s^F(x \to x')e^{-k\ln\left[\omega_s^F(x \to x')/\omega_s^F(x' \to x)\right]}.$$
(2.50)

The diagonal elements are thus unchanged by this weighting, i.e. $\left[\hat{\Omega}_s(k)\right]_{x,x} = \omega_s(x \to x).$

The generating function of \mathcal{R}_F can then be written

$$\langle e^{-k\mathcal{R}_F} \rangle_F = \langle 1 | \, \bar{g}^k \left(\prod_{s=0}^{t-1} \hat{\Omega}_s(k) \right) \bar{f}^{-k} \, | f \rangle \,, \tag{2.51}$$

where $\langle ... \rangle_F$ denotes the ensemble average in the forward process. \overline{f} and \overline{g} are diagonal matrices with f(x) and g(x) as their diagonal elements and $|f\rangle$ is a column vector with f(x) as its elements.

To measure \mathcal{R}_B in the backward process (the process with time-reversed protocol), we use the same modified transition matrix but indexed backwards, i.e. $\hat{\Omega}_{t-s}(k)$. The generating function for the backward process is

$$\langle e^{-k\mathcal{R}_B} \rangle_B = \langle 1 | \, \bar{f}^k \left(\prod_{s=0}^{t-1} \hat{\Omega}_{t-s}(k) \right) \bar{g}^{-k} \, |g\rangle \,, \tag{2.52}$$

where the terms have the same meanings as (2.51). Note the subscript B on the ensemble average denoting that this average is calculated with regards to the statistical weights associated to trajectories generated by the backward process.

Next, we note that the modified transition matrix satisfies the symmetry

 $\hat{\Omega}_s^{\mathrm{tr}}(k) = \hat{\Omega}_s(1-k)$, which is proved by noting that

$$\begin{bmatrix} \hat{\Omega}_{s}^{\mathrm{tr}}(k) \end{bmatrix}_{x',x} = \begin{bmatrix} \hat{\Omega}_{s}(k) \end{bmatrix}_{x,x'},$$

$$= \omega_{s}^{F}(x' \to x) \exp\left[-k \ln\left[\frac{\omega_{s}^{F}(x' \to x)}{\omega_{s}^{F}(x \to x')}\right]\right],$$

$$= \omega_{s}^{F}(x \to x') \exp\left[(1-k) \ln\left[\frac{\omega_{s}^{F}(x' \to x)}{\omega_{s}^{F}(x \to x')}\right]\right],$$

$$= \omega_{s}^{F}(x \to x') \exp\left[-(1-k) \ln\left[\frac{\omega_{s}^{F}(x \to x')}{\omega_{s}^{F}(x' \to x)}\right]\right],$$

$$= \left[\hat{\Omega}_{s}(1-k)\right]_{x',x},$$
(2.53)

which can then be used to transform (2.51):

$$\langle e^{-k\mathcal{R}_F} \rangle_F = \langle 1 | \bar{g}^k \left(\prod_{s=0}^{t-1} \hat{\Omega}_s(k) \right) \bar{f}^{-k} | f \rangle ,$$

$$= \langle f | \bar{f}^{-k} \left(\prod_{s=0}^{t-1} \hat{\Omega}_{t-s}^{\mathrm{tr}}(k) \right) \bar{g}^k | 1 \rangle ,$$

$$= \langle 1 | \bar{f}^{1-k} \left(\prod_{s=0}^{t-1} \hat{\Omega}_{t-s}(1-k) \right) \bar{g}^{-(1-k)} | g \rangle ,$$

$$= \langle e^{-(1-k)\mathcal{R}_B} \rangle_B.$$

$$(2.54)$$

This then gives us the fundamental fluctuation relationship, rewritten for clarity

$$\langle e^{-k\mathcal{R}_F} \rangle_F = \langle e^{-(1-k)\mathcal{R}_B} \rangle_B. \tag{2.55}$$

2.5.2 Fundamental relation for reversal of dynamics

If we weight the transition matrix with increments (2.31) then we obtain a different weighted matrix $\overline{\Omega}(k)$, then the matrix turns out to have the symmetry property.

$$\bar{\Omega}_s(k) = \bar{\Omega}_s^{\dagger}(k-1), \qquad (2.56)$$

which is proved by

$$\begin{bmatrix} \bar{\Omega}_{s}(k) \end{bmatrix}_{x',x} = \omega_{s}(x \to x') \exp\left[-k \ln\left[\frac{\omega_{s}(x \to x')}{\omega_{s}^{\dagger}(x \to x')}\right]\right],$$

$$= \omega_{s}(x \to x') \frac{\omega_{s}^{\dagger}(x \to x')}{\omega_{s}^{\dagger}(x \to x')} \exp\left[-k \ln\left[\frac{\omega_{s}(x \to x')}{\omega_{s}^{\dagger}(x \to x')}\right]\right],$$

$$= \omega_{s}^{\dagger}(x \to x') \exp\left[\ln\frac{\omega_{s}(x \to x')}{\omega_{s}^{\dagger}(x \to x')}\right] \exp\left[-k \ln\left[\frac{\omega_{s}(x \to x')}{\omega_{s}^{\dagger}(x \to x')}\right]\right],$$

$$= \omega_{s}^{\dagger}(x \to x') \exp\left[-(k-1) \ln\left[\frac{\omega_{s}(x \to x')}{\omega_{s}^{\dagger}(x \to x')}\right]\right],$$

$$= \left[\hat{\Omega}_{s}^{\dagger}(k-1)\right]_{x',x}.$$
(2.57)

The generating function for \mathcal{K} is given by

$$\langle e^{-k\mathcal{K}} \rangle = \langle 1 | \prod_{s=0}^{t-1} \hat{\Omega}_s(k) | f \rangle, \qquad (2.58)$$

We then use (2.57) to obtain

$$\langle e^{-k\mathcal{K}} \rangle = \langle e^{-(k-1)\mathcal{K}^{\dagger}} \rangle,$$
 (2.59)

where \mathcal{K}^{\dagger} is the value of the functional \mathcal{K} when evaluated along a trajectory of the adjoint process.

2.5.3 Fundamental relation for reversal of protocol and dynamics

Similarly to the above, the generating function for \mathcal{T} is given

$$\langle e^{-k\mathcal{T}} \rangle = \langle 1 | \prod_{s=0}^{t-1} \check{\Omega}_s(k) | f \rangle, \qquad (2.60)$$

where the transition matrix $\check{\Omega}_s(k)$ results from weighting Ω_s with increments (2.34), and obeys the symmetry

$$\check{\Omega}_s^{\rm tr}(k) = \check{\Omega}_s^{\dagger}(k-1), \qquad (2.61)$$

proved by noting

ſ

$$\begin{split} \check{\Omega}_{s}^{\mathrm{tr}}(k) \big]_{x',x} &= \left[\check{\Omega}_{s}(k) \right]_{x,x'}, \\ &= \omega_{s}(x' \to x) \exp\left[-k \ln\left[\frac{\mu_{s}(x)}{\mu_{s}(x')} \right] \right], \\ &= \omega_{s}(x' \to x) \frac{\mu_{s}(x')}{\mu_{s}(x)} \exp\left[-k \ln\left[\frac{\mu_{s}(x)}{\mu_{s}(x')} \right] \right] \frac{\mu_{s}(x)}{\mu_{s}(x')}, \\ &= \omega_{s}^{\dagger}(x' \to x) \exp\left[-k \ln\left[\frac{\mu_{s}(x)}{\mu_{s}(x')} \right] \right] \exp\left[\ln\left[\frac{\mu_{s}(x)}{\mu_{s}(x')} \right] \right], \\ &= \omega_{s}^{\dagger}(x' \to x) \exp\left[-(1-k) \ln\left[\frac{\mu_{s}(x')}{\mu_{s}(x)} \right] \right], \\ &= \left[\check{\Omega}_{s}^{\dagger}(1-k) \right]_{x',x}, \end{split}$$
(2.62)

which gives us

$$\langle e^{-k\mathcal{T}} \rangle_F = \langle e^{-(1-k)\mathcal{T}^{\dagger}} \rangle_B.$$
 (2.63)

2.6 Large deviations

In this section we introduce the relevant aspects large deviation theory before detailing how to obtain the large deviation rate function. The theory of large deviations is an extension of Cramer's theorem Cramér [1936, 1938] that details the exponential decay in probability of fluctuations of a stochastic system around the most probably state/trajectory. We provide a rough sketch of the theory for functions of discrete variables. See Dembo and Zeitouni [2009] for a far more rigorous treatment.

2.6.1 Definition of the large deviation principle

Consider a discrete random variable A_t where the subscript t is an index. For example, A_t could be the sum or product of t random variables. Let $\mathbb{P}[A_t = at]$ be the probability that A_t takes some value at. If the limit

$$-\lim_{t \to \infty} \frac{1}{t} \ln \mathbb{P}\left[A_t = at\right],\tag{2.64}$$

exists then we say that $\mathbb{P}[A_t = at]$ satisfies a large deviation principle. In general this limit could be 0 or ∞ in cases where $\mathbb{P}[A_t = at]$ decays to zero with increasing t slower or faster (respectively) than exponentially. The cases of interest are when this limit exists and is different from 0 and ∞ . In these cases we call the resulting function of a the 'rate function' Varadhan [1966]. In general t can be any variable appropriate for taking a thermodynamic limit (e.g. particle number or system size),

but here we will consider t is a number of discrete time steps, hence the label t. We will express the above slightly more loosely as

$$\mathbb{P}\left[A_t = at\right] \sim e^{-E(a)t},\tag{2.65}$$

as t approaches infinity. The rate function is now explicitly included on the RHS as E(a). The rate function is always positive and has a minimum around the concentration points of the distribution of A_t . The rate function tells us about the fluctuation properties of the variable A_t/t , and allows us to quantify how exponentially unlikely a given fluctuation away from the mean is, in the long-time limit.

2.6.2 Obtaining the large deviation rate function

In order to calculate the rate function - provided that the limit (2.64) exists - we first define the scaled-cumulant generating function (SCGF) as

$$\xi_{A_t}(k) \coloneqq \lim_{t \to \infty} \frac{1}{t} \ln G_{A_t}(k), \qquad (2.66)$$

where $k \in \mathbb{R}$ and $G_{A_t}(k)$ is the generating function given by

$$G_{A_t}(k) \coloneqq \langle e^{-kA_t} \rangle = \sum_{at \in \mathbb{R}} e^{-kat} \mathbb{P} \left[A_t = at \right].$$
(2.67)

The Gärtner-Ellis theorem Gärtner [1977]; Ellis [1984] states that if $\xi_{A_t}(k)$ exists and is differentiable for all $k \in \mathbb{R}$, then the process A_t satisfies a large deviation principle of the form of (2.65) with a rate function E(a) that is the Legendre-Fenchel transform of the SCGF

$$E_{A_t}(a) = \inf_{k \in \mathbb{R}} \left[ka + \xi_{A_t}(k) \right].$$
 (2.68)

The Legendre-Fenchel transform is described and explained in Rockafellar [2015].

To see why one can obtain the rate function from the SCGF, let us suppose the process obeys a large-deviation principle such that

$$\mathbb{P}\left[A_t = at\right] \sim e^{-E(a)t},\tag{2.69}$$

and substitute this into the generating function (2.67)

$$\langle e^{-kA_t} \rangle \approx \sum_{at \in \mathbb{R}} e^{-kat} \mathbb{P}(A_t = at) = \sum_{at \in \mathbb{R}} e^{-t[ka + E(a)]}.$$
 (2.70)

To approximate this sum, we consider its largest summand, found by locating the

minimum of ka + E(a). We are able to make this approximation as the error associated with this approximation is the same order as the error in the large deviation approximation itself. Hence, we can write

$$\langle e^{-kA_t} \rangle \approx e^{t \inf_a [ka + E(a)]}.$$
 (2.71)

We then substitute (2.71) into (2.66) to write

$$\xi_{A_t}(k) = \lim_{t \to \infty} \frac{1}{t} \ln \langle e^{-kA_t} \rangle = \inf_a \left[ka + E(a) \right], \qquad (2.72)$$

where we have dropped the 'approximate' notation on the RHS as the limit will mean the terms of (2.65) which are not time-extensive will vanish. We then recognise that the final expression is the Legendre-Fenchel transform of E(a). As the Legendre-Fenchel transform is *self-involutive*, we can write the desired result:

$$E(a) = \inf_{a} \left[ka + \xi_{A_t}(a) \right].$$
(2.73)

Note, that if we had differently defined the generating function (4.21) with a positive 'dummy parameter' k rather than -k, then we would instead locate the maximum of ka - E(a) in (2.70), meaning that the rate function E(a) would be obtained as the supremum of $ka - \xi(a)$ in (2.73).

The theorem allows us obtain the rate function E(a) without direct computation of $\mathbb{P}[A_t = at]$. As long as we can compute the SCGF and show that it is differentiable, then this is sufficient to prove that $\mathbb{P}[A_t = at]$ satisfies a LDP. In cases where the SCGF $\xi(k)$ is not differentiable, then it can be proven that the Legendre-Fenchel transform of the SCGF yields the convex envelope of the rate function E(a) Touchette [2009]. This is only a brief and simplified introduction, for a more complete and rigorous presentation of the Gärtner-Ellis theorem, see Ellis [1995]; Dembo and Zeitouni [2009].

2.6.3 Rate function from modified Markov matrix

By rewriting an expectation of the form (2.67), the SCGF can often be obtained as the logarithm of the principal eigenvalue of the Markov transition matrix which is weighted to count the relevant quantity Touchette and Harris. Consider an additive quantity $\mathcal{R}(\mathbf{x}_t) = \sum_{s=0}^{t-1} r(x_s \to x_{s+1})$ with time independent increments $r(x_s \to x_{s+1})$ measured along a trajectory $\mathbf{x}_t = \{x_s\}_{s=0}^t$. The generating function for this quantity is written:

$$G_{\mathcal{R}}(k) = \langle e^{-k\mathcal{R}} \rangle = \sum_{\mathbf{x}_{t}} e^{-k\mathcal{R}} \mathbb{P} \left[\mathbf{X}_{t} = \mathbf{x}_{t} \right],$$

$$= \sum_{\mathbf{x}_{t}} e^{-k[r(x_{0} \to x_{1}) + r(x_{1} \to x_{2}) + \dots + r(x_{t-1} \to x_{t})]} \pi_{0}(x_{0}) \omega(x_{0} \to x_{1}) \omega(x_{1} \to x_{2}) \dots \omega(x_{t-1} \to x_{t})$$

$$= \sum_{\mathbf{x}_{t}} \pi_{0}(x_{0}) \omega(x_{0} \to x_{1}) e^{-kr(x_{0} \to x_{1})} \omega(x_{1} \to x_{2}) e^{-kr(x_{1} \to x_{2})} \dots \omega(x_{t-1} \to x_{t}) e^{-kr(x_{t-1} \to x_{t})}$$

$$(2.74)$$

where the ω terms are the transition probabilities as above. In this expression we can recognise terms $\omega(x \to x')e^{-kr(x\to x')}$, which we can write as $\hat{\omega}(x \to x')$ and consider as elements of a weighted transition matrix

$$\left[\hat{\Omega}(k)\right]_{x',x} = \hat{\omega}(x \to x'): \qquad (2.75)$$

We can condense the notion now and write the sum as a vector operation (including the summation vector $\langle 1 |$ as defined in (2.11))

$$G_{\mathcal{R}}(k) = \langle 1 | \,\hat{\Omega}(k)^t \, | \pi_0 \rangle \,. \tag{2.76}$$

The initial state $|\pi_0\rangle$ can also be rewritten in terms of the basis described by the matrix $\hat{\Omega}(k)$ i.e. $|\pi_0\rangle = \sum_i a_i |\hat{\psi}_i\rangle$, where a_i are some constants and $|\hat{\psi}_i\rangle$ are the eigenvectors of $\hat{\Omega}(k)$. This allows us to finally rewrite the generating function as

$$G_{\mathcal{R}}(k) = -\sum_{i} a_{i} |\hat{\psi}_{i}\rangle \,\hat{\Omega}(k)^{t} = -\sum_{i} a_{i} |\hat{\psi}\rangle_{i} \,\hat{\lambda}_{i}^{t}, \qquad (2.77)$$

where we have now replaced the matrix with its eigenvalues $\hat{\lambda}_i \operatorname{via} \hat{\Omega}(k) |\hat{\psi}_i\rangle = \hat{\lambda}_i |\hat{\psi}_i\rangle$. We now approximate the sum by the largest contribution

$$G_{\mathcal{R}}(k) = -a_{\max} \left| \hat{\psi}_i \right\rangle \hat{\lambda}_{\max}^t.$$
(2.78)

as all other eigenvalues will approach zero faster than the maximum eigenvalue as t increases. As above, taking the limit $-\lim_{t\to\infty} \frac{1}{t} \ln G_{\mathcal{R}}(k)$ yields the SCGF

$$\xi_{\mathcal{R}}(k) = -\lim_{t \to \infty} \frac{1}{t} \ln G(k) = \lim_{t \to \infty} \left[\frac{\ln a_{\max} |\hat{\psi}_{\max}\rangle}{t} + \frac{\ln \hat{\lambda}_{\max}^t}{t} \right] = \ln \hat{\lambda}_{\max}, \quad (2.79)$$

which is transformed into the rate function via the Legendre-Fenchel transform as described above. That is, we can obtain the rate-function by the Legendre-Fenchel transform of the log largest eigenvalue of a modified Markov transition matrix, weighted to count the relevant quantity \mathcal{R}

2.7 Fluctuation relations

In this section, we discuss how a fluctuation relation (FR) can be classified phenomenologically according to the types of relationship that a functional of the form of (2.21) satisfies before detailing some examples of these various types of FT.

Throughout this section we consider \mathcal{R}_F to be a nondimensional functional taking values in some discrete set \square with distribution $\mathbb{P}[\mathcal{R}_F = R]$ to take some value R in the forward process,

2.7.1 Integral fluctuation relations

A nondimensional discrete functional \mathcal{R}_F with distribution $\mathbb{P}[\mathcal{R}_F = R]$ to take some value R in the forward process, satisfies an integral fluctuation relation (IFR) if

$$\langle e^{-\mathcal{R}_F} \rangle = \sum_{R \in \mathfrak{m}} \mathbb{P} \left[\mathcal{R}_F = R \right] e^{-R} = 1.$$
 (2.80)

This can be easily obtained by setting k = 0 in (2.55) (and dropping the subscript F or B as a similar relation applies in both processes, forward and reverse). The corresponding relationship for \mathcal{R}_B is obtained by setting k = 1 in (2.55). By Jensen's inequality (the convexity of the exponential function implies $e^{\langle A \rangle} \leq \langle e^A \rangle$), (2.80) implies the inequality

$$\langle \mathcal{R} \rangle \ge 0. \tag{2.81}$$

The IFR implies that there could exist trajectories for which the value of \mathcal{R} is negative. We will revisit this shortly in , but for now we note that (2.81) has the same form as the 2nd law of thermodynamics. A relationship of the form of (2.80) always holds for functionals of the type described by (2.21) for any normalised choice of f and g.

Jarzynski equality

Let us take a process that satisfies time-dependent detailed balance (2.15). We initialise the system with $\mu_0(x_0)$ and then allow the system to evolve for t time steps and measure a quantity given by (2.21) where $f(x_0) = \mu_0(x_0)$ and $g(x_t) = \mu_t(x_t)$.

Using a Boltzmann distribution $\mu_s(x) = e^{-U_s(x)}/Z_s$, where $U_s(x)$ is the potential energy of configuration x and Z_s is the partition function, the boundary term of (2.23) becomes

$$\ln [f(x_0)] - \ln [g(x_t)] = \ln [\mu_0(x_0)] - \ln [\mu_t(x_t)]$$

= $\frac{\Delta U}{T} - \Delta(\ln Z).$ (2.82)

We can choose the transition weights $r(x \to x')$ to be equal to (2.42) and associate the current term with the heat dissipated Q as in Sec. 2.4. Using this choice of $r(x \to x')$, the expression for free energy $F = -T \ln Z$ and the first law of thermodynamics $\mathcal{W} = Q + \Delta U$ to relate the heat dissipated to the work done, we can write

$$\mathcal{R}_F = \frac{\mathcal{W} - \Delta F}{T},\tag{2.83}$$

and hence (2.80) in this case is written

$$\langle e^{-\mathcal{R}_F} \rangle = \langle e^{-(\mathcal{W} - \Delta F)/T} \rangle,$$

= 1, (2.84)

or equivalently

$$\langle e^{-\mathcal{W}/T} \rangle = e^{-\Delta F/T},$$
 (2.85)

since ΔF is trajectory independent and can be taken outside of the ensemble average. Eq. (2.85) is known as the Jarzynski equality (JE) Jarzynski [1997b,a]. Superficially, the JE seems similar to a previously obtained result known as the Bochkov-Kuzovlev equality (BKE) Bochkov and Kuzovlev [1977, 1981] which we would write as

$$\langle e^{-\mathcal{W}_0/T} \rangle = 1 \tag{2.86}$$

However, as is explained in Jarzynski [2007], the two results apply in slightly different scenarios as suggested above by the use of W_0 rather than W in the BKE. Whereas (2.85) was originally derived for processes with a time-dependent potential, the Bochkov-Kuzovlev equality applies to systems with a time-independent potential but subject to a time-varying driving force. If the driving force results from the time-varying potential, then both equalities can hold for a given system. While the most obvious difference is that that the JE relates differences in free-energy to the dissipated work, ultimately the difference between the two equalities derives from the fact that the work done is defined differently (noted in Jarzynski [2007] and Horowitz and Jarzynski [2007]).

Integral fluctuation relation for entropy

Let us consider a system with initial distribution π_0 . We measure a quantity defined by (2.21), with f and g chosen to correspond to the initial and final distribution of the process. That is, if $|\pi_t\rangle$ is determined by a master equation (2.9), the boundary part of (2.23) can be written,

$$B(f,g) = \ln \frac{f}{g} = \ln \frac{\pi_0(x_0)}{\pi_t(x_t)}.$$
(2.87)

The current and boundary terms of \mathcal{R}_F then correspond to changes in environment and system entropy c.f. (2.40), we have

$$\mathcal{R}_F = \Delta S_{\text{env}}(\mathbf{x}_t) + \Delta S_{\text{sys}}(x_0, x_t) = \Delta \mathcal{S}_{\text{tot}}, \qquad (2.88)$$

and hence (2.80) becomes an integral relationship for the total entropy change

$$\langle e^{-\Delta S_{\rm tot}} \rangle = 1.$$
 (2.89)

As mentioned earlier, by Jensen's inequality we can obtain a rather suggestive expression from such an IFR. Namely (2.89) implies

$$\langle \Delta S_{\rm tot} \rangle \ge 0,$$
 (2.90)

which is exactly a statement of the 2^{nd} law of thermodynamics. The fluctuation relation is consistent with the 2^{nd} law of thermodynamics as long as the latter is properly interpreted as a statement about averages, not individual trajectories.

Eq. (2.90) is valid regardless of whether the dyanmics satisfy time-dependent detailed balance. If time-dependent detailed balance *does* hold and the system begins and ends in equilibrium, then (2.84) and (2.89) are equivalent.

Housekeeping heat

If we again consider a functional of the form (2.32), but this time consider a process in which time-dependent detailed balance is not satisfied, then the conjugate process is a time-reversal of the dynamics but *not* a reversal of the protocol. In this case, \mathcal{K} and \mathcal{R}_F measure the same quantity and only differ by boundary terms.

Associating \mathcal{K} with the 'housekeeping heat' as before in Sec. 2.4.2, and using (2.56) and (2.58), we can obtain an integral fluctuation relation

$$\langle e^{-\mathcal{S}_{\rm hk}/T} \rangle = 1,$$
 (2.91)

which is the integral fluctuation relationship for housekeeping heat obtained by Speck and Seifert [2005].

Hatano-Sasa fluctuation relation

Again similarly to the above, associating \mathcal{T} to the 'excess heat' as in (2.49) allows us to obtain the IFR

$$\langle e^{-\mathcal{S}_{\text{ex}}/T + \Delta S_{\text{sys}}} \rangle = 1,$$
 (2.92)

which is known as the Hatano-Sasa relation Hatano and Sasa [2001]. This is most often considered to be a extension of the 2^{nd} law for stochastic systems as it details the expected behaviour of the dissipated heat and the system entropy which are exactly the quantities originally considered by Carnot in his formulation of the 2^{nd} law.

2.7.2 Detailed fluctuation relations

The detailed fluctuation relation (DFR) represents a stronger relationship for \mathcal{R}_F , namely

$$\frac{\mathbb{P}^B_{\pi_t}\left[\mathcal{R}_B = -R\right]}{\mathbb{P}_{\pi_0}\left[\mathcal{R}_F = R\right]} = e^{-R},\tag{2.93}$$

where $\mathbb{P}^B_{\pi_t}$ is the process with a reversed protocol seen earlier in (2.30). It is trivial to show that a DFR implies the corresponding IFR. The DFR and IFR represent constraints on the distribution $\mathbb{P}[\mathcal{R}_F = R]$. The DFT provides the easiest way to explain why the fluctuation relation and (2.90) are consistent with the 2nd law of thermodynamics: while trajectories with negative entropy production *are* possible, they are exponentially less likely than trajectories with positive entropy production.

We note that a very broad class of nonequilibrium dynamics satisfy a detailed fluctuation theorem of the form (2.93) with a functional similar to (2.40) Jarzynski [2000]; Evans and Searles [2002].

Further statistical properties of $\mathbb{P}[\mathcal{R}_F = R]$ can be derived from the validity of DFRs and IFR as shown in Merhav and Kafri [2010].

The relation (2.93) can be obtained from the generating function symmetry (2.55) by rewriting the latter as

$$\sum_{R \in \mathbb{n}} \mathbb{P} \left[\mathcal{R}_F = R \right] e^{-kR} = \sum_{R \in \mathbb{n}} \mathbb{P}^B \left[\mathcal{R}_B = -R \right] e^{(1-k)R}, \tag{2.94}$$

and then noting that to be satisfied for all k requires (2.93).

Evans-Searles fluctuation relation

If we consider a time-independent protocol, the forward and backward process are identical, i.e. $\mathbb{P} = \mathbb{P}^B$. If we further set $f = g = \pi_0$ in (2.21) and (2.22), then we have $\mathcal{R}_F = \mathcal{R}_B = \mathcal{Z}$, where \mathcal{Z} is the 'Dissipation function' of Evans and Searles Evans and Searles [1994]; Searles and Evans [1999]. In this case (2.93) is written as

$$\frac{\mathbb{P}_{\pi_0}\left[\mathcal{Z}=-z\right]}{\mathbb{P}_{\pi_0}\left[\mathcal{Z}=z\right]} = e^{-z},\tag{2.95}$$

which holds for any choice of initial distribution π_0 Evans and Searles [2002].

Steady state fluctuation relation

A special case of the Evans-Searles FR is obtained when we set $f = g = \mu$, i.e. the system is allowed to approach stationarity before the functional is evaluated. As above in 2.7.1, \mathcal{R}_F can be identified with the total entropy change of the system and we obtain the stronger DFR for total entropy production in the steady state

$$\frac{\mathbb{P}_{\mu}\left[\Delta S_{\text{tot}} = -S\right]}{\mathbb{P}_{\mu}\left[\Delta S_{\text{tot}} = S\right]} = e^{-S}.$$
(2.96)

This is a stochastic form of the original FR proposed in Evans et al. [1993]. It has been proven for systems with chaotic dynamics Gallavotti and Cohen [1995a], and for stochastic diffusive systems Kurchan [1998]; Lebowitz and Spohn [1999]. In these studies, the SSFT strictly only applies in the long-time limit as only the environmental entropy change is considered. With the inclusion of the system entropy change, the SSFT holds for finite times in the steady state Seifert [2005].

Crooks fluctuation relation

Now we consider a process that begins and ends in stationary states μ_0 and μ_t where $\mu_0 \neq \mu_t$, with time-dependent transition probabilities that satisfy time-dependent detailed balance (2.15). We set $f = \mu_0$ and $g = \mu_t$. As before with the JE (2.85), making these choices for f, g and choosing $r_s(x \to x') = \ln \left[\frac{\omega_s(x \to x')}{\omega_s(x' \to x)}\right]$ allows us to identify \mathcal{R}_F as the work dissipated, i.e. $\mathcal{R}_F = \mathcal{W} - \Delta F$. Since the initial and final stationary states are not necessarily the same, there can be a nonzero free energy difference between Eq. (2.93) becomes

$$\mathbb{P}^{B}\left[\mathcal{W} - \Delta F = -w_{d}\right] = \mathbb{P}\left[\mathcal{W} - \Delta F = w_{d}\right]e^{-w_{d}},\tag{2.97}$$

which is the Crooks nonequilibrium work relation for systems with a free-energy difference Crooks [1999, 2000]. We here stress again that although the process begins and ends in stationary states, this will not be true for s > 0, i.e. the system will not be in a stationary state at intermediate times during the evolution. The Crooks nonequilibrium work relation allows more reliable calculation of ΔF than the JE in many applications. Since \mathbb{P}^B is normalised, (2.97) implies the JE, which can be seen by taking the ensemble average over both sides.

2.7.3 Crooks-type fluctuation relations

The set of relationships known as Crooks-type fluctuation relations compare the distribution of the process $\mathbb{P}[\mathcal{R}=R]$ with the distribution of the same quantity \mathcal{R} in some some conjugate process. We drop the subscript on \mathcal{R} as the same functional is evaluated in both the forward and conjugate process. The conjugate process most frequently picked is the time-reversed process obtained from the adjoint dynamics. Taking the adjoint dynamics as an example, the statement of the Crooks-type fluctuation relation is

$$\mathbb{P}^{\dagger}\left[\mathcal{R}=R\right] = \mathbb{P}\left[\mathcal{R}=R\right]e^{-R},\tag{2.98}$$

which implies an IFR for \mathcal{R} since \mathbb{P} and \mathbb{P}^{\dagger} are both normalised. These IFRs can also be obtained from symmetries of the generating function as discussed above. The Crooks-type FR does not imply a DFR for \mathcal{R} since \mathbb{P}^{\dagger} refers to a different process. One can consider (2.97) to be a relation of this type where the conjugate process is the time-reversed protocol.

2.7.4 Fluctuation relations in the long-time limit

One final calss of fluctuation relation worth noting is the class of FRs that hold in the long time limit, as $t \to \infty$. Recall from Sec. 2.3.1 that the functional \mathcal{R}_F (and of course \mathcal{R}_B) can be split into two contributions: a boundary term \mathcal{B} and a current term \mathcal{J}_t . If we pick $f = g = \pi_0$ and r to be antisymmetric transition weights, then \mathcal{R}_F is the dissipation function \mathcal{Z} as in Sec. 2.7.2 and can be split into two terms,

$$\mathcal{Z} = \mathcal{J}_t + \ln \frac{\pi_0(x_0)}{\pi_0(x_t)}.$$
(2.99)

Note that the first term on the right is extensive in time, whereas the second term is not. One would expect that the current-term would grow linearly with increasing time, and hence we might also expect that in the limit as $t \to \infty$, one can neglect

the boundary term altogether. In that case, one could rewrite (2.95) as

$$\frac{\mathbb{P}\left[\mathcal{J}_t = -jt\right]}{\mathbb{P}\left[\mathcal{J}_t = jt\right]} \sim e^{-jt}.$$
(2.100)

Recalling from Sec. 2.4 that \mathcal{J}_t can be identified with the environmental entropy production, Eq. (2.100) is the stochastic equivalent of the Gallavotti-Cohen fluctuation theorem found in Gallavotti and Cohen [1995a,b] for the steady-state of deterministic systems (in which entropy is related to the contraction/expansion of phase-space).

2.8 Conclusion

We have now covered all of the concepts that are required to further discuss stochastic dynamics for feedback systems. Specifically we have defined trajectory functionals and described how these functionals are related to thermodynamic concepts and derived some known results. However, we have so far only discussed 'autonomous' systems. In Ch. 3 we will show how to extend the above framework for systems subject to external control. We will then discuss 'feedback', whereby the specific control protocol is chosen in response to the specific trajectory of the system. "We have modified our environment so radically that we must now modify ourselves in order to exist in this new environment."

> Norbert Wiener, The Human Use of Human Beings

Control and feedback

Having detailed the some established fluctuation relations, in this chapter we introduce the notion of feedback control and consider how to treat controlled systems. In Sec. 3.1 we introduce concepts from the branch of engineering and mathematics known as 'control theory' and integrate this into the framework detailed in Ch. 2. Before we continue, Sec. 3.2 completes the literature review content by detailing some basic information theory, the relevance of which becomes clear nearer the end of the chapter. In Sec. 3.3 we discuss how to introduce functionals for feedback systems. Specifically, we want to introduce functionals that can be related to timereversal as before. We discuss the various ways in which the conjugate process can be formulated for feedback systems, with comments on their use and applicability, including the celebrated Sagawa-Ueda fluctuation relationship.

3.1 Control theory

Control theory is a framework that straddles a boundary between engineering and mathematics. Some of the fundamental concepts appear as early as the 1800s in the works of Airy and Maxwell (as described in Fuller [1976]). In particular, Maxwell discussed extensively the feedback effects of a steam-regulating device known as the Watt's governor Maxwell [1867b], attempting to create a framework for sensibly discussing such devices. Many of those ideas were later expanded to apply to general linear systems Routh [1877]; Hurwitz [1964]. Since then, concepts relating to control of systems have been established as their own field, including the area of 'cybernetics', developed by Norbert Wiener to describe the constraints and possibilities for regulated or controlled systems Wiener et al. [1948].

By identifying the 'inputs' and 'outputs' of a system, a framework for control can be devised whereby the system output can be controlled to match some desired target called the 'reference'. In feedback controlled systems, the output is monitored and the difference between the actual output and the reference used to vary the input so as to match the output to the reference. Contemporary control theory largely deals with the effects of control or feedback on dynamical systems but also encompasses A number of other topics. These include study of the efficiency of effectiveness of control, known as the 'controllability' of a system Ogata and Yang [1970]; Isidori [1995]; the 'observability' of a system, to what extent those are representative of the internal system state Kalman [1959]; and the stability of the control, detailing the effect of fluctuations on the output and controllability of a system.

A strong bridge between control theory and thermodynamics, as well as a pre-saging of one of the key results of information-thermodynamics is found in the consideration of the effect of feedback on the distribution of the system in state space Touchette and Lloyd [2000]. Control techniques are most often expected to match the system output to the reference from any general state. That is, although the system can begin anywhere in state space, the objective of the control is to reach some pre-determined target state (a subset of the total state space); this achieves a reduction of entropy, as the uncertainty on the system state should be decreasing in time as it approaches the target state. Conversely, control might be used to increase the entropy of a system, by altering the control variables to allow a system to evolve from a known initial state to a randomised final state.

3.1.1 Control protocols

A control protocol is a procedure for changing the values of control parameters of a system. Those control parameters are a subset of the 'inputs' of a given system¹ and can be altered over the course of the system evolution. Here we discuss two methods of control. Open-loop and closed-loop control methods for generating and applying a control protocol.

Bayesian networks (introduced by Neapolitan [1990]; Pearl [1985]) are graphs

¹The other input being the initial condition.



Figure 3.1: Bayesian causal diagram of a Markov chain. The nodes represent states of the system. Edges between vertices represent causal relationships between those states. For example, here, the state at time s, x_s , is causally influential over the state at time s + 1, x_{s+1} . The arrow is labelled with the transition probability, as the transition is the mechanism through which x_s causes X_{s+1} to take the value x_{s+1} , and $\omega_s(x_s \to x_{s+1})$ is the conditional probability of this transition.

used to represent the conditional relationships between variables in a model. Each vertex on the graph represents a variable and the directed edges describe conditional dependence between them. Bayesian networks are always directed acyclic graphs (DAGs) if they represent causal relationships between variables. An undirected graph has no clear 'causal' relationships and a graph that is not acyclic would not satisfy the local Markov property Russell et al. [1995]. The local Markov property states that each variable should be conditionally independent of any variables that are not causally influenced by it. Graphs that are cyclic or undirected can still be used in the study of stochastic processes, although they are called 'Markov random fields' Kindermann et al. [1980]. An example of a Markov random field is the Ising model.

To map a Markov chain as a Bayesian network, each state is represented as a node. Edges connect successive variables with the direction of causation running from past to future states. In the case of the Markov chains discussed above, the vertices represent the system states along a sample path and the directed edges represent the dependency of the states on one another. Successive states are related via the transition probability, which can be seen as a kind of conditional probability; given that $X_s = x_s$ the probability that $X_{s+1} = x_{s+1}$ is $\omega(x_s \to x_{s+1})$. i.e. $\mathbb{P}[X_{s+1} = x_{s+1} | X_s = x_s] = \omega(x_s \to x_{s+1})$. Fig. 3.1 shows a generic section of a causal diagram for a specific realisation of a Markov chain. For specific paths, the graph is a 'tree-like' graph, with a unique path between any two vertices.

3.1.2 Open-loop control

In the previous chapter we discussed the possibility for the evolution of a system to be governed by a variable known as the 'control parameter'. Up until now we have discussed a time-dependent protocol, whereby the transition probabilities are time-dependent and thus the control parameter is just the running time s. This is an example of what is known as 'open-loop' control Kuo [1981]. This refers to the fact that the system is 'open' to some external interference. The control parameter evolves completely independently of the controlled system. An example of open-loop control is a washing machine. The washing machine is programmed to spin and rest at pre-determined times regardless of the state of the clothes it is washing. Clean and dirty objects placed in the machine are subject to the same cycle regardless of their input state or indeed their state at any instant in the process.

The process represented by Fig. 3.1 could be thought of as an example openloop control. The control protocol here is simply the time-dependence of the transition probabilities which evolve independently of the system state. To be more explicit and include an external control protocol into the system, we now extend our notation. First, we restrict our interest to time-independent transition probabilities. In analogy to the system state and trajectory we define a variable C which takes values $c \in \gamma$ which is the control protocol. The full control protocol is written in the same manner as a trajectory, i.e. $\mathbf{c}_t = \{c_s\}_{s=0}^t$. We now write the transition probabilities as

$$\omega(x \to x' \mid c) = \mathbb{P}\left[X_{s+1} = x_{s+1} \mid X_s = x_s, C_s = c_s\right].$$
(3.1)

Since the control parameter is used as an extra condition in the probability we carry over this notation to the transition probability. However it is important to remember that this notation only reminds us that the transition probability is dependent on the value of the control parameter and has no other special mathematical meaning.

The probability of a trajectory through the system is then written as

$$\mathbb{P}\left[\mathbf{X}_{t} = \mathbf{x}_{t}, \mathbf{C}_{t} = \mathbf{c}_{t}\right] = \mathbb{P}\left[\mathbf{X}_{t} = \mathbf{x}_{t} \mid \mathbf{C}_{t} = \mathbf{c}_{t}\right] \mathbb{P}\left[\mathbf{C}_{t} = \mathbf{c}_{t}\right]$$
$$= \mathbb{P}\left[\mathbf{C}_{t} = \mathbf{c}_{t}\right] \pi_{0}(x_{0}) \prod_{s=0}^{t-1} \omega(x_{s} \to x_{s+1} \mid c_{s})$$
(3.2)

where $\mathbb{P}\left[\mathbf{C}_{t}=\mathbf{c}_{t}\right]$ is the distribution of control protocols.

As mentioned earlier in Sec. 2.2.1, in order to be able to define meaningful fluctuation relations, we must take care to preserve the validity of (2.16) and (2.17). To do this, we generalise (2.6) to read

$$\omega(x \to x' \mid c) > 0 \implies \omega(x' \to x \mid c) > 0 \quad \forall x, x' \in \chi \; \forall c \in \gamma,$$
(3.3)

which allows us to make a given transition impossible for some value of the control

parameter without breaking the conditions of ergodic consistency.

3.1.3 Closed loop control

For systems with some form of feedback, we talk of 'closed loop' control Aström and Murray [2010], whereby the evolution of the control parameter is dependent on the evolution of the controlled system. The most explicit context for this is measurement-based feedback. Consider a system that evolves according to some transition probabilities which are determined by a control protocol. This control protocol is in turn determined by some measurement performed on the system state/trajectory, with the measurement history given as $\mathbf{y}_t = \{y_s\}_{s=0}^t$. Without loss of generality, we will always assume that \mathbf{X}_s and \mathbf{Y}_s both take values in the same state space χ . That is, if the variable X_s represents the system state at time s, then y_s also the state that the measurement device believes² the system to be in. As a contrast to the example given above of an open-loop controller, consider an air-conditioner (AC) as an example of a closed-loop controller. A (sufficiently sophisticated!!) AC will attempt to match the temperature of its environment to a pre-set 'reference' point. The AC can measure the current temperature of its environment and then output heated or cooled air in response. The control protocol of the AC can then be different depending on the specific evolution of the room temperature i.e. the same AC will perform different protocols in a room filled with ice and a room filled with coal-burners.

As an aside, we note that feedback control can be notationally and conceptually treated in several different ways. Throughout, we shall consider feedback whereby the value of the control protocol is uniquely determined by the measurement outcome. That is, the feedback controller maps measurement outcomes to control actions so c = c(y). We condense the notation for transition rates and write $\omega(x \to x' \mid c(y)) = \omega(x \to x' \mid y)$. In general, it is possible to consider stochastic control such that the value of the control parameter is related to the measurement outcome through some conditional probability $\mathbb{P}[C = c \mid Y = y]$. However, in our case this generalisation is not necessary and any randomness in choosing a control action could be absorbed into the measurement process. That is, we could write $p(c \mid x) \coloneqq p(c \mid y)p(y \mid x)$ and then call c our measurement variable and look directly at the correlation between x and c.

In systems where the state of the control system is dependent on the state of

²We use the loaded word 'believe' here, although we really mean a kind of naive guess based solely on the outcome of the current measurement. If all you know is that your measurement device is in state y_s then you must rationally conclude that the system is most likely in state y_s .



Figure 3.2: Bayesian causal diagram of a single repeated feedback process. Here, the nodes representing the system state are connected to not only the node representing the previous state, but also the node that represents the control parameter value, which is itself causally dependent on the previous state.

the controlled system, and in turn the evolution of the controlled system is dependent on the control system, the feedback is said to form a 'closed-loop', as no reference to any event outside the system is needed to determine the control protocol. Fig. 3.2 shows the Bayesian causal diagram for this kind of system. Here the control parameter C_s is taken to be exactly the variable Y_s , which depends on X_s through a process characterised by the conditional probability $p(y_s | x_s) := \mathbb{P}[Y_s = y_s | X_s = x_s]$. In the diagram, the control protocol is explicitly included as a set of red nodes. These red nodes have two edges representing their causal relationships. The first is a green arrow that represents the conditional probability p(y | x). We will call this the 'measurement' process, as Y takes a value that is conditionally dependent on the value of X, and hence the outcome of Y could be used to make an inference about the value of X.

The purple arrow denotes the influence of the control parameter on the transition probability. This arrow reflects the fact that the outcome of Y_s is a causal determinant of the outcome of X_{s+1} . In order to obtain an outcome of X_{s+1} , it is necessary that X_s and Y_s are determined, which is reflected by the direction and placement of the arrows in Fig. 3.2. The probability of an individual trajectory is thus written

$$\mathbb{P}\left[\mathbf{X}_{t} = \mathbf{x}_{t}, \mathbf{Y}_{t} = \mathbf{y}_{t}\right] = \pi_{0}(x_{0})p(y_{0} \mid x_{0})\prod_{s=0}^{t-1}\omega(x_{s} \to x_{s+1} \mid y_{s})p(y_{s} \mid x_{s}).$$
(3.4)

In a given time step, we have $\omega(x \to x' \mid y)p(y \mid x)$ for closed-loop as opposed to $\omega(x \to x' \mid c)p(c)$ for open-loop control.

Clearly, if we consider the composite system formed of both X and Y, i.e. Z = (X, Y), then we can obtain a Markov chain $\mathbf{Z}_t = \{X_s, Y_s\}_{s=0}^t$. In this section we will not do this as it essentially circumvents the interesting facet of the system, namely that the system is bi-partite or even multi-partite in general Hartich et al. [2014]. The conditional dependency of the interaction of the two subsystems is precisely the behaviour of interest, and the terms used to quantify this interaction have some significance when written explicitly, rather than hidden by the construction of a new Markov chain on a larger state space.

Impacts on previous concepts

It is important to note that in this section we have not introduced any concepts that fundamentally alter anything we have previously considered. In Sec. 2.2.1, we considered that the control parameter was just the time variable s and that the control protocol was the time-dependence of the transition probabilities. All we have done in this section is generalised the notation to account for general control parameters. This is made clear by realising that in Sec. 2.2.1, we could have written the time-dependent transition probabilities as $\omega(x \to x' \mid s)$ instead of $\omega_s(x \to x')$.

For a fixed control protocol \mathbf{c}_t , the fundamental fluctuation relationships from Sec. 2.5 still hold and the connection to thermodynamic entropy made in Sec. 2.4 still holds, with the stochastic entropy production for a controlled system being given by

$$\Delta S_s = \ln \frac{\omega(x_s \to x_{s+1} \mid y_s)}{\omega(x_{s+1} \to x_s \mid y_s)}.$$
(3.5)

3.2 Information theory

In this section we give a brief introduction to some key concepts from information theory that will be relevant near the end of this chapter. First, we introduce the concept of 'Shannon entropy' and 'mutual information', before describing a measure of information known as the 'change in uncertainty' and how it relates to the mutual information. Specifically, we will describe why information theoretic concepts will be



Figure 3.3: The Shannon entropy H(X) plotted against p. H(X) is maximised at p = 1/2 and is zero when p = 0 or p = 1.

useful in analysing the behaviour of the measurement device, as they are intuitively applicable to the framework of measurement. A more thorough derivation of these concepts is found in Cover and Thomas [2012] which provides a general and near exhaustive overview of information theory as a whole, as well as its applications. Here we consider only discrete random variables, although all of the relations in this section have analogues for continuous random variables. Mostly, the sums are replaced with integrals, i.e. $\sum_{x \in \mathbb{Z}}$ for an integer valued variable would become $\int_{x \in \mathbb{R}} dx$ for a real valued variable.

3.2.1 Shannon entropy

The Shannon entropy, named Claude Shannon and also known as the Shannon information, measures the disorder or randomness of a random variable. Consider a random variable X that takes integer values, i.e. $x \in \mathbb{Z}$. We write as shorthand $P(x) \coloneqq \mathbb{P}[X = x]$, for the probability that x is realised.

The particular form for the Shannon information was chosen by Shannon to satisfy several criteria; information is non-negative; events that always occur do not contain information; information due to independent events should be additive. The 'information content' associated with outcome x is $-\ln P(x)^3$. The Shannon Information is the average information content associated with observing outcomes of the random variable X Shannon [1948]; Jaynes [1957]. That is, the Shannon information is given by the expectation value of the information content

$$H(X) \coloneqq -\sum_{x \in \mathbb{Z}} P(x) \ln P(x).$$
(3.6)

³This really could be argued to be the key idea. Events that occur often carry little information.

The Shannon information satisfies $0 \le H(X) \le \ln N$, where N is the number of possible outcomes (i.e. it is the size of set X). The lower bound is obtained if some outcome of X has P(x) = 1, whereas the upper bound is satisfied when P(x) = 1/N.

As an example of how the Shannon information measures 'randomness', consider the example of a variable X which takes values a or b. We set P(a) = p and thus P(b) = 1 - P(a) = 1 - p. The Shannon information is then H(X) = H(p) = $-p \ln p - (1 - p) \ln(1 - p)$. If p = 1 or p = 0 then H(X) = 0 and X contains no information and is not random, since X only ever gives a single outcome. When either outcome is equally as likely then p = 1 - p = 1/2 and $H(X) = \ln 2$. Hence, when X is as random as it can possibly be, its Shannon information achieves a maximum. Fig. 3.3 shows how H(X) varies with p.

3.2.2 Mutual information

If the Shannon entropy quantifies the randomness of a variable, then the mutual information quantifies the correlation between two random variables. To see this, we take two random variables X and Y (which take values $x \in \mathbb{Z}$ and $y \in \mathbb{Z}$ respectively) with joint distribution $P(x, y) = \mathbb{P}[X = x, Y = y]$ from which we can obtain marginal distributions via $P(x) = \sum_{y \in Y} P(x, y)$ and $P(y) = \sum_{x \in X} P(x, y)$.

In the case of independent random variables, the joint probability is P(x, y) = P(x)P(y). In this case, the conditional probability becomes P(x|y) = P(x,y)/P(y) = P(x), meaning that one cannot obtain any greater certainty on the value of X by observing an outcome of Y. An example of this would be independent coin-flips; one cannot guess the outcome of one coin-flip by looking at the outcome of the other. On the other hand, if X and Y are perfectly correlated then really only of the variables is 'necessary'; if the outcome of Y is known, then the outcome of X is also known exactly.

For multiple random variables, the Shannon entropy is extended to the joint Shannon entropy by

$$H(X,Y) \coloneqq -\sum_{x,y \in \mathbb{Z}} P(x,y) \ln P(x,y).$$
(3.7)

The Shannon entropy for the conditional probability $P(x \mid y)$ is given by

$$H(X \mid y) \coloneqq -\sum_{x \in \mathbb{Z}} P(x \mid y) \ln P(x \mid y), \tag{3.8}$$

and by averaging $H(X \mid y)$ over all outcomes y, the conditional Shannon entropy is

defined

$$H(X \mid Y) \coloneqq \sum_{y \in \mathbb{Z}} P(y) H(X \mid y) = -\sum_{x,y \in \mathbb{Z}} P(x,y) \ln P(x \mid y).$$
(3.9)

The conditional Shannon entropy satisfies:

$$H(X \mid Y) = H(X, Y) - H(Y), \quad H(Y \mid X) = H(X, Y) - H(X).$$
(3.10)

By definition, $H(X \mid y) \ge 0$ and $H(X \mid Y) \ge 0$. Hence

$$H(X,Y) \ge H(Y), \quad H(X,Y) \ge H(X), \tag{3.11}$$

which implies that the addition of a second random variable can never decrease the overall randomness of the system.

The mutual information is defined as

$$I(X,Y) := H(X) + H(Y) - H(X,Y),$$

= $\sum_{x,y} P(x,y) \ln \frac{P(x,y)}{P(x)P(y)}.$ (3.12)

As we demonstrate below, the mutual information satisfies

$$0 \le I(X, Y) \le H(X), \quad 0 \le I(X, Y) \le H(Y).$$
 (3.13)

I(X,Y) = 0 is obtained when X and Y are independent which can be verified by substituting P(x,y) = P(x)P(y) into (3.12). On the other hand, I(X,Y) = H(X)is obtained when H(X | Y) = 0 (or equivalently, in the case where H(X | y) = 0 for any y). This condition states that for any y, there is an x for which P(x | y) = 1 and hence we can always estimate x precisely from an observation of y. By symmetry, similar statements can be obtained to show $I(X,Y) \leq H(Y)$.

It is possible to use a Venn diagram to represent the relationship between H(X), H(Y), H(X,Y), H(X | Y), H(Y | X) and I(X,Y). Fig. 3.4 represents the Shannon entropy of the two variables, X and Y, as two circles where the size of the circle represents the value of the Shannon entropy. The joint entropy is the union of the two circles. When the two variables are correlated, there is an overlap between the circles and which represents their mutual information, the amount of uncertainty reduced in one variable when the other is known. The area of each circle left over when the mutual information is removed represents the conditional Shannon information. A correlation between the circles reduces the joint entropy of the two variables by exactly the mutual information. If there is any correlation



Figure 3.4: Venn diagram illustrating the relationship between Shannon entropy, conditional Shannon entropy, joint Shannon entropy and Mutual information. The red circle represents H(X) and the blue circle is H(Y). Their union is the joint entropy H(X, Y).

between the variables, their joint entropy is less than if there were no correlation⁴.

If X and Y are perfectly correlated then I(X,Y) = H(X) = H(Y) = H(X,Y), and the circles in Fig. 3.4 would be totally overlapping, leaving no conditional information. In general, the mutual information describes the degree of correlation between X and Y with highly correlated random variables having a higher mutual information.

To prove (3.13), we use the fact that $\ln a - 1 \ge 1 - a$ when a > 0 to write

$$-\sum_{x\in\mathbb{Z}}P(x,y)\ln\frac{P(x)P(y)}{P(x,y)} \ge -\sum_{x\in\mathbb{Z}}P(x,y)\left(1-\frac{P(x)P(y)}{P(x,y)}\right) = 0,$$
 (3.14)

which implies $I(X, Y) \ge 0$ for the lower bounds of (3.13). For the upper bounds of the inequalities, we have I(X, Y) = H(X) - H(X | Y) = H(Y) - H(Y | X) which implies $I(X, Y) \le H(X)$ and $I(X, Y) \le H(Y)$. It is possible to use a Venn diagram to represent the relationship between H(X), H(Y), H(X, Y), H(X | Y), H(Y | X)and I(X, Y).

3.2.3 Change in uncertainty

In the context of measurements, it will be of interest to not only characterise the amount of information entropy of the outcome of a random variable, but also quantify the change in this information upon adding extra conditions. For example, we wish to quantify how much the uncertainty of the system state changes in a particular realisation or on average upon making a new measurement of a system.

⁴This statement follows from reversing (3.12) to read H(X, Y) = H(X) + H(Y) - I(X, Y).

As in Ch. 2, consider that \mathbf{x}_t and \mathbf{y}_t are trajectories of Markov chains and that at each time s, the values y_s and x_s are related through the conditional probability $p(y_s | x_s)$ as above in Sec. 3.1.1. Assume that we have access only to the y_s s and do not have access to the x_s s directly, thus must infer the value x_s through our observations of y_s s. That is, we wish to quantify the reduction in uncertainty of x_s before and after observing y_s . Prior to observing y_s , the only source of information we have on the current value of x_s is the trajectory \mathbf{y}_{s-1} . The question we ask is, upon observing y_s , by how much does one's uncertainty on x_s change?

Bayes Rule and Uncertainty

Let us consider a system evolving under some discrete-time Markovian dynamics. Its trajectory through state space up to time t can be written $\mathbf{x}_t = \{x_s\}_{s=0}^t$. If we now have some measurement process running parallel to this, we will obtain a series of measurements $\mathbf{y}_t = \{y_s\}_{s=0}^t$.

Given that $-\ln p$ expresses the information content of an individual event that occurs with probability p, then it follows that the change in uncertainty after witnessing such an event is $\ln p^5$. We define the change in uncertainty to be the difference in the information gained (i.e. entropy 'decreased') when conditioning on an additional variable. For the system described above, the change in uncertainty on x_s upon observing y_s is written

$$\Delta I_s = \ln p(x_s \mid y_0, \dots, y_s) - \ln p(x_s \mid y_0, \dots, y_{s-1})$$

= $\ln \frac{p(x_s \mid y_0, \dots, y_s)}{p(x_s \mid y_0, \dots, y_{s-1})}.$ (3.15)

Writing the change in uncertainty in this way makes explicit that ΔI_s is the difference in information gained by observing x_s conditioned on the measurement history and when conditioned on the measurement history *plus* the most recent measurement. That is, it is the amount of additional information gained when the most recent measurement is taken into account. In advance of Sec. 3.3.5, we now show that (3.15) is also expressed as

$$\Delta I_s = \ln \frac{p(y_s \mid x_s)}{p(y_s \mid y_0, \dots, y_{s-1})},$$
(3.16)

⁵Since one's uncertainty *decreases* by exactly the amount of information stored in the event.

To show this, we need to demonstrate that

$$\frac{p(y_s \mid x_s)}{p(y_s \mid y_0, \dots, y_{s-1})} = \frac{p(x_s \mid y_0, \dots, y_s)}{p(x_s \mid y_0, \dots, y_{s-1})}.$$
(3.17)

First, we add an extra condition to the numerator of (3.16). Namely, we condition the numerator on the antecedent values of y_s . We are free to do this because in a simple independent measurement, the outcome y_s is only directly dependent on x_s . y_s is dependent on its antecedent values indirectly through x_s . Hence, y_s is conditionally independent of y_0, \ldots, y_{s-1} given x_s . Thus, we write

$$\frac{p(y_s \mid x_s)}{p(y_s \mid y_0, \dots, y_{s-1})} = \frac{p(y_s \mid x_s, y_0, \dots, y_{s-1})}{p(y_s \mid y_0, \dots, y_{s-1})}.$$
(3.18)

We can now use the conditional Bayes rule to re-write the numerator. The conditional Bayes rule for events A = a and B = b, with some condition C = c is given by

$$P(A = a \mid B = b, C = c) = \frac{P(A = a, B = b \mid C = c)}{P(B = b \mid C = c)}.$$
(3.19)

And so, the LHS of (3.17) becomes

$$\frac{p(y_s \mid x_s, y_0, \dots, y_{s-1})}{p(y_s \mid y_0, \dots, y_{s-1})} = \frac{p(x_s, y_s \mid y_0, \dots, y_{s-1})/p(x_s \mid y_0, \dots, y_{s-1})}{p(y_s \mid y_0, \dots, y_{s-1})}.$$
 (3.20)

We then reapply Bayes' rule in order to write the RHS of (3.20) as

$$\frac{p(y_s \mid x_s, y_0, \dots, y_{s-1})}{p(y_s \mid y_0, \dots, y_{s-1})} = \frac{p(x_s \mid y_0, \dots, y_s)}{p(x_s \mid y_0, \dots, y_{s-1})}.$$
(3.21)

which implies (3.17).

Relationship to Mutual Information

The ensemble average of the time-integrated change in uncertainty is exactly the mutual information between the two variables. First, we write that the time-integrated change in uncertainty is

$$\Delta \mathcal{I}_t \left[\mathbf{x}_t, \mathbf{y}_t \right] = \sum_{s=0}^t \ln \frac{p(y_s \mid x_s)}{p(y_s \mid y_{s-1}, \dots, y_0)},$$
$$= \sum_{s=0}^t \ln \frac{p(y_s \mid x_s)}{p(y_s \mid \mathbf{y}_{s-1})},$$
$$= \ln \frac{P(\mathbf{y}_t \mid \mathbf{x}_t)}{P(\mathbf{y}_t)},$$
(3.22)

where we have used the general 'product rule' for probability quoted in Schum [2001]

$$\mathbb{P}\left[\mathbf{Y}_t = \mathbf{y}_t\right] = P(\mathbf{y}_t) = \prod_{s=0}^t p(y_s \mid \mathbf{y}_{s-1}), \qquad (3.23)$$

with $p(y_0 | \mathbf{y}_{-1}) \coloneqq 1$. The ensemble average of (3.22) is then

$$\langle \Delta \mathcal{I}_t \left[\mathbf{X}_t, \mathbf{Y}_t \right] \rangle = \sum_{x \in X, y \in Y} P(\mathbf{x}_t, \mathbf{y}_t) \ln \frac{P(\mathbf{y}_t \mid \mathbf{x}_t)}{P(\mathbf{y}_t)},$$

$$= \sum_{x \in X, y \in Y} P(\mathbf{x}_t, \mathbf{y}_t) \ln \frac{P(\mathbf{y}_t, \mathbf{x}_t)}{P(\mathbf{x}_t)P(\mathbf{y}_t)},$$
(3.24)

where in going from the first to the second line we have used the definition of conditional probability. The right hand side of Eq.(3.24) is then exactly the mutual information between \mathbf{x}_t and \mathbf{y}_t , c.f. Eq. (3.12). Hence, the average total change in uncertainty upon making a series of measurements is exactly the mutual information between the two variables.

3.3 Reversal of feedback protocol

Now that we have introduced control and feedback to our stochastic systems, we need to consider how these inclusions impact the form of the trajectory functionals. Specifically, we want to know whether it is still possible to discuss trajectory functionals that relate 'forward' and 'backward' trajectories, since we can then obtain functionals that relate to stochastic entropy production. In this section we will discuss the general form of the trajectory functionals for feedback systems and in what sense they can be 'reversed' and what the resultant fluctuation relations mean (if they mean anything at all).

3.3.1 General form of fluctuation relationship

Whereas previously we considered functionals of the form 2.21 for simple Markov chains, we now consider general functionals of two trajectories of the form

$$\mathcal{G}_t\left[\mathbf{x}_t, \mathbf{y}_t\right] \coloneqq \ln \frac{\mathbb{P}\left[\mathbf{X}_t = \mathbf{x}_t, \mathbf{Y}_t = \mathbf{y}_t\right]}{\mathbb{Q}\left[\tilde{\mathbf{X}}_t = \tilde{\mathbf{x}}_t, \tilde{\mathbf{Y}}_t = \tilde{\mathbf{y}}_t\right]}$$
(3.25)

where \mathbb{Q} is the distribution of the reversed trajectories in some conjugate process. There are several criteria we would want \mathcal{G}_t to fulfil. Namely, those are

- 1. The functional should be well-defined for all possible trajectories
- 2. The functional should ideally have some relation to physical quantities (e.g. heat, entropy, current etc)
- 3. The functional should be asymptotically time-extensive
- 4. The functional should satisfy some time reversal parity such that $\mathcal{G}_t[\mathbf{x}_t, \mathbf{y}_t] = -\mathcal{G}_t[\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t]$ (recalling that $\tilde{\mathbf{x}}_t$ is the time reversed trajectory).

Our conjugate process should be chosen such that the functional \mathcal{G}_t can meet these requirements. In order to relate these functionals to physical concepts (such as entropy or particle current) and to the previously obtained FRs, we wish to consider a conjugate process that is in some sense a 'time reversal' of the forward process. In the previous chapter, we obtained conjugate processes by time-reversing the evolution of the control protocol or by defining new dynamics. For feedback systems, the conjugate process, and how it could relate to time-reversal is not always clear. Such systems involve the interaction of the main 'controlled' system with some secondary 'controller' system, and this interaction cannot always be simply reversed. As before, different choices for the 'conjugate' process will yield different fluctuation relationships. The functional essentially defines a relationship between two processes which could be used to obtain a detailed fluctuation relation of the form

$$\frac{\mathbb{P}\left[\mathcal{G}_t = gt\right]}{\mathbb{Q}\left[\mathcal{G}_t = -gt\right]} = e^{gt}.$$
(3.26)

Eq. (3.26) is obtained by using the definition of \mathcal{G}_t in (3.25) and the final criteria listed above, $\mathcal{G}_t[\mathbf{x}_t, \mathbf{y}_t] = -\mathcal{G}_t[\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t]$, and writing

$$\mathbb{P}\left[\mathcal{G}_{t} = gt\right] = \sum_{\mathbf{x}_{t}, \mathbf{y}_{t}} \mathbb{P}\left[\mathbf{X}_{t} = \mathbf{x}_{t}, \mathbf{Y}_{t} = \mathbf{y}_{t}\right] \delta\left(\mathcal{G}_{t}\left[\mathbf{x}_{t}, \mathbf{y}_{t}\right] = gt\right)$$

$$= \sum_{\mathbf{x}_{t}, \mathbf{y}_{t}} \mathbb{Q}\left[\tilde{\mathbf{X}}_{t} = \tilde{\mathbf{x}}_{t}, \tilde{\mathbf{Y}}_{t} = \tilde{\mathbf{y}}_{t}\right] e^{\mathcal{G}_{t}[\mathbf{x}_{t}, \mathbf{y}_{t}]} \delta\left(\mathcal{G}_{t}\left[\mathbf{x}_{t}, \mathbf{y}_{t}\right] = gt\right)$$

$$= \sum_{\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}} \mathbb{Q}\left[\tilde{\mathbf{X}}_{t} = \tilde{\mathbf{x}}_{t}, \tilde{\mathbf{Y}}_{t} = \tilde{\mathbf{y}}_{t}\right] e^{-\mathcal{G}[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}]} \delta\left(\mathcal{G}_{t}\left[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}\right] = -gt\right)$$

$$= \mathbb{Q}\left[\mathcal{G}_{t} = -gt\right] e^{gt}.$$
(3.27)

3.3.2 Trajectory through state space

The process shown in Fig. 3.2 generates trajectories through state space that can be thought of in several different ways. As mentioned above at the end of Sec. 3.1.3, we can gather the (X_s, Y_s) pairs and treat them as states in a new Markov process. This would generate the trajectory shown by the red line in Fig. 3.5. Conversely, we could group the pairs (X_{s+1}, Y_s) to generate a Markov process shown by the purple line in Fig. 3.5. Both of these are reasonable choices, and both Markov models are easily realisable in simulation. However, if we wish to implement the feedback system in simulation we are forced to consider the selection of X_s and Y_s as separate processes. That is, we cannot select a Y_s according to the measurement process characterised by $p(y_s | x_s)$ until x_s is known, and hence we cannot transition directly from (x_s, y_s) to (x_{s+1}, y_{s+1}) , but must move through a transitory state (x_{s+1}, y_s) and then enact the measurement process on this new state to transition from (x_{s+1}, y_s) to (x_{s+1}, y_{s+1}) . We will call the process shown with the blue line in Fig. 3.5 a "move-then-measure" process, as the subsystem X must move to a new state before the measurement on that state can be performed and the realise the next value of Y.

3.3.3 Reverse trajectory

Consider that the conjugate process to generate the reverse trajectory is the same as the forward process, i.e. $\mathbb{Q} = \mathbb{P}$ on a fixed time interval $\{0, \ldots, t\}$. Using the forward process to generate the reverse trajectory is still clearly a "move-then-measure" process. The probability of obtaining that reverse trajectory is then written

$$\mathbb{P}\left[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}\right] = \pi_{0}(x_{t})p(y_{t} \mid x_{t})\prod_{s=0}^{t-1}\omega(x_{s+1} \to x_{s} \mid y_{s+1})p(y_{s} \mid x_{s}), \qquad (3.28)$$



Figure 3.5: State-space plot showing three different interpretations of a trajectory through the system. Blue circles show the states (x_s, y_s) , purple squares mark the transitionary states (x_{s+1}, y_s) that occur when the system has evolved but a new measurement has not yet been performed. The red line shows the Markov trajectory through the (X, Y) pairs, whereas the blue line shows the path through the system that is taken in simulation, in which the state (x_s, y_s) transitions to (x_{s+1}, y_s) before a new measurement is performed which then transitions the system from (x_{s+1}, y_s) to (x_{s+1}, y_{s+1}) . The purple line show another Markov process that links the transitionary states (x_{s+1}, y_s)



Figure 3.6: Bayesian causal diagram for generating the time-reversed trajectory when $\mathbb{Q} = \mathbb{P}$. In this process, whilst the trajectory maybe be an exact reversal of a forward trajectory, the transition $x' \to x$ in this process occurs under a different control parameter value y'.



Figure 3.7: State-space plot showing two different trajectories. This plot compares the implementation of the forward process described by Eq. (3.4) (blue arrows) to the conjugate process as it may be realised by implementing the suggested by Eq. (3.28) (green arrows).

and is represented by a causal diagram as shown in Fig. 3.6.

Substituting 3.28 into 3.25 yields

$$\frac{\mathbb{P}\left[\mathbf{x}_{t}, \mathbf{y}_{t}\right]}{\mathbb{P}\left[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}\right]} = \frac{\pi_{0}(x_{0})p(y_{0} \mid x_{0})}{\pi_{0}(x_{t})p(y_{t} \mid x_{t})} \prod_{s=0}^{t-1} \frac{\omega(x_{s} \to x_{s+1} \mid y_{s})p(y_{s+1} \mid x_{s+1})}{\omega(x_{s+1} \to x_{s} \mid y_{s+1})p(y_{s} \mid x_{s})},$$

$$= \frac{\pi_{0}(x_{0})}{\pi_{0}(x_{t})} \prod_{s=0}^{t-1} \frac{\omega(x_{s} \to x_{s+1} \mid y_{s})}{\omega(x_{s+1} \to x_{s} \mid y_{s+1})},$$
(3.29)

where the terms $p(y \mid x)$ all cancel owing to the telescoping product. It is not clear what (3.29) represents. Firstly, the RHS appears to only depend on \mathbf{y}_s as a condition rather than providing any term relating to its probability of being realised. The terms $\frac{\omega(x_s \to x_{s+1} \mid y_s)}{\omega(x_{s+1} \to x_s \mid y_{s+1})}$ are a ratio between transition rates across a given bond under different control parameters and not obviously related to any meaningful quantity.

Fig. 3.7 shows the path taken through the state space by the forward process and the conjugate process. We see from this a detail that is not immediately obvious from (3.29). That is, that the two paths are different, only sharing the same (x_s, y_s) states, but moving through different transitory states. Of course, if we combine the


Figure 3.8: Bayesian causal diagram for generating the time-reversed trajectory and preserving the trajectory through path-space. The causal diagram features a cycle which means that the process cannot be simulated as there is no 'input' state.

transition and measurement into a single process as in Sec. 3.1.3, then the reverse process would move backwards along the red line in Fig. 3.5 and we would obtain fluctuation relations as before, as the functional is essentially the same as (2.27), only on a different state space.

Perhaps the biggest problem with this scheme is that it violates the first (and probably most important) requirement we placed on the functional above, that the functional should be well-defined for all realisable trajectories. Here we have $\mathbb{P}[\mathbf{x}_t, \mathbf{y}_t] > 0 \implies \mathbb{P}[\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t] > 0$ as $\omega(x_s \to x_{s+1} \mid y_s) > 0 \implies \omega(x_{s+1} \to x_s \mid y_{s+1}) > 0$. We could have a situation where $\mathbb{P}[\mathbf{x}_t, \mathbf{y}_t] > 0$ and $\mathbb{P}[\tilde{\mathbf{x}}_t, \tilde{\mathbf{y}}_t] = 0$, hence $\mathcal{G}_t[\mathbf{x}_t, \mathbf{y}_t] = \infty$.

3.3.4 Modified-dynamics

In an attempt to make sure that the functional is defined for all trajectories realisable by the forward process and that the reverse process produces the reverse trajectory exactly, we could choose a conjugate process such that,

$$\mathbb{Q}\left[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}\right] = \pi_{0}(x_{t})p(y_{0} \mid x_{0}) \prod_{s=0}^{t-1} \omega(x_{s+1} \to x_{s} \mid y_{s})p(y_{s+1} \mid x_{s+1}),$$
(3.30)

where the terms labelled $e^{\Delta S_{\text{sys}}}$ and $e^{\Delta S_{\text{env}}}$ are generalisations of (2.41) and (2.45) when the system is explicitly conditioned on some control protocol. This conjugate process is obtained by first considering the control protocol as generated by the forward trajectory, which occurs with probability $\sum_{s=0}^{t} p(y_s \mid x_s)$. We then assume that this protocol is employed in the reverse process so that the reverse trajectory is exactly the reverse of the forward trajectory, that is, it follows the blue path in Fig. 3.5 in reverse. There are two ways of rationalising this process; firstly as a "move-then-measure" procedure as described above, or instead considering the conjugate process to be a "measure-then-move" process. Both of these rationalisations encounter problems in their interpretation as real physical processes for reasons we will describe below.

Placing 3.30 into 3.25 leads us to the FR

$$\frac{\mathbb{P}\left[\mathbf{x}_{t}, \mathbf{y}_{t}\right]}{\mathbb{Q}\left[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}\right]} = \underbrace{\frac{e^{\Delta S_{\text{sys}}}}{\pi_{0}(x_{0})p(y_{0} \mid x_{0})}}_{q_{0}(x_{t})p(y_{t} \mid x_{t})} \underbrace{\prod_{s=0}^{t-1} \frac{\omega(x_{s} \rightarrow x_{s+1} \mid y_{s})}{\omega(x_{s+1} \rightarrow x_{s} \mid y_{s})}}_{q_{0}(x_{s+1} \mid x_{s+1})} \underbrace{\frac{e^{\Delta S_{\text{env}}}}{p(y_{s+1} \mid x_{s+1})}}_{p(y_{s+1} \mid x_{s+1})},$$

$$= e^{\mathcal{R}'_{F}}. \tag{3.31}$$

The latter term in the product on the LHS of (3.31) cancels and we are left with ratios of transition probabilities forward and backward across the same bond under the same value of the control parameter. While notationally different from before, this term is essentially the exponential of the counting increments (2.26), albeit generalised to include the control protocol. As such, we can associate such a term with the change in system entropy as before in Sec. 2.4. The first term on the LHS of (3.31) (outside the product) could be thought of as a generalisation of the change in environmental entropy as it is essentially the ratio of probability of finding the system in the initial state (x_0, y_0) and the final state (x_t, y_t) .

A recent paper by Ponmurugan [2010] obtains a fluctuation relation the same as (3.31) for a feedback controlled system by starting from a different consideration. Their starting point is that in the conjugate process, the feedback device processes an exactly equal amount of information as in the forward process. Essentially, their approach is 'top-down' assuming that the information obtained by the feedback device is the same in the conjugate process, whereas our approach is 'bottom-up', fixing the microscopic dynamics so that each transition can occur, and so that the reverse process moves along the trajectory of the forward process. However, we now argue that this fluctuation relation is not physically relevant as the conjugate process is not physically realisable.

Let us first consider that the conjugate process is a "move-then-measure" process as in Sec. 3.3.2. While not immediately obvious from the form of the FR, the causal diagram shown in Fig. 3.8 provides insight on a problem with this formulation. The first point of interest for this process is that it is clearly acausal. In fact, referring

to Fig. 3.8 as a 'causal' diagram is something of a misnomer, as it does not fulfil the 'acyclic' criteria of a Bayesian causal diagram owing to the presence of a loop. The loop implies that a value x_s is a descendant of itself. Stated in words, the outcome of X_s is conditionally dependent on the outcome of X_{s+1} which is in turn conditionally dependent on the outcome of Y_s , itself conditionally dependent on the outcome of X_s . The result of this is that variables conditionally depend on variables which are not determined yet as they depend on the first variable. That is, if we attempted to implement this process in an algorithm, we would find ourselves at the first step having to reference variables that will not be determined until later steps and which themselves will be determined based on the variable currently being determined. Stated simply, there is no 'first cause' in this causal chain. Finally, even if we could somehow understand this acausal conjugate process, imagining it as a "move-then-measure" process means the conjugate trajectory follows the green path from Fig. 3.7, which we already suggested was a problem in Sec. 3.3.3.

Instead, if we consider the conjugate process to be a "measure-then-move" process then we again run into conceptual problems stemming from the causal diagram. In order to move backwards along the same trajectory as the forward process, the state of the measurement in the conjugate process needs always to reflect the *next* state to be transitioned to. That is, to move backwards along the blue line in Fig. 3.7, the measurement must be performed to transition from y_s to y_{s-1} so that the system can then transition from x_s to x_{s-1} under the control parameter y_{s-1} , mirroring how the system transitions from x_{s-1} to x_s under y_{s-1} in the forward process. This could be understood in multiple ways. We could be selecting a trajectory of the "measure-then-move" process in which the measurement is not correlated with the same variable as it is in the forward process, i.e. in the forward process Y_s is correlated with X_s , whereas in the backward process Y_s is correlated with X_{s-1} . Alternatively, the 'measurement' process may actually be an entirely different mechanism (albeit one still characterised by the conditional distribution p(y|x)). This conditional probability either describes a prediction of the next state, or it represents a measurement of an as yet undetermined state, the latter of which presents causal problems. We are unable to give a clear interpretation to the measurement terms in this reverse process, and have no other reason to believe that the ratio of these two probabilities describing different processes is of any direct interest.

In this case then, the functional is not really any different from 2.27. However, owing to the issue discussed above, this functional is of limited interest as the reverse protocol is either not physically realisable or not obviously of any specific relevance to experimental situations.



Figure 3.9: Bayesian causal diagram for a trajectory subject to open-loop feedback control. A measurement history \mathbf{y}_t is sampled from $\mathbb{P}[\mathbf{Y}_t = \mathbf{y}_t]$ and then used as a control parameter. For (3.33), we consider a \mathbf{y}_t such that each transition $x' \to x$ occurs under the same y that it would in the forward process.

3.3.5 Open-loop control

Another approach, discussed often in the literature (for example Horowitz and Vaikuntanathan [2010]; Sagawa and Ueda [2010]) involves a conjugate process with trajectories given by

$$\mathbb{Q}\left[\tilde{\mathbf{x}}_{t}, \tilde{\mathbf{y}}_{t}\right] = \mathbb{P}\left[\tilde{\mathbf{x}}_{t} \mid \tilde{\mathbf{y}}_{t}\right] \mathbb{P}\left[\tilde{\mathbf{y}}_{t}\right]$$
$$= \mathbb{P}\left[\mathbf{y}_{t}\right] \pi_{0}(x_{t}) \prod_{s=0}^{t-1} \omega(x_{s+1} \to x_{s} \mid y_{s}).$$
(3.32)

In this process, the \mathbf{y}_t is provided running the forward dynamics many times and recording \mathbf{y}_t and then reversing that protocol and enacting open-loop control. Thus, in this conjugate process, no 'feedback' actually occurs as shown in the associated Bayesian diagram in Fig. 3.9. As the diagram shows, the measurement history is fed into the process as the system evolves, but each value of y_s is actually sampled from $\mathbb{P}[\mathbf{Y}_t = \mathbf{y}_t]$.

Again, we stress that no feedback takes place. The conjugate process is an open-loop control process, contrasting with the forward process's closed-loop structure. This is reflected in the causal diagram which features a strictly tree-like structure with only a single path connecting any two nodes, again contrasting with Fig. 3.2, which has non-unique paths connecting nodes. Eq. (3.32) leads to the FR

$$\frac{\mathbb{P}\left[\mathbf{x}_{t}, \mathbf{y}_{t}\right]}{\mathbb{P}\left[\tilde{\mathbf{x}}_{t} \mid \tilde{\mathbf{y}}_{t}\right] \mathbb{P}\left[\tilde{\mathbf{y}}_{t}\right]} = \underbrace{\frac{e^{\Delta S_{\text{sys}}}}{\pi_{0}(x_{0})}}_{e^{\Delta S_{\text{tot}}}} \underbrace{\frac{e^{\Delta S_{\text{env}}}}{\frac{e^{\Delta S_{\text{env}}}}{\omega(x_{s+1} \to x_{s} \mid y_{s})}}_{i_{s=0}} \underbrace{\prod_{s=0}^{e^{\Delta T_{t}}} \frac{e^{\Delta T_{t}}}{p(y_{s} \mid x_{s})}}_{p(y_{s} \mid y_{s-1})},$$

$$= e^{\Delta S_{\text{tot}} + \Delta \mathcal{I}_{t}}.$$
(3.33)

In Eq. (3.33), we can recognise the final term on the right hand side (second product) as $e^{\Delta \mathcal{I}_t}$, the exponential of the time integrated change in uncertainty (3.22). We know from Sec. 3.2.3 that this term quantifies the information gained by the controller as it measures the state of the controlled system, and hence by considering this conjugate process, we obtain a functional that is the sum of the stochastic entropy production ΔS_{tot} and change in uncertainty of the feedback controller $\Delta \mathcal{I}_t$ along the trajectory.

Aside from resulting in physically meaningful quantities, this functional also has the benefit that it is well defined (not singular) for all possible trajectories in forward and backward processes, as the protocol dependent microscopic reversibility (3.3) applies and so both the entropy term ΔS_{tot} and the information term ΔI_t are non-singular for all trajectories that are realisable by the forward *and* conjugate processes.

The ensemble average of the right hand side of (3.33) and evaluating that average gives exactly the 'Sagawa-Ueda' fluctuation relation Sagawa and Ueda [2010]

$$\langle e^{-\Delta S_{\rm tot} - \Delta \mathcal{I}_t} \rangle = 1,$$
 (3.34)

from which the 'generalised 2nd law' can be immediately obtained by using the convexity of the logarithm (Jensen's inequality) i.e. $\langle e^A \rangle \geq e^{\langle A \rangle}$ to obtain

$$\langle \Delta S_{\text{tot}} \rangle \ge -\langle \Delta \mathcal{I}_t \rangle.$$
 (3.35)

3.4 Generalised second law

3.4.1 Interpretation of the generalised law

The 'traditional' second law as expressed by (2.90) is not obeyed in the case of $\Delta I_t > 0$ in (3.35). Eq. (3.35) suggests that in the presence of feedback, it is possible to achieve negative values for the entropy production. This expression is also known

as the 'second law of information thermodynamics'⁶ and quantifies the reduction in entropy possible because of devices like Maxwell's dæmon. However, such reductions in entropy are not considered in this framework to be violations of the second law, since it is clear that we have neglected the cost of operating the dæmon.

Through the measurement process, the dæmon's uncertainty is reduced by exactly $\Delta \mathcal{I}_t$, which presumably corresponds to a change in the state of the dæmon. Neglecting the changed state of the dæmon, one might think the second law had been violated by this action. However, in order for the total system (that is, the system and the dæmon) to be operated in a full cycle, the dæmon's state must eventually be reset. That is, the dæmon must regain uncertainty equal to $\Delta \mathcal{I}_t$. That is, despite having reduced the entropy of the controlled system by $\Delta \mathcal{S}_t$, we must reset the dæmon, and in so doing will generate entropy equal to $\Delta \mathcal{I}_t$. Hence, even if we reach the upper bound of (3.35), and achieve $\Delta \mathcal{S}_t = -\Delta \mathcal{I}_t$, we must also generate entropy $\Delta \mathcal{I}_t$ in the dæmon Bennett [1973]. The Landauer principle also means that if we wish to erase the information $\Delta \mathcal{I}_t$ to reset the dæmon, this process cannot be done without generating at least $\Delta \mathcal{I}_t$ worth of entropy Landauer [1961].

A dæmon must possess information about the state of the system in order to be able to manipulate the system such that work can be extracted without dissipating heat. We here note that, while the dæmon may be in possession of an amount of information ΔI_t , (3.35) does not guarantee that the entropy production will be altered by the mere possession of this information. In the case of Maxwell's dæmon, the dæmon must not only possess information about the position and momentum of the particles on either side of the partition, but must use this information intelligently to inform his control actions (the opening and closing of the partition). Clearly, if a dæmon disregards the information about the system state and instead decides to do nothing to the partition, then the entropy production will not be affected. That is, while (3.35) places a boundary on what a dæmon *can* achieve through feedback on a given measurement, it does not tell us what *will* be achieved by a given protocol. To quantify the effectiveness that a specific feedback protocol will have on the entropy production, one must introduce an 'efficacy parameter' as in Sagawa and Ueda.

3.4.2 Szilárd engine

As a concrete example of a system that is described by (3.35), we consider the Szilárd engine previously mentioned in Sec. 1.2. As described before, Szilárd's thought

⁶Although information thermodynamics doesn't really have first, third or zeroth laws distinct



Figure 3.10: Control cycle of the Szilárd engine. The system starts in the top left, labelled 1 and proceeds clockwise until it ends in the same state.

experiment considers a single-particle gas in contact with a heat-bath at temperature T. For the purposes of this illustration, we set $k_{\rm B} = T = 1$. Fig. 3.10 shows the cycle of the Szilárd engine. The engine works by operating the following procedure in a cycle:

- 1. The particle is placed in a container of volume V_0 . The particle is allowed to reach thermal equilibrium with the heat bath.
- 2. A partition is inserted in the centre dividing the container into two equal volumes $V_0/2$
- 3. An errorless measurement is made to determine which side on the partition the particle is on.
- 4. The partition is allowed to quasistatically move to the opposite side of the container, which is recognised as isothermal expansion of an ideal gas. The system does⁷ work

$$\mathcal{W} = -\int_{V_0/2}^{V_0} \frac{k_{\rm B}T}{V} \mathrm{d}V = -k_{\rm B}T\ln 2.$$
(3.36)

from those that already exist in the framework of stochastic thermodynamics.

 $^{^{7}}$ The negative sign in (3.36) comes from the fact that the system is *performing* work

5. With the volume of the container now returned to V_0 , the system is allowed to reequilibrate and the cycle completes in the same state is begins.

The engine operates in a closed cycle and so there is no change in internal energy. The system does work during the expansion phase and during reequilibration exhanges an amount of heat $\Delta Q = W < 0$ with the heat bath. Hence, heat is extracted from the heat bath and used by the system to do work, violating the Kelvin-Planck formulation of the second law. We can also note that $\Delta S_{env} = \Delta Q/T < 0$, and $\Delta S_{sys} = 0$, hence $\Delta S_{tot} = -\ln 2$ always, and so by extension $\langle \Delta S_{tot} \rangle = -\ln 2 < 0$, and so (2.90) is also not obeyed by the Szilárd engine.

For a single cycle, let us consider that the particle position is modelled by a random variable X that can take the values $\{L, R\}$ depending on whether the particle in in the left or the right side of the container. The measurement variable Y also takes the values $\{L, R\}$. The control protocol is the decision to allow the partition to expand to the left or the right. Since the particle position is binary, there is one 'bit' of information in the position variable X. Measured in the natural logarithm, the Shannon information of X is $H(X) = \ln 2$. The measurement process is errorless and hence the mutual information $\langle \Delta \mathcal{I} \rangle = H(X) = \ln 2$. The generalised second law reads $\langle \Delta S_{tot} \rangle \geq -\langle \Delta \mathcal{I} \rangle$, which is obeyed by the Szilárd engine as $\langle \Delta S_{tot} \rangle = -\langle \Delta \mathcal{I} \rangle = -\ln 2$. More generally it can be shown Sagawa and Ueda that no matter what volume ratio the two sides are divided into and with a non error-free measurement, the Szilárd engine still obeys the generalised second law (3.35).

3.5 Conclusion

Having examined three separate conjugate processes in Sec. 3.3, we are left with the conclusion that for feedback processes, only an open-loop control conjugate process as described in Sec. 3.3.5 provides an experimentally relevant or testable fluctuation relation. In Ch. 4, we study a concrete example of a feedback controlled device known as an 'information heat engine' that is able to convert an information resource created through measurement into a work resource. Specifically we shall consider a feedback controller that uses the information resource to reduce the entropy production of the system directly in precisely a manner that is bounded by Eq. (3.35).

"We all behave like Maxwell's demon. Organisms organize. In everyday experience lies the reason sober physicists across two centuries kept this cartoon fantasy alive."

> James Gleick, The Information: A History, a Theory, a Flood

Fluctuations in an Information engine

In the present chapter, we will discuss a concrete model of an information engine and study its fluctuation properties by making use of the relationships described and obtained above in Ch. 2 and Ch. 3. The chapter is structured as follows. Sec. 4.1 describes information engines in general. Sec. 4.2 describes the model we will use of an information engine that operates using a Maxwell's Dæmon type feedback system. In Sec. 4.3 we briefly discuss the applicability/relevance of the fluctuation relations obtained in Ch. 2 and Ch. 3 for this system. In Sec. 4.4 we consider the fluctuations of information on the level of individual trajectories and how to obtain large deviation rate functions. In Sec. 4.5 we show how a two-site version of the model can be used to obtain exact expressions for the large deviation rate function. In Sec. 4.6 we then discuss the difficulties in studying the large deviations of information in systems with a larger state space. We give the details of an approximation and compare to numerical results for these larger systems. In Sec. 4.7 we summarise these results. This chapter is based on the paper Maitland et al. [2015] with some differences in notation.

4.1 Information engines

In this section we introduce the concept of information engines in general before presenting our own model of an information engine inspired by contemporary experimental work on small-scale¹ feedback mechanisms.

4.1.1 Types of information engine

Information engines are a class of thermodynamic system that use information to extract heat from the environment and use this to do work, without the need for a free-energy difference between the initial and final states of the system in contrast to traditional heat engines. Interest in these systems mostly stems from the prospect of creating thermodynamic devices that are capable of extracting work from information reservoirs or information processes such as measurements. However, understanding the theoretical importance of information in physics is also a motivation.

Feedback-driven engine

A feedback-driven engine is a device whereby a measurement is performed on the controlled system. The measurement outcomes are then fed back into the system as described above in Ch. 3 and used to drive the system towards the desired reference by varying the control parameters accordingly. The 'information' in this context is the correlation between the controlled system and the measurement device. The information here can be thought of a resource created by the interaction of two systems, i.e. the measurement device interacts with the measured system and a correlation is created between their states. In this chapter we will study a feedback-driven engine. Measurements are performed which generate correlations between the measurement and controlled systems.

Information-reservoir driven engine

The term 'information reservoir' Barato and Seifert [2014b] is used to denote a a memory-register or tape that has its Shannon entropy modified by interaction with another system. Interaction with an information reservoir allows one to reservoirdriven engine uses an 'information reservoir' to do work. This can be an information register that is fed into the system and used to vary the control parameters according to the predetermined input from the reservoir. Another approach described

¹Small time and spatial scales.

in Mandal et al. [2013] uses an empty memory register and extracts heat from the environment by writing into the memory register. In this context, the information reservoir is seen as a generalisation of the concept of a 'heat bath', it is a resource that already exists rather than a resource created by the interaction of the two systems.

Bi-partite systems

Another type of information engine is realised in a bipartite system Barato et al. [2013a,b]. Bipartite systems are a Markov process described by two variables such that in any transition between states only one variable is allowed to change. Since each component of the system is fully described by this type of model, the entropy production of each component is also readily accessible. By viewing the controlled and measurement systems as a bipartite system, it can be seen that the entropy decrease realised by a Maxwell's dæmon system comes at the cost of an increase in the entropy of the measurement system. Particularly, Hartich et al. [2014] derives bounds for the entropy reduction of the controlled system that can be achieved from its coupling to the measurement system in a Maxwell's dæmon.

4.1.2 Model motivation

We consider a model that abstractly resembles an experimental set-up involving a colloidal particle rotating in an electric field Toyabe et al. [2010]. The experimental system demonstrated a type of particle ratchet where some block or obstacle is placed in the system in order to prevent it from evolving in an undesired manner. By ratcheting in this way, the particle's own thermal motion can be used to do work. Reducing the entropy of a system on average (or extracting heat from the environment and using that heat to do an amount of work²) without a change in free-energy is forbidden by the traditional second law of thermodynamics, as mentioned in Sec. 1.1 and (2.90). Hence, experiments of this type that act as a 'proof of concept' for measurement-based feedback allowing for heat extraction from the environment have drawn a lot of attention Lopez et al. [2008]; Toyabe et al. [2010]; Bérut et al. [2012]; Koski et al. [2014a,b]; Camati et al. [2016].

Specifically, the experiment cited above used a dimeric polystyrene particle pinned on a surface so that it can exhibit rotational Brownian motion. The particle's motion is purely thermal, the particle being buffeted by the thermal fluctuations of the buffer solution it is suspended in. A rotating electric field creates a constant

²The prohibition against this can be seen from (2.85), where $\Delta F = 0 \implies \langle \mathcal{W} \rangle \ge 0$. \mathcal{W} is the 'work done' and must be negative if we are *extracting* work.

torque on the particle, and switching of the phase of the field modulates the torque. This constant torque is balanced so as to hold the particle at some equilibrium position but allow it to fluctuate around this position. If the particle rotates far enough away from its equilibrium position in the desired direction, the phase of the field is switched and the particle adopts a new equilibrium position. The switching protocol is based on measurement of the particle's orientation so as to switch the phase in such a way that the particle rotates preferentially in a single direction.

A cartoon diagram in Toyabe et al. [2010] of the experimental set-up represents the system as a staircase upon which a particle is hopping, and the switching of the field phase as barriers being placed behind the particle as it hops 'up' the staircase. The staircase is an abstraction of the potential well that the particle is in^3 , and the barrier represents the switching phase of the potential to prevent the particle from returning to its previous state. We take this cartoon diagram and model it explicitly, with the difference that we consider a single movable barrier which is repositioned by the feedback controller at every time step.

Our model is comprised of a random walk on a one dimensional lattice with a single movable barrier. The random walk acts as an analogy to the colloidal particle in the experiment, with the walker's random motion modelling the thermal motion of the experimental particle. We model in discrete space, as in the real experiment there was a single coarse-grained measurement, essentially allowing the identification of discrete 'states'.

In experiment, the measurements are performed at regular intervals which allows the whole feedback process to be thought of in discrete time steps. This is clearly an idealisation of an actual physical system. The experiment in Toyabe et al. [2010] studies how the efficacy of the control is affected by the frequency of measurement/control events. Here we have one measurement/control event for each particle jump.

In our model, the random walker transitions in discrete time between sites on a lattice of size L with periodic boundary conditions and probabilities q and pof jumping respectively left and right at each time step, such that p + q = 1. The random walker's motion is then described by a single parameter p, which in the case where $p \neq q$ describes a system with a bias in one direction. Without loss of generality, we consider q > p for biased systems. We label jumps left as 'down', and jumps right as 'up' as though the particle were moving in a tilted potential as per the experiment the model is inspired by.

 $^{^{3}}$ Very abstract. The height up the staircase represents higher potential energy. The fact that there are 'steps' is to represent the sinusoidal shape of the well.



Figure 4.1: Schematic of the model. x and y represent the state of the particle and the measurement respectively. The feedback mechanism makes a measurement y of the particle position x and places the barrier to the left of the particle, according to the outcome of the measurement. If y = x or y = x + 1 then the barrier influences the particle movement by preventing certain jumps as in the top and middle right. For all other measurement values, the barrier has no effect and the particle moves freely as in the bottom right.

In the spirit of the feedback process described in Toyabe et al. [2010], at each time step a measurement is made of the particle position and the barrier is moved to the measured location of the particle in an attempt to prevent it from moving down. When the particle attempts to jump 'through' the barrier, it is blocked and remains at the same site. If the measurement is always correct, the particle can only ever jump up or stay still and the model is identical to a totally asymmetric random walk process. However, if the measurement is incorrect then the barrier will be placed incorrectly and will either not affect the particle's motion or will act as a blockade for jumps up. The system is initialised by first choosing a site uniformly from the lattice sites and then performing the measurement process to obtain a first measurement to determine the control protocol. Fig. 4.1 is a schematic diagram of the system where the bias is represented by showing the lattice as a staircase. The three possible results of the action of the feedback device are shown, including the situation where the feedback has no effect on the particle motion.

4.2 Model description

In this section we detail the model we will be working with for the remainder of this chapter. The position of the particle on the lattice at time s is represented as a random variable X_s with a specific realisation denoted by x_s . As above in Ch. 2 and, a trajectory of the system is written $\mathbf{X}_t = \{X_s\}_{s=0}^t$ with a specific trajectory denoted by \mathbf{x}_t . The initial distribution of the particle position on the lattice is described by $\pi_0(x_0) = \mathbb{P}[X_0 = x_0]$.

We use the feedback-control framework described above in Sec. 3.1. That is, we consider that the transition probabilities are determined by the outcome of the measurement variable (which acts directly as a control parameter). The measurement is represented in a similar fashion to the system trajectory, and is written as $\mathbf{Y}_t = \{Y_s\}_{s=0}^t$. The measurement Y_s at time s only depends on the current state X_s and we denote,

$$p(y \mid x) = \mathbb{P}\left[Y_s = y \mid X_s = x\right] \ \forall s, \tag{4.1}$$

as the probability of obtaining outcome y given that the system is in state x. We again assume a given measurement y_s determines a unique control parameter and thus along with the departure state x_s determines the probability of transitions to the next state x_{s+1} . Again, as above we write the transition probabilities as

$$\omega(x_s \to x_{s+1} \mid y_s). \tag{4.2}$$

The conditional distribution in (4.1) is derived from \mathbb{P} , the path space measure of the full process $\{(\mathbf{X}_s, \mathbf{Y}_s)\}_{s=0}^t$. This process is a Markov chain on the state space given by X and Y pairs and can be described by the transition matrix

$$[\Omega]_{(x',y'),(x,y)} \coloneqq \omega(x \to x' \mid y) p(y' \mid x').$$

$$(4.3)$$

We also write $P(\mathbf{x}_s) = \mathbb{P}[\mathbf{X}_s = \mathbf{x}_s]$ and $P(\mathbf{y}_s) = \mathbb{P}[\mathbf{Y}_s = \mathbf{y}_s]$ for the marginal distributions of the process and measurement trajectories respectively. Note that, while the measurements Y_s are conditionally independent given the path \mathbf{X}_t , the marginal measurement process \mathbf{Y}_t exhibits correlations after integrating out \mathbf{X}_t and is not a sequence of independent identically distributed random variables. The probability of obtaining \mathbf{y}_t given \mathbf{x}_t for this model is

$$P(\mathbf{y}_t \mid \mathbf{x}_t) = \prod_{s=0}^t p(y_s \mid x_s).$$
(4.4)

Entropy and Information

From Sec. 3.1 we use the definition of entropy production for controlled systems. That is, the entropy production at time s in a feedback system is given by (3.5), reprinted here

$$\Delta S_s = \ln \frac{\omega(x_s \to x_{s+1} \mid y_s)}{\omega(x_{s+1} \to x_s \mid y_s)}.$$
(4.5)

For a system with feedback it is also necessary to quantify and study the information gained through measurement Horowitz and Vaikuntanathan [2010]. The change in uncertainty of the feedback device over the full trajectory, defined earlier in Sec. 3.2.3, is reprinted here

$$\Delta \mathcal{I}_t \left[\mathbf{x}_t, \mathbf{y}_t \right] = \sum_{s=0}^t \ln \frac{p(y_s \mid x_s)}{P(y_s \mid \mathbf{y}_{s-1})},$$
$$= \ln \frac{P(\mathbf{y}_t \mid \mathbf{x}_t)}{P(\mathbf{y}_t)}.$$
(4.6)

Recall that this is the sum of the individual changes in uncertainty at each time step as given by (3.16), also reprinted here

$$\Delta I_s = \ln \frac{p(y_s \mid x_s)}{p(y_s \mid y_0, \dots, y_{s-1})}.$$
(4.7)

The denominator term in (4.6) is given by

$$P(\mathbf{y}_t) = \sum_{\mathbf{x}_t} \mathbb{P} \left[\mathbf{X}_t = \mathbf{x}_t, \mathbf{Y}_t = \mathbf{y}_t \right],$$

= $\sum_{\mathbf{x}_t} \pi_0(x_0) p(y_0 \mid x_0) \prod_{s=0}^{t-1} \omega(x_s \to x_{s+1} \mid y_s) p(y_{s+1} \mid x_{s+1}).$ (4.8)

Model parameters

The model can be described by three parameters: the lattice size L, the motion bias p and the parameter r which characterises the measurement accuracy $p(y \mid x)$,

$$p(y \mid x) = \begin{cases} r & \text{if } y = x\\ w = \frac{1-r}{L-1} & \text{if } y \neq x \end{cases},$$
(4.9)

where $0 \le r \le 1$ and r + (L-1)w = 1. The measurement error is independent of which site x the walker is at and all incorrect measurements (i.e. any $y \ne x$) are equally likely.

We consider the case of 'accurate' measurements (r > w) as in this regime the measurements can be used to make useful inferences about the system state and the information gained through measurement can be used in the operation of the information engine. Contrastingly, in the special case r = w = 1/L, the joint probability in (4.8) factorises as each measurement is statistically independent. In that case

$$P(y_s \mid y_{s-1}, \dots, y_0) = p(y_s \mid x_s) = \frac{1}{L},$$
(4.10)

and (4.7) is always zero; no information is ever gained by the dæmon. Interestingly, we note that even when r < w it is possible for the Dæmon to gain information. For a system with only two measurement outcomes, a daemon that measures 'wrongly' more often than 'correctly' is still useful as whenever the daemon gives an output, one can guess that the system is more likely to be in the opposite state. In general we do not study the r < w regime as it is not clear how to utilise the information gained from a single measurement by a device that measures wrongly more frequently than correctly for larger systems.

In the case of r = 1, we would have a perfect measurement device. The measurements would never be never wrong and in this model, the entropy production could *never* be positive. In this case, the mutual information between the measurement and the system would be equal to the Shannon entropy of the system state, i.e. $\langle \Delta \mathcal{I}_t \rangle = I[\mathbf{X}_t; \mathbf{Y}_t] = I[\mathbf{X}_t; \mathbf{X}_t] = H(\mathbf{X}_t)$. Whilst arbitrarily accurate measurements are theoretically possible, completely infallible measurements are of limited experimental relevance⁴ and so we do not test the r = 1 case. However, we will revisit this concept of perfect measurement later in Sec. 5.4.3

The three parameters, p, r and L completely characterise the model. At each time step the contribution to the entropy production is given by (4.5) and is $\Delta S_s \in \{-\ln (p/q), 0, \ln (p/q)\}$, where the non-zero terms correspond to successful jumps down or up and 0 is the entropy produced if the particle attempts to move through the barrier. All three parameters determine the average particle current

$$\frac{\langle J_t \rangle}{t} = pr + pw(L-2) - qw(L-1) \ \forall t > 0,$$
(4.11)

where positive current is defined to be to the right. The first two terms in (4.11) correspond to jumps to the right in the case of a correct measurement or an incorrect measurement that does not block the particle. The final term corresponds to jumps to the left when the particle position has not been measured correctly. Our choice of parameters $(q \ge p \text{ and } r > w)$ can produce a positive average current (current

⁴At least to mere mortals.

against the bias), while in the absence of feedback one would expect a negative current. As an example, consider a three site system where p = 0.3 r = 0.9. The bias is in the negative direction (q = 1 - p > p), but (4.11) shows that the average current should be positive $(\frac{\langle J_t \rangle}{t} = 0.16)$.

4.2.1 Calculating $P(\mathbf{y}_t)$

The change in uncertainty at a given time step cannot be counted by observing individual transitions on a finite-state Markov chain. Eq. (4.7) shows that we need the entire measurement history up until that time in order to calculate the change in uncertainty for a transition at a given time. Aside from this fact, the term $P(y_s | \mathbf{y}_{s-1})$ is difficult to calculate in general, as it is heavily dependent on the specific causal structure of the feedback process Ito and Sagawa [2013].

Evaluating the total information gained along an individual trajectory $\Delta \mathcal{I}_t [\mathbf{x}_t, \mathbf{y}_t]$ also requires the full measurement history, but is easier to calculate as we will see presently. Subtracting $\Delta \mathcal{I}_s$ at two consecutive time steps, we see that (4.7) can be written as

$$\Delta I_s = \Delta \mathcal{I}_s - \Delta \mathcal{I}_{s-1},$$

= $\ln \frac{P(\mathbf{y}_s | \mathbf{x}_s)}{P(\mathbf{y}_s)} - \ln \frac{P(\mathbf{y}_{s-1} | \mathbf{x}_{s-1})}{P(\mathbf{y}_{s-1})},$
= $\ln \frac{P(y_s | x_s) P(\mathbf{y}_{s-1})}{P(\mathbf{y}_s)},$ (4.12)

where we have used the conditional independence of the measurements (4.4). Eq. (4.12) only requires calculation of $P(\mathbf{y}_t)$ (4.8).

We can evaluate (4.8) by writing the sum over trajectories \mathbf{x}_t as a matrix product, representing the terms in the sum as the elements of matrices

$$[M_y]_{x',x} \coloneqq \omega(x \to x' \mid y) p(y \mid x). \tag{4.13}$$

These M_y matrices can be thought of as transition matrices for a Markov process under a fixed value of the control parameter. It is also worth noting that these matrices are similarity transforms of one another and thus have the same spectrum⁵. With the initial probability distribution vector $|\pi_0\rangle$ and the summation vector $\langle 1|$

⁵This is true for any lattice size L.

(as defined in Sec. 2.2.1), we can write $P(\mathbf{y}_t)$ as

$$P(\mathbf{y}_t) = \langle 1 | \prod_{s=0}^t M_{y_s} | \pi_0 \rangle.$$
(4.14)

For any given series of measurements \mathbf{y}_t we can calculate $P(\mathbf{y}_t)$. Each of the matrices has a size $L \times L$ and we must multiply t + 1 of these matrices to calculate the information obtained up to time step t. Eq. (4.12) can then be written as

$$\Delta I_s = \ln p(y_s \mid x_s) \frac{\langle 1 \mid \prod_{s'=0}^{s-1} M_{y_{s'}} \mid \pi_0 \rangle}{\langle 1 \mid \prod_{s'=0}^{s} M_{y_{s'}} \mid \pi_0 \rangle}.$$
(4.15)

We will see in Sec. 4.5 that this representation of ΔI_s also clarifies potential cancellation of terms in the products.

4.3 Entropy production and integral fluctuation relations

As a preliminary to Sec. 4.4, We numerically check the generalised integral fluctuation relation (3.34). We also show that the integral fluctuation relation for entropy (2.89) does not hold, and check the equality,

$$\langle e^{-\Delta \mathcal{I}_t} \rangle = 1, \tag{4.16}$$

which follows from the definition (4.6), which can be rearranged to

$$e^{-\Delta \mathcal{I}_t} = \frac{P(\mathbf{y}_t)}{P(\mathbf{y}_t \mid \mathbf{x}_t)}.$$
(4.17)

The expectation of (4.17) over all trajectories is given by

$$\langle e^{-\Delta \mathcal{I}_t} \rangle = \sum_{\mathbf{x}_t, \mathbf{y}_t} P(\mathbf{x}_t, \mathbf{y}_t) \frac{P(\mathbf{y}_t)}{P(\mathbf{y}_t \mid \mathbf{x}_t)},$$

$$= \sum_{\mathbf{x}_t, \mathbf{y}_t} P(\mathbf{y}_t \mid \mathbf{x}_t) P(\mathbf{x}_t) \frac{P(\mathbf{y}_t)}{P(\mathbf{y}_t \mid \mathbf{x}_t)},$$

$$= \sum_{\mathbf{x}_t, \mathbf{y}_t} P(\mathbf{x}_t) P(\mathbf{y}_t),$$

$$= 1,$$

$$(4.18)$$



Figure 4.2: Numerical tests of (2.89), (3.34) and (4.16) in semi log scale. Parameter values are L = 3, p = 0.2, r = 0.9. Points are plotted from averages over 10^7 realisations. Error bars indicate the standard error on mean of $\langle e^{-S_t - \mathcal{I}_t} \rangle$ at representative points.

where we have used the definition of conditional probability in going from the first line to the second. Fig. 4.2 confirms the (in)equalities (2.89), (3.34) and (4.16) at different times t with numerical data obtained from Monte Carlo simulation of a biased three-site system. Since the means of exponential quantities are determined by rare events, fluctuations are large as indicated by the error bars on representative points.

4.4 Fluctuations of information

It is possible to capture the fluctuations of the entropy production in this system by computing the large deviation rate function. Since entropy production is not the only quantity of interest in a feedback system, we here ask whether it is also possible to analyse the fluctuations of the information flow between the system and feedback controller by obtaining a large deviation rate function for information.

4.4.1 Large Deviation Analysis

In order to study the fluctuation properties of $\Delta \mathcal{I}_t/t$, we follow the framework described in Sec. 2.6 and start by assuming that $\Delta \mathcal{I}_t/t$ obeys a large deviation principle of the form,

$$\mathbb{P}\left[\Delta \mathcal{I}_t = it\right] \sim e^{-E(i)t},\tag{4.19}$$

as t approaches infinity, where E(i) is the 'large deviation rate function'. Recall from Sec. 2.6 that the rate function tells us about the fluctuation properties of the variable $\Delta I_t/t$, and allows us to quantify how exponentially unlikely a given fluctuation away from the mean is in the long-time limit.

As before, to calculate the rate function, we first obtain the scaled cumulant generating function (SCGF),

$$\xi(k) \coloneqq \lim_{t \to \infty} -\frac{1}{t} \ln G(k), \tag{4.20}$$

where G(k) is the moment generating function,

$$G(k) \coloneqq \langle e^{-k\Delta \mathcal{I}_t} \rangle = \int_{-\infty}^{\infty} e^{-ku} \mathbb{P}\left[\Delta \mathcal{I}_t = u\right] \mathrm{d}u, \qquad (4.21)$$

and obtain the rate function from the SCGF via the Legendre-Fenchel transform

$$E(i) = \sup_{k \in \mathbb{R}} \left[\xi(k) - ki \right].$$
(4.22)

In practice, the SCGF (and hence the rate function) can be obtained from a modified transition matrix as described in Sec. 2.6.3. However, this approach can only be applied if the measured quantity depends only on the system state or the transitions occurring at a given time, as is the case with particle current. However, (4.12) shows that this is not the case for information, since the amount of information gained in a single measurement depends on the entire measurement history up until that point⁶. Hence individual states or transitions on the enlarged state space of pairs (x_s, y_s) cannot be associated with specific values of ΔI_s .

However, in the next section we show that for L = 2, ΔI_s can be simplified using (4.15) to an expression that only depends on the departure and target states in a single transition, and so the SCGF can be obtained by weighting the transition matrix (4.3) to count information gain leading to an exact analytical expression of the large deviation rate function. This simplification does not hold for $L \geq 3$, even if the system is fully connected, i.e. every state is reachable from every other state

⁶That is, the information gained is not a state function.

in a single transition. In Sec. 4.6 we describe approximate methods for obtaining the rate function by approximating the sequence of ΔI_s s as a Markov chain.

4.4.2 Fluctuations of ΔI_s

The variables $\Delta I_t/t$ and ΔI_s are of interest, being the time-averaged information gain and instantaneous change in uncertainty, respectively. On the level of an individual trajectory the model allows for three types of events that influence ΔI_s . These are roughly described as follows,

- Correct measurements,
- Incorrect measurements that are not consistent with previous measurements,
- Incorrect measurements that are consistent with previous measurements.

Correct measurements are fairly intuitive to understand, the particle position is measured correctly. We will observe below in Sec. 4.4.3 that correct measurements yield positive amounts of information⁷. We will also see later that correctly observed jumps against the bias yield more information than correctly observed jumps in the direction of the bias that are blocked⁸. For unbiased systems the difference in ΔI_s between jumps up and down is small. Series of correct measurements yield one of two baseline values for information gain ΔI_s that correspond to jumps up and blocked jumps down.

However, incorrect measurements are slightly more complicated. Some incorrect measurements are obvious; if a particle appears to have suddenly jumped to a site it should not be able to access, then the ratio $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ in (4.12) changes and the amount of information gained is zero. Further incorrect measurements can lead to negative values in information gain, interpreted as a change towards greater uncertainty of the system state.

Whenever a series of incorrect measurements is made, the next series of correct measurements gains large positive amounts of information that partially retrieve the 'lost' information of the incorrect measurements. As these correct measurements are made, the ratio $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ relaxes and ΔI_s returns to a baseline value which we explain in detail in Sec. 4.4.3. The information theoretic interpretation is that successive correct measurements allow the observer to infer which measurement was

⁷We note here that this accords with the information theoretic intuition. Correct measurements give positive quantities of information, since they reduce uncertainty.

⁸One might wonder how a 'blocked' jump can be observed. Here we are referring to a jump event where the particle doesn't move (it jumps into the site it is currently occupying), and the measurement history reflects this, i.e. $y_s = x_s$ and $y_{s+1} = x_{s+1}$.

incorrect.

However, an incorrect measurement could be made that nevertheless looks like a plausible measurement (one could measure the particle to have jumped one step forward when in fact it has not moved). These type of measurements do not alter $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ and when made consecutively lose information that cannot be regained by subsequent positive measurements.

The consequences of this is that large positive deviations of $\Delta I_t/t$ are not generated by an accumulation of the largest values of ΔI_s as these are necessarily preceded by large negative values as described above. They are instead generated by strings of correct measurements which each generate less information than the largest values of ΔI_s . In contrast, large negative deviations are generated by sequences of incorrect measurements which happen to represent a possible system trajectory. In this case large negative values of ΔI_s can accumulate and are not compensated by subsequent large positive values. In general, since the baseline values discussed in the next section depend on the trajectory in case of correct measurements, atypical trajectories also play a role in the realisation of large deviations of $\Delta I_t/t$.

4.4.3 Trajectory Analysis

To understand the various events that occur in the system, in Fig. 4.4 we plot trajectories of ΔI_s , $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$, X_s and Y_s . ΔI_s is given by (4.12) (reproduced again here slightly rewritten for clarity)

$$\Delta I_s = \ln p(y_s \mid x_s) + \ln \frac{P(\mathbf{y}_{s-1})}{P(\mathbf{y}_s)}, \qquad (4.23)$$

and hence ΔI_s will always differ from $\ln P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ by $\ln r$ or $\ln w$ depending on whether the measurement is correct or incorrect.

The value of $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ is determined by the measurement history itself and is not directly dependent on the particle trajectory \mathbf{x}_t . We observe that there are two 'baseline' values for this ratio that are obtained when the measurements represent a possible trajectory for the particle, i.e. the particle does not appear to jump more than one site in a single time step and does not appear to move 'through' the barrier. These two baseline values correspond to the particle jumping up or being blocked attempting to jump down.

If the particle is blocked at site x for successive steps starting, for example, at time s - 3, the probability of the correct measurement history is calculated via a

matrix product as in Eq. (4.14),

$$P(\mathbf{y}_s) = \langle 1 | M_y M_y M_y | v_{s-3} \rangle, \qquad (4.24)$$

where y = x and where $|v_{s-3}\rangle = \prod_{k=0}^{s-3} M_{y_k} |\pi_0\rangle$. The ratio $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ is then dominated by the leading eigenvalue λ_{\max}^L of the matrices, and therefore should approach $1/\lambda_{\max}^L$ in a small number of time steps. Recall that all matrices have the same eigenvalues since they are related by translations of rows and columns. The corresponding lower baseline value for ΔI_s is given by

$$\Delta I^b = \ln \frac{r}{\lambda_{\max}^L}.$$
(4.25)

On the other hand, let us assume that the particle jumps up for successive time steps, starting in site x at time s - 3, and we measure this correctly. Then the probability of the measurement history $P(\mathbf{y}_s)$ is given by a product of matrices with increasing index

$$P(\mathbf{y}_s) = \langle 1 | M_{y+2} M_{y+1} M_y | v_{s-3} \rangle, \qquad (4.26)$$

where y = x, and y + n is understood with periodic boundary conditions. The successive matrices are translated by one column and one row, i.e. $M_{y+1} = T^{-1}M_yT$, where the translation T is such that

$$(T|v\rangle)_k = v_{k+1} \text{ and } (\langle w|T\rangle_k = w_{k-1}.$$
 (4.27)

Similar to the eigenvalue case, the expression (4.26) then is dominated by vectors $\langle v \rangle$ and a scale factor α such that

$$\left(M_{y}\left|v\right\rangle\right)_{k} = \alpha v_{k+1},\tag{4.28}$$

with periodic boundaries (i.e. $(M_y | v \rangle)_1 = \alpha v_L$). So the upper baseline value for information gain ΔI_s should be given by

$$\Delta I^u = \ln \frac{r}{\alpha},\tag{4.29}$$

where α can be found numerically from (4.28) for any given system.

Fig. 4.3 shows data from simulation confirming our predictions for the baseline values and demonstrate their scaling with L in unbiased and biased systems, respectively. In these simulations, we force specific events (repeated up jumps and



Figure 4.3: Baseline ΔI values for unbiased and biased systems. Symbols show the value of ΔI_s for consecutive up jumps and blocked down jumps obtained from simulation. Lines show the predictions from (4.25) and (4.29).

repeated up jumps that are blocked) and sample ΔI_s after 100 time steps (in order to avoid any transient behaviour from the initial conditions). In an unbiased system, the difference in $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ between an up and blocked down jump shrinks with increasing L, whereas the values are roughly constant for biased systems. Fig. 4.4a shows a typical section of a trajectory in a biased system. The baseline values of $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ and ΔI_s for blocked down and up jumps are seen around t = 310and t = 315 respectively.

Isolated incorrect measurements as seen in Fig. 4.4a (at times t = 304 and t = 307) and Fig. 4.4b (at time t = 44) cause $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ to increase while we observe that the $\ln w$ term is negative such that $\Delta I_s = 0$. In the following measurements, $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ is still larger than its baseline value and as the $\ln r$ contribution is small in comparison, ΔI_s is also larger than the baseline.

The information theoretic interpretation of these observations is that upon making a measurement that is not compatible with the previous measurements, no new information is gained. This is because it is not clear to the observer whether the current measurement is incorrect, or the previous measurements were erroneous (or both). It is only on subsequent measurements that information is gained, as more measurements allow the observer to make inferences about which measurements were incorrect. After an incorrect measurement, $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ (and hence ΔI_s) returns quickly to a baseline value.

To observe very large values of ΔI_s and instantaneously gain large amounts of information, it is necessary to first lose larger amounts of information through incorrect measurements. An example of this is shown in Fig. 4.4b, where the large amounts of information gained between t = 57 and t = 59 cannot balance the losses



(a) Trajectory of a biased system for L = 10, (b) Trajectory of a biased system for L = 10, p = 0.2, r = 0.9.



Figure 4.4: Trajectories of \mathbf{x}_t (gold line) and \mathbf{y}_t (square boxes) plotted with the information gain ΔI_s , (blue line), and $\ln(P(\mathbf{y}_{s-1})/P(\mathbf{y}_s))$ (red line). Note that these quantities have different units. Note also that for x_s and y_s there are periodic boundary conditions between 0 and 10.

between t = 53 and t = 55. Hence it is not possible to gain additional information by making strategically 'wrong' measurements, as a series of correct measurements would yield more total information.

In the case of a wrong measurement that still represents a possible trajectory for the particle, $P(\mathbf{y}_{s-1})/P(\mathbf{y}_s)$ does not change but the ln w contribution from the incorrect measurement means that ΔI_s takes a negative value. If the following measurements are correct and also compatible with the previous wrong measurement, then subsequent measurements will only gain a baseline amount of information. An example of this is shown in Fig. 4.4c. Here, a wrong measurement occurs at time t = 50 which is compatible with the previous history. Subsequent measurements do not allow the observer to ascertain that any of the previous measurements were incorrect and hence this information loss is not recovered.

4.5 2-site model

A special case of the model occurs when L = 2. In this case, there are just two sites, and thus the state space is fully connected. The system can transition from any state to any other state. Here we will describe how this fact allows for analytic calculation of the large deviation rate function, and compare analytic results to data obtained from simulation.

4.5.1 System matrices

For a system with L = 2 sites labeled 1 and 2, the matrices used in (4.14) are

$$M_1 = \begin{pmatrix} qr & qw \\ pr & pw \end{pmatrix}$$
(4.30)

$$M_2 = \begin{pmatrix} pw & pr \\ qw & qr \end{pmatrix}$$
(4.31)

corresponding to the two measurement outcomes. For the L = 2 case the matrices have only one non-zero eigenvalue $\lambda = qr + pw$, allowing them both to be written as tensor products on the eigenspace of the corresponding eigenvector. That is, we can write $M_y = |v_y^r\rangle \langle v_y^l|$ where $\langle v_y^l|$ and $|v_y^r\rangle$ are the left and right eigenvectors of M_y with respect to λ . The matrices can be rewritten as

$$M_1 = |v_1^r\rangle \langle v_1^l| = \begin{pmatrix} q\\ p \end{pmatrix} \begin{pmatrix} r & w \end{pmatrix}$$
(4.32)

and

$$M_2 = |v_2^r\rangle \langle v_2^l| = \begin{pmatrix} p\\ q \end{pmatrix} \begin{pmatrix} w & r \end{pmatrix}.$$
(4.33)

Writing the matrices in this way shows that all but the final term of the upper product in (4.15) cancel with terms in the lower product, leaving an inner product between two vectors:

$$\frac{P(\mathbf{y}_{s})}{P(\mathbf{y}_{s-1})} = \frac{\langle 1 | M_{y_{s}} M_{y_{s-1}} \dots M_{y_{1}} M_{0} | \pi_{0} \rangle}{\langle 1 | M_{y_{s-1}} \dots M_{y_{0}} M_{y_{0}} | \pi_{0} \rangle}, \\
= \frac{\langle 1 | v_{y_{s}}^{r} \rangle \langle v_{y_{s}}^{l} | v_{y_{s-1}}^{r} \rangle \langle v_{y_{s-1}}^{l} | \dots | \pi_{0} \rangle}{\langle 1 | v_{y_{s-1}}^{r} \rangle \langle v_{y_{s-1}}^{l} | \dots | \pi_{0} \rangle}, \\
= \langle v_{y_{s}}^{l} | v_{y_{s-1}}^{r} \rangle.$$
(4.34)

Eq. (4.34) shows us that $P(\mathbf{y}_s)/P(\mathbf{y}_{s-1})$ can take four values corresponding to the values taken by y_{s-1} and y_s , i.e., $y_{s-1}, y_s \in \{1, 2\}$. However, as the two M_y matrices are permutations (meaning in this case that their eigenvectors are also permutations), only two distinct values can be obtained. These correspond to the cases when $y_s \neq y_{s-1}$ or when $y_s = y_{s-1}$. The change in uncertainty upon making a measurement in the two-site system is

$$\Delta I_s = \ln \frac{p(y_s \mid x_s)}{\langle v_{y_s}^l \mid v_{y_{s-1}}^r \rangle},\tag{4.35}$$

which for a given x_s depends only on the previous and current measurements y_{s-1} and y_s . From (4.35) it can further be deduced that the process $\{\Delta I_s\}_{s=0}^t$ is a sequence of i.i.d. random variables. This simplification only holds for L = 2 and is demonstrated explicitly below in Sec.4.5.2.

4.5.2 Distribution of ΔI_s

As $p(y_s | x_s)$ has two possible values and $\langle v_{y_s}^l | v_{y_{s-1}}^r \rangle$ also has two possible values, ΔI_s takes four possible values. Specifically, these are

$$\ln \frac{r}{pr+qw} \coloneqq a, \ \ln \frac{r}{qr+pw} \coloneqq b, \tag{4.36}$$

for correct measurements made after jumps in the up and down directions (whether blocked or not), respectively. For the same cases followed by incorrect measurements,



Figure 4.5: Cumulative density function for ΔI_s for different system sizes L with p = 0.2 and r = 0.9. Dashed lines give theoretical values for L = 2 as given by (4.36) and(4.37).

 ΔI_s takes the values

$$\ln \frac{w}{pr+qw} \coloneqq c, \ \ln \frac{w}{qr+pw} \coloneqq d. \tag{4.37}$$

The values taken by ΔI_s can be associated with elements of the transition matrix (4.3), for transitions on the state space (x, y). For L = 2 this matrix is

$$\Omega = \begin{pmatrix} qr & qw & pw & pr \\ pr & pw & qw & qr \\ qr & qw & pw & pr \\ pr & pw & qw & qr \end{pmatrix}.$$
(4.38)

The values in (4.36) and (4.37), along with the probabilities of these events occurring, given by the transition matrix (4.38), are enough to determine the cumulative density function (CDF) of the random variable ΔI_s for all $s \geq 1$ since we can easily count all possible events

$$\operatorname{CDF}_{\Delta I}(\Delta i) \coloneqq \mathbb{P}\left[\Delta I \le \Delta i\right] \quad \forall \ \Delta i \in \mathbb{R},$$

$$(4.39)$$

which is plotted in Fig. 4.5 and compared with numerical data.

Independence of ΔI_s for L = 2

To demonstrate that ΔI_s are i.i.d. random variables for L = 2, we want to show that the distribution of ΔI_{s+1} is independent of ΔI_s and identical for all s. We can do this by exhaustively considering all possible transitions, as there are only 16 possible transitions and much degeneracy between the cases. Let us first assume that the system is in the state $(x_s, y_s) = (1, 1)$. The probabilities for ΔI_{s+1} taking the values from (4.36) and (4.37) are as follows:

ΔI_{s+1}	Transition	Probability
a	$(1,1) \to (2,2)$	pr
b	$(1,1) \to (1,1)$	qr
с	$(1,1) \to (1,2)$	qw
d	$(1,1) \rightarrow (2,1)$	pw

By the translation invariance in the system, the same probabilities apply if the system was in state $(x_s, y_s) = (2, 2)$.

If the system starts in the state $(x_s, y_s) = (1, 2)$, then the probabilities are:

ΔI_{s+1}	Transition	Probability
a	$(1,2) \to (1,1)$	pr
b	$(1,2) \to (2,2)$	qr
С	$(1,2) \to (2,1)$	qw
d	$(1,2) \to (1,2)$	pw

which again by translation invariance also holds for beginning in the state (2, 1). We can see then that the probability to obtain a particular value of ΔI_{s+1} is in fact independent of the state (x, y) of the system. In particular, it does not depend on the previous value ΔI_s . Therefore, $\{\Delta I_s\}_{s=0}^t$ is indeed a sequence of i.i.d. random variables.

4.5.3 Large deviations of information

It is possible to weight the Markov transition matrix (4.38) with the values of ΔI_s from (4.36) and (4.37) to obtain the tilted matrix,

$$\Omega'(k) = \begin{pmatrix} e^{-bk}qr & e^{-ck}qw & e^{-dk}pw & e^{-ak}pr \\ e^{-ak}pr & e^{-dk}pw & e^{-ck}qw & e^{-bk}qr \\ e^{-bk}qr & e^{-ck}qw & e^{-dk}pw & e^{-ak}pr \\ e^{-ak}pr & e^{-dk}pw & e^{-ck}qw & e^{-bk}qr \end{pmatrix},$$
(4.40)

which has principal eigenvalue

$$\lambda(k) = e^{-ak}pr + e^{-bk}qr + e^{-ck}qw + e^{-dk}pw.$$
(4.41)

The logarithm of (4.41) is taken as the SCGF $\xi(k)$ and Legendre transformed according to (4.22) into the rate function E(i). Fig. 4.6 shows this rate function plotted with data from simulation for a two-site system. The datapoints in Fig. 4.6 are $-1/t \ln p(i)$ where p(i) is the probability density function, estimated from the data using kernel density smoothing. In the long-time limit the data converge well to the rate function.

4.6 *L*-site models

For systems with three or more sites, the process $\{(X_s, Y_s)\}_{s=0}^t$ is still a Markov chain with a stationary state. However, unlike the L = 2 case, the information gained in each measurement along a trajectory $\{\Delta I_s\}_{s=0}^t$ is not a sequence of i.i.d. random variables. Eq. (4.15) cannot be reduced to a simpler form as the M_y matrices for $L \ge 3$ have more than one non-zero eigenvalue, and cannot in general be written in a form that allows cancellations like (4.34). Indeed, $\{\Delta I_s\}_{s=0}^t$ is also not a Markov chain because each value depends on the entire measurement history \mathbf{y}_s up to that point, which is a larger object at each successive value of s.

4.6.1 General behaviour

From simulation we see that the information $\Delta \mathcal{I}_t$ reaches linear growth after some initial transient, and so we expect the information increments ΔI_s to also converge to a stationary distribution. As the matrices used to calculate ΔI_s all have principal eigenvalue $\lambda_{\max}^L < 1$ (for all parameters except 1 - p = r = 1), the system should exhibit an exponential decay of correlations. This allows us to assume that, at



Figure 4.6: Large deviation rate function (4.22) for L = 2, p = 0.2, and r = 0.9. The solid black line shows the theoretical curve obtained from Eq. (4.41) and (4.22). Symbols represent data sampled at different finishing times t. Data obtained from an ensemble of 10^7 realisations.

long times, ΔI_s only has significant dependence on a finite number of the previous measurement outcomes / events. This observation that only a finite number of previous events are relevant will be the basis of the approximation in Sec. 4.6.3.

We numerically investigate the behaviour of ΔI_s for systems with $L \geq 3$, specifically we focus here on results for L = 10. Other L values (larger and smaller) do not significantly differ in their general behaviour or numerics. Fig. 4.5 shows the CDF for ΔI_s , for various system sizes up to L = 20. The shape of the function and the position of the minimum, maximum and most likely intermediate values do not vary significantly. The scaling of the most likely, minimum and maximum of ΔI_s with L is detailed in Sec. 4.4.3 and 4.6.2.

Fig. 4.7 shows a scatter plot of ΔI_{s+1} against ΔI_s for L = 10. Independent random variables plotted this way would produce a symmetric cloud or grid of points. However this plot features diagonal patterns which correspond to correlation between the two variables. The projected probability densities are shown as histograms along the axes. The histograms suggest that ΔI_s and ΔI_{s+1} are identically distributed as expected. To check how successive values of ΔI_s are correlated and hopefully verify that the ΔI_s s are independent, we compute the average sample autocorrelation function (ACF), defined as

$$\operatorname{ACF}_{\Delta I_s}(\tau) = \langle \frac{t \sum_{s=1}^{t-\tau} (\Delta I_s - \langle \Delta I \rangle) (\Delta I_{s+\tau} - \langle \Delta I \rangle)}{(t-\tau) \sum_{s=1}^{t} (\Delta I_s - \langle \Delta I \rangle)^2} \rangle.$$
(4.42)

Fig. 4.8 shows the ACF of ΔI_s after an initial transient period for biased and unbiased cases with 95% white-noise confidence intervals. The autocorrelation functions show significant negative correlation between ΔI_s and ΔI_{s+1} , but beyond this no significant correlation. That is, the information gained at successive time steps is anti-correlated. This is because when an incorrect measurement is made the next measurement is likely to be correct, as correct measurements are more probable, and the amount of information gained will be positive (see Sec. 4.4.3).

While successive correct measurements do each contribute positive amounts of information, they do not differ as radically as the change from a negative to positive amount of information. Fig. 4.8 also suggests that biased systems are less strongly anti-correlated than unbiased. In Sec. 4.6.3 we use the one time-step correlation and distribution of ΔI_s as grounds for constructing a single-step Markov chain model of $\{\Delta I_s\}_{s=0}^t$ that we believe captures most of the relevant features.



Figure 4.7: Scatter plot of ΔI_{s+1} against ΔI_s for L = 10, p = 0.5, r = 0.9 to illustrate correlations as explained in the text. Histograms on the axes show the density of points on the plot. Datapoints are from consecutive time steps sampled from late times in 10^7 independent realisations.



Figure 4.8: The sample average autocorrelation function of ΔI_s (4.42) for biased and unbiased systems. Data obtained from 10⁷ realisations. The ACF shows significant negative correlations between successive time steps. 95% confidence intervals plotted as blue dashed lines. L = 10 and r = 0.9 in both cases.



Figure 4.9: Points show numerical results for $Max(\Delta I_s)$ and $-Min(\Delta I_s)$ (from 10⁷ trajectory realisations of length t = 500) and predicted values given by (4.43) and (4.44) for varying L and for p = 0.5, r = 0.9.

4.6.2 Maximum and Minimum of ΔI_s

Fig. 4.9 shows numerical results for the maximum and minimum of ΔI_s from 10⁷ realisations up to t = 1000 for varying L. The minimum amount of information per time step is obtained when an incorrect measurement is made and we argue that, as in the L = 2 case (see Eq. (4.37)) its value is given by

$$\operatorname{Min}(\Delta I_s) = \ln \frac{w}{\lambda_{\max}^L}.$$
(4.43)

The maximum value is observed to be exactly the minimum value reflected across the lower baseline value for information gain (4.25), i.e. it is given by,

$$\operatorname{Max}(\Delta I_s) = \ln \frac{r}{\lambda_{\max}^L} - \operatorname{Min}(\Delta I_s) = \ln \frac{r}{w}.$$
(4.44)

This holds for all $L \ge 4$ as shown by the numerical results in Fig. 4.9. For $L \le 3$, the observed discrepancy is probably due to the fact that incorrect measurements are more constrained e.g. any barrier placement in a two-site system will interfere



Figure 4.10: The large deviation rate function (4.22) for L = 10, p = 0.5, r = 0.9. Data obtained from 10^7 realisations. As in Fig. 4.6, the data points are obtained by using kernel density smoothing to estimate the probability density function. Lines show rate functions obtained from the Markov approximation with different numbers of bins. The lines for 12 and 24 bins are not distinguishable at this scale. Points represent data sampled at different finishing times t. The cutoff for positive deviations is explained in the text and given in Eq. (4.45), and is represented by a vertical line.

with the particle motion.

4.6.3 Markov chain model

To obtain an approximate rate function, we assume that after an initial transient, $\{\Delta I_s\}_{s=0}^t$ is described by a stationary Markov process taking values in a continuous range. We define a finite state space Σ by coarse-graining this range and replacing ΔI_s by its expected value in each bin. The transition matrix for this chain is then obtained by binning data from a scatter plot such as Fig. 4.7 into these states and normalising the number of counts in each bin. To count the information gain, the new Markov matrix on the state space Σ is then weighted with the value of ΔI_s in the target state. The SCGF (and thus the rate function) can then be obtained from the largest eigenvalue of this tilted transition matrix.

To check the method it can be shown that for L = 2, as the number of bins


Figure 4.11: The large deviation rate function (4.22) for L = 10, p = 0.2, r = 0.9. Data obtained from 10^7 realisations. As in Fig. 4.10, the data points are obtained by using kernel density smoothing to estimate the probability density function. The solid line shows the rate function obtained by using 24 bins. Here there is no cut-off present in the data, although presumably one still exists

is increased and more data is used in the scatter plot, the method converges to the analytically obtained rate function for that case. Fig. 4.10 and Fig. 4.11 show the rate functions obtained from data through this method for L = 10 plotted alongside the same data displayed as points. Fig. 4.10 shows convergence of the estimated rate functions with increasing number of bins, which appears to be consistent with the data. This confirms the validity of the one-step Markov approximation for a wide range of fluctuations. Fig. 4.11 shows the same plot for a different value of p in order to confirm that the method applies to unbiased systems. However, as can be seen in Fig. 4.10, for certain parameters the data indicate a cut-off in the rate function for positive deviations that the Markov approximation does not predict. In the next subsection we explain this feature by noting that large fluctuations of ΔI_s do not accumulate in the way that the Markov model allows.

4.6.4 Beyond Markovian analysis

To correct the numerically obtained rate functions shown by solid lines in Fig. 4.10 we must consider the maximum possible value for I_t/t . Fig. 4.7 suggests that it is possible to obtain large amounts of information on consecutive time steps (the top right corner of this scatter-plot has a small but non-zero population). However, investigation of individual trajectories of the system reveals that consecutive large positive amounts can only be obtained after consecutive large negative amounts (see Sec. 4.4.3). This is not reflected in the ACF in Fig. 4.8 owing to the fact that these events occur very rarely, as can be seen from the marginal histograms in Fig. 4.7.

The maximum value of $\Delta I_t/t$ is obtained by measuring correctly every time step. This maximum value is given by

$$\frac{\Delta \mathcal{I}_{\max}^u}{t} = \ln \frac{r}{\alpha},\tag{4.45}$$

where α is numerically obtained from the M_y matrices for that system. This is discussed in Sec. 4.4.3.

Fig. 4.10 shows the cut-off value for an unbiased system with a black vertical line. Unlikely trajectories in biased systems that always step against the bias can generate large positive deviations of $\Delta I_t/t$ and so the rate function cuts off at higher values. The cut-off is therefore less relevant when comparing data for biased systems to the predicted rate function. This is corroborated by Fig. 4.11 which shows no cut-off in range of the data. There is still a cut-off for unbiased systems that can be obtained via the same method, but it does not appear in the simulation data since the trajectories that generate the corresponding values are even rarer than the analogous trajectories in unbiased systems.

The ΔI_s process is clearly not a one-step Markov chain, and the rate function obtained this way is also limited by finite sampling of transitions and limitations on the number of bins used. However, the one-step Markov model gives a rate function that converges reasonably quickly with increasing number of data points and together with the cut-off, captures well the shape of the sampled data in a way that a Gaussian or i.i.d. approximation would not. An *n*-step Markov chain model might also capture this but it appears that simply including the cut-off at the maximum value is sufficient to obtain a good approximate rate function.

4.7 Discussion

4.7.1 Summary

The information gain $\Delta \mathcal{I}_t$ is a quantity recently introduced in the analysis of feedback systems Parrondo et al. [2015]; Horowitz and Vaikuntanathan [2010] and studied as a component in the development of information thermodynamics and information engines Horowitz and Parrondo [2011]. The fluctuation properties of this quantity are relevant when considering information processing in feedback devices; the quantity of information gained is directly proportional to the work required to delete that information from the feedback device's memory.

In this chapter we have studied a simple model of an information engine and obtained an exact analytical expression of the large deviation rate function for information gain $\Delta \mathcal{I}_t$ in a two-site system. For larger systems we have shown that a one-step Markov approximation captures most of the relevant details of the large deviations. We are also able to predict the cut-off of this rate function by considering the maximum amount of information that can be obtained.

Significantly, the one-step Markov approximation allows us to easily obtain an approximation of the rate function from data by sampling the information gain at consecutive time steps. This is computationally easier than directly sampling the distribution of $\Delta I_t/t$ (which requires very long times or cloning-type algorithms Rohwer et al. [2014]) but together with the theoretically predicted cut-off seems to provide a consistent estimate of the shape of the large deviation rate function.

4.7.2 Possible extensions

The work presented in this chapter is easily amenable to two extentions. Firstly, the methods discussed here could be used to analyse different control systems, or the trajectory sampling could be augmented in order to obtain more accurate large deviation rate functions so as to better understand the fluctuations of information.

Information reservoir

In Sec. 4.1, we mentioned that instead of measurement based feedback, the information resource exploited by the feedback mechanism could be an 'information reservoir' such as a memory register than can be written into. Such reservoir based models are well studied Mandal and Jarzynski [2012]; Barato and Seifert [2013]; Deffner [2013]; Mandal et al. [2013], Recent work Barato and Seifert [2014a]; Shiraishi et al. [2016] has shown that the reservoir-based and measurement-based models are essentially the same and that the different expressions of the second law for information thermodynamics can obtained and described by a single 'master' second law. It would be possible to alter our model here for open-loop control mediated by an information reservoir and study the fluctuations of entropy and information flow. This may be potentially easier for a reservoir-driven model since there are no non-Markovian effects to complicate the computation of the rate function.

Population sampling

Cloning algorithms Lecomte and Tailleur [2007]; Giardina et al. [2011] operate by creating a population of simulations and only allowing simulations that show the behaviour of interest to continue. Simulations that are following uninteresting trajectories or behaviour are replaced with 'clones' of the simulations that are showing more interesting behaviour. By counting whenever a simulation is culled, the uninteresting behaviour can still be weighted and 'observed', but without wasting further computational resources (time and memory). This technique can be used to observe rare events without having to wait long times or use unreasonably large ensembles. Using a cloning algorithm, the rare fluctuations or trajectories of a system can be observed via a biasing of the system that allows more thorough sampling of the trajectory-space without artificially influencing the statistical weight of the rare events.

4.7.3 Conclusion

The rate functions obtained here demonstrate that the information gained by the measuring device in a simple feedback system shows a strong asymmetry around the mean. The cut-off value for this rate function is sensitive to the dynamics of the system, namely whether the particle motion is symmetric or asymmetric. It would be of interest to study other information engines to check whether these findings are generic and to what extent the Markov approximation is applicable in other systems. The large deviation rate function offers the possibility to explore detailed fluctuation relationships beyond (3.34) for ΔI_t . To obtain a detailed fluctuation relationship, care must be taken in deciding how to meaningfully time-reverse a feedback system.

In this context, we remark that the feedback process included in our model considers a measurement process that is fully integrated into the system, meaning that neither the measurement process nor the action of the feedback affects the instantaneous state of the controlled system. Real-world measurement and feedback processes may influence the state of the system by causing some 're-action' e.g., 'nudging' individual particles when the measurement is made or when the barrier is moved. However, we believe the framework of a reaction-less measurement process is useful as a first step towards understanding the fluctuations of information in closed-loop feedback devices.

Consideration of more complicated feedback systems where the system reacts to the measurement process or has more complex dependence on multiple control parameters, could help to establish which details are important for real physical systems. In fact, the relationship between the controlled system and the controller could even potentially be mirrored allowing two systems to enact control on each other. In the next chapter, we will consider a model for such a feedback device with mirrored or 'mutual' feedback. Contrastingly to the model in the current chapter, the distinction between the controlled and controller system is less obvious in the case of mutual feedback⁹.

⁹In fact, such a distinction proves to not really exist.

Alice laughed: "There's no use trying" she said; "one can't believe impossible things."

"I daresay you haven't had much practice," said the Queen. "When I was younger, I always did it for half an hour a day. Why, sometimes I've believed as many as six impossible things before breakfast."

Lewis Carroll, Alice in Wonderland

5

Modelling of the measurement system

In this chapter, we discuss how feedback is more concretely modelled and explore the limitations of this framework. We discuss how one can explicitly model the measurement system and verify that the gains from feedback control are compensated by the cost of performing measurements. In Sec. 5.1, we restate the relationship of our framework to thermodynamic concepts like work and how a controlled system can do work by extracting heat from the environment whilst a measurement system dissipates heat. In Sec. 5.2 we describe an explicit model of measurement and feedback and restate the functionals that can be used to measure the relevant quantities. In Sec. 5.3 we show data from simulation that confirms the intuition that work is extracted at the cost of work being spent elsewhere. In Sec. 5.4 we discuss a recently proposed 'double-dæmon' Ford [2016] and how the use of feedback on both the controlled system and the measurement system might push past the bounds stated in Sec. 5.1. Finally, in Sec. 5.5 we compare these findings to related frameworks for explicitly modelling feedback and comment on experimental set-ups.

5.1 Work and the 2nd law

It is known that feedback control is capable of enabling the extraction of work from an isolated system Abreu and Seifert [2011]; Koski et al. [2014a,b]. By using information obtained about the system, control protocols can be employed which exploit that reduction in uncertainty of the system. Indeed, it is exactly this statement that the generalised second law (3.35) represents. If one considers the second law to only relate to the entropy production S_{tot} , then feedback 'violates' the second law. However, if one generalises the law to include the information ΔI_t possessed by the feedback device, then the total quantity still obeys a 'second law-like' inequality (3.35). Interpreted simply, the maximum 'entropy change' (be it work or stochastic entropy production, depending on the choice of boundary terms) that can be obtained is bounded by the amount by which the uncertainty about the system can be reduced by the control device.

5.1.1 Extraction of work through feedback

Consider a system evolving along a trajectory up until some measurement time t_m , $\mathbf{x}_{t_m} = \{x_s\}_{s=0}^{t_m}$. Prior to measurement, there will be some uncertainty on the specific state of this chain. A single measurement y is made of the state, and let us assume that both x and y have the same state space χ . Through the measurement, a correlation is established between y and the state of the system at the instant of measurement, x_{t_m} . This reduction of uncertainty (or information gain) can be used to inform some protocol for shifting the parameters to extract energy from the system (or otherwise drive it towards some desired outcome) in the period $t_m < s < t$. The average information gained about the system is quantified by the mutual information between the system and measurement device at the instant of measurement, that is

$$I[X_{t_m}; Y_{t_m}] = \langle \Delta \mathcal{I}[X_{t_m}, Y] \rangle, \qquad (5.1)$$

where X_{t_m} and Y_{t_m} are random variables representing the system state (at the moment of measurement t_m) and the measurement outcome respectively, and

$$\Delta \mathcal{I}[x,y] = \ln \frac{p(y \mid x)}{p(y)}, \quad \forall x, y \in \chi.$$
(5.2)

By exploiting the information gained about the system state according to some 'exploitation' protocol a negative amount of 'work' W_x can be done by the system which translates to a positive amount of energy that can be extracted from the system. The system can be allowed to relax back to equilibrium, meaning that there is zero change in internal energy and the only nonzero quantities will be the work done (and entropy production). Abreu and Seifert [2011] show that the amount of work that can be extracted after the measurement is bounded by the mutual information between the system and the measurement at the instant of measurement, i.e.

$$\langle \mathcal{W}^X \rangle \ge -k_B T \langle \Delta \mathcal{I} \left[X_{t_m}, Y_{t_m} \right] \rangle.$$
 (5.3)

In a stochastic system, the work done \mathcal{W}^X is given by the current-like term of \mathcal{R}_F as in Sec. 2.7.1. Eq. (5.3) tells us that the maximum energy we can extract from a system about which we possess some information is bounded by the mutual information, similar to (3.35). Further, (5.3) suggests that in the presence of measurement based feedback, the work done can be *negative*, meaning that we can *gain* useful work from such a system. For our models we can set $k_BT = 1$, and hence (3.35) and (5.3) are nearly equivalent, with the only differences being the choice of boundary terms for the functional as described earlier in Ch. 2 and the fact that previously, $\Delta \mathcal{I}_t$ was the total information gained by a series of measurements whereas here $\Delta \mathcal{I}$ is the mutual information between the two systems after a single measurement. Since for us, entropy production and work done are both based on the functional \mathcal{R} but with different boundary conditions, during this chapter we will refer to 'entropy production' and 'work done' somewhat interchangeably. However, it is important to note that in other models/frameworks, the two quantities cannot necessarily be simply conflated this way.

Eq. (5.3) shows that if we are in possession of an information resource, we can extract energy through feedback. However, we do not get this information for free. A measurement essentially creates a correlation between the state of the measurement device and the state of the system that is to be measured. This correlation is quantified by the mutual information, and creates an information resource. Since every 'bit' of information costs energy to create, according to the views of Szilard [1929] and Brillouin [1951]¹, then the process of measurement costs energy² proportional to the amount of information gained in the measurement. The finding that the process of measurement necessarily generates entropy or costs work, is illustrated in Granger and Kantz [2011]. The bounds on the work done in measurement are Sagawa and Ueda [2009]

$$\langle \mathcal{W}_m \rangle \ge k_B T \langle \Delta \mathcal{I} \left[X_{t_m}, Y_{t_m} \right] \rangle,$$
(5.4)

¹Or erase according to Landauer [1961].

 $^{^2\}mathrm{Either}$ in the process of measurement or when the information is 'deleted'.

where $\langle \Delta \mathcal{I} [X_{t_m}, Y_{t_m}] \rangle$ is the mutual information that quantifies the correlation between the system and measurement device after the measurement process ends at t_m .

If we consider the entire process then, on average we can extract $\langle \mathcal{W}^X \rangle$ by using the information from measurement, but we must on average pay $\langle \mathcal{W}_m \rangle$ in order to create that information. Combining (5.3) and (5.4), we obtain,

$$\langle \mathcal{W}^X \rangle + \langle \mathcal{W}_m \rangle \ge 0.$$
 (5.5)

Even operating optimal feedback control (i.e. we achieve the bounds in (5.3) and (5.4)) leaves us at best empty handed. It costs as much energy to make the measurement as we can extract through feedback control. In the following sections, we model the measurement device and measurement process explicitly in order to demonstrate this.

5.2 Modelling the measurement process

In Chapters 3 and 4, we never explicitly described the measurement process. That is, we simply considered that the measurement creates a correlation between the two variables X_s and Y_s without modelling the dynamics of the system represented by Y or how it becomes correlated with X. In this section we first briefly describe a recently studied Langevin model that can be used to model the measurement process. We then present a discrete-space and -time version of the model dæmon that allows us to use our previous framework to study its operation.

5.2.1 Langevin model

We here detail the model proposed in Ford [2016], although we will use slightly different notation. The key benefit of this model is that it explicitly models the measurement process as a phase in the evolution of the total system with its own set of dynamics, similar to models in Granger and Kantz [2011]; Mandal and Jarzynski [2012]; Barato and Seifert [2013]; Strasberg et al. [2013]. The model we will now discuss includes the measurement device, measurement process and feedback as explicit components of the system and phases of its evolution. The measurement process refers to a process that results in creating a correlation between the states of the two components, and 'control' refers to the change in parameters that is enacted thereafter. One component of the system has control enacted on it and so can be called the 'controlled' system, whereas the other component is used as a

'measurement device' or 'daemon'.

The model consists of two 1-D oscillators connected to tethers via Hookean springs. Both oscillators respond to noise in the environment. The oscillators are coupled via a third Hookean spring. The oscillator positions take values in \mathbb{R} . A trajectory in this system is written $(\mathbf{x}_t, \mathbf{y}_t) = \{x_s, y_s\}_{s \in [0,t]}^3$. The dynamics are modelled with a pair of coupled stochastic differential equations (SDEs). Specifically, a pair of coupled and over-damped Langevin equations

$$\frac{\mathrm{d}X_s}{\mathrm{d}s} = -K_x(s) \left[X_s - \lambda_x(s) \right] - K_{x,y}(s) \left[X_s - Y_s \right] + \sqrt{2}\xi_x(s), \tag{5.6}$$

$$\frac{\mathrm{d}Y_s}{\mathrm{d}s} = -K_y \left[Y_s - \lambda_y\right] - K_{x,y}(s) \left[Y_s - X_s\right] + \sqrt{2}\xi_y(s),\tag{5.7}$$

where $K_{x,y}(s)$, $K_x(s)$, $\lambda_x(s)$ are the control parameters of the model which characterise the spring constants and tether-position of the x oscillator. K_y and λ_y remain constant throughout. The $\xi(s)$ terms are white noise with properties $\langle \xi_x(s) \rangle = 0$ and $\langle \xi_x(s)\xi_x(s') \rangle = \delta(s-s')$ Gardiner et al. [1985]; Van Kampen [1992]. When the control parameters are fixed, the system corresponds to two coupled Ornstein-Uhlenbeck processes and therefore the stationary distributions of the oscillators under these dynamics are normal distributions characterised by the control parameters. The stationary distribution of the entire system is a bivariate Gaussian distribution.

The system evolves through four phases. In the first phase, the oscillators are not coupled and allowed to approach stationarity independently. The second 'measurement' phase follows in which the coupling strength between the oscillators is slowly and linearly raised before being instantaneously taken back to zero at the end of this phase, uncoupling the oscillators but leaving their positions correlated. The y oscillator position is used to inform the parameter values for the next phase. In the third 'exploitation' phase, the parameters $K_x(s)$ and $\lambda_x(s)$ instantaneously shift to exploitation values and then slowly return to their initial values. The final phase is a relaxation, allowing the system to reach the same stationary distribution from before the measurement phase and allowing the system to be operated again in a cycle.

We can use the established framework from Ch. 2 to construct related model in discrete time and space, which we now proceed to detail.

³Note the difference in indexing since this is a continuous time process.

5.2.2 Discretised model

In analogy with oscillators moving freely in space, we consider two particles moving randomly on discrete lattices. Our model is a discrete-time process with discrete state space $\chi = \{ma : m \in \mathbb{Z}\}$ where a > 0 is some lattice spacing. A trajectory of this process is written $(\mathbf{x}_t, \mathbf{y}_t) = \{x_s, y_s\}_{s=0}^t$. The particle dynamics are formed of two components. The two particles are subject to 'forces' acting on them in a similar fashion to the Langevin model, which is described by a deterministic drift term

$$d_x(x, y, s) = -K_x(s)(x - \lambda_x(s)) - K_{x,y}(s)(x - y) \quad \in \mathbb{R},$$
 (5.8)

$$d_y(y,x,s) = -K_y(y-\lambda_y) - K_{x,y}(s)(y-x) \in \mathbb{R}.$$
(5.9)

It should be noted that the drift term generates real numbers \mathbb{R} , whereas the particle can only take positions on the lattice. Under purely deterministic dynamics, the particles would move to new sites based on these 'drift' terms, i.e. the X particle would move from site x_s to x_{s+1} if $x_{s+1} - a/2 \leq x_s + d_x(x, y, s) < x_{s+1} + a/2$. However, in analogy with the stochastic environmental noise, we add a gaussian noise centred on the location of the drift term. The transition probabilities are then given by

$$\omega_X(x \to x' \mid y, s) = \Pr\left[x' - \frac{a}{2} \le \mathcal{N}(x + d_x(x, y, s), 1) < x' + \frac{a}{2}\right], \quad (5.10)$$

$$\omega_Y(y \to y' \mid x, s) = \Pr\left[y' - \frac{a}{2} \le \mathcal{N}(y + d_y(y, x, s), 1) < y' + \frac{a}{2}\right], \quad (5.11)$$

where \mathcal{N} is a Gaussian noise centred around the drift term. To be clear, the probability of a transition x to x', for example, is the probability that the value of the Gaussian noise with mean $x+d_x$ is between the lattice positions x'-a/2 and x'+a/2. The position of the other particle and the time s acts as a control parameter for each system as they determine the drift term. Since the jumps are independent, we can write the transition probability of transitions from states (x, y) to (x', y') as

$$\omega((x,y) \to (x',y')) = \omega(x \to x')\omega(y \to y').$$
(5.12)

A cartoon diagram of a single jump using these dynamics (using the X particle as an example) is shown in Fig. 5.1

As in the Langevin model, this discrete model evolves through four distinct 'phases'. Firstly a phase in which x and y are allowed to approach stationarity independently. This is followed by a measurement phase $0 \le s \le t_m$, in which the



Figure 5.1: Cartoon diagram depicting the dynamics for the X particle described by (5.8) and (5.10). The X particle occupies a site on the lattice shown by the blue circle. The drift term (5.8) is shown as a black arrow. The current site plus the drift term (i.e. $x + d_x(x, y, s)$) determines the centre of the noise term shown as a pink line. The noise is shown as a pale blue Gaussian PDF around that centre. The transition probability (5.10) is shown by highlighting in red the area under the distribution (and centred around a lattice site x') that corresponds to the probability of transitioning to the site x'.

particles are coupled; an exploitation phase $t_m < s \leq t_{ex}$, in which the parameters $K_x(s)$ and $\lambda_x(s)$ are varied and a final relaxation phase $t_{ex} < s \leq t$, during which the system is allowed to re-equilibrate. The specifics of these phases are discussed in Sec. 5.2.4.

These dynamics are slightly contrived to produce a model analogous to the Langevin model, but in discrete time and space so that we can use the framework we have established for functionals of Markov chains in previous chapters. However, since the dynamics are so closely related to the original system, we also know that they share some important properties. Namely, for each fixed set of parameters $K_x, K_y > 0, K_{x,y} \ge 0$, there is a unique stationary distribution. Both particles are always localised despite the state space being infinite. We also know that the

stationary distribution for fixed parameters is unique since the Gauss noise means that process is irreducible⁴. Finally, since the dynamics are so closely related to the Langevin model, the stationary distributions are expected to be close to a discrete version of the multivariate Gaussian.

5.2.3 Functional for the process

In order to track the entropy production of the two oscillators, we use the trajectory functionals defined in Sec. 2.3, but generalised for control as in Sec. 3.1. We write

$$\mathcal{R}_{F}^{X} = \ln \frac{\mu_{0}^{x}(x_{0})}{\mu_{t}^{x}(x_{t})} + \sum_{s=0}^{t-1} \ln \frac{\omega(x_{s} \to x_{s+1} \mid y_{s}, s)}{\omega(x_{s+1} \to x_{s} \mid y_{s}, s)},$$
(5.13)

$$\mathcal{R}_{F}^{Y} = \ln \frac{\mu_{0}^{y}(y_{0})}{\mu_{t}^{y}(y_{t})} + \sum_{s=0}^{t-1} \ln \frac{\omega(y_{s} \to y_{s+1} \mid x_{s}, s)}{\omega(y_{s+1} \to y_{s} \mid x_{s}, s)},$$
(5.14)

where μ_0^i and μ_t^i are the stationary distribution of subsystem *i* at the beginning and end of the process. We are interested in control protocols that have the same control parameter values at the beginning and end of the protocol, and so starting from stationarity and given enough time to relax after the protocol, the system will begin and end in the same distribution. That is, $\mu_0^i = \mu_t^i$. In the initial condition the oscillators are not coupled to one another $(K_{x,y}(0) = 0)$ and are only coupled to their respective tethers with the same strength, which are both placed at the origin ($\lambda_x = \lambda_y = 0$ and $K_x(0) = K_y = 1$), so the stationary distributions of both oscillators will be identical and independent. The functionals \mathcal{R}_F^X and \mathcal{R}_F^Y track the entropy production of each oscillator as it evolves, and hence their sum

$$\mathcal{R}_F = \mathcal{R}_F^X + \mathcal{R}_F^Y, \tag{5.15}$$

tracks the entropy production of the entire system. We can sum the contributions to entropy production linearly⁵ like this because the functionals are independent owing to the independence of the transition probabilities (5.12).

Since the entire system begins and ends in the same stationary distribution, we can interpret this functional as related to the 'work done' as we did before in Sec. 2.7.1, $\mathcal{R}_F = \mathcal{W} - \Delta F$ where the current term is the 'work done' and the boundary term is the change in free energy ΔF . Of course, the system can only ever *approach* stationarity, but we assume that the discrepancy between the initial

 $^{^{4}}$ Every state is accessible from every other state, and since the normal distribution has full support, in this system every state is accessible from any other state in a single transition.

⁵Even though the processes are correlated.

and final state is negligible and so the change in free energy should be zero in the ensemble average, i.e. $\langle \Delta F \rangle = 0$. Thus, we only need to look at the current-like terms of the functional, \mathcal{R}_F . We expect that, given a good control protocol, $\langle \mathcal{R}_F^X \rangle$ should be negative after exploitation whereas $\langle \mathcal{R}_F^Y \rangle$ will be positive and should at least compensate the negative $\langle \mathcal{R}_F^X \rangle$ such that $\langle \mathcal{R}_F \rangle \geq 0$. The protocols for exploitation are detailed in Sec. 5.2.4.

We can also write for the final value of \mathcal{R}_F (which we write as $\mathcal{R}_F(t)$)

$$\mathcal{R}_F(t) = \mathcal{W}_m + \mathcal{W}^X, \tag{5.16}$$

where \mathcal{W}_m is the work done during the measurement phase and \mathcal{W}^X is the work gained during the exploitation phase. \mathcal{W}_m is given by $\mathcal{W}_m = \mathcal{R}_F(t_m) - \mathcal{R}_F(0) = \mathcal{R}_F(t_m)$, whereas $\mathcal{W}^X = \mathcal{R}_F(t) - \mathcal{R}_F(t_m)$.

5.2.4 Protocol for Measurement and Exploitation

We now describe the specifics of the measurement and exploitation of this system. We will here use notation and concepts applicable to the discretised model. Similar expressions and concepts are obtained for the continuous system.

Throughout this chapter we use the word 'protocol' to refer to the evolution of all parameters over the entire process. We will use 'measurement phase' to refer to the time during which the oscillators are coupled and correlated with one another and 'measurement protocol' to refer to the evolution of the parameters during this time. Similarly, we use 'exploitation phase' and 'exploitation protocol' to refer to the duration of time and evolution of parameters after the measurement phase, that ideally results in extraction of 'energy' from the system, i.e. negative work done. The specific values used for parameter evolution are summarised in Sec. 5.2.5.

Equilibration

Prior to the measurement and exploitation phase, we first initialise both oscillators at the origin on the lattice and allow time⁶ for the system to approach a stationary state. During this time, the parameters are $K_{x,y} = 0$ and $K_x = K_y = 1$ and $\lambda_x = \lambda_y = 0$. The oscillators are uncoupled and will approach independent identical stationary Gaussian distributions.

⁶Ideally an infinitely long period of time.

Measurement

The oscillators will be at equilibrium following the previous equilibration phase. Over the duration of the measurement phase, $K_{x,y}$ is raised from 0 to some finite positive value k over the time s = 0 to $s = t_m$. All other parameters remain constant during this time. At the end of the measurement phase, the oscillator position variables X_s and Y_s will be correlated. At the end of this phase, we instantly switch $K_{x,y} = 0$ in preparation for the next phase.

During the measurement phase, we expect R_F^X , R_F^Y and hence R_F to be positive and increasing up until t_m . This is owing to the fact that for the process to create a correlation between the two systems, it must necessarily be entropy producing Granger and Kantz [2011]; Sagawa and Ueda [2012]. The smallest average entropy production is in the limit of a quasistatic coupling, where the coupling is introduced infinitely slowly (i.e. $t_m \to \infty$) Spinney and Ford.

Exploitation and Relaxation

After the measurement phase, we begin the 'exploitation phase'. To select a feedback protocol, the value of y is used to determine what values K_x and λ_x should take. Instantaneously at the beginning of this phase, we switch $K_x = 1+k$, $\lambda_x = ky_{t_m}/(1+k)$. The y oscillator is allowed to relax and returns to its original state. The exploitation takes place during the phase $t_m < s \leq t_{\text{ex}}$. During this time, the changed control parameters are returned to their original values.

Finally, in the phase $t_{\text{ex}} < s \leq t$, we allow the system to equilibrate. Since all control parameters have been returned to their original values, the system will relax back to the distribution it was initialised in. In reality, the system will only ever asymptotically approach the stationary distribution and hence the final distribution will not be exactly the same as the initial distribution. In simulation, we use a long enough relaxation period that the difference will be assumed to be so small as to be negligible. The system thus begins and ends in the same state and the total change in system entropy (and free energy) should be zero.

During exploitation and relaxation, we expect that \mathcal{R}_F^X can reach negative values on average owing to the exploitation, while \mathcal{R}_F^Y will be positive such that the final value of \mathcal{R}_F is positive on average.

5.2.5 Optimal protocols

We now discuss the optimal protocols for measurement and exploitation. These protocols are considered 'optimal' as they minimise the work that must be done in

s	$K_{x,y}$	K_x	λ_x
s = 0	$K_{x,y} = 0$	$K_x = 1$	$\lambda_x = 0$
$0 \le s < t_m$	$K_{x,y} \to k$	$K_x = 1$	$\lambda_x = 0$
$s = t_m$	$K_{x,y} = 0$	$K_x = 1 + k$	$\lambda_x = \frac{ky_{t_m}}{1+k}$
$t_m < s \le t_{\rm ex}$	$K_{x,y} = 0$	$K_x \to 1$	$\lambda_x = \frac{ky_{t_m}}{1+k}$
$t_{\rm ex} < s \le t$	$K_{x,y} = 0$	$K_x = 1$	$\lambda_x = \frac{ky_{t_m}}{1+k}$

Table 5.1: Evolution of control parameters. Here the ' \rightarrow ' notation is to be read that a parameter changes from its previous value (its value from the row above) to the value written after the \rightarrow , and does so linearly over the time period specified in the first column. For example, $K_x \rightarrow 1$ in the 4th row and 3rd column means that K_x changes linearly from 1 + k to 1 over the time period t_m to t_{ex} .

measurement and maximise the work that can be extracted from the system. That is, they achieve the bounds in (5.3), (5.4) and (5.5). For measurement, it has been shown in Granger and Kantz [2011]; Sagawa and Ueda [2009] that raising the coupling linearly and infinitely slowly reaches the lower bound in (5.4). Optimal control protocols for the exploitation of a single oscillator modelled by Langevin equations such as (5.6) and (5.7) have been previously studied in Granger and Kantz [2011]; Abreu and Seifert [2011]; Sagawa and Ueda [2012]. Since our model is a discretisation of the scheme discussed in the literature, we use the same exploitation protocol suggested in Abreu and Seifert [2011]. Specifically, the protocol we employ for our system is shown in Table 5.1. There is not necessarily any reason to suspect that this protocol will be 'optimal' for our discretised model, save for the fact that in appropriate limits (lattice spacing $a \to 0$ and time stepping $\Delta s \to 0$) our model should scale to the Langevin model ((5.6) and (5.7)) for which the protocol is optimal.

5.3 Simulation

In this section we present data obtained from numerical simulation of the coupled oscillator system. We first check the values of the trajectory functionals $\langle \mathcal{R}_F \rangle$, $\langle \mathcal{R}_F^X \rangle$ and $\langle \mathcal{R}_F^Y \rangle$ over the course of the entire protocol. We then examine the limiting behaviour and verify that longer timescales for measurement and exploitation approach the bound of (5.5). That is, we provide evidence that in the quasistatic limit (all timescales $\rightarrow \infty$), one could achieve $\langle \mathcal{R}_F \rangle = 0$ or equivalently from (5.16), $\langle \mathcal{W}_m \rangle + \langle \mathcal{W}^X \rangle = 0$.



Figure 5.2: The work done over the course of a trajectory, with several different time-scales for the evolution of the control parameter. Data for the solid lines was obtained by averaging values from 10^7 simulations. The blue, black and red lines show the evolution of $\langle \mathcal{R}_F \rangle$ (5.15) for $t_m = t_e = 2000, 3000$, and 6000 time steps respectively, with relaxation periods of 4000 time steps afterwards. The total times are thus t = 8000, 10000 and 16000.

5.3.1 Data

Fig. 5.2 shows averages of $\langle \mathcal{R}_F \rangle$ (5.15) for three different timescales of the control protocol and Fig. 5.3 shows the values of $\langle \mathcal{R}_F^X \rangle$ and $\langle \mathcal{R}_F^Y \rangle$ for a single timescale. Both plots show data obtained from averaging over 10⁷ realisations. In each simulation, the oscillators' initial positions are the origin. The oscillators are then allowed to approach stationarity before the process begins properly (i.e. the control parameters start evolving). After that, the measurement phase is carried out over t_m time steps. At the end of the measurement phase, the exploitation protocol of X is carried out over $t_e = t_{\text{ex}} - t_m$ time steps, before the system is allowed to re-equilibrate during a relaxation phase. The control parameters during the relaxation period have the same values as they do during the initial equilibration. For the data in Fig. 5.2, we consider $t_m = t_e$ and plot $\langle \mathcal{R}_F \rangle$ for $t_m = 4000, 6000$, and 12000. Consistently with our expectations in Sec. 5.2.3, Fig. 5.2 shows that the functional (5.15) can never be negative on average but can approach zero as the timescales for measurement and



Figure 5.3: The values of $\langle \mathcal{R}_F^X \rangle$ and $\langle \mathcal{R}_F^Y \rangle$ for $t_m = t_e = 3000$. Data obtained by averaging values from 10^7 simulations. While $\langle \mathcal{R}_F^X \rangle$ can be negative after the process, $\langle \mathcal{R}_F^Y \rangle$ is much larger and positive.

exploitation are increased. Fig. 5.3 shows that while $\langle \mathcal{R}_F^X \rangle$ can be negative after exploitation, $\langle \mathcal{R}_F^Y \rangle$ is positive and of a much larger magnitude than $\langle \mathcal{R}_F^Y \rangle$.

The measurement phase for Fig. 5.3 takes place between time steps s = 0and s = 3000. During this time, the coupling between the oscillators is linearly increasing from K(0) = 0 to K(s) = k. During this phase, both $\langle \mathcal{R}_F^X \rangle$ and $\langle \mathcal{R}_F^Y \rangle$ are positive as expected, since the measurement process costs work. In a physical system, this corresponds to the energy cost of changing the system Hamiltonian, since changing the parameters will change the internal energy of the system. The exploitation phase follows after s = 3000 and continues until s = 6000 after which the system is allowed to relax. During the exploitation and relaxation phase, work is extracted from the x oscillator and hence $\langle \mathcal{R}_F^X \rangle$ is negative at the end of the relaxation. Interestingly, here it is not negative at the end of the exploitation phase but requires some relaxation time since the parameter has changed too quickly. Longer exploitation phases allow $\langle \mathcal{R}_F^X \rangle$ to become more negative before relaxing. During the exploitation phase, the y oscillator is simply allowed to relax and $\langle \mathcal{R}_F^X \rangle$ decreases but still remains positive overall.



Figure 5.4: Convergence of $\langle \mathcal{W}_m \rangle - \langle \Delta \mathcal{I} \rangle$ and $\langle \mathcal{W}^X \rangle + \langle \Delta \mathcal{I} \rangle$ to zero with increasing timescale. Red circles show $\langle \mathcal{W}_m \rangle - \langle \Delta \mathcal{I} \rangle$ and $\langle \mathcal{W}^X \rangle + \langle \Delta \mathcal{I} \rangle$. Error bars indicate the relative error on data points. Black diamonds and purple squares show the quantity plotted in red ±0.5 respectively to support that $\langle \mathcal{W}_m \rangle - \langle \Delta \mathcal{I} \rangle$ and $\langle \mathcal{W}^X \rangle + \langle \Delta \mathcal{I} \rangle$ indeed vanish exponentially in time. The solid line is $\propto e^{-t}$ and is included as a visual aid. As the timescales increase, $\langle \mathcal{W}_m \rangle - \langle \Delta \mathcal{I} \rangle$ and $\langle \mathcal{W}^X \rangle + \langle \Delta \mathcal{I} \rangle$ appear to show an exponential decay to zero with a log-scaled vertical axis. Data from 10⁷ simulations.

5.3.2 Limiting behaviour

From Fig. 5.2, we can see that when the overall time scale t is longer, the final value of $\langle \mathcal{R}_F \rangle$ is smaller. This is consistent with the prediction that slower protocols can achieve 'better' results. That is, if we couple the oscillators slowly enough, we can introduce the correlation with a minimum of entropy production or work done and when the exploitation timescale is long enough we can extract more work from the system.

Recall from (5.3) that the work extracted from the system is bounded by the mutual information created by the measurement. Fig. 5.4 shows data from simulation that suggests that both $\langle \mathcal{W}^X \rangle + \langle \Delta \mathcal{I} [X_{t_m}, Y_{t_m}] \rangle$ and $\langle \mathcal{W}_m \rangle - \langle \Delta \mathcal{I} [X_{t_m}, Y_{t_m}] \rangle$ asymptotically approach zero as the appropriate timescale t is increased. The straight lines followed by the data in log-scale is fairly convincing evidence that in the limits $t_m \to \infty$ and $t_{\text{ex}} - t_m \to \infty^7$, the difference between $\langle \mathcal{R}_F \rangle = \langle \mathcal{W}_m \rangle + \langle \mathcal{W}^X \rangle$ and $\langle \Delta \mathcal{I} [X_{t_m}, Y_{t_m}] \rangle$ vanishes exponentially and the process would achieve its lower bound.

⁷That is, as the measurement *and* exploitation time scales are simultaneously taken to infinity.



Figure 5.5: Control flow diagrams illustrating various feedback scenarios.a) Ordinary feedback control. The 'system' state is correlated with a 'dæmon'. The dæmon selects a feedback protocol which is then used to control the system.b) The same scenario as a) but reversed so that the dæmon benefits from a feedback protocol selected by the system.

c) The double dæmon. Both system and dæmon are correlated and their states are used to select feedback protocols for each another.

5.4 Double feedback

A model recently proposed by Ford [2016] asks why we cannot simultaneously enact an exploitation protocol on the y oscillator. This rationale behind this proposal can be understood as an observation that since the dæmon is a physical system, 'feedback' can just as equally be performed on the dæmon by selecting feedback protocols based on the state of the controlled system. If one can perform feedback in both directions, then one might ask why this cannot be done simultaneously. Fig. 5.5 shows this rationale in a control-flow diagram. Previously it was possible to extract energy from a single system equal to the energy input in creating the measurement. Now, it should be possible to extract additional energy from the second system. In effect, we 'pay' a one-off cost to create a correlation between the two systems, but can extract energy equal to that cost twice. Since both systems are acting as a feedback controller for each other, we refer to this kind of system as a 'double dæmon'.

5.4.1 Exploitation of the measurement device

As before in Sec. 5.1.1, we set $k_BT = 1$ for convenience. The work obtained from exploitation of the x oscillator is bounded by (5.3). Similarly, if we were to exploit the 'measurement' system (the y oscillator) in exactly the same way, the bounds on the work done by that exploitation should be

$$\langle \mathcal{W}^Y \rangle \ge -\langle \Delta \mathcal{I} \left[X_{t_m}, Y_{t_m} \right] \rangle.$$
 (5.17)

By combining (5.3), (5.4) and (5.17), we can write

$$\langle \mathcal{W}^X \rangle + \langle \mathcal{W}^Y \rangle + \langle \mathcal{W}_m \rangle \ge -\langle \Delta \mathcal{I} \left[X_{t_m}, Y_{t_m} \right] \rangle,$$
 (5.18)

which suggests that the overall work done on the system can be negative. To clarify, $\langle \mathcal{W}^Y \rangle$ and $\langle \mathcal{W}^Y \rangle$ are expected to be negative quantities⁸, whereas $\langle \mathcal{W}_m \rangle$ is a positive quantity equal to the work done in correlating the two systems. $\langle \Delta \mathcal{I} [X_{t_m}, Y_{t_m}] \rangle$ is the mutual information between the two systems at the end of this correlating process. This is to say that if we enact feedback symmetrically, we should be able to extract work, and that work extracted is bounded by the mutual information.

5.4.2 Double dæmon simulation

The results of simulation are shown in Fig. 5.6 and Fig. 5.7. We simulate the double dæmon by operating the same protocol as before, but instead of allowing the y oscillator to relax after measurement, we select exploitation protocols in a symmetrical manner to the selection of a protocol for x. That is, where we switch the tether of the x oscillator to $\lambda_x(t_m) = y_{t_m}k/(1+k)$, we also switch $\lambda_y(t_m) = x_{t_m}k/(1+k)$, as well as changing both spring constants to (1+k). We then return these parameters to their initial values and allow relaxation. The first observation is that $\langle \mathcal{R}_F \rangle$ can be negative when the timescale is long enough. When we check the individual components of $\langle \mathcal{R}_F \rangle$, $\langle \mathcal{R}_F^X \rangle$ and $\langle \mathcal{R}_F^Y \rangle$, we note that they are both the same shape as $\langle \mathcal{R}_F^X \rangle$ in Fig. 5.3.

 $^{^{8}}$ They will be negative since they are the 'work *done*' by a process that expects to gain rather



Figure 5.6: The work done over the course of a trajectory for the double dæmon, with several different time-scales for the evolution of the control parameter. Data for the solid lines was obtained by averaging values from 10^7 simulations. The blue, black and red lines show the evolution of $\langle \mathcal{R}_F \rangle$ (5.15) for $t_m = t_e = 2000, 3000$, and 6000 time steps, with relaxation periods of 4000 time steps afterwards. The purple dashed line shows the mutual information $\langle \Delta I \rangle$ since it is the lower bound of (5.18).

5.4.3 Second-law violating?

It is tempting to say that the double dæmon's behaviour constitutes a 'violation' of the second law since it appears that work is extracted from the system despite having counted cost of creating the information resource. Before we claim that this behaviour is a violation it is important to note that (5.18) has the same bound as (5.3). Indeed, both of these inequalities express the same idea as (3.35). That is, in the presence of feedback control, quantities like the entropy production or work extracted can be negative and this effect is bounded by the mutual information. Since this fact is not considered a 'violation' of the second law in the case of ordinary feedback control, there is no reason to suspect it is a 'violation' here either.

In the case of single oscillator exploitation, the model of coupled oscillators and treatment of the measurement process allows us to see that work extraction from feedback control in one system is compensated by a work cost in another

than expend work.



Figure 5.7: The values of $\langle \mathcal{R}_F^X \rangle$ and $\langle \mathcal{R}_F^X \rangle$ for $t_m = t_e = 3000$. Data obtained by averaging values from 10^7 simulations. Here, both oscillators benefit from exploitation and so the final value of both functionals is negative.

system. Here, the position variable of the y oscillator acts as the measurement variable used to select a control protocol for the x oscillator. The cost of creating the correlation between the variables balances the gains from extraction. We consider our description of the system is 'complete' in so far as we have modelled all the relevant components and processes.

For double exploitation, we essentially have a system that appears to be 'self-exploiting'. That is, the system uses its own state to select a feedback protocol. Fig. 5.8 shows what 'self feedback' is and how the double dæmon can be considered as a self feedback system. Seen this way, we are not looking at two systems selecting control protocols for each other but instead seeing a system select a control protocol for itself. For the double dæmon, the variable to be measured is (X, Y) and the control variable that is used to select a feedback protocol is also (X, Y). This is the same as picking a feedback protocol for a single system X based directly on X rather than some measurement variable Y. When we consider this kind of 'self-feedback' we essentially neglect the cost of correlating the control variable for the system with the system state. Earlier in Sec. 4.2 we mentioned the case of perfect measurements for the 1D



Figure 5.8: Control flow diagrams illustrating various feedback scenarios.

a) Ordinary feedback control. The 'system' state is correlated with a 'dæmon'. The dæmon then selects a feedback protocol which is then used to control the system. The system benefits from the feedback control but the dæmon does not.

b) Self feedback. The feedback protocol is generated directly by the state without including an intermediate 'dæmon'.

c) The double dæmon. If one considers both the system and demon as a single composite system, then the situation in c) appears identical to b).

information engine; the control variable is perfectly correlated with the system state, but we have not counted the cost of creating that correlation since we do not model the control variable.

To pick a feedback protocol, we must gain information about the state of the system through some process. To exploit just a single oscillator, we couple the X oscillator to the Y oscillator and then can use the state of Y to infer the state of X without disturbing the state of X. The Y system acts as the control variable for X. When feedback is enacted in both directions, the 'system' is (X, Y) and we do not explicitly have a control variable. To be able to exploit the state of the system (X, Y), we would have to correlate it with some additional third system and use the state of this third system to pick exploitation protocols for X and Y. Although the

trajectory functional \mathcal{R}_F counts the cost of creating a correlation between the X and Y variables, it does not count the cost of creating the correlation between (X, Y) and the variable used to select a control protocol. This essentially means that the double dæmon model is not 'complete'. To fully count the entropy production of the entire system we would need to include the entropy production from the operation of this additional third 'proxy' system that is used to select a feedback protocol for (X, Y).

5.5 Conclusion

In this chapter we have included dynamics for the measurement system of a feedback controller and demonstrated that for measurement-based feedback control, the work extracted through feedback control is at least balanced by the work cost of creating an information resource used to choose the feedback protocol. Rather than consider that we gain information about the state of the system through some unspecified process, we have modelled the process by which the a correlation is created between the measured and measurement systems. The correlation between the two systems is then exploited by one system and work is extracted, having paid a cost during the measurement phase that can at best be fully 'recouped' from exploitation, consistent with intuitions of the second law of thermodynamics. Overall, we note that this view accords with the argument of Ford [2016]; one cannot 'borrow' from the bank of entropy and pay back later, we must 'pay' first and benefit afterwards. That is, to be able to enact a protocol that decreases entropy on average, we must first produce entropy. Similarly here, to enact feedback control we must exploit information from an information resource, and to create that information resource costs work up front.

We have also considered what happens when feedback is enacted symmetrically between two systems and identified this as 'self-feedback' whereby a system uses its own state in order to select a control protocol. In this case the work that can be extracted is bounded by the mutual information, although we note that this does not allow us to 'violate' the second law since we have now not properly accounted for the feedback process. For double feedback, the bound on it's performance is given by (5.3). If one wanted to fully count the entropy of the total feedback process for double feedback, then one should include an additional 'measurement' system that is used to select a control protocol for the whole system.

5.5.1 Related frameworks

In this chapter we studied a system of coupled oscillators in order to model the process of measurement and feedback. We are able to consider the coupled oscillators as a 'system' and 'dæmon' and show that in order to be able to extract work in a cycle from an information engine, prior work must be done in creating the information resource that is then exploited in feedback. However, this is not the only way of modelling a Maxwell's dæmon system. Earlier in Sec. 4.1 we discussed two other frameworks that are used to study Maxwell's dæmon. Neither of these models necessarily employ explicit measurement or feedback processes. Nevertheless we discuss how our results are similar to the findings from these other frameworks.

Bipartite systems

The intuition that entropy reduction in one system is compensated for by entropy production in another is confirmed more directly by the study of bipartite systems Barato et al. [2013a,b]. These are processes modelled as two variables but only allow one variable can change in each transition. Previous work has also considered feedback and information processing in bipartite systems modelled by Langevin dynamics Horowitz and Sandberg [2014]; Hartich et al. [2016]. By considering each variable as representing the state of a subsystem, the process can be considered as two interacting systems. The interactions and evolution of the whole system is 'autonomous' and operates without control or time-varying parameters.

Since all the dynamics are explicitly modelled in these systems, the entropy production from each transition is easily accessible and is shown to obey a second-law like inequality. Hartich et al. [2014] uses this framework to make a direct analogy to Maxwell's dæmon, considering one variable to represent the 'system' and the other variable to represent the 'dæmon'. The study confirms that the work extracted from one system is compensated for by heat dissipation in the other subsystem and places bounds on those quantities. Our model can be seen as a bipartite system since, while both variables change in a single transition, this is only because we model in discrete time. Our model shows a similar result in Sec 5.3 to previous work on bipartite systems; the dæmon can use feedback control to extract work from one system, but must dissipate heat through its own operation with the result that the entire system cannot benefit on average.

Information registers/tapes

Other 'autononous' daemons are realised in models involving a system interacting with a binary tape Mandal et al. [2013]; Barato and Seifert [2014b]. In fact many of these system can be modelled as bipartite systems and so there is some overlap with the above. These types of Maxwell daemons use sets of dynamics that write information into an empty memory register. The dynamics and interaction with the tape can be chosen such that thermodynamic work is done by the system whilst the information entropy of the register increases. In this framework, one could think of information and entropy being equivalent, and that the generation of information in the memory register is what balances the entropy reduction in the system. Alternatively, one can consider the Landauer perspective; that the information must be 'deleted' at a later time which will cost work at least equal to the work that could be extracted by the system. Our model is slightly different in that that the work costs must be 'paid up front'. That is, in our model the daemon must pay the cost of measurement before work can be extracted from the system.

5.5.2 Experimental considerations

Since the model is based on a Brownian particle connected to a tether via a Hookean spring then constructing exactly that system may be the most appropriate way to test the framework (i.e verify the bounds in (5.3), (5.4) and (5.5)). The difficulty with this system is that the protocol requires the alteration of the spring constant which is not necessarily easy to do. The experiment should ideally take place at constant temperature so as not to alter the thermal fluctuations that the protocol is trying to exploit. However, the only way of changing the spring constant appears to be changing the temperature of the spring Manninen and Säynäjäkangas [2012] which will be difficult to do without affecting the temperature of the medium. Given these considerations, it is likely that the experimental apparatus has to be somewhat abstracted from the model description.

Many groups have previously used colloidal particles in optical-tweezers to test the concepts of stochastic thermodynamics Wang et al. [2002]; Carberry et al. [2004]; Trepagnier et al. [2004]; Speck et al. [2007]; Blickle et al. [2006]; Gomez-Solano et al. [2010]. By using a laser (specifically, the momentum exchange between the laser beam photons and the particle), it is possible to confine a particle to a specific region Ashkin et al. [1986]. Simple laser-traps can be used to establish the harmonic potentials that are similar to the potential energy of the Hookean spring. To couple particles together using laser-trapping might require a slightly more sophisticated set-up. Ideally the particles should be coupled to one another, but establishing a *controllable* coupling force between the two particles may be difficult.

Finally, it is also possible to construct electronic system as per Averin et al. [2011]; Koski et al. [2014a,b, 2015] that are capable of using feedback control. Such systems use feedback control to manipulate and cool single electrons in electronic circuits. Since these kinds of systems have already been used to verify the previous information thermodynamic results such as the Sagawa-Ueda FR, it may also be possible to use them to study the measurement and feedback system in more detail.

"I think that the task of philosophy is not to provide answers, but to show how the way we perceive a problem can be itself part of a problem."

Slavoj Žižek

Outlook and Discussion

In this final chapter, we will provide an overview and summary of the research presented in the previous chapters. Sec. 6.1 summarises our findings and points towards other contemporary research and directions for future study. Finally, Sec. 6.2 provides some concluding remarks on the status of information thermodynamics and the prognosis for the second law and Maxwell's dæmon.

As stated in Ch. 1 and evidenced by the large body of literature Bekenstein [1972]; Bennett [1987]; Evans et al. [1993]; Evans and Searles [1994]; Sheehan and Means [1998]; Ritort [2004]; Cavina et al. [2016]; challenging the second law is something of an enduring fascination for Physicists. It is worth noting (as in Sec. 2.7) that the second law of stochastic thermodynamics does *not* forbid decreases in any kind of entropy or negative values of trajectory functionals. Rather, the second law inequalities are statements of expectation. They refer to the average behaviour of the systems under consideration, and hence individual 'violations' of the second law are always possible¹.

¹No matter how large the system or how long the time period, in an infinite ensemble the probability of seeing at least one trajectory that reduces entropy is unity.

6.1 Summary and outlook

In this thesis we have described how feedback can be integrated into the framework of stochastic thermodynamics and using discrete-time and -space Markov chains as a basis we have shown how some key results can be obtained. Our contribution to this framework has been an analysis of the notion of 'time reversal' from the perspective of causal diagrams. Particularly, in Ch. 3 we explored the alternatives for 'time reversed' conjugate processes and found that only open-loop control as the conjugate process results in a meaningful fluctuation relation. Specifically, we provided a conceptual justification for use of open-loop feedback in the derivation of the Sagawa-Ueda fluctuation relation.

In Ch. 4 we presented a concrete model of an information engine which is analytically solvable for small systems. Since impact of information flow in these systems is well studied Ito [2016], we provide an analysis of the fluctuations of that information flow. Particularly, we obtained an exact expression of the large deviation rate function for information flow in a two site system. For general systems with more than two sites, we provided an approximate analysis of the fluctuations of information by constructing a Markov model in order to be able to infer a large deviation rate function from transitions on a coarse-grained state space. Recent work has explored the fluctuations of information flow for systems with continuous information flow Rosinberg and Horowitz [2016] as opposed to our discrete model. It would be of potential interest to explore the differences between information fluctuations in discrete and continuous systems. Our research may also be connected to work on using large deviation theory to optimise feedback protocols themselves Gagliardi et al. [2016].

Finally, we included the measurement process explicitly in Ch. 5. By modelling the system and dæmon as explicitly and considering the measurement process as a process that creates a correlation between their states, we were able to see that entropy reduction in one system is compensated by commensurate increases in other system. We confirmed previously found bounds on the work that can be extracted through feedback control as well as the work that must be paid to beforehand in measurement. Our work is part of broader interest in full models of Maxwell's dæmon Barato et al. [2013a,b]; Mandal et al. [2013]; Barato and Seifert [2014b]; Hartich et al. [2014], other contemporary studies that have found limits on the outcomes of feedback control Barato et al. [2014]; Machta [2015] and clarifications of existing bounds Wächtler et al. [2016]; Boyd et al. [2016]. Models of Maxwell's dæmon attempt to extract work on average from the controlled system. We note that it is possible to design protocols aimed at maximising the possibility for positive fluctuations in the extraction of work Cavina et al. [2016]. There is also work on developing optimal protocols for memory erasure Gavrilov and Bechhoefer [2016], since many views of Maxwell's dæmon view information erasure as being the key component to the operation of the dæmon.

In the next section, we will provide a brief history of the second law and place our work within the narrative of ongoing effort to fully explore and understand its meaning and ramifications.

6.2 The second law II

The culmination of nearly 150 years of discussions about Maxwell's Dæmon has been the framework of information thermodynamics Sagawa and Ueda [2010], which illuminates the profound importance of information in thermodynamics Seifert [2012]; Abreu and Seifert [2012]. By including the effect of system memory Zhou and Segal [2010] and/or information processing Maruyama et al. [2009] the 2nd law can be reformulated to include information entropy and used to study the operation of finite-time information processing systems systems Berry et al. [1999] such as information heat engines Jayannavar [1996] and refrigerators Mandal et al. [2013].

6.2.1 Exorcising the dæmon

As was alluded to in Sec. 1.3, the second law has a long and storied history of being challenged and reformulated in response to those challenges. The challenge posed by Maxwell's 'neat-fingered being' has been so powerful, that devices, systems and processes that provide apparent 'violaions' of the second law are all granted the title of 'dæmons', taking on the role of spiritual successors to that first challenge. The common spelling 'demon' perhaps belies a fact about the dæmon, that it has 'bedeviled' physics by continually challenging a law thought to be fundamental and forcing reinterpretation of that law. As each dæmon has challenged the law, new ideas and frameworks have been developed in order to 'save' the second law and 'exorcise' the dæmon.

In its original form, the dæmon itself did not *violate* the second law, rather, the dæmon simply observed and exploited the rare and unexpected events that occur on microscopic scales. The dæmon's ability was accepted as a consequence of the kind of judgements and actions that humans themselves are capable of: observing a situation and responding accordingly to bring about desired and possibly otherwise unlikely outcomes. The dæmon constitutes less of a violation of the second law, than a clarification of it. That is, in the presence of some special circumstances such as interference from a 'dæmon', the second law does not hold. However, in the pursuit of precision, the information thermodynamical framework has been able to augment this statement. Additional to pointing out that special circumstances can violate the second law, we have been able to place bounds on the degree to which the second law can be broken. Rather than being a statement about the average behaviour of the entropy production $\langle \Delta S_t \rangle$, the second law becomes a statement about a new quantity with a similar inequality, $\langle \Delta S_t \rangle + \langle \Delta I_t \rangle \geq 0$. That is, the second law describes the behaviour of a more general quantity with terms that refer to the system and its interaction with the dæmon. The second law is thus generalised and the action of the dæmon understood to only violate an incomplete version of the law.

However, this can be augmented still. By modelling the dæmon itself and including the process by which the dæmon's state becomes correlated with the system, we can show that the entropy reduction in the exploited system is compensated by the operating cost of the dæmon. The dæmon is thus fully understood not to constitute a 'violation' of the second law. When we consider the entropy of the entire scenario – the system *and* the dæmon – we see that this quantity is greater than or equal to zero on average. Thus feedback control and dæmons cannot provide us with sources of unlimited work extraction of indefinite entropy reduction. However, we are fully aware of the limitations on what can be achieved by exploiting systems using feedback or dæmonic control, or by creating systems with autonomous dæmonic dynamics. While we cannot decrease the total entropy on average, we can be clever about where it is produced and reduced.

6.2.2 The nature of information

Much of the framework of feedback and control theory involves talk about information and information processing. Several thinkers have even made direct equivalences between information and thermodynamic entropy. 'Information' is an incredibly loaded term. Loaded not just with different methodological interpretations, but also philosophical baggage. On the methodological side, whilst Shannon's 'information entropy' and the Boltzmann entropy have similar mathematical definitions, they do not share the same foundations. Thermodynamic entropy is a property of probability distributions, whereas Shannon information was originally developed as a property of a communication system. Thermodynamic entropy is fundamentally a relationship between temperature and heat which have no clear analogue in communications systems. Divorced of concepts like energy, heat and temperature, many have argued that the connection between the concepts is theoretical rather than actual Rapoport [1976]; Morowitz [1986]; Müller [2007]; Ben-Naim [2008]. That said, stochastic thermodynamics can already bridge the gap between abstract mathematical systems and real physical systems and so this criticism is perhaps not so crippling.

The much harder problem to address is the interpretation of what information itself actually *is.* In Sec. 3.2.3 we described an interpretation of the change in uncertainty as an amount of information 'gained' by an observer. However, we were referring to a quantity calculated as $\ln p$, where p is some probability. Given that $p \ge 0$, it is clear that $\ln p \le 0$. Rather than describing an amount of information 'gained' by an observer, it is perhaps more appropriate to think of it as an amount of uncertainty that is 'removed' by the act of observation². Uncertainty is more naturally thought of as a property of an agent's belief; it has an epistemic character. On the other hand, Landauer's "information is physical", seems to suggest that not only does the existence of information have an impact on the physics of a given system, but can also be read literally and taken to mean that information is a physical property of a process or system and hence has an ontological character.

6.2.3 Information as perspective

Whether information/entropy represents a property of the universe or the uncertainty of the agents within it is something of a metaphysical issue, but it has an impact on the way in which we think about the operation of information thermodynamic devices. If entropy is only a quantification of an agent's uncertainty, then we must adduce reasons as to why one's uncertainty must increase. Indeed, the entire endeavour of science can be seen as a somewhat successful reduction in human uncertainty.

As an example; in performing a simulation, it is possible to keep track of a complete description of the system state and hence never see any increase in entropy/uncertainty. It is only in discarding that detailed information that uncertainty on the future state of the system appears³. Increases in entropy could thus be seen as a consequence of not modelling a system comprehensively enough. In essence, the second law could be related to the central anxiety of the statistical/probabilistic

 $^{^{2}}$ This view isn't particularly flattering to the observers; they are never really in possession of information, they are in possession of uncertainty.

 $^{^{3}}$ Compare this with another famous physical 'dæmon'; Laplace's dæmon, who possesses knowledge of the position and momentum of every particle in the universe at a single instant. In a deterministic universe, the entire past and future would be known to such a dæmon and hence can never be more or less certain.

methodology; 'in discarding a full microscopic description of a system, what might have been lost that was worth keeping?'. That is to say, if increasing uncertainty is a consequence of the act of approximation, then perhaps entropy is something humanity simply has to live with, since "truth ... is much too complicated to allow anything but approximations" von Neumann [1947].

Bibliography

- David Abreu and Udo Seifert. Extracting work from a single heat bath through feedback. *Europhysics Letters*, 94(1):10001, 2011.
- David Abreu and Udo Seifert. Thermodynamics of genuine nonequilibrium states under feedback control. *Physical Review Letters*, 108(3):030601, 2012.
- Arthur Ashkin, JM Dziedzic, JE Bjorkholm, and Steven Chu. Observation of a single-beam gradient force optical trap for dielectric particles. *Optics letters*, 11 (5):288–290, 1986.
- Karl Johan Aström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton University Press, 2010.
- Dmitri V Averin, Mikko Möttönen, and Jukka P Pekola. Maxwell's demon based on a single-electron pump. *Phys. Rev. B*, 84(24):245448, 2011.
- AC Barato and U Seifert. Unifying three perspectives on information processing in stochastic thermodynamics. *Physical Review Letters*, 112(9):090601, 2014a.
- AC Barato, D Hartich, and U Seifert. Information-theoretic versus thermodynamic entropy production in autonomous sensory networks. *Physical Review E*, 87(4): 042104, 2013a.
- Andre C Barato and Udo Seifert. Stochastic thermodynamics with information reservoirs. *Phys. Rev. E*, 90(4):042150, 2014b.
- Andre C Barato, David Hartich, and Udo Seifert. Rate of mutual information between coarse-grained non-markovian variables. *Journal of Statistical Physics*, 153(3):460–478, 2013b.
- Andre C Barato, David Hartich, and Udo Seifert. Efficiency of cellular information processing. New Journal of Physics, 16(10):103024, 2014.

- Andre Cardoso Barato and Udo Seifert. An autonomous and reversible maxwell's demon. *Europhysics Letters*, 101(6):60001, 2013.
- Jacob D Bekenstein. Black holes and the second law. Lettere Al Nuovo Cimento (1971–1985), 4(15):737–740, 1972.
- Arieh Ben-Naim. A Farewell to Entropy: Statistical Thermodynamics Based on Information. World Scientific, 2008.
- Charles H Bennett. Logical reversibility of computation. Maxwell's Demon. Entropy, Information, Computing, pages 197–204, 1973.
- Charles H Bennett. The thermodynamics of computation a review. Int. J. Theor. Phys., 21(12):905–940, 1982.
- Charles H Bennett. Demons, engines and the second law. *Scientific American*, 257 (5):108–116, 1987.
- R Stephen Berry, Vladimir A Kazakov, Stanislaw Sieniutycz, Z Szwast, and Anatoly M Tsirlin. Thermodynamic Optimization of Finite Time Processes. Wiley, 1999.
- Antoine Bérut, 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, 2012.
- Valentin Blickle, Thomas Speck, Laurent Helden, Udo Seifert, and Clemens Bechinger. Thermodynamics of a colloidal particle in a time-dependent nonharmonic potential. *Phys. Rev. Lett.*, 96(7):070603, 2006.
- GN Bochkov and Yu E Kuzovlev. General theory of thermal fluctuations in nonlinear systems. *Zh. Eksp. Teor. Fiz*, 72:238–243, 1977.
- GN Bochkov and Yu E Kuzovlev. Nonlinear fluctuation-dissipation relations and stochastic models in nonequilibrium thermodynamics: I. generalized fluctuationdissipation theorem. *Physica A: Statistical Mechanics and its Applications*, 106 (3):443–479, 1981.
- Alexander B Boyd, Dibyendu Mandal, and James P Crutchfield. Correlationpowered information engines and the thermodynamics of self-correction. arXiv preprint arXiv:1606.08506, 2016.
- Leon Brillouin. Maxwell's demon cannot operate: Information and entropy. i. J.
App. Phys., 22(3):334-337, 1951.

- Patrice A Camati, John PS Peterson, Tiago B Batalhao, Kaonan Micadei, Alexandre M Souza, Roberto S Sarthour, Ivan S Oliveira, and Roberto M Serra. Experimental rectification of entropy production by a maxwell's demon in a quantum system. *Scientific Reports*, 2016.
- DM Carberry, James Cowie Reid, GM Wang, Edith M Sevick, Debra J Searles, and Denis J Evans. Fluctuations and irreversibility: an experimental demonstration of a second-law-like theorem using a colloidal particle held in an optical trap. *Phys. Rev. Let.*, 92(14):140601, 2004.
- Sadi Carnot, Hippolyte Carnot, and William Thomson Lord Kelvin. Reflections on the motive power of heat and on machines fitted to develop that power. J. Wiley, 1890.
- Sean Carroll. From eternity to here: the quest for the ultimate theory of time. Penguin, 2010.
- Vasco Cavina, Andrea Mari, and Vittorio Giovannetti. Optimal processes for probabilistic work extraction beyond the second law. arXiv preprint arXiv:1604.08094, 2016.
- Rudolf Clausius. Über die bewegende kraft der wärme und die gesetze, welche sich daraus für die wärmelehre selbst ableiten lassen. Annalen der Physik, 155(3): 368–397, 1850.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Harald Cramér. Über eine eigenschaft der normalen verteilungsfunktion. Mathematische Zeitschrift, 41(1):405–414, 1936.
- Harald Cramér. Sur un nouveau théoreme-limite de la théorie des probabilités. Actualités scientifiques et industrielles, 736(5-23):115, 1938.
- Gavin E Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721, 1999.
- Gavin E Crooks. Path-ensemble averages in systems driven far from equilibrium. *Physical review E*, 61(3):2361, 2000.
- Sebastian Deffner. Information-driven current in a quantum maxwell demon. Phys-

ical Review E, 88(6):062128, 2013.

- Amir Dembo and Ofer Zeitouni. Large deviations techniques and applications, volume 38. Springer Science & Business Media, 2009.
- Arthur Eddington. The Nature of the Physical World: Gifford Lectures (1927). Cambridge University Press, 2012.
- Paul Ehrenfest and Tatiana Ehrenfest. The conceptual foundations of the statistical approach in mechanics. Courier Corporation, 2002.
- Richard S Ellis. Large deviations for a general class of random vectors. *The Annals of Probability*, pages 1–12, 1984.
- Richard S Ellis. An overview of the theory of large deviations and applications to statistical mechanics. *Scandinavian Actuarial Journal*, 1995(1):97–142, 1995.
- Denis J Evans and Debra J Searles. Equilibrium microstates which generate second law violating steady states. *Physical Review E*, 50(2):1645, 1994.
- Denis J Evans and Debra J Searles. The fluctuation theorem. Advances in Physics, 51(7):1529–1585, 2002.
- Denis J Evans, Ezechiel Godert David Cohen, and Gary P Morriss. Probability of second law violations in shearing steady states. *Physical Review Letters*, 71(15): 2401, 1993.
- Ian J Ford. Maxwell's demon and the management of ignorance in stochastic thermodynamics. *Contemporary Physics*, pages 1–22, 2016.
- Ian J Ford and Michael Maitland. Entropy-reducing dynamics of a double demon. arXiv preprint arXiv:1605.02202, 2016.
- Anthony Thomas Fuller. The early development of control theory. Journal of Dynamic Systems, Measurement, and Control, 98(2):109–118, 1976.
- Alessio Gagliardi, Alessandro Pecchia, and Aldo Di Carlo. Large deviation theory to model systems under an external feedback. arXiv preprint arXiv:1603.03786, 2016.
- Giovanni Gallavotti and Ezechiel Godert David Cohen. Dynamical ensembles in nonequilibrium statistical mechanics. *Physical Review Letters*, 74(14):2694, 1995a.

- Giovanni Gallavotti and Ezechiel Godert David Cohen. Dynamical ensembles in stationary states. *Journal of Statistical Physics*, 80(5-6):931–970, 1995b.
- Crispin W Gardiner et al. *Handbook of stochastic methods*, volume 3. Springer Berlin, 1985.
- Jürgen Gärtner. On large deviations from the invariant measure. Theory of Probability & Its Applications, 22(1):24–39, 1977.
- Momčilo Gavrilov and John Bechhoefer. Erasure without work in an asymmetric, double-well potential. arXiv preprint arXiv:1607.02458, 2016.
- Cristian Giardina, Jorge Kurchan, Vivien Lecomte, and Julien Tailleur. Simulating rare events in dynamical processes. *Journal of statistical physics*, 145(4):787–811, 2011.
- Roy J Glauber. Time-dependent statistics of the ising model. Journal of mathematical physics, 4(2):294–307, 1963.
- James Gleick. The Information: A History, a Theory, a Flood. New York: Pantheon Books, 2011.
- Juan Ruben Gomez-Solano, Ludovic Bellon, Artyom Petrosyan, and Sergio Ciliberto. Steady-state fluctuation relations for systems driven by an external random force. *Euro. Phys. Lett.*, 89(6):60003, 2010.
- Léo Granger and Holger Kantz. Thermodynamic cost of measurements. *Physical Review E*, 84(6):061110, 2011.
- Jonathan J Halliwell, Juan Pérez-Mercader, and Wojciech Hubert Zurek. *Physical origins of time asymmetry*. Cambridge University Press, 1996.
- Rosemary J Harris and Gunther M Schütz. Fluctuation theorems for stochastic dynamics. Journal of Statistical Mechanics: Theory and Experiment, 2007(07): P07020, 2007.
- David Hartich, Andre C Barato, and Udo Seifert. Stochastic thermodynamics of bipartite systems: transfer entropy inequalities and a maxwell's demon interpretation. J. Stat. Mech: Theory and Experiment, 2014(2):P02016, 2014.
- David Hartich, Andre C Barato, and Udo Seifert. Sensory capacity: An information theoretical measure of the performance of a sensor. *Physical Review E*, 93(2): 022116, 2016.

- Takahiro Hatano and Shin-ichi Sasa. Steady-state thermodynamics of langevin systems. *Physical review letters*, 86(16):3463, 2001.
- Jordan Horowitz and Christopher Jarzynski. Comparison of work fluctuation relations. Journal of Statistical Mechanics: Theory and Experiment, 2007(11):P11002, 2007.
- Jordan M Horowitz and Massimiliano Esposito. Thermodynamics with continuous information flow. *Physical Review X*, 4(3):031015, 2014.
- Jordan M Horowitz and Juan MR Parrondo. Thermodynamic reversibility in feedback processes. *Europhysics Letters*, 95(1):10005, 2011.
- Jordan M Horowitz and Henrik Sandberg. Second-law-like inequalities with information and their interpretations. New J. Phys., 16(12):125007, 2014.
- Jordan M Horowitz and Suriyanarayanan Vaikuntanathan. Nonequilibrium detailed fluctuation theorem for repeated discrete feedback. *Physical Review E*, 82(6): 061120, 2010.
- Jordan M Horowitz, Takahiro Sagawa, and Juan MR Parrondo. Imitating chemical motors with optimal information motors. *Physical Review Letters*, 111(1):010602, 2013.
- Adolf Hurwitz. On the conditions under which an equation has only roots with negative real parts. *Selected papers on mathematical trends in control theory*, 65: 273–284, 1964.
- Alberto Isidori. Nonlinear control systems. Springer Science & Business Media, 1995.
- Sosuke Ito. Information Thermodynamics on Causal Networks and its Application to Biochemical Signal Transduction. Springer, 2016.
- Sosuke Ito and Takahiro Sagawa. Information thermodynamics on causal networks. *Physical Review Letters*, 111(18):180603, 2013.
- Christopher Jarzynski. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Physical Review E*, 56(5):5018, 1997a.
- Christopher Jarzynski. Nonequilibrium equality for free energy differences. *Physical Review Letters*, 78(14):2690, 1997b.

- Christopher Jarzynski. Hamiltonian derivation of a detailed fluctuation theorem. Journal of Statistical Physics, 98(1-2):77–102, 2000.
- Christopher Jarzynski. Comparison of far-from-equilibrium work relations. *Comptes Rendus Physique*, 8(5):495–506, 2007.
- Arun M Jayannavar. Simple model for maxwell's-demon-type information engine. Physical Review E, 53(3):2957, 1996.
- Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106 (4):620, 1957.
- Rudolf Kalman. On the general theory of control systems. *IRE Transactions on* Automatic Control, 4(3):110–110, 1959.
- William Thomson Lord Kelvin. On the age of the sun's heat. Macmillan's Magazine, 5:388–393, 1862a.
- William Thomson Lord Kelvin. Physical considerations regarding the possible age of the sun's heat. 1862b.
- William Thomson Lord Kelvin. The sorting demon of maxwell. Proc. R. Soc. London, 9(1):13–1, 1879.
- William Thomson Lord Kelvin, James Prescott Joule, and Victor Regnault. On the Dynamical Theory of Heat: With Numerical Results Deduced from Mr. Joule's Equivalent of a Thermal Unit and M. Regnault's Observations on Steam. Society, 1851.
- Ross Kindermann, James Laurie Snell, et al. *Markov random fields and their applications*, volume 1. American Mathematical Society Providence, RI, 1980.
- Jonne V Koski, Ville F Maisi, Jukka P Pekola, and Dmitri V Averin. Experimental realization of a szilard engine with a single electron. *Proceedings of the National Academy of Sciences*, 111(38):13786–13789, 2014a.
- Jonne V Koski, Ville F Maisi, Takahiro Sagawa, and Jukka P Pekola. Experimental observation of the role of mutual information in the nonequilibrium dynamics of a maxwell demon. *Physical Review Letters*, 113(3):030601, 2014b.
- Jonne V Koski, Aki Kutvonen, Ivan M Khaymovich, Tapio Ala-Nissila, and Jukka P Pekola. On-chip maxwell's demon as an information-powered refrigerator. *Phys. Rev. Lett.*, 115(26):260602, 2015.

Benjamin C Kuo. Automatic control systems. Prentice Hall PTR, 1981.

- Jorge Kurchan. Fluctuation theorem for stochastic dynamics. Journal of Physics A: Mathematical and General, 31(16):3719, 1998.
- Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.
- Rolf Landauer. The physical nature of information. *Physics letters A*, 217(4):188–193, 1996.
- Joel L Lebowitz. Macroscopic laws, microscopic dynamics, time's arrow and boltzmann's entropy. Physica A: Statistical Mechanics and its Applications, 194(1-4): 1–27, 1993.
- Joel L Lebowitz and Herbert Spohn. A gallavotti-cohen-type symmetry in the large deviation functional for stochastic dynamics. J. Stat. Phys., 95(1-2):333-365, 1999.
- Vivien Lecomte and Julien Tailleur. A numerical approach to large deviations in continuous time. Journal of Statistical Mechanics: Theory and Experiment, 2007 (03):P03004, 2007.
- Harvey Leff and Andrew F Rex. Maxwell's Demon 2 Entropy, Classical and Quantum Information, Computing. CRC Press, 2010.
- Gilbert N Lewis. A new principle of equilibrium. Proceedings of the National Academy of Sciences of the United States of America, 11(3):179, 1925.
- Seth Lloyd. Quantum-mechanical maxwell's demon. *Physical Review A*, 56(5):3374, 1997.
- Benjamin J Lopez, Nathan J Kuwada, Erin M Craig, Brian R Long, and Heiner Linke. Realization of a feedback controlled flashing ratchet. *Physical Review Letters*, 101(22):220601, 2008.
- Benjamin B Machta. Dissipation bound for thermodynamic control. Physical review letters, 115(26):260603, 2015.
- Michael C Mackey. *Time's arrow: the origins of thermodynamic behavior*. Courier Corporation, 2003.
- Christian Maes and Karel Netočný. Time-reversal and entropy. J. Stat. Phys., 110 (1-2):269–310, 2003.

- Michael Maitland, Stefan Grosskinsky, and Rosemary J Harris. Large deviation analysis of a simple information engine. *Phys. Rev. E*, 92(5):052136, 2015.
- 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.
- Dibyendu Mandal, Haitao Quan, and Christopher Jarzynski. Maxwell's refrigerator: An exactly solvable model. *Physical review letters*, 111(3):030602, 2013.
- Timo Manninen and Jukka Säynäjäkangas. Mechanical properties of ferritic stainless steels at elevated temperature. In *Proceedings of the Fourth International Experts Seminar on Stainless Steel in Structures*, 2012.
- Koji Maruyama, Franco Nori, and Vlatko Vedral. Colloquium: The physics of maxwell's demon and information. *Reviews of Modern Physics*, 81(1):1, 2009.
- Daniel C Mattis and M Lawrence Glasser. The uses of quantum field theory in diffusion-limited reactions. *Reviews of Modern Physics*, 70(3):979, 1998.
- James Clerk Maxwell. On the dynamical theory of gases. Philosophical transactions of the Royal Society of London, 157:49–88, 1867a.
- James Clerk Maxwell. On governors. Proceedings of the Royal Society of London, 16:270–283, 1867b.
- Neri Merhav and Yariv Kafri. Statistical properties of entropy production derived from fluctuation theorems. Journal of Statistical Mechanics: Theory and Experiment, 2010(12):P12022, 2010.
- Leonor Michaelis and Maud L Menten. Die kinetik der invertinwirkung. *Biochem.* z, 49(333-369):352, 1913.
- Harold Morowitz. Entropy and nonsense. *Biology and Philosophy*, 1(4):473–476, 1986.
- Ingo Müller. A history of thermodynamics: the doctrine of energy and entropy. Springer Science & Business Media, 2007.
- Richard E Neapolitan. Probabilistic reasoning in expert systems: theory and algorithms. Wiley, NY, 1990.
- Katsuhiko Ogata and Yanjuan Yang. Modern control engineering. 1970.

- Yoshitsugu Oono and Marco Paniconi. Steady state thermodynamics. Progress of Theoretical Physics Supplement, 130:29–44, 1998.
- Jung Jun Park, Kang-Hwan Kim, Takahiro Sagawa, and Sang Wook Kim. Heat engine driven by purely quantum information. *Physical Review Letters*, 111(23): 230402, 2013.
- Juan MR Parrondo, Jordan M Horowitz, and Takahiro Sagawa. Thermodynamics of information. *Nature Physics*, 11(2):131–139, 2015.
- Judea Pearl. Bayesian networks: A model of self-activated memory for evidential reasoning. University of California (Los Angeles). Computer Science Department, 1985.
- Barbara Piechocinska. Information erasure. Physical Review A, 61(6):062314, 2000.
- Max Planck. Vorlesungen über die Theorie der Wärmestrahlung. Leipzig, 1921.
- M Ponmurugan. Generalized detailed fluctuation theorem under nonequilibrium feedback control. *Physical Review E*, 82(3):031129, 2010.
- Ilya Prigogine and Dilip Kondepudi. Modern thermodynamics: From heat engines to dissipative structures. West Sussex, England: John Wiley & Sons Ltd, 1998.
- Anatol Rapoport. General systems theory: A bridge between two cultures. third annual ludwig von bertalanffy memorial lecture. *Behavioral science*, 21(4):228– 239, 1976.
- Félix Ritort. Work fluctuations, transient violations of the second law and freeenergy recovery methods: Perspectives in theory and experiments. In *Poincaré Seminar 2003*, pages 193–226. Springer, 2004.
- Ralph Tyrell Rockafellar. Convex analysis. Princeton university press, 2015.
- Christian M Rohwer, Florian Angeletti, and Hugo Touchette. Convergence of free energy and large deviation estimators. arXiv preprint arXiv:1409.8531, 2014.
- Martin Luc Rosinberg and Jordan M Horowitz. Continuous information flow fluctuations. arXiv preprint arXiv:1607.05438, 2016.
- Edward John Routh. A treatise on the stability of a given state of motion: particularly steady motion. Macmillan and Company, 1877.

- Stuart Russell, Peter Norvig, and Artificial Intelligence. A modern approach. Artificial Intelligence. Prentice-Hall, Egnlewood Cliffs, 25:27, 1995.
- Takahiro Sagawa and Masahito Ueda. Information thermodynamics: Maxwell's demon in nonequilibrium dynamics. In Heinz Georg Schuster, Rainer Klages, Wolfram Just, and Christopher Jarzynski, editors, Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond. John Wiley & Sons.
- Takahiro Sagawa and Masahito Ueda. Second law of thermodynamics with discrete quantum feedback control. *Physical Review Letters*, 100(8):080403, 2008.
- Takahiro Sagawa and Masahito Ueda. Minimal energy cost for thermodynamic information processing: measurement and information erasure. *Physical review letters*, 102(25):250602, 2009.
- Takahiro Sagawa and Masahito Ueda. Generalized jarzynski equality under nonequilibrium feedback control. *Physical Review Letters*, 104(9):090602, 2010.
- Takahiro Sagawa and Masahito Ueda. Nonequilibrium thermodynamics of feedback control. *Physical Review E*, 85(2):021104, 2012.
- Shin-ichi Sasa and Hal Tasaki. Steady state thermodynamics. *Journal of statistical physics*, 125(1):125–224, 2006.
- David A Schum. The evidential foundations of probabilistic reasoning. Northwestern University Press, 2001.
- Gunther M. Schütz. Phase transitions and critical phenomena, vol 19, c. Domb and J. Lebowitz eds.(Academic, London, 2001), 2001.
- Debra J Searles and Denis J Evans. Fluctuation theorem for stochastic systems. *Physical Review E*, 60(1):159, 1999.
- Udo Seifert. Entropy production along a stochastic trajectory and an integral fluctuation theorem. *Physical Review Letters*, 95(4):040602, 2005.
- Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012.
- Ken Sekimoto. Kinetic characterization of heat bath and the energetics of thermal ratchet models. *Journal of the physical society of Japan*, 66(5):1234–1237, 1997.
- Ken Sekimoto. Langevin equation and thermodynamics. Progress of Theoretical Physics Supplement, 130:17–27, 1998.

- Claude E Shannon. A mathematical theory of communication. Bell System Technical Journal, 27:623, 1948.
- DP Sheehan and JD Means. Minimum requirement for second law violation: A paradox revisited. *Physics of Plasmas (1994-present)*, 5(6):2469–2471, 1998.
- Naoto Shiraishi, Takumi Matsumoto, and Takahiro Sagawa. Measurement-feedback formalism meets information reservoirs. New Journal of Physics, 18(1):013044, 2016.
- T Speck, V Blickle, C Bechinger, and Udo Seifert. Distribution of entropy production for a colloidal particle in a nonequilibrium steady state. *Euro. Phys. Lett.*, 79(3): 30002, 2007.
- Thomas Speck and Udo Seifert. Integral fluctuation theorem for the housekeeping heat. Journal of Physics A: Mathematical and General, 38(34):L581, 2005.
- Richard Spinney and Ian Ford. Fluctuation relations: A pedagogical overview. In Heinz Georg Schuster, Rainer Klages, Wolfram Just, and Christopher Jarzynski, editors, Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond. John Wiley & Sons.
- Philipp Strasberg, Gernot Schaller, Tobias Brandes, and Massimiliano Esposito. Thermodynamics of a physical model implementing a maxwell demon. *Physical review letters*, 110(4):040601, 2013.
- Leo Szilard. Über die entropieverminderung in einem thermodynamischen system bei eingriffen intelligenter wesen. Z. Phys., 53(11-12):840–856, 1929.
- Richard Chace Tolman. The principles of statistical mechanics. Courier Corporation, 1938.
- Hugo Touchette. The large deviation approach to statistical mechanics. *Physics Reports*, 478(1):1–69, 2009.
- Hugo Touchette and Rosemary J Harris. Large deviation approach to nonequilibrium systems. In Heinz Georg Schuster, Rainer Klages, Wolfram Just, and Christopher Jarzynski, editors, Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond. John Wiley & Sons, 2013.
- Hugo Touchette and Seth Lloyd. Information-theoretic limits of control. *Physical Review Letters*, 84(6):1156, 2000.

- Shoichi Toyabe, Takahiro Sagawa, Masahito Ueda, Eiro Muneyuki, and Masaki Sano. Experimental demonstration of information-to-energy conversion and validation of the generalized jarzynski equality. *Nature Physics*, 6(12):988–992, 2010.
- EH Trepagnier, Christopher Jarzynski, Felix Ritort, Gavin E Crooks, CJ Bustamante, and J Liphardt. Experimental test of hatano and sasa's nonequilibrium steady-state equality. *Proc. Nat. Acad. Sci. U.S.A.*, 101(42):15038–15041, 2004.
- George Eugene Uhlenbeck, George Ford, and Elliot Waters Montroll. Lectures in statistical mechanics. 1963.
- Nicolaas Godfried Van Kampen. Stochastic processes in physics and chemistry, volume 1. Elsevier, 1992.
- SR Srinivasa Varadhan. Asymptotic probabilities and differential equations. Commun. Pure App. Math., 19(3):261–286, 1966.
- Gore Verbinski. *Pirates of the Caribbean: The Curse of the Black Pearl*. Walt Disney Pictures and Jerry Bruckheimer Films, 2003.
- John von Neumann. The mathematician. *The Works of the Mind*, 1(1):180–196, 1947.
- Christopher W Wächtler, Philipp Strasberg, and Tobias Brandes. Stochastic thermodynamics based on incomplete information: Generalized jarzynski equality with measurement errors with or without feedback. arXiv preprint arXiv:1608.01574, 2016.
- GM Wang, Edith M Sevick, Emil Mittag, Debra J Searles, and Denis J Evans. Experimental demonstration of violations of the second law of thermodynamics for small systems and short time scales. *Phys. Review Letters*, 89(5):050601, 2002.
- Norbert Wiener et al. Cybernetics. Hermann Paris, 1948.
- Yun Zhou and Dvira Segal. Minimal model of a heat engine: Information theory approach. *Physical Review E*, 82(1):011120, 2010.