Imperial College London

Mathematical modelling of long-term progression of atopic dermatitis

Harley Day ([redacted]**)**

Thesis submitted for the Degree of Doctor of Philosophy of Imperial College London

Department of Bioengineering

December 2023

Copyright statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Statement of originality

I hereby declare that this thesis and the work reported herein is the product of my own research, except where explicitly stated, and that I have acknowledged and referenced the work of others.

Harley Day (2024)

This thesis is dedicated to my wife, parents, and sister. I love you all.

Gubing,

Liz and Russ,

Clarissa,

and my bicycles.

Thanks for getting me here.

Acknowledgements

My completion of this doctoral thesis has been made possible by the love and support of family, friends, and colleagues, alongside the invaluable care received from the NHS and Imperial College's Student Counselling Service. Their positive impact on my life reaches beyond the academic and shall extend beyond my time at Imperial College. Their compassion has facilitated a transformative improvement to my mental health, turning a challenging period of my life into a therapeutic and gratifying experience that I look back on with a sense of accomplishment.

My heartfelt thanks and deep admiration go to my academic supervisor, Dr Reiko Tanaka. Your guidance, expertise and intellectual generosity have been instrumental in shaping my research and personal growth. Your sensitivity to the state of mind of your students is a credit to you, and in personally ensuring I received timely professional help, you initiated a life-changing process of recovery. I am forever grateful.

To the wonderful Dr Elisa Domínguez-Hüttinger, your patience and endless encouragement make you a joy to work with. Your friendly and enthusiastic style made it easy to share new or half-formed ideas. Watching you nurture and develop those ideas during our discussions gave me a newfound confidence that has had a profound impact on me. Our many fruitful discussions were the highlights of this PhD. It has been my privilege to work with you.

To Dr Guillem Hurault, Dr Tara Hameed, Prakrati Dangarh, Jamie Lee, Ariane Duverdier, Gemma Pitotti, Dr Yuxin Hao, Jennifer Carus, Dr Bryony Halcrow, Dr Ashwin Kumar, Dr Enrico Tosoratti, Dr Luke Heaton, Dr Alina Păduraru, Clare Pomery and others who I have been fortunate to befriend during my university journey, you are likely unaware of the enormous positive influence you have had on my life. Your friendship has been a treasured comfort and support, far greater in personal significance than any academic works. Particular thanks to Guillem, Tara, Prakrati, Jamie, and Ariane who shared their considerable expertise, and checked my working and writing.

I extend my deep gratitude to Dr Claudia Clopath, Dr Spyros Masouros, Dr Peter Weinberg, and Dr Darryl Overby for their invaluable contributions to my academic journey from undergraduate engineering student to a much calmer, wiser and happier PhD student. Each of you has generously shared your expertise in writing techniques and provided crucial support regarding mental health. Your guidance and care have equipped me with the necessary tools to progress in my PhD journey and navigate the process of personal recovery. I am truly thankful that you were there for me.

To the fantastic GPs and pharmacists who first took on my case and lifted me from the depths of depression, I remain acutely aware of the privilege it was to freely receive your care when I most needed it. The NHS facilitated my recovery as it does for so many in the UK. I count myself incredibly lucky to have benefited from its continued existence in a time when it comes under ever-increasing strain, a fact that became even more apparent when the COVID-19 pandemic struck mid-way through my PhD. I sincerely hope the NHS will be there for future generations, supported by proper public health measures and funding.

My heartfelt appreciation goes to Dr Sandra Roscoe, Bioengineering Department Wellbeing Advisor. Our conversations about mental health, writing and cycling (along with Rus Kesley) were a lifeline and a great help during an extremely cold TCR No.9. Your unwavering support and onward referrals helped me stay afloat in the sometimes-stormy sea of my mind.

The Imperial College Student Counselling Service provided the safe space for difficult discussions, teaching me the cognitive behavioural therapy (CBT) techniques that have allowed me to swim those

seas with confidence. The impact of our conversations and ongoing practice of CBT exercises continues to facilitate my wellbeing by calming the waters.

My beloved wife, parents, and sister, your love and support has brought a lasting calm to that stormy sea by restoring my sense of joy. With your help, I have been able to reach the shore and can now simply watch the waves standing alongside you all.

Gubing, whose intelligence, grace, and kindness I have admired since we first met in our undergraduate days, you pulled me from the water. Dr Gubing Wang, my love, I am overjoyed to be your husband. I look forward to our adventures together.

Thank you all from the bottom of my heart.

Abstract

The skin is the largest organ of the human body, and performs a vital defensive and physiological role, protecting the delicate tissues beneath from pathogens and physical damage while preventing water loss and regulating body temperature. Many of these protective functions are performed by the epidermis, a paper-thin layer of cells rich in structural proteins and lipids which produce antimicrobial peptides and plays host to cells of the immune system. Failure of this epidermal layer of the skin to maintain a competent barrier against pathogens can trigger the activation of allergic immune responses, causing the most common form of skin disease, atopic dermatitis (AD).

While mild-to-moderate AD can usually be well controlled by emollient and corticosteroid therapies, some AD patients go on to develop a more treatment-resistant form of chronic AD. The pathophysiological mechanisms of chronic AD are poorly understood, making development of new therapies and management of chronic AD difficult.

In this thesis, we develop a mechanistic mathematical model of chronic AD with the aim of elucidating how the complicated disease dynamics of chronic AD emerge as the result of dysregulated interactions between the barrier and immune system. We show how different genotypes can lead to different disease dynamics and propose how various therapies might affect these dynamics.

The model we develop is structured as a hybrid dynamical system, consisting of both continuous-valued and discrete-valued components which interact to produce counterintuitive dynamical properties. We develop a series of novel methods designed to elucidate the range of dynamics available in our model. Finally, we investigate how the methods we develop here can be applied to hybrid dynamical systems more broadly.

Contents

Acknowledgements 1
Abstract 3
Contents 4
List of Figures 8
List of tables 12
Chapter 1. Introduction 15
1.1. Background 15
1.2. A mathematical modelling approach to understand chronic atopic dermatitis (AD) 16
1.3. Aims and objectives 17
Chapter 2. Previous work on modelling atopic dermatitis 18
2.1. Previously published “double-switch” model of AD onset 18
2.1.1. Model equations 18
2.1.2. Parameter values 20
2.1.3. Qualitative phenotypes of AD onset 20
2.2. Reasons why the double-switch model does not reflect chronic AD 20
2.3. Understanding the four behaviours of the double-switch 22
2.3.1. Finding the focal point values of the double-switch model 22
2.3.2. Inferring switch transition dynamics from focal point values 25
2.3.3. Calculating BE bifurcations 28
2.4. Conclusion 31
Chapter 3. Modelling the effect of type-2 T-helper cells in chronic AD 32
3.1. Background 32
3.2. Modelling Th2 population dynamics 32
3.2.1. Model assumptions on timescale and bistability 33
3.2.2. The Yates model 33
3.2.3. Model 1: Systemic IL4/TGFβ feedback 35
3.2.4. Model 2: Systemic IL4 feedback + Th1/Th2 mutual inhibition 39
3.3. Effect of Dupilumab treatment on both models 46
3.3.1. Effect of dupilumab on critical values 47
3.3.2. Reversal of inflammation under Dupilumab treatment 48
3.4. Comparing Models 1 and 2 49
3.5. Implementing the reversible T-switch in the double-switch model 50
3.6. Conclusion 51
Chapter 4. Mathematical modelling of the regulatory mechanisms of keratinocyte differentiation for epidermal homeostasis 52
4.1. Model description 52

4.2. Model behaviour 55
4.2.1. Switch-like and history-dependent keratinocyte differentiation in response to change in the extracellular calcium level. 55
4.2.2. Modulation of keratinocyte differentiation by immune responses and inflammation 56
4.2.3. Summarising keratinocyte differentiation as a bistable switch 58
4.3. Conclusions 58
Chapter 5. Modelling chronic AD 60
5.1. Triple-switch model 60
5.1.1. RK-switch 61
5.1.2. G-switch 61
5.1.3. T-switch 62
5.1.4. Mechanisms included in the triple-switch model 62
5.1.5. Parameters and units 65
5.2. Bifurcation analysis 67
5.2.1. Focal points and stability of the triple switch model 71
5.2.2. Boundary-equilibrium bifurcations 73
5.2.3. Connectivity and presence of the behaviour regions for different parameter values . 79
5.2.4. Crosstalk between switches 86
5.2.5. Summary of our parameter sampling method 94
5.2.6. Parameter sensitivity of bifurcation diagrams 94
5.3. Conclusion 98
Chapter 6. Dynamical analysis of switch transitions 100
6.1. Vocabulary and definitions 101
6.1.1. Switch configurations 101
6.1.2. Switch transitions 102
6.1.3. Transition networks or trees 104
6.1.4. Families of trees 105
6.2. Constructing configuration-behaviour tables 105
6.3. Switch transitions 110
6.4. Limitations of focal point analysis 112
6.4.1. Initial condition dependence of switch behaviours 112
6.4.2. Initial condition dependence of transitions 115
6.5. Proposed solution for multiple predicted transitions 116
6.5.1. Sequential switch state configurations with multiple predicted transitions – the full solution 117
6.5.2. How to proceed without calculating the separatrix 118
6.6. Presenting the networks as families of trees 122

6.7. Conclusion 124
Chapter 7. Transition networks and treatments 125
7.1. Recap of the analysis steps to find the tree diagrams 125
7.2. Categories of switch configuration 127
7.2.1. Absorbing configurations 127
7.2.2. Absorbing loops 127
7.2.3. Transient loops 129
7.2.4. Transient configurations 130
7.3. Categories of disease network 130
7.3.1. Networks with one absorbing configuration 130
7.3.2. Networks with multiple absorbing configurations 131
7.3.3. Networks with absorbing loops 132
7.3.4. Networks with transient loops 135
7.4. Clinical interpretation of our networks 135
7.4.1. Pathogen clearing therapies. 136
7.4.2. Barrier reinforcing therapies 137
7.4.3. Dupilumab therapy 138
7.5. Therapies to bring about permanent recovery of symptoms 138
7.6. Conclusion 139
Chapter 8. Summary and Conclusions 140
8.1. Finding all possible connectivity of the regions in bifurcation diagrams 140
8.1.1. Requirements and limitations 141
8.1.2. General applicability to hybrid dynamical systems 143
8.2. Sensitivity analysis for BE bifurcation diagrams 143
8.2.1. Requirements and limitations 143
8.2.2. Possible extensions of this method to more complicated hybrid dynamical systems 144
8.3. “Racing switches” method to infer switch transitions from behaviours 145
8.3.1. Requirements and limitations 146
8.3.2. Extension of this method to other hybrid dynamical systems 146
8.4. Calculating the basins of attraction for each transition networks 147
8.4.1. Requirements and limitations 147
8.4.2. Extension to other hybrid dynamical systems 147
8.5. Future work 147
8.5.1. Specifying parameter ranges 148
8.5.2. Calculating separatrices to find the extent of each basin of attraction 148
8.5.3. Checking the effects of common AD therapies 149

8.5.4. Stratification of AD patients based on their transition tree 149
8.6. Final remarks 149
References 153
Appendix A Modifying the K-switch 1
Introduction 1
Two models for kallikrein regulation 1
Previously published, “old” model 1
New kallikrein model 3
K-switch conclusion 5
Effect of flattening the K-switch on the model’s bifurcation diagram 7
Conclusion 8
Appendix B Focal points and stability of the triple-switch model ODEs 9
Appendix C No bistability in the TGFβ feedback model using 1st-order cooperative inhibition 12
Appendix D Reparameterization of the Th1/Th2 mutual inhibition model 14
Appendix E Calculating threshold values as a function of dupilumab concentration for the systemic IL4 feedback+Th1/Th2 mutual inhibition model 16
Appendix F Calculating threshold values as a function of dupilumab concentration for the systemic IL4 feedback model 17
Appendix G Derivations of the boundary values for the stability regions in Figure 3.6 18
Appendix H Derivation of T-switch behaviours 21
Appendix I Focal point values and BE bifurcation lines with the revisable T-switch 26
Appendix J Standard parameter sets 29
Appendix K Behaviour-state dependence of the on the G- and RK-switch bifurcation diagrams 30
Appendix L Sensitivity analyses 42
Appendix M Numerical integration of the triple-switch model 46
Appendix N Numerically checking the bifurcation diagrams 51
Appendix O Calculating the shape of the basins of attraction 54
What to do in cases where multiple switch transitions occur in sequence 58
Appendix P Dupilumab therapy for treatment-resistant chronic AD 60
Concluding remarks on Dupilumab therapy 61

List of Figures

Figure 2.1 Cartoon of the RK- and T-switches of the double-switch model. 19
Figure 2.2 Relationship between the continuous- and discrete-valued state variables of the double-switch model. 20
Figure 2.3 A transition network. 21
Figure 2.4 Demonstration of the history-dependence of the switch state. 26
Figure 2.5 Predicting future switch transitions using the focal point value (a) predicting a switch will turn from off to on, and (b) on to off. 27
Figure 2.6 Bifurcation diagram for the double-switch model of AD 30
Figure 3.1 Yates model schematic. The cytokine signals from Th1 and Th2 cells govern T-helper cell activation, proliferation, and death. 34
Figure 3.2 A schematic of the IL4 feedback model 36
Figure 3.3 Bifurcation diagram for the model of equation (3.5) with $\pi 2 = 0.2$. The blue lines represent positive stable steady states, while the red lines represent positive unstable steady states. The system only has a pitchfork bifurcation at D=1 37
Figure 3.4 Hysteretic bistability produced by assuming order 2 inhibition with $\pi 2 = 1.5$. Blue lines represent stable steady states, while red lines are unstable steady states. The blue asterisks are limit points where a branch of steady states loses its bistability. The Th2 bistable region is between $D-$ and $D+$. $\pi 2$ represents the rate of increase of proliferation of Th2 cells in response to TGFβ stimulation. The threshold values $D-$ and $D+$ are caused by a transcritical and a pitchfork bifurcation, respectively. 37
Figure 3.5 Proposed assumption of dimerization of a signalling molecule TGFβ which promotes Th2 proliferation. 37
Figure 3.6 One parameter bifurcation diagrams obtained by MatCont, showing the four possible qualitative behaviours of the system. (a) monostable, (b) bistable with $D+ > 1$, (c) bistable with $D+ = 1$ and (d) irreversible bistability with $D- < 0$. (e) The values of $D-$ (black line) and $D+$ (thick green line) for different values of $\pi 2$ in the systemic IL4 feedback model. (f) Zoomed-in section showing the region where $D+ > 1$. 39
Figure 3.7 A schematic of the Th1-Th2 mutual inhibition model. 40
Figure 3.8 Regions of the $D1, D2$ plane in which each of the four steady states arise (a – d). Bistable regions are shown in (e) 42
Figure 3.9 Transitions between the four cases of stable positive steady states by changing the value $D2$ while keeping the ratio $D1/D2$constant and greater than 1. 43
Figure 3.10 One-parameter bifurcation diagrams of the Th1/Th2 mutual inhibition model when $D1/D2 = 2$. (a) for the Th1 population and (b) for the Th2 population. 43
Figure 3.11 One-parameter bifurcation diagrams of the Th1/Th2 mutual inhibition model when $D1/D2 = 0.5$. 45
Figure 3.12 Transitions between the two cases of stable positive steady states by changing the value $D2$ while keeping the ratio $D1/D2$ constant and less than 1. 45
Figure 3.13 Width Δ of the bistable region for Th2 cell number as we vary $\pi 1$, $\pi 2$ and δ. 46
Figure 3.14 Width of the bistable regions as we increase the ratio $C = D1/D2$ from 1 upwards. (a) Th1 only bistable region $\Delta 0$ and, (b) Th1/Th2 bistable region Δ. 46
Figure 3.15 Effect of dupilumab treatment on the critical values of the Th2 bistability $D-$ and $D+$ 47
Figure 3.16 One-parameter bifurcation diagrams when the dupilumab treatment has reduced the width of the bistable region to zero $\Delta = D+ - D- = 0$ at a cusp bifurcation for (a) systemic IL4/TGFβ feedback model of section 3.2.3, when we set $d = 1$ and (b) Th1 cell population for the systemic IL4+Th1/Th2 mutual inhibition model where the inset plot shows a zoomed-in region of the plot and

(c) Th2 cell population for the systemic IL4 feedback + Th1/Th2 mutual inhibition of section 3.2.4when we set $d = 1.2877$. 48
Figure 3.17 Th2 switch after dupilumab treatment. 49
Figure 3.18 Th2 switch functions. 50
Figure 3.19 Bifurcation diagrams for the double-switch model with modified K- and T-switches 51
Figure 4.1: The keratinocyte differentiation motif comprises positive feedback with cooperativity. . 53
Figure 4.2: Bistability observed in the kinetic model. 56
Figure 4.3 The effect of increasing NFκB on the threshold values of differentiation marker Stat3. 56
Figure 4.4 The effect of IL4 on the threshold values of differentiation marker Stat3. 58
Figure 5.1 Diagram of the triple-switch model. The state variables defined by switch functions are represented by squares, while state variables of the ODE are represented by circles. 61
Figure 5.2 Mechanisms of pathogen proliferation and killing. 62
Figure 5.3 Mechanisms of barrier formation and degradation. 63
Figure 5.4 Mechanisms of AMP production and degradation. 64
Figure 5.5 Dendritic cells are produced in response to innate immune response R at rate κD and degrade at rate δD 65
Figure 5.6 Transitions of each switch in our toy example. 69
Figure 5.7 Relationship between the continuous-valued and discrete-valued state variables. 72
Figure 5.8 Notation describing the focal points as a function of the configuration of the three switches. 72
Figure 5.9 Cartoon boundary equilibrium bifurcation curves for the RK-switch in the $\kappa P, \alpha P$ plane. 76
Figure 5.10 Behaviours of the T-switch as we vary the duty cycle and amplitude of the R-switch. 77
Figure 5.11 Boundary-equilibrium bifurcation curves for the G-switch in the $\kappa B, \delta B$ plane 78
Figure 5.12 Example bifurcation diagrams where the connectivity of behaviour regions differs. 79
Figure 5.13 The four ways to divide up the plane while both BE bifurcation lines maintain positive gradient: (a) The oscillation phenotype can be abolished entirely making it occur only for negative parameters (i.e. inadmissible values) but (b) recovered if the intersection point is in the positive plane. (c) The bistable phenotype can be abolished in the same way if the relative gradients of the lines are reversed but (d) recovered if the intersection point moves into the positive plane. 80
Figure 5.14 Possible arrangements of the BE bifurcation diagrams for the G-switch. (a) when condition 5.37 is not satisfied, and (b) when 5.37 is satisfied. 84
Figure 5.15 Number of bifurcation diagrams of each possible layout. 85
Figure 5.16 G-switch bifurcation diagrams with the R-switch and T-switch in different states. 88
Figure 5.17 Bifurcation diagrams for the RK-switch with the G-switch off (a), the G-switch on (b) and the difference between the G-switch off versus on causing five regions to form (c). 92
Figure 5.18 G-switch bifurcation overlaps under the p4 parameterisation. 93
Figure 5.19 Bifurcation diagrams for the RK-switch with the G-switch off (a), the G-switch on (b) and the nine regions produced by considering the changes brought about by turning on the G-switch (c) 93
Figure 5.20 The area of an arbitrary n-sided polygon can be calculated given the coordinates of the vertices. 95
Figure 5.21 Contiguous polygons share vertices, simplifying our calculations somewhat. 95
Figure 5.22 Local sensitivities of the RK-switch for the “p1” parameterisation 98
Figure 6.1 Switch state network, with every potential transition marked by an arrow. 100
Figure 6.2 Example transition networks 103
Figure 6.3 Transition diagrams arranged as trees, with transient configurations on branches leading to absorbing configurations and loops. 105
Figure 6.4 Illustration of the set of switch states $\boldsymbol{C}$ and set of behaviours $\boldsymbol{b}$ whose mapping we identify. 106

Figure 6.5 Regions of the G-switch bifurcation diagram which produces the state-behaviour relationship shown in Table 6.5. 107
Figure 6.6 Regions of the RK-switch bifurcation diagram where the state-behaviour relationship between the RK- and G-switch are uniform across each region. 108
Figure 6.7 Mapping from C to b expressed in the configuration-behaviour table in Table 6.7. 109
Figure 6.8 A stable absorbing configuration. 110
Figure 6.9 Switch transition from [0,0,1] to an absorbing configuration [0,0,0]. 111
Figure 6.10 Transition network with every predicted transition. 112
Figure 6.11 Vector field of the ODE showing regions where exotic behaviours emerge (a) In the red region below the black solution curve, the ODE dynamics cause the RK-switch to turn on. (b) In the red region above the black solution curve, the ODE dynamics cause the RK-switch to turn off. (c) The emergent behaviours. The blue region shows where the system is bistable as normal, but the green regions show a "bistable after transition" behaviour B* and B**. 113
Figure 6.12 Possible formation of an oscillating behaviour for focal points that would otherwise be predicted to form a bistable behaviour. The red region corresponds to values of the state variables which lead to oscillating switch behaviour. Blue regions produce bistable behaviour as expected while green regions produce bistable behaviour after the switch changes state. 114
Figure 6.13 Three-dimensional representation of the continuous-valued state space (neglecting the state variable A for visual simplicity). 115
Figure 6.14 Basins of attraction for both predicted transitions from [0,1,1]→[0,1,0] and from [0,1,1]→[0,0,1] 116
Figure 6.15 Transition diagram showing every predicted transition. Multiple-prediction switch configurations are highlighted in blue. Each transition is labelled with the switch that undergoes the transition. 117
Figure 6.16 Six possible paths through the network in Figure 6.15. 117
Figure 6.17 Timescale of dynamics of P and B can determine which switch transitions. If the dynamic of P is faster than that of B, the RK-switch wins the race, causing the RK-switch to transition. If the dynamic of B is faster than that of P, the G-switch wins the race, causing the G-switch to transition. 119
Figure 6.18 Transition network for the example parameter region with the assumed switch ranking {RK=1st, T=3rd, G=2nd}. 121
Figure 6.19 All six transition networks for the example parameter region, each corresponding to a permutation of the switch ordering. 122
Figure 6.20 Example family of networks all sharing the same parameter values but differing in the basin of attraction the into which the state variables (P, B, A, D) fall. 123
Figure 7.1 Complete workflow of our inference of switch state networks from the triple-switch model. Stages 1 through to 5 allow us to estimate the relative frequencies of each network of switch configuration produced by our model. 126
Figure 7.2 An absorbing configuration. 127
Figure 7.3 Hypothetical absorbing loop with an odd number of switches. This loop cannot occur. 127
Figure 7.4 Period-2 absorbing loops. 128
Figure 7.5 Period-4 absorbing loops. 128
Figure 7.6 Period-6 absorbing loops. 129
Figure 7.7 Transient loops. 129
Figure 7.8 Exit transitions from the transient loop, represented by blue arrows. 130
Figure 7.9 A transient configuration. 130
Figure 7.10 A family of transition networks where all siblings converge to the same absorbing configuration. 130

Figure 7.11 Distribution of absorbing configurations across all networks which have just one absorbing configuration. 131
Figure 7.12 A family of transition networks featuring multiple absorbing configurations. ..... **Error! Bookmark not defined.**
Figure 7.13 Number of absorbing configurations of each network, weighted by size of parameter region producing each type. 132
Figure 7.14 Family of transition networks with absorbing loops only. 133
Figure 7.15 Family of transition networks where some siblings feature absorbing loops alongside absorbing configurations. 134
Figure 7.16 Proportion of parameter space giving rise to multiple absorbing loops. 134
Figure 7.17 Proportion of parameter space giving rise to each maximum loop length in networks. 135
Figure 7.18 A family of transition networks, some of which feature transient loops. 135
Figure 7.19 Tree diagram with a squirrel scampering down the branches to the roots. 136
Figure 7.20 Tree diagram with squirrel jumping around to demonstrate chronic AD dynamics. Forcing the RK-switch to turn off, (a) external stimulation causes a jump transition from the root to the branch and then (b) the disease dynamics play out as the squirrel scampering back down that branch. 137
Figure 7.21 Tree diagram with squirrel jumping around to demonstrate chronic AD dynamics. Forcing the G-switch to turn off, (a) external stimulation causes a jump transition from the root to the branch and then (b) the disease dynamics play out as the squirrel scampering back down that branch. 138
Figure 7.22 Jumping from one tree to another for a genotype where multiple absorbing configurations exist. 138
Figure 8.1 Connectivity networks of bifurcation diagrams. 141
Figure 8.2 Four ways to arrange the BE bifurcation lines in the plane (reproduced from Figure 5.13). 142
Figure 8.3 Cartoon representation of the sensitivity analysis method 143
Figure 8.4 Example configuration-behaviour mapping 146

List of tables

Table 2.1 Nominal parameter values for the double-switch model. Reproduced from [18]. 21
Table 2.2 Notation used to specify focal point values for the continuous-valued state variables in our hybrid dynamical system. We use a letter as an argument when the state of that switch is yet to be specified. 23
Table 2.3 Summary of the inferred switch state based on knowledge of the current value of the continuous-valued state variable. 26
Table 2.4 Transition conditions: conditions where we can infer the switch must change state in the future. 27
Table 2.5 Switch behaviours as a result of the turning on and turning off conditions of the RK-switch. 28
Table 2.6 Boundary-equilibrium bifurcation lines defining the edges of the regions.. 28
Table 3.1 Meaning and nominal values of the nondimensionalized model parameters of the Th1/Th2 mutual inhibition model. 40
Table 3.2 Table summarising the steady states which exist for the six parameter ranges of $D2$. $x1$ and $x2$ represent the focal point solutions of the ODE when $D1/D2 > 1$. 44
Table 3.3 Nominal parameters for the modified T-switch 51
Table 4.1 Nominal parameters of the kinetic model. 54
Table 4.2 Definition of the threshold value $B-$ and $B+$ based on the modulation of the thresholds found in section 4.2.2. 58
Table 5.1 State variables of the triple-switch model 65
Table 5.2 Nominal parameter values and ranges considered in our analysis. 65
Table 5.3 Configuration-behaviour table of our toy example. 69
Table 5.4 Five stages of bifurcation analysis 71
Table 5.5 Transition conditions: conditions where we can infer the switch must change state in the future. 73
Table 5.6 Switch behaviours as a result of the turning on and turning off conditions. 74
Table 5.7 Boundary-equilibrium bifurcation curves defining the edges of the regions satisfying the transition conditions defined in Table 5.5. 74
Table 5.8 Only two possible bifurcation diagram layouts exist for the RK-switch behaviours. The letters represent the behaviours present in the positive $\kappa P, \alpha P$ plane in each bifurcation diagram. Only the green cell and the blue cell are possible. The white cells are impossible. 82
Table 5.9 Two possible scenarios for bifurcation diagram layouts arise from the gradient condition. The letters represent the behaviours present in the positive $\kappa B, \delta B$ plane. 84
Table 5.10 Enumerated possibilities for parameter regions with different behavioural regions shown in the bifurcation diagrams for the RK- and G-switches. 85
Table 5.11 BE bifurcation lines for all the switches' behaviours. 86
Table 5.12 Dependence of the BE bifurcation diagrams on the switch states. 87
Table 5.13 Effect of the R- and T-switches on the threshold values of the G-switch, $B-$ and $B+$. 87
Table 5.14 Table summarising mapping of R- and T-switch state to G-switch behaviour 89
Table 5.15 Number of G-switch behaviour regions produced by each parameterisation. 89
Table 5.16 Switch configuration-behaviour relationship for the G-switch bifurcation diagram regions in Figure 5.16 90
Table 5.17 Regions of parameter space where the RK-switch shows consistent configuration-behaviour relationships 92
Table 5.18 Behaviour regions produced by the RK-switch when the G-switch changes state under parameterisation p4. 93
Table 6.1 Switch configuration 101

Table 6.2 Ways of representing the switch configuration graphically and in mathematical notation in the rest of this chapter. 102
Table 6.3 Switch behaviours represented by the vector $\boldsymbol{b} = B, R, Os$ 105
Table 6.4 Behaviour of the T-switch for different RK-switch states 107
Table 6.5 Example behaviour of the G-switch for different RK-switch and T-switch states for region 1 of the parameter plane $\kappa B, \delta B$ shown in Figure 6.5. 107
Table 6.6 Example behaviour of the RK-switch for different G-switch states. 108
Table 6.7 Example of a configuration-behaviour table summarising the effects of crosstalk between the three switches in the example region of parameter space. 109
Table 6.8 Transition table, summarising the inferences we can make about future switch transitions. 110
Table 6.9 Predicted transitions (or in this case lack of transitions) of the switches for the switch state configuration [0,0,0]. 110
Table 6.10 Predicted transition of the switches for the switch states [0,0,1]. 111
Table 6.11 Predicted transition of the switches for the switch states [0,1,1]. 111
Table 6.12 Predicted transitions for the configuration-behaviour Table 6.7 111
Table 6.13 Permutations of switch rankings in races where all three switches compete to reach their threshold values. 120
Table 6.14 Table of predicted transitions given the assumption that the state variable P has the fasted time scale, followed by B, followed by D. This means the switches win head-to-head races in the order RK-switch 1st, G-switch 2nd, T-switch 3rd. 121
Table 6.15 Number of transition trees 124
Table 8.1 Two possible ways in which the BE bifurcation lines can arrange in the plane given all parameter values must be positive. 142
Table 8.2 Configurations of the steady states of the ODE for various ranges of $\pi 2$ and D. 13

Acronyms and terminology

AD	Atopic dermatitis
APCs	Antigen-presenting cells
DCs	Dendritic cells
TCRs	T-cell receptors
PAR2s	Protease-activated receptors
KLK	Kallikrein enzymes
LEKTI	Lympho-epithelial kazal-type-related inhibitor, which competitively inhibits KLK
Th1, Th2	T-helper cells of type 1, and type 2
IL4, IL13	Interleukin-4, Interleukin-13
ODEs	Ordinary differential equations

Chapter 1. Introduction

1.1. Background

The skin is our shield from and physical connection to the outside world. It provides not only a physical barrier but also provides our first line of defence against microbial and viral pathogens as well as chemical assaults. Skin is organized into three layers which differ in structure and function: the epidermis, dermis, and hypodermis.

The uppermost layer, the epidermis, provides the majority of protective functions of the skin. This layer acts as a barrier preventing water loss [1], can resist physical abrasion and chemical damage, and is sufficiently flexible to bend and stretch, all the while accommodating the constant scrutinising presence of the immune system. The surface of this layer is populated by commensal bacteria, viruses and fungi which help to displace or repel pathogenic species. The epidermis is around the thickness of this piece of paper, varying from around 50 microns on the eyelids to 500 microns on the sole of the foot [2]. This thin layer, and the complex array of immune cells lurking in and beneath its surface, provide our first line of defence against the diverse and changeable environments we occupy. When this defensive system goes wrong, the impacts can be severe, leading to the development of a range of diseases, including those under the umbrella term “dermatitis”.

Atopic dermatitis (AD) is a chronic inflammatory skin condition characterized by dry, itchy, and inflamed skin. AD is one of the most common inflammatory skin disease, affecting up to 20% of children and 10% of adults in high income countries [3]–[5] and with a globally increasing incidence [6], [7]. It can occur at any age, but often starts in childhood. An onset before the age of 6 is estimated for 80% of AD patients [8]. After initial development, symptoms can continue for several years but can go on to develop a chronic relapsing-remitting nature stretching into adulthood [9], [10].

While this form of dermatitis falls under the umbrella of atopic diseases, which also includes asthma and allergic rhinitis (hay fever), the term “atopic” can be a little confusing in this context because it commonly refers to the tendency to produce IgE antibody or to develop comorbid atopic diseases such as allergic rhinitis or asthma. Despite this, a significant proportion of patients with AD, particularly those with mild disease (20-30%) do not show signs of IgE-mediated sensitisation or display those comorbidities [11]. As a result of this confusion, the terms atopic dermatitis, atopic eczema or simply eczema are commonly used interchangeably. Throughout this thesis we do not focus on IgE production, but nevertheless refer to the disease as “atopic dermatitis”.

AD is thought to be caused by a combination of genetic and environmental factors [12]. It is clinically heterogeneous, with observed lifetime prevalence variable between different countries, and even between populations with high genetic similarity (such as individuals from similar ethnic groups), suggesting that disease dynamics are affected by environmental factors [7]. While epidemiological studies have identified a range of environmental factors such as living in an urban environment, dry climatic conditions, and typical Western diet high in sugars and polyunsaturated fatty acids [13], effect sizes for these are generally small. The strongest known genetic risk factor is in the FLG gene coding for the structural epidermal protein filaggrin. Loss of function of this gene is known to increase risk of developing AD by around threefold [14], suggesting structural deficiencies are essential to the pathophysiology. Genetic association studies have identified a total of 31 loci associated with AD, which include genes associated with innate immunity, T-helper cell regulation and Th2 cytokines IL-4 and IL-13 [15], [16]. Variations in the relative magnitude and timing of these environmental and genetic influences are thought to explain the observed clinical heterogeneity [12].

Given what is known of the complex interplay between environmental factors and genetic risk factors, it is perhaps unsurprising that AD pathophysiology remains incompletely understood. The disease is

thought to arise as the result of interplay between multiple dysfunctional systems, namely epidermal barrier regulation, predominantly type-2-skewed immune dysregulation, and abnormalities in the skin microbiome [17]. While a vast array of individual mechanisms relating to each system in the context of AD are known, it remains unclear how their interactions lead to the relapsing-remitting disease course and sometimes treatment-resistant form of chronic AD.

1.2. A mathematical modelling approach to understand chronic atopic dermatitis (AD)

In this thesis, we develop a mechanistic mathematical model for chronic AD. With this model, we aim to understand how chronic AD is sustained and why it has a relapsing-remitting phenotype.

We take a systems biology approach, integrating observed or proposed biological mechanisms into mathematical models representing regulatory systems thought to play a role in the pathophysiology of the disease. We integrate each of the mechanisms together into one coherent mechanistic model which provides insight into how their interactions bring about the complex dynamics of chronic AD.

A feature of AD pathophysiology, and indeed many biological systems, is the range of timescales over which it operates. Multiple tissues made up of different cell types, each transcribing a variety of genes in response to signalling molecules and detected pathogens all interact together to give rise to the disease dynamics we see. The multitude of interactions occurring at various timescales makes these dynamics challenging to predict and explain. It is the job of a mathematical model to distil this messy network of interactions to a more tractable and comprehensible model which can inform future experimentation and treatment decisions.

We use a few of different methods to distil the essential regulatory interactions from the messy biology. Firstly, where our mathematical model includes unnecessary detail and complexity, we undergo a process of removing or simplifying various parts and then checking the effects of the change. If we are able to verify the performance of the model against experimental data or knowledge, we consider keeping the change we made.

Secondly, we break down the problem into smaller sub-models considering particular fast-timescale regulatory motifs such as particular parts of the immune response of known importance to AD. We can then examine these component parts in isolation and work out how each part functions, searching for experimental observations of these systems in isolation if possible, and using these to guide our model development and verification. Once we understand the component parts, we can build up a model of their interactions while ensuring we are still able to interpret what each mechanism actually does.

Finally, we explicitly separate different timescales in our model using a quasi-steady-state assumption. This means we assume the fast timescale processes have been allowed enough time that their short-lived dynamical behaviour has died away, leaving the variables we are interested in close to their steady-state values. The steady-state behaviours of our fast-timescale sub-models are all incorporated into a hybrid dynamical system.

Hybrid dynamical systems are so-called because they constitute a hybrid of both continuous- and discrete-valued variables. For our purposes, the discrete valued variables represent the various quasi-steady-states of the fast-timescale systems while the continuous-valued variables represent the slower timescale. The discrete-valued variables are able to instantaneously switch from one value to another as the model's dynamics evolve. This instantaneous switch represents the rapid transition of the fast-timescale submodel from one steady-state to another under particular conditions.

Meanwhile, the slow-timescale processes of tissue repair and growth, cell migration and pathogen growth are represented by the continuous-valued variables using an ordinary differential equation. This thesis will discuss the process of wiring together our fast- and slow-timescale mechanisms into this hybrid dynamical system. This approach to timescale separation is identical to that used by Domínguez-Hüttinger et al. [18] where a mathematical model of the development of AD from health to allergic sensitisation was described.

1.3. Aims and objectives

In this thesis, we build a model of the long-term dynamics of AD, reflecting the pathophysiology of patients experiencing the chronic and potentially treatment-resistant form of this disease. We aim to use this mathematical model to infer the reasons for the emergence of hard-to-predict relapsing-remitting disease phenotypes.

Aim: To develop a mechanistic model for long-term dynamics of AD and verify model performance by elucidating the range of dynamics it predicts and checking it reflects known chronic AD behaviour and reproduces chronic AD risk factors as expected.

To achieve this aim, we carry out a process of iterative model development, where we test and verify components of our model for chronic AD. The model we develop will be structured as a hybrid dynamical system, which presents certain challenges to us in pursuing the second aim of elucidating its range of dynamics. We wish to carry out something like a bifurcation analysis, identifying how parameter changes can bring about changes to the model's dynamics. No ready-made methods exist for analysing hybrid dynamical systems in the way we require. For this reason, we pursue the following objectives:

Objective 1: Examine the previous mathematical model for the onset of AD [18], and identify which modifications are required to extend its scope to the long-term dynamics of chronic AD.

Objective 2: Modify and extend components of the previous mathematical model for the onset of AD [18], carrying out an iterative process of model modification and verification.

Objective 3: Develop analytical methods required to identify the disease dynamics arising from the proposed physiological mechanisms. These methods include bifurcation analysis tools and parameter sensitivity tools.

By pursuing each of our objectives in turn, we will develop a model for chronic AD, and a suite of mathematical tools for its analysis. The model takes the form of a hybrid dynamical system; chosen to represent the differing timescales of the biological processes in AD aetiology. Methods for the analysis of hybrid dynamical systems are relatively underdeveloped, so we fill this research gap by developing a selection of novel methods for analysing hybrid dynamical systems. We will then apply these methods to verify our model's performance and infer the impact of dupilumab therapy on the AD aetiology. Finally, we will discuss how these methods can generalise to other hybrid dynamical systems that could conceivably be applied related diseases such as psoriasis.

Chapter 2. Previous work on modelling atopic dermatitis

The pathogenic mechanisms of chronic AD are poorly understood, specifically how interactions between known biochemical processes can lead to the complicated long-term dynamics observed clinically. In this chapter, we reexamine a previous attempt at modelling AD and identify why this model was not able to capture the long-term dynamics of AD. We discuss the mathematical methods used to analyse the previous AD model, which introduces the methods we develop throughout this thesis. Finally, we summarise which parts of the model would benefit from modification or extension, helping us form a model later in this thesis which fills the knowledge gap by elucidating how known biochemical mechanisms can lead to long-term AD dynamics.

2.1. Previously published "double-switch" model of AD onset

Domínguez-Hüttinger et al. developed a multi-scale mathematical model of AD onset [18] which considers system-level interactions between immune system and skin barrier. Systems operating at differing timescales were brought together into a single model by using a "hybrid dynamical system" which allowed the model to summarise faster-timescale dynamics using a series of "switches". The model contains both continuous-valued and discrete-valued state variables, coupled together by a set of ordinary differential equations (ODEs). The dynamics of the continuous-valued variables are governed by the set of ODEs, which is in turn dependent on the values of the discrete-valued variables. The values of the discrete-valued variables depend on the dynamics of the continuous-valued variables.

The relatively slow dynamics of, for example, barrier repair and growth, which occur on the timescale of days, are modelled by the ODE on a continuous scale. Meanwhile, the instantaneous "switching" between the discrete-valued states represents the faster-timescale dynamics of the immune responses, which occur on the timescale of minutes to hours. The model considers two fast timescale phenomena: the innate immune response, and the Th2 response, each of which is designated a "switch". For this reason, we generally refer to the model as the "double-switch model".

2.1.1. Model equations

The double-switch model of AD consists of five state variables. Three of them continuous-valued and two discrete-valued (each taking one of two possible values). The three continuous-valued variables' dynamics are governed by equation 2.1.

$$\begin{aligned}\frac{dP(t)}{dt} &= P_{\mathrm{env}}\frac{\kappa_{\mathrm{P}}}{1+\gamma_{\mathrm{B}}B(t)} - \alpha_{\mathrm{I}}R(t)P(t) - \delta_{\mathrm{P}}P(t),\\ \frac{dB(t)}{dt} &= \kappa_{\mathrm{B}}\frac{1}{1+\gamma_{\mathrm{R}}R(t)}\frac{1}{1+\gamma_{\mathrm{T}}T(t)}(1-B(t)) - \delta_{\mathrm{B}}K(t)B(t),\\ \frac{dD(t)}{dt} &= \kappa_{\mathrm{D}}R(t) - \delta_{\mathrm{D}}D(t).\end{aligned} \tag{2.1}$$

The variables of this ODE, P(t), B(t) and D(t) correspond to latent properties of AD, which may not be immediately apparent to a patient or clinician.

The infiltrated pathogen load (P) represents the density of pathogenic organisms on the skin surface. It increases by the penetration of environmental stressors (P_{env}) through the barrier (B) at rate κ_{P} and inhibited by barrier integrity at rate γ_{B}. Infiltrated pathogens (P) are eradicated by innate immune responses (R) at rate α_{I} and also naturally degrade at rate δ_{P}.

Barrier integrity (B) is determined by the balance between its production and degradation. Barrier production is described by a number between 0 (no ability to repel pathogens or defend against chemical or physical assault) and 1 (healthy barrier). This phenomenological representation is built

into the structure of the ODE through the $\left(1 - B(t)\right)$ term. The capacity to restore barrier function at rate κ_{B} is impaired by inflammation associated with innate immune responses (R) at rate γ_{R}, and inflammatory cytokines IL4 and IL13 produced by differentiated Th2 cells (T[1]) at rate γ_{T}.Degradation of the barrier occurs as a result of desquamation mediated by active kallikreins (K) at rate δ_{B}.

The concentration of activated dendritic cells (D) increases at rate κ_{D} while inflammatory signals from innate immune responses (R) persist, but dendritic cells degrade naturally at rate δ_{D}. The concentration of dendritic cells (D) brings about the activation of Th2 cells, which form the adaptive immune branch of this model.

The two discrete-valued state variables R(t) and T(t), correspond to symptoms of AD which are more immediately apparent to patients and clinicians in the form of an inflammatory response. R(t) denotes the level of activated immune receptors, and T(t) the level of GATA3 transcription (relative to the maximum transcription level).

An additional discrete-valued variable K(t) represents the concentration of active kallikrein enzymes in the epidermis. Note that K(t) is simply a scaled version of R(t), so we don't consider it as strictly speaking an additional independent variable. The values of R(t) and K(t) always move in concert, related by a simple equation that does not vary with time, and so they may be considered together as a single dynamical variable. We often lump these two together in discussions, referring to them as the "RK" variable.

Equation 2.1 formalises the model assumptions using law of mass action with zero-order production and first-order degradation terms, except for the barrier production term, $\kappa_{\mathrm{B}}(1 - B(t))$, which represents the phenomenological recovery of skin barrier in response to disruption. This barrier production term is designed to be self-limiting, causing the barrier integrity B to be naturally limited to the range 0 (no barrier integrity) to 1 (fully competent barrier). The inhibitory rates are described in the form of 1/(1 + x).

The dynamics of the discrete-valued RK and T variables are described as bistable switches:

$$(R(t), K(t)) = \begin{cases} (R_{\text{off}}, K_{\text{off}}) \rightarrow (R_{\text{on}}, K_{\text{on}}) & \text{if } \dot{P}(t) > 0 \text{ and } P(t) = P^{+}, \\ (R_{\text{on}}, K_{\text{on}}) \rightarrow (R_{\text{off}}, K_{\text{off}}) & \text{if } \dot{P}(t) < 0 \text{ and } P(t) = P^{-}, \end{cases} \tag{2.2}$$

where R_{off}, K_{off} and R_{on} are constants, and $K_{\text{on}} = m_{\text{on}} P(t) - \beta_{\text{on}}$ and

$$T(t) = T_{\text{off}} \rightarrow T_{\text{on}} \text{ if } \dot{D}(t) > 0 \text{ and } D(t) = D^{+} \tag{2.3}$$

The R and K values (which always change in concert) are determined by the dynamics of P, and the T value is determined by the dynamics of D. This relationship between the continuous-valued and discrete-valued state variables is summarised in Figure 2.2 by the black arrows, while the feedback from the discrete-valued to continuous-valued variables is shown by the red arrows.

The dynamics of the continuous-valued variable P determines the values of R and K, while the dynamics of the continuous-valued

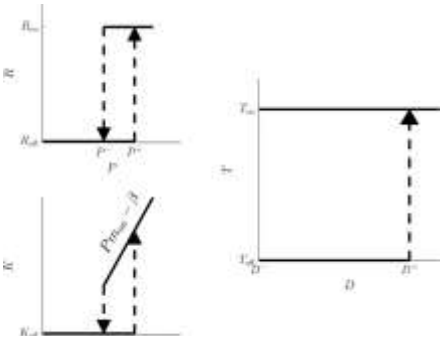


Figure 2.1 Cartoon of the RK- and T-switches of the double-switch model.

[1] Note that in the published definition of the double-switch model [18], what we refer as the "T-switch" was originally called the "G-switch". The letter G was chosen to refer to the GATA3 gene in Th2 cells. We have chosen to replace the G with a T to refer to Th2-cells. We make this decision because, later in this thesis, we introduce a new switch representing the granular layer of the epidermis, which we call the "G-switch". This new G-switch is not related in any way to the "G-switch" in the double-switch model.

variable D determines the discrete-valued variable T. These discrete-valued variables in turn affect the dynamics of the continuous-valued variables.

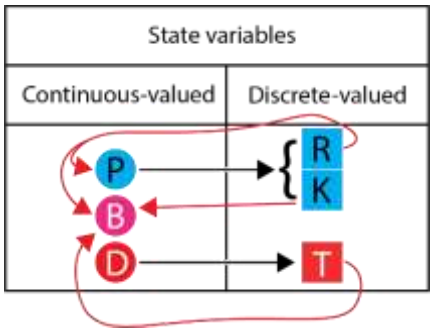


Figure 2.2 Relationship between the continuous- and discrete-valued state variables of the double-switch model.

2.1.2. Parameter values

When the double-switch model was developed, parameter values were estimated from values found in literature, taken from approximations of bifurcation diagrams to clinical data, or carried over from previous publications. Nominal values for the double-switch model shown in [18] are duplicated in Table 2.1 along with references justifying their origin.

In this thesis, the parameter values we set in our model development are similarly fluid, taking values based on our bifurcation analyses rather than being the result of comparisons with clinical data. The models we present are not intended as quantitative predictors of disease dynamics, but rather as qualitative representations of the patterns of behaviour observed in chronic AD.

2.1.3. Qualitative phenotypes of AD onset

Previous analysis of the double-switch model revealed that the dynamics always fell into four distinct categories of behaviour.

- recovery (where the barrier was able to recover to a healthy state even after a pathogen challenge)
- bistability (where it could remain in either a healthy or unhealthy state)
- oscillation (where it would repeatedly switch between healthy and unhealthy)
- damage (where the barrier would always become unhealthy)

The locations in the parameter space which results in each of the categories are separated by boundary-equilibrium bifurcations. The mechanistic nature of the model makes it possible to interpret the reasons each behaviour emerged.

In the next section, we first describe why these four categories of behaviour are not sufficient to represent chronic AD dynamics, and then show how we calculate analytical expressions for the boundary-equilibrium bifurcation lines. This examination of the bifurcation methods will give us a good grounding for the methods we develop for our chronic AD model later in this thesis.

2.2. Reasons why the double-switch model does not reflect chronic AD

To understand how the double-switch model is limited to describing allergic sensitisation and why it cannot represent the dynamics of chronic AD, it is useful to think about the hybrid dynamical system as a "transition network" as illustrated in Figure 2.3. Thinking about the hybrid dynamical system in this way will help us make sense of the dynamics and discuss our methods and results throughout this thesis.

Table 2.1 Nominal parameter values for the double-switch model. Reproduced from [18].

SYMBOL	Name	Nominal value	Reference
P_{env}	Environmental stress load	95 (mg/mL)	[19]
γ_B	Barrier-mediated inhibition of pathogen infiltration	1	[19]
κ_P	Nominal skin permeability	0.6 (1/day)	[19]
α_I	Rate of pathogen eradication by innate immune responses	0.25 (1/day)	[19]
δ_P	Basal pathogen death rate	1 (1/day)	[19]
κ_B	Barrier production rate	0.5 (1/day)	[19], [20]
γ_R	Innate immunity-mediated inhibition of barrier production	10	[19]
δ_B	Rate of kallikrein-dependent barrier degradation	0.1	[19]
γ_T	Adaptive immunity-mediated inhibition of barrier production	1	[21]
κ_D	Rate of DC activation by receptors	4 cells/(mL × day)	[22], [23]
δ_D	Rate of DC degradation	0.5 (1/day)	[24]
P^-	Receptor inactivation threshold	26.6 (mg/mL)	[25]
P^+	Receptor activation threshold	40 (mg/mL)	[25]
D^+	GATA-3 activation threshold	85 (cells/mL)	[23], [26]
R_{off}	Receptor off level	0	[25]
R_{on}	Receptor on level	16.7	[25]
T_{off}	GATA-3 off level	0	[26]
T_{on}	GATA-3 on level	1	[26]
K_{off}	Kallikrein off level	0	[25]
m_{on}	Slope of the linear relation between P(t) and Kon	0.45	[25]
β_{on}	Y-intercept of the linear relation between P(t) and Kon	6.71	[25]

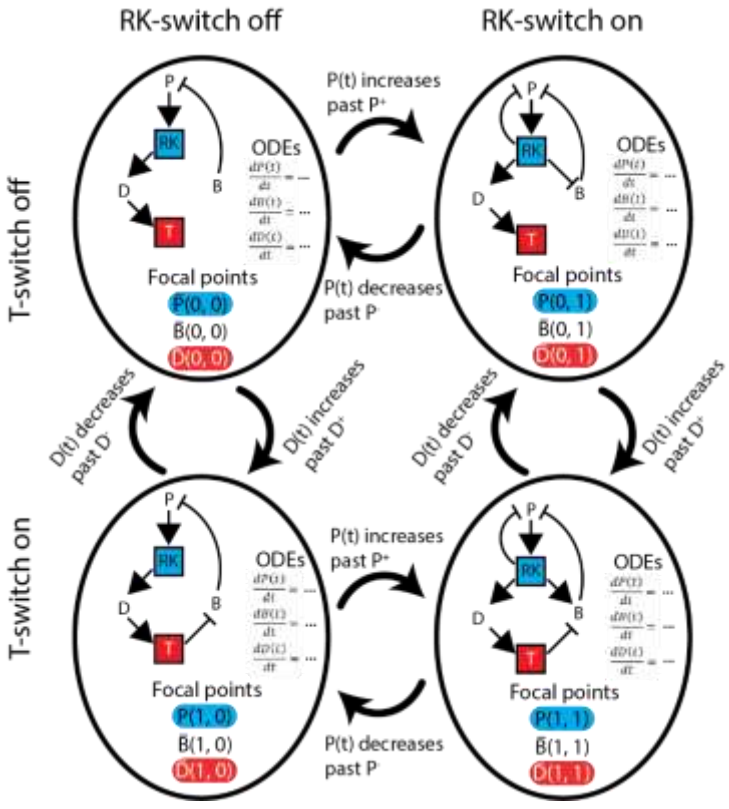


Figure 2.3 A transition network. Each node represents a particular configuration of switch configurations. Each has a slightly different schematic of mechanisms, with a different ODE and associated focal points. Transitions between these nodes are annotated with the threshold value that must be crossed to bring about that transition.

In Figure 2.3, each node in the transition network represents a particular configuration of the switches, and the arrows between the nodes represent a switch transitioning state. The dynamics obey the ODE for the current switch state until the state variables cross one of the switches' threshold values, at which point the switch state instantaneously transitions to the new state, and the dynamics

start following the new ODE. We will make heavy use of this "transition network" view throughout this thesis.

The switch configurations represent AD symptoms which are immediately perceptible to a patient, such as redness, inflammation, and itch. The continuous-valued variables of the ODE represent the latent features of the disease, such as the pathogen load, integrity of the barrier and concentration of dendritic cells in the draining lymph-nodes. Because it focusses on the state of the symptoms of AD, the transition network presents us with a map of the disease progression, giving us some idea of how particular combinations of symptoms may follow one another.

Previous analysis in [18] showed that the structure of the double-switch model limited the potential transitions shown in Figure 2.3, limiting the range of motion around the transition network in two ways. Firstly, the T-switch is irreversible meaning we can never transition back to the "T-switch off" states once it's been activation. Secondly, activating the T-switch has no effect on the qualitative dynamics which were identified. The underlying problem is that a transition of the T-switch can never cause the RK-switch to undergo a transition. This effectively limits all the dynamics of the model to merely the behaviours we listed; recovery, bistability, oscillation and damage, all of which only involve the RK-switch.

Chronic AD, by contrast, is characterised by long-lasting and often difficult-to-predict dynamics which are often relapsing-remitting in nature [27]. This feature is not adequately explained by the four behaviours of which the double-switch model is capable.

We propose that a mechanistic model for chronic AD should produce a sequence of transitions of switch state, occurring in a potentially long series that characterises the hard-to-predict nature of the disease dynamics. We imagine a model where one switch transitions state causing the ODE dynamics to trigger a different switch to transition state, which in turn causes another transition and so on. The double-switch model did not do this due to some structural limitations which we describe in the next section 2.3.

2.3. Understanding the four behaviours of the double-switch

To understand why there are only four qualitative phenotypes and why sequences of switch transitions are not possible for the double switch model, it is necessary to first discuss the method used to carry out bifurcation analysis. We use this method to work out which parameter values cause each qualitative phenotype. The method also shows us why the activation of the T-switch had no effect on the qualitative phenotype, making it impossible for series of transitions to occur as we expect for chronic AD.

2.3.1. Finding the focal point values of the double-switch model

We can calculate the boundary-equilibrium bifurcation lines analytically by first calculating the focal points of the ODE for each of the four switch configurations ([0,0], [0,1], [1,0] and [1,1]). We then rearrange these expressions to work out when the focal point values are above or below each switch's threshold values.

When writing the focal point values, it's important to keep track of the state of the RK- and T-switch states because the value of the focal point can be different for each configuration of the switches. We write the switch states as either 0 (off) or 1 (on) within brackets next to our expressions for the focal point. We will always write the RK-switch value as the first argument and the T-switch value as the second argument. For example, $\bar{P}(0,1)$ represents the focal point value for state variable P when the RK-switch is off, and the T-switch is on. In our calculations, we sometimes want to hold off from specifying the state of a switch until we actually evaluate the focal point value. In such instances, we

specify the switch state by the letter corresponding to that switch. For example, $\bar{P}(1,T)$ represents the focal point value of state variable P when the RK-switch is on, and the T-switch is in state T (yet to be specified). This notation is summarised in Table 2.2.

Table 2.2 Notation used to specify focal point values for the continuous-valued state variables in our hybrid dynamical system. We use a letter as an argument when the state of that switch is yet to be specified.

Notation for focal points

$\bar{X}$({RK-switch},{T-switch})

Focal point value	R- and K-switch	T-switch
$\bar{P}(0,0)$ $\bar{D}(0,0)$	R_{off} K_{off}	T_{off}
$\bar{P}(0,1)$ $\bar{D}(0,1)$	R_{off} K_{off}	T_{on}
$\bar{P}(1,0)$ $\bar{D}(1,0)$	R_{on} K_{on}	T_{off}
$\bar{P}(1,1)$ $\bar{D}(1,1)$	R_{on} K_{on}	T_{on}
$\bar{P}(1,T)$ $\bar{D}(1,T)$	R_{on} K_{on}	T

4 configurations

We can produce expressions for all four configurations of the switches by simply solving the ODEs after setting the derivatives to zero. Then, when we substitute in the values for the R, K and T values, we obtain the focal point value for that switch configuration.

When the RK-switch is off, the value of the K-switch is K_{off}, but when the RK-switch is on, the value becomes $K_{\text{on}} = m_{\text{on}}P(t) - \beta$ and since $P(t)$ is a state variable of the ODE, this changes the structure of the ODE somewhat and requires different treatment. The state of the T-switch causes no such complication because both the T_{off} and T_{on} states are constants, so the structure of the ODEs remains the same in both states. For this reason, we solve the problem once with the K-switch off, and then solve it again with the K-switch on and our results apply to both configurations of the T-switch.

When the RK-switch is off, the focal points are calculated by setting $R(t) = R_{\text{off}}$ and $K(t) = K_{\text{off}}$. We represent the T-switch value by the letter T, and can simply substitute either $T = T_{\text{off}}$ or $T = T_{\text{on}}$ into the resulting equation 2.4 for the focal points:

$$\begin{aligned} 0 &= P_{\text{env}}\frac{\kappa_{\text{P}}}{1+\gamma_{\text{B}}\bar{B}(0,T)} - \alpha_{\text{I}}R_{\text{off}}\bar{P}(0,T) - \delta_{\text{P}}\bar{P}(0,T) \\ 0 &= \kappa_{\text{B}}\frac{1}{1+\gamma_{\text{R}}R_{\text{off}}}\frac{1}{1+\gamma_{\text{T}}T}(1-\bar{B}(0,T)) - \delta_{\text{B}}K_{\text{off}}\bar{B}(0,T) \\ 0 &= \kappa_{\text{D}}R_{\text{off}} - \delta_{\text{D}}\bar{D}(0,T) \end{aligned} \tag{2.4}$$

We can rearrange these equations to find the focal point values $\bar{P}, \bar{B}, \bar{D}$

$$\bar{P}(0,T) = P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1+\gamma_{\text{B}}\bar{B}(0,T)\right)(\delta_{\text{P}}+\alpha_I R_{\text{off}})}$$
$$\bar{B}(0,T) = \frac{\kappa_B}{\kappa_B + \delta_B K_{\text{off}}(1+\gamma_R R_{\text{off}})(1+\gamma_T T)} \tag{2.5}$$
$$\bar{D}(0,T) = \frac{\kappa_D R_{\text{off}}}{\delta_D}$$

In these expressions, we must first evaluate the value of $\bar{B}(0,T)$ before using this result to calculate $\bar{P}(0,T)$.

When the RK-switch is on, the focal points are the solutions to the slightly more complicated (note the change highlighted in yellow).

$$0 = P_{\text{env}} \frac{\kappa_{\text{P}}}{1+\gamma_{\text{B}}\bar{B}} - \alpha_{\text{I}} R_{\text{on}} \bar{P} - \delta_{\text{P}} \bar{P}$$
$$0 = \kappa_{\text{B}} \frac{1}{1+\gamma_{\text{R}} R_{\text{on}}} \frac{1}{1+\gamma_{\text{T}} T} (1-\bar{B}) - \delta_{\text{B}} (m_{on}\bar{P} - \beta)\bar{B} \tag{2.6}$$
$$0 = \kappa_{\text{D}} R_{\text{on}} - \delta_{\text{D}} \bar{D}$$

Again, we need to find the solution for $\bar{B}$ before we can evaluate the $\bar{P}$ and $\bar{D}$ solutions. The solution for $\bar{B}$ takes the form of roots of a quadratic equation (2.7), with some fairly complicated coefficients.

$$a\bar{B}^2 + b\bar{B} + c = 0 \tag{2.7}$$

To help define the coefficients concisely, we're going to lump common terms together, representing them by the letters X and Q as

$$Q = \frac{P_{\text{env}}\kappa_P}{\delta_P + \alpha_I R_{\text{on}}}, \quad X = (1+\gamma_R R_{\text{on}})(1+\gamma_T T) \tag{2.8}$$

The quadratic coefficients can now be written as

$$a = \kappa_B\gamma_B - \delta_B X\beta\gamma_B$$
$$b = \kappa_B + \delta_B X(m_{on}Q - \beta) \tag{2.9}$$
$$c = -\kappa_B$$

We can denote the two solutions as $\bar{B}_{\text{left}}$ and $\bar{B}_{\text{right}}$ where $\bar{B}_{\text{left}} < \bar{B}_{\text{right}}$. Firstly, we can note that the y-intercept of the quadratic (the c coefficient) is always negative since all the parameters, including κ_B, are always positive-valued by assumption. This means that so long as the a coefficient is positive, we are guaranteed to have one negative real and one positive real root, such that $\bar{B}_{\text{left}} < 0$ and $\bar{B}_{\text{right}} > 0$. Since the state variables must be positive, this means that the system would have only one focal point in the admissible range. We can write down the condition for this situation to occur as $a > 0$, which gives us an inequality condition for our parameter values as

$$\begin{aligned} a &> 0 \\ \kappa_B\gamma_B - \delta_B X\beta\gamma_B &> 0 \\ \beta &> \frac{\kappa_B}{\delta_B X} \end{aligned} \tag{2.10}$$

This condition effectively tells us that the K-switch's sloped section can only be a maximum height above the P-axis for the system to have a positive focal point value. Throughout our analysis, we set β such that this condition is met, and our system has just one positive focal point value.

With this condition met, we can simply calculate what we can now guarantee is a positive-valued solution using the quadratic formula, ensuring we take only the right-hand solution,

$$\bar{B}(1,T) = \frac{-b + \sqrt{b^2 - 4ac}}{2a} \tag{2.11}$$

by disregarding the negative solution as inadmissible. Now we have $\bar{B}(1,T)$, we can calculate the remaining focal point values for $\bar{P}(1,T)$ and $\bar{D}(1,T)$ as before (equation 2.12).

$$\begin{aligned} \bar{P}(1,T) &= P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1 + \gamma_{\text{B}} \bar{B}(1,T)\right)(\delta_{\text{P}} + \alpha_I R_{\text{on}})} \\ \bar{D}(1,T) &= \frac{\kappa_D R_{\text{on}}}{\delta_D} \end{aligned} \tag{2.12}$$

Armed with these expressions, we can calculate the equation of the boundary-equilibrium bifurcation lines in the κ_P, α_P plane, reproducing the bifurcation diagrams found in the previous analysis of this model.

2.3.2. Inferring switch transition dynamics from focal point values

The state of each switch is determined by the value of one of the continuous-valued state variables of the system. For example, the RK-switch state is determined by the dynamics of $P(t)$, as defined by the solutions to equation 2.1, while the T-switch state is determined by the dynamics of $D(t)$.

We can work out which state (off or on) a switch must be in by comparing the current value of the corresponding state variable to the threshold values of that switch. For example, for the RK-switch, if the value $P(t) > P^+$, we know the switch must be in the “on” state, while conversely if the value $P(t) < P^-$, we know the switch must be “off”. However, if the current value lies between the two thresholds, $P^- < P(t) < P^+$, the state of the switch is ambiguous unless we know something about the dynamics of the state variable leading up to the current time t. This situation is illustrated in Figure 2.4.

Consider an arbitrary time t where the value $P(t)$ is marked in Figure 2.4 by a red circle. Notice in our example $P^- < P(t) < P^+$, so the state of the switch is ambiguous. We can use the history of the dynamics of $P(t)$ shown in Figure 2.4 to work out the current state of the RK-switch. If we can identify a time $t^* < t$ where the dynamic satisfied either the switching off or on conditions, when we can work out if the switch is off or on.

In Figure 2.4 (a) we know the switch must be in the "off" state, since the previous dynamics crossed the switching off threshold P^- at time t^* (red star) while decreasing in value and the dynamic has not increased above the switching on threshold since that time t^*, we know the switch must still be off. (b) we know that the switch must be on because the dynamic crossed the switching on threshold P^+ at time t^* (red star) while increasing in value and the dynamic has not decreased below P^- since that time.

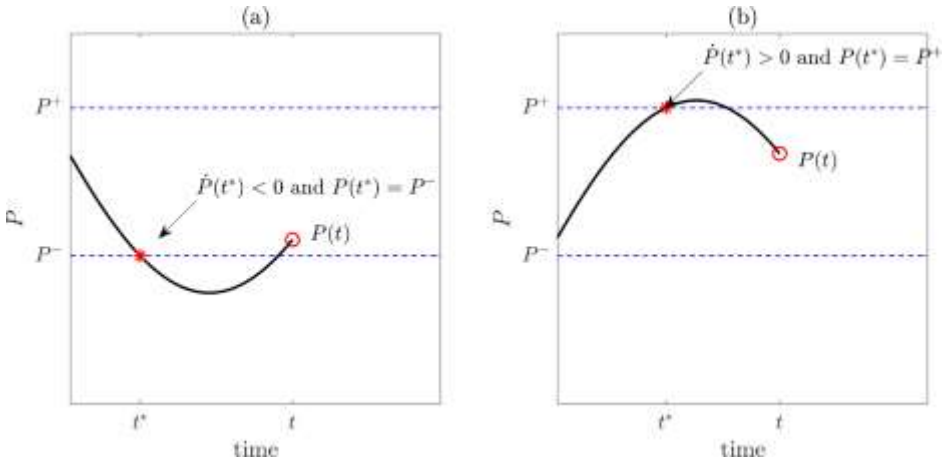


Figure 2.4 Demonstration of the history-dependence of the switch state.

This knowledge of the state of the switch is summarised in Table 2.3. Note that the example we give here features the state variable P, and we discuss the RK-switch, but that the concept applies identically to the state variable D and the T-switch.

Table 2.3 Summary of the inferred switch state based on knowledge of the current value of the continuous-valued state variable.

State variable value	**Inference about the state of the switch**
$P(t) < P^-$	Off
$P(t) > P^+$	On
$P^- < P(t) < P^+$	May be either off or on. We can only work out if it's off or on if we know the previous dynamics of $P(t)$ crossed one of the threshold values in the appropriate direction.

We can extend this same idea to consider how the dynamics of the ODE may cause the switches to change state in the future. When the focal point value of a state variable is arranged such that the dynamics of the continuous-valued state variables must cross a switch threshold value to reach that focal point, we can infer that the associated switch will change state at some time t^+ in the future. To make this prediction, we need to know the current state of the switch and compare the focal point values to the switch thresholds P^-, P^+, D^- and D^+.

For example, consider the situation in Figure 2.5 (a) where both the RK- and T-switch are off, but the focal point value for P is above the switching on threshold for the RK-switch $\bar{P}(0,0) > P^+$, then we know the future dynamics of the model must cause the RK-switch to turn on. The dynamics required to get from the current value of $P(t)$ to the focal point value $\bar{P}(0,0)$ must be such that at a time in the future $t^+ > t$, the switching on condition $\dot{P}(t^+) > 0$ and $P(t^+) = P^+$ must be satisfied. Conversely, in Figure 2.5 (b), the RK-switch is initially on, but the focal point value is below the switching off threshold for the RK-switch $\bar{P}(1,0) < P^-$, so we know the RK-switch will turn off at some time t^+ in the future.

In both examples, we won't know exactly when the switching time t^+ occurs unless we integrate our ODEs, but we know it must occur. The dynamic predicted by our ODEs after the switching time t^+, marked with a dashed black line in Figure 2.5, does not actually occur in our model dynamics because

the value of R and K immediately changes at the switching time t^+, causing the dynamic to instantly change direction and converge towards a different focal point value.

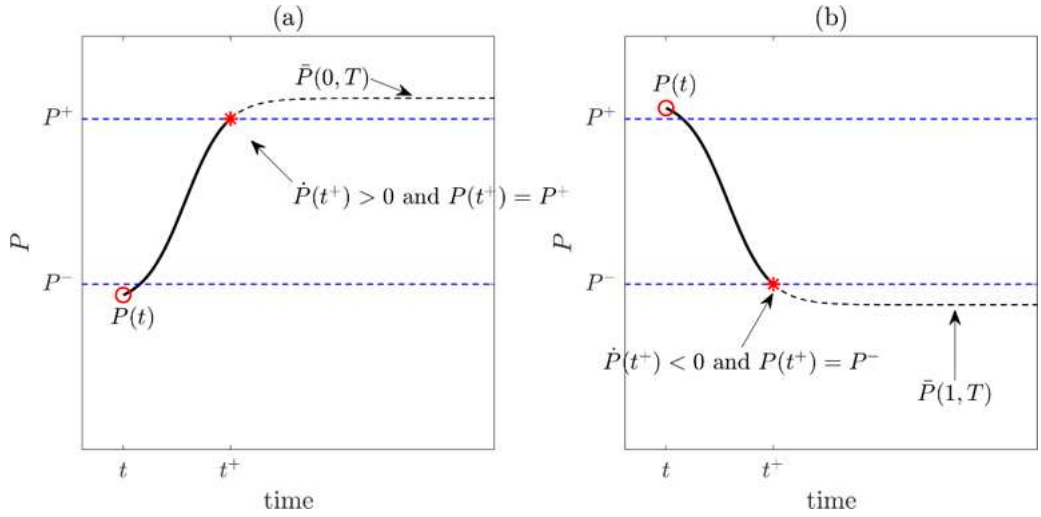


Figure 2.5 Predicting future switch transitions using the focal point value (a) predicting a switch will turn from off to on, and (b) on to off.

Both of the examples shown above are for the RK-switch, but the same logic applies to the T-switch. Table 2.4 summarises the conditions under which we can predict a transition will take place in the future for each switch.

Table 2.4 Transition conditions: conditions where we can infer the switch must change state in the future.

Switch	**Threshold values of the switch**	**Switching conditions**	
		Region where the switch turns from **off** to **on** **"Turn on condition"**	Region where the switch turns from **on** to **off** **"Turn off condition"**
RK-switch	P^-, P^+	$\bar{P}(0,T) > P^+$	$\bar{P}(1,T) < P^-$
T-switch	D^-, D^+	$\bar{D}(R,0) > D^+$	$\bar{D}(R,1) < D^-$

The switching conditions identified in Table 2.4 define the regions of the parameter space where the switches turn on or off. Shading the area of the parameter plane that satisfies each condition allows us to visualise regions of the parameter space where we expect the switch to behave in particular ways.

There are four different combinations of ways these regions can satisfy the switching off and on conditions for each switch (Table 2.5), causing that switch to behave in one of four distinct ways:

- Remaining off, which we refer to as "recovery" since a patient who has recovered from AD will have their switch off. We abbreviate this to "R" in our figures. In this case, only the turning off condition is satisfied, so the switch either remains off or stays off.
- Bistable, which we abbreviate to "B" in our figures. In this case, neither transition condition is satisfied, so the switch remains in whichever state it currently holds.
- Oscillating, which we abbreviate to "Os" in our figures. In this case, both transition conditions are satisfied, causing the switch to transition from off to on in sequence forever.
- Remaining on, which we refer to as "damage" because the switch remaining on constitutes damage. We abbreviate this to "D" in our figures. In this case, only the turning on condition is satisfied, so the switch either remains on, or stays on.

These four behaviours are of course the "qualitative phenotypes" we discussed the double-switch model producing in section 2.1.3 above.

Table 2.5 summarises these four behaviours and shows the switching condition for the RK-switch specifically, however the same idea also applies to the T-switch. To work out the behaviour of the T-switch we simply replace the "turn on" and "turn off" conditions with those corresponding the T-switch; $\bar{D}(R,0) > D^+$ and $\bar{D}(R,1) < D^-$ respectively. Note how we place the "yet to be specified" letter notation as the argument for the other switch. For the RK-switching conditions, we simply specify T, rather than a state for the T-switch. We will address the effect (or more accurately lack of effect) of the T-switch on these switching conditions later in this chapter.

Table 2.5 Switch behaviours as a result of the turning on and turning off conditions of the RK-switch.

		Turn off condition e.g. $\bar{\boldsymbol{P}}(\mathbf{1},\boldsymbol{T}) < \boldsymbol{P}^-$	
		Satisfied	**No satisfied**
Turn on condition e.g. $\bar{\boldsymbol{P}}(\mathbf{0},\boldsymbol{T}) > \boldsymbol{P}^+$	**Satisfied**	Oscillation (Os)	Switch remains on (D)
	Not satisfied	Switch remains off (R)	Bistable (B)

The boundaries of these regions are where the focal point value of a state variable collides with, in other words equates to, the threshold value of one of the switches, and are called boundary-equilibrium bifurcations. The boundary equilibrium bifurcation curves separating these regions are thus defined by the equalities in Table 2.6.

Table 2.6 Boundary-equilibrium bifurcation lines defining the edges of the regions..

Switch	Edge of the switching off region	Edge of the switching on region
RK-switch	$\bar{P}(1,T) = P^-$	$\bar{P}(0,T) = P^+$
T-switch	$\bar{D}(R,T) = D^-$	$\bar{D}(R,0) = D^+$

When we plot the regions separated by the BE bifurcation lines defined in Table 2.6, we denote which side of the boundary line satisfies the switching condition by including a coloured arrow, red for "turn off" and blue for "turn on" in all the bifurcation diagram figures throughout the rest of this thesis. This allows us to visually denote which regions of the plane correspond to each of the four possible switch behaviours described in Table 2.5.

2.3.3. Calculating BE bifurcations

Rearranging the expressions for the BE bifurcation lines for the RK-switch, we can show that the "switching on condition" is given by

$$\bar{P}(0,T) > P^+ \quad \Rightarrow \quad P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1+\gamma_{\text{B}}\bar{B}(0,T)\right)(\delta_P + \alpha_I R_{\text{off}})} > P^+ \tag{2.13}$$

If we set $R_{\text{off}} > 0$, we can easily rearrange this to give us the inequality in the α_I, κ_P plane as

$$\alpha_I > \frac{P_{\text{env}}}{P^+\left(1+\gamma_B\bar{B}(0,T)\right)R_{\text{off}}} k_p - \frac{\delta_p}{R_{\text{off}}} \tag{2.14}$$

The parameter R_{off} is usually set to zero, which means the equation for the BE bifurcation line is undefined since $R_{\text{off}} = 0$ appears in the denominator. Looking at equation (2.14), we can see that the R_{off} value multiples with α_I (our y-axis parameter). Thus, when $R = R_{\text{off}}$, we can expect the BE bifurcation line to have no dependence on the α_I parameter and form a vertical line in the (κ_P, α_I) plane. We can work out the κ_P value of the edge of the switching on condition by returning to equation (2.13) again and setting $R_{\text{off}} = 0$ in

$$P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1+\gamma_{\text{B}}\bar{B}(0,T)\right)\delta_P} > P^+ \tag{2.15}$$

Rearranging this, we can write an inequality in terms of κ_P as

$$\kappa_P > \frac{P^+\left(1+\gamma_{\text{B}}\bar{B}(0,T)\right)}{P_{\text{env}}} \tag{2.16}$$

The switching off condition follows similar derivation from $\bar{P}(1,T) < P^-$ to give

$$\alpha_I < \frac{P_{\text{env}}}{P^-\left(1+\gamma_B\bar{B}(1,T)\right)R_{\text{on}}} k_p - \frac{\delta_p}{R_{\text{on}}} \tag{2.17}$$

We always set $R_{\text{on}} > 0$, so we know this BE bifurcation line won't be vertical and we don't need to worry about the divide by zero problem.

In summary, we may write the equations defining the edges of these regions as

Edge of switching off region:
$$\alpha_I = \frac{P_{\text{env}}}{P^-\left(1+\gamma_B\bar{B}(1,T)\right)R_{\text{on}}} k_p - \frac{\delta_p}{R_{\text{on}}} \tag{2.18}$$

Edge of switching on region:
$$\kappa_P = \frac{P^+\left(1+\gamma_{\text{B}}\bar{B}(0,T)\right)}{P_{\text{env}}} \tag{2.19}$$

Both of these equations are functions of T, that is they depend on the state of the T-switch. In principle, the gradient of 2.18 should change when we change the state of the T-switch, and the location of 2.19 on the κ_P axis should also change.

Plotting these two regions on the κ_P, α_I plane, we obtain the bifurcation diagram showing the locations which induce each of the four dynamic phenotypes. As expected, the "turn on" condition results in a vertical line (drawn in blue in Figure 2.6). Note the blue arrows point to the right to denote that the region to the right of the blue line satisfies the "turn on" condition. Similarly, the "turn off" condition, and the corresponding region satisfying that inequality are denoted in red with red arrows in Figure 2.6.

We can choose to evaluate the BE bifurcation lines with either the T-switch off ($T = T_{\text{off}}$) or T-switch on ($T = T_{\text{on}}$). Applying the parameter values originally shown for the model (Table 2.1), we find that the bifurcation diagrams look the same with either the T-switch off and T-switch on. As we mentioned in section 2.2, this presents a problem for us because we expect to see Th2 activation making the disease phenotype become more severe.

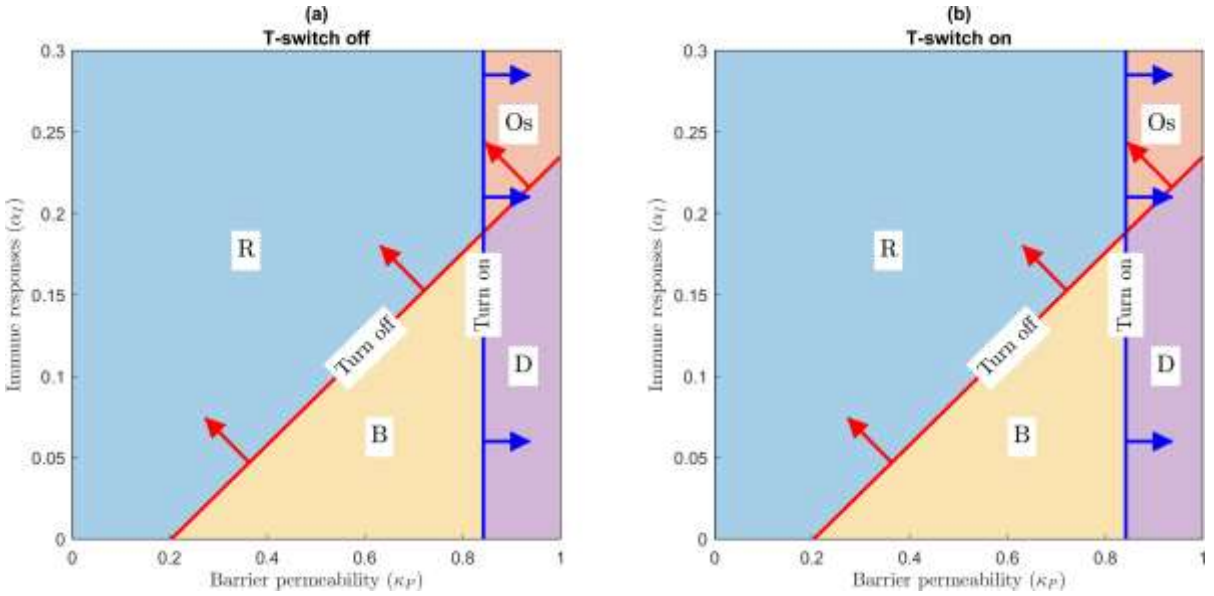


Figure 2.6 Bifurcation diagram for the double-switch model of AD
The two BE bifurcation lines represent the "turn off" condition (red) and "turn on" condition (blue). These BE bifurcation lines separate the four regions R (recovery), B (bistable), Os (oscillation) and (D) damage. The arrows emerging from each BE bifurcation line show which side of the line satisfies the switching condition.

Why do the bifurcation diagrams look the same when the T-switch is off versus on? To answer this question, we can look at the focal point value for $\bar{B}(0,T)$ (the focal point value of B when the RK-switch is off and the T-switch is in state T),

$$\bar{B}(0,T) = \frac{\kappa_B}{\delta_B + \delta_B K_{\text{off}}(1+\gamma_R R_{\text{off}})(1+\gamma_T T)}$$

We can note first of all that the expression contains a term with T in it (representing the current state of the T-switch), but that this term is multiplied by $K_{\text{off}} = 0$, as specified in Table 2.1. This means that the value of this focal point may be simplified to just κ_B/δ_B, and so it's clear that it is unaffected by the state of the T-switch, and we can write

$$\bar{B}(0,0) = \bar{B}(0,1) = \frac{\kappa_B}{\delta_B} \tag{2.20}$$

This equality of the barrier focal point value with either T-switch off or on when the RK-switch is off means that the edge of the switching on region (equation 2.19) must remain unmoved by the transition of the T-witch. Thus, the blue "switching on" BE bifurcation lines in Figure 2.6 (a) and (b) are in identical positions on the axes.

Now, to examine the red "switching off" BE bifurcation line, we must consider the focal point values when the RK-switch is on. Under these conditions the focal point value for the barrier integrity B is a little more complicated, being specified as the solution to a quadratic equation (2.7), whose coefficients are dependent on the state of the T-switch through one of the "lumped terms" we've defined, $X = (1+\gamma_R\, R_{\text{on}})(1+\gamma_T T)$ (equations 2.8).

In this term X, the dependence on T acts through a term $(1+\gamma_T T)$, where we have specified the parameter value $\gamma_T = 1$ (see Table 2.1). Thus, turning the T-switch from $T_{\text{off}} = 1$ to $T_{\text{on}} = 1$ results in a doubling of the value of X, but the effect of this is not very strong due to the remaining parameter values in the expressions for a and b. Thus, we can conclude that turning on the T-switch does indeed change the value of $\bar{B}(1,T)$, but not enough to be noticed on the bifurcation diagrams for the parameter values found for the double-switch model.

Taken together, these results mean that the BE bifurcation lines remain unmoved by changes to the T-switch state, so the bifurcation diagram in Figure 2.6 (a) where the T-switch is off looks identical to Figure 2.6 (b) where the T-switch is on. This has important implications for the double-switch's ability to reproduce chronic AD phenotypes. We summarise these shortcomings in the next section, 0.

2.4. Conclusion

In this Chapter, we have presented the double-switch model, originally developed by Domínguez-Hüttinger et al. [18]. We have drawn the model in the form of a "transition network" diagram in Figure 2.3, which provides us with a good foundation for the methods we develop in the remainder of this thesis. We'll come back to this transition network idea in Chapter 6.

We introduced the idea of the four qualitative phenotypes, which reflected the dynamics of the early stages of AD, with relapsing-remitting disease course, but were insufficient to represent the more variable dynamics of chronic AD due to some limitations built into the double-switch model. We also showed how we can calculate the parameter regions where these qualitative phenotypes emerge.

In the coming chapters, we take apart the double-switch model, and reexamine its founding assumptions. We make modifications with the aim of addressing the limitations of the double-switch model. These are:

1. The double-switch model assumes that the T-switch was irreversible. Once it turns on, it can never turn off. This means that, once allergically sensitised, it is theoretically impossible for an AD patient to recover. In Chapter 3 we revisit the founding assumptions leading to this conclusion, and introduce a reversible form of the T-switch, where these paths become possible for patients with chronic AD.
2. The double-switch model could not produce a transition network in which sequential transitions occur. This limitation is imposed by the way the T-switch state has no impact on the bifurcation diagrams. In Chapter 4 and Chapter 5, we make new assumptions about the formation of the skin barrier and incorporate these into the ODE of a new model. We show that these modifications allow us to overcome the problems we identified in the double-switch model, namely the lack of sequential switching in the transition network.

We also look at the model which gives us the "K" part of the RK-switch (i.e. the part of the model that considers active kallikrein enzymes in the epidermis), and identify an assumption we wish to modify. This process is detailed in Appendix A. This modification is less impactful on the rest of the hybrid dynamical system, but it is nevertheless important in that it improves our confidence in our founding assumptions.

Chapter 3. Modelling the effect of type-2 T-helper cells in chronic AD

Understanding the Th2 response is critical to describing the long-term dynamics of AD, since it is known that interfering with Th2 cytokine signalling using immunotherapy treatment (e.g. dupilumab) can have a significant impact on chronic AD symptoms [1]. For this reason, we wish to understand how this Th2 cytokine signalling is affected by antigen stimulation within the epidermis.

In this chapter, we examine the type-2 immune response, forming a model of the regulation of Th2 cell population, which release Th2 cytokines. While the previous modelling approach to this problem (adopted by Hutinger et al. for the double-switch model [2]) focussed on the genetic regulation within individual Th2 cells, we are looking to understand how Th2 cytokines (the target of dupilumab therapy) are affected by the varying stimulation of dendritic cells (DCs) by antigens in the epidermis. To obtain this picture, we carry out a literature search, identifying existing models of T-helper cell regulation which we then take apart and simplify to elucidate what it tells us about Th2 population dynamics specifically.

3.1. Background

The release of alarmins in AD lesional skin triggers the activation of dendritic cells (DCs). Once these DCs find antigens, they migrate to the draining lymph-node and present the antigens to naïve T-cells. Some of these T-cells will go on to differentiate into T-helper cells, whose role is to orchestrate adaptive immune responses to that antigen.

Stimulation of naïve T-cells by DCs in the presence of IL-4 cytokine gives rise to the production of GATA3 transcription factor [28], which drives the cells to differentiate into T-helper 2 cells (Th2). In AD skin, the DC stimulation generally gives rise to this type-2 response [29], whereupon the Th2 cells migrate back to the lesional site and release inflammatory cytokines IL4 and IL13 [30]. Once naïve T-cells are activated and differentiate into Th2 cells, they live for around one month in humans [31].

The Th2 cytokines also trigger production of IgE antibody by B-cells [28]. In fact, the “atopic” part of AD refers to the tendency to produce IgE antibodies in response to harmless pathogens [32]. In IgE-sensitised patients, subsequent exposure the allergens can contribute to emergence of AD flares. Type-2 immune responses generally target extracellular pathogens such as helminths, and stimulate tissue repair but also contributes to allergic diseases including AD [28].

The double-switch model considers the effect of T-helper type 2 (Th2) cells using a model of GATA3 transcription developed by Höfer et al. [26] who modelled the molecular kinetics of expression of GATA3 transcription factor in naïve T-cells in response to IL-4 stimulation by DCs, i.e. how much Th2 cytokines a single T-cell produces in response to IL-4 stimulation by DCs. The model demonstrated that GATA3 transcription, once transcribed at a certain minimum rate, is able to maintain its own expression through a positive feedback loop. This consigns the naïve T-helper cell to differentiate irreversibly into a Th2 cell and produce the Th2 cytokines (IL4 and IL13). For this reason, the double-switch model features an irreversible T-switch[2]. However, we need a model of the population dynamics of Th2 cells to describe how the concentration of the Th2 cytokines they produce in the lesional skin is affected by the stimulation of DCs.

3.2. Modelling Th2 population dynamics

We are looking to identify the relationship between DC stimulation of naïve T-cells in the draining lymph-node, and the expression of Th2 cytokines in the AD lesional tissue.

[2] Note that throughout the publication describing the double-switch model, the switch summarising Th2 activation is referred to as the “G-switch” after the GATA3 gene. In this thesis, we name it the “T-switch” to ensure the naming convention is consistent for all chapters.

3.2.1. Model assumptions on timescale and bistability

We only look at the steady states of Th2 cytokine levels (which is assumed to be proportional to the number of Th2 cells), by assuming that the number of Th2 cells is at a quasi-steady state with respect to the number of DCs. This assumption holds if the turnover rate of DCs is much slower than that of Th2 cells. However, mouse data indicated the opposite, i.e. the turnover rate of DCs is about twice faster than that of Th2 cells [24]. Hence, we cannot assume that the T-helper cell number will remain at the steady state as the DC number changes. However, the number of Th2 cells will reach a steady state if the number of DCs is stable for more than a month. We therefore implicitly assume here that the number of DCs is stable over this month timescale, the time required for the number of Th2 cells to reach a steady state.

Naïve T-helper cells are stimulated by DCs and differentiate into either Th1 or Th2 effector cells, depending on the nature of co-stimulatory signals and the cytokine environment. Excitatory and inhibitory effects of cytokines regulate activation, proliferation, and death of the Th cells. IFNγ and IL-12 cytokine produced by Th1 cells can induce activation of naïve T-cells into new Th1 cells, and promote their growth [33] while IL-4 cytokine released by Th2 cells can do the same for the Th2 cell population [34]. These dual positive feedback mechanisms and mutual inhibition of the Th1 and Th2 effectors suggest that the Th cell numbers will exhibit bistable behaviours [35], [36].

Another potential mechanism for bistability is positive feedback from the existing population of T-cells. T-helper cells produce chemokines which attract T-helper cells from elsewhere in the body. For example, TARC/CCL17, one of the major chemokines attracting Th2 [37], is released from keratinocytes, which is induced by IFNγ (Th1 cytokine) and TNFα (produced by mainly macrophages, but also other immune cells) [38]. Release of TARC from T cells is induced by the Th2 cytokine IL-4, while IL-4 inhibits release of TARC from keratinocyte [38], [39]. Other chemokines attracting Th2 include PARC/CCL18 [40], whose release from monocyte/macrophage is induced by IL-4 and IL-13 (Th2 cytokines) [41]. Hence, inflammation with upregulating Th1 and/or Th2 cytokines would further strengthen the activation of naïve T-cells and attraction of existing T-helper cells from other tissues, reinforcing the bistability in T-helper cell number.

3.2.2. The Yates model

We searched the literature and identified several models of the dynamics of Th2 cell number [42]–[44], and selected Yates et al. [44] for further consideration for our purpose, because it was biologically realistic yet relatively simple among these three models. This Yates model forms the basis of our work for the remainder of this chapter.

The full Yates model is complex, featuring inputs from DCs and also costimulatory signals from Th1 and Th2 cytokines, which bias the naïve T-cells to differentiate into one cell type or the other. They examine the two most dominant phenotypes of T-helper cells, i.e., Th1 and Th2 cells. Th1 cells are required for effective immunity against bacterial, virus and protozoa infections, whereas Th2 cells promote antibody production, defence against parasitic infections such as helminths and play a key role in the development of atopic sensitisation.

Naïve T-helper cells are stimulated by DCs and differentiate into either Th1 or Th2 effector cells, depending on the nature of co-stimulatory signals and the cytokine environment. Excitatory and inhibitory effects of cytokines regulate the activation, proliferation, and death of the Th cells (Figure 3.1).

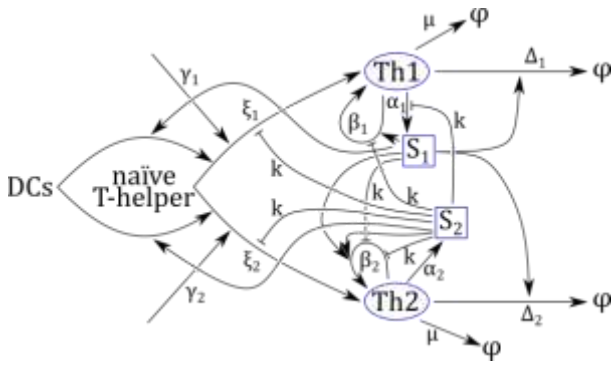


Figure 3.1 Yates model schematic. The cytokine signals from Th1 and Th2 cells govern T-helper cell activation, proliferation, and death.

The dynamics of the numbers of Th1 (T_1) and Th2 (T_2) cells are described in Yates *et al.* [44] by four terms representing activation, proliferation, antigen-induced cell death (AICD) and natural cell death, in order:

$$\begin{aligned}\frac{\mathrm{d}T_1}{\mathrm{d}t} &= \frac{(\gamma_1+S_1)}{1+kS_2}\xi_1 + \frac{\beta_1 S_1}{1+kS_2}T_1 - \Delta_1 S_1 T_1^2 - \mu T_1,\\ \frac{\mathrm{d}T_2}{\mathrm{d}t} &= \frac{(\gamma_2+S_2)}{1+kS_2}\xi_2 + \frac{\left(\beta_2^{IL4}S_2+\beta_2^{IL2/4}S_1S_2\right)}{(1+kS_1)(1+kS_2)}T_2 - \Delta_2 S_1 T_1 T_2 - \mu T_2,\end{aligned} \tag{3.1}$$

S_1 and S_2 represent Th1 cytokines (including IL-2, IL-12 and $IFN\gamma$) produced by Th1 cells and Th2 cytokines (including IL-10, IL-4 and $TGF\beta$) produced by Th2 cells, respectively. The cytokine concentrations are described by $S_1 = \frac{\alpha_1 T_1}{1+kS_2}$ and $S_2 = \alpha_2 T_2$, respectively, which correspond to quasi-steady state solutions of the kinetic equations for cytokines, $\dot{S_1} = \frac{\alpha_1}{1+kS_2}T_1 - S_1$ and $\dot{S_2} = \alpha_2 T_2 - S_2$, where α_1 and α_2 are the rate of cytokine production by Th1 and Th2 cells, respectively and k is the rate of inhibition of cytokine production in Th1 cells by Th2 cytokines. The Th2 cytokines inhibit the production of the Th1 cytokines, but the Th1 cytokines do not inhibit Th2 cytokine production. We use the cytokine steady-states in the dynamics of T-helper cells in equation 3.1 because the timescale of cytokine production and degradation is far shorter than the timescale of T-helper cell production and death rate.

Activation of Th1 and Th2 cells occurs as a result of successful encounters between naïve T-cells and DCs in the lymph nodes, with the strength of activation towards either Th1 or Th2 being described by ξ_1 and ξ_2, respectively. Th1 activation is enhanced by IL-12 produced either by Th1 cells (S_1) or by other sources (γ_1). Th2 activation is enhanced by IL-4 produced either by Th2 cells (S_2) or by other sources (γ_2). Activation of both Th1 and Th2 cells is inhibited by IL-10 cytokine produced by Th2 cells (S_2), since it downregulates APC activity [45] and MHC class II expression [46] that are required for successful activation. The inhibitory effects of Th1 and Th2 cytokines on Th cell activation and proliferation are assumed to be the same and are represented phenomenologically by denominators of the form $(1+kS_i)$. Note that we use the same parameter k for the inhibition of S_1 production by S_2 ($\dot{S_1} = \frac{\alpha_1}{1+kS_2}T_1 - S_1$) for simplicity.

Once activated, both Th1 and Th2 cells proliferate. The proliferation of Th1 cells is enhanced by IL-2 (S_1, with the proliferation rate β_1) and inhibited by $TGF\beta$ from Th2 cells (S_2). The proliferation of Th2 cells is enhanced by IL-4 both alone ($\beta_2^{IL4}S_2$) and in strict combination with IL-2 ($\beta_2^{IL2/4}S_1S_2$), and is inhibited by $IFN\gamma$ (S_1) from Th1 cells and by $TGF\beta$ (S_2) from Th2 cells. The model assumes that the external sources of Th1 and Th2 cytokines (γ_1 and γ_2) only affect the activation of Th-cells, but not proliferation and all the intercellular signalling after the activation of Th-cells.

Th1 and Th2 cells disappear by two cell death mechanisms, namely, antigen-induced cell death (AICD) and constant per capita removal including peripheral diffusion, necrosis and other apoptosis pathways. AICD is brought about by FasL signalling. As FasL receptors are expressed only by Th1 cells, Th1 AICD occurs by Th1-Th1 interactions (T_1^2) and Th2 AICD occurs as a result of Th1-Th2 interactions (T_1T_2). IL-2 (S_1) increases the transcription rate and expression levels of FasL receptors which are required for AICD [47]–[49]. The model assumes the same death rate, μ, for Th1 and Th2 cells for simplicity. When their population isn't replenished by production of fresh T-helper cells, the population number depletes to around half after 8.4 days [31], so we assume that T-cell numbers change on a similar timescale in the models presented here.

We assume that the concentration of Th2 cytokines (T in the triple-switch model) is proportional to the number of Th2 cells (T_2 in the Yates model above), i.e. each Th2 cell produces the same amount of Th2 cytokines. We also assume that the number of available naïve T-cells is unchanged over time and that the activation rates of naïve T-cells (both ξ_1 and ξ_2 [$cells\ (nM)^{-1}\ hour^{-1}$]) are proportional to the number of activated DCs in lymph nodes ($D(t)$ in the triple-switch model),

$$c_1\xi_1(t) = c_2\xi_2(t) = D(t) \tag{3.2}$$

where c_1 and c_2 are arbitrary constants.

The relative sizes of $\xi_1(t)$ and $\xi_2(t)$ represent the probability that a naïve T-cell will differentiate into a Th1 or Th2 cell respectively after stimulation by a DC. The ratio of these probabilities, $\frac{\xi_1(t)}{\xi_2(t)} = \frac{c_1}{c_2}$, remains constant so that when a naïve T-cell is stimulated by a DC, the probability of it becoming a Th1 or a Th2 cell remains constant over time.

Untangling the complexity of the Yates model

The Yates model provides a complicated picture of the T-helper cell regulation. We propose some simplifications which might help in our specific application.

Simplification 1. We are interested in obtaining practical insight about AD in the absence of other potentially distracting immune responses. We therefore assume that there are no additional sources of Th1 or Th2 cytokines ($\gamma_1 = \gamma_2 = 0$). We consider the dynamics of the Th2 response in the absence of other significant simultaneous immune responses.

Simplification 2. We set the rate of Fas-induced cell death for Th2 cells to zero ($\Delta_2 = 0$), i.e. Th1 cells preferentially undergo activation-induced cell death through the Fas pathway, but not Th2 cells. This is based on empirical observation [50], and allows us to remove a term from the original ODEs.

We consider two hypothetical mechanisms that could give rise to the bistability with hysteresis which we expect to see. Firstly, feedback of IL4 from the skin to the lymph-nodes (covered in section 3.2.3) and secondly Th1/Th2 mutual inhibition combined with IL4 feedback (covered in section 3.2.4).

We develop models corresponding to each hypothesis by further simplifying the Yates model in slightly different ways, and conduct bifurcation analyses on both. Our analysis demonstrates that the model with systemic IL4 feedback does not demonstrate bistability unless dimerization in TGFβ signalling of order two or higher is assumed. The model with both systemic IL4 feedback and Th1/Th2 mutual inhibition also demonstrates bistability with hysteresis without such additional assumption.

3.2.3. Model 1: Systemic IL4/TGFβ feedback

Model description

Our first hypothesis was that the bistability of Th2 response results from systemic positive feedback via IL4 ($S_2 = \alpha_2 T_2$) from the lesional tissue to the lymph-node (where fresh Th2 cells are being

produced) without the involvement of Th1 activation. IL4 produced by Th2 cells in the skin is detected by DCs, which produce their own IL4 [51] , which in turn promotes differentiation of newly activated T-helper cells into Th2 cells [52]. Therefore, we can assume that the activation rate of naïve T-cells into Th2 cells is increased by IL4 produced by already activated Th2 cells.

This set of assumptions is represented by the model schematic (Figure 3.2).

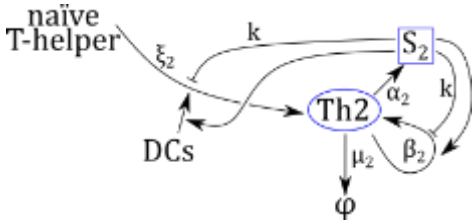


Figure 3.2 A schematic of the IL4 feedback model

The dynamics of the number of Th2 cells is described by

$$\frac{dT_2}{dt} = \frac{S_2}{1+kS_2}\xi_2 + \beta_2^{IL4}\frac{S_2}{1+kS_2}T_2 - \mu T_2 \tag{3.3}$$

by removing Th1-derived cytokines (S_1) in the T_2 ODE in (1). Th2 cells are activated at rate ξ_2, which is enhanced by the Th2 cytokine IL4 (hence the factor S_2 in the numerator) and inhibited by the Th2 cytokine IL-10 (hence the $1+kS_2$ in the denominator of the first term). Once activated, Th2 cells proliferate at rate β_2^{IL4}. Th2 proliferation is enhanced by IL4 signalling (hence the factor S_2 in the numerator) and inhibited by $TGF\beta$ signalling (hence the $1+kS_2$ in the denominator of the second term).

Substituting the scaling factors (obtained by adapting the scaling factors used in [44]) $\mathrm{T} = \mathrm{k}\alpha_2\mathrm{T}_2 = \mathrm{k}\mathrm{S}_2$, $\tau = \mu t$, $\mathrm{D} = \xi_2\alpha_2/\mu$, and $\pi_2 = \beta_2^{\mathrm{IL4}}/(\mu\mathrm{k})$, we obtain

$$\frac{d\mathrm{T}}{d\tau}\frac{\mu}{k\alpha_2} = \frac{\mu}{k\alpha_2}D\frac{\mathrm{T}}{1+\mathrm{T}} + \frac{\mu}{k\alpha_2}\pi_2\frac{T^2}{1+T} - \frac{\mu}{k\alpha_2}\mathrm{T}. \tag{3.4}$$

The factor $\frac{\mu}{k\alpha_2}$ cancels out, leaving us with the simplified nondimensional ODE,

$$\frac{dT}{d\tau} = T\left(\frac{D}{1+T} + \pi_2\frac{T}{1+T} - 1\right) \tag{3.5}$$

Now that we have a relatively simple expression for the ODE, we are in a good position to analyse the focal point values of this model as we vary the dendritic cell stimulation D. We show that bistable behaviour is not possible for the ODE in equation 3.5 and then propose a modification which is able to restore a bistable behaviour.

Focal point analysis and bistability

The steady states of equation 3.5 are at $T = 0$ and $T = \frac{1-D}{k-1}$ $(\pi_1 \neq 1)$, the two solutions of the equation for the steady state values $0 = T\left(\frac{\mathrm{D}}{1+T} + \pi_2\frac{\mathrm{T}}{1+T} - 1\right)$.

The resulting bifurcation diagram for T with respect to the parameter D does not demonstrate a bistable hysteresis (Figure 3.3). This result can be seen because there are only two possible steady states, and it is not possible for both to be stable simultaneously. The manifolds defining these two steady states are straight lines, which meet at a pitchfork bifurcation. This is discussed in more detail and proved analytically in Appendix C. Figure 3.3 was produced using numerical continuation methods in MatCont.

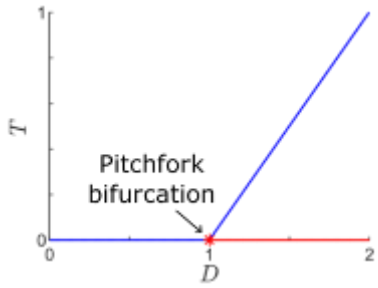


Figure 3.3 Bifurcation diagram for the model of equation (3.5) with $\pi_2 = 0.2$. The blue lines represent positive stable steady states, while the red lines represent positive unstable steady states. The system only has a pitchfork bifurcation at D=1.

A bistable hysteretic switch may be observed if we slightly modify the original model to the one described by equation 3.6 (Figure 3.4). We modify the second term to have a power 2 exponents in the denominator due to an assumed dimerization of the TGFβ signalling molecule (Figure 3.5). Figure 3.4 was produced using numerical continuation in MatCont. In this bifurcation diagram, the threshold values D^- and D^+ of the bistable region are defined by a transcritical and pitchfork bifurcation, respectively. The power-2 exponent in the denominator due to dimerised TGFβ signalling causes sufficient nonlinearity in the system to create the transitional bifurcation."

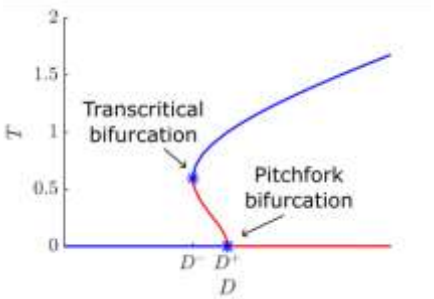


Figure 3.4 Hysteretic bistability produced by assuming order 2 inhibition with $\pi_2 = 1.5$. Blue lines represent stable steady states, while red lines are unstable steady states. The blue asterisks are limit points where a branch of steady states loses its bistability. The Th2 bistable region is between D^- and D^+. π_2 represents the rate of increase of proliferation of Th2 cells in response to TGFβ stimulation. The threshold values D^- and D^+ are caused by a transcritical and a pitchfork bifurcation, respectively.

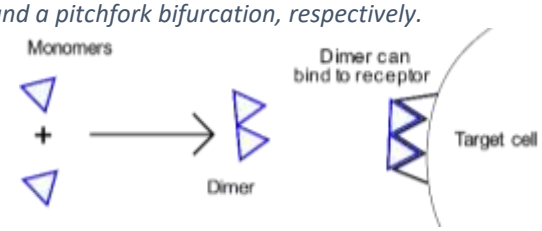


Figure 3.5 Proposed assumption of dimerization of a signalling molecule TGFβ which promotes Th2 proliferation.

$$\frac{dT_2}{dt} = \frac{S_2}{1 + \mathrm{k}S_2}\xi_2 + \beta_2^{IL4}\frac{S_2}{1 + \mathrm{k}^2 S_2^2}T_2 - \mu T_2 \tag{3.6}$$

Introducing scaling factors in the same way as above, but this time setting a new factor $\pi_2 = \beta_2^{\mathrm{IL4}}/(\mu \mathrm{k}^2)$, we can write the simplified ODEs 3.7.

$$\frac{dT}{d\tau} = T\left(D\frac{1}{1+T} + \pi_2\frac{T}{1+T^2} - 1\right) \tag{3.7}$$

with the second term in the form of $\frac{T^2}{1+T^2}$. In this term, the numerator T^2 results from the dependence of Th2 proliferation on T_2 (rate of growth of the Th2 population increases in proportion to the size of

the population) and on IL4 cytokine S_2 (the IL4 cytokine acts as a growth factor). The T^2 in the numerator is therefore independent of the TGFβ signalling.

The $1 + T^2$ in the denominator qualitatively models the effect of dimerization of TGFβ or downstream signalling molecules. TGFβ receptors and some molecules downstream in the signalling pathway (such as SMADs) dimerise to activate target genes of the TGFβ signalling pathway [53]–[55]. It is therefore plausible that such effects could cause the inhibition of Th2 proliferation by TGFβ to take the form $\frac{1}{1+T^2}$, where the inhibitory effect of T is weaker for $T \ll 1$. However, this model does not reflect the dominant view in the literature that Th1/Th2 mutual inhibition is important in the pathophysiology of AD, instead presenting an alternative mechanism of TGFβ dimerization as the cause of bistable behaviours. It is interesting to note that the known qualitative behaviours of the system can be reproduced with a simpler and less complex model through TGFβ cooperative inhibition, but that the dominant view in the experimental literature is on the more complex hypothesis of Th1/Th2 mutual inhibition.

Parameter sensitivity: Effect of rate of Th2 proliferation π_2 on bistability

To assess the robustness of the bistable behaviours produced by the model described by equation (3.7) and compare this to that of the more complex Th1/Th2 mutual inhibition model described in section 3.2.4 below, we evaluated how the bistability we observe in Figure 3.6 changes as we vary the only free parameter, π_2. An exchange of places between the pitchfork bifurcation and another transcritical bifurcation gives rise to four behaviours characterised by different placement of the threshold values D^+ and D^- (Figure 3.6 (a-d)).

For values of $\pi_2 < 0.78$, there is only one transcritical bifurcation at $D = 1$ marking the value of D when the steady state T can become greater than zero. In this case, the bifurcation diagram shows a monostable nature (Figure 3.6 a). For values of $0.78 < \pi_2 < 1$, a transcritical bifurcation moves from negative values of T to take values of $T > 0$, replacing the transcritical bifurcation as the source of the threshold value D^+. For this reason, the D^+ value increases slightly above 1 in this parameter range (Figure 3.6 b). This slight increase in the value of D+ isn't particularly important to us in terms of our interpretation of this model, but it is an interesting complication to arise from an otherwise relatively simple ODE model. For values $1 < \pi_2 < 2$, the bifurcation diagram features threshold values D^- and D^+ defined by a transcritical bifurcation and a pitchfork bifurcation (Figure 3.6 c). For values $\pi_2 > 2$, the transcritical bifurcation defining the D^- threshold moves into the negative domain, rendering the bistability irreversible (Figure 3.6 d).

The movement of the threshold values D^+ and D^- on the D axis as we changed the value of $\pi_2 = \beta_2^{IL4}/(\mu k^2)$, is calculated by numerical continuation in MatCont and summarised in Figure 3.6 (e). To calculate this diagram, we selected the D^- "limit point" in MatCont, and then carried out a numerical continuation of the codimension-2 object, allowing the parameters D and π_2 to vary.

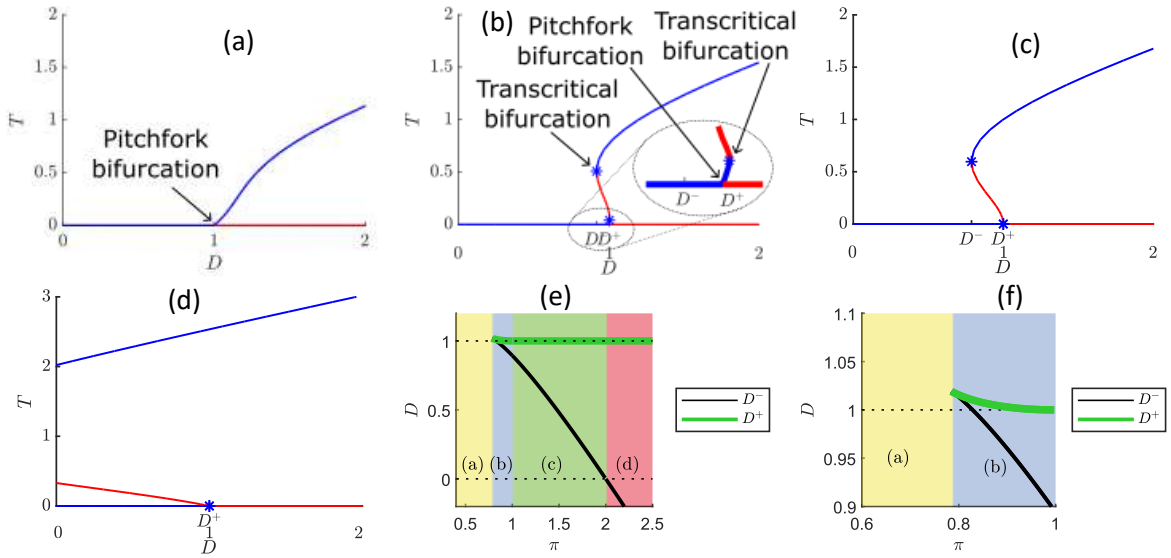


Figure 3.6 One parameter bifurcation diagrams obtained by MatCont, showing the four possible qualitative behaviours of the system. (a) monostable, (b) bistable with $D^+ > 1$, (c) bistable with $D^+ = 1$ and (d) irreversible bistability with $D^- < 0$. (e) The values of D^- (black line) and D^+ (thick green line) for different values of π_2 in the systemic IL4 feedback model. (f) Zoomed-in section showing the region where $D^+ > 1$.

The system does not show bistable behaviour (Figure 3.6 (a)) if $\pi_2 < 0.78$. The system is bistable and reversible (Figure 3.6 (b, c)) whilst $0.78 < \pi_2 < 2$, and the bistability becomes irreversible as the switching off threshold value D^- becomes negative (Figure 3.6 (d)) if $\pi_2 > 2$. For the reversible bistable behaviours, the upper threshold D^+ can either be $D^+ > 1$ (Figure 3.6 (b)) or $D^+ = 1$ (Figure 3.6 (c)).

For the bistable and reversible case (b), there is a very small non-zero Th2 response while the switch is still off (shown in the zoomed-in section of Figure 3.6 (b)), relative to the magnitude of the Th2 response when the switch is on. Since the magnitude of the non-zero Th2 response shown zoomed-in in Figure 3.6 (b) is so small relative to the stable steady state values above D^+, it seems likely that we cannot distinguish between cases (b) and (c) in our larger triple-switch model, and it is reasonable to consider them as equivalent for our purposes.

3.2.4. Model 2: Systemic IL4 feedback + Th1/Th2 mutual inhibition

Model description

As an alternative hypothetical mechanism that leads to the bistability of Th2 response, we investigated the effects of mutual inhibition between Th1 and Th2 cells on activation and growth of Th1 and Th2 cell numbers through cytokine signalling, in addition to the systemic IL4 feedback with assumed monomeric feedback by TGFβ described by equation 3.7. The resulting model (schematic in Figure 3.7) is described by

$$\begin{aligned} \frac{dT_1}{dt} &= \xi_1 \frac{S_1}{1+kS_2} + T_1 \frac{\beta_1 S_1}{1+kS_2} - \Delta_1 S_1 T_1^2 - \mu T_1, \\ \frac{dT_2}{dt} &= \xi_2 \frac{S_2}{1+kS_2} + T_2 \frac{\left(\beta_2^{IL4} S_2 + \beta_2^{IL2/4} S_1 S_2\right)}{(1+kS_1)(1+kS_2)} - \mu T_2. \end{aligned} \tag{3.8}$$

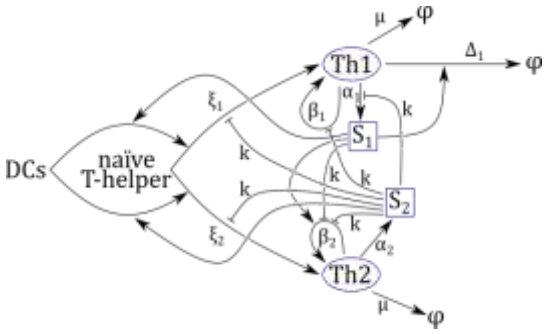


Figure 3.7 A schematic of the Th1-Th2 mutual inhibition model.

We nondimensionalised with the following scaling factors: $x_1 = k\alpha_1 T_1$, $x_2 = k\alpha_2 T_2 = kS_2$, $S_1 = \frac{\alpha_1 T_1}{1+kS_2}$, $\tau = \mu t$, $D_i = \xi_i \alpha_i/\mu$, $\delta_1 = \frac{\Delta_1}{\mu k^2 \alpha_1}$, $\pi_1 = \beta_1/(\mu k)$, $\pi_2 = \beta_2^{IL2/4}/(\mu k^2)$, $\rho = \beta_2^{IL4} k/\beta_2^{IL2/4}$ and obtained the nondimensionalised ODEs,

$$\begin{aligned} \frac{\mathrm{d}x_1}{\mathrm{d}\tau} &= D_1 \frac{x_1}{(1+x_2)^2} + \pi_1 \frac{x_1^2}{(1+x_2)^2} - \delta_1 \frac{x_1^3}{1+x_2} - x_1, \\ \frac{\mathrm{d}x_2}{\mathrm{d}\tau} &= D_2 \frac{x_2}{1+x_2} + \pi_2 \left(\rho + \frac{x_1}{1+x_2}\right) \frac{x_2^2}{1+x_1+x_2} - x_2, \end{aligned} \tag{3.9}$$

where $\rho = 0.1$, since the potency of IL4 alone as a growth factor for Th2 cells is approximately one-tenth the potency of IL4 in combination with IL-2 [56]. A derivation of this nondimensionalisation process may be found in Appendix D.

The variables, $D_1(t)$ and $D_2(t)$, representing the propensity by DCs for Th1 activation and Th2 activation respectively, are described as $D_1(t) = \frac{\alpha_1}{\mu c_1} D(t)$ and $D_2(t) = \frac{\alpha_2}{\mu c_2} D(t)$, using $c_1 \xi_1(t) = c_2 \xi_2(t) = D(t)$ described in equation 3.2 above. We further assume the ratio of rates of Th1 to Th2 activation $\frac{D_1(t)}{D_2(t)} = \frac{c_1 \alpha_1}{c_2 \alpha_2}$ to be constant. The scaling factors, and their respective values are specified in Table 3.1. The values specified in this table are chosen partly because they are of similar magnitude, and thus plausible guesses for growth and death rates, and so that the regions illustrated in Figure 3.8 are of a suitable size to be visible.

Table 3.1 Meaning and nominal values of the nondimensionalized model parameters of the Th1/Th2 mutual inhibition model.

Parameter/ state variable ($\boldsymbol{i} = \mathbf{1}, \mathbf{2}$)	Interpretation	Value
$\boldsymbol{x_i}$	Strength of the Th(i) response	
$\boldsymbol{D_i}$	Strength of activation by dendritic cells of naïve T-helper cells towards Th(i) phenotype	$\frac{D_1}{D_2} > 1$
$\boldsymbol{\pi_i}$	Efficacy of growth factors at stimulating proliferation of Th(i) cells	$\pi_1 = 1.5$ $\pi_2 = 3$
$\boldsymbol{\delta_1}$	Susceptibility of Th1 cells to Fas-induced cell death	$\delta_1 = 1$
$\boldsymbol{\rho}$	The ratio of the strength of IL4 as a growth factor for Th2 cells to the strength of IL-2 in combination with IL4	$\rho = 0.1$

Focal point analysis and bistability

We are interested in how the nondimensionalised number of Th2 cells x_2 (which is proportional to Th2 cytokine concentrations) varies as D_1 and D_2 (that are proportional to *D* in the triple-switch model) change while keeping their ratio D_1/D_2 constant. Since our AD triple switch model will consider timescales longer than the lifespan of a T-helper cell, we will investigate the steady states for the number of Th1 and Th2 cells ($\overline{x_1}$ and $\overline{x_2}$).

There are four possible cases for the focal points of the system[3], which we denote with different superscripts of $*$:

a. **Completely off** at $(\overline{x_1}, \overline{x_2}) = (0,0)$,
b. **Th1 response only** at $(\overline{x_1}, \overline{x_2}) = (x_1^*, 0)$,
c. **Th2 response only** at $(\overline{x_1}, \overline{x_2}) = (0, x_2^*)$ and
d. **Th1/Th2 "mixed response"** where both Th1 and Th2 response is nonzero, $(\overline{x_1}, \overline{x_2}) = (x_1^{**}, x_2^{**})$.

The range of the parameters (D_1, D_2) corresponding to each case was calculated by evaluating the eigenvalues of the Jacobian of the system at each of these four steady states. The regions of the (D_1, D_2) plane (**light blue – completely off**, **yellow – Th2 response only**, **orange – Th2 response only** and **pink – Th1/The** areas) corresponding to the four cases of stable and positive steady states are shown in Figure 3.8 (a – d). The model demonstrates bistability on the (D_1, D_2) plane where two stability regions overlap (Figure 3.8 (e)) and the corresponding steady states are stable simultaneously. The boundaries of these four regions in the (D_1, D_2) plane are found by finding equations describing the eigenvalues of each steady state as a function of D_1 and D_2 or using numerical continuation algorithms in MatCont. A derivation of each boundary is detailed in Appendix G.

The blue and red lines in Figure 3.8 represent a ratio of $D_1/D_2 = 2$ and $D_1/D_2 = 0.5$, respectively. Changing the value of D_2 by moving along the blue and red lines results in transitioning between cases, with the threshold values for each transition given by the intersection of the boundaries with the D_1/D_2 lines annotated on the D_2 axis and circled in blue or red respectively. The boundaries of the overlapping regions, where the system transitions from one focal point to another in a history-dependent way (i.e. a switch with hysteresis), are marked by the thick border lines shown in Figure 3.8 (a – d). The boundary values for the ratio $D_1/D_2 = 2$ are shown in blue on the D_2 axis, while the boundary values for the ratio $D_1/D_2 = 0.5$ are shown in red on the D_2 axis. For diagram Figure 3.8 (e), only the boundary values for the blue line $D_1/D_2 = 2$ are shown.

We consider the cases when $\frac{D_1}{D_2} > 1$ and $\frac{D_1}{D_2} < 1$ (corresponding to the blue and red lines on Figure 3.8 respectively), separately since they produce qualitatively different behaviours.

[3] This can be deduced by looking at the structure of the ODE. Setting both $x_1 = 0$ and $x_2 = 0$ it is immediately clear that $\frac{dx_1}{d\tau} = 0$ and $\frac{dx_2}{d\tau} = 0$, meaning the origin is a steady state $(\overline{x_1}, \overline{x_2}) = (0,0)$. Setting $x_1 = 0$, we obtain a cubic equation for x_2. On the other hand, setting $x_2 = 0$, we obtain a cubic for x_1. A cubic polynomial always has at least one real solution, so we can see that the steady state will always exist for the cases $(\overline{x_1}, \overline{x_2}) = (x_1^*, 0)$ and $(\overline{x_1}, \overline{x_2}) = (0, x_2^*)$. Finally, we have numerically confirmed that there is a fourth steady state with both $x_1 > 0$ and $x_2 > 0$, making the fourth steady state $(\overline{x_1}, \overline{x_2}) = (x_1^{**}, x_2^{**})$.

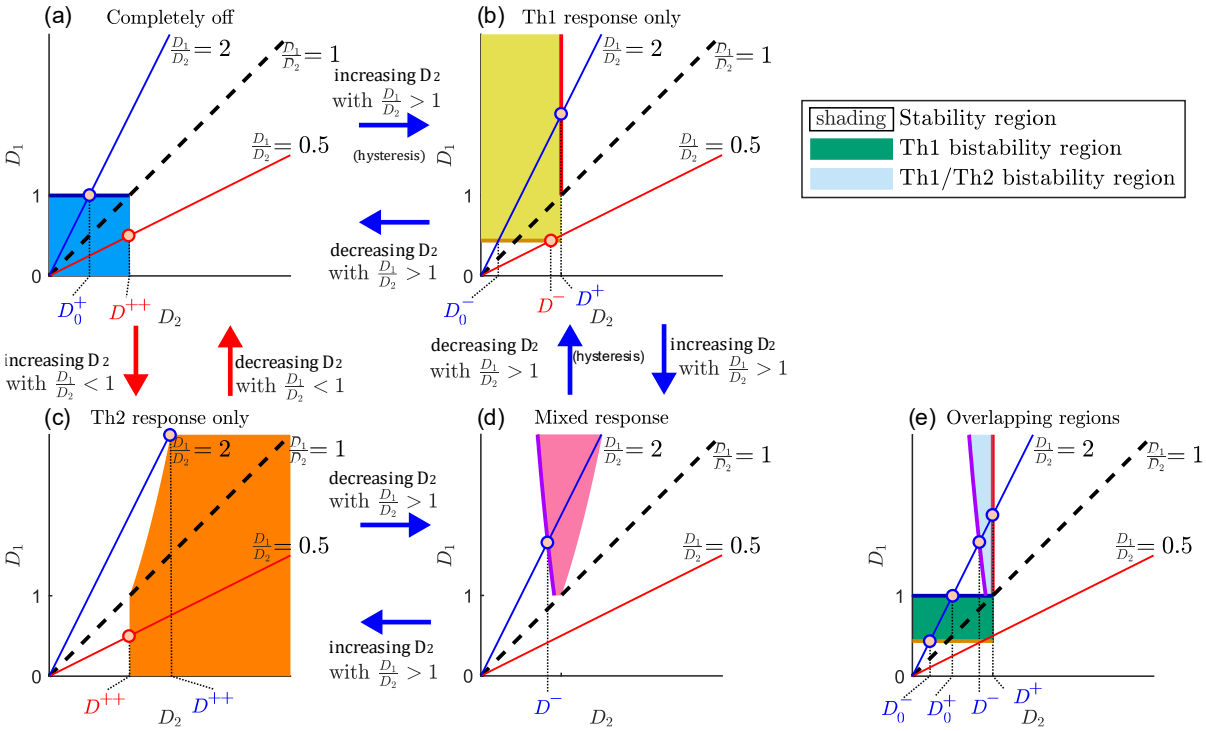


Figure 3.8 Regions of the (D_1, D_2) plane in which each of the four steady states arise (a – d). Bistable regions are shown in (e).

Case 1: $\frac{D_1}{D_2} > 1$ (blue line)

We illustrate this case by the blue line on Figure 3.8, representing the ratio $D_1/D_2 = 2$. When we change D_2 while keeping the ratio $D_1/D_2 = 2$, the focal point moves along the blue line of $D_1/D_2 = 2$ from (a) to (b) to (d) to (c) (sequence shown in Figure 3.9). Parts of the blue line cross bistable regions (Figure 3.8 (e)), where two steady states are stable simultaneously.

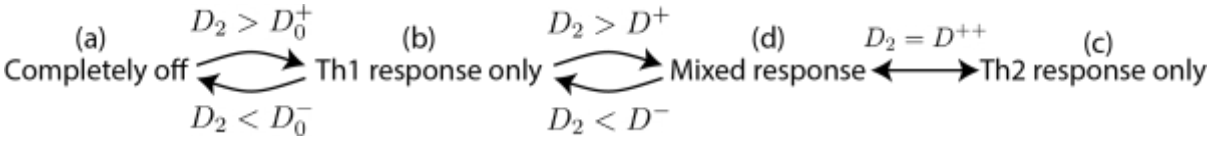


Figure 3.9 Transitions between the four cases of stable positive steady states by changing the value D_2 while keeping the ratio D_1/D_2constant and greater than 1.

The region in Figure 3.8 (c) does not overlap with the other regions meaning that Th2 only responses are always monostable (not bistable). In other words, for there to be a bistable Th2 response, this model requires an active Th1 response as well.

The one-parameter bifurcation diagrams in Figure 3.10, created by MatCont, show the Th1/Th2 mutual inhibition model's steady states of $\overline{x_1}$ and $\overline{x_2}$ as we vary D_2 with $D_1/D_2 = 2$ (moving along the blue lines in Figure 3.8 (a-e)). The x-axis in Figure 3.10 shows the value of D_2 below the blue line in Figure 3.8 (a-e) showing us the projection of that blue line onto the D_2 axis.

The curve representing the steady state value $\overline{x_1}$ or $\overline{x_2}$ is coloured blue if every eigenvalue of the Jacobian has a negative real part (meaning the focal point is stable), and coloured red if any of the eigenvalues has positive real part (meaning that the focal point is unstable). In the region between D_0^- and D_0^+, there are two stable and one unstable steady state. However, for x_2, both stable steady states, and the unstable steady state overlap at $\overline{x_2} = 0$, making it look like only one steady state is present between D_0^- and D_0^+. In the region between D^- and D^+, there are two unstable and two stable steady states, all of which can be seen for on the plot for $\overline{x_1}$, but not for $\overline{x_2}$ where one stable focal point and one unstable focal point at $\overline{x_2} = 0$.

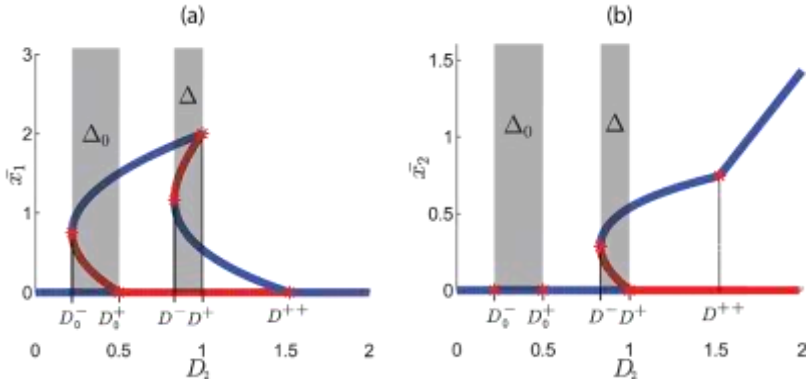


Figure 3.10 One-parameter bifurcation diagrams of the Th1/Th2 mutual inhibition model when $D_1/D_2 = 2$. (a) for the Th1 population and (b) for the Th2 population.

Bistability is observed in the range $[D_0^-, D_0^+]$ (we call this the "Th1 bistability region" since only the Th1 response are activated, i.e. with $\overline{x_1} > 0$), and also in the range$[D^-, D^+]$ (we call this the "Th1/Th2 bistability region" since both the Th1 and Th2 responses are active in this region). The changes that occur to the steady states as we cross the threshold values D_0^-, D_0^+, D^-, D^+ and D^{++} are summarised in Table 3.2. Since we are interested in Th2 bistability, our main concern is for the threshold values relating to the Th1/Th2 bistability region, D^- and D^+, and the width of this bistable region which we define as $\Delta = D^+ - D^-$.

Table 3.2 Table summarising the steady states which exist for the six parameter ranges of D_2. $\overline{x_1}$ and $\overline{x_2}$ represent the focal point solutions of the ODE when $D_1/D_2 > 1$.

Threshold value of D_2	$\boldsymbol{D_0^-}$	$\boldsymbol{D_0^+}$	$\mathbf{D^-}$	$\boldsymbol{D^+}$	$\boldsymbol{D^{++}}$
Steady state $\overline{x_1}$ changes as we increase D_2 past the threshold value	Stable positive state appears $\overline{x_1} = 0$ (stable) $\overline{x_1}$>0 (unstable) **$\overline{x_1}$>0 (stable)**	Steady state at 0 becomes unstable Stable positive state remains $\overline{x_1} = 0$ (unstable) **$\overline{x_1} > 0$ (stable)**	A new positive stable state (smaller than the previous one) appears. $\overline{x_1} = 0$ (unstable) **$\overline{x_1} > 0$ (stable)** $\overline{x_1} > 0$ (unstable) **$\overline{x_1} > 0$ (stable)**	The larger positive steady state becomes unstable $\overline{x_1} = 0$ (unstable) **$\overline{x_1} > 0$ (stable)**	The remaining positive steady state disappears; 0 becomes stable. $\overline{x_1} = 0$ (unstable) **$\overline{x_1} = 0$ (stable)**
Steady state $\overline{x_2}$ changes as we increase D_2 past the threshold value	A second stable steady state at 0 appears **$\overline{x_2} = 0$ (stable)** $\overline{x_2}$=0 (unstable) **$\overline{x_2} = 0$ (stable)**	One of the steady states at 0 becomes unstable, but the other remains stable $\overline{x_2} = 0$ (unstable) **$\overline{x_2} = 0$ (stable)**	Stable positive state appears Steady state at 0 remains stable $\overline{x_2} = 0$ (unstable) **$\overline{x_2} = 0$ (stable)** $\overline{x_2} > 0$ (unstable) **$\overline{x_2} > 0$ (stable)**	0 becomes unstable Positive focal point remains stable $\overline{x_2} = 0$ (unstable) **$\overline{x_2} > 0$ (stable)**	$\overline{x_2} = 0$ (unstable) **$\overline{x_2} > 0$ (stable)**

Case 2: $D_1/D_2 < 1$ (red line)

The steady state moves between (a) and (c) along the red line while changing D_2 Figure 3.8 (a) and (c)) and keeping the ratio $D_1/D_2 = 0.5$. The pink region in Figure 3.8 (e) is always above the line $D_1/D_2 = 1$, meaning that the "mixed response" steady state does not exist when $D_1/D_2 < 1$.

Since regions (a) and (c) do not overlap at all, but instead share a boundary at $D_2 = 1$, the system shows no bistable behaviours around this transcritical bifurcation point, which we label D^{++}. Although the "Th1 only" focal point $(\overline{x_1}, \overline{x_2}) = (x_1^*, 0)$ can be stable for values $D_1/D_2 < 1$, none of the flow lines which converge to $(x_1^*, 0)$ (the encircled branch in Figure 3.11 (a)) originate from points near to the other steady states $(0,0)$ or $(0, x_2^*)$. Therefore, although the system may be technically bistable for some values of D_2 while $D_1/D_2 < 1$, we can only observe monostable behaviour by varying the value of D_2 while $D_1/D_2 < 1$.

The resulting behaviour is shown in the bifurcation diagram (Figure 3.11) which is the projection of the steady states of the system encountered on the red line in Figure 3.8 (a-e) onto the D_2 axis. As before, we compute the one-parameter bifurcation diagrams using MatCont, calculating the steady states of $\overline{x_1}$ and $\overline{x_2}$ as we vary D_2, but now we set$D_1/D_2 = 0.5$ (moving along the red lines in Figure 3.8). The resulting one-parameter bifurcation diagram is shown in Figure 3.11. As before, the x-axis of this figure is the projection of the red line onto the D_2 axis.

For any value $D_1/D_2 < 1$, the "completely off" steady state remains stable as we vary D_2 (Figure 3.11). Indeed, under the condition $D_1/D_2 < 1$, only a transcritical bifurcation from the "completely off" state to the "Th2 response only" state occurs at the value $D_2 = D^{++} = 1$. This contradicts the observation that Th1 responses are observed when the number of DCs <citation here> is low <citation here>. Therefore, we will not consider the case where $D_1/D_2 < 1$. The "Th1 only" steady state $(\overline{x_1}, \overline{x_2}) = (x_1^*, 0)$ does form a branch of solutions which is stable for a range of values between D^- and D^{++}, but the basin of attraction to this steady state is not accessible from the basin for the "completely off" steady state $(\overline{x_1}, \overline{x_2}) = (0, 0)$, so it will not be seen unless some forcing function is applied to the system. Such a forcing function would require some additional assumptions, and so we neglect this scenario in favour of the case where $D_1/D_2 > 1$.

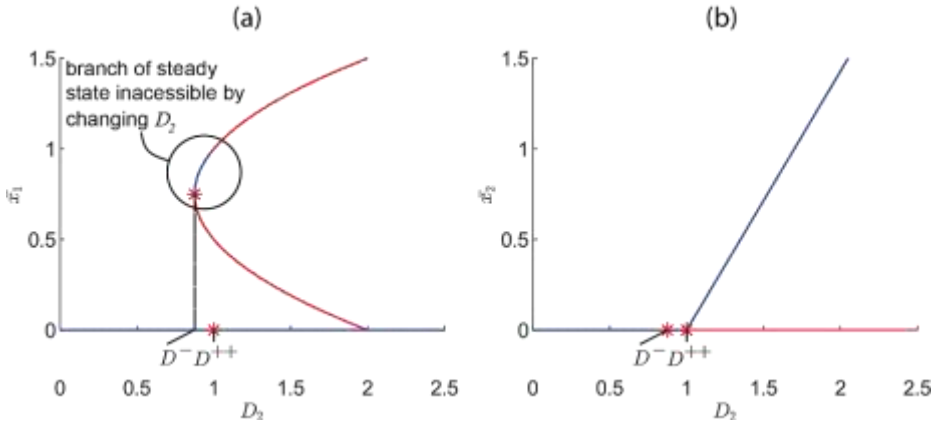


Figure 3.11 One-parameter bifurcation diagrams of the Th1/Th2 mutual inhibition model when $D_1/D_2 = 0.5$.

We summarise the transitions between the different steady states of the system for $D_1/D_2 < 1$ in Figure 3.12.

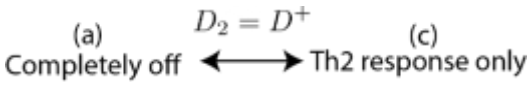


Figure 3.12 Transitions between the two cases of stable positive steady states by changing the value D_2 while keeping the ratio D_1/D_2 constant and less than 1.

To transition from (a) to (c), we must increase the value D_2 above $D^{\wedge} +$ (shown in red on Figure 3.8 (c)).

Robustness of bistability

To check that the bistability is robustly observed for different parameter sets rather thn an aberration of the parameter values we've chosen, we analysed how the width of the bistable regions, $\Delta_0 = D_0^+ - D_0^-$ and $\Delta = D^+ - D^-$, changes according to changes in the remaining parameters π_1 and π_2 and δ, and increasing the activation ratio $\frac{D_1}{D_2} > 1$.

(1) Robustness against changes in proliferation rates π_1 and π_2 and FAS-induced death rate δ

The width of the bistable region, $\Delta = D^+ - D^-$, varies by changing π_1, π_2 and δ (Figure 3.13). The bistability exists for only some values of π_1, π_2 and δ, and the width of this region varies. The white regions show where the bistability is abolished, and replaced with monostable behaviour (in other words, the width of bistability drops to 0 in these white regions). Bistability disappears (i.e. the width

of the Th2 bistable region Δ reduces to zero) if FAS-mediated cell death δ is large relative to the rate of reproduction of the Th1 cells π_1. High δ and low π_1 means that Th1 cells die off quicker than they replicate, so Th1 cells can no longer sustain their own population purely by replication. Additionally, when π_2 is low, the rate of Th2 cell death is more rapid than the rate of Th2 cell reproduction causing the bistable region to disappear.

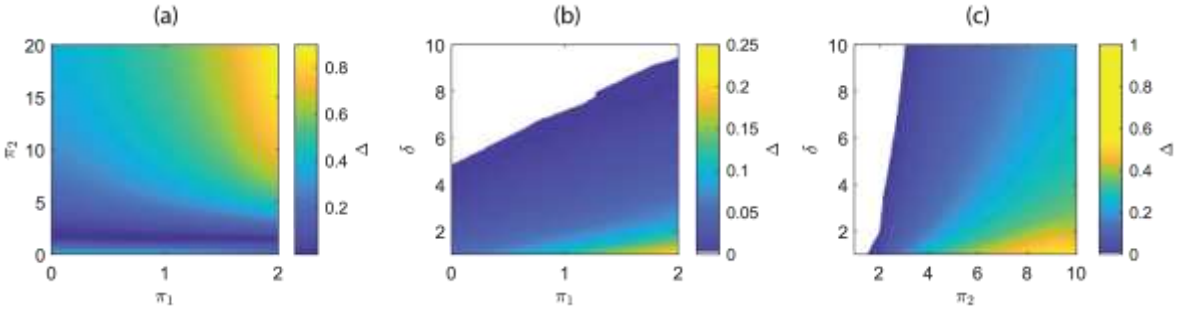


Figure 3.13 Width Δ of the bistable region for Th2 cell number as we vary π_1, π_2 and δ. Parameters: $D_1/D_2 = 2$, $\rho = 0.1$ The white region corresponds to a width of zero. Where the parameter is not varied, it has the value $\pi_1 = 1.5, \pi_2 = 3$ and $\delta = 1$. (a) Width Δ as π_1 and π_2 are varied. (b) Width Δ as π_1 and δ are varied. (c) Width Δ as π_2 and δ are varied.

(2) Robustness against changes in $D_1/D_2(>1)$

The bistable hysteretic behaviour persists as we increase D_1/D_2 (Figure 3.14). Increasing the ratio $\frac{D_1}{D_2} > 1$ simply increases the widths Δ_0 and Δ as the threshold values D_0^-, D_0^+, D^- and D^+ of the bistable region change, but the bistable hysteretic behaviour is maintained.

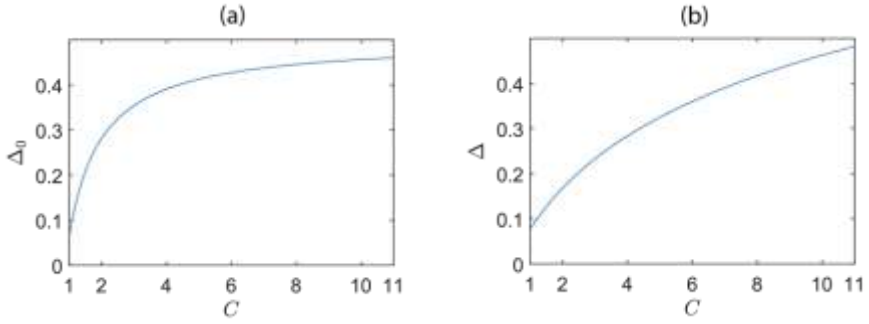


Figure 3.14 Width of the bistable regions as we increase the ratio $C = D_1/D_2$ from 1 upwards. (a) Th1 only bistable region Δ_0 and, (b) Th1/Th2 bistable region Δ.

3.3. Effect of Dupilumab treatment on both models

Now that we have formed two candidate models for Th2 cell population dynamics and confirmed that they both robustly express bistable behaviour, we can start to work out the potential effects of dupilumab treatment on Th2 cell numbers. Dupilumab is a monoclonal antibody with competitive inhibitory activity to the IL4 receptor. It is able to block both IL4 and IL13 signalling to T-cells. Dupilumab is known to be effective in some cases of severe AD where other therapies have been ineffective.

Here, we extend our models to incorporate the effects of competitive IL4 inhibition by dupilumab to investigate the mechanisms of dupilumab treatment. We describe the effect of dupilumab by

replacing the term representing the Th2 cytokines S_2 with a modified $\frac{S_2}{1+[d]}$ in any terms relevant to IL4 signalling. The new denominator represents the effect of competitive inhibition by dupilumab.

For the systemic IL4 feedback model with TGFβ dimerization (section 3.2.3) the original equation (3.6) becomes

$$\frac{dT}{d\tau} = \frac{D}{1+[d]}\frac{T}{1+T} + \frac{\pi_2}{1+[d]}\frac{T^2}{1+T^2} - T \tag{3.10}$$

For the systemic IL4 feedback + Th1/Th2 mutual inhibition model (section 3.2.4) the original equation (3.9) becomes

$$\begin{aligned} \frac{dx_1}{d\tau} &= D_1\frac{x_1}{(1+x_2)^2} + \pi_1\frac{x_1^2}{(1+x_2)^2} - \delta_1\frac{x_1^3}{1+x_2} - x_1, \\ \frac{dx_2}{d\tau} &= \frac{D_2}{1+[d]}\frac{x_2}{1+x_2} + \frac{\pi_2}{1+[d]}\left(\rho + \frac{x_1}{1+x_2}\right)\frac{x_2^2}{1+x_1+x_2} - x_2. \end{aligned} \tag{3.11}$$

Note that IL4 signalling is only represented by the S_2 in the numerator of the first two terms of ODEs, with the S_2 terms in the denominators representing the effect of IL-10 and $TGF\beta$ signalling which remain unaffected by Dupilumab. For this reason, the $\frac{1}{1+[d]}$ term appears only in the denominator of the first two terms in equations 3.10 and 3.11.

3.3.1. Effect of dupilumab on critical values

As the dupilumab concentration is increased, the threshold values D^- and D^+both increase both with different slope and eventually become equal (Figure 3.15). Thus, the width of the Th2 bistable region, $\Delta = D^+ - D^-$ gradually decreases until the bistability completely disappears when $\Delta = 0$.

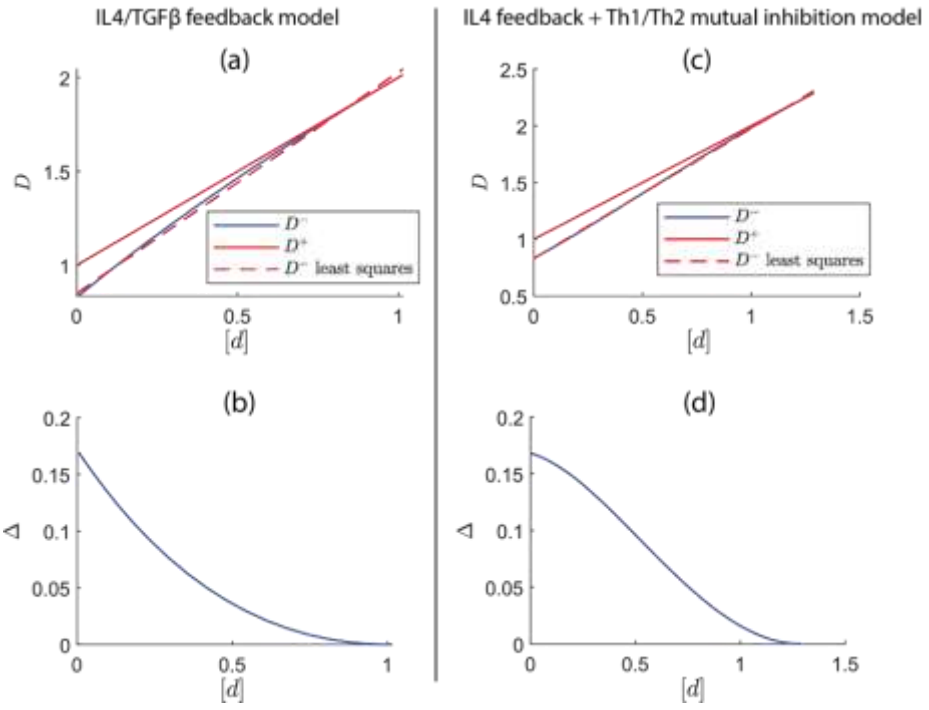


Figure 3.15 Effect of dupilumab treatment on the critical values of the Th2 bistability D^- and D^+ and width of the bistable region $\Delta = D^+ - D^-$ for (a-b) the "systemic IL4/TGFβ feedback model" described in section 3.2.3, (c-d) the "systemic IL4 feedback + Th1/Th2 mutual inhibition model" described in section 3.2.4.

While D^+ is analytically described by $D^+ = 1 + [d]$ (Appendix F), D^- has no easily derived analytical form. We therefore describe D^-phenomenologically by $D^- = A + B[d]$. The coefficients A and B are obtained by a least-squares regression against the numerically calculated line for D^- (see Appendix F). The regression line is the red dashed line in Figure 3.15 (a) and (c) while the numerically calculated

line is in solid blue. To ensure that the regression line for D^- meets with D^+, we requre that $\mathrm{A} < 1$ and $\mathrm{B} > 1$ so that $D^- < D^+$ when $[d] = 0$. Under this parameterisation, the width of the bistable region may be defined by equation 3.12.

$$D^+ - D^- = (1 - A) - (1 - B)[d] \tag{3.12}$$

becomes zero at equation 3.13.

$$[d] = \frac{1 - A}{1 - B} \tag{3.13}$$

where bistability is no longer observed (Figure 3.16 d). We assume that for all $[d] > \frac{1-A}{1-B}$, the width remains at zero (i.e. it can't become negative). We can define this as equation 3.14.

$$\Delta = \max(0, (1 - A) - (1 - B)[d]) \tag{3.14}$$

Note that for the Th1/Th2 mutual inhibition model (Figure 3.16 (b) and (c)), the Th1 response still shows bistable behaviour (Figure 3.16 (b) grey shading at around $D_2 = 0.5$ visible in the zoomed subplot), but the Th2 response loses the bistable behaviour (Figure 3.16 (c)). Since we are only interested in the Th2 bistability, the Th1 bistability being persistent is not important to us.

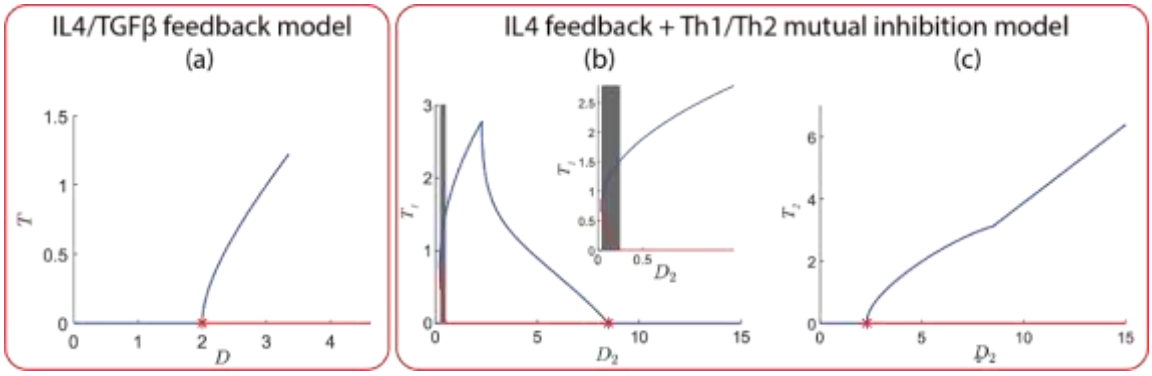


Figure 3.16 One-parameter bifurcation diagrams when the dupilumab treatment has reduced the width of the bistable region to zero $\Delta = D^+ - D^- = 0$ at a cusp bifurcation for (a) systemic IL4/TGFβ feedback model of section 3.2.3, when we set $[d] = 1$ and (b) Th1 cell population for the systemic IL4+Th1/Th2 mutual inhibition model where the inset plot shows a zoomed-in region of the plot and (c) Th2 cell population for the systemic IL4 feedback + Th1/Th2 mutual inhibition of section 3.2.4when we set $[d] = 1.2877$.

3.3.2. Reversal of inflammation under Dupilumab treatment

Analysis of this dupilumab model demonstrates that inhibition of IL4 signalling can achieve two things in both of our candidate models (IL4 feedback model and the IL4 feedback + Th1/Th2 mutual inhibition model)

(1) increases D^- towards D^+, making the bistable region narrower and easier to switch off when we decrease the DC number below $D = D^-$. A decrease of $\Delta = D^+ - D^-$ means that a patient can more easily achieve the healthy steady state by reducing the antigen exposure of their lesional skin.
(2) Increases D^+, making the Th2 response more difficult to be activated by an increase of $D > D^+$.

The implications of these two points are that when dupilumab treatment is applied, we require the DC number to be decreased less to switch off the Th2 response (red lines in Figure 3.17) relative to the untreated case (blue lines in Figure 3.17) since the switching off threshold value D^- has been increased by the presence of dupilumab.

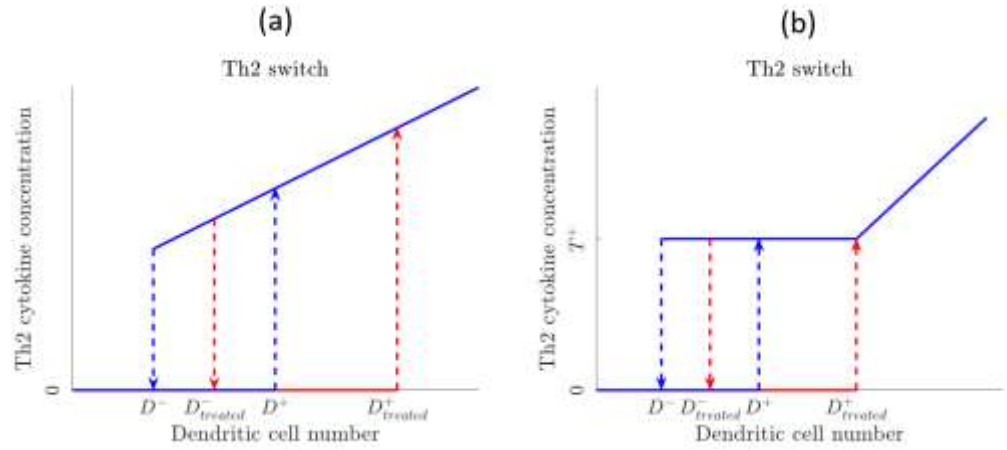


Figure 3.17 Th2 switch after dupilumab treatment.
Blue lines represent the untreated switch, while the red lines represent the switch after dupilumab treatment.
(a) systemic IL4 feedback model of section 3.2.3,
(b) Systemic IL4 feedback + Th1/Th2 mutual inhibition of section 3.2.4.

3.4. Comparing Models 1 and 2

We produced two models (equations 3.7 and 3.11) which both demonstrate that the Th2 responses can be described by a bistable switch with hysteresis with respect to the number of activated DCs in the lymph nodes. For the purpose of this discussion, we call refer to the two models as follows:

Model 1: TGFβ negative feedback model, described by the nondimensionalised ODEs in equation 3.7.

Model 2: IL4-feedback + Th1/Th2 mutual inhibition model, described by nondimensionalised ODEs in equation 3.9.

Both models produce similar results, where bistable hysteresis is observed, although the models have different mechanisms to bring about that bistability and Model 2 is quite a bit more complicated than model 1. Model 1 produces bistability as the result of an order-2 feedback term, while Model 2's bistability results from that mutual inhibition motif.

Model 1 requires that we assume TGFβ negative feedback to be of order-2 (i.e. dimerization) in order to produce bistable behaviours. The more complicated model 2 does not require the assumption of order-2 TGFβ feedback but instead relies on the mutual inhibitory relationship between Th1 and Th2 cells to produce bistable behaviours.

Model 2 is supported by extensive literature proposing that Th1/Th2 mutual inhibition generates polarising immune environments as seen here. It is interesting to note that the known qualitative behaviours of the system can be reproduced with a more parsimonious model through TGFβ cooperative inhibition, but that the experimental emphasis is still on Th1/Th2 mutual inhibition.

The effect of dupilumab on both models was found to be very similar; namely, dupilumab treatment has the effect of increasing critical value D^- until it reaches D^+, making the bistable region narrower, and thus easier to reverse the Th2 response. The two models, therefore, predict the same mechanism of action for dupilumab, despite the mechanism producing bistability being different for the two models.

Since both of the models described here provide a reasonable explanation for the bistable behaviour of the Th2 cytokine response, we cannot reasonably disregard either model. Fortunately, both models provide a very similar result, featuring a bistable Th2 cell number where the threshold values of that bistability increase when dupilumab is applied. For the time being, we proceed with both models, and remain agnostic as to the underlying mechanism producing the bistability.

The shape of the one-parameter bifurcation diagrams for the systemic IL4 feedback model (Figure 3.4) and the IL4-feeback + Th1/Th2 model (Figure 3.10) can be simplified by piecewise switch definitions.

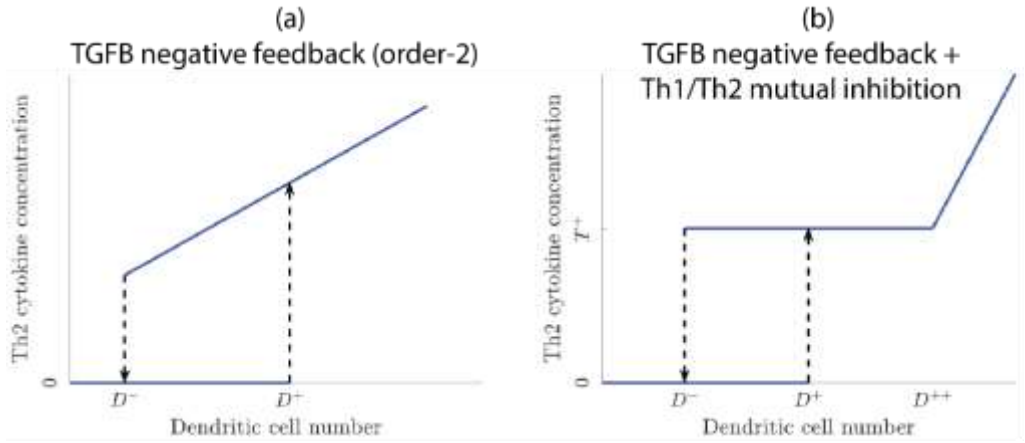


Figure 3.18 Th2 switch functions.

For the systemic IL4 feedback model (Figure 3.18 (a)) we use

$$T(t) = \begin{cases} T_{\text{off}} \to m \cdot (D(t) - D^-) + T_{\text{on}}, & \text{if } \dot{D}(t) > 0 \text{ and } D(t) = D^+ \\ m \cdot (D(t) - D^-) + T_{\text{on}} \to T_{\text{off}}, & \text{if } \dot{D}(t) < 0 \text{ and } D(t) = D^- \end{cases} \tag{3.15}$$

For the systemic IL4 feedback + Th1/Th2 mutual inhibition model (Figure 3.18 (b)) we use

$$T(t) = \begin{cases} T_{\text{off}} \to T_{\text{on}} & if \quad D(t) = D^+ \; and \; \dot{D}(t) > 0 \\ T_{\text{on}} \to T_{\text{off}} & if \quad D(t) = D^- \; and \; \dot{D}(t) < 0 \\ m \cdot (D(t) - D^{++}) + T_{\text{on}} & if \quad D(t) > D^{++} \end{cases} \tag{3.16}$$

We will use Figure 3.18 and equations 3.15 and 3.16 to describe the population numbers of Th2 cells in our model development from this point on. Since we have no good way to determine parameter values for this model, so we simply set the threshold values, height T_{on} and gradient m of this function arbitrarily. We will assume that T_{off} (i.e. the level of Th2 activation when there is no dendritic cell signal to activate the system) is zero because we do not expect to see an active Th2 population under healthy conditions.

3.5. Implementing the reversible T-switch in the double-switch model

In this chapter, we proposed a new T-switch model on T-helper cell population dynamics. We implemented several simplifying assumptions to this population dynamics model, tested two hypotheses about Th2 cell regulation and considered both the IL4/TGFβ negative feedback model and the IL4 feedback + Th1/Th2 mutual inhibition model. We concluded that both our proposed regulatory hypotheses are able to generate bistable hysteretic switching behaviour in the concentration of Th2 cytokines. Whilst we have undoubtedly increased the complexity of our model through our modifications here, we are more confident in our founding assumptions. It is therefore a necessary complication.

We can now incorporate the reversible T-switch in the double-switch model (equation 2.1) by simply replacing the original T-switch definition (equation 2.3) with the modified definitions in either equation 3.15 or 3.16. We check the effect of our modification by reproducing the analysis for the bifurcation diagram of the double-switch, using the T-switch definition in equations 3.15 and 3.16.

We set the parameters for the new T-switch as in Table 3.3. The new parameters D^- and m were chosen arbitrarily. D^- to be approximately half of D^+, and m so that the gradient would resemble the bifurcation diagrams in Figure 3.18 (a) and (b).

Table 3.3 Nominal parameters for the modified T-switch

Symbol	Name	Nominal value	Reference
D^-	Th2 deactivation threshold	40 (cells/mL)	
D^+	Th2 activation threshold	85 (cells/mL)	[23], [26]
T_{off}	Th2 population when not activated	0	[26]
T_{on}	Th2 population when activated	1	[26]
m	Gradient of Th2 switch	0.1	

While our new T-switch definitions bring about a dependence on the D state variable, it doesn't cause great complication to our analysis because the ODE for dendritic cell number D is decoupled from the other ODEs in the system, being only affected by the R-switch. Thus, the analysis to find the bifurcation lines for both versions of the T-switch is nearly identical for the IL4/TGFβ feedback model and the IL4 feedback+Th1/Th2 mutual inhibition model. Both analyses are described in Appendix I. The culmination of this analyses are the new expressions for the boundary-equilibrium bifurcation lines for this modified double-switch model.

Armed with our new expressions for the boundary-equilibrium lines, we can once again plot the bifurcation diagram for this model with the T-switch off and then again with the T-switch on (Figure 3.19).

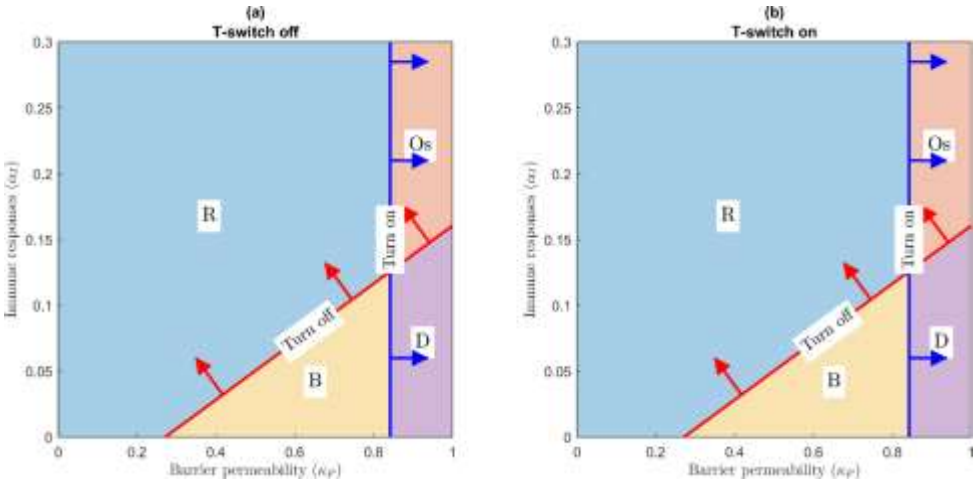


Figure 3.19 Bifurcation diagrams for the double-switch model with modified K- and T-switches

We can note again that the bifurcation diagrams remain identical when the T-switch off and on. This is expected because the inequalities that define the locations of the BE bifurcation lines depend only on the value of the T-switch (i.e. the magnitude of T_{off} and T_{on}) and not on the reversibility or otherwise of the switch.

3.6. Conclusion

In this chapter, we have formed a mathematical of the regulation of Th2 cell population. Bifurcation analyses of the model have shown that Th2 population numbers show bistability when the positive feedback effects of TGFβ cytokine signalling and mutual inhibition between the Th1 and Th2 population are considered. We analysed these two sources of positive feedback separately, and found that while either mechanism alone is sufficient to bring about bistability, the Th1 mutual inhibition mechanism results in a more robust form of bistability in the sense that the bistability is retained even if we vary the remaining hyperparameters (π_1,π_2 and δ). We conclude that we may summarise the Th2 cytokine concentration as we vary the DC stimulation by a bistable hysteretic switch. This bistable

switch is similar to the form used in the previous “double-switch” model in that it can be simplified to a bistable hysteretic switch, but our new switch is now reversible so that the Th2 response can be deactivated as well as activated.

Chapter 4. Mathematical modelling of the regulatory mechanisms of keratinocyte differentiation for epidermal homeostasis

In the double-switch model of AD, the skin barrier functions are simplified to a single number between zero representing badly damaged skin, and one representing a fully competent barrier. This description of the skin with a single value was baked into the ODE representing the barrier by structuring it as a logistic growth equation. In the context of early-stages of AD, this simplification was sufficient, but it begins to appear insufficient for a model considering the long-term dynamics of chronic AD, where these individual features of the barrier function can begin to fail independently of one-another.

“Skin barrier integrity” is a term used in the double-switch model to describe the combination of features exhibited by healthy skin, namely the ability to resist physical, chemical, and pathogenic assaults all of which come about as a result of various products within the epidermis. In this chapter, we begin to unpack how these products are regulated and come up with a more nuanced model for barrier integrity.

The physical resilience of the barrier is achieved through expression of connective-tissue proteins such as filaggrin and corneodesmosomes which connect corneocytes at their faces, forming a physically tough surface consisting of dead cells resistant to abrasion. Keratinocytes begin to produce lipid-filled bags called lamellae as they migrate upwards through the skin barrier towards the stratum corneum. Eventually, these lamellae form a densely packed impermeable layer, impermeable to water.

Pathogenic resistance is brought about by a combination of anti-microbial peptides produced by the keratinocytes themselves [57], which are able to directly kill potentially pathogenic microbes, and bacteriocins produced by commensal species of microbe living on the skin surface [58]. Part of the healthy function of the skin is to maintain a population of friendly commensal species on the skin surface, which act both to actively kill off potentially pathogenic species, but also displace them by competing with them for resources [59].

Although many of the intercellular signalling pathways and the genetic regulation of these barrier components have been described experimentally, it is not yet clear how these individual mechanisms interact to bring about the dysfunctional skin barrier on timescales pertinent to chronic AD. In this chapter, we build a mathematical model to bridge the gap between the biological knowledge of signalling pathways and the higher-level description of keratinocyte differentiation in the granular layer of the epidermis. We consider how this model could be incorporated into our hybrid dynamical model of chronic AD as a switch, and we name it the G-switch (G referring to the granular layer). This chapter summarises work produced in collaboration with Dr Elisa Domínguez-Hüttinger. The manuscript is under preparation for submission for a journal publication.

4.1. Model description

Model development for this keratinocyte differentiation model was carried out by Elisa Domínguez-Hüttinger et al. who carried out a literature search to find regulatory mechanisms pertinent to the processing of calcium signals and the regulation of the EDC, curating a network of proposed regulatory pathways and carrying out a process of model reduction that led to the simplified network in Figure 4.1. This simplified network was found to reproduce the behaviour of the larger network. We consider

Stat3 and Np63 as state variables of an ODE system modelling the genetic regulation of keratinocyte differentiation.

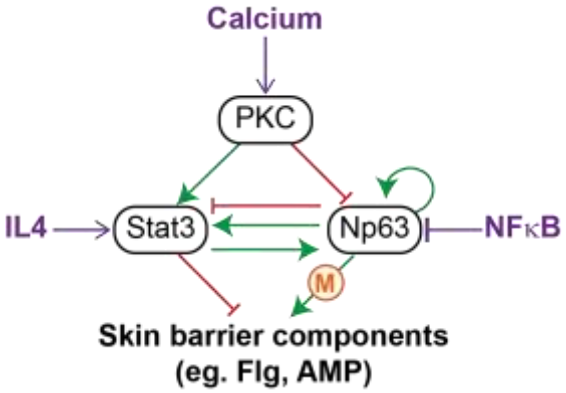


Figure 4.1: The keratinocyte differentiation motif comprises positive feedback with cooperativity.

The resulting regulatory structure, which we refer as ***keratinocyte differentiation motif*** (Figure 4.1), represents the dynamic interplay between molecules (Stat3 and Np63) that directly regulate the expression of the epidermal differentiation complex, a family of around 50 genes which produce structural proteins such as filaggrin, and a range of AMPs [68]. The input to the network is the external calcium level (increased upon barrier damage) that activates protein kinase C (PKC), which activates Stat3 and inhibits Np63 expression. Stat3 and Np63 activate with each other forming a positive feedback loop. Np63 induces its own expression, forming a second positive feedback loop. Np63 also inhibits Stat3, forming an additional negative feedback loop. The dual effect of Np63 on Stat3 in the keratinocyte differentiation motif is a consequence of conflicting signals found by combing the literature, where Np63 induces expression of Stat3 but also induces degradation of the products.

The expression of terminal differentiation markers (Flg, AMPs) is induced by Np63 and inhibited by Stat3. Expression of Stat3 and Np63 are also induced by pro-inflammatory cytokines (IL4) and activation of innate immune responses (NFκB), respectively.

The kinetic model of the keratinocyte differentiation motif is given by the dynamic coupling between $Stat3$ and $Np63$ and the output epidermal differentiation complex EDCas the system of equations 4.1.

$$\begin{aligned}
\frac{dStat3(t)}{dt} &= PKC + IL4 + v_{Np63} \cdot Np63(t) - Stat3(t) \cdot \left(d_{Stat3} + \varphi \cdot Np63(t)\right) \\
\frac{dNp63(t)}{dt} &= va_{Np63} \cdot \frac{Np63(t)^{n_H}}{{k_{Np63}}^{n_H} + Np63(t)^{n_H}} + v_{Stat3} \cdot Stat3(t) - Np63(t) \cdot \left(d_{Np63} + NFkB + PKC \cdot d_{PKC}\right) \quad (4.1) \\
\frac{dEDC(t)}{dt} &= \frac{a_{EDC} \cdot (\alpha \cdot Np63(t) \cdot e^{-\beta t})}{1 + i_{EDC}\, Stat3(t)} - d_{EDC} \cdot EDC(t)
\end{aligned}$$

$Stat3(t)$, $Np63(t)$ and $EDC(t)$ represent the expression levels of Stat3, Np63 and the terminal differentiation markers. Stat3 is produced at a constant rate, PKC and by $IL4$, representing the calcium- and inflammation-mediated induction of Stat3, respectively. Stat3 is also induced by $Np63$ with rate v_{Np63}, comprising a positive feedback loop as $Np63$ production is also induced by Stat3. The basal degradation rate of Stat3 is described by d_{Stat3}, and Stat3 is further degraded by Np63 with rate φ.

Np63's auto-induced expression, forming an additional positive feedback loop, is described by a Hill function with a maximal velocity va_{Np63}, a half maximal inductor concentration k_{Np63} and a Hill-coefficient n_H that determines the degree of steepness of the sigmoidal function.

$$va_{Np63} \cdot \frac{Np63(t)^{n_H}}{{k_{Np63}}^{n_H} + Np63(t)^{n_H}}$$

Np63 is produced via a Stat3-dependent pathway as well, with a production rate of v_{Stat3}. Degradation of Np63 has a basal rate d_{Np63}, an $NFkB$ -dependent rate and a PKC-dependent rate weighted by the parameter d_{PKC}.

EDC expression is induced by Np63 with a decaying exponential ($\alpha \cdot Np63(t) \cdot e^{-\beta t}$) with rate a_{EDC}, representing the long-term effects of methylation of EDC promoter by Np63 [69]–[71]. The integral of Np63 is weighted by parameter α and the decaying exponential decreases with parameter β. EDC expression is inhibited by Stat3 with rate i_{EDC} and degraded linearly with rate d_{EDC}.

Note that this parameter optimisation was carried out by Elisa Domínguez-Hüttinger and colleagues, an article describing this process is in preparation.

The nominal parameter values resulting from this parameter optimisation are given in Table 4.1.

Symbol	Nominal values	Units	Description
		Inputs	
$\boldsymbol{PKC}$	0.1 to 5 (low to high Calcium switch)	$[y]/[t]$	Ca-mediated induction of Stat3
$\boldsymbol{IL4}$	0	$[y]/[t]$	Inflammation-mediated Stat3 induction
$\boldsymbol{NFkB}$	0	$1/[t]$	Immune response-mediated degradation of Np63
		Stat3-specific parameters	
$\boldsymbol{v_{Np63}}$	2	$1/[t]$	Strength of Np63 on Stat3 production (positive feedback #1)
$\boldsymbol{d_{Stat3}}$	1	$1/[t]$	Stat3 degradation rate
$\boldsymbol{\varphi}$		$1/([t] \cdot [y])$	Strength of Np63 on Stat3 degradation
		Np63-specific parameters	
$\boldsymbol{va_{Np63}}$	10	$1/[t]$	Maximal strength of Np63 on Np63 (positive feedback #2)
$\boldsymbol{n_H}$	3	NA	Hill-coefficient for the positive feedback of Np63 on Np63
$\boldsymbol{k_{Np63}}$	1	$[y]$	Half-maximal Np63 concentration for Np63 activation
$\boldsymbol{v_{Stat3}}$	1	$1/[t]$	Effect of Stat3 on Np63 production (positive feedback #3)
$\boldsymbol{d_{Np63}}$	6	$1/[t]$	Np63 degradation rate
$\boldsymbol{d_{PKC}}$	1	$1/[y]$	Strength of PKC-mediated Np63 degradation
		Epidermal differentiation markers-specific parameters	
$\boldsymbol{a_{EDC}}$	100	$1/[t]$	Strength of Np63-mediated induction of EDC expression
$\boldsymbol{i_{EDC}}$	57	$1/[y]$	Strength of Stat3-mediated inhibition of EDC expression
$\boldsymbol{d_{EDC}}$	1	$1/[t]$	EDC degradation
$\boldsymbol{\alpha}$	8.2	$1/[t]$	Weighting factor for Np63 integral-over-time
$\boldsymbol{\beta}$	2	$1/[t]$	Exponent of decaying exponential weighting historical Np63

Table 4.1 Nominal parameters of the kinetic model.

4.2. Model behaviour

4.2.1. Switch-like and history-dependent keratinocyte differentiation in response to change in the extracellular calcium level.

The key regulatory structure of our kinetic model is mutual positive feedback between Stat3 and NP63. Note that this bistable behaviour emerging from a system featuring positive-feedback loops is similar to the situation for the T-switch in section 3.2.4, where mutual negative-feedback of Th1 and Th2 populations leads to a bistable outcome. This kind of mutual feedback is a common feature known to result in bistability of a system [36], [73].

Indeed, our kinetic model with the nominal parameter values (Table 4.1) shows a bistable behaviour. This bistability is a structural feature of our model as it occurs through intersection of the Stat3 nullcline (a first order polynomial) and the Np63 nullcline (a sigmoidal function) (Figure 4.2 A).

Projection of our model dynamics into the 2D Np63-Stat3 phase plane can visualise the effects of changing PKC, the main input parameter of our model. For low PKC values, there are two steady states, at a low and a high Np63 (and Stat3), corresponding to low and high expression of epidermal differentiation markers. As PKC increases, the Stat3 nullcline shifts upwards while the Np63 nullcline straightens, eventually resulting in a loss of the low Np63 steady state (Figure 4.2 B).

Bifurcation analysis confirms bistable focal point levels of the terminal differentiation markers ($\overline{Np\ 63}$) with the calcium level (PKC) as a bifurcation parameter (Figure 4.2 C). The observed bistability is consistent with a switch-like history-dependent keratinocyte differentiation in response to the extracellular calcium level changes. The expression levels of terminal differentiation markers are stabilised eventually at a low or high state [74]–[76] after a short-lived transient change triggered by a change in the extracellular calcium level. When it is stabilised at a high state due to a sufficiently long-lasting increase of the calcium level, it remains at a high state for days after the extracellular calcium level decreases [77].

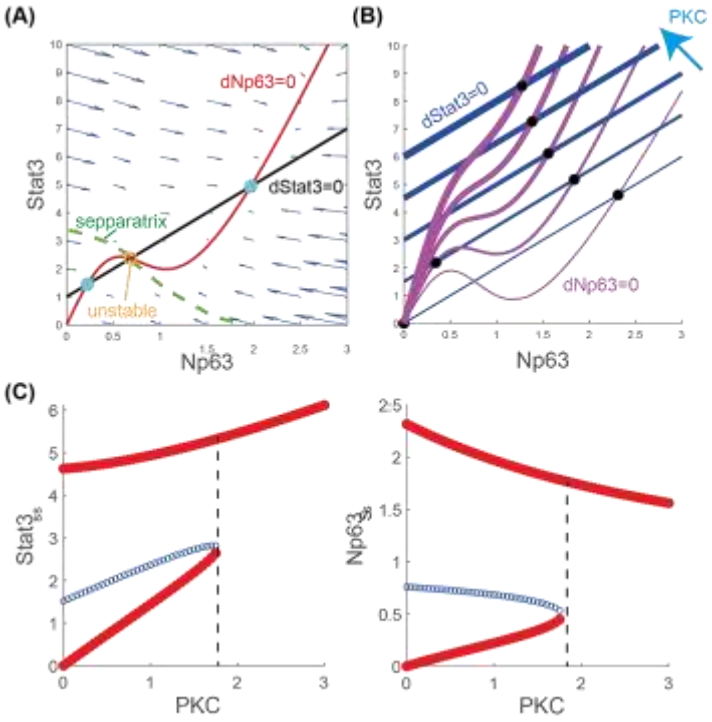


Figure 4.2: Bistability observed in the kinetic model.

4.2.2. Modulation of keratinocyte differentiation by immune responses and inflammation

To investigate the effects of innate immune responses and inflammation on keratinocyte differentiation, we conducted a similar bifurcation analysis of the model with an increased level of NFκB (to mimic innate immune responses) and of IL4 (to mimic inflammation).

Effects of innate immune responses

Innate immune responses (an increase in NFκB value) were found to lower the activation and deactivation thresholds of PKC (B^+ and B^- both decrease) for keratinocyte differentiation and increase the width of the region of bistability ($B^+ - B^-$ increases). The deactivation threshold B^- eventually moves to below 0, making the bistability irreversible. This is illustrated in

The observed changes represent the skin being more sensitive to barrier damage. This result is consistent with an increased expression of terminal differentiation markers triggered by innate immune responses [78]–[86].

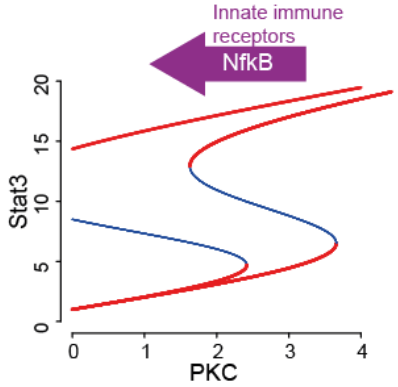


Figure 4.3 The effect of increasing NFκB on the threshold values of differentiation marker Stat3.

Effects of inflammation

When the type-2 adaptive immune response is activated, Th2 cells locate to the lesional skin site and release inflammatory cytokines IL4 and IL13. Here, we assess the effect of IL4 cytokine on the threshold values of the Stat3-Np63 system.

Release of inflammatory cytokine IL4 increases the sensitivity of the keratinocytes to calcium stimulation; the thresholds B^+ and B^- both decrease linearly as IL4 increases (Figure 4.4). Their decrease occurs with the same slope hence the width of the bistability region is unaltered. Both high and low focal point values of Stat3 and Np63 are unaffected by IL4.

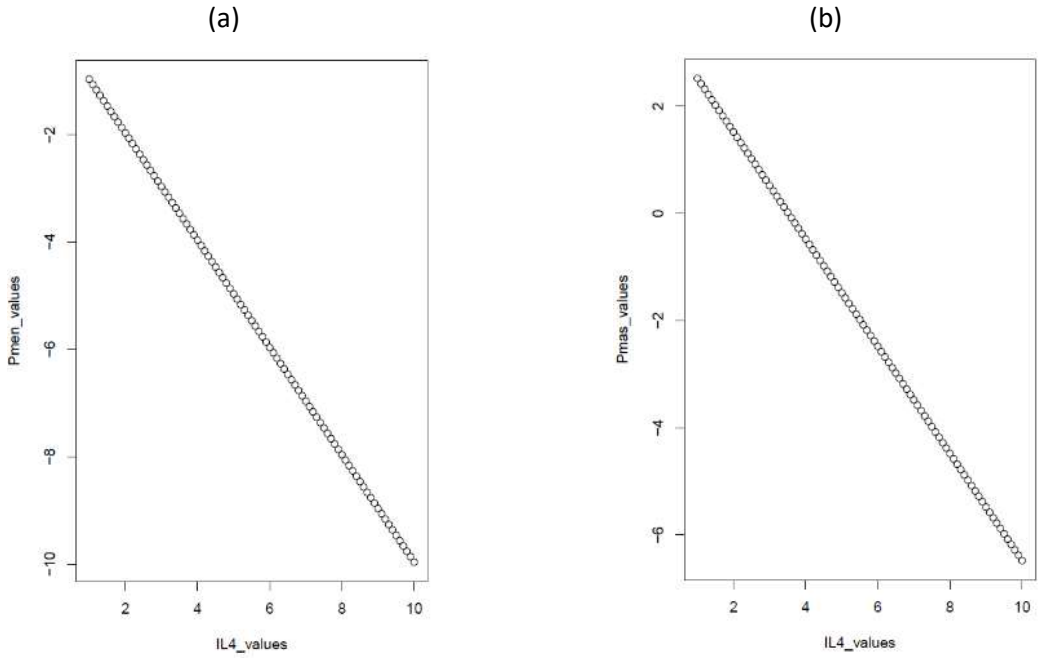


Figure 4.4 The effect of IL4 on the threshold values of differentiation marker Stat3.
Threshold B^- decreases linearly as IL4 is increased.
Threshold B^+ decreases linearly as IL4 is increased.

4.2.3. Summarising keratinocyte differentiation as a bistable switch

Finally, we simplify this model of keratinocyte differentiation into a "switch" similar to the R-, K- and T-switches of the double-switch model. We define a piecewise function as equation 4.2 where $G_{\text{on}} > G_{\text{off}}$ and $B^+ > B^-$.

$$G(t) = \begin{cases} G_{\text{on}} \to G_{\text{off}}, & \text{if } \dot{B}(t) > 0 \text{ and } B(t) = B^+ \\ G_{\text{off}} \to G_{\text{on}}, & \text{if } \dot{B}(t) < 0 \text{ and } B(t) = B^- \end{cases} \tag{4.2}$$

Given our analysis of the modulation of the threshold values by innate immune response, and IL4-mediated inflammatory signals, we define the threshold values B^+ and B^- as functions of the R- and T-switches. The R-switch represents the activity of the innate immune response, while the T-switch represent the activation of Th2 cells which release IL4.

Table 4.2 Definition of the threshold value B^- and B^+ based on the modulation of the thresholds found in section 4.2.2.

B^-		R-SWITCH	
		off	on
T-SWITCH	off	0.2	0.3
	on	0.1	0.2

B^+		R-SWITCH	
		off	on
T-SWITCH	off	0.8	0.9
	on	0.7	0.8

In the coming chapters, we will incorporate this new "G-switch" into the hybrid dynamical system model.

4.3. Conclusions

In this chapter, we proposed the first mechanistic model of keratinocyte differentiation to our knowledge. First, we assembled a regulatory network of biochemical interactions involved in keratinocyte differentiation based on our literature search. These interactions were converted into

mathematical assumptions and compiled into an ODE. Nominal parameter values were obtained by fitting our model through global optimisation to time-course data of calcium switch experiments. Bifurcation analysis of our model at these parameter values was found to reproduce the abrupt and history-dependent keratinocyte differentiation in response to calcium that has been reported experimentally and showed that immune responses and inflammation shape this decision-making process by shifting the thresholds of activation and inactivation of keratinocyte differentiation.

Given these bistable properties of the mathematical model for keratinocyte differentiation, we choose to represent the rapid responses of this genetic regulation by a bistable switch whose threshold values are modulated by signals from the innate and adaptive immune systems (the R- and T-switch, respectively). This bistable switch summarises the effect of damage to the topmost layer of the epidermis: Calcium released by epidermal damage signals to the granular layer of keratinocytes that they should enter the “active” mode of the bistable response, enhancing their production of barrier components. Upon successful repair of the barrier, the calcium signalling may drop sufficiently to cause the granular layer keratinocytes to exit this “active” state, reverting to their homeostatic level of production of barrier components.

Chapter 5. Modelling chronic AD

Having defined the new model for keratinocyte differentiation in Chapter 4, we now need to incorporate this feature into the existing model of AD. In this chapter, we explain how the innate immune system, adaptive immune system and epidermal tissue growth and repair interact, incorporating these systems into a single mechanistic model which we call the "triple-switch" model.

The model is structured as a hybrid dynamical system to help incorporate the drastically different timescales of each system into a single model. In our case, we find that genetic-level regulation in the RK- and G-switches occur more rapidly relative to the tissue-level dynamics of epidermal barrier growth (with RK- and G-switch dynamics occurring on timescale of hours, and tissue-level dynamics on the timescale of days). Similarly, the population dynamics of Th2 cells is assumed to be rapid relative to barrier growth. The slower dynamics of pathogen load (perturbations are assumed slow due to gradual exposure to pathogen from the environment, and the robustness of microbial populations due to the dense web of inter-species interactions), barrier growth, production of AMPs and population of dendritic cells are represented by a system of ordinary differential equations (ODEs).

Whilst it is possible to represent this diversity of timescale using a conventional ODE model, our hybrid-dynamical system approach has the advantage of explicitly separating the two to draw attention to the parts that matter to us. The fast-timescale systems are assumed to be symptomatic of AD (such as inflammation or flakiness of the skin) while the slower timescales refer to the hidden or less tangible variables which patient and clinician may not be aware of unless specific tests are carried out such as pathogen level or barrier integrity.

With this new "triple-switch" mechanistic model, we aim to explain the poorly understood dynamics of chronic AD. The model we present is a hybrid dynamical system similar to the double-switch model [18]. The tools available to analyse this type of model are relatively underdeveloped, and mainly relate to optimisation rather than bifurcation analysis. In this chapter, we develop bifurcation analysis tools specifically designed for this class of model and apply them to identify the dynamics of our new model.

5.1. Triple-switch model

The triple-switch model (Figure 5.1) is defined by

$$\begin{aligned}
\frac{dP(t)}{dt} &= (P_{env} + P_{ext})\frac{\kappa_P}{1+\gamma_{PB}\cdot B} - \alpha_P \cdot A \cdot P - \delta_P \cdot P \\
\frac{dB(t)}{dt} &= \kappa_B \cdot \frac{1}{1+\gamma_{BB}\cdot B}\cdot\frac{1}{1+\gamma_{BR}\cdot R(P^-,P^+)} G(B^-,B^+) - \delta_B \cdot K(P^-,P^+)\cdot B - \delta_{Bext}(t)\cdot B \\
\frac{dA(t)}{dt} &= \kappa_A \cdot \frac{1}{1+\gamma_{AB}\cdot B}\cdot G(B^-,B^+) - \delta_A \cdot A \\
\frac{dD(t)}{dt} &= \kappa_D \cdot R(P^-,P^+) + D_{ext}(t) - \delta_D \cdot D
\end{aligned} \tag{5.1}$$

The states of the switches, R, K, G, and T, are discrete-valued (bottom of Figure 5.1), and defined by a piecewise-smooth function,

$$\big(R(P^-,P^+),K(P^-,P^+)\big) = \begin{cases} (R_{\text{off}},K_{\text{off}}) \to (R_{\text{on}},K_{\text{on}}), & \text{if } \dot{P}(t) > 0 \text{ and } P(t) = P^+ \\ (R_{\text{on}},K_{\text{on}}) \to (R_{\text{off}},K_{\text{off}}), & \text{if } \dot{P}(t) < 0 \text{ and } P(t) = P^- \end{cases} \tag{5.2}$$

$$T(D^-,D^+) = \begin{cases} T_{\text{off}} \to m\cdot(D(t)-D^-)+T_{\text{on}}, & \text{if } \dot{D}(t) > 0 \text{ and } D(t) = D^+ \\ m\cdot(D(t)-D^-)+T_{\text{on}} \to T_{\text{off}}, & \text{if } \dot{D}(t) < 0 \text{ and } D(t) = D^- \end{cases} \tag{5.3}$$

$$T(D^-,D^+,D^{++}) = \begin{cases} 0 \to T_{\text{on}} & \text{if} \quad D(t) = D^+ \text{ and } \dot{D}(t) > 0 \\ T_{\text{on}} \to 0 & \text{if} \quad D(t) = D^- \text{ and } \dot{D}(t) < 0 \\ m(D-D^{++})+T_{\text{on}} & \text{if} \quad D(t) > D^{++} \end{cases} \tag{5.4}$$

$$G(B^-, B^+) = \begin{cases} G_{\text{off}} \to G_{\text{on}}, & \text{if } \dot{B}(t) > 0 \text{ and } B(t) = B^+(R,T) \\ G_{\text{on}} \to G_{\text{off}}, & \text{if } \dot{B}(t) < 0 \text{ and } B(t) = B^-(R,T) \end{cases} \tag{5.5}$$

The three switches are each justified by a mechanistic model detailed in Appendix A, Chapter 3 and Chapter 4, respectively.

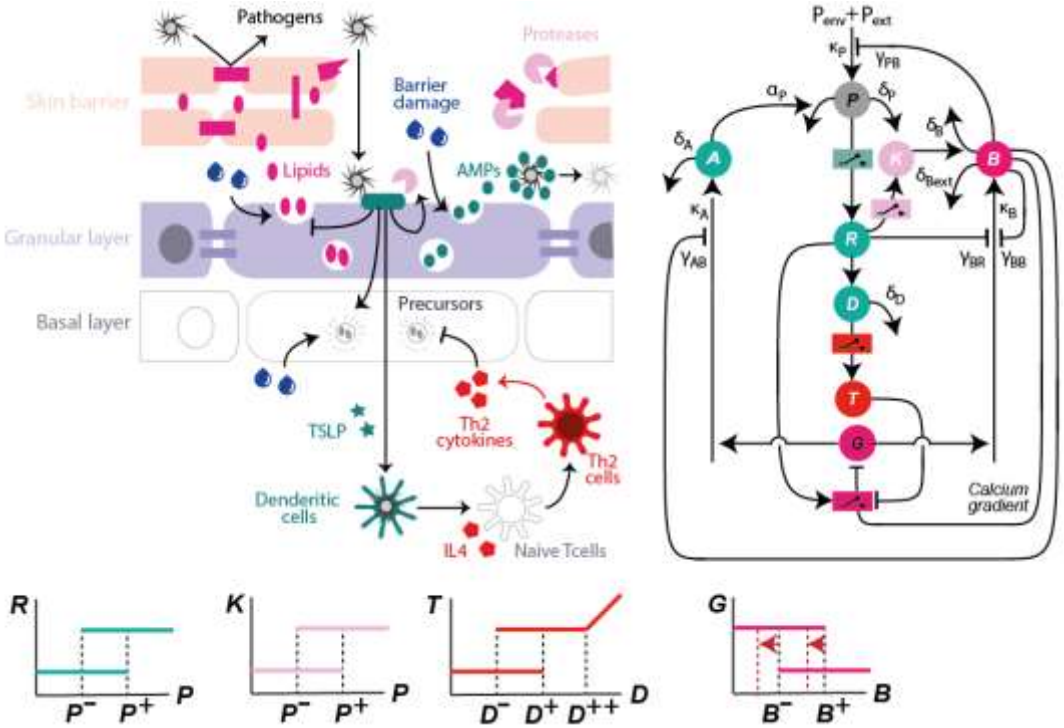


Figure 5.1 Diagram of the triple-switch model. The state variables defined by switch functions are represented by squares, while state variables of the ODE are represented by circles.

5.1.1. RK-switch

The R and K state variables are defined by the R- and K-switches, whose discontinuities occur at the same threshold values, so they effectively act as one discontinuous state variable, K being simply a scaled version of R. For this reason, we generally refer to this switch as the "RK -switch", only separating the "RK-" prefix into the "R-" and "K-" switches in circumstances where our discussion demands we distinguish their respective effects on our analysis.

The RK-switch, representing innate immune responses and kallikrein activation is defined by two piecewise functions $R(t)$ and $K(t)$ in equation 5.2. The RK-switch is activated by the presence of pathogens, which are detected by toll-like-receptors (TLRs) and PAR2 receptors which activate innate immune responses and upregulate kallikrein activity respectively. The K-switch definition emerges from the slightly modified model of kallikrein regulation detailed in Appendix A.

5.1.2. G-switch

There are a couple of peculiar points about the G-switch. First, the threshold values of the G-switch, B^- and B^+ are themselves functions of R and T. Turning on the RK-switch causes the threshold values B^- and B^+ to increase, while turning on the T-switch causes them to decrease. Second, the G-switch has the "off" state higher than the "on" state $G_{\text{off}} > G_{\text{on}}$, while both of the other switches are the other way around $R_{\text{off}} < R_{\text{on}}$, $K_{\text{off}} < K_{\text{on}}$ and $T_{\text{off}} < T_{\text{on}}$. The G-switch is defined by equation 5.5, and results from the modelling in Chapter 4.

5.1.3. T-switch

The T-switch results from the modelling discussed in Chapter 3, and can be defined by either equation 5.3 or 5.4 depending on which of the two models of Th2 regulation we consider, as we discussed in section 3.4.

Our model for chronic AD is structured in such a way that it is insensitive to which of these two candidate models we chose. The T-switch is coupled to the rest of the system not through any effect on the smooth dynamics produced by the ODE (the discrete-variable T doesn't appear in the ODEs at all), but instead by changing the threshold values of the G-switch. We don't use any of the details of the piecewise definitions 5.3 or 5.4 for the T-switch, but rather only care about whether there is an active Th2 response or not.

For this reason, the T-switch can effectively be simplified to

$$T(D^-, D^+) = \begin{cases} 0 \to T_{\text{on}} & \text{if} \quad D(t) = D^+ \text{ and } \dot{D}(t) > 0 \\ T_{\text{on}} \to 0 & \text{if} \quad D(t) = D^- \text{ and } \dot{D}(t) < 0 \end{cases} \tag{5.6}$$

without loss of detail. We can now simply say the "T_{on}" value represents an active Th2 response, and we don't need to quantify it. The T-switch definitions are both reversible bistable switches, with an upper threshold D^+ and lower threshold D^- as shown in Figure 5.1.

Since the T-switch can only be activated by the RK-switch, we must assume that the steady state of D when the RK-switch is activated is greater than the activation threshold D^+ of the T-switch, i.e.

$$\frac{R_{\text{on}}\kappa_D}{\delta_D} > D^+ \tag{5.7}$$

so that the RK-switch is able to cause the T-switch to turn on.

Conversely, we must also assume that when the RK-switch is off the steady state of D is less than the deactivation threshold, i.e.

$$\frac{R_{\text{off}}\kappa_D}{\delta_D} < D^- \tag{5.8}$$

so that the RK-switch is able to turn off the T-switch.

These two conditions ensure that it is possible to activate and deactivate the T-switch respectively.

5.1.4. Mechanisms included in the triple-switch model

The triple-switch model considers the mechanisms described below. The construction of this ODE draws heavily on the previous "double-switch model" work by Domínguez-Hüttinger et al. [18], with a few modifications chosen to introduce features we regard as important to chronic AD. Where new assumptions have resulted in changes or additional features, we clearly describe what we change relative to the previous work.

5.1.4.1. Infiltrated pathogen load P

The infiltrated pathogen load P is increased by the penetration of external stressors P_{env} at a constant rate, and we also provide a variable input $P_{\text{ext}}(t)$. The penetration of these pathogens is inhibited at rate γ_{PB} by the barrier B which has permeability κ_P, as in the double-switch model and described as

$$\left(\mathrm{P_{env}} + \mathrm{P_{ext}(t)}\right)\frac{\kappa_{\mathrm{P}}}{1 + \gamma_{\mathrm{PB}} \cdot \mathrm{B(t)}}$$

The infiltrated pathogens are killed by AMPs at rate α_P and die naturally at rate δ_P. These are summarised as

$$-\alpha_{\mathrm{P}} \cdot \mathrm{A(t)} \cdot \mathrm{P(t)} - \delta_{\mathrm{P}} \cdot \mathrm{P(t)}$$

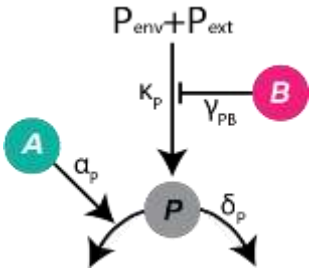


Figure 5.2 Mechanisms of pathogen proliferation and killing.

These mechanisms are illustrated in Figure 5.2.

5.1.4.2. Barrier integrity B

The strength of barrier integrity B, or the permeability barrier of stratum corneum, is determined by its structural components such as filaggrin, lipids and corneodesmosomes.

Stratum corneum (the uppermost layer of the skin barrier) is produced by lamellar body secretion of barrier components by keratinocytes in the granular layer G at rate κ_B. It is this layer that physically repels abrasive and chemical assaults, and prevents penetration of most pathogens, so the integrity of this layer is our definition for the state variable B.

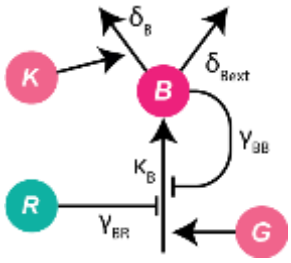


Figure 5.3 Mechanisms of barrier formation and degradation.

The barrier-related mechanisms we consider in our ODE are illustrated in Figure 5.3. Each arrow on the diagram is explained below.

R inhibits production of B

Innate immune activation R inhibits secretion of lipids into the barrier [87] at rate γ_{BR}. For example, [88] show that corneocytes treated with PAR2 agonist show sequestration of incomplete secretion of lamellar bodies; furthermore, PAR2 knock-out leads to enhanced LB secretion in mice [88]. Along the same lines, ([89], Fig. 8) show that acidification of the epidermis, which leads to increased kallikrein and hence PAR2 activity, leads to accumulation of un-secreted lamellar bodies in the stratum granulosum of hapten (Ox)-induced AD mice.

B inhibits production of B

This mechanism appeared in the double-witch model in a very different form, being incorporated as a "phenomenological" self-limiting rate of barrier production where the production rate decreases linearly in strength as barrier B increases. Here, we incorporate this mechanism with more rigorous justification, instead of using a Michaelis-Menten kinetics type term $\frac{1}{1+B}$.

We assume cornification is inhibited by $B(t)$ itself, at rate γ_{BB}, since barrier disruption results in an immediate reduction in the calcium concentration surrounding the cells in the granular layer, triggering lamellar body secretion [90]. For instance, Fig1 in [91] shows that LB secretion occurs immediately after tape stripping. Similarly [92] shows that barrier perturbation with either acetone or tape stripping leads to an immediate rapid secretion of LB contents from the uppermost granular cell layer. Damage to the skin barrier leads to an impairment of the epidermal calcium gradient, specifically to a rise in extracellular calcium levels in the basal layer. This signal can be mimicked experimentally by *in vitro* calcium switch experiments [93].

Taken together, these two inhibitory effects of B and R on B production, and promotion of barrier production by calcium signalling are modelled by

$$\kappa_{\mathrm{B}} \frac{1}{1+\gamma_{\mathrm{BB}} B(t)} \frac{1}{1+\gamma_{\mathrm{BR}} R(\mathrm{P}^{-}, \mathrm{P}^{+})} G(t) \tag{5.9}$$

Noteworthy, this term builds up barrier integrity B only, and not the state of differentiation of keratinocytes in the granular layer (which we represent by the G state variable). Additionally, we must note that we assume that the release of barrier precursors from G does not affect the levels of G because G simply tells us the state of differentiation of keratinocytes, which are able to replace the secreted liposomes at a sufficient rate that they don't run out. Further, this term (multiplicative in G) reflects that, since G is a precursor for B, it is necessary for B production.

K causes degradation of B

This mechanism appeared in the double-switch model in exactly the same form. In our triple-switch model, we have modified the K-switch (see Appendix A), but the form of this term in the ODE is identical.

The barrier is degraded as a result of desquamation mediated by active kallikreins $K(\mathrm{P}^-, \mathrm{P}^+)$ at rate δ_{B}. The barrier is also damaged due to scratching, which is accelerated by itch-scratch cycle with $T(\mathrm{t})$ as direct neuromodulators [94]. However, we do not include this process (from T to B degradation) explicitly in our current model, as the effects of T on B are already included by inhibition of barrier precursors (G) by T. Scratching or physical damage can also lead to barrier degradation, which we represent by the forcing function $-\delta_{\mathrm{Bext}}(t)B(t)$.

The rate of degradation of barrier is therefore summarised by

$$-\delta_{\mathrm{B}} \cdot \mathrm{K}(\mathrm{P}^-, \mathrm{P}^+) \cdot \mathrm{B}(\mathrm{t}) - \delta_{\mathrm{Bext}}(\mathrm{t}) \cdot \mathrm{B}(\mathrm{t}) \tag{5.10}$$

5.1.4.3. Antimicrobial peptides $\boldsymbol{A}$

In the double-switch model, the mechanism of pathogen clearance by AMPs was included only indirectly in the double-switch model as pathogens were killed by innate immune responses R but did not have its own state variable in the ODE. In our new "triple-switch model", we introduce the AMPs as a state variable of their own, allowing AMP dynamics to play out separately from barrier integrity, coupled through our ODE system. These mechanisms are illustrated in Figure 5.4.

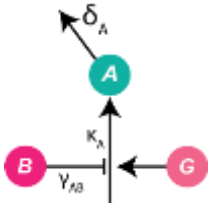


Figure 5.4 Mechanisms of AMP production and degradation.

Here, we assume that AMPs A are derived from AMP precursors G at rate κ_{A} is induced by barrier damage (i.e. inhibited by the barrier B) at rate γ_{AB} . For example, ([95], Fig. 1) show an transient increase of CRAMP and mBD3 peaking at 8 hours after barrier damage, and ([96], Fig. 4b) show that RNase 7 secretion is increased 1h post barrier disruption; this effect is reduced and delayed if skin is occluded.

We model this rate of AMP production and inhibition by barrier by

$$\kappa_{\mathrm{A}} \frac{1}{1 + \gamma_{\mathrm{AB}} B(t)} G(t) \tag{5.11}$$

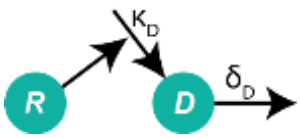
Note that this term in rate of change of AMPs reflects our assumption that the release of AMPs from G does not affect the levels of G (because G reflects state of differentiation of keratinocytes rather than some quantity that use used up). Further, this term (multiplicative in G) reflects that since G is a precursor for AMP, it is necessary for AMP.

[97] suggest that AMP release is inhibited by IL4 and IL13 cytokines produced by Th2 cells T. However, we do not include this regulation in the model since there is no further experimental evidence for T-interference with AMP release. Further, the effects of T on A are also already included by inhibition of AMP precursors (G) by T.

We assume AMPs degrade at rate δ_{A}.

5.1.4.4. Dendritic cells $\boldsymbol{D}$

The concentration of dendritic cells ($D(t)$) in draining lymph-nodes increase when inflammation (R) persists, but DCs degrade naturally.



Activation of innate immune responses TLR receptors induces production of thymic stromal lymphopoietin (TSLP) signalling molecule, which activates dendritic cells at rate κ_D [98]. TSLP-activated dendritic cells then carry antigens from the epidermis through the lymphatic system to the draining lymph-node [99], [100].

Figure 5.5 Dendritic cells are produced in response to innate immune response R at rate κ_D and degrade at rate δ_D.

In draining lymph-nodes, dendritic cells D present antigens to naïve T cells, causing them to differentiate into Th2 cells through stimulation of the OX40 ligand at rate k_D [26], [101] in the presence of IL4 (Th2 cytokines, T) [102]. Dendritic cells are also assumed to decay naturally at rate δ_D. These mechanisms are illustrated in the schematic in Figure 5.5.

Further, IL13 (Th2 cytokines) accelerates the rate of migration of DCs to the draining lymph-nodes [28] and degradation of DCs is impaired by IL4 (another Th2 cytokine) [103]. Within the T-switch model (see section 3.1), similar positive feedback mechanisms involving these Th2 cytokines are already considered, albeit not connected to dendritic cell migration or degradation. We do not include these proposed migration and degradation of DCs affects in the ODE for D since we assume the effects of such positive feedback to be already adequately represented in the T-switch model.

5.1.5. Parameters and units

Our model consists of two parts: an ODE of four continuous-valued state variables (P, B, A, D), and three switches. These three switches produce four discrete-valued state variables (R, K, T, G). We might conclude that there are eight state variables in total but for the complication of the RK-switch. We have already described how the R and K states are related through simple proportionality, so there are in effect only seven state variables in total rather than eight. The units of these variables are defined in Table 5.1.

Table 5.1 State variables of the triple-switch model

Model variables		
Symbol	**Name**	**Units**
$\boldsymbol{t}$	time	Days
$\boldsymbol{P(t)}$	Infiltrated pathogen	mg/ml
$\boldsymbol{B(t)}$	Skin barrier function	dimensionless
$\boldsymbol{A(t)}$	Antimicrobial peptides	mg/ml
$\boldsymbol{D(t)}$	Dendritic cells in lymph nodes	cells/ml
$\boldsymbol{R(t)}$	Activity of innate immune receptors	dimensionless
$\boldsymbol{K(t)}$	Activity of kallikreins	dimensionless
$\boldsymbol{G(t)}$	Level of differentiation of granular layer	dimensionless
$\boldsymbol{T(t)}$	Th2 inflammation	dimensionless

The model has 28 parameters in total, 18 of which describe the ODE, and the other 14 describe the discontinuous switches (see Table 5.2). There are also two parameters which we use to suggest locations in the ODE where external perturbations (such as scratching or increased pathogen load) could be considered, but we do not make use of these in our analysis.

Table 5.2 Nominal parameter values and ranges considered in our analysis.

Parameters				
Symbol	**Name**	**Nominal value**	**Permissible range**	**Units**
$\boldsymbol{P_{\mathrm{env}}}$	Levels of environmental pathogen load	1	0.1 - 10	mg/ml

Parameters

Symbol	Name	Nominal value	Permissible range	Units
$\boldsymbol{\kappa_P}$	Nominal skin barrier permeability	0 - 35 (bifurcation range)	0 - 35	
$\boldsymbol{\gamma_{PB}}$	Strength of inhibition of pathogen infiltration by barrier	0.1	0.01 - 1	$1/[B]$
$\boldsymbol{\alpha_P}$	Killing efficacy of AMPs for *S. aureus*	0 to 35 (bifurcation range)	0 - 35	$1([A][t])$
$\boldsymbol{\delta_P}$	Natural pathogen death rate	1	0.1 - 10	mg/(ml day)
$\boldsymbol{\kappa_B}$	Nominal rate of secretion of skin barrier precursors during cornification / LB secretion	0 to 1 (bifurcation range)	0 - 1	$1/[t]$
$\boldsymbol{\gamma_{BB}}$	Strength of inhibition of LB secretion/cornification (barrier components) by barrier	1	0.1 - 10	$1/[B]$
$\boldsymbol{\gamma_{BR}}$	Strength of inhibition of cornification/LB secretion (barrier components) by innate immune responses	1	0.1 - 10	$1/[R]$
$\boldsymbol{\delta_B}$	Barrier degradation rate (by active kallikreins)	0 - 1 (bifurcation range)	0 - 1	$\frac{1}{[t][K]}$
$\boldsymbol{\kappa_B}$	Nominal rate of secretion of skin barrier precursors during cornification / LB secretion	0 - 1 (bifurcation range)	0 - 1	$1/[t]$
$\boldsymbol{\gamma_{BB}}$	Strength of inhibition of LB secretion/cornification (barrier components) by barrier	1	0.1 - 10	$1/[B]$
$\boldsymbol{\gamma_{BR}}$	Strength of inhibition of cornification/LB secretion (barrier components) by innate immune responses	1	0.1 - 10	$1/[R]$
$\boldsymbol{\delta_B}$	Barrier degradation rate (by active kallikreins)	0 - 1 (bifurcation range)	0 - 1	$\frac{1}{[t][K]}$
$\boldsymbol{\kappa_A}$	Nominal rate of secretion of AMPs during cornification / LB secretion	1	0.1 - 10	$1/[t]$
$\boldsymbol{\gamma_{AB}}$	Strength of inhibition of AMP release by barrier	10	1 - 100	$1/[B]$
$\boldsymbol{\delta_A}$	Natural degradation of AMPs	1	0.1 - 10	$1/[t]$
$\boldsymbol{\kappa_D}$	Rate of Dendritic cells recruitment by innate immune responses	1	0.1 - 10	$\frac{[D]}{[t][R]}$
$\boldsymbol{\delta_D}$	Natural decay of dendritic cells	1	0.1 - 10	$1/[t]$
Parameters defining external perturbations				
$\boldsymbol{P_{ex}(t)}$	Fluctuations in environmental pathogen load	0	0	$[P]$

Parameters

Symbol	Name	Nominal value	Permissible range	Units
$\boldsymbol{\delta_{\mathrm{Bext}}(t)}$	Barrier degradation due to external perturbations such as scratching, tape stripping, or use of soaps and detergents.	0	0	$1/[t]$
Parameters defining the discontinuities of the switches				
$\boldsymbol{R_{off}}$	Level of innate immune response when healthy	0	0 – 0.1	$[R]$
$\boldsymbol{R_{on}}$	Level of innate immune response when reacting to pathogens	1	0.1 - 10	$[R]$
$\boldsymbol{K_{off}}$	Concentration of active kallikreins when healthy	1	0.1 – 1.5	$[K]$
$\boldsymbol{K_{on}}$	Concentration of active kallikreins when reacting to pathogens	2	1.5 - 10	$[K]$
$\boldsymbol{T_{off}}$	Strength of Th2 response when T-cells are not being activated by DCs	0	-	$[T]$
$\boldsymbol{T_{on}}$	Strength of Th2 response when DCs are activating naïve T-cells	1	0.1 - 10	$[T]$
$\boldsymbol{G_{off}}$	Level of expression of terminal differentiation markers of poorly differentiated keratinocytes	2	1.5 - 10	$[G]$
$\boldsymbol{G_{on}}$	Level of expression of terminal differentiation markers of differentiated keratinocytes	1	0.1 – 1.5	$[G]$
$\boldsymbol{P^{-}}$	Pathogen load below which innate immune responses become deactivated	0.2		$[P]$
$\boldsymbol{P^{+}}$	Pathogen load above which innate immune responses become activated	0.8		$[P]$
$\boldsymbol{D^{-}}$	Minimum number of DCs required to sustain a Th2 response	0.2		$[D]$
$\boldsymbol{D^{+}}$	Minimum number of DCs required to activate a Th2 response	0.8		$[D]$
$\boldsymbol{B^{-}}$	Minimum barrier damage to trigger expression of terminal differentiation markers	$\diagdown$ | R_{off} | R_{on} T_{off} | 0.2 | 0.3 T_{on} | 0.1 | 0.2		$[B]$
$\boldsymbol{B^{+}}$	Minimum barrier damage to sustain expression of terminal differentiation markers	$\diagdown$ | R_{off} | R_{on} T_{off} | 0.8 | 0.9 T_{on} | 0.7 | 0.8		$[B]$

5.2. Bifurcation analysis

As a first step of the model analysis for the triple-switch model, we conduct a bifurcation analysis. In the triple-switch model, the three switches represent the symptoms of AD, while the continuous-

valued variables of the ODE represent the latent properties of the disease which go largely unnoticed except possibly in specific clinical settings. The parameter values in our model correspond to the genotype of a patient. Two patients with slightly different genotypes can thus be represented by two different points in the parameter space. We wish to identify how the clinical phenotype emerges as a result of the mechanisms we include in our model, and how this is affected by genotype. For this reason, our bifurcation analysis is concerned with identifying how the dynamics of the switches (i.e., symptoms) are affected by the parameter values (i.e., genotype).

In the previous analysis of the double-switch model, the dynamics were inferred by calculating bifurcation diagrams summarising the behaviour of the RK-switch, of which four types were possible; recovery (R), bistability (B), oscillation (Os) and damage (D). Those four categories were sufficient to fully describe the double-switch model's dynamics because the dynamics of the RK-switch were unaffected by the state of the T-switch. It made sense to say, for example, that the model was behaving in an "oscillatory" manner because the RK-switch state would oscillate on and off regardless of the state of the T-switch.

By contrast, the triple-switch model presents us with more complicated interactions between the switches. Crosstalk between the three switches results in the behaviour of one switch being affected by the state of another switch. This in turn leads to situations where the transition of one switch can induce a transition of another, and so on in potentially long sequences.

Toy example to demonstrate sequential switch transitions due to crosstalk.

Imagine a very simple hybrid dynamical system featuring just two switches. Each switch can be either off (represented by a "0" in Table 5.3) or on (represented by a "1"). This means there are four possible configurations of the switches. Assume that we know the behaviour of both switches at every one of these configurations (Table 5.3). The "behaviour" defines where the focal point value of the system lies in relation to the switch's threshold value, which can be used to infer how that switch will change state in the future. We call this table a "configuration-behaviour table" since it tells us how the configuration of the switches map to their behaviour.

Table 5.3 Configuration-behaviour table of our toy example.

	Configuration		***Behaviour***	
	Switch 1	*Switch 2*	*Switch 1*	*Switch 2*
row 1	*0*	*0*	*R*	*D*
row 2	*0*	*1*	*D*	*D*
row 3	*1*	*1*	*D*	*R*
row 4	*1*	*0*	*R*	*R*

*Given the information in this configuration-behaviour table, we can work out how the crosstalk between the two switches will cause a sequence of transitions. In row 1, the switch 1 is off and in the **R** (recovery) behaviour while switch 2 is off and in the **D** (damage) behaviour. From this, we can deduce that switch 2 will turn on because the **D** behaviour means the focal point value is above the switching on threshold for switch 2. Thus, to reach the focal point value, the model's dynamics must result in switch 2 turning on. This process was explained in greater detail in section 2.3.2. This transitions us to row 2.*

*Now we're in row 2 of the configuration-behaviour table (*Table 5.3*), where switch 1 is off and in the **D** behaviour while switch 2 is on and in the **D** behaviour. From this, we can deduce that switch 1 will turn on because switch 1's **D** behaviour means the focal point value is above the switch on threshold for switch 1. To reach the focal point value, the model's dynamics cause switch 1 to turn on. This transitions us to row 3.*

*Now we're in in row 3 where switch 1 is on and in the **D** behaviour while switch 2 is on and in the **R** behaviour. We can deduce that switch 2 will turn off because the **R** behaviour means the focal point value is below the turning off threshold for switch 2. To reach the focal point value, the model's dynamics cause switch 2 to turn off. This transitions us to row 4.*

*Now we're in in row 4 where switch 1 is on and in the **R** behaviour while switch 2 is off and in the **R** behaviour. We can deduce that switch 1 will turn off because the **R** behaviour means the focal point value is below the turning off threshold for switch 1. To reach the focal point value, the model's dynamics cause switch 1 to turn off. This transitions us back to row 1, where the whole process repeats*

This sequence of transitions is illustrated in Figure 5.6. The resulting dynamic is cyclical in nature, despite the behaviours of each of the switches never having the "oscillation" behaviour.

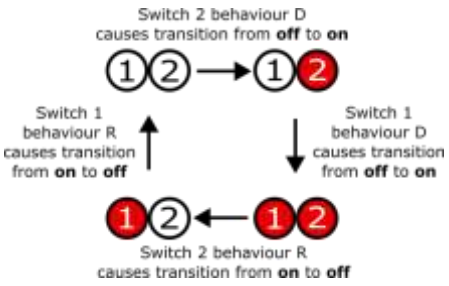


Figure 5.6 Transitions of each switch in our toy example.

In this toy example, we wrote the configuration-behaviour table which summarises the crosstalk between the switches and used this information to infer the dynamics of the switches. This demonstrates how we can summarise the crosstalk between the switches by constructing this configuration-behaviour table, and how that table can be used to infer the switch transition dynamics.

For the remainder of this chapter, we construct bifurcation diagrams with the aim of defining the configuration-behaviour table for all points in the parameter ranges we consider. We develop a series of analytical methods which allow us to search the parameter space extensively, identifying all the possible configuration-behaviour tables, and defining the regions in the parameter space which gives rise to each configuration-behaviour table. Finally, we carry out a sensitivity analysis to assess the regions that are affected by changes to the parameter values to assess the robustness of our result.

The first step of this process is to identify parameters of particular interest to us and consider how the three switches change behaviour as we vary these interesting parameters. We identify these parameters of interest as κ_P, α_P, κ_B and δ_B, which represent particularly interesting genetic predisposing risk factors for AD. Their ranges are defined in Table 5.2.

κ_p	High infiltration of pathogens, κ_P could occur when the filaggrin gene (Flg) is mutated, a known genetic risk factor for AD [14].
α_P	Rate of pathogen killing α_P could represent dysregulated expression of antimicrobial peptides [96], [104]–[106].
κ_B	Impaired calcium sensing or damaged basal skin layer decreases this parameter.
δ_B	Rate of desquamation. This can be increased by scratching, increased KLK expression, high pH or impaired expression of LEKTI (see Appendix A for the effect of high pH on KLK expression).

By selecting these four parameters, we are effectively reducing our domain of interest from the full parameter space to just the 4D space over which these four vary.

We consider variation of all the remaining parameters (all parameters in Table 5.2 except for κ_P, α_P and κ_B, δ_B) by exhaustively enumerating how the bifurcation diagrams can be laid out as these remaining parameters are varied, and ensuring we sample the parameter space for each possibility we identify. We are able to show analytically that there are only four variants of bifurcation diagrams produced by our model as we vary the remaining parameters.

We proceed with our analysis in 5 stages detailed in Table 5.4. Every stage of our analysis derives results analytically rather than numerically. It is possible to numerically integrate the ODEs of our model, simulating the dynamics and switch transitions for any specified parameter values and initial conditions. We lay out a method to achieve this kind of simulation in Appendix M. However, it is preferable to first use analytical tools to get an overall idea of the possible qualitative behaviours the model can display. An analytical approach means we are sensitive to potentially subtle model behaviours like multi-stability, which may not be readily apparent when simply simulating the system. Analytical tools also allow us to identify the reasons for changes in model behaviour in response to parameter and initial condition variation.

By the end of this chapter, we do obtain a list of parameter regions, each associated with a configuration-behaviour table. The next chapter, Chapter 6, focusses on taking the information in these configuration-behaviour tables and turning it into dynamics we can interpret in terms of disease symptoms.

Table 5.4 Five stages of bifurcation analysis

Section number and title	Summary of analysis steps
5.2.1 Focal points and stability of the triple switch model	Calculate the focal points of the ODE and assess their stability.
5.2.2 Boundary-equilibrium bifurcations	Use our focal point results to calculate the bifurcation diagrams for each of the switches considered in isolation, temporarily neglecting the effects of interactions with the other switches.
5.2.3 Connectivity and presence of the	Consider how variation of parameters not already on the bifurcation diagram axes can affect the behaviours represented on the diagram.
5.2.4 Crosstalk between switches	Describe how crosstalk between the switches affects the dynamics of the model.
5.2.6 Parameter sensitivity of	Carry out a sensitivity analysis, calculating the rate of change of the size of the parameter regions as the parameter values of our model vary. All the bifurcation diagrams we draw consist of polygon-shaped regions within which the switches had a particular behaviour. With this sensitivity analysis technique, we are able to evaluate how rapidly the size of these polygon regions changes as we vary the model's parameters.

5.2.1. Focal points and stability of the triple switch model

Before proceeding, it is worth noting that we define "focal points" as equilibrium values of our ODE which are globally stable. They may therefore be understood as the value which the continuous-valued variables would reach if the ODE dynamics were allowed to proceed while all of the switches are held in their current state. However, our hybrid dynamical system does not necessarily reach the focal points which we calculate in this section because the ODE dynamics sometimes cause one of the three switches to transition from one state to another, thereby interrupting the ODE dynamics by changing the focal point value the system is heading towards. The implications of these interrupted dynamics will be discussed extensively in chapter 6.

A similar term, used later in this thesis, is "absorbing state". "Focal point" values refer to the continuous-valued part of the model (i.e. the ODE) only. The "absorbing state", in contrast, will be used to refer to the discrete-valued part of this model (i.e. the switches).

When we calculate the value of the focal points, we must consider that they are dependent not only on the parameters of the ODE, but also on the value of the discrete-valued state variables R, K, T and G (this dependence represented by the red arrows in Figure 5.7), each of which can take one of two possible values depending on the state of the switches. The R and K state variables effectively act as one, since they are both defined by a single switch (the "RK-switch"). Thus, rather than four degrees of freedom, we have just three. This leaves us with 2^3=8 configurations of the switches, so we can see that there must be eight focal point values, one for each of these configurations.

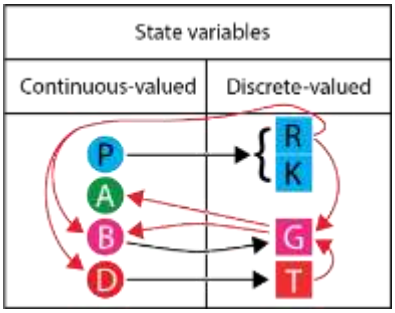


Figure 5.7 Relationship between the continuous-valued and discrete-valued state variables.

To help us distinguish between the focal point values we calculate for each configuration, we denote each of the focal points as a function with the three arguments representing the RK-switch, T-switch, and G-switch in that order. A "0" represents the switch state being off, "1" represents on, and "*" is a wildcard meaning either off or on (used to represent situations where the state is not relevant). For example, to express the steady state of the state variable P when the RK-switch is off, the T-switch either on or off, and the G-switch off, we write $\bar{P}(1,*,0)$. For the steady state of the variable B with the RK-switch in either state, the T-swich in either state, and the G-witch off, we write $\bar{B}(*,*,0)$. Some examples of this notation are shown in

Figure 5.8. This form of notation will become useful in the next subsection as we begin to discuss bifurcations caused by the discontinuous switches.

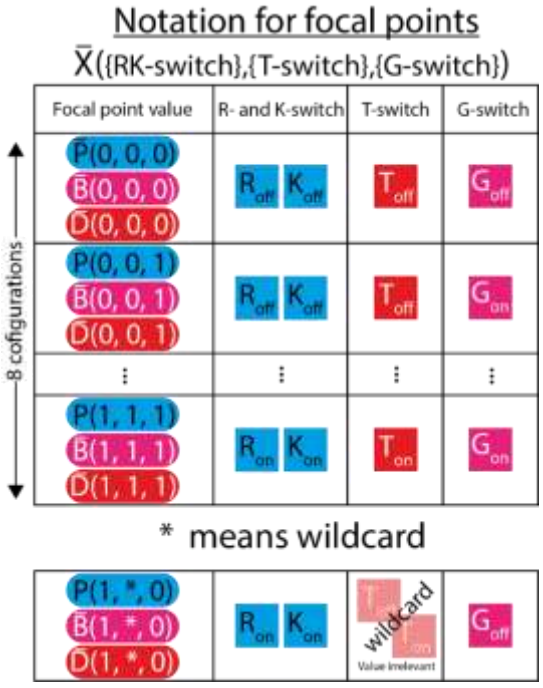


Figure 5.8 Notation describing the focal points as a function of the configuration of the three switches.

The focal points of the triple-switch ODE (equation 5.1) are written as follows.

$$
\begin{aligned}
\bar{P}(R,T,G) &= \frac{(P_{env}+\overline{P_{ext}})\cdot\kappa_P}{(\alpha_P\cdot\bar{A}(R,T,G)+\delta_P)\cdot(1+\gamma_{PB}\cdot\bar{B}(R,T,G))} \\
\bar{A}(R,T,G) &= \frac{\kappa_A}{\delta_A}\cdot\frac{G}{1+\gamma_{AB}\cdot\bar{B}(R,T,G)} \\
\bar{B}(R,T,G) &= \frac{-1+\sqrt{1+4\cdot\gamma_{BB}\cdot\frac{\kappa_B}{(\delta_B\cdot K+\overline{\delta_{Bext}})}\cdot\frac{G}{(1+\gamma_{BR}\cdot R)}}}{2\cdot\gamma_{BB}} \\
\bar{D}(R,T,G) &= \frac{\kappa_D\cdot R+\overline{D_{ext}}}{\delta_D}
\end{aligned}
\tag{5.12}
$$

The analytical steps to find these expressions and prove that the focal point value is always stable are found in Appendix B.

In the ODEs of the triple-switch model (equation 5.1), we can see that none of the expressions for the rates of change have any dependence on the state of the T-switch (denoted by the letter T). As a consequence, none of the focal point values depend on T either. This is because we assume the state of the T-switch to affect only the threshold values of the G-switch rather than act on the terms of the ODE. For this reason, the second argument of our equilibrium value expressions (the position representing T) is usually filled by the wildcard *, acknowledging that it has no impact on the value of the focal point. For the case of the switching conditions for the T-switch, this second argument is filled by either a 0 or 1 representing the current T-switch state, which is relevant to determining if the switch will change state. There is also a dependence on both R and T in our expression for the threshold values of the G-switch, so we explicitly acknowledge this dependence by writing $B^-(R,T)$ and $B^+(R,T)$.

5.2.2. Boundary-equilibrium bifurcations

The switching conditions in Table 5.5 define the regions of the parameter space where the switches turn on or off. These inequalities are obtained by examining the definitions of the switches, and identifying when the focal point value of the ODE crosses the threshold values of each switch. When the "turn off condition" is satisfied, it means that we can guarantee the switch in question will turn off because the ODE dynamics must cross the switching off threshold on the way to the focal point value. Conversely, if the "turn on condition" is satisfied, it means we can guarantee that the switch in question will turn on because the ODE dynamics must cross the switching on threshold on the way to the focal point value.

Table 5.5 Transition conditions: conditions where we can infer the switch must change state in the future.

Switch	**Threshold values of the switch**	**Switching conditions**	
		Region where the switch turns from **off** to **on** **"Turn on condition"**	Region where the switch turns from **on** to **off** **"Turn off condition"**
RK-switch	P^-, P^+	$\bar{P}(0,*,G) > P^+$	$\bar{P}(1,*,G) < P^-$
T-switch	D^-, D^+	$\bar{D}(R,0,*) > D^+$	$\bar{D}(R,1,*) < D^-$
G-switch	$B^-(R,T)$, $B^+(R,T)$ (dependent on the RK- and T-switch state)	$\bar{B}(R,*,0) > B^+(R,T)$	$\bar{B}(R,*,1) < B^-(R,T)$

We examine the two transitions conditions for each switch, drawing bifurcation diagrams showing the region of the parameter place in which each condition is satisfied. For any arbitrary point in the

parameter space, we can work out the switch's behaviour by checking these two conditions, resulting in four possible behaviours. The conditions under which each of the four behaviours arise are summarised in Table 5.6. This inference about the behaviour applies to each of the three switches based in their specific expressions for the "turn off condition" and "turn on condition".

Table 5.6 Switch behaviours as a result of the turning on and turning off conditions.

		Turn off condition	
		Satisfied	**No satisfied**
Turn on condition	**Satisfied**	Oscillation (Os)	Switch remains on (D)
	Not satisfied	Switch remains off (R)	Bistable (B)

- Recovery "R". In this case, the turning off condition is satisfied but the turning on condition is not satisfied. So the switch turns off and stays off.
- Bistable "B". In this case, neither transition condition is satisfied, so the switch remains in the current state.
- Oscillating "Os". In this case, both the switching off and switching on conditions are satisfied, causing the switch to transition from off to on and then off again in sequence forever.
- Damage "D". In this case, the turning on condition is satisfied but the turning off condition is not satisfied. So the switch turns on, and stays on.

Having identified that these four behaviours should be possible for each switch, we can turn our attention to identifying where in the parameter space each behaviour will occur.

The boundaries of these regions are where the focal point of a state variable collides with, in other words equates to, the threshold value of one of the switches, and are called **boundary-equilibrium bifurcations** since they represent the boundary where the equilibrium value collides with the threshold value. The boundary-equilibrium bifurcation curves separating these regions are thus defined by the equations in Table 5.7. These represent the critical points in the parameter space where the focal point value of the ODE collides with the threshold value of each switch. These lines separate the four behaviours R, B, Os and D we have just identified.

Table 5.7 Boundary-equilibrium bifurcation curves defining the edges of the regions satisfying the transition conditions defined in Table 5.5.

Switch	Edge of the switching off region	Edge of the switching on region
RK-switch	$\bar{P}(0,*,G) = P^+$	$\bar{P}(1,*,G) = P^-$
T-switch	$\bar{D}(R,0,*) = D^+$	$\bar{D}(R,1,*) = D^-$
G-switch	$\bar{B}(R,*,0) = B^+(R,T)$	$\bar{B}(R,*,1) = B^-(R,T)$

When we plot the regions separated by the BE bifurcation lines defined in Table 5.7, we denote which side of the boundary line satisfies the switching condition by including a coloured arrow, red for "turn off" and blue for "turn on". This allows us to visually denote which regions of the plane correspond to each of the four possible switch behaviours described in Table 5.5.

5.2.2.1. RK-switch behaviour bifurcations

We calculate the boundary-equilibrium bifurcation lines for switching on and off the RK-switch in κ_P, α_P space since both of these parameters directly affect the pathogen level P which is the continuous-valued variable input to the RK-switch.

The RK-switch turns from off to on when $\bar{P}(0,*,G) > P^+$. Substituting the expression for the focal point values from 5.12 gives us

$$\frac{(P_{\mathrm{env}} + \overline{P_{\mathrm{ext}}}) \cdot \kappa_{\mathrm{P}}}{(\alpha_P \cdot \bar{A}(0,*,G) + \delta_P)(1 + \gamma_{\mathrm{PB}} \cdot \bar{B}(0,*,\mathrm{G}))} > P^{+} \tag{5.13}$$

Rearranging the inequality to isolate α_P leads to

$$\alpha_P < \frac{(P_{env} + \overline{P_{ext}})}{P^{+} \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(0,*,G)\right) \cdot \bar{A}(0,*,\mathrm{G})} \cdot \kappa_P - \frac{\delta_P}{\bar{A}(0,*,\mathrm{G})} \tag{5.14}$$

where $\bar{B}(0,*,G) = \frac{-1+\sqrt{1+4\cdot\gamma_{BB}\frac{\kappa_B}{(\delta_B\cdot K_{\mathrm{off}}+\delta_{Bext})}\cdot\frac{\mathrm{G}}{(1+\gamma_{BR}R_{\mathrm{off}})}}}{2\cdot\gamma_{BB}}$ and $\bar{A}(0,*,G) = \frac{\kappa_A}{\delta_A}\cdot\frac{G}{1+\gamma_{AB}\cdot\bar{B}(0,*,\mathrm{G})}$.

The RK-switch turns from on to off when $\bar{P}(1,*,G) < P^{-}$. By the same process of substitution and rearranging, we can define this condition with inequality

$$\alpha_P > \frac{(P_{env} + \overline{P_{ext}})}{P^{-} \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(1,*,G)\right) \cdot \bar{A}(1,*,G)} \cdot \kappa_P - \frac{\delta_P}{\bar{A}(1,*,G)} \tag{5.15}$$

where $\bar{B}(1,*,G) = \frac{-1+\sqrt{1+4\cdot\gamma_{BB}\frac{\kappa_B}{(\delta_B\cdot K_{\mathrm{on}}+\delta_{Bext})}\cdot\frac{G}{(1+\gamma_{BR}\cdot R_{\mathrm{on}})}}}{2\cdot\gamma_{BB}}$, and $\bar{A}(1,*,G) = \frac{\kappa_A}{\delta_A}\cdot\frac{G}{1+\gamma_{AB}\cdot\bar{B}(1,*,G)}$

These inequalities take the form of first-order polynomials (i.e. a straight line) with respect to κ_P, with (5.14) taking the form $\alpha_P < m_{on} \cdot \kappa_P + c_{on}$, and (5.15) the form $\alpha_P > m_{off} \cdot \kappa_P + c_{off}$, where m_{on} and m_{off} are the gradients, and c_{on} and c_{off} the y-intercepts of the inequalities in the κ_P, α_P plane. We can therefore expect straight lines in the κ_P, α_P plane as shown in Figure 5.9.

The gradient and y-intercept of these BE-bifurcation lines are summarised as

$$\begin{aligned} m_{\mathrm{on}} &= \frac{(P_{env} + \overline{P_{ext}})}{P^{+} \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(0,*,G)\right) \cdot \bar{A}(0,*,\mathrm{G})} \\ m_{\mathrm{off}} &= \frac{(P_{env} + \overline{P_{ext}})}{P^{-} \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(1,*,G)\right) \cdot \bar{A}(1,*,G)} \\ c_{\mathrm{on}} &= -\frac{\delta_P}{\bar{A}(0,*,\mathrm{G})} \\ c_{\mathrm{off}} &= -\frac{\delta_P}{\bar{A}(1,*,\mathrm{G})} \end{aligned} \tag{5.16}$$

Since all of our parameters take positive values by assumption, and all of the focal point values are also positive-valued as a result, we can see from equation 5.16 that c_{on} and c_{off} must both take negative values. Further, the gradients m_{on} and m_{off} must also both take positive values since both numerator and denominator are products of positive values in equation 5.16, so the BE bifurcation lines must always slope upwards. Taken together, the positive gradient, and negative y-intercept of both the off and on BE-bifurcation lines mean that the BE bifurcation lines must slope upwards and intercept the x-axis. A cartoon example of such an arrangement is shown in Figure 5.9 (a).

The four regions of the plane separated by the example BE bifurcation curves in Figure 5.9 (a) each exhibit one of the four possible switch behaviours since they each satisfy the two inequality conditions in different ways described in Table 5.6. The four behaviours in their respective regions are labelled in

Figure 5.9 (b). It is possible to arrange the BE bifurcation lines in such a way that only three of the behaviours (R, B and D) are produced in the positive quadrant we consider as shown in Figure 5.9 (c). We will address why this can occur and define conditions on the remaining model parameters to define when this alternative layout occurs in section 5.2.3.

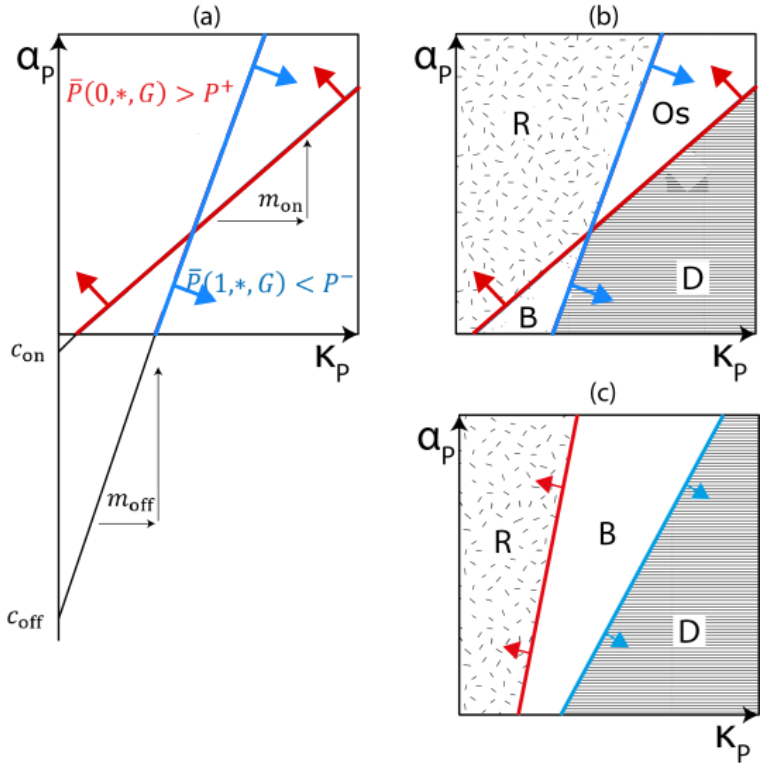


Figure 5.9 Cartoon boundary equilibrium bifurcation curves for the RK-switch in the κ_P, α_P plane.

5.2.2.2. T-switch behaviour bifurcations

The T-switch is also able to exhibit the four behaviours **R**, **B**, **Os** and **D** described in Table 5.6, and we can follow a similar process of calculating the BE bifurcation curves to work out the regions of parameter space where these behaviours occur.

The T-switch state depends on the state variable D (number of DCs in draining lymph-node), and D is coupled only the R-switch state *not* the other variables. This means the dynamics of D depend solely on the dynamics of the RK-switch. This sole dependence can be seen by examining the ODE for D,

$$\frac{dD(t)}{dt} = \kappa_{\mathrm{D}} \cdot R(P^-, P^+) - \delta_{\mathrm{D}} \cdot D(t) \tag{5.17}$$

where the rate of increase of D depends on the state of the switch R, and none of the other variables. We assume for simplicity that the stimulation of DCs by superantigens is zero, $D_{\mathrm{ext}} = 0$.

This means that once we specify the dynamics of the R-switch, we can fully define the behaviour of the T-switch and safely neglect the rest of the system. We can, without loss of generality, assume the R-switch is oscillating with a particular duty cycle defined by t_{off} and t_{on} between the off and on values $R_{\mathrm{off}} = 0$ and R_{on}. With this notional dynamic for the R-switch, we can work out how the T-switch behaves in response to various duty cycles of R-switch dynamics as shown in Figure 5.10. The analytical steps to reach this result are described in Appendix H.

In Figure 5.10 (a), oscillating R-switch dynamics with various duty cycles are represented on the plane. R-switch states which are constant over time are represented along the axes of (a). When the R-switch is constantly on, the resulting behaviours of the T-switch is read off the t_{on} axis and we see that tells us the T-switch is in the D (damage) behaviour meaning that it turns on and remains on. Conversely, when the R-switch is constantly off, we read the T-switch behaviour of the t_{off} axis which tells us the T-switch behaviour is R (recovery) meaning it turns off and remains off. For all the intermediate duty cycles, we can the other behaviours emerge for particular ratios of t_{on} and t_{off}. The Os (oscillation) behaviour only emerges if the duty cycle of the R-switch gives it a low enough frequency to allow the D dynamics to cross both the D^- and D^+ thresholds, which is why the Os region appears on the upper right of the plane. Conversely, if the R-switch frequency is too high (i.e. low t_{off} and t_{on} values), the D dynamics oscillate with a low amplitude between D^- and D^+ causing the T-switch to be bistable.

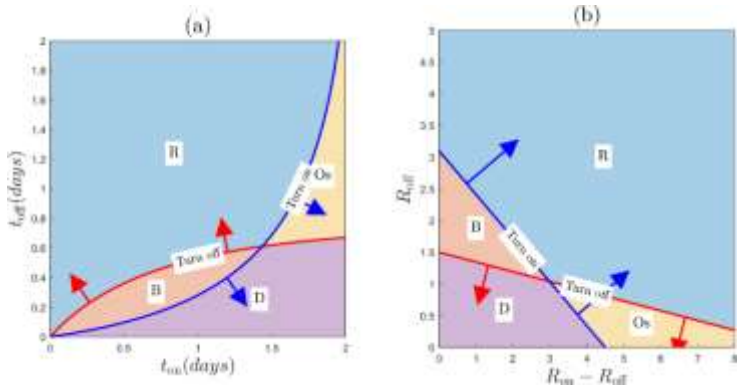


*Figure 5.10 Behaviours of the T-switch as we vary the duty cycle and amplitude of the R-switch. (a) t_{off} and t_{on} or **(b)** R_{off} and $R_{on} - R_{off}$ (chosen so that all points on the diagram satisfy $R_{on} - R_{off} > 0$). Parameters (unless varied on axes): $\kappa_D = 1, \delta_D = 1, t_{off} = 0.5, t_{on} = 0.5, R_{off} = 1, R_{on} = 5, D^- = 2, D^+ = 3.1$*

In this thesis, we simply assume that the state of the T-switch follows that of the R-switch, we means we implicitly assume the T-switch falls in the Os (oscillation) region of Figure 5.10 (a). We do this because we are not usually in the position of knowing the details of the R-switch's duty cycle or frequency. We simply assume it to oscillate slowly enough that the T-switch is able oscillate, turning off and on during each oscillation.

5.2.2.3. G-switch behaviour bifurcations

We calculate the boundary-equilibrium bifurcation lines for switching on and off the G-switch in the (κ_B, δ_B) plane since both of these parameters directly affect the barrier integrity B (i.e. the input state for the G-switch), which is the continuous-valued state variable at the input of the G-switch, and describe the effects of two risk factors. κ_B describes the effects of impaired calcium sensing or damaged basal skin layer δ_B the strength of KLK expression, increased by high pH or impaired expression of LEKTI.

The condition to turn on the G-switch $\bar{B}(*,*,0) > B^+(R,T)$ is written as

$$\frac{-1 + \sqrt{1 + 4\gamma_{BB} \cdot \dfrac{\kappa_B}{\left(\delta_B \cdot K + \overline{\delta_{Bext}}\right)} \cdot \dfrac{G_{\mathrm{off}}}{(1 + \gamma_{BR} \cdot R)}}}{2 \cdot \gamma_{BB}} > B^+(R,T) \tag{5.18}$$

by substituting in the solution for $\bar{B}(*,*,0)$. Rearranging for the variable δ_B, we obtain

$$\delta_B < \frac{4 \cdot \gamma_{BB} \cdot G_{\mathrm{off}}}{((2 \cdot B^+(R,T) \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K} \cdot \kappa_B - \frac{\overline{\delta_{Bext}}}{K} \tag{5.19}$$

Inequality 5.19 takes linear form with respect to κ_B,

$$\delta_B < m_{\mathrm{B\ on}} \cdot \kappa_B + c_{\mathrm{B\ on}} \tag{5.20}$$

where $m_{B\ \mathrm{on}} = \frac{4 \cdot \gamma_{BB} \cdot G_{\mathrm{off}}}{((2 \cdot B^{+} \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K}$ is the gradient, and $c_{B\ \mathrm{on}} = -\frac{\overline{\delta_{Bext}}}{K}$ the y-intercept of the inequality in the κ_B, δ_B plane. We can therefore expect straight lines in the κ_B, δ_B plane.

Similarly, the condition to turn off the G-switch $\bar{B}(*,*,1) < B^{-}$ is described as

$$\delta_B > \frac{4 \cdot \gamma_{BB} \cdot G_{\mathrm{on}}}{((2 \cdot B^{-}(R,T) \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K} \cdot \kappa_B - \frac{\overline{\delta_{Bext}}}{K} \tag{5.21}$$

Inequality (5.21) takes the linear form in

$$\delta_B > m_{\mathrm{B\ off}} \cdot \kappa_B + c_{\mathrm{B\ off}} \tag{5.22}$$

where $m_{B\ \mathrm{off}} = \frac{4 \cdot \gamma_{BB} \cdot G_{\mathrm{on}}}{((2 \cdot B^{-} \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K}$ is the gradient, and $c_{B\ \mathrm{off}} = -\frac{\overline{\delta_{Bext}}}{K}$ the y-intercept of the inequality in the κ_B, δ_B plane. We can therefore expect straight lines in the κ_B, δ_B plane which pass through the origin as shown in Figure 5.11.

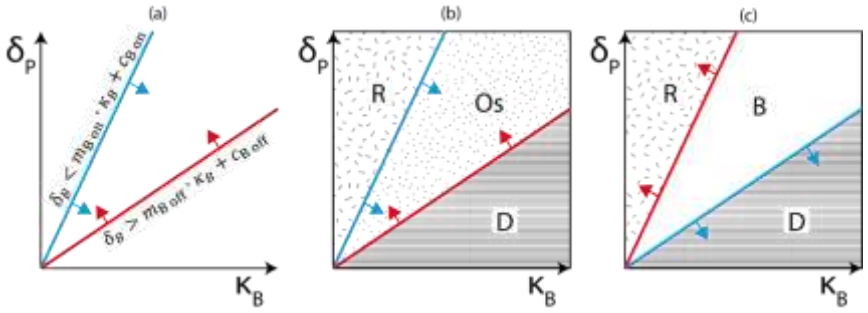


Figure 5.11 Boundary-equilibrium bifurcation curves for the G-switch in the κ_B, δ_B plane

Once again, we can infer the behaviour of the G-switch within the regions of the plane separated by these BE bifurcation lines. In this case, since the intersection point always lies at the origin, we can only produce a maximum of three regions in the positive quadrant. The recovery and damage behaviours must always be present because we know the gradient $m_{B\ \mathrm{on}}$ and $m_{B\ \mathrm{off}}$ are always positive and cannot be infinite. The third behaviour located between the recovery and damage regions is determined by the relationship between the BE line gradients. In the example in Figure 5.11 (a), the behaviours would be laid out as in Figure 5.11 (b). We will show in section 3.3.3 that it is also possible to parameterise the model such that the behaviours shown in Figure 5.11 (c) are produced. We will also define a condition on the parameters which can tell us which of the two variety of diagram occurs.

5.2.2.4. BE-bifurcation conclusion

We have shown above that all of the three switches of the triple-switch model can exhibit four possible behaviours, recovery, bistability, oscillation and damage. We labelled these four behaviours R, B, Os and D respectively. We calculated the analytical forms of the BE bifurcation lines separating these regions and observed that the RK- and G-switches can each produce two varieties of bifurcation diagram where different selections of these behaviours are present across the regions.

We placed some assumptions on the range of the parameters κ_D, δ_D, R_{off} and R_{on} so that the T-switch would behave in a fairly simple way: simply following the state of the T-switch and oscillating

when the T-switch oscillates. This is equivalent to assuming the duty cycle t_{off} and t_{on} of the R-switch oscillations are much larger than the rate of DC decay δ_D.

5.2.3. Connectivity and presence of the behaviour regions for different parameter values

In section 5.2.2 above, we obtained two inequalities to describe how each of the switches turn on and off. The two inequalities for each switch left us with four possible dynamical phenotypes recovery, bistability, oscillation and damage. These behaviours are found by working how the BE bifurcation lines slice up the parameter plane into polygon shaped regions. Sometimes the BE lines slice up the plane so that three regions are produced, but sometimes we obtain four regions, or different behaviours are present. For example, the G-switch bifurcation diagram shown in Figure 5.12 (a) has no "B" (bistable) region and Figure 5.12 (b) has no "Os" (oscillation) region.

We can define the difference between these variants of bifurcation diagram more rigorously by drawing a "connectivity network" for each bifurcation diagram we produce. The nodes of the network represent the behaviours present on the bifurcation diagram. The nodes are linked if those behaviour regions share an edge on the diagram. Two such examples are shown at the bottom of Figure 5.12. Two bifurcation diagrams are considered different if their connectivity networks differ. This test is only sensitive to bifurcation diagrams which are fundamentally different in the sense that it is impossible to apply a transformation to the axes which maps one to the other. For example Figure 5.12 (a) cannot be stretched or squeezed to look like Figure 5.12 (b) because it contains different behaviour regions.

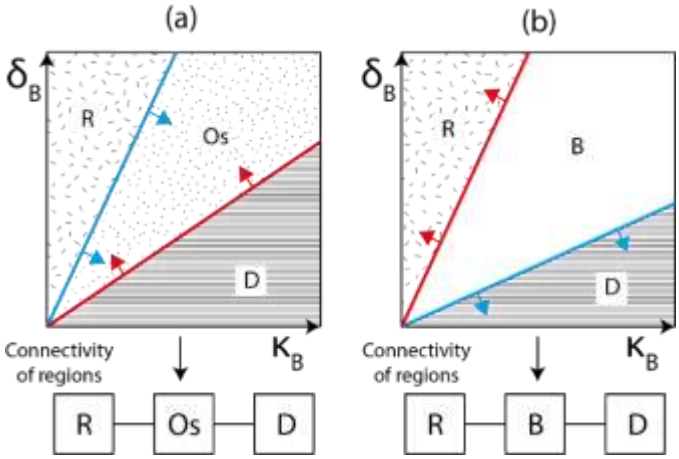


Figure 5.12 Example bifurcation diagrams where the connectivity of behaviour regions differs.

Each variety of connectivity network can lead to a different combination of behaviours from the switches of the mode, which in turn can lead to different switch dynamics. We therefore need to find all the variations of connectivity network that can be produced, and how much of the parameter space of the model results in each possibility. To achieve this, we work out how the parameter value of the model affect the connectivity of the regions.

We first examine the bifurcation diagrams to deduce which connectivity networks are possible given the geometry of the BE lines. Then we create such a method to work out when each of these connectivity networks is produced and write down conditions on our parameters for each case.

5.2.3.1. Connectivity analysis of the RK-switch behaviours

Given that the BE bifurcation curves for the RK-switch behaviours are straight lines in the κ_P, α_P plane, and that the gradients of these are always positive and the α_P-intercepts always negative, there are only four ways they can be arranged in the positive quadrant (Figure 5.13).

For the RK-switch, the two BE bifurcation lines may be defined as

$$\alpha_P = m_{\text{on}}\kappa_P + c_{\text{on}} \qquad \text{Edge of turning on condition.} \qquad (5.23)$$

$$\alpha_P = m_{\text{off}}\kappa_P + c_{\text{off}} \qquad \text{Edge of turning off condition.} \qquad (5.24)$$

Where the gradients and y-intercepts of these lines are given by equations 5.25.

$$m_{\text{on}} = \frac{(P_{env} + \overline{P_{ext}})}{P^{+}\left(1 + \gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)} \qquad c_{\text{on}} = -\frac{\delta_P}{\bar{A}(0,*,*)}$$

$$m_{\text{off}} = \frac{(P_{env} + \overline{P_{ext}})}{P^{-}\left(1 + \gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*)} \qquad c_{\text{off}} = -\frac{\delta_P}{\bar{A}(1,*,*)} \qquad (5.25)$$

These two BE bifurcation lines intersect in the κ_P, α_P plane. Figure 5.13 shows the only four ways in which two nonparallel straight lines with positive gradient and negative y-intercept can divide up the positive quadrant.

We have laid out the four scenarios shown in Figure 5.13 such that they can be defined by whether (1) the intercept, $(\widehat{\kappa_P}, \widehat{\alpha_P})$, of the two BE bifurcation lines is in the positive quadrant and (2) the gradient of the "turn on" BE bifurcation line is greater than the "turn off" BE bifurcation line. By working out when each of these conditions are met, we can write an inequality defining the parameter ranges where each type of bifurcation diagram is produced. We tackle each of these conditions in turn below.

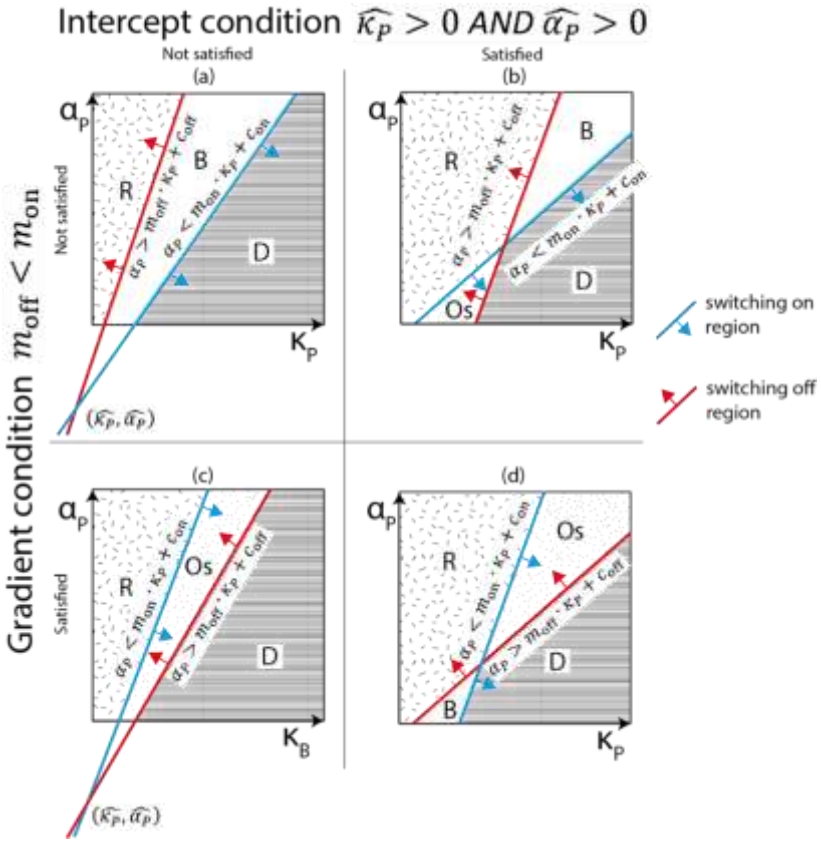


Figure 5.13 The four ways to divide up the plane while both BE bifurcation lines maintain positive gradient: (a) The oscillation phenotype can be abolished entirely making it occur only for negative parameters (i.e. inadmissible values) but (b) recovered if the intersection point is in the positive plane. (c) The bistable phenotype can be abolished in the same way if the relative gradients of the lines are reversed but (d) recovered if the intersection point moves into the positive plane.

The intercept conditions.

The intercept $(\widehat{\kappa_P}, \widehat{\alpha_P})$ of the two BE bifurcation lines may be expressed as

$$\begin{aligned}\widehat{\kappa_P} &= \frac{c_{\text{off}} - c_{\text{on}}}{m_{\text{on}} - m_{\text{off}}} \\ &= \frac{\delta_P}{(P_{env} + \overline{P_{ext}})}\left(\frac{P^- P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\left(\bar{A}(1,*,*) - \bar{A}(0,*,*)\right)}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)}\right)\end{aligned} \tag{5.26}$$

$$\begin{aligned}\widehat{\alpha_P} &= \frac{c_{\text{off}} m_{\text{on}} - c_{\text{on}} m_{\text{off}}}{m_{\text{on}} - m_{\text{off}}} \\ &= \delta_P\left(\frac{P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right) - P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)}\right)\end{aligned} \tag{5.27}$$

The intercept $(\widehat{\kappa_P}, \widehat{\alpha_P})$ of the two lines is in the positive quadrant if $\widehat{\kappa_P} > 0$ and $\widehat{\alpha_P} > 0$ simultaneously. This scenario is shown in Figure 5.13 (b) and (d). These scenarios may be written as conditions 5.28 and 5.29, which should be satisfied simultaneously.

Definition of the intercept conditions:

$$\widehat{\kappa_P} = \frac{\delta_P}{(P_{env} + \overline{P_{ext}})}\left[\frac{P^- P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\left(\bar{A}(1,*,*) - \bar{A}(0,*,*)\right)}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)}\right] > 0 \tag{5.28}$$

$$\widehat{\alpha_P} = \delta_P\left[\frac{P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right) - P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)}\right] > 0 \tag{5.29}$$

The gradient condition

The gradient of the "turn on" BE bifurcation line m_{on} (sketched in blue in Figure 5.13) is greater than the gradient of the "turn off" BE bifurcation line m_{off} (sketched in red in Figure 5.13) if $m_{\text{off}} < m_{\text{on}}$. Substituting in the expressions for each gradient (equations 5.25), we obtain equation 5.30.

$$\frac{(P_{env} + \overline{P_{ext}})}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*)} < \frac{(P_{env} + \overline{P_{ext}})}{P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)} \tag{5.30}$$

Bringing the denominators up, and comparing them to 0, we obtain the gradient condition 5.31.

Definition of the gradient condition:

$$P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*) > 0 \tag{5.31}$$

Comparing the intercept and gradient conditions, we find that they reduce to a single inequality condition.

If the gradient condition 5.31 is satisfied, the RK-switch bifurcation diagram will take the form sketched in Figure 5.13 (c) or (d), while if it is not satisfied, it will resemble either (a) or (b). If the intercept conditions 5.28 and 5.29 are both satisfied simultaneously, the bifurcation diagram will resemble the form sketched in Figure 5.13 (b) or (d), while if they are not satisfied it will resemble (a) and (c).

Here we show that the conditions 5.28, 5.29 and 5.31 are all equivalent to a single inequality condition which is satisfied for part of our admissible parameter range, but not satisfied for the rest. Thus, variation of the remaining parameters leads to just two of the outcomes we drew in Figure 5.13. The behaviours available in each of these scenarios are highlighted in green and blue in Table 5.8. The white cells are not possible to achieve for positive parameter values.

Table 5.8 Only two possible bifurcation diagram layouts exist for the RK-switch behaviours. The letters represent the behaviours present in the positive (κ_P, α_P) plane in each bifurcation diagram. Only the green cell and the blue cell are possible. The white cells are impossible.

	Intercept condition $\widehat{\kappa_P} > 0$ and $\widehat{\alpha_P} > 0$		
Gradient condition $m_{\text{off}} < m_{\text{on}}$		**Not satisfied**	**Satisfied**
	Not satisfied	R, B, D	R, B, Os, D
	Satisfied	R, Os, D	R, B, Os, D

To prove that conditions 5.28, 5.29 and 5.31 are all equivalent, firstly we note that the numerator of condition 5.28 is positive only if $\bar{A}(1,*,*) > \bar{A}(0,*,*)$. Barrier level B is always lower when the RK-switch is on (this can be seen from the equilibrium value expression in equation B.7 (reproduced below) and noting that $R_{\text{off}} < R_{\text{on}}$:

$$\begin{array}{ccc} \dfrac{-1+\sqrt{1+4\cdot\gamma_{BB}\cdot\dfrac{\kappa_B}{(\delta_B\cdot K+\overline{\delta_{Bext}})}\cdot\dfrac{G}{(1+\gamma_{BR}\cdot R_{\text{on}})}}}{2\cdot\gamma_{BB}} & < & \dfrac{-1+\sqrt{1+4\cdot\gamma_{BB}\cdot\dfrac{\kappa_B}{(\delta_B\cdot K+\overline{\delta_{Bext}})}\cdot\dfrac{G}{(1+\gamma_{BR}\cdot R_{\text{off}})}}}{2\cdot\gamma_{BB}} \\ B(1,*,*) & < & B(0,*,*) \end{array} \quad (5.32)$$

The barrier B appears in the denominator of the expression for the steady state of AMP expression A in equation 5.12 (and reproduced below in equation 5.33). This reverses the direction of the inequality for the AMP steady state, which means the AMP level with the RK-switch on is always greater than when the RK-switch is off as summarised in equation 5.33 below. Again, this makes intuitive sense because AMP production is enhanced by activation of the innate immune receptors (the R-switch).

$$\begin{array}{ccc} \dfrac{\kappa_A}{\delta_A}\cdot\dfrac{G}{1+\gamma_{AB}\cdot\bar{B}(1,*,*)} & > & \dfrac{\kappa_A}{\delta_A}\cdot\dfrac{G}{1+\gamma_{AB}\cdot\bar{B}(0,*,*)} \\ \bar{A}(1,*,*) & > & \bar{A}(0,*,*) \end{array} \quad (5.33)$$

in turn guarantees that the numerator of condition 5.28 is always positive.

Further, since $P^+ > P^-$ by design and, as aforementioned in equation 5.32, $B(1,*,*) < B(0,*,*)$, we can also guarantee that the numerator of condition 5.29,

$$P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right) - P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right) \quad (5.34)$$

is always positive since the first term is always larger than the second.

Having established that the numerators of conditions 5.28 and 5.29 are both positive, we can conclude that these conditions can only be violated if the denominator becomes negative. Therefore, we can simplify conditions 5.28 and 5.29 to just their denominator, which is common to both. Now both conditions are summarised by the new condition 5.35.

RK-switch bifurcation diagram condition:

Not satisfied → {R, B, D}

Satisfied → {R, B, Os, D}

$$P^{-}\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*)-P^{+}\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)>0 \quad (5.35)$$

This denominator condition 5.35 is identical to the gradient condition 5.31. Thus, we may conclude that conditions 5.28, 5.29 and 5.31 must all be satisfied simultaneously or violated simultaneously.

This means that for positive valued ODE parameters, and RK-switch parameters which satisfy our assumptions $P^{+}>P^{-}$ and $R_{\mathrm{off}}<R_{\mathrm{on}}$, it is not ever possible to produce RK-switch bifurcation diagrams of the form shown in Figure 5.13 (b) or (c), they can only take the form in Figure 5.13 (a) or (d). If 5.35 is satisfied, the bifurcation diagram will resemble Figure 5.13 (d). If condition 5.35 is *not* satisfied, the bifurcation diagram will resemble Figure 5.13 (a).

Randomly sampling the parameter space with unform distributions within our specified plausible ranges (see Table 5.2 for parameter ranges) for each parameter reveals that condition 5.35 can indeed by violated in our plausible parameter value ranges. This means that, for positive-valued parameters, it is possible to produce bifurcation both diagrams like Figure 5.13 (a) and Figure 5.13 (d). In section 5.2.3.4 below, we combine this results with a similar result for the G-switch to assess how much of the admissible parameter space produces each form of bifurcation diagram.

5.2.3.2. Connectivity analysis of the T-switch behaviours

The T-switch is something of a special case in this step of our analysis since its behaviour depends only on the parameter values κ_D and δ_D, the T-switch parameters D^{-}, D^{+}, and R-switch values R_{off} and R_{on} which we have already considered in our analysis in section 5.2.2.2 above. Thus, we know that the remaining parameters of the model have no direct effect on the behaviours of the T-switch available in the positive quadrant of the κ_D, δ_D plane except through changing the state of the RK-switch as described in section 5.2.2.2. We therefore skip this analysis step for the T-switch as we have already fully described how this switch behaves.

5.2.3.3. Connectivity analysis of the G-switch behaviours

In section 5.2.2.3, we derived the bifurcation lines for the G-switch in the κ_B, δ_B plane, showing that it is only possible to see at most three of the behaviours for a given parameterisation of the model. These three behaviours always include R (recovery) and D (damage), with the third behaviour being either B (bistable), or Os (oscillation) as illustrated in Figure 5.11. The obvious question arises; which of these two potential arrangements of the bifurcation diagrams is produced for admissible parameterisations of our model? To find out, we can examine the expressions we derived for the BE bifurcation lines of the G-switch behaviours. These were equations 5.19 and 5.21.

These BE bifurcation lines are defined as straight lines of the form:

$$\delta_B = m_{\mathrm{B\,on}}\kappa_B + c_{\mathrm{B\,on}} \qquad \text{(edge of parameter region where the G-switch turns on)}$$

$$\delta_B = m_{\mathrm{B\,off}}\kappa_B + c_{\mathrm{B\,off}} \qquad \text{(edge of parameter region where the G-switch turns off)}$$

Where the gradients and δ_B-intercepts are

$$m_{B\,\mathrm{on}} = \frac{4\cdot\gamma_{BB}\cdot G_{\mathrm{off}}}{((2\cdot B^{+}\cdot\gamma_{BB}+1)^2-1)\cdot(1+\gamma_{BR}\cdot R)\cdot K} \qquad c_{B\,\mathrm{on}} = -\frac{\overline{\delta_{Bext}}}{K}$$

$$m_{B\,\mathrm{off}} = \frac{4\cdot\gamma_{BB}\cdot G_{\mathrm{on}}}{((2\cdot B^{-}\cdot\gamma_{BB}+1)^2-1)\cdot(1+\gamma_{BR}\cdot R)\cdot K} \qquad c_{B\,\mathrm{off}} = -\frac{\overline{\delta_{Bext}}}{K} \quad (5.36)$$

We assume in this analysis that the concentration of superantigens $\overline{\delta_{Bext}}$ is zero (more complicated situations where this parameter is considered may be developed in future based on the results we produce here).

In this case, the y-intercept of both lines is zero because we assume $\overline{\delta_{Bext}} = 0$, so both lines always meet at the origin. We are left with only the relationship between the gradients to produce the different layouts of bifurcation diagrams. We can simply compare the "turns on" gradient with the "turns off" gradient to work out if the bistable or oscillation region is present in addition to recovery and damage,

$$m_{B\,\mathrm{off}} > m_{B\,\mathrm{on}} \tag{5.37}$$

Substituting our expressions for these gradients, we may rewrite the condition 5.37 as

$$\frac{4 \cdot \gamma_{BB} \cdot G_{\mathrm{on}}}{((2 \cdot B^{-} \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K} - \frac{4 \cdot \gamma_{BB} \cdot G_{\mathrm{off}}}{((2 \cdot B^{+} \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K} > 0 \tag{5.38}$$

Which may be simplified to condition 5.39 below.

G-switch bifurcation diagram condition: Not satisfied → {R, Os, D} Satisfied → {R, B, D}	
$\frac{G_{\mathrm{on}}}{((2 \cdot B^{-} \cdot \gamma_{BB} + 1)^2 - 1)} - \frac{G_{\mathrm{off}}}{((2 \cdot B^{+} \cdot \gamma_{BB} + 1)^2 - 1)} > 0$	(5.39)

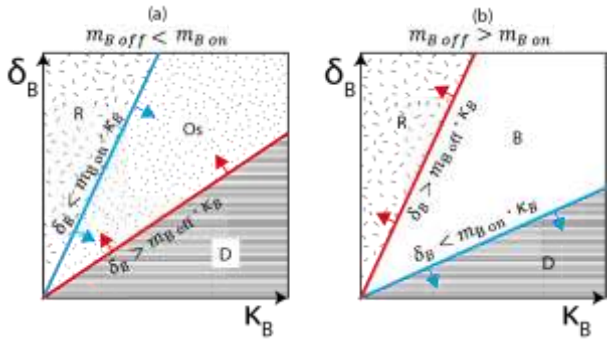


Figure 5.14 Possible arrangements of the BE bifurcation diagrams for the G-switch. (a) when condition 5.37 is not satisfied, and (b) when 5.37 is satisfied.

In each of these cases, only three behaviours are possible by varying κ_B and δ_B. These are summarised in Table 5.9.

Table 5.9 Two possible scenarios for bifurcation diagram layouts arise from the gradient condition. The letters represent the behaviours present in the positive (κ_B, δ_B) plane.

Gradient condition $m_{B\,\mathrm{off}} > m_{B\,\mathrm{on}}$	Behaviours produced
Not satisfied	**R, Os, D**
Satisfied	**R, B, D**

In Figure 5.14 (a), the gradient condition 5.39 is not satisfied and so the oscillation phenotype is produced in the region where both the switching on and switching off condition are satisfied. In Figure 5.14 (b), the opposite occurs because the gradient condition 5.39 is satisfied, causing a bistable region to be produced. Thus, it is not possible to produce both the oscillation and bistable regions simultaneously.

5.2.3.4. Enumerating the possible bifurcation layouts for the whole model

Having summarised the possible bifurcation diagram layouts for the RK-, T- and G-switches, we can combine the information from each to deduce the overall bifurcation layouts possible for the model.

We found that the RK-switch has two varieties of bifurcation diagram, and also that the G-switch has two varieties of bifurcation diagram. In section 5.2.2.2, we calculated the bifurcation diagram for the T-switch, but limited the behaviours it can exhibit by applying conditions on the parameter values for $\kappa_B, \delta_D, D^-, D^+, R_{\text{off}}$ and R_{on} and further assuming the oscillations of the R-switch occur on a timescale longer than δ_D. Thus, the T-switch bifurcation diagram is a moot point and can be safely ignored in the context of analysing its connectivity.

The parameters of the model can be selected such that the varieties of bifurcation diagram for the RK- and G-switch may be formed in any combination. Four combinations of bifurcation diagrams emerge, listed in Table 5.10.

Table 5.10 Enumerated possibilities for parameter regions with different behavioural regions shown in the bifurcation diagrams for the RK- and G-switches.

1000 samples of parameter region specified in Table 5.2.		RK-switch behaviours available	
		R, B, D	**R, B, Os, D**
G-switch available	**R, B, D**	Possibility 1: 138 13.8%	Possibility 3: 486 48.6%
	R, Os, D	Possibility 2: 72 7.2%	Possibility 4: 304 30.4%

We can work out roughly how much of the admissible parameter range gives rise to each combination by randomly sampling the parameter space using uniform distributions ranging from the minimum and maximum values we specified in Table 5.2. For each random parameter set, we check the conditions (5.35) and (5.39) we worked out for the RK- and G-switch bifurcation diagrams respectively. Using this sampling regime to generate 1000 parameter sets, we produced the distribution of behaviour-varieties shown in Figure 5.15. This approach allowed us to find parameter sets corresponding to each of the possibilities listed in table J.1 (found in Appendix J).

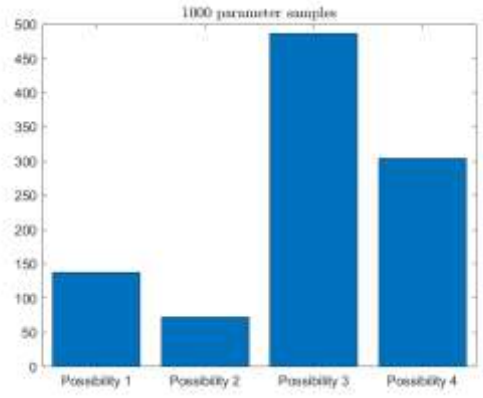


Figure 5.15 Number of bifurcation diagrams of each possible layout.
Produced by sampling the parameter ranges specified in Table 5.2 with a uniform distribution.

Our analysis in this section has demonstrated that these four possibilities cover the whole range of bifurcation diagram layouts that it is possible to produce using only positive-valued parameters. Throughout the coming sections, we analyse each of these four parameterisations, calling them P1, P2, P3 and P4 respectively.

5.2.4. Crosstalk between switches

So far, we have considered the bifurcation diagrams for each switch in isolation without considering how the state of the other switches affects the location of the BE bifurcation lines. However, the variables R, K and G all appear in the expression for the focal points (equations B.7). This means that the focal points depend on the states of the RK- and G-switches.

Given that the focal points are dependent on the switch state, we can also expect the BE bifurcation lines to be dependent on the switches because their equations are defined using the focal point values. In addition, the BE bifurcation lines also include the threshold values of the switches. In the case of the G-switch, the threshold values B^- and B^+ are themselves functions of the T-switch (they decrease when the T-switch turns on and increase when the RK-switch turns on), so the G-switch bifurcation diagram is dependent on the RK- and T-switch states. These effects on the BE bifurcation lines are summarised in Table 5.12.

So, why is this interesting to us? The BE bifurcation lines define the regions of the parameter space where the switches behave in particular ways. If these BE bifurcation lines are themselves dependent on switch state, that means there is a relationship between switch state and switch behaviour.

The switch behaviours cause the switches to change state, so there's a mechanism whereby changes in one switch's state can bring about a change in the behaviour of another, which in turn causes that other switch to change state. This process may occur repeatedly, one switch transition leading to another, and then another and so on in a sequence. These sequences of switch transitions constitute the dynamics of our triple-switch model. Our objective is to characterise all of the various forms these dynamics can take, and the corresponding parameter values for each dynamic. The first step toward this objective is to work out all the various ways that switch state and switch behaviour can be related, and this is what we discuss in this section.

Method to define the behaviour-state relationship.

In our notation, we acknowledge this dependence on the state of other switches by specifying the switch states as function arguments in our formulae. For example, the "turning off" BE bifurcation line for the RK-switch are written as $\bar{P}(1,*,G) = P^-$, where the argument G tells us that the focal point value $\bar{P}$ depends on the state of the G-switch. The "wildcard" $*$ symbol tells us that the T-switch has no effect on the BE line. All of the BE bifurcation lines are summarised in Table 5.11.

Table 5.11 BE bifurcation lines for all the switches' behaviours.

Switch	Edge of the switching off region	Edge of the switching on region
RK-switch	$\bar{P}(0,*,G) = P^+$	$\bar{P}(1,*,G) = P^-$
T-switch	$\bar{D}(R,0,*) = D^+$	$\bar{D}(R,1,*) = D^-$
G-switch	$\bar{B}(R,*,0) = B^+(R,T)$	$\bar{B}(R,*,1) = B^-(R,T)$

Note that the bifurcation lines for the G-switch show the threshold values B^+ and B^- are functions of R and T. This shows us that the G-switch bifurcation diagrams depend on the RK-switch state and the T-switch state. We summarise all of the dependencies of the RK- and G-switch behaviours in Table 5.12.

Throughout this thesis, we make the simplifying assumption that T-switch state follows the R-switch state as described in section 5.2.3.2 above. This boils down to "if the R-switch is off the T-switch will turn off soon afterwards, and likewise if the R-switch is on the T-switch will turn on soon afterwards". For this reason, we do not include the T-switch in our analysis of behaviour-state dependence in this section. It is sufficient for us to just look at the RK- and G-switches.

Now we're going to show these effects on our bifurcation diagram. To achieve this, we're going to deal with the G-switch bifurcation diagrams first, and then move onto the RK-switch bifurcation diagrams.

Table 5.12 Dependence of the BE bifurcation diagrams on the switch states.

Bifurcation diagram parameters	Switch whose behaviour is shown	Switch dependence of the focal point expression	Switch dependence of the threshold values	Switches on which the bifurcation diagram depends
κ_P, α_P	RK-switch	$\bar{P}$ depends on G (the state of the G-switch)	No dependence	G-switch
κ_B, δ_B	G-switch	$\bar{B}$ depends on R, K and G (the state of the RK- and G-switches)	B^- and B^+ depend on R and T (the state of the RK- and T-switches)	RK-switch, T-switch

5.2.4.1. G-switch behaviour-state mapping

In the case of the G-switch, we must consider the effects of the RK-switch and the T-switch. We have repeated the BE bifurcation lines for the G-switch behaviours below in 5.40 and 5.41. We've also reproduced the parameter values for B^- and B^+ as Table 5.13 to show how these values are affected by changes of the RK- and T-switches.

Turn on BE bifurcation line for the G-switch

$$\delta_B = \frac{4 \cdot \gamma_{BB} \cdot G_{\text{off}}}{((2 \cdot B^+(R,T) \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K} \cdot \kappa_B - \frac{\overline{\delta_{Bext}}}{K} \tag{5.40}$$

Turn off BE bifurcation line for the G-switch

$$\delta_B = \frac{4 \cdot \gamma_{BB} \cdot G_{\text{on}}}{((2 \cdot B^-(R,T) \cdot \gamma_{BB} + 1)^2 - 1) \cdot (1 + \gamma_{BR} \cdot R) \cdot K} \cdot \kappa_B - \frac{\overline{\delta_{Bext}}}{K} \tag{5.41}$$

Table 5.13 Effect of the R- and T-switches on the threshold values of the G-switch, B^- and B^+.

$\boldsymbol{B^-(R,T)}$		R-switch	
		off	on
T-switch	off	0.2	0.3
	on	0.1	0.2

$\boldsymbol{B^+(R,T)}$		R-switch	
		off	on
T-switch	off	0.8	0.9
	on	0.7	0.8

Turning on the RK-switch causes the values of B^- and B^+ to increase, while the switch values R and K also increase (since $R_{on} > R_{off}$ and $K_{on} > K_{off}$) both appear in the denominator of equations for the bifurcation lines in 5.40 and 5.41. Thus, turning on the RK-switch cases the gradient of these lines to decrease. Turning on the T-switch causes the values of B^- and B^+ to decrease, which has the effect of increasing the gradient of the bifurcation lines in 5.40 and 5.41.

Now we've worked out what effect the changing switch states will have on the gradients of the BE bifurcation lines; we can plot the bifurcation diagrams under each of the four combinations of switch states for the R- and T-switch. For this discussion, we'll use the P1 parameterisation listed in our standard parameter sets (see Appendix J) and neglect to show the results for the other

parameterisations in this chapter for brevity. The results from the P2, P3 and P4 parameterisations may be found in Appendix K.

First, we plot the bifurcation diagram for the G-switch behaviour under each of the four combinations of R- and T-switch states. These four are shown in Figure 5.16 (a) through (d).

Note how, in Figure 5.16, the BE bifurcation lines in the "R-switch on" column have lower gradient than those in the "R-switch off" column. Similarly, the BE bifurcation lines in the "T-switch on" column have higher gradient than those in the "T-switch off" column as we predicted.

Having plotted these four diagrams, we may next consider them all in combination, drawing every BE bifurcation line from the other diagrams Figure 5.16 (a) through (d) onto one graph, they produce diagram (e). Diagram (e) can be thought of as the result of merging diagrams (a) through (d), finding all of the overlapping regions, and assigning each a number. For this parameterisation, we find nine regions in the "overlaps" diagram (e).

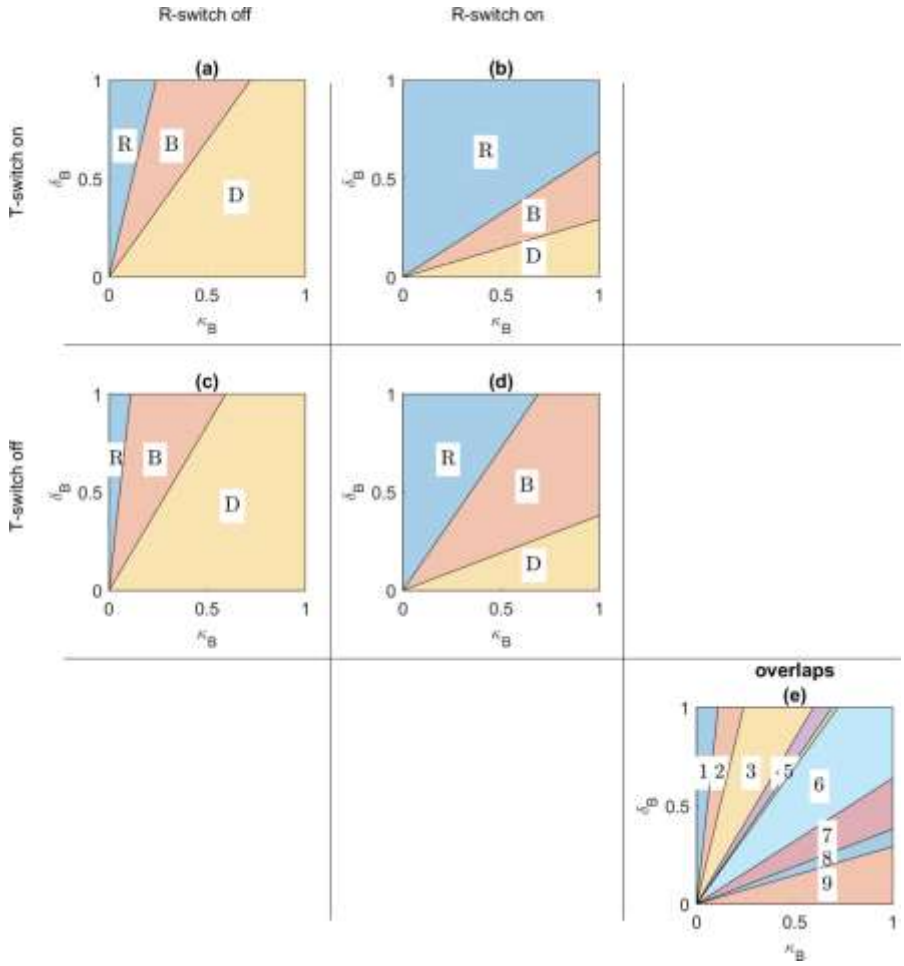


Figure 5.16 G-switch bifurcation diagrams with the R-switch and T-switch in different states.

- *(a) shows the bifurcation diagram with the R-switch off and T-switch off,*
- *(b) R-switch on and T-switch off,*
- *(c) R-switch off and T-switch on,*
- *(d) R-switch on and T-switch on,*
- *(e) shows the overlaps of all the other bifurcation diagrams a through d.*

Within each of the regions in Figure 5.16 (e) each of the four combinations of R- and T-switch states gives us a particular G-switch behaviour. Each such region can be summarised by a table like Table

5.14, where each row in the "G-switch behaviour" column is filled in with the behaviour of the G-switch when the R- and T-switches are in the corresponding states.

Table 5.14 Table summarising mapping of R- and T-switch state to G-switch behaviour

R-switch state	T-switch state	G-switch behaviour
0	0	R
0	1	B
1	0	R
1	1	R

In Figure 5.16 (e), each region produces a different table. In other words, the mapping from R- and T-switch state to G-switch behaviour is one-to-one, and no two regions have the same mapping. The process of separating the parameter plane into the regions in Figure 5.16 (e) by overlapping all of the other bifurcation diagrams Figure 5.16 (a-d) has resulted an exhaustive list of the various mappings from R- and T-switch state to G-switch behaviour which exist for the parameter values in the range of the axes κ_B and δ_B.

Table 5.16 summarises this mapping for every region shown in Figure 5.16 (e). We have highlighted in green each of the seven regions which brings about a change in the G-switch behaviour as a result of a change in one or both of the switch states. These changes in behaviour will become very important in driving the model's dynamics as we discuss in Chapter 6.

The bifurcation diagrams we drew in Figure 5.16 was constructed using the parameters p1 (see Appendix J for parameter values). However, we produced four different parameter sets, which each induces a different combination of bifurcation diagrams from the RK- and G-switches. We can repeat our analysis of the G-switch's behaviour-state dependence using these other parameter sets P2, P3 and P4. The resulting grids of bifurcation diagrams and tables summarising the behaviour-state relationship are all shown in Appendix K.

To conclude our analysis of the G-switch, we have found that each of our four parameterisations P1, P2, P3 and P4 produces a selection of regions in the κ_B, δ_B plane, each with a unique behaviour-state mapping. The number of regions produced by each parameterisation is recorded in Table 5.15.

Table 5.15 Number of G-switch behaviour regions produced by each parameterisation.

Parameterisation	P1	P2	P3	P4
Number of G-switch behaviour regions	9	9	9	9
Number of regions where the behaviours change	7	6	7	7

Table 5.16 Switch configuration-behaviour relationship for the G-switch bifurcation diagram regions in Figure 5.16

Region number	RK-switch	T-switch	G-switch
1	0	0	R
1	1	0	R
1	0	1	R
1	1	1	R
2	0	0	R
2	1	0	R
2	0	1	B
2	1	1	R
3	0	0	B
3	1	0	R
3	0	1	B
3	1	1	R
4	0	0	B
4	1	0	R
4	0	1	D
4	1	1	R
5	0	0	B
5	1	0	R
5	0	1	D
5	1	1	B
6	0	0	D
6	1	0	R
6	0	1	D
6	1	1	B
7	0	0	D
7	1	0	B
7	0	1	D
7	1	1	B
8	0	0	D
8	1	0	B
8	0	1	D
8	1	1	D
9	0	0	D
9	1	0	D
9	0	1	D
9	1	1	D

5.2.4.2. RK-switch behaviour-state mapping

Having found the mapping from R- and T-switch states to G-switch behaviour using the G-switch bifurcation diagrams above, we now turn our attention to the behaviours of the RK-switch.

In the case of the RK-switch, turning on the G-switch causes the gradient of the BE bifurcation lines to increase, causing an area of the bifurcation diagram to be “swept” by the motion of the line. As we

showed in section 5.2.3.1, there are two possible varieties of bifurcation diagram for the RK-switch, one with the {**R**, **B**, **D**} behaviours, and the other with the {**R**, **B**, **Os**, **D**} behaviours. We consider each of these two possibilities in turn below. Once again, we approach the problem by examining the bifurcation lines, for which we have the following expressions.

The BE bifurcation line defining the edge of the turning on region is written as 5.42 and the turn off line as 5.43 (reproduced from section 5.2.2.1 above to save to having to go and look for them).

$$\alpha_P = \frac{(P_{env} + \overline{P_{ext}})}{P^+ \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(0,*,G)\right) \cdot \bar{A}(0,*,\mathrm{G})} \cdot \kappa_P - \frac{\delta_P}{\bar{A}(0,*,\mathrm{G})} \tag{5.42}$$

$$\alpha_P = \frac{(P_{env} + \overline{P_{ext}})}{P^- \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(1,*,G)\right) \cdot \bar{A}(1,*,G)} \cdot \kappa_P - \frac{\delta_P}{\bar{A}(1,*,G)} \tag{5.43}$$

Recall, the focal point values are given by

$$\bar{B}(0,*,G) = \frac{-1 + \sqrt{1 + 4 \cdot \gamma_{BB} \frac{\kappa_B}{\left(\delta_B \cdot K_{\mathrm{off}} + \overline{\delta_{Bext}}\right)} \cdot \frac{\mathrm{G}}{(1 + \gamma_{BR} R_{\mathrm{off}})}}}{2 \cdot \gamma_{BB}} \tag{5.44}$$

$$\bar{A}(0,*,G) = \frac{\kappa_A}{\delta_A} \cdot \frac{G}{1 + \gamma_{AB} \cdot \bar{B}(0,*,\mathrm{G})} \tag{5.45}$$

$$\bar{B}(1,*,G) = \frac{-1 + \sqrt{1 + 4 \cdot \gamma_{BB} \frac{\kappa_B}{\left(\delta_B \cdot K_{\mathrm{on}} + \overline{\delta_{Bext}}\right)} \cdot \frac{G}{(1 + \gamma_{BR} \cdot R_{\mathrm{on}})}}}{2 \cdot \gamma_{BB}} \tag{5.46}$$

$$\bar{A}(1,*,G) = \frac{\kappa_A}{\delta_A} \cdot \frac{G}{1 + \gamma_{AB} \cdot \bar{B}(1,*,G)} \tag{5.47}$$

We can note a couple of things about these BE bifurcation line equations. Firstly, they depend on the parameter values κ_B and δ_B (highlighted in blue in equations 5.44 to 5.47 above) which appeared on the axes of the bifurcation diagrams for the G-switch. This means that, when we set our model parameters to be within a particular region on the G-switch bifurcation diagram, the RK-switch bifurcation diagram must be calculated with those specific values of κ_B and δ_B.

Secondly, the BE bifurcation line equations depend on the state of the G-switch, but not the state of the T-switch. For this reason, we need only consider two configurations of our switches; when the "G-switch is off" and "G-switch is on". The state of the T-switch is irrelevant in this case so we can safely ignore it.

Swept regions when {R, B, D} behaviours are produced

We first examine the RK-switch bifurcation diagrams for the P1 parameterisation, where the behaviours on the diagram are {R, B, D}. The same analysis for the other three parameterisations P2, P3 and P4 can be found in Appendix K. Once we've chosen one of the parameterisations P1, P2, P3 or P4, we still need to work out the values for κ_B and δ_B because the RK-switch BE lines are sensitive to these. To set these parameter values, we pick a region from the G-switch bifurcation diagram Figure 5.16 (e) and take the values of κ_B and δ_B at the centroid. There are 9 regions in the G-switch bifurcation diagram, so we end up producing nine different bifurcation diagrams for the RK-switch, each with slightly different values for κ_B and δ_B. For example, Figure 5.17 is produced using the values of κ_B and δ_B taken from the centroid of region 1 in Figure 5.16 (e).

In the P1 parameterisation, the behaviours produced by the RK-switch in the κ_P, α_P plane are R, B and D. When the G-switch turns on, the BE lines increase in gradient, sweeping out a region of the graph where the switch behaviour changes as a result of the G-switch transition.

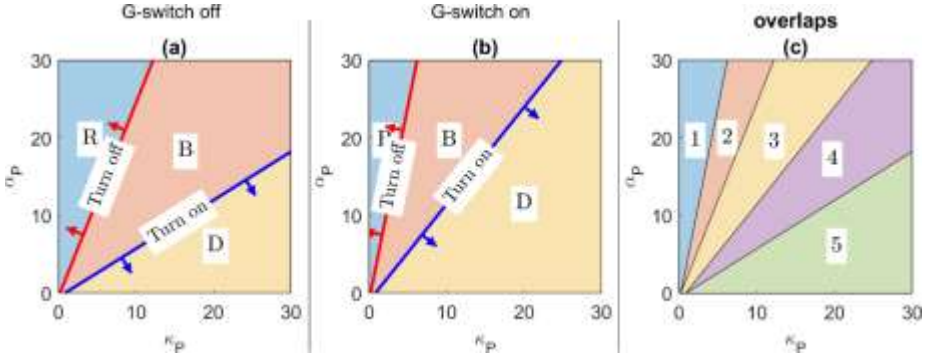


Figure 5.17 Bifurcation diagrams for the RK-switch with the G-switch off (a), the G-switch on (b) and the difference between the G-switch off versus on causing five regions to form (c).

For these parameter values, the area swept by the motion of the lines due to the G-switch turning on leads to five regions being produced (Figure 5.17 c) in the bifurcation diagram, each of which produces a different combination of behaviours as the G-switch changes state. These are summarised in Table 5.17. Note how in regions 2 and 4, the behaviour of the RK-switch changes due to the change of state of the G-switch. This means that if we set the triple-switch model parameters to be anywhere in that region, the behaviour of the RK-switch will change as a result of the G-switch changing state. This relationship can give rise to some interesting dynamics where the switches cause each other to transition in predictable sequences, but this is dealt with in Chapter 6. For the rest of this chapter, we simply show how the other variations of bifurcation diagram look when the switches change state.

Table 5.17 Regions of parameter space where the RK-switch shows consistent configuration-behaviour relationships.

Region number in Figure 5.17 (c)	Behaviour of the RK-switch	
	With G-switch off	With G-switch on
1	R	R
2	R	B
3	B	B
4	B	D
5	D	D

Swept regions when {R, B, Os, D} behaviours are produced.

Let's show another example to demonstrate some different behaviour-state relationships. In our next example, we use the p4 parameterisation, where condition 5.35 is met so the RK-switch bifurcation diagram has four behaviours {R, B, Os, D}, and condition 5.39 is not satisfied so the G-switch bifurcation diagram has four behaviours {R, Os, D}. For this example, we take κ_B and δ_B values from the centroid of region 5 of Figure 5.18.

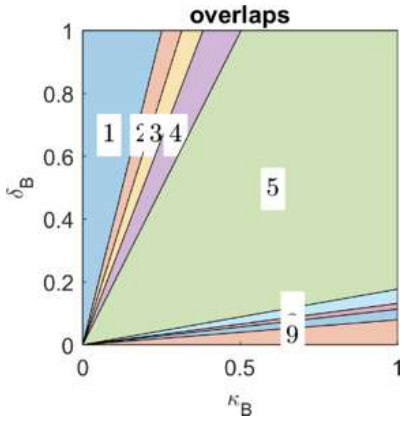


Figure 5.18 G-switch bifurcation overlaps under the p4 parameterisation.

Under these conditions, when the G-switch is off, these regions are laid out as in Figure 5.19 (a). Once again, when the G-switch turns on, the gradient of the BE bifurcation lines increases, resulting in a slightly different bifurcation diagram Figure 5.19 (b). The area swept by the motion of these lines produces Figure 5.19 (c).

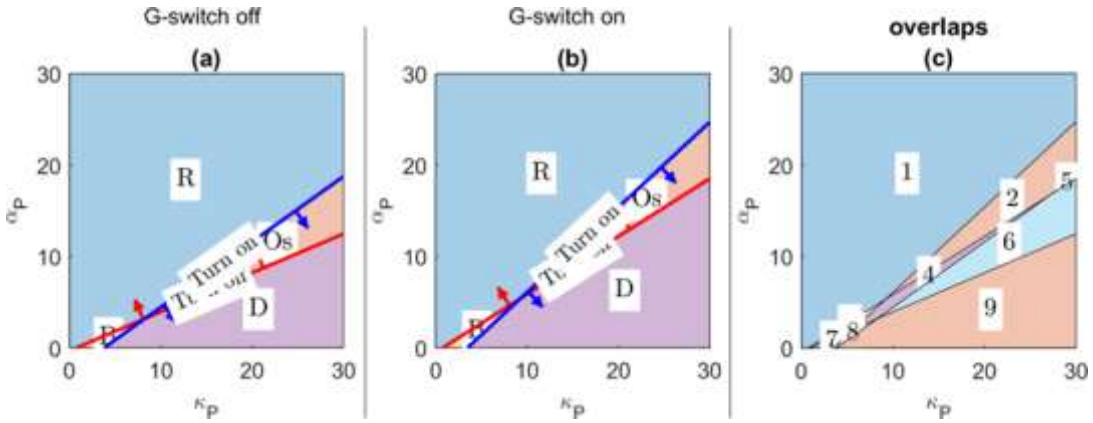


Figure 5.19 Bifurcation diagrams for the RK-switch with the G-switch off (a), the G-switch on (b) and the nine regions produced by considering the changes brought about by turning on the G-switch (c)

In this case, the motion of the BE bifurcation lines leads to nine regions being produced. These are listed in Table 5.18. Five of these nine regions induce a change in the behaviour of the RK-switch when the G-switch changes state.

Table 5.18 Behaviour regions produced by the RK-switch when the G-switch changes state under parameterisation p4.

Region number in Figure 5.19 (c)	RK-switch behaviour	
	With the G-switch off	With the G-switch on
1	R	R
2	R	Os
3	R	B
4	R	D
5	Os	Os
6	Os	D
7	B	B
8	B	D
9	D	D

In both of the examples we have given in Figure 5.17 and Figure 5.19, we had to select the values for κ_B and δ_B from the G-switch bifurcation diagram. We used values from the centroid of regions we

had already identified on the G-switch bifurcation diagram. In general, each of these centroidal values of κ_B and δ_B will affect the RK-switch bifurcation diagram slightly differently, leading to potentially different numbers of regions being produced by the G-switch changing state. In Appendix K, we list the number of RK-switch bifurcation diagram regions produced by all possible combinations of parameterisation and choice of G-switch bifurcation diagram region. In each case, the RK-switch bifurcation diagram produces between five and nine regions as a result of the "overlap" process we describe here.

5.2.5. Summary of our parameter sampling method

Our method for sampling the parameter space has proceeded with the aim of identifying all the possible dynamics which our model can produce. To achieve this, we kept track of how the parameter values, and the state of the other switches, affects the layout of the bifurcation diagrams. We found that the bifurcation diagrams we draw for the RK- and G-switches take on certain common layouts which we were able to categorise according to the switch behaviours present on the diagram and were able to write inequalities on our parameters defining when each layout of behaviours would occur. We found that there are only four possible combinations of layouts of behaviours on the RK- and G-switch bifurcation diagrams, meaning that the model always produces one of these four combinations. We identified one set of parameters for each of these, and called them P1, P2, P3 and P4.

Each of these four parameterisations summarises a different overall behaviour from the triple-switch model, producing different combinations of behaviours from the switches as we vary the values of the bifurcation parameters $\kappa_P, \alpha_P, \kappa_B$ and δ_B. For example, the P4 parameterisation, which causes the RK-switch's bifurcation diagram to contain the behaviours {R, B, Os, D} and the G-switch's bifurcation diagram to contain the behaviours {R, Os, D} is the only parameterisation where the triple-switch model is able to exhibit Os (oscillation) behaviours from both the RK- and G-switch simultaneously. The resulting dynamics, as we shall see in the next chapter, are different to those produced by the P1 parameterisation where none of the switches can exhibit the Os behaviour.

Having defined these four parameterisations, we next considered how the states of the other switches would affect the layout of the bifurcation diagrams. We found that the boundary-equilibrium lines on the bifurcation diagrams are affected by the switch states causing the behaviour of one switch to be affected by the others. This crosstalk was then exhaustively enumerated for each of the parametrisations P1, P2, P3 and P4. In this way, we have been able to produce an exhaustive list of regions in the parameter space where the three switches have differing behaviour-state mappings.

However, a shortcoming of our selection of just four parameter sets is that we are not considering how parameter variation within the four parameter regions we found can potentially bring about more subtle changes in the switches' crosstalk by altering how the BE bifurcation lines move relative to each other. Variations of these remaining parameters within our four parameter categories can lead to differences in the extent to which the boundary-equilibrium bifurcation lines move, and this in turn can lead to differences in the configuration-behaviour relationships that emerge. For this reason, we recommend that future work should sample within each of these parameter regions more carefully.

5.2.6. Parameter sensitivity of bifurcation diagrams

Now that we've defined our method to calculate the various regions in the parameter space where our three switches interact in interesting ways, we can now ask how robust these regions are to variation of the model parameters. Do some of the regions change size very rapidly compared to others as we vary the parameters? Are there particular parameters which cause the regions to change size more rapidly than others? To answer these questions, it is helpful to have a method to calculate the derivative of the areas of each region with respect to each of the model's parameters.

We are able to calculate these derivatives analytically, making use of the particularly simple equations we have obtained for the BE bifurcation lines, and combining this with some nice geometry relating to the polygon-shaped regions we have obtained.

Let us denote the area of one of an arbitrary polygonal region as A, while the parameters of the model are summarised in a vector $\theta = \begin{bmatrix} P_{\text{env}} \\ \kappa_P \\ \vdots \end{bmatrix}$. We wish to calculate the derivative of the area of the region A with respect to the parameters θ, $\begin{bmatrix} {dA}/{dP_{\text{env}}} \\ {dA}/{d\kappa_P} \\ \vdots \end{bmatrix}$. We take the following five steps to organise the process of calculating this area analytically.

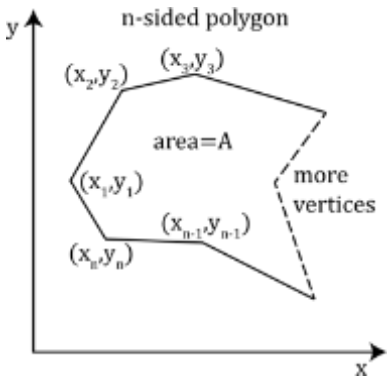


Figure 5.20 The area of an arbitrary n-sided polygon can be calculated given the coordinates of the vertices.

Stage 1. Calculate the area of each region.

To calculate the area of an arbitrary polygon such as those we tend to find (Figure 5.20), we use the general formula for the area A using the formula

$$A = \left| \frac{(x_1 y_2 - x_2 y_1) + (x_2 y_3 - x_3 y_2) + \cdots (x_{n-1} y_n - x_n y_{n-1}) + (x_n y_1 - x_1 y_n)}{2} \right| \tag{5.48}$$

Stage 2. Calculate the derivate of the area with respect to the vertices.

We organise the vertices of the areas into a vector $v = [x_1 \quad y_1 \quad x_2 \quad y_2 \quad \cdots \quad x_n \quad y_n]$. The derivative is

$$\begin{aligned} \frac{dA}{dv} &= \left[\frac{dA}{dx_1} \quad \frac{dA}{dy_1} \quad \frac{dA}{dx_2} \quad \frac{dA}{y_2} \quad \cdots \quad \frac{dA}{dx_n} \quad \frac{dA}{dy_n} \right] \\ &= \left[\frac{y_2 - y_n}{2} \quad \frac{-x_2 + x_n}{2} \quad \frac{y_3 - y_1}{2} \quad \frac{-x_3 + x_1}{2} \quad \cdots \quad \frac{y_1 - y_{n-1}}{2} \quad \frac{-x_1 + x_{n-1}}{2} \right] \end{aligned} \tag{5.49}$$

For our application, we're interested in the areas of multiple polynomials which have continuous edges, where they share vertices. For example, in Figure 5.21, there are four regions, but they all share the middle vertex (x_5, y_5). In fact, there are only 9 vertices in the diagram, four of them (4, 6, 7 and 8) shared by two areas each, and one (the middle vertex 5) shared by all four.

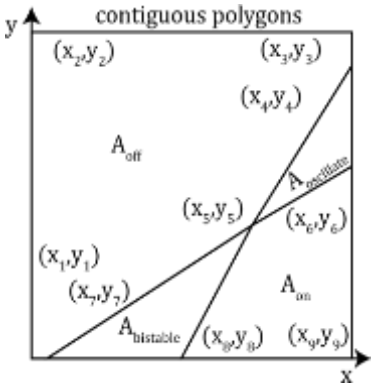


Figure 5.21 Contiguous polygons share vertices, simplifying our calculations somewhat.

To take advantage of this shared information between the contiguous areas, we can organise our calculation as a matrix to include each of the regions by defining a vector of regions

$$A = [A_{\text{off}} \quad A_{\text{bistable}} \quad A_{\text{oscillate}} \quad A_{\text{on}}] \tag{5.50}$$

Then we can write the derivative of these areas with respect to the vertices by

$$\frac{dA}{dv} = \begin{bmatrix} \frac{\partial A_{\text{off}}}{\partial x_1} & \frac{\partial A_{\text{bistable}}}{\partial x_1} & \frac{\partial A_{\text{oscillate}}}{\partial x_1} & \frac{\partial A_{\text{on}}}{\partial x_1} \\ \frac{\partial A_{\text{off}}}{\partial y_1} & \frac{\partial A_{\text{bistable}}}{\partial y_1} & \frac{\partial A_{\text{oscillate}}}{\partial y_1} & \frac{\partial A_{\text{on}}}{\partial y_1} \\ \frac{\partial A_{\text{off}}}{\partial x_2} & \frac{\partial A_{\text{bistable}}}{\partial x_2} & \frac{\partial A_{\text{oscillate}}}{\partial x_2} & \frac{\partial A_{\text{on}}}{\partial x_2} \\ \frac{\partial A_{\text{off}}}{\partial y_2} & \frac{\partial A_{\text{bistable}}}{\partial y_2} & \frac{dA_{\text{oscillate}}}{\partial y_2} & \frac{\partial A_{\text{on}}}{\partial y_2} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial A_{\text{off}}}{\partial x_n} & \frac{\partial A_{\text{bistable}}}{\partial x_n} & \frac{\partial A_{\text{oscaillte}}}{\partial x_n} & \frac{\partial A_{\text{on}}}{\partial x_n} \\ \frac{\partial A_{\text{off}}}{\partial y_n} & \frac{\partial A_{\text{bistable}}}{\partial y_n} & \frac{\partial A_{\text{oscaillate}}}{\partial y_n} & \frac{\partial A_{\text{on}}}{\partial y_n} \end{bmatrix} \tag{5.51}$$

This matrix will contain quite a few zeros where the area represented in the column does not include the vertex represented in the row. For example in Figure 5.21, the derivative with respect to vertex 2 (x_2, y_2) (rows 3 and 4 of the matrix), $\frac{dA}{dv}$ is non-zero only for the A_{off} area, but zero otherwise as it only appears on the edge of the A_{off} area.

Stage 3. Calculate the rate of change of the vertices with respect to the parameters.

Next, we can calculate formulae for each of the vertices of the regions in our bifurcation plots for the RK-switch and G-switch (figures 3, 5, 6 and 7), and calculate the derivatives of these with respect to the parameter vector θ. We can do this by examining the equations we derived for the BE bifurcation lines. For example, taking the equation defining the turning off edge for the RK-switch, we can differentiate with respect to P_{env} as shown in 5.52.

$$\begin{aligned} \alpha_P &= \frac{(P_{env} + \overline{P_{ext}})}{P^+ \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(0,*,G)\right) \cdot \bar{A}(0,*,\text{G})} \cdot \kappa_P - \frac{\delta_P}{\bar{A}(0,*,\text{G})} \\ \frac{d\alpha_P}{dP_{\text{env}}} &= \frac{\kappa_P}{P^+ \cdot \left(1 + \gamma_{PB} \cdot \bar{B}(0,*,G)\right) \cdot \bar{A}(0,*,\text{G})} \end{aligned} \tag{5.52}$$

We can evaluate these derivatives at each of the vertices of our polygons to obtain $\frac{dv}{d\theta}$ defined in equation 5.53.

$$\frac{dv}{d\theta} = \begin{bmatrix} \frac{\partial x_1}{\partial P_{\text{env}}} & \frac{\partial y_1}{\partial P_{\text{env}}} & \cdots & \frac{\partial y_n}{\partial P_{\text{env}}} \\ \frac{\partial x_1}{\partial \kappa_p} & \frac{\partial y_1}{\partial \kappa_P} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ \frac{\partial x_1}{\partial \delta_D} & \cdots & \cdots & \frac{\partial y_n}{\partial \delta_D} \end{bmatrix} \tag{5.53}$$

Stage 4. Calculate the sensitivity by multiplying the matrices, using the chain rule to recover the desired derivatives.

Now we can calculate the rate of change of the area of each region with respect to the parameters by the matrix multiplication

$$\frac{dA}{d\theta} = \frac{dA}{dv} \cdot \frac{dv}{d\theta} \tag{5.54}$$

The results is a matrix of derivatives, each column of which represents the rate of change of area as we vary the parameters θ for each one of the regions. We can write this as follows in equation 5.55.

$$\frac{dA}{d\theta} = \begin{bmatrix} \frac{\partial A_{\text{off}}}{\partial P_{\text{env}}} & \frac{\partial A_{\text{bistable}}}{\partial P_{\text{env}}} & \frac{\partial A_{\text{oscillate}}}{\partial P_{\text{env}}} & \frac{\partial A_{\text{on}}}{\partial P_{\text{env}}} \\ \frac{\partial A_{\text{off}}}{\partial \kappa_p} & \frac{\partial A_{\text{bistable}}}{\partial \kappa_P} & \frac{\partial A_{\text{oscillate}}}{\partial \kappa_P} & \frac{\partial A_{\text{on}}}{\partial \kappa_P} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial A_{\text{off}}}{\partial \delta_D} & \frac{\partial A_{\text{bistable}}}{\partial \delta_D} & \frac{\partial A_{\text{oscillate}}}{\partial \delta_D} & \frac{\partial A_{\text{on}}}{\partial \delta_D} \end{bmatrix} \tag{5.55}$$

Stage 5. Normalise the sensitivities by the area of each region so we can compare each region's sensitivity more easily.

Normalisation is carried out by dividing the sensitivity matrix by the vector of areas of each behaviour region as follows

$$\frac{dA}{d\theta} \cdot \frac{1}{A} = \begin{bmatrix} \frac{\partial A_{\text{off}}}{\partial P_{\text{env}}} \cdot \frac{1}{A_{\text{off}}} & \frac{\partial A_{\text{bistable}}}{\partial P_{\text{env}}} \cdot \frac{1}{A_{\text{bistable}}} & \frac{\partial A_{\text{oscillate}}}{\partial P_{\text{env}}} \cdot \frac{1}{A_{\text{oscillate}}} & \frac{\partial A_{\text{on}}}{\partial P_{\text{env}}} \cdot \frac{1}{A_{\text{on}}} \\ \frac{\partial A_{\text{off}}}{\partial \kappa_p} \cdot \frac{1}{A_{\text{off}}} & \frac{\partial A_{\text{bistable}}}{\partial \kappa_P} \cdot \frac{1}{A_{\text{bistable}}} & \frac{\partial A_{\text{oscillate}}}{\partial \kappa_P} \cdot \frac{1}{A_{\text{oscillate}}} & \frac{\partial A_{\text{on}}}{\partial \kappa_P} \cdot \frac{1}{A_{\text{on}}} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial A_{\text{off}}}{\partial \delta_D} \cdot \frac{1}{A_{\text{off}}} & \frac{\partial A_{\text{bistable}}}{\partial \delta_D} \cdot \frac{1}{A_{\text{bistable}}} & \frac{\partial A_{\text{oscillate}}}{\partial \delta_D} \cdot \frac{1}{A_{\text{oscillate}}} & \frac{\partial A_{\text{on}}}{\partial \delta_D} \cdot \frac{1}{A_{\text{on}}} \end{bmatrix} \tag{5.56}$$

By this means, we obtain a measure of how much effect variation of each parameter has on the rate of expansion or contraction of each of the behaviour regions we're interested in. Furthermore, this measure of rate of change reports the rate of expansion relative to the area itself. A doubling in area as we vary a parameter would produce the same value for two regions of different sizes.

We wish to evaluate the impact of variation of our parameters on the output of the model (in this case, the size of the parameter region we've identified as interesting to us). With the method described above, it is very straightforward to evaluate the analytical local sensitivities of the areas on the bifurcation plots. We must first specify the parameter values around which we calculate these derivatives. Then, we can evaluate the derivatives analytically, and substitute in the parameter values to obtain the local sensitivities. We can then ascertain which parameters of the model most strongly affect the size of the areas of the behaviour regions for the parameter values we've investigated so far.

The calculation itself just involves a couple of matrix multiplications. The first matrix we calculate, $\frac{dA}{dv}$ just summarises the geometry of the bifurcation diagram and is usually sparce in structure due to the absence of certain vertices from some of the areas we're considering. The second matrix $\frac{dv}{d\theta}$ (equation 5.53) summarises our model structure and whilst this is a little more involved to calculate, this need only be done once and the results stored. It's also just simple differentiation, so it's well within the abilities of a symbolic algebra software such as MATLAB's Symbolic Algebra Toolbox, which we use in our software rather than running out of marker pens. This approach is also readily adapted to other hybrid dynamical systems where the BE bifurcation curves can be symbolically defined.

Should we wish to consider a different location in the parameter space, we need only substitute in different parameter values to our analytical expressions for $\frac{dv}{d\theta}$, and potentially update the $\frac{dA}{dv}$ matrix if the geometry of the bifurcation diagram changes for our new parameter values. For this reason, it takes few computations to obtain our sensitivity figures, so it's reasonable to use this method to sample the local parameter sensitivity at many locations in the parameter space, producing a global sensitivity analysis.

Using the method described above we can directly evaluate the local sensitivities of each of the behaviours of the RK-switch when on the bifurcation diagram for which only the {R, B, D} behaviours are produced. With this information, we can detect which parts of the model most strongly affect the dependence of the RK-switch on the AD risk factors represented by κ_P and α_P. We also gain understanding of which parts of the model affect the RK-switch behaviour by seeing which parameters are sensitive.

These sensitivities are shown in Figure 5.22. Since we are evaluating the sensitivities of the areas of these behaviour regions in the κ_P, α_P plane, it does not make sense to evaluate the sensitivity to the κ_P or α_P parameters. They are therefore omitted from the list on the x-axis of Figure 5.22.

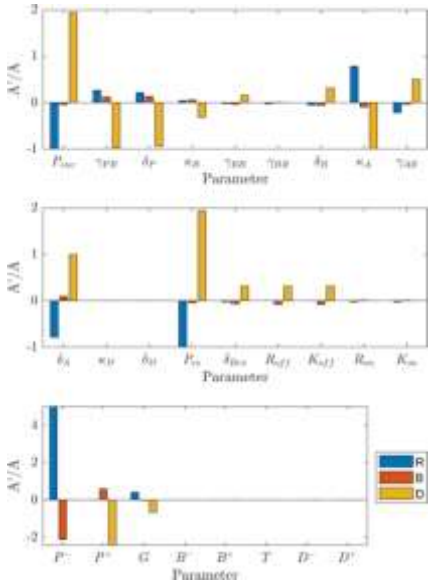


Figure 5.22 Local sensitivities of the RK-switch for the "p1" parameterisation (parameter values in Appendix J).

In this chapter, we omit the remaining sensitivity analysis diagrams for the G-switch, and the other types of RK-switch bifurcation diagram in the interests of brevity. These sensitivity analysis diagrams may be found in Appendix L.

5.3. Conclusion

In this chapter, we defined the triple-switch model of chronic AD. The model consists of a hybrid dynamical system with three discrete-valued variables (RK-, T- and G-switches), and four continuous-valued state variables (P, B, A and D). This hybrid dynamical system structure was chosen to help us combine the slow-timescale and fast-timescale disease mechanisms into one mechanistic model. The discrete-valued variables represent the fast-timescale innate immune response, inflammatory adaptive immune response, and differentiation of the barrier, respectively. These correspond to the visible symptoms of AD such as redness, inflammation, and flakiness of the skin. The continuous-valued variables represent the slow-timescale mechanisms which are less immediately apparent to a patient, such as pathogen load, barrier integrity, level of anti-microbial peptide on the skin surface

and dendritic cell concentration in the lymph-nodes. These features are certainly measurable by taking skin swabs of patients or applying imaging techniques in-vivo, so their dynamics usually remain invisible as the disease phenotype emerges, particularly in humans.

The dynamics of this hybrid dynamical system are driven by crosstalk between the three switches. We identified how this crosstalk occurs by applying the novel bifurcation analysis methods developed in this chapter which allowed us to obtain an exhaustive list of ways in which the switches may interact to bring about AD dynamics. This helps us fill the knowledge gap in our understanding of how complicated dynamics of AD emerge from physiological mechanisms.

Finally, we carried out a sensitivity analysis on the bifurcation diagrams we produced, observing how variation of the parameters brings about changes to the size of the various regions corresponding to different configurations of the crosstalk we identified. We were able to evaluate the rate of change of the regions with respect to each parameter analytically.

In the next chapter, we present novel methods for dynamic analysis, which build on the results from this chapter to infer how the model's dynamics proceed over time. We make use of the switch behaviours, the bifurcation diagrams showing us the parameter dependence of these behaviours, and finally the important observation that the bifurcation diagrams are dependent on the switch states.

Chapter 6. Dynamical analysis of switch transitions

Now that we have defined a model for chronic AD, we wish to verify that the model can reproduce known AD dynamics. To achieve this, we use the bifurcation analysis tools we developed in chapter 5 to infer the range of dynamical behaviours we can produce. The first step in our analysis is to verify that the model is producing dynamics corresponding to the poorly understood relapsing-remitting dynamics of chronic AD. After this, we wish to fill in the knowledge gap by identifying how the mechanisms included in our model actually bring about these dynamics. In short, we wish to leverage the explanatory power of mechanistic modelling to inform our knowledge of chronic AD pathogenesis.

We can draw a diagram summarising every possible switch configuration of the model, and every possible transition between switch configurations as in Figure 6.1. The further to the right in Figure 6.1, the more switches in our model are in their "active" state, meaning that there is more inflammation or poorer keratinocyte differentiation. The leftmost configuration, where all three switches are off corresponds to a completely healthy state, while the rightmost configuration corresponds to the worst possible scenario in our model. The model's dynamics may be visualised as a series of transitions between the nodes in Figure 6.1.

We show in this chapter that transitions can occur sequentially, forming long chains of transitions from state to state until either the dynamics reach a switch configuration which is sustained, or a loop of switch configurations is formed. By the end of this chapter, we aim to explain how we can infer these model dynamics, the limitations of the method we identify, and present our results in a way that will help us to understand the biological implications in Chapter 7.

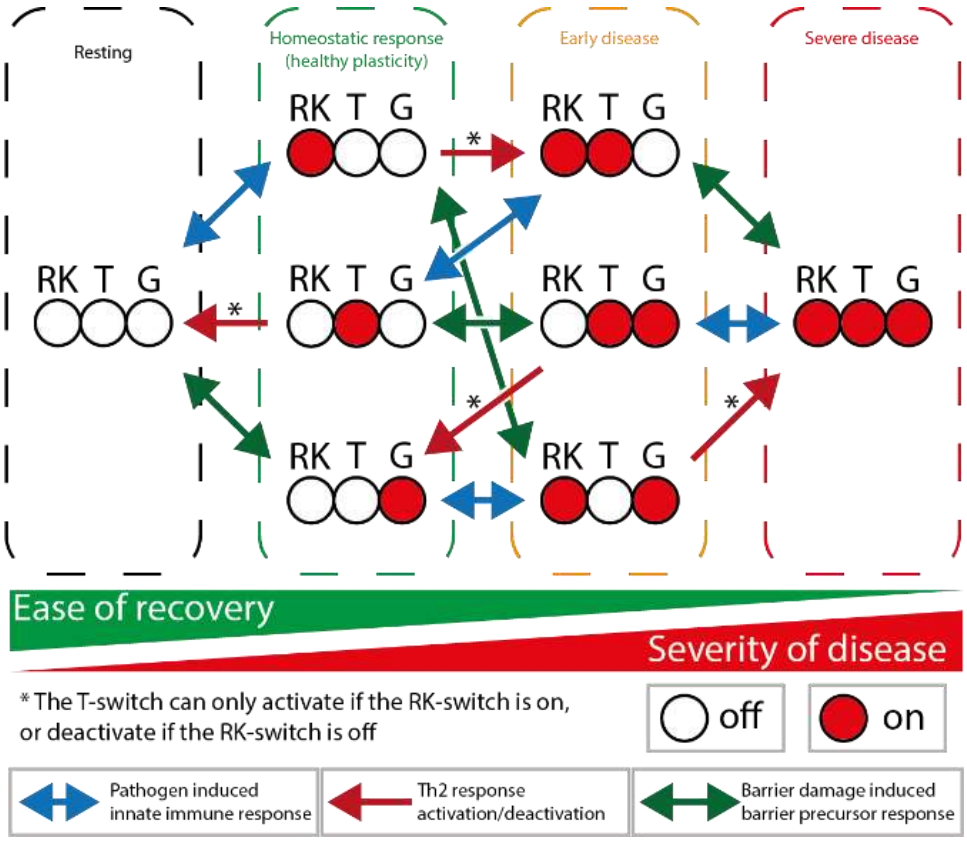


Figure 6.1 Switch state network, with every potential transition marked by an arrow.

6.1. Vocabulary and definitions

In this chapter, the concepts will be explained using some consistent vocabulary. We define it all here both to prepare the reader for what is to come, and to act as a reference.

6.1.1. Switch configurations

The states of the three switches (RK-, T- and G-switch) of our model (equations 5.2, 5.6 and 5.5) can each be either on or off. We call the on/off states of all three switches collectively the "**switch configuration**" or simply "**configuration**". We can write the switch configurations as a vector $\boldsymbol{C}$ of ones and zeros (defined in 6.1), representing the off and on states using 0 and 1 respectively for the RK-, T- and G-switches in that order. This is often useful in notation where we wish to explicitly keep track of which switch is undergoing a change of state.

$$\boldsymbol{C} \in \{0,1\}^3 \tag{6.1}$$

We always express the states of the switches in the same order: RK-switch, T-switch, G-switch. For example, the vector $\boldsymbol{C} = [0, 0, 1]$ would represent the switch configuration specified in Table 6.1.

Table 6.1 Switch configuration

RK-switch	T-switch	G-switch
Off	Off	On

For visual simplicity in our diagrams, we denote the switch configuration by using three circles, coloured white if the switch is off or red if the switch is on and omit any additional labelling. For example, we could denote the switch state [0,0,1] by the circles . We occasionally include letters within these circles to remind us of which switch is being represented, for example 

. When a swich is on (i.e. a red circle), this corresponds to the unhealthy state for that switch. Thus, the more red circles are present, the unhealthier the skin.

Indexing the switch configuration for ease of computation

Our system consists of three switches, each of which can take on two states (off or on), so we know our system can take on a total of $2^3 = 8$ possible states. In the computer code written in this PhD, it is often more convenient for us to represent these states using a vector $\boldsymbol{S}$ of length 8, where each element is a zero or one corresponding to each of the possible configurations of $\boldsymbol{C}$. We define $\boldsymbol{S}$ as follows.

$$\boldsymbol{S} \in \{0,1\}^8 \tag{6.2}$$

When representing a switch state, the vector $\boldsymbol{S}$ always consists of all zeros with a single entry of value one which represents the current switch configuration. We order the vector $\boldsymbol{S}$ in ascending binary order as in Table 6.2.

Table 6.2 Ways of representing the switch configuration graphically and in mathematical notation in the rest of this chapter.

RK-switch	T-switch	G-switch	$\boldsymbol{C}$ configuration representation	Circles representation	Index of the "1" value in the $\boldsymbol{S}$ vector
Off	Off	Off	[0,0,0]		1
Off	Off	On	[0,0,1]		2
Off	On	Off	[0,1,0]		3
Off	On	On	[0,1,1]		4
On	Off	Off	[1,0,0]		5
On	Off	On	[1,0,1]		6
On	On	Off	[1,1,0]		7
On	On	On	[1,1,1]		8

Thus, to represent the switch state [0, 0, 1] using the vector $\boldsymbol{S}$ we would write as below in 6.3.

$$\boldsymbol{S} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \tag{6.3}$$

We choose to represent the switch configuration using the vector $\boldsymbol{S}$ since it leads to a convenient way to summarise switch state networks as an 8-by-8 matrix, and an easy way to compute the next switch configuration given the current switch configuration. In all of the MATLAB code associated with this chapter, the $\boldsymbol{S}$ vector is used for internal calculations, but we recognise that this notation is less interpretable for a human reader so the code always translates the result into the $\boldsymbol{C}$ format for the user. Whenever we show figures discussing specific switch configurations, we always illustrate these configurations using the three circles format for visual clarity. Where possible, we label each circle with RK, T and G to remind the reader of which switch it represents. However, where this would become visually cluttered, we omit these labels.

6.1.2. Switch transitions

The dynamics of the triple-switch's ODE can progress in such a way that the state variables P, D and B may cross the RK-, T- and G-switches' respective threshold values required to cause a change of state. For example, if the RK-switch is initially in the "off" state, the RK-switch would turn "on" if the P state variable increases past P^+. The moment of "transition" from one state to the other occurs instantaneously as the state variable crosses the switch's threshold. Since our model has three switches, it is possible for three different transitions to occur at a particular moment in time.

In general, we are more interested in charactering which switch transition occurs in a particular timeframe and the implications for subsequent transitions than working out exactly when these transitions occur. This is because we're interested in characterising the qualitative behaviour of the model for various parameterisations, and this does not require that we precisely characterise the

switch timing. For this reason, we represent the sequence of switch configuration in order by giving them a subscript $\boldsymbol{S}_n$, which represents the nth switch state of the model from a defined initial switch state. We represent transitions from $\boldsymbol{S}_n$ to $\boldsymbol{S}_{n+1}$ by

$$\boldsymbol{S}_{n+1} = F \cdot \boldsymbol{S}_n \tag{6.4}$$

where a matrix F=(F_{ij}) with $i \in [1,8]$ and $j \in [1,8]$ and F_{ij} represents whether or not the model will transition from state i to state j in the switch state vector $\boldsymbol{S}n$, taking a value of zero if this transition does not occur, and one if it does occur. We may deduce the nth switch state by

$$\boldsymbol{S}_n = F^n \cdot \boldsymbol{S}_0 \tag{6.5}$$

where $\boldsymbol{S}_0$ is the initial state.

Our model is deterministic, so it follows that for a given parameter vector θ and set of initial conditions of the continuous-valued state variables (P_0, B_0, A_0, D_0) there is a single corresponding transition matrix F, where each switch state leads to at most one predictable transition. This means each switch state in our network cannot have out-degree>1 (i.e. only one arrow going *from* a specific switch state), though it is possible to have in-degree>1 (i.e. multiple arrows going *to* a specific switch state). This guarantees that every $\boldsymbol{S}_n$ vector we calculate contains just a single "1" digit, representing just a single state after n transitions, so equation 6.5 always returns an interpretable value for $\boldsymbol{S}_n$.

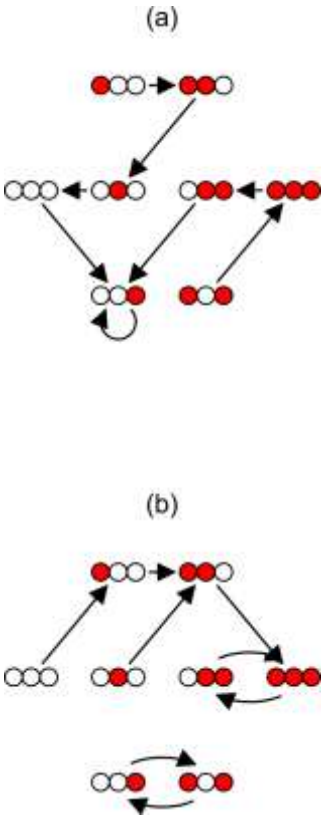


Figure 6.2 Example transition networks

In this chapter, we describe how to calculate the transition matrix F for any parameter vector θ of the triple-switch model. We identify 215 regions of parameter space separated by boundary equilibrium bifurcation lines in the bifurcation diagrams and separatrices in the state space.

The BE bifurcation lines are themselves dependent on the parameters of the model, and we use the methods developed in Chapter 5 to track the impact of variation of all the parameters. We also account for the effect of initial conditions (P_0, B_0, A_0, D_0) on the transitions by considering the basins of attraction for each type of network. In each of these parameter regions and basins of attraction, the transition matrix F is the same for all parameter values and initial conditions contained within the region.

The resulting transition matrix F is then represented in the form of a network diagram, consisting of arrows representing individual transitions from one state to the next. Two examples of such network diagrams are shown in Figure 6.2.

Where a switch state does not lead to a transition, we draw a curved arrow pointing back at this switch configuration. For example, all the sequences of transitions in Figure 6.2 (a) end at [0,0,1], which is an absorbing switch configuration leading to no further transitions, so it has an arrow pointing at itself. It can also occur that we reach an absorbing sequence of switch configurations which transition between each other in a loop as shown in Figure 6.2 (b) with [0,1,1] and [1,1,1], and another loop with [0,0,1] and [1,0,1].

In all network diagrams, we divide the switch configurations into two distinct groups: absorbing configurations and transient configurations. Absorbing configurations are those which have no outgoing transition as represented by a self-referential arrow as in Figure 6.2 (a) for state [0,0,1] or

appear in a loop sequence as in Figure 6.2 (b) for configurations [0,1,1] and [1,1,1]. At absorbing loop configurations, we can write $\boldsymbol{S}_{n+k} = \boldsymbol{S}_n$where k is any positive even number[4].

Transient configurations are all other configurations, which by definition occur only once in the transient part of the model dynamics. At such transient configurations $\boldsymbol{S}_{n+k} \neq \boldsymbol{S}_n$ for any $k \neq 0$.

6.1.3. Transition networks or trees

To help us visually distinguish transient versus absorbing configurations and loops in each network we produce, we arrange the absorbing configuration or loop at the bottom of the diagram and arrange all of the transient configurations leading to them as branches, rather like a tree. The absorbing configurations or loops act as roots to each tree, and each transient configurations appears on the branch of exactly one tree in each network. Aside from organising the switch states in a way that allows us to spot patterns and similarities between networks, the tree representation also has the effect of preventing the transition arrows from overlapping, ensuring visual clarity. We revisit this tree analogy in Chapter 7 where we consider the effect of treatments as analogous to squirrels jumping between the branches of the trees.

Arranging the network in this way, we can immediately see how many absorbing configurations and loops are present, how the transient configurations relate to each of these and any special properties of each "tree" of transient configurations. In Figure 6.3 (a) we see a tree diagram where there is a single absorbing configuration, so all the transient configurations states eventually lead to it. Figure 6.3 (b) shows a similar situation, except that a second absorbing configurations exists so we see two trees.

Note that because the model is deterministic, each switch state can only have one outgoing transition arrow. This means it is not possible for transitions to occur between the branches of different trees, so each tree remains separate from the others. In other words, in networks such as Figure 6.3 (b) where multiple absorbing configurations are observed, each absorbing configuration has its own set of transient switch states, and there is no overlap between the set for each absorbing configuration.

It is possible for absorbing loops to act as the root to multiple branches, as shown in Figure 6.3 (c) where the two transient configurations lead to the same absorbing loop, though they join at different points in the loop. We consider such cases as a single tree since the transient configurations share the same "root" absorbing loop. This is perhaps analogous to a tree where multiple trunks are visible, but which share the same root system, as often occurs with birch trees as shown in Figure 6.3 (d).

[4] For any given switch state containing an even number of switches in the "on" position, the adjacent switch states must have an odd number of switches in the "on" position. Conversely, for any switch state containing an odd number of switches in the "on" position, the adjacent switch states must contain an even number of switches in the "on" position. This means that any loop of switch states must connect through alternating odd- and even-numbered switch states, and the only way to form a loop of odd- and even-numbered nodes in a loop is by having an even number of nodes.

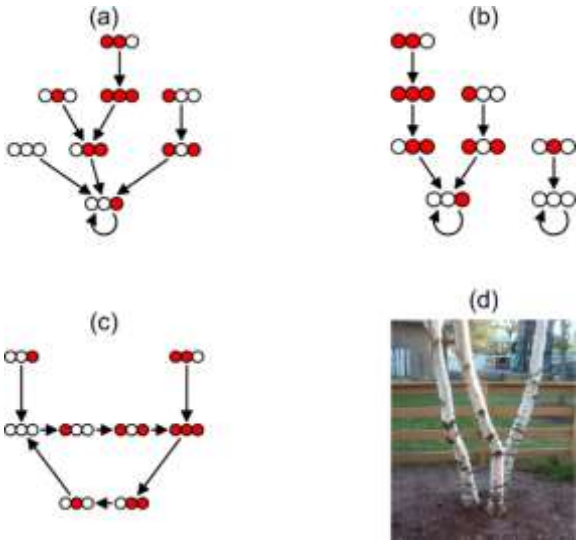


Figure 6.3 Transition diagrams arranged as trees, with transient configurations on branches leading to absorbing configurations and loops.
(a) There is one global absorbing switch configuration, (b) two absorbing switch configurations and (c) with a single absorbing loop of 6 switch configurations. (d) A paper birch tree, where multiple trunks often originate from the same roots giving the misleading appearance of multiple closely spaced trees rather than a single tree with multiple trunks.

6.1.4. Families of trees

Later in this chapter, we will show that the parameter regions we investigate can in principle produce multiple networks, each corresponding to a different basin of attraction for the continuous-valued state variables $\{P, B, A, D\}$. These networks all share the same parameter values but differ only in the initial conditions of the state variables. Due to this connection, we group these networks together, referring to them as "families" of trees. We explain this process in section 6.6.

6.2. Constructing configuration-behaviour tables

To identify if a spontaneous transition of switch state occurs, we make use of the concept of each switch's behaviour, described in section 5.2.2. We showed how each switch may exhibit one of the four possible behaviours; recovery (R), bistable (B), oscillation (Os [5]) or damage (D), depending on how their two "switching conditions" are satisfied. We define the behaviour states as a vector $\boldsymbol{b}$ of letters from the set $\{R, B, Os, D\}$, representing the behaviours of the RK-, T- and G-switches in that order.

$$\boldsymbol{b} \in \{R, B, Os, D\}^3 \tag{6.6}$$

For example, the vector $\boldsymbol{b} = [B, R, Os]$ would represent the switch states specified in Table 6.3.

Table 6.3 Switch behaviours represented by the vector $\boldsymbol{b} = [B, R, Os]$

RK-switch	T-switch	G-switch
B	R	Os

Given that each switch has four possible behaviours to select, and we have three switches, we might initially conclude that there are 4^3=64 possible values that the vector $\boldsymbol{b}$ can take. However, the behaviour of the T-switch is limited due to the structure of the ODE for D, where we assume that the

[5] Note, we choose to denote oscillation using "Os" rather than "O" to prevent it looking similar to the symbol for zero (which we already use to denote a switch being in the "off" configuration).

timescale of the RK-switch transitions is always slower than the timescale of dendritic cell dynamics (see section 5.2.2.2). The behaviour of the T-switch is "R" if the RK-switch is off, and "D" if the RK-switch is on, so we can see the T-switch is actually limited to just these two behaviours. This halves the number of possible values for $\boldsymbol{b}$ down to 4 x 2 x 4 = 32. The switch configuration $\boldsymbol{C}$ has just 2^3=8 possible values.

In section 5.2.4, we showed how each configuration of the states (on or off) of the three switches gives us a distinct combination of behaviours from each switch. We used our bifurcation analysis to chop up the parameter space into regions where the relationship between the switch state configuration $\boldsymbol{C}$ and the induced behaviours $\boldsymbol{b}$ is homogeneous across each region. Thus, for each region we identified, the switch state configuration $\boldsymbol{C}$ maps to a single value of the behaviour vector $\boldsymbol{b}$. The set of possible configurations for $\boldsymbol{C}$ and $\boldsymbol{b}$ is illustrated in Figure 6.4.

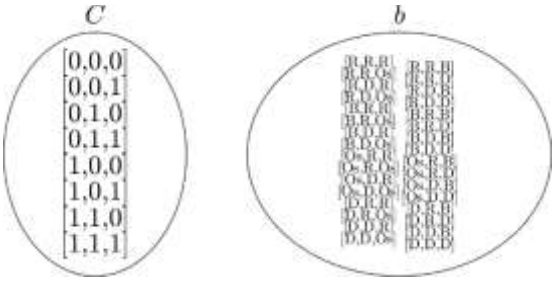


Figure 6.4 Illustration of the set of switch states $\boldsymbol{C}$ and set of behaviours $\boldsymbol{b}$ whose mapping we identify.

This mapping may be many-to-one. In other words, there may be multiple values of $\boldsymbol{C}$ which yield the same value of $\boldsymbol{b}$. However, each value of $\boldsymbol{C}$ yields just one value of $\boldsymbol{b}$ within each of the parameter regions we identified in Chapter 5. Each combination of RK-switch region and G-switch region has one mapping $\boldsymbol{C} \mapsto \boldsymbol{b}$ which we use to infer the spontaneous transitions between switch state configurations.

This configuration-behaviour mapping $\boldsymbol{C} \mapsto \boldsymbol{b}$ summarises the crosstalk between the three switches, telling us when a change of state can bring about a change in behaviour, and in what way. It tells us what kind of dynamics we should expect from each switch for any given configuration of the switches, and crucially how any changes to the switch configuration will affect those dynamics. By combining this information about each switch's behaviour, we are able to reconstruct the emergent dynamics of the entire model. This process of reconstructing the model's dynamics from the configuration-behaviour mapping is discussed in the next section 6.3, but for now we will simply define how we can calculate the mapping for any region in the parameter space.

We summarise this mapping in a table listing the behaviours $\boldsymbol{b}$ for each switch state configuration $\boldsymbol{C}$. This so-called "configuration-behaviour table" is built by combining the information from the RK-switch bifurcation diagram and G-switch bifurcation diagram, along with our assumption about the T-switch's timescale of change. This process is the first step in our dynamical analysis, and proceeds in stages, taking this information from the T-switch, then the G-switch and finally the RK-switch before we combine it all into one configuration-behaviour table.

T-switch state-behaviour mapping

Addressing the T-switch first, we summarise its behaviour in response to the state of the RK-switch in Table 6.4. Throughout this PhD, we assume that the timescale of the RK-switch transitions is slow enough that the T-switch always has enough time to transition state before the RK-switch transitions again. This is equivalent to assuming that the T-switch will always turn off if the RK-switch is off, and conversely that the T-switch will always turn on if the RK-switch is on. Note how the state of the G-

switch does not affect the behaviour of the T-switch due to the nature of the ODE for the state variable D, which depends solely on the RK-switch state.

If we were to remove this simplifying assumption, we'd have to consider the dependence of the T-switch on the frequency of the RK-switch, which we discussed in section 5.2.2.2.

Table 6.4 Behaviour of the T-switch for different RK-switch states

RK-switch state	T-switch behaviour
0	R
1	D

G-switch state-behaviour mapping

Next, we can take on the G-switch. We identified regions of the κ_B, δ_B plane in section 5.2.4 which show a particular behaviour in response to the RK- and T-switch states. For the parameter values we're using in our example (P1 parameterisation, defined in Appendix J), and considering the four different configurations of the RK- and T-switches, we obtain the nine regions shown in Figure 6.5. In the example we're building here we will focus on the first region only, highlighted in blue with a thick border in Figure 6.5.

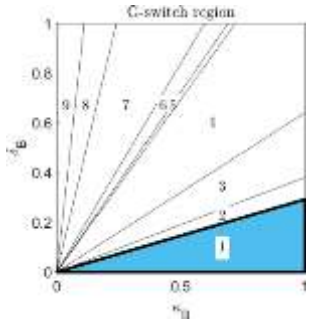


Figure 6.5 Regions of the G-switch bifurcation diagram which produces the state-behaviour relationship shown in Table 6.5.

We can summarise the behaviour of the G-switch in this particular region of the κ_B, δ_B plane in response to the states of the RK- and T-switches as in Table 6.5. Since the G-switch depends on the state of two other switches, we need to specify all four possible configurations. In this example case, the table isn't particularly exciting as the behaviour is the same for all the various switch states listed. In general, the behaviour may vary for each of the switch states listed.

Table 6.5 Example behaviour of the G-switch for different RK-switch and T-switch states for region 1 of the parameter plane κ_B, δ_B shown in Figure 6.5.

RK-switch state	T-switch state	G-switch behaviour
0	0	R
0	1	R
1	0	R
1	1	R

RK-switch state-behaviour mapping

For all values of the parameters κ_B and δ_B which fall within the region labelled "1" in Figure 6.5, the G-switch will behave as specified in Table 6.5. However, these parameter values κ_B and δ_B also have an effect on the RK-switch behaviours, since they are able to move the BE bifurcation lines in the κ_P, α_P plane. This means that, as we move around the κ_B, δ_B plane, the regions on the RK-switch

bifurcation diagram can change size. Therefore, before we turn our attention to working out how the switch states map to the RK-switch behaviours, we must define values for κ_B and δ_B.

The most sensible way to define these values is simply to take the centroid value of κ_B and δ_B (i.e. the point on the bifurcation diagram which is inside the polygon region we're considering, but also furthest from all of the edges). Using these centroid values for κ_B and δ_B, we calculate the bifurcation lines for the RK-switch, which separate the κ_P, α_P plane into regions as we described in section 5.2.4. In our example, the RK-switch bifurcation diagram features 5 regions as shown in Figure 6.6. For our example, we're going to focus on region 1, highlighted in blue with a thick border in Figure 6.6.

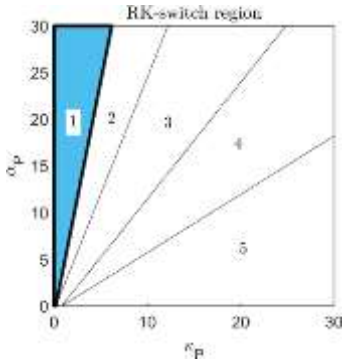


Figure 6.6 Regions of the RK-switch bifurcation diagram where the state-behaviour relationship between the RK- and G-switch are uniform across each region.
Region 1 produces the state-behaviour relationship shown in Table 6.6.

For each region, we may summarise this mapping between the G-switch state and the RK-switch behaviour for each region of the κ_P, α_P plane using a table containing a column for the G-switch state, and another for the RK-switch behaviour produced for that configuration. Note that since the state of the T-switch has no effect on the behaviour of the RK-switch, there is no need to include a column for the T-switch state. The result is a table consisting of two rows, listing the RK-switch behaviour for G-switch off and G-switch on.

In the example we're building, the table corresponding to region 1 in Figure 6.5 is summarised in Table 6.6. In this particular example, the RK-switch behaviour remains "recovery" no matter the state of the G-switch. In general, the RK-switch behaviour can change when the state of the G-switch changes.

Table 6.6 Example behaviour of the RK-switch for different G-switch states.
Corresponds to region 1 of the parameter plane κ_P, α_P shown in Figure 6.6.

G-switch state	RK-switch behaviour
0	R
1	R

When we take any pair of the regions which we identified from the κ_P, α_P and κ_B, δ_B planes, we obtain a complete picture of the behaviour of all the switches for any configuration of the switch states. In this example, we're going to combine the region labelled "1" in Figure 6.5 and the region labelled "1" in Figure 6.6.

We can summarise all of the state-behaviour relationships into a single table by simply merging Table 6.4, Table 6.5 and Table 6.6 into one longer table, which we call the "configuration-behaviour table". In our example, we obtain Table 6.7.

Each row of the table simply combines the information from each of the smaller tables. For example, the switch states {RK-switch=0, T-switch=1, G-switch=1} tells us that the RK-switch is "recovery" R, corresponding to the first row of Table 6.6, while the T-switch behaviour is R, corresponding to the

first row of Table 6.4. Finally, the G-switch behaviour is R, corresponding to the second row of Table 6.5.

Table 6.7 Example of a configuration-behaviour table summarising the effects of crosstalk between the three switches in the example region of parameter space.

$\boldsymbol{C}$			$\boldsymbol{b}$		
RK-switch state	T-switch state	G-switch state	RK-switch behaviour	T-switch behaviour	G-switch behaviour
0	0	0	R	R	R
0	0	1	R	R	R
0	1	0	R	R	R
0	1	1	R	R	R
1	0	0	R	D	R
1	0	1	R	D	R
1	1	0	R	D	R
1	1	1	R	D	R

The configuration-behaviour table defines the mapping $\boldsymbol{C} \mapsto \boldsymbol{b}$ for a particular combination of regions in the G-switch and RK-switch bifurcation diagrams. We can express the information in this table as a set diagram (Figure 6.7). In this example, which is a particularly simple case where the switches all tend towards full recovery, most of the possible behaviour configurations $\boldsymbol{b}$ are not mapped onto, and so they simply pile up unclaimed on the right of the figure. The remaining two behaviour configurations $[R, R, R]$ and $[R, D, R]$ are each mapped onto by four switch state configurations.

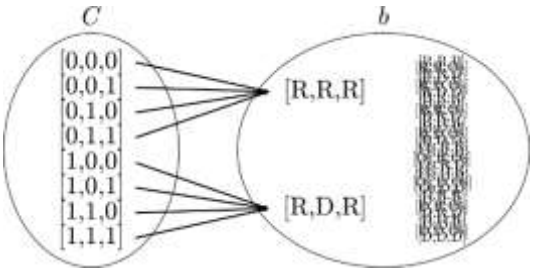


Figure 6.7 Mapping from C to b expressed in the configuration-behaviour table in Table 6.7.

This mapping summarises the effect of the crosstalk between the three switches, condensing all of the information about how the location of the stationary points of the ODE move across switch thresholds as the switch states change. We assume that the mapping we define is uniform across the volume of parameter space spanned by the regions used to construct it.

There is a slight complication caused by the effect of the parameter values κ_B and δ_B on the RK-switch behaviour region, which will move location slightly as these parameters are varied. A more rigorous approach would sample lots of values for κ_B and δ_B inside the G-switch behaviour polygon and use each sample to generate its own RK-switch behaviour polygons. For the remainder of this thesis, we make the simplifying assumption that the centroid values of κ_B and δ_B to be representative of the whole region.

In the next section, we describe how we use this condensed information of the state-behaviour mapping to deduce which switch transitions occur spontaneously due to ODE dynamics, and when this mapping doesn't give us enough information to completely characterize the model's dynamics. We analyse the specific example in Table 6.7, and plot the resulting networks of switch state configurations.

6.3. Switch transitions

In this section, we examine how the crosstalk between the three switches, summarised by the configuration-behaviour tables we construct for each of the 215 parameter regions we have identified, can induce series of spontaneous transitions of switch state. We describe how to use the information in the configuration-behaviour table to work out which transitions take place, and what the initial conditions of the switches must be to induce each of the dynamics we predict. We describe how the initial conditions of the continuous-valued state variables affect the dynamics of the system, and how we account for this effect. Finally, we summarise our results as a network diagram, mapping all of the transitions in a way that can tell us something about the experience of an AD patient.

Predicting switch state transitions from configuration-behaviour tables

The four behaviour categories we defined in Chapter 5 summarise the location of the focal points of the ODE relative to their switch thresholds. Each category of behaviour (Recovery, Bistable, Oscillation and Damage) tells us how the switch will transition state in the future given its current state. We use exactly the same logic we described in section 2.3.2: "Inferring switch transition dynamics from focal point values" to infer the transition dynamics.

We examine the current state, and the current behaviour of each switch, and use this to infer the nature of any future transitions as summarised in Table 6.8. This table stell us the type of transition a switch must undergo to allow the ODE dynamics to reach their focal point.

Table 6.8 Transition table, summarising the inferences we can make about future switch transitions.

		Behaviour			
		Recover	Bistable	Oscillation	Damage
State	off	No transition	No transition	off → on	off → on
	on	on → off	No transition	on → off	No transition

We consider the configuration-behaviour table we have constructed and look up the particular combination of state and behaviour for each switch in each row of the table to construct a new table we call a "transition table". This table summarises all of the transitions which should take place for every possible configuration C of the switches in the model.

Continuing our example with the configuration-behaviour table in Table 6.7, and for just examining the first row where the switch states are [0, 0, 0] we can see that no transitions are predicted. We summarise this in Table 6.9.

Table 6.9 Predicted transitions (or in this case lack of transitions) of the switches for the switch state configuration [0,0,0].

Switch name	State	Behaviour	Predicted transition
RK-switch	0	R	No transition
T-switch	0	R	No transition
G-switch	0	R	No transition

This tells us that for the switch state configuration [0, 0, 0], the behaviour of every switch is "recovery". We conclude that the switch state configuration [0, 0, 0] is an absorbing configuration for the parameter region we're looking at. We illustrate this absorbing configuration with the diagram in Figure 6.8. We can conclude that if the model is parameterised with the values from this parameter region, and we set the switches to be off, then the focal point will lie in the region of state space such that all the switch states will remain in that configuration.

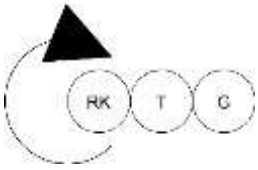


Figure 6.8 A stable absorbing configuration.

We repeat this process for the next switch state configuration in our table [0,0,1]. This means the RK-switch is off, the T-switch is off and the G-switch is on. In this case, we know the G-switch should transition from on to off because it is still in "recovery" behaviour. This means the switch configuration [0,0,1] must transition to [0,0,0] (which we just showed is a stable absorbing configuration). We can add this transition to the network diagram we're building up, shown in Figure 6.9.

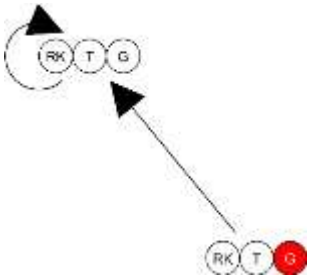


Figure 6.9 Switch transition from [0,0,1] to an absorbing configuration [0,0,0].

Table 6.10 Predicted transition of the switches for the switch states [0,0,1].

Switch name	State	Behaviour	Predicted transition
RK-switch	0	R	No transition
T-switch	0	R	No transition
G-switch	1	R	1→0

Let's skip forward a row to the switch state [0,1,1] to aid our explanation, and work out the predicted transitions. In this case, the RK-switch is off, T-switch on and the G-switch on. Checking the appropriate row of the configuration-behaviour table, we see that all three switches are in the "recovery" behaviour. This means that both the T-switch and G-switch are predicted to turn from on to off, as summarised in Table 6.11. This is a problem since it appears to contradict the deterministic nature of this model. One switch should change state before the other, and our predictions do not help us determine which transition occurs first. We address this point in the next section, but for now we simply note down all predictions and move on.

Table 6.11 Predicted transition of the switches for the switch states [0,1,1].

Switch name	State	Behaviour	Predicted transition
RK-switch	0	R	No transition
T-switch	1	R	1→0
G-switch	1	R	1→0

We can continue to work our way through the switch state configurations, noting the predicted transitions, producing Table 6.12, and adding each transition to the growing network diagram. At the end of this process, we end up with a table summarising the predicted transitions for every possible configuration of the switch states, and a network diagram summarising all of the switch transitions.

Table 6.12 Predicted transitions for the configuration-behaviour Table 6.7

Row number	RK-switch state	T-switch state	G-switch state	**Predicted transition**		
				RK-switch	T-switch	G-switch
1	0	0	0	-	-	-
2	0	0	1	-	-	1→0
3	0	1	0	-	1→0	-
4	0	1	1	-	1→0	1→0
5	1	0	0	1→0	0→1	-
6	1	0	1	1→0	0→1	1→0
7	1	1	0	1→0	-	-
8	1	1	1	1→0	-	1→0

Some switch state configurations in Table 6.12 have just one predicted transition while others have two or even three predicted transitions. The network diagram showing all of these transitions is drawn below (Figure 6.10). Clearly this transition diagram contradicts the fact that the model is deterministic, since we know that each switch configuration should have only one outgoing transition arrow.

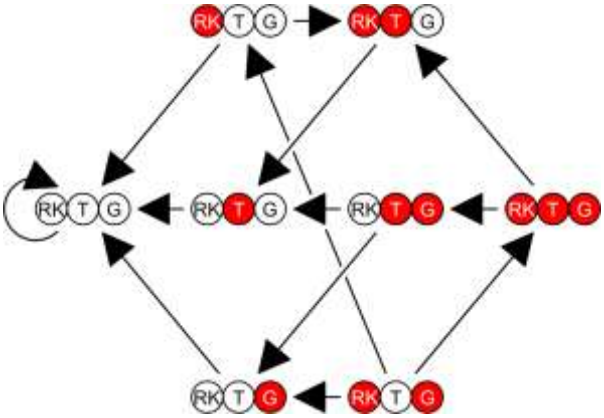


Figure 6.10 Transition network with every predicted transition.

6.4. Limitations of focal point analysis

The classification of each switch into the four categories; recovery, bistable, oscillation or damage are determined by focal point analysis, where we compare the focal point values $\bar{P}$, $\bar{D}$ and $\bar{B}$ relative to their respective switch thresholds. In general, however, this focal point analysis is not sufficient to find all regions of parameter space in which each behaviour can occur. Furthermore, once these behaviours have been determined for each switch and we predict multiple transitions to occur, we must consider which of the proposed switch transitions occurs. Both of these limitations emerge as a result of dependency on the initial conditions of the system. We discuss each limitation in the coming sub-sections 6.4.1, 6.4.2 and 6.5, where we lay out the geometry of the problem and potential numerical approach to solve it. We discuss this numerical solution to the problem in Appendix O. In this thesis, we work around this problem using a couple of simplifying assumptions, which we describe in section 6.5.2.

6.4.1. Initial condition dependence of switch behaviours

The trajectory of the ODE system for certain initial conditions may cross some of the switch thresholds as the system converges towards the equilibrium point, allowing an additional behaviour to coexist alongside that which we already identify. In this section, we give two hypothetical examples where the initial conditions would cause the focal point analysis to miss some more exotic types of switch behaviour. Neither example is backed up by numerical analysis – they're intended merely as a thought experiment to demonstrate how some more exotic behaviours can emerge.

Example 1 – Bistable after transition

Figure 6.11 shows sketches of the vector fields defined by the ODE in the (P, B) plane when the RK-switch is off (a) and on (b). In each case, a region of the plane consists of ODE solutions which cross one of the switch thresholds for the RK-switch, inducing an exotic behaviour which we would not spot by simply using our focal point analysis method.

We are able to identify these exotic behaviours by drawing the ODE solution that just skims the switch thresholds P^+ and P^-, intersecting them at a point where the vector field's gradient is parallel to the plane representing the threshold value of the switch. In Figure 6.11 (a) and (b), these special solutions are illustrated by a black curve, which converges towards the focal point represented by a green circle. Note how the intersect the threshold lines P^- and P^+ at parallels.

In Figure 6.11 (a), the focal point analysis would predict that if the RK-switch is already off, it would stay off because the value of the focal point (green circle in Figure 6.11 (a)) with the RK-switch off lies below the "switching on" threshold $\bar{P}(0,*,*) < P^+$. However, we can see that for any initial conditions in the green region below the solution curve, the RK-switch would actually transition from off to on because the ODE dynamics from these points must cross the "switching on" threshold P^+ on their way to the focal point $\bar{P}(0,*,*)$.

Similarly, in Figure 6.11 (b), focal point analysis would predict that if the RK-switch is already on, it would stay on because the focal point value (green circle in Figure 6.11 (b)) with the RK-switch on lies above the "switching off" threshold $\bar{P}(1,*,*) > P^-$. However, for the initial conditions in the green region above the solution curve, the RK-switch would actually transition from on to off because the ODE dynamics for these points must cross the switching off threshold on the way to the focal point $\bar{P}(1,*,*)$.

This type of initial-condition-dependent switch transition can be detected by looking for points on the threshold lines P^- and P^+ where the vector field of the ODE lines parallel to the threshold line. Such points are to the location in the vector field where the ODE solutions just skim the threshold, defining the top and bottom corners of the green regions in Figure 6.11 (a) and (b) respectively. We could then reconstruct the shape of these green regions by evaluating the ODE solution as we have illustrated in Figure 6.11 (a) and (b), integrating backwards in time, moving away from the focal point value until we collide with the limits of our axes. This would give us a series of points defining the curved edge of the green regions. This process would allow us to numerically define the shape of the green regions in Figure 6.11 (a) and (b), from which we can reconstruct Figure 6.11 (c) where we know the full extent of the exotic behaviours B* and B** in the PB plane. Throughout this thesis, we do not attempt to solve this problem by integrating the ODE to find these theorised problematic regions, but rely on the focal point analysis, simply bearing in mind the limitation we describe here.

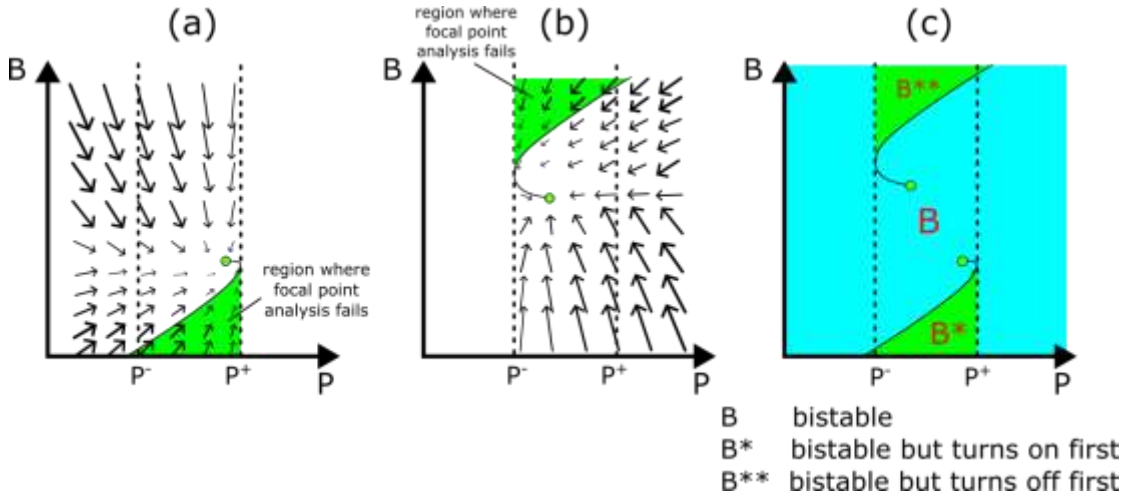


Figure 6.11 Vector field of the ODE showing regions where exotic behaviours emerge (a) In the red region below the black solution curve, the ODE dynamics cause the RK-switch to turn on. (b) In the red region above the black solution curve, the ODE dynamics cause the RK-switch to turn off. (c) The emergent behaviours. The blue region shows where the system is bistable as normal, but the green regions show a "bistable after transition" behaviour B and B**.*

In the example shown in Figure 6.11, the behaviour predicted by focal point analysis would be "bistable", but would be subtly incorrect only for the range of initial conditions that fall within the green shaded regions in Figure 6.11 (c). Rather, these two green regions are "bistable after transition". For the B* behaviour, when the switch is off, it will first turn on. After this transition is complete, the switch becomes bistable because the ODE dynamics must have entered the blue region of Figure 6.11 (c), where the rest of the ODE dynamic reaches the focal point without triggering another transition. For the B** behaviour, when the switch is on, it will first turn off. After this transition is complete, the switch becomes bistable for the same reason.

Example 2 Oscillation emerging from bistability

We can imagine a slightly more complex case, where the shaded regions feed into each other as shown in Figure 6.12. In this scenario the ODE dynamics overshoot the switch's threshold, causing a switch transition. This switch transition in turn leads to further ODE dynamics which cross the other switch threshold, transitioning the switch back to its original state where this oscillating process starts again. This effectively causes the "oscillation" Os behaviour to be produced, but in a way that would not be

detected by focal-point analysis because in this case it's caused by overshoot in the ODE dynamics rather than the location of the focal point.

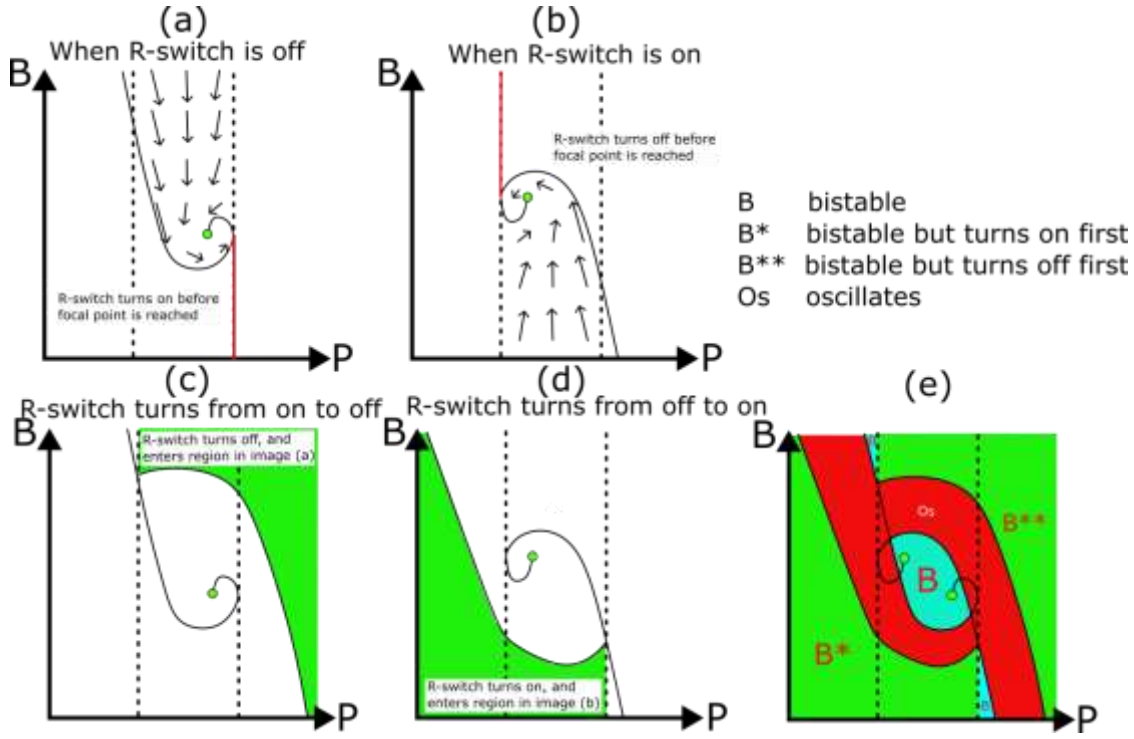


Figure 6.12 Possible formation of an oscillating behaviour for focal points that would otherwise be predicted to form a bistable behaviour. The red region corresponds to values of the state variables which lead to oscillating switch behaviour. Blue regions produce bistable behaviour as expected while green regions produce bistable behaviour after the switch changes state.

In Figure 6.12 (a), we have drawn the vector field of the ODE when the RK-switch is off. We can see that the region above the black line will converge to the focal point value (green circle) without crossing the switching on threshold P^+. Meanwhile, the region below the black line must cross the P^+ threshold before reaching the focal point value. This causes the R-switch to turn on.

Figure 6.12 (b), shows a similar drawing corresponding to the vector field of the ODE when the R-switch is on. Similarly, the region below the black line will converge to the focal point without crossing the switching off threshold P^-. Meanwhile, the region above the black line must cross the P^- threshold before reaching the focal point value. This causes the R-switch to turn off.

In Figure 6.12 (c), we have added a black line representing the ODE solution with the R-switch on where this solution collides with the switching off threshold and the black line from Figure 6.12 (a). All points above this new line must first cross the switching off threshold before the ODE dynamics proceed towards the focal point by following the vector field in Figure 6.12 (a). This special region is highlighted in green in Figure 6.12 (c).

Similarly, in Figure 6.12 (d), we have added a black line representing the ODE solution with the R-switch off, where this solution collides with the switching on threshold and the black line from Figure 6.12 (b). All points below this new line must first cross the switching on threshold before the ODE dynamics proceed towards the focal point by following the vector field in Figure 6.12 (b). Once again, this special region is highlighted in green in Figure 6.12 (d).

By combining the images from Figure 6.12 (a-d), we can construct the "oscillation doughnut" of Figure 6.12 (e), which shows us where the unexpected behaviours B*, B** and Os must occur in addition to the conventional bistable B behaviour predicted by focal point analysis.

The examples we have just described are both characterised by ODE solutions which cross the switch thresholds before they reach the focal point. In cases where the ODE solutions have this property, the question of how the switches behave becomes dependent on the initial-condition of the continuous-valued state variables, and as we have just shown, additional exotic switch behaviours such as

"bistable after transition" (B* and B**) and the unexpected "doughnut" of oscillation (Os) can be induced as a result. These additional switch behaviours can only ever occur in addition to the already identified behaviours (such as the bistable B phenotype predicted by focal point analysis in our examples in Figure 6.11 and Figure 6.12).

Even in cases where these additional behaviours (B*, B**, doughnut Os) occur, the region of the state space where initial conditions induce them only represent a fraction of the possible initial conditions, so we may at worst miss an additional behaviour which occurs for a certain set of initial conditions. Additionally, the region of the state space containing these initial conditions which lead to the alternative behaviour may be strangely shaped (see the "doughnut of oscillation" in Figure 6.12 (e)), and thus not straightforward to calculate.

Previous work [3] showed that for the double-switch model, the region of parameter space where multiple switch behaviours coexist was fairly small in extent. The triple-switch model is very similar in structure with regard to the RK-switch and G-switch, so we are confident that equilibrium point analysis will miss only a small extent of parameter space in which these more exotic behaviours occur in addition to the behaviours identified by our focal point method. The information about switch behaviours which we obtain using focal point analysis is likely close enough to the complete version that it makes no difference in our inferences here.

Future work could identify the extent of this region featuring coexistence of oscillation and bistable behaviours for this triple-switch model. For the remainder of this thesis, we use the focal point analysis method, but bear in mind that this is a simplification.

6.4.2. Initial condition dependence of transitions

In the example we've been building so far, all of the predicted swich transitions were based on knowledge of each switch's behaviour. We summarised all of our predicted transitions in our transition table (Table 6.8), which shows four switch configurations where multiple switches are predicted to transition.

The triple-switch model is deterministic, so any set of initial conditions and parameters must lead to a single solution to the model's ODEs, which we can think of as a line of points through the continuous-valued state space (P, B, A, D). The switch transitions occur when this line of solution points collides with the switch threshold values $P^-, P^+, B^-, B^+, D^-, D^+$.

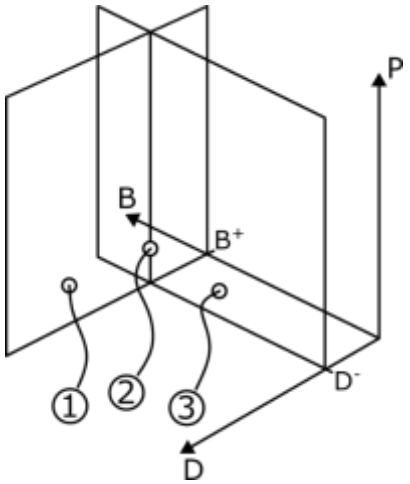


Figure 6.13 Three-dimensional representation of the continuous-valued state space (neglecting the state variable A for visual simplicity).

To gain a little intuition for the switch transitions are determined by the ODE solution, we can illustrate the solution in three-dimensional (P, B, D) space, (neglecting to draw the remaining continuous-valued

state variable A). In this view, the threshold values of each switch can be seen as 2-dimensional planes, and ODE solution will cause a switch transition when it collides with one of these planes.

Returning to the example in Table 6.8 where the switch state configuration results in two predicted transitions, we can visualise how the two possible switch transitions (either T- or G-switch) is selected depending on the initial conditions of the ODE. In Figure 6.13, we have drawn the (P,B,D) space with the two relevant switch threshold planes, the green plane representing the B^+ threshold to switch off the G-switch and the red plane representing the D^- threshold to switch off the T-switch.

Initially, our switch configuration is [0,1,1], meaning that our continuous-valued state variables must be in the region where $B < B^+$ and $D > D^-$. We have illustrated three trajectories of solutions through this state space in Figure 6.13. For trajectory ①, the solution intersects only with the plane for B^+ (which turns off the G-switch), so the G-switch must turn from on to off. For trajectory ②, the trajectory intersects both planes at the line of points where the planes themselves intersect each other. In this scenario, both the G-switch and D-switch simultaneously transition from on to off. Finally, trajectory ③ intersects only the plane where $D = D^-$, causing the T-switch to turn off. In each case, the transition is affected by the initial conditions of the continuous-valued state variables.

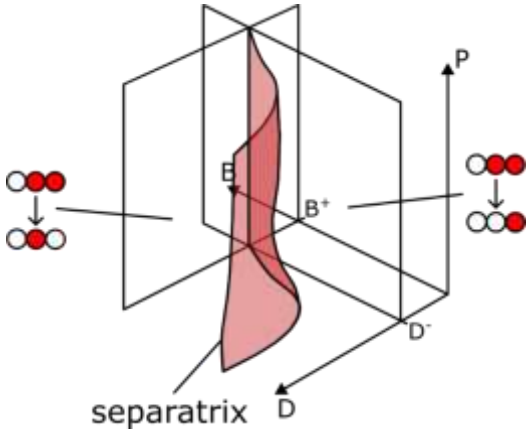


Figure 6.14 Basins of attraction for both predicted transitions from [0,1,1] →[0,1,0] and from [0,1,1] →[0,0,1].

The sets of points from which the ODE trajectory intersects both planes simultaneously are a separatrix, separating the basins of attraction for each predicted transition. These basins of attraction and their separatrix are illustrated in Figure 6.14. The points for which the trajectory only intersects only the B^+ or D^- plane form a three-dimensional volume in the (P, B, D) space, but form a four-dimensional volume when we also consider the continuous-valued state variable A. The points from which the trajectory intercepts both the B^+ and D^- planes simultaneously would form a sheet in this three-dimensional space, and form a three-dimensional volume when we also consider the state variable A. It is in principle possible to numerically calculate the shape of this separatrix as detailed in Appendix O, but we do not make the calculation in this thesis.

6.5. Proposed solution for multiple predicted transitions

In section 6.4.2. we describe how the state space splits into basins of attraction for each of our predicted transitions, the shapes of which are determined by the solution of the ODEs. In this section, we consider how the shapes of these basins of attraction interact when a sequence of switch

transitions occurs. We will be able to form a picture of what the basins of attraction look like for each of the transition networks we predict.

To give ourselves a concrete example to work with, consider the configuration-behaviour table we've been building throughout this chapter. The transition network is illustrated in Figure 6.15, with all the switch configurations predicting multiple transitions highlighted in blue.

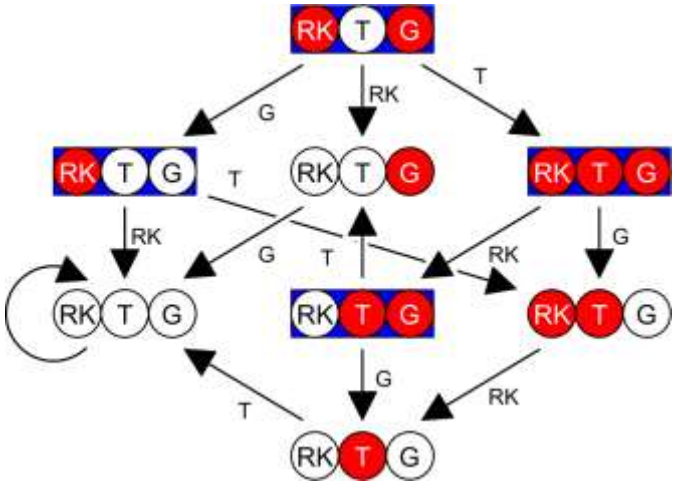


Figure 6.15 Transition diagram showing every predicted transition. Multiple-prediction switch configurations are highlighted in blue. Each transition is labelled with the switch that undergoes the transition.

6.5.1. Sequential switch state configurations with multiple predicted transitions – the full solution

We can imagine passing through the network shown in Figure 6.15, travelling from the switch state configuration with no entry transitions [1,0,1]. We chose this switch state configuration as our start point because it's the furthest removed from the absorbing configuration [0,0,0]. Then we can imagine transitioning through all the intermediate states in the direction of the arrows until we reach the absorbing configuration [0,0,0].

Clearly whenever we reach one of the multiple-prediction switch state configurations (highlighted in blue), we effectively reach a fork in the road and end up having to decide which of the outward transition arrows to take, which increases the number of possible ways to travel through the network. We can work out every combination of different decisions at these forks in the road to build a proposed list of the ways the ODEs might propel us through this network of switch state configurations. Taking every possible combination of decisions at forks in the path, we find that there are six different proposed ways to make these decisions at forks in the path and these are illustrated in Figure 6.16.

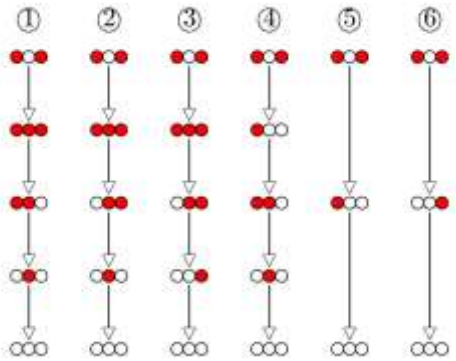


Figure 6.16 Six possible paths through the network in Figure 6.15.

Given that our model is deterministic in nature, we might naïvely assume that each of these possible pathways must be associated with a basin of attraction in the state space of the model. This logic seems quite appealing if we imagine the basin of attraction to be simply a volume of the continuous-valued state space, in which we could consider a sort of jigsaw puzzle of basins, which share boundaries at the separatrices, like the illustration in Figure 6.14. However, this is actually an

incomplete picture of the basin of attraction. Since our model is a hybrid dynamical system, the basins of attraction look a little more complicated than this picture.

Each basin of attraction is actually a sequence of regions in the continuous-valued state space, each connecting at an interface surface on one of the switch threshold planes. We have no guarantee that the basin of attraction for consecutive switch states in the paths we predict in Figure 6.16 actually do meet on these planes. Thus, it is possible that some of the predicted paths through the network are not possible in practice. To work out which paths through the network occur, and which do not, it would be necessary to numerically calculate the basins of attraction to check where they intersect at the switch threshold values.

In Appendix O, we propose a method to calculate the basins of attraction for arbitrary paths through a transition network like those in Figure 6.16. However, this method is so complicated in practice that we do not apply it here. Appendix O is intended as a start point for future work which may wish to calculate the relative size of each basin of attraction and rigorously check which pathways through a transition network can occur.

In considering the full complexity of the basins of attraction problem, defining what these separatrices could look like and then laying out a method in Appendix O to calculate the separatrices, we catch a glimpse of how complicated a problem this is. The next section discusses how we proceed without the detailed information about the separatrices of the system. The assumptions we use to work around this problem allows us to infer the approximate dynamics predicted by our model. We can use our approximated dynamics to assess whether it is necessary to calculate the separatrices in detail and show us where to apply the more complex method if required.

6.5.2. How to proceed without calculating the separatrix

In this thesis, we do not compute the shape or size of each of these basins of attraction (as described in Appendix O) due to the restrictive complexity of that task. It is prudent to first assess what we would gain from going into that level of detail. We make this assessment by building a list of the networks that are theoretically possible if the basins of attraction fulfil two assumptions. These assumptions are as follows.

Assumption 1: Each predicted transition has a basin of attraction, each of which occupies approximately equal volume in the state space.

Assumption 2: The separatrices chop up the continuous-valued state space in such a way that each possible sequence of transitions occurs with approximately equal frequency for a randomly chosen set of initial conditions in our specified admissible range.

Taken together, these assumptions mean that each switch state network we predict occurs with approximately equal frequency for randomly chosen initial conditions within the admissible range we specify.

These assumptions have been made because they are the most parsimonious given our lack of knowledge to the contrary.

Future work could calculate the separatrices numerically using the procedure described in Appendix O, allowing us to drop these heavy assumptions. These separatrix shapes could then show how the relative sizes and shapes of the basins of attraction for each possible transition might affect the

likelihood of seeing each of the switch state networks we produce here, which could subtly change the pathways each patient takes through the network of switch states.

6.5.2.1. Racing switches method

For each switch state where we predict more than one transition, we can think of the problem as that of "racing" the state variables for those switches towards their corresponding threshold values. The continuous-valued state variables corresponding to each switch (RK-switch corresponds to P, T-switch to D, and G-switch to B) race each other towards their respective threshold values, with the winner determining which switch changes state first.

Returning to the example we've been building up throughout this chapter, we can illustrate this "racing" idea for one of the multiple-prediction switch configuration, namely the state [1,1,1]. In this scenario, both the RK-switch and G-switch are predicted to transition from on to off so there is an ambiguity as to which one actually makes the transition.

We consider a hypothetical point in the continuous-valued state space, where we numerically integrate forwards in time. The outcome of this numerical integration will tell us which of the state variables reaches their respective switch threshold first. We can imagine a scenario where the dynamics of P are more rapid than those of B, causing P to cross the switch off threshold for the RK-switch (P^-) before the state variable B crosses the switch off threshold for the G-switch (B^+). In this case, the RK-switch would transition from on to off, and we'd end up with a new switch configuration of [0,1,1]. Alternatively, if the timescale of the dynamics of B are faster than those of P, then the opposite can occur, causing the G-switch to transition from on to off, transitioning us to the new switch configuration of [1,1,0]. Both of these scenarios are illustrated in Figure 6.17. Effectively, the state variables P and B can be thought of as "racing" each other towards their respective threshold values. The winner of the race gets to transition first.

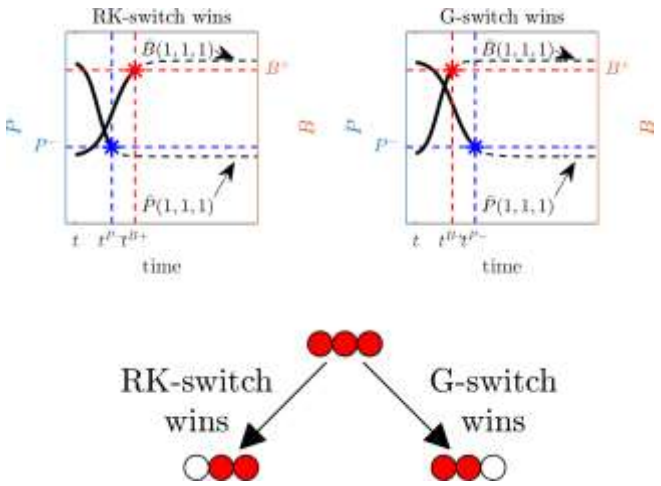


Figure 6.17 Timescale of dynamics of P and B can determine which switch transitions. If the dynamic of P is faster than that of B, the RK-switch wins the race, causing the RK-switch to transition. If the dynamic of B is faster than that of P, the G-switch wins the race, causing the G-switch to transition.

Framing the problem as a "race" in this way allows us to consider the potential outcomes more easily for networks which feature multiple-prediction switch configuration. Assuming each continuous-valued state variable is associated with a particular timescale and that these timescales determine which switch wins race, we can enumerate the permutations of outcomes in all the races found in a particular network.

In situations where just two switches race each other, clearly there are only two different orders in which those two switches can be ranked in the races (2P2=2). In cases where all three switches are

involved in races, the number of ways in which those switches can be ranked increases to 6 (3P3=6). Each of these six variations is listed in Table 6.13.

Table 6.13 Permutations of switch rankings in races where all three switches compete to reach their threshold values.

Variant	Placement in race to threshold		
	RK	T	G
1	1	2	3
2	2	3	1
3	3	1	2
4	1	3	2
5	3	2	1
6	2	1	3

For each parameter region, we construct a configuration-behaviour table such as that shown in Table 6.8. Then we count how many of the switches appear in races anywhere in the table and use this to calculate the number of permutations of ways the switches can rank in head-to-head races to their thresholds.[6] Each configuration-behaviour table therefore generates either 1, 2 or 6 switch state networks.

In the example we've been building through this chapter (Table 6.8), all three switches compete in races, so we can see the configuration-behaviour table will generate 3P3=6 networks. Notice this same example also ended up generating 6 pathways through the network in Figure 6.16. This is no accident. Each pathway through the network is associated with a separate basin of attraction, each of which causes the switches to rank differently in head-to-head.

Having worked out how many switch state networks we can produce, we work our way through the different permutations of switch rankings we've worked out, noting down the rank of each switch in our transition table. For each multiple-transition switch state (i.e., switch state producing a "race"), we remove predicted transitions which do not come "first" in their race. The winning switch in any multiple-transition switch state is simply the switch with the best rank, not necessarily the switch ranking 1st. For example, we have reproduced our table of predicted transitions below in Table 6.14 and assumed for this example the following switch ranking: RK-switch 1st, T-switch 3rd, G-switch 2nd. Note how for the switch configuration [0,1,1], the G-switch "wins" the race because the RK-switch is not competing. Therefore, 2nd place allows the G-switch to win because it's a better rank than the T-switch's 3rd place.

[6] Note that it is not necessary for all three switches to appear in the same head-to-head race for our behaviour-state table to count as having three competing switches. It is sufficient for each of the three switches to appear in separate races in pairs. For example, one switch state produces a race between the RK- and T-switch, a second switch state produce a race between the T- and G-switch, and then a third switch state a race between the G- and RK-switch. This configuration would produce 3P3=6 variations, despite none of these races involving all three switches at once.

Table 6.14 Table of predicted transitions given the assumption that the state variable P has the fasted time scale, followed by B, followed by D. This means the switches win head-to-head races in the order RK-switch 1st, G-switch 2nd, T-switch 3rd.

RK-switch state	T-switch state	G-switch state	**Predicted transition**		
			RK-switch (1st)	T-switch (3rd)	G-switch (2nd)
0	0	0	-	-	-
0	0	1	-	-	1→0
0	1	0	-	1→0	-
0	1	1	-	1→0 (loser)	1→0 (winner)
1	0	0	1→0 (winner)	0→1 (loser)	-
1	0	1	1→0 (winner)	0→1 (loser)	1→0 (loser)
1	1	0	1→0	-	-
1	1	1	1→0 (winner)	-	1→0 (loser)

Finally, after this assumption about the timescales of each switch is applied, we are in a position to draw our fully defined transition network. For the predictions listed in Table 6.14, the resulting network of switch configuration is as in Figure 6.18.

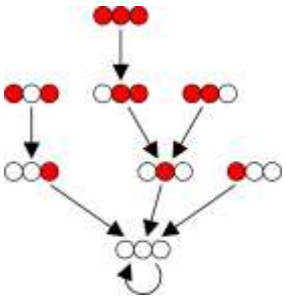


Figure 6.18 Transition network for the example parameter region with the assumed switch ranking {RK=1st, T=3rd, G=2nd}.

The example we've been working through has three switches competing, so there are 3P3=6 permutations of orders in which they can win races. Each of these permutations corresponds to its own basin of attraction and leads to a transition network. All six resulting networks are shown in Figure 6.19.

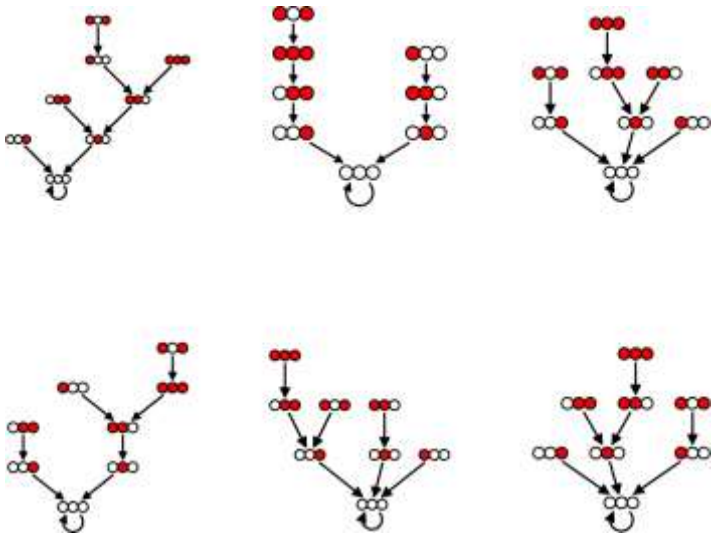


Figure 6.19 All six transition networks for the example parameter region, each corresponding to a permutation of the switch ordering.

Each parameter region we identify is associated with one configuration-behaviour table, each of which produces a "family" of transition networks like that shown in Figure 6.19. Each of these families consists of either 1, 2 or 6 "sibling" networks, each sharing identical parameterisations but differing in their initial conditions for the continuous-valued state variables P, B, A and D since they all possess their own basin of attraction.

6.5.2.2. Two ways of seeing the same problem: timescale versus geometry of the state space

A final note for this section to make absolutely clear the limitations of our assumptions. The assumptions we've made here allow us to work out what the transition networks look like for each of the basins of attraction present in a parameter region. We've reframed the problem as one of "racing" simply to help us form the analogy of differing timescales so we can rank our switches in the order that they cross their respective thresholds. However, while this analogy is useful, it's helpful to remember that each of these "race" scenarios is actually associated with a basin of attraction, each transition diagram being associated with a series of separatrices in the continuous-valued state space.

We are assuming the separatrices cut up the continuous-valued state space in such a way that each basin of attraction has approximately equal volume and that as a result each of the transition networks we draw here has equal likelihood of being produced given a randomly chosen set of initial conditions for our state variables. We have no guarantee that this assumption will hold, and so we suggest that future work should use the methods we describe in the abstract to determine the relative volumes of these basins.

6.6. Presenting the networks as families of trees

Now we've defined how to calculate the switch state networks for each parameter region, we can begin to organise the results and classify them to help us interpret some biological implications.

We gather "families" of networks which originate from the same parameter region (i.e., they all share the same parameter values but differ in the initial conditions of the continuous-valued state variables). Each network in this family has its own basin of attraction in the continuous-valued state space, which we assume to have nonzero volume. We have already shown that families can have either one, two or six members depending on the number of switches competing in races.

As was described in Figure 6.3, we find it helpful to plot all of our networks in the form of trees, where each tree consists of branches of transient switch configuration leading to roots consisting of an

absorbing configuration or loop. We group these families of networks together so that we can compare how these trees differ as we vary the basin of attraction into which the continuous-valued state variables fall while keeping the parameter values the same for each sibling of the family. One such family of networks is illustrated in Figure 6.20. In this example, each family member has two absorbing configurations, and the trees differ slightly in structure, but the absorbing configurations are always the same between siblings of the same family. This is a general result and can be deduced since the absorbing configurations are defined solely by the focal points of the model, which remain unchanged no matter which switches win in the "races".

Another interesting point to note is that sometimes, a pair of siblings can be identical. In Figure 6.20, trees (a) and (f) are identical, despite having theoretically different basins of attraction. From this, we can deduce that our method of "racing switches" described in section 6.5.2.1 in some way overgeneralises the problem and finds a separatrix where there is none. Sometimes, the transitions in a network are arranged such that the order of the switches timescales is irrelevant in determining the sequence of transitions to reach the absorbing configuration. We have not investigated this phenomenon further, but it's interesting to note the limitation of the "racing switches" method.

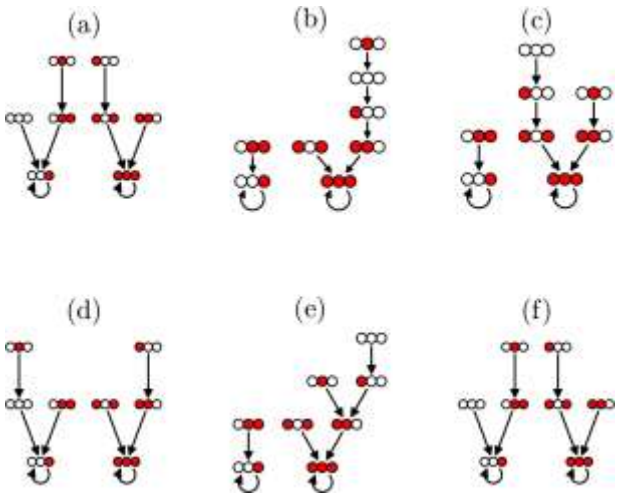


Figure 6.20 Example family of networks all sharing the same parameter values but differing in the basin of attraction the into which the state variables (P, B, A, D) fall.

We found 215 parameter regions in Chapter 5, each producing one family of trees. The number of siblings in each family was usually six, but some had only two siblings. The total number of transition trees produced by each parameterisation is listed in Table 6.15.

The 215 parameter regions we found had a total of 107 *different* configuration-behaviour mappings between them. This tells us that some of the configuration-behaviour mappings could be produced by multiple parameterisations. These 107 different configuration-behaviour mappings produced a total of 178 different transition trees. That's an average of 1.66 different transition trees per configuration-behaviour mapping, which tells us that there was significant overlap between the various structures of transition tree produced by each mapping.

Table 6.15 Number of transition trees

Parameterisation	P1	P2	P3	P4	**Total**	**Number of different types**
Number of configuration-behaviour mappings	45	54	58	58	**215**	**107**
Number of trees	234	288	312	312	**1,146**	**178**

6.7. Conclusion

In this chapter, we inferred how the three switches of our triple-switch model spontaneously transition state over time due to the dynamics of the ODE. We represented these transitions between different switch configurations using novel "transition tree" diagrams, which summarise the future dynamics of the model for all possible starting configuration of the system.

We discussed the limitations of our novel method of analysis, which is able to show us all possible transition trees for a given set of parameters but is unable to tell us how each transition tree relates to the initial conditions of our continuous-valued variables. Specifically, using our methods, we do not calculate the size of the basin of attraction corresponding to each transition tree we identify. In Appendix O, we propose a method to calculate these basins of attraction, but we do not carry out the calculation in this thesis.

The transition trees we presented in this chapter provide a map of the dynamics of the disease for a particular parameterisation. Like all maps, it represents a compromise between certain details, designed to offer the best view for a particular task. In our case, we neglect information about the timescale of the dynamics, but in doing so gain a picture of how the dynamics evolve from every possible starting configuration of the system, both for the switch state configuration (i.e. the different nodes in each tree) and for the continuous-valued variables (each sibling in a family of trees represents a different basin of attraction).

Chapter 7. Transition networks and treatments

In this chapter, we aim to identify how common treatments for AD are able to perturb the system towards healthier outcomes, and why in some cases the disease proves resistant to treatment. We aim to explain how these treatments work on chronic AD, and why they sometimes prove ineffective. To achieve this, we draw "transition networks", which characterise how the current state of the disease will evolve in the future, effectively drawing a map of all possible longitudinal dynamics predicted by our model. These diagrams help us fill a knowledge gap by providing a first mechanistic explanation of how the pathophysiological mechanisms are responsible for the complex disease dynamics.

Armed with these diagrams, we will interpret what these predicted dynamics mean for AD patients, and how various treatments affect the development of symptoms, We are able to provide an explanation for the variable efficacy of dupilumab treatment in cases of chronic AD.

7.1. Recap of the analysis steps to find the tree diagrams

The process of analysis we use to obtain the trees of switch state transitions has been gradually explained through Chapter 5 and Chapter 6. In this chapter, we begin by zooming out so that we can see the whole process from start to finish in one diagram (Figure 7.1). Armed with this macroscopic view, we discuss how we can normalise our results to obtain the relative likelihoods of seeing each of the tree diagrams given the parameter ranges we consider in Table 5.2.

We have broken down this analysis process into five stages as follows:

Stage 1. Separate the parameter space into four separate categories, based on the behaviours displayed in the bifurcation diagrams for the RK- and G-switches. Weight each of these four groups according to their relative frequency in a uniform random sample of the parameter ranges we defined in Table 5.2. This sampling procedure was described in section 5.2.3.4.

Stage 2. Calculate the bifurcation diagram for the G-switch and note the number of regions in the κ_B, δ_B plane that have uniform configuration-behaviour relationship. This process was described in section 5.2.4.1.

Stage 3. For each region in the G-switch bifurcation diagram, we calculate the centroid value for κ_B and δ_B and use these in our calculations for the RK-switch bifurcation diagram. We note the number of regions in the κ_P, α_P plane which have uniform configuration-behaviour relationship. This process was described in section 5.2.4.2.

Stage 4. Combine the individual configuration-behaviour relationships from each combination of regions from the RK- and G-switch bifurcation diagrams into one configuration-behaviour table. This process was described in section 6.2.

Stage 5. Calculate the number of competing switches and write the transition network for each permutation of winners in these switch races. We weight each network by the group weighting, region areas for the RK- and G-switch regions and divide by the number of siblings in the network family. This process was described in section 6.5.2.

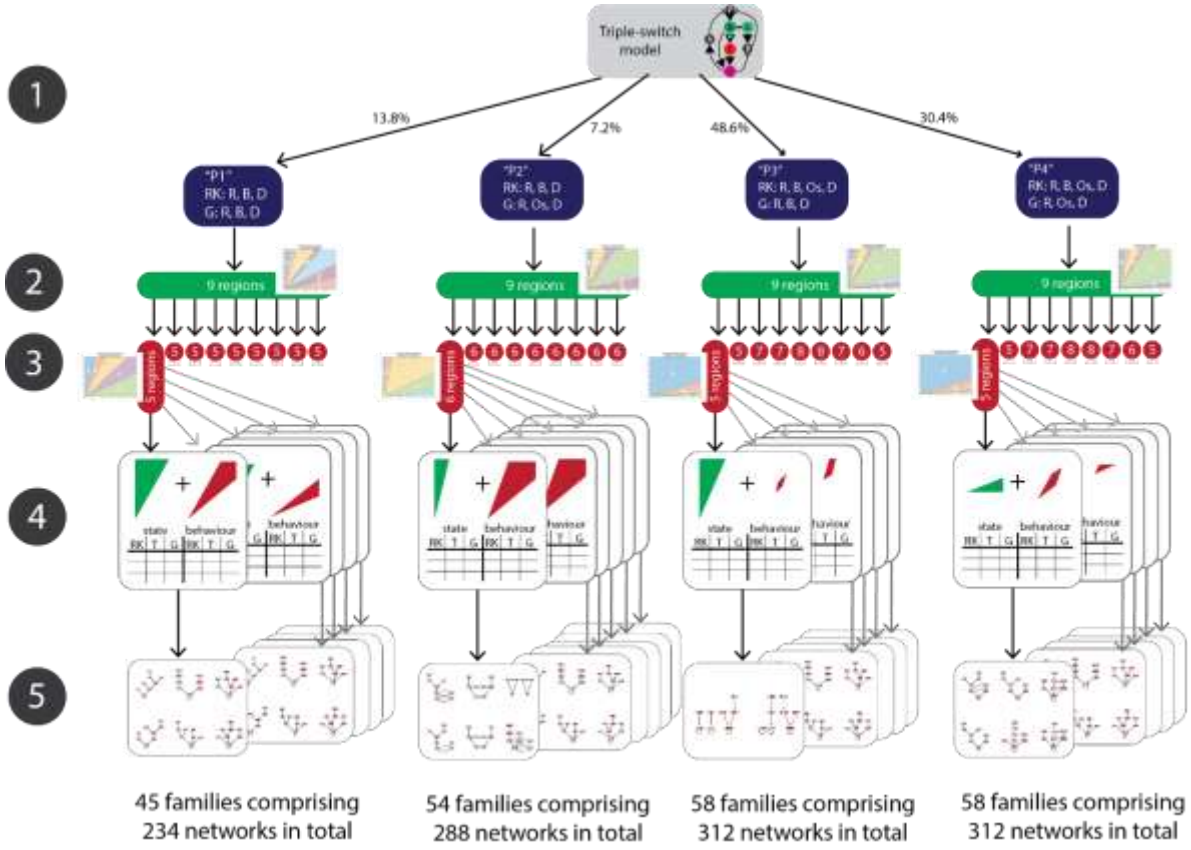


Figure 7.1 Complete workflow of our inference of switch state networks from the triple-switch model. Stages 1 through to 5 allow us to estimate the relative frequencies of each network of switch configuration produced by our model.

We wish to identify how the clinical phenotype emerges as a result of the mechanisms we include in our model, and how this is affected by genotype (i.e. the parameter values). For this reason, we keep track of which regions of parameter space generate each transition network, and the size of each of these regions. If we assume the genotypes of patients are uniformly distributed across the parameter ranges we consider, and so the volume of parameter space leading to each network approximately tells us how commonly that network should occur.

To place a number on the amount of parameter space corresponding to each network, we simply calculate the area of each polygonal region which we recorded for the RK- and G-switch bifurcation diagrams, which we call W_{RK} and W_G respectively. We then take the product of these two areas to obtain a sort of hypervolume, representing how much of the parameter space produces this family of networks. Next, we multiply this hypervolume by the percentage of the parameter space which comprises that parameterisation family, which we call W_p . These percentages for each parameterisation were calculated by randomly sampling the parameter space as detailed in section 5.2.3.4, Table 5.10. Finally, we divide this net weight for each region equally between the sibling networks produced in each family. We call the size of the family $N_{\text{family members}}$.

$$W_{\text{network}} = \frac{W_p W_{RK} W_G}{N_{\text{family members}}} \tag{7.1}$$

This rather crude measure is used to normalise all of the frequencies in the bar charts that follow in this chapter, where we compare the relative frequencies of different types of network.

7.2. Categories of switch configuration

In this section, we define three categories into which every switch configuration in our transition networks may be organised. Once we can define these features within each network, we can then break down our list of transition networks into descriptive categories.

7.2.1. Absorbing configurations

Any switch configuration where no more switch transitions occur. In our transition networks/trees, we represent these configurations with a curved arrow pointing back at the configuration, as the example in Figure 7.2.

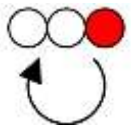


Figure 7.2 An absorbing configuration.

7.2.2. Absorbing loops

Any sequence of configurations that are linked by transitions to form a closed loop. Such a loop of configurations will repeat indefinitely once any of these configurations is encountered. An example is in Figure 7.9 (b).

The absorbing loops can consist of either two, four or six states. It is not possible to produce a period of any other length since this will always result in two switches being different for adjacent switch configuration in the cycle. Figure 7.3 shows a hypothetical loop with an odd number of configurations, where we end up with two switch configuration mismatching for one of the transitions. This violates our definition of a transition, where we assume only one switch changes at any instant in time.

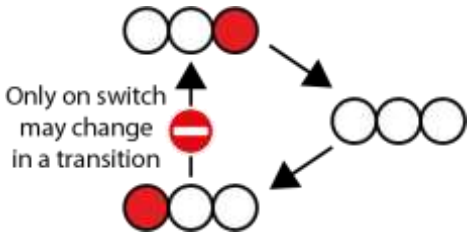


Figure 7.3 Hypothetical absorbing loop with an odd number of switches. This loop cannot occur.

There are multiple ways in which these absorbing loops can emerge, depending on the period of the loop. For loops of period 2 (i.e., which repeat after 2 transitions), one of the three switches may change state, with the other two remaining fixed in state. For oscillations of period 4, two switches alternate changing sate with the remaining switch fixed in state, while for oscillations of period 6, all three switches must be involved, transitioning in a repeating sequence. Each of these possible modes described below.

Period-2 loops

There are six ways this can occur, each one shown in Figure 7.4. Either the RK- or G-switch may oscillate independently of the other switches. The T-switch cannot oscillate independently due to the structure of the ODEs as discussed in section 5.2.2.2.

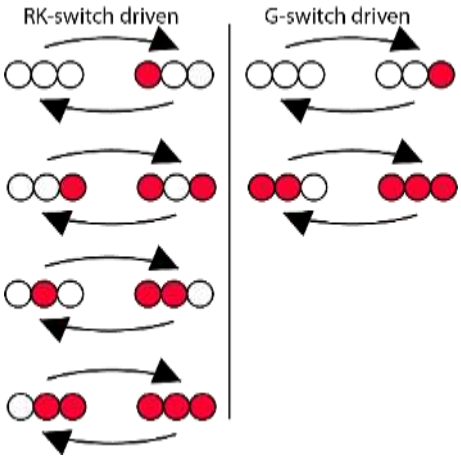


Figure 7.4 Period-2 absorbing loops.

Period-4 loops

There are four possible loops of period 4, each illustrated in Figure 7.5.

RK- and T-switch may trigger each other's transitions. In this situation, the G-switch may be either on, or off, contributing the 2 possibilities on the left of Figure 7.5.

It is also possible for the RK- and G-switch to trigger each other's transitions since the RK- and G-switch configuration each have the effect of moving the bifurcation lines for the other switch. This situation however is reliant on the frequency of the RK-switch oscillation being faster than the timescale of the dendritic cell death δ_D (i.e., months) which tells us how quickly the T-switch reacts to a transition of the R-switch. In this situation, the T-switch may be either on or off, contributing the 2 possibilities on the right of Figure 7.5.

It is not possible for the T- and G-switches to oscillate together since the state of the G-switch has no effect on the value of the state variable D, which determines the state of the T-switch.

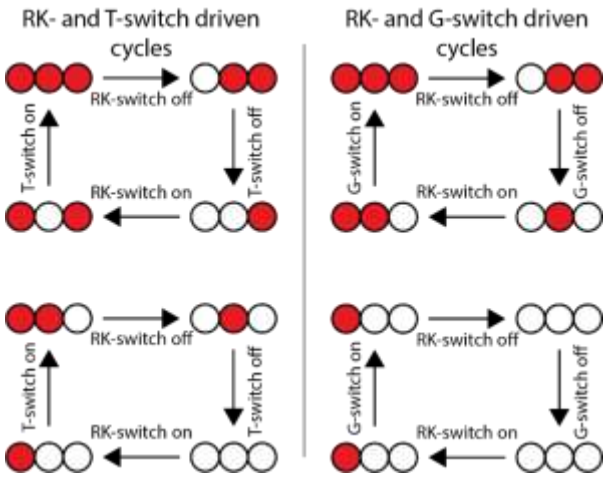


Figure 7.5 Period-4 absorbing loops.

Period-6 loops

There are two possible loops of period six, illustrated in Figure 7.6.

In these period-3 loops, all three switches trigger one-another in sequence. We can see that there are only two possible ways this can occur because there are only two possible ways that the three switches

can induce each other to transition, as shown in Figure 7.6. (a) shows the loop created when the RK-switch triggers the G-switch to transition, which triggers the T-switch to transition. (b) shows the loop when the RK-switch triggers the T-switch, which triggers the G-switch.

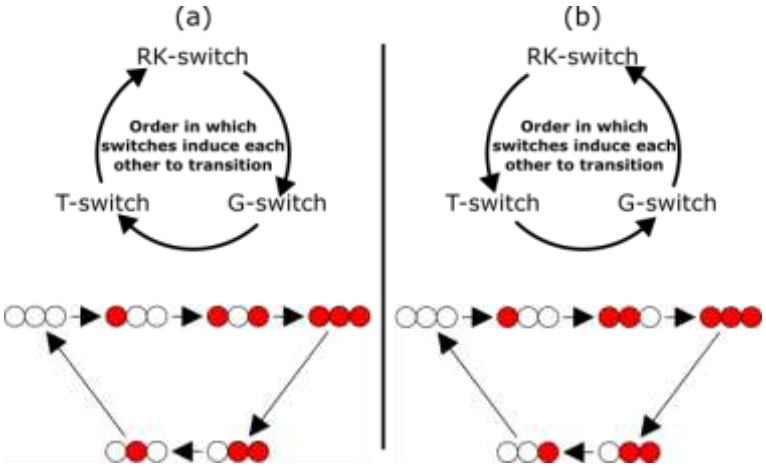


Figure 7.6 Period-6 absorbing loops.

In the loop of Figure 7.6 (a), we are assuming the G-switch responds more quickly than the T-switch, which means we are implicitly assuming the rate of differentiation of the keratinocytes in the granular layer of the epidermis is more rapid than the rate of dendritic cell buildup in the draining lymph-node.

7.2.3. Transient loops

There are two remaining loops of period 2, which are a special case due to the structure of the ODEs of this model. As discussed in section 5.2.2.2, the T-switch state must always transition to match the state of the RK-switch. In these remaining two loops of period 2, the state of the RK- and T-switch are different, so we know that eventually there must be a T-switch transition which takes us out of the loop.

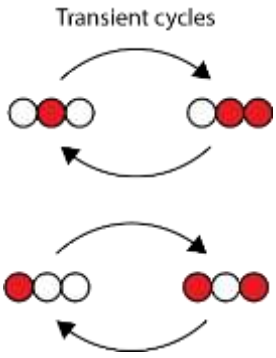


Figure 7.7 Transient loops.

There are also two period 2 oscillations which are unstable due to this effect, illustrated in Figure 7.7. These special varieties of loop can only occur if the timescale of the G-switch oscillation is shorter than the timescale of the T-switch change, δ_D. This relationship between the G- and T-switch timescales allows the G-switch to undergo transitions back and forth between the two states in the loop before the T-switch has a chance to make the transition out of the loop.

We know that the system must exit the loop on the timescale of the dendritic cell dynamics, δ_D since these determine how rapidly the T-switch responds to the RK-switch state. Using our focal point analysis method, we have no way of knowing which of the two configurations in the loop will produce the transition out of the loop. For this reason, we consider both possible transitions out of transient loops.

Where we find a transient loop in our transition networks, we denote the possible transitions out of the loop using blue arrows, as shown in Figure 7.8.

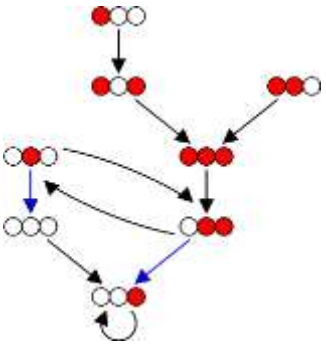


Figure 7.8 Exit transitions from the transient loop, represented by blue arrows.

7.2.4. Transient configurations



Figure 7.9 A transient configuration.

Any configuration which occurs in the transition dynamics before they reach an absorbing configuration or loop is a transient configuration. These configurations always appear in the "branches" of the transition trees. An example is in Figure 7.9.

7.3. Categories of disease network

Each of the transition diagrams we identified has particular features which are pertinent to our model's predictions about the lived experience of chronic AD, and our inferences of mechanisms and potential treatment strategies. We can use the categories of switch configuration we just defined to break down our list of transition diagrams into a few useful categories.

7.3.1. Networks with one absorbing configuration

The majority of the networks we produce have just one absorbing configuration upon which all of the transient dynamics converge. Of the 178 different transition networks we found, 76 of these feature a single absorbing configuration. Siblings in a family of transition trees always have the same absorbing configuration. Each sibling within the family has potentially different transient configurations leading to the absorbing configuration, although sometimes siblings within one family are identical. On such family is shown in Figure 7.10.

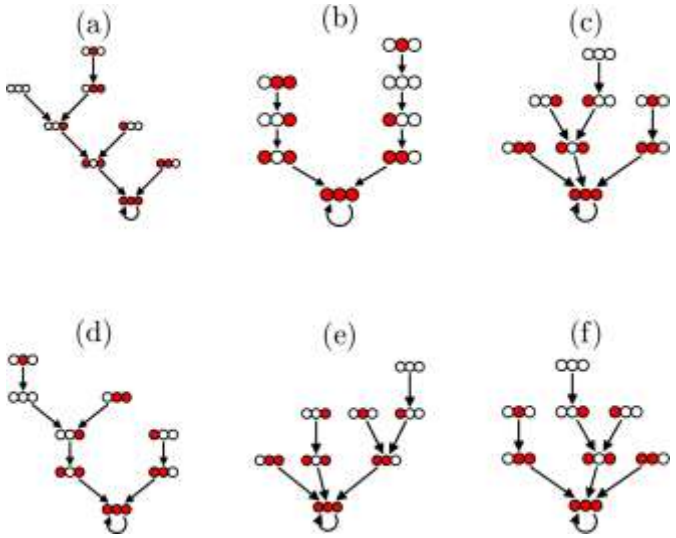


Figure 7.10 A family of transition networks where all siblings converge to the same absorbing configuration.

Recall that the siblings in a family all share the same parameterisation (i.e., the same genotype) but differ in the initial conditions of the latent variables P, B, A and D. For genotypes in this category, differences in the latent continuous-valued state variables cause the transient dynamics of chronic AD to differ even though the genotype does not change. From the perspective of a patient, events which

trigger a temporary flare or even treatments which cause a temporary respite of symptoms can induce different dynamics back to the absorbing configuration even in the same individual.

The absorbing configuration featured in this category of networks can be any of the following: [0,0,0] (full health), [0,0,1] (no AD flare, but impaired skin barrier), [1,1,0] (AD flare, but otherwise competent skin barrier), [1,1,1] (AD flare with impaired skin barrier – the worst-case scenario). Accounting for the are of the parameter space producing each network, the distribution of single absorbing configuration networks is shown in Figure 7.11

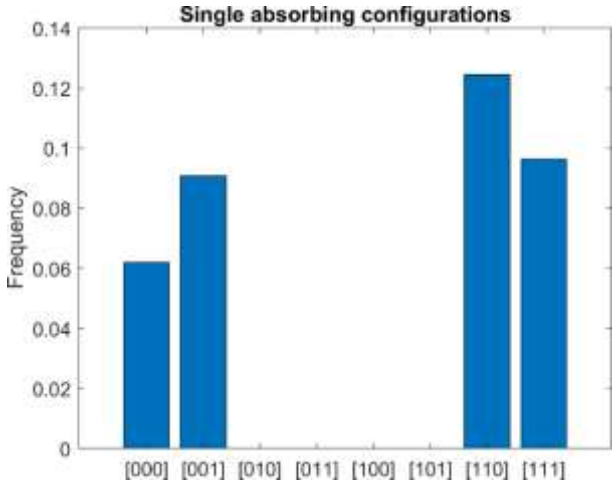


Figure 7.11 Distribution of absorbing configurations across all networks which have just one absorbing configuration.

The switch configurations [0,1,0], [0,1,1], [1,0,0] and [1,0,1] are never found to be absorbing configurations because for each of these configuration, the RK-switch and T-switch have opposite states. We described in section 5.2.2.2 why we assume the T-switch state always "follows" the state of the RK-switch. This means any configuration where the RK- and T-switches have different states is always guaranteed to lead to a transition so it can never be an absorbing configuration.

7.3.2. Networks with multiple absorbing configurations

Some networks feature more than one absorbing configuration. Of the 178 different transition networks we found, 62 of these feature multiple absorbing configurations. Siblings in a family of transition trees in this category always have the same absorbing configurations.

Every switch configuration can lead to at most one transition, so each transient configuration appears in only one tree leading one absorbing configuration. For this reason, each of the absorbing configurations has its own tree of transient configurations which all transition towards that final absorbing configuration. However, sibling networks in a family may have these transient configurations laid out differently. For example, Figure 7.12 (a) the transient configuration [1,0,0] leads eventually to the absorbing configuration [0,0,0] (full recovery) while in its sibling network Figure 7.12 (b) [1,0,0] leads to the absorbing configuration [1,1,0] (active inflammatory response). These two outcomes are quite different, and the only difference between the two sibling networks is the initial conditions of the latent variables (P, B, A and D).

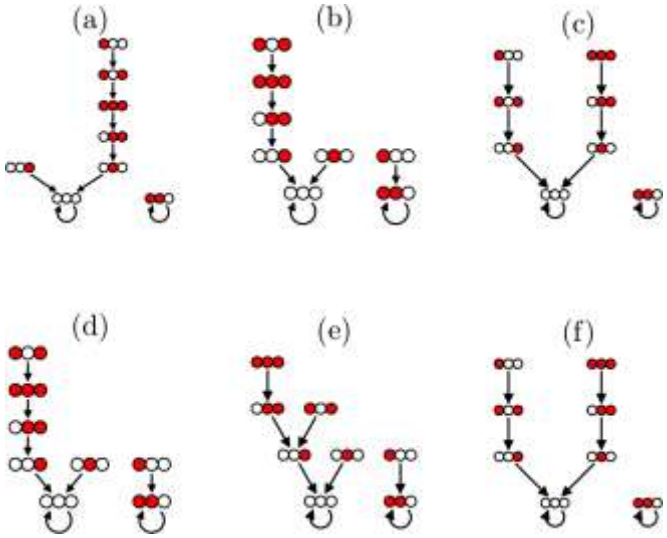


Figure 7.12 A family of transition networks featuring multiple absorbing configurations.

From the perspective of an AD patient with this genotype, this means the values of the latent variables can have a significant effect on the long-term prognosis of the AD symptoms. This may go some way to explaining why chronic AD symptoms can be so hard to predict and control.

The distribution of number of absorbing states in each network is distributed as in Figure 7.13. This distribution shows the proportion of the parameter space we searched which gives rise to each type of network. The weight for each network is calculated as described in equation 7.1. The networks with "0" absorbing configurations only produce absorbing loops and no absorbing configurations.

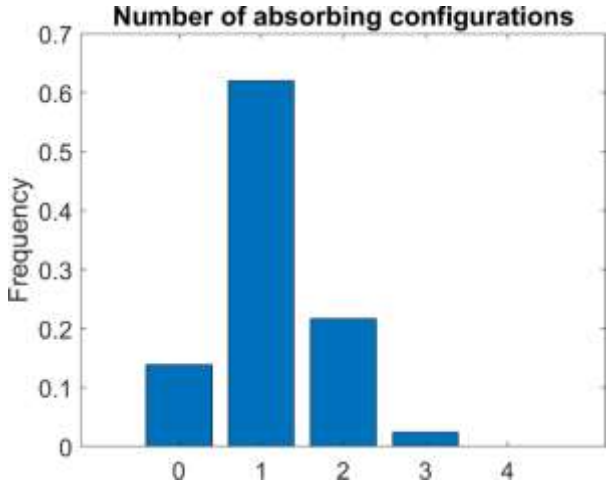


Figure 7.13 Number of absorbing configurations of each network, weighted by size of parameter region producing each type.

7.3.3. Networks with absorbing loops

A minority of networks include loops of absorbing configurations, where our model predicts dynamics that slowly change over time, producing such features as recurrent flare events and intermittent periods of impaired skin growth. Such predictions are consistent with the lived experience of chronic AD patients. Of the 178 different transition networks we found, 51 of them feature absorbing loops. Of these, 40 contained only absorbing loops, while the other 11 contained both absorbing loops and absorbing configurations simultaneously.

In each of the 40 networks where absorbing loops were the only possible outcome, siblings in the same family always have different absorbing loops as shown in Figure 7.14. In this example, the six

siblings only contain absorbing loops, but the length of those loops varies from 2 to 6 configurations, and some of the siblings feature two period-2 absorbing loops.

For AD patients with genotype in this category, we may expect their disease dynamics to be relapsing remitting in nature, and extremely difficult to predict due to the variation between sibling networks produced by differences in the latent variables (P, B, A and D).

The triple-switch model, being a hybrid dynamical system, features two types of variable: discrete-valued variables which we represent as the switch configurations, and continuous-valued variables whose dynamics are determined by the ODE. The discrete-valued variables are immediately apparent on AD skin, since they represent inflammation, reddening and flakiness of the skin. Meanwhile the continuous-valued variables represent the hidden information such as the amount of pathogen on the skin surface, the integrity of the barrier, quantity of AMP on the surface and concentration of antigen-presenting dendritic cells in the draining lymph-nodes. Since it is these continuous-valued latent variables that determine which of the siblings in a family of transition trees will best represent a particular patient's disease, it is likely difficult to predict a patient's future disease dynamics.

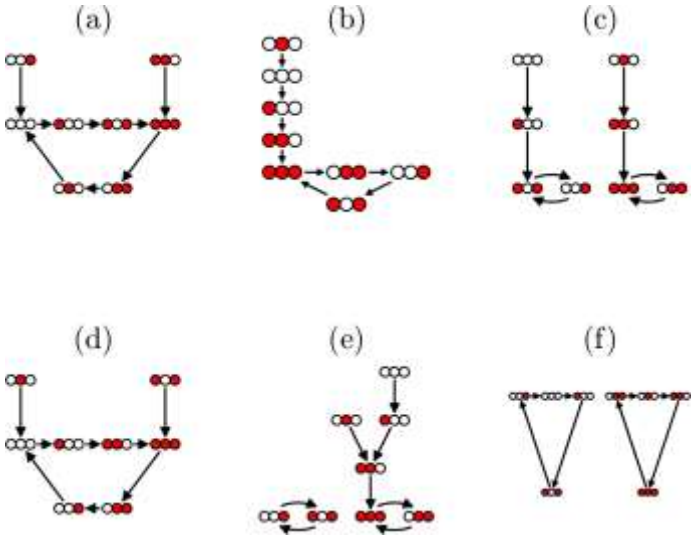


Figure 7.14 Family of transition networks with absorbing loops only.

For the 11 networks where both absorbing loops and absorbing states were present simultaneously, sibling networks in a family always have the transient states leading to each absorbing loop or state arranged differently. In addition, sibling networks in the same family don't always possess the absorbing loop – sometimes only the absorbing configuration is present. For example, in Figure 7.15 (a) the transient configuration [0,0,1] leads to an absorbing loop of six configurations while in its sibling network Figure 7.15 (b) that same transient configuration leads to an absorbing configuration of [1,1,0]. The only difference between these sibling networks is the initial conditions of the latent state variables (P, B, A and D). This shows us how the initial conditions of these latent variables can play a role in the long-term dynamics of chronic AD, changing whether the long-term prognosis is that of a relapsing-remitting disease course, or that of symptoms which do not change over time.

A patient with a genotype producing this category of transition network may both experience relapsing-remitting symptoms and also periods of time where their symptoms settle down to become more settled as random fluctuations in environmentally derived pathogens or skin damage events

cause the patient to move between the basins of attraction corresponding to each sibling network in the family. This can explain why chronic AD dynamics are so hard to predict.

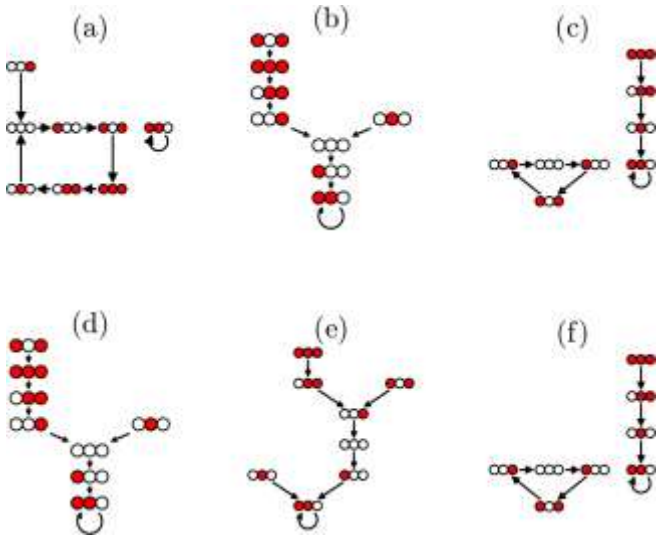


Figure 7.15 Family of transition networks where some siblings feature absorbing loops alongside absorbing configurations.

The proportion of parameter space producing networks with 2 absorbing loops is roughly as much as produced 3 absorbing loops (Figure 7.16). This suggests these two phenotypes may appear with roughly equal frequency.

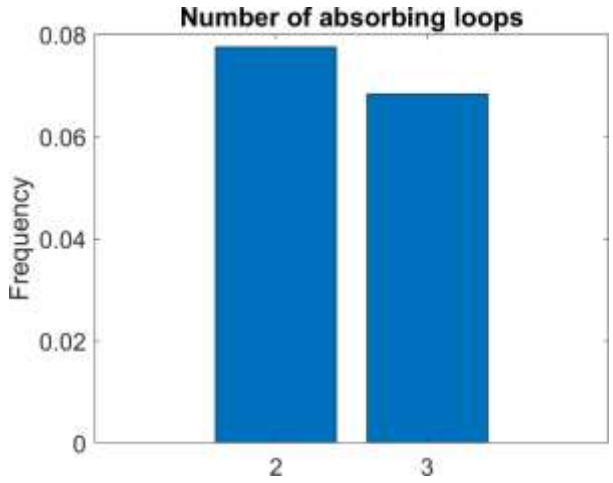


Figure 7.16 Proportion of parameter space giving rise to multiple absorbing loops.

The absorbing loops featured have various lengths. The length of the longest absorbing loop in each network is distributed as in Figure 7.17. Most of the networks have their longest loop of length 2, but around half of all networks with loops have maximum loop lengths of 4 or 6. These longer loops aren't rarities.

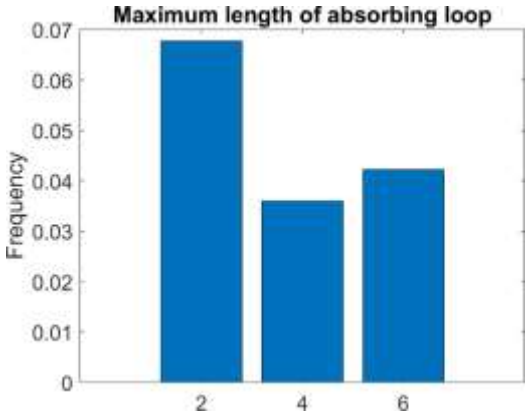


Figure 7.17 Proportion of parameter space giving rise to each maximum loop length in networks.

7.3.4. Networks with transient loops

Of the 178 networks we found, 23 of these featured at least one absorbing loop. Every single one of these was produced by the "p2" parameterisation since they require both the RK-switch and G-switch to have the behaviour **Os** to be created. Figure 7.18 contains three siblings featuring transient loops (a), (c) and (f), while the other three siblings have no transient loop (b), (d) and (e).

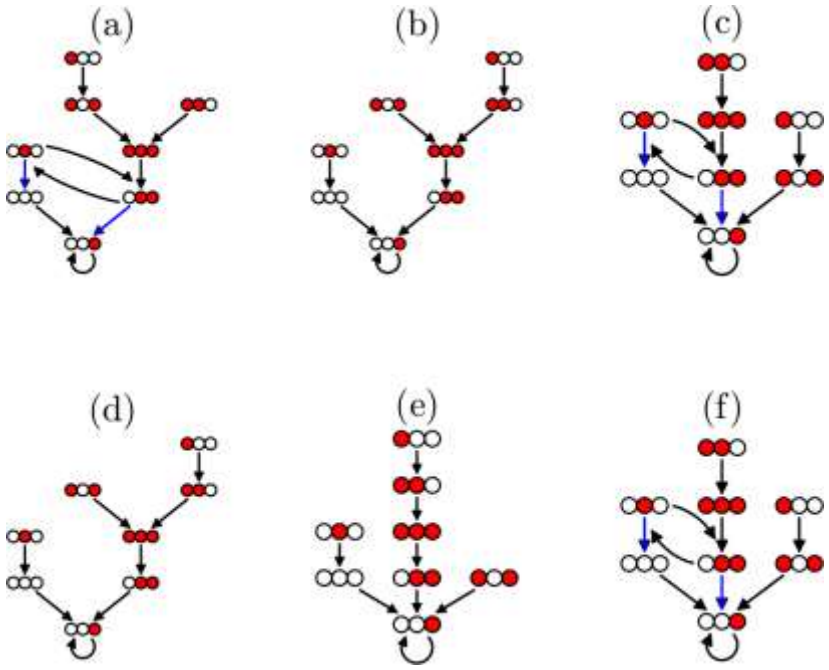


Figure 7.18 A family of transition networks, some of which feature transient loops.

The presence of these transient loops in our predicted networks suggests that chronic AD patients may undergo periods of relapsing-remitting disease course, followed by periods of relative stability as they enter an absorbing configuration. In cases where the genotype of the patient generates transient loops, these periods of relapsing remitting disease course can occur in the absence of any external perturbations by treatment or pathogen challenge.

Networks with this transient loop property occurred for only 2.5% of the parameter space.

7.4. Clinical interpretation of our networks

So far, we've considered the dynamics of chronic AD as represented by our transition networks in the absence of any environmental factors. We propose modelling these environmental factors using the

forcing functions $P_{\mathrm{env}}(t)$, $\delta_{\mathrm{Bext}}(t)$ and $D_{\mathrm{ext}}(t)$ (corresponding to challenge by pathogens, barrier damage such as scratching and presentation of superantigens to the dendritic cells in the site of the skin lesion).

In addition to the stressors originating from the environment, we can represent common treatments for AD by similar forcing function terms, representing application of topical probiotics with the potential to displace pathogenic species and emollient therapy which can improve the structural integrity and water retention of the skin barrier. Similarly, application of dupilumab (a relatively new biologic therapy for chronic AD) can be represented by moving the threshold values for the T-switch as discussed in Appendix P.

The arrows representing switch state transitions show us how the disease will evolve in the future if no environmental factors or treatments are applied. To view the future dynamics, we simply follow the arrows down the tree from the branches to the root, rather like a squirrel descending a tree from the canopy. Under normal conditions, this squirrel doesn't jump, but rather scampers along the branches, following the transition arrows downwards towards the root, as illustrated by the cartoon in Figure 7.19.

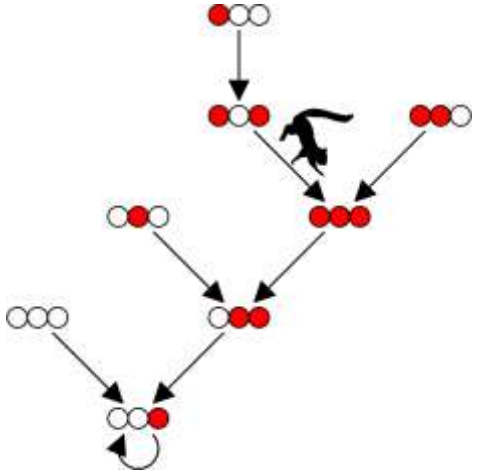


Figure 7.19 Tree diagram with a squirrel scampering down the branches to the roots.

In this chapter, we show that the impact of treatments or environmental stressors can be thought of as causing the squirrel to "jump" between branches of the same tree, or even between trees in the same family. Finally, we summarise the lessons we can glean about controlling chronic AD through application of treatments.

When a therapy is applied to the triple-switch model, we can consider its effect on the latent variables such as pathogen load. If a therapy manages to reduce the pathogen load sufficiently to cross one the the switches' threshold values, it can bring about a transition of switch configuration which would not otherwise occur. This would be seen on our transition diagrams as a "jump" by the squirrel from one branch to another.

In this thesis, we have considered the effect of dupilumab therapy on the T-switch (see Appendix P for a discussion of the implications of dupilumab therapy), but we do not consider the effects of other therapies with explicit terms in our ODEs. In this section, we discuss the potential implications such an extension to the model might have, and how it could be represented on our transition diagrams.

7.4.1. Pathogen clearing therapies.

We can imagine the effects of applying an external perturbation to the pathogen P state variable, for example by washing the skin surface of pathogens using a water or bleach bath [108], applying

antibiotic therapies to the skin [109] or applying topical probiotics [110] to displace or attack the pathogenic species by producing bacteriocins [111] or otherwise interfering with their virulence [112]. These types of treatment could have the effect of decreasing the rate of environmental loading of pathogens $P_{\mathrm{env}}(t)$ described in the triple-switch model ODEs (see section 5.1), causing the dynamics of $P(t)$ to dip downward for a time.

Clearly, such a perturbation would have ripple-effects on the other state variables of the ODE, since they are all coupled. However, for simplicity let us imagine that the perturbation occurs on a much shorter timescale than the ODE dynamics, so we may safely assume that in the short term, only the $P(t)$ dynamic is affected. Sufficient application of such treatments could have the effect of causing the value of P to eventually cross the switching off threshold for the RK-switch (P^-), leading to transition from one switch state configuration to another.

This transition would be "forced" in the sense that it would not occur in the absence of treatment (i.e., would not appear as a black arrow in our trees). We represent a forced transition using a red arrow added to our tree diagram. We can imagine our notional squirrel "jumping" along this red arrow, rather than just scampering along the branches as we previously illustrated in Figure 7.19. An example of this forced transition is illustrated by the red arrow and jumping squirrel in Figure 7.20 (a) after which the disease dynamics play out as the squirrel scampers back down the branch Figure 7.20 (b).

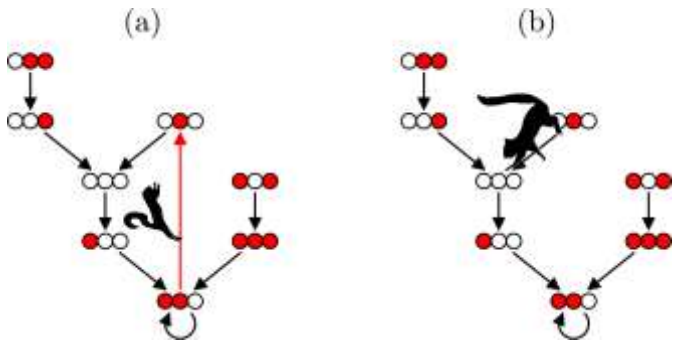


Figure 7.20 Tree diagram with squirrel jumping around to demonstrate chronic AD dynamics. Forcing the RK-switch to turn off, (a) external stimulation causes a jump transition from the root to the branch and then (b) the disease dynamics play out as the squirrel scampering back down that branch.

In this instance, the treatment applied has only a short-term effect on the disease. We assume that the patient spends most of their time at the root state [1,1,0], which corresponds to inflamed, but otherwise well differentiated skin barrier. Upon application of the treatment to decrease their pathogen load, the notional squirrel jumps up to state [0,1,0], which corresponds to a less inflamed state but with a persistent Th2 response. After this, the squirrel scampers back down the branch to the root, passing briefly through the "full health" state where the epidermis is both competent and not inflamed. However, as time passes, the pathogen simply returns, causing first the innate immune responses (RK-switch) to reactivate, and then the Th2 response to return (T-switch).

7.4.2. Barrier reinforcing therapies.

The effect of the most commonly used form of treatment for AD, emollients, could be added to the ODEs of the mode in the form of a forcing term to the ODE for the state variable B. Applying emollient therapy would have the effect of temporarily increasing the barrier integrity (though not the level of differentiation of the keratinocytes therein[7]). Such treatment could thereby cause the G-switch to

[7] The subtle difference between barrier integrity (as represented by a value in the range [0,1] by the double-switch model), and level of differentiation of the keratinocytes (which is what B means in the triple-switch model)

transition from on to off. Once again, we can refer to the tree diagrams to work out the potential implications of this form of treatment. Figure 7.21 shows one such scenario for a tree chosen for illustrative purposes. In Figure 7.21 (a), the treatment is applied, causing the G-switch to turn off and the squirrel to jump up to the branch of the tree via the red "forcing" arrow. Figure 7.21 (b) shows the squirrel subsequently beginning to descend from that branch towards the root once the treatment is removed.

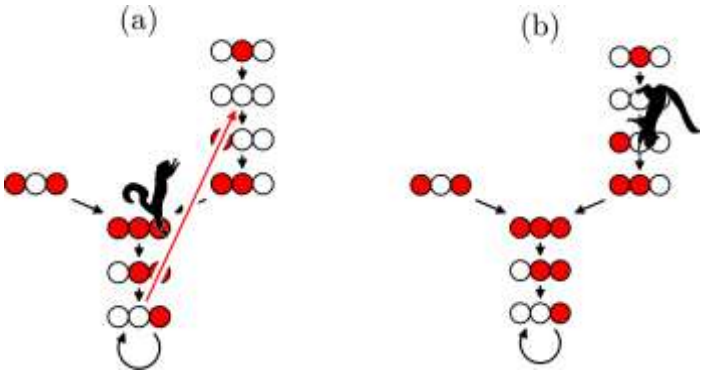


Figure 7.21 Tree diagram with squirrel jumping around to demonstrate chronic AD dynamics. Forcing the G-switch to turn off, (a) external stimulation causes a jump transition from the root to the branch and then (b) the disease dynamics play out as the squirrel scampering back down that branch.

7.4.3. Dupilumab therapy

Finally, application of some of the newest and most sophisticated therapies for AD, the biologic therapy dupilumab, may be considered. We discuss this idea in greater detail in Appendix P.

7.5. Therapies to bring about permanent recovery of symptoms

When we apply a therapy in chronic AD, we make the squirrel "jump" from one switch configuration to another. When we consider the transition networks featuring multiple absorbing configurations or loops, this jump could push an AD patient's dynamics from one tree to another, causing them to converge towards a different absorbing configuration or loop, shown in Figure 7.22

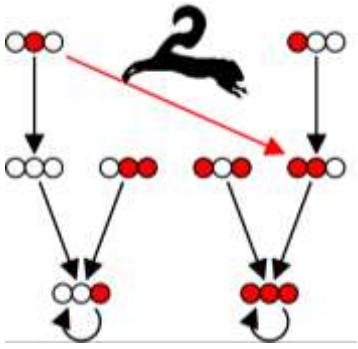


Figure 7.22 Jumping from one tree to another for a genotype where multiple absorbing configurations exist.

In this scenario, it would be very important to be able to work out what the transition networks look like for a particular patient's genotype. If would be particularly important to first determine that the network corresponding to a particular patient actually has more than one absorbing configuration. If the person has only one absorbing configuration, then no amount of forced switching transition is going to cause a sustained change to the symptoms of disease because there would be no other trees

means adding emollient therapy to the triple-switch model is a little more complicated than implied here. In this chapter, we assume a very simple perturbation to the B state variable to aid the discussion, but more careful consideration of this problem should be made in future work.

in the network to transition into. Some caution should be taken to ensure that we don't misidentify the patient's network, inadvertently making the symptoms get worse as a result of the treatment.

7.6. Conclusion

In this chapter, we categorise the transition networks produced by our triple-switch model according to some common features. We split our networks into four categories based on their absorbing configurations or loops, and discussed how frequently each type of feature is found in the networks.

A key finding in this chapter is that our model provides a compelling explanation for why it is so difficult to predict the dynamics of chronic AD. Our model suggests that the difficulty arises due to the strongly timescale-dependent dynamics of the disease. Slight differences in the relative speed of the slow-timescale variables can lead to different switch transitions taking place, and this in turn leads to dramatically different outcomes such as relapsing-remitting disease course or persistent and apparently stable symptoms. This is seen in our results as siblings in the same family of transition trees having distinct features, such as absorbing loops in some instances, and absorbing states in another.

Another key finding, discussed in Appendix P, is that dupilumab treatment is able to change the structure of the transition trees produced by our model, sometimes facilitating an improvement in the disease symptoms. The dupilumab treatment causes the threshold values of the T-switch to increase, which in turn limits the ability of the RK-switch (innate immune response) to activate the T-switch (type-2 adaptive immune response). We discuss in Appendix P how the effect of dupilumab treatment varies depending on the specific topology of the transition trees, but we do not provide a detailed analysis of how the effect of dupilumab therapy varies depending on model parameters.

Finally, we discussed how to interpret the networks in the context of external factors such as pathogen loading, or various therapies that can be applied to chronic AD.

Chapter 8. Summary and Conclusions

In this thesis, we developed novel analytical methods for the analysis of hybrid dynamical systems. Hybrid dynamical systems have a relatively underdeveloped suite of analytical tools, and most that do exist are aimed at simulation and optimisation of these systems. Our methods have taken a more analytical path, focussing on identifying the full extent of dynamics the model can produce. We applied these novel methods to our mechanistic model of chronic AD, using them to infer the range of dynamics predicted by the model. In this chapter, we summarise the novelties in this workflow, and propose generalisations and further applications for these methods in hybrid dynamical systems.

A straightforward method to calculate the model dynamics would be to use numerical integration techniques to solve the ODE and detect when the switches transition from one state to the next by examining that ODE solution. The method of simulation, along with our example MATLAB code are discussed in Appendix M where it is used to double-check our analysis. While this method may be used to compute solutions, it does not allow us to gain insight into qualitative properties of those solutions as we vary the parameter values and initial conditions.

Bifurcation analysis techniques for ODE models to obtain this kind of qualitative description of dynamics are well-established with a number of powerful numerical methods and software packages available. For example, in Appendix A and Chapter 3, we make use of numerical continuation methods for bifurcation analysis using the well-established MATLAB package MatCont.

However, the study of hybrid dynamical systems has not yet advanced to the same degree, with no equivalently mature methodologies and software packages available for our hybrid dynamical system model. The only software package we found for hybrid dynamical systems in MATLAB, HyEQ_Toolbox [113], features only the ability to simulate these systems, but no tools for bifurcation analysis. Where necessary we developed novel methods and wrote our own MATLAB code specific to our application.

The methods we developed in this thesis were designed with our particular model in mind but can be applied to hybrid dynamical systems more generally. Some of these methods are subject to a few caveats, requiring the hybrid dynamical system to have the particular properties which we leverage in our application. In each of the sections below, we summarise the method, discuss its requirements and limitations, and then consider how it can be applied to hybrid dynamical systems more broadly.

8.1. Finding all possible connectivity of the regions in bifurcation diagrams

This method is developed in section 5.2.3 where it is used to work out all the variations of layout for the bifurcation diagrams produced by the triple-switch model. For example, Figure 8.1 (a) and (b) demonstrate two hypothetical examples of bifurcation diagrams with different parameter values. While the actual placement of the BE bifurcation lines differ between the two diagrams, the "connectivity of the regions" are topologically identical, as shown by the identical summarising networks shown at the bottom of (a) and (b). In this network, each node represents a region present on the bifurcation diagram, and each edge represents a BE bifurcation line which separates those regions. A different change of parameters shown in Figure 8.1 (c) and (d) has moved the BE bifurcation lines such that it brings about a change to the connectivity network.7

The method we developed uses some reasoning about the geometry of the BE bifurcation lines, translating this reasoning into a set of inequality conditions on the parameters which can tell us the connectivity of bifurcation diagram. For a more detailed explanation of the process (see section 5.2.3).

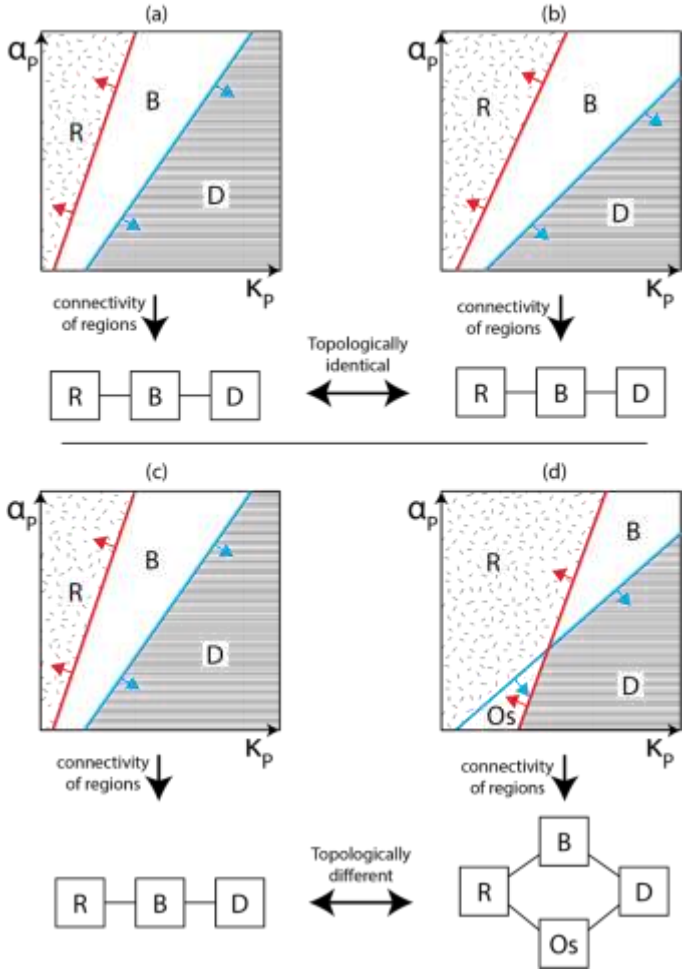


Figure 8.1 Connectivity networks of bifurcation diagrams.
(a) and (b) represent two bifurcation diagrams from different parameterisations, but which have topologically identical connectivity networks. (c) and (d) represent two bifurcation diagrams with topologically different connectivity networks.

8.1.1. Requirements and limitations

The following points are required to get this method to work.

1. BE bifurcation lines must be defined by algebraic expressions.

   One can imagine a situation where a hybrid dynamical model yields BE bifurcation lines which can only be approximated through numerical methods. This would arise if the ODE of the system had no closed-form solution for its equilibrium points. In such a scenario, it would be impossible to write down conditions on the parameters and therefore this method would not work.

   We were fortunate in our model that the ODEs are well behaved, giving us just one equilibrium value which we can write down as an algebraic expression as found in section 0. Should a hybrid dynamical system contain an ODE with more complicated features, such as multiple equilibrium values with no algebraic expression, limit cycles or more even chaotic attractors, the method we propose would quickly break down due to the difficulty in defining the BE bifurcation lines.

2. BE bifurcation lines should be sufficiently simple that it is possible to obtain a clear picture of the connectivity of the regions.

   This point is rather harder to define than the first. Connectivity here means the way the behaviour regions are arranged on the bifurcation diagram. For example, in Figure 8.1 (a) where the BE

bifurcation lines do not cross, and it is easy to see the connectivity network. Where the lines *do* cross in Figure 8.1 (d), it is again easy to see that there are four connections produced between the behaviour regions.

We were fortunate in our model that the BE bifurcations always produced straight lines through the plane of bifurcation parameters. This meant the lines could only intersect in the diagram a maximum of once, which limited the variety of ways these lines could be arranged relative to each other to separate the plane into behaviour regions. We were able to define just four ways in which the two BE bifurcation lines could be arranged on the bifurcation diagram, illustrated in Figure 8.2, and then rule out two of these possibilities due to the constraint that all parameter values must be positive (Table 8.1).

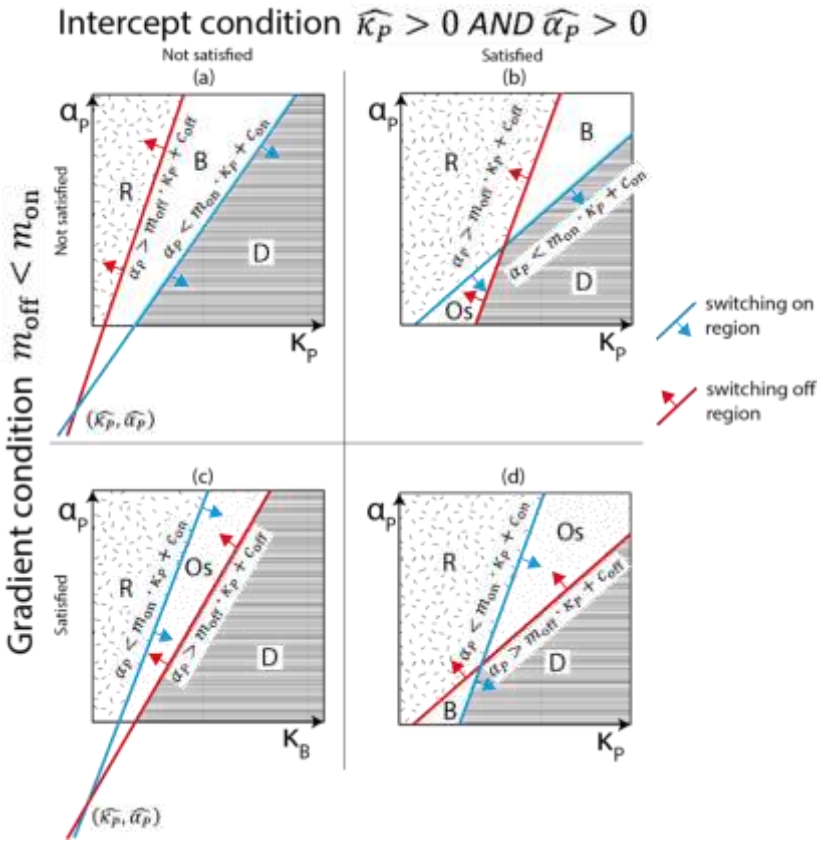


Figure 8.2 Four ways to arrange the BE bifurcation lines in the plane (reproduced from Figure 5.13).

Table 8.1 Two possible ways in which the BE bifurcation lines can arrange in the plane given all parameter values must be positive.

	Intercept condition $\widehat{\kappa_P} > 0$ and $\widehat{\alpha_P} > 0$		
Gradient condition $m_{\text{off}} < m_{\text{on}}$		**Not satisfied**	**Satisfied**
	Not satisfied	R, B, D	R, B, Os, D
	Satisfied	R, Os, D	R, B, Os, D

If the BE bifurcation lines are curves with rather complicated expressions, it may become restrictively difficult to work out how the geometry of the bifurcation diagrams would be affected by changing parameter values.

8.1.2. General applicability to hybrid dynamical systems

This method relies heavily on the ability to define the BE bifurcation lines by relatively simple algebraic expressions. I would argue that this requirement isn't as restrictive as it first appears. The Occam's razor principle that recommends understanding phenomena using the least complex model possible should generally act to disincentivise the building of models that fail to meet this requirement. As such, I would expect this method to be practical in most hybrid dynamical systems applications.

In our application of this method, we always considered bifurcation diagrams with two parameters (κ_P, α_P for the RK-switch behaviours, and κ_B, δ_B for the G-switch behaviours). We could consider bifurcation diagrams with three or more parameters, which would lead to more complicated geometrical relationships between the BE bifurcation curves. In a three-dimensional space, these BE bifurcation curves would form planes, and their intersections would be lines. Under these conditions, it should still be possible to develop some idea of the geometrical relationship between the planes and write down these as conditions on the parameter, but it would be more difficult to visualise and lead to more potential cases than we identified in 2D.

We advise first attempts at this method be in as few dimensions as possible. Once some understanding of the system has been developed, more dimensions can be considered if necessary.

8.2. Sensitivity analysis for BE bifurcation diagrams

This method was developed in section 5.2.6 to calculate how changes to the parameter values of the model affect the area taken up on the bifurcation diagrams by each switch behaviour. We analytically calculated the derivative of each BE bifurcation line with respect to each parameter and used these to work out the rate of motion of each vertex with respect to each parameter as represented by the cartoon Figure 8.3. We then applied these results to a general formula for the area of a polygon, tracing how the motion of each vertex causes the polygon area to change. By mixing some simple calculus with computational geometry, we were able to calculate the rate of change of each behaviour area without too much difficulty. The whole calculation can be neatly carried out by arranging the relevant derivative terms into matrices and carrying out a couple of matrix multiplications. This makes it fairly easy to implement in software and easily scalable.

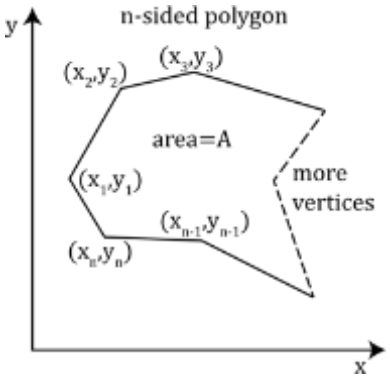


Figure 8.3 Cartoon representation of the sensitivity analysis method

8.2.1. Requirements and limitations

Our sensitivity analysis method requires the ability to evaluate derivatives of the BE bifurcation curves with respect to each parameter at a finite number of points in the bifurcation plane. In this thesis, we were able to analytically calculate these derivates since we have algebraic expressions for the curves. For example, for the RK-switch bifurcation diagram, we calculated the vertex defined by the intersection of the two BE bifurcation lines as in equations 5.14 and 5.15.

$$\widehat{\kappa_P} = \frac{c_{\text{off}} - c_{\text{on}}}{m_{\text{on}} - m_{\text{off}}} = \frac{\delta_P}{(P_{env} + \overline{P_{ext}})}\left(\frac{P^- P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\left(\bar{A}(1,*,*) - \bar{A}(0,*,*)\right)}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)}\right) \quad (8.1)$$

$$\widehat{\alpha_P} = \frac{c_{\text{off}} m_{\text{on}} - c_{\text{on}} m_{\text{off}}}{m_{\text{on}} - m_{\text{off}}} = \delta_P\left(\frac{P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right) - P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)}{P^-\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*) - P^+\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)}\right) \quad (8.2)$$

We calculate how rapidly this vertex moves as we vary each model parameter by simply differentiating these expressions, offloading the tedious differentiation task to MATLAB's symbolic algebra toolbox. We end up with a vector of symbolic expressions for the derivatives for the $\widehat{\kappa_P}$ component, and another for the $\widehat{\alpha_P}$ component, which we concatenate to form a matrix summarising the motion of the vertices v with respect to the parameter values θ.

$$\frac{dv}{d\theta} = \begin{pmatrix} \frac{\partial \widehat{\kappa_P}}{\partial P_{env}} & \frac{\partial \widehat{\alpha_P}}{\partial P_{env}} \\ \frac{\partial \widehat{\kappa_P}}{\partial P_{env}} & \frac{\partial \widehat{\alpha_P}}{\partial P_{env}} \\ \vdots & \vdots \end{pmatrix}$$

This matrix of derivatives gives us the horizontal and vertical components of the motion of the BE line for every parameter in our model. Once this matrix of symbolic expressions was calculated, we simply evaluated it by substituting values for the parameters.

We can imagine a situation where we do not have algebraic expressions for the BE bifurcation curves, and so this matrix of symbolic expressions cannot be obtained. We may instead have a numerically calculated series of points along the BE line. Under these conditions, it is still possible to apply numerical continuation methods (using a software package such as auto or MatCont) to obtain these values. Once the rates of change are obtained and arranged into a matrix $\frac{dv}{d\theta}$ as specified in equation 5.53, the rest of this sensitivity analysis method proceeds as normal.

8.2.2. Possible extensions of this method to more complicated hybrid dynamical systems

1. When behaviour regions are not polygons

To obtain an exact result, we relied on the behaviour areas all being 2D polygons. For n-sided polygons, there is a relatively simple general expression for the area, reproduced here as equation 8.3.

$$A = \left|\frac{(x_1 y_2 - x_2 y_1) + (x_2 y_3 - x_3 y_2) + \cdots (x_{n-1} y_n - x_n y_{n-1}) + (x_n y_1 - x_1 y_n)}{2}\right| \quad (8.3)$$

We could imagine a situation where a hybrid dynamical system gives us BE bifurcation *curves* (i.e. not straight lines). Under such conditions, clearly this n-sided polygon idea does not represent the regions we'd need to consider. However, we could simply place many vertices along the curved edges of each region to approximate area by a polygon with a large number of sides. It would be advisable in these cases to try a few different levels of resolution and work out when the calculated area nears convergence to the true area.

The use of large numbers of polygon vertices would increase the length of the vector specifying their locations $v = [x_1 \quad y_1 \quad x_2 \quad y_2 \quad \cdots \quad x_n \quad y_n]$, and thus also increase the size of the matrix of

derivatives of each area with respect to these vertices. Care should be taken to ensure the matrices remain of sensible size when attempting this calculation. It should be possible to write some robust and fast code to carry out these calculations automatically for any given bifurcation diagram with curved BE bifurcation curves.

2. When behaviour regions are considered in higher dimensions

In our application of this method in this thesis, we only considered 2D regions for each behaviour. However, it should be possible to generalise this method to multiple dimensions by using a higher-dimensional version of equation 8.3. The expression for the n-volume of an n-dimensional simplex[8] is in equation 8.4.

$$A = \frac{1}{n!}|v_1 - v_0 \quad v_2 - v_0 \quad \dots \quad v_n - v_0|, \tag{8.4}$$

where $v_0 \dots v_n$ are column vectors specifying the location of each vertex of the simplex. If we place this n-dimensional version of the equation into the workflow we describe in section 0, the calculus remains pretty simple, and the organisation of the calculation into matrices remains the same.

In general, behaviour regions are not all simplexes. However, more complex polyhedrons could be broken down into a collection of simplexes for which equation (8.4) may be applied to work out their n-volume and rate of change of n-volume. Once this is achieved, we need to add up all the areas and rates of change to evaluate the sensitivity of the whole region.

I have not tried to apply this method in three-dimensions or greater in this thesis, as the complexity appears high. This technique is not dissimilar to those applied in finite-element analysis, so there are potentially some transferrable techniques there.

8.3. "Racing switches" method to infer switch transitions from behaviours

When we start trying to work out how the switches of the hybrid dynamical system transition state over time as a result of the ODE dynamics, we come across a dilemma. To properly solve the problem, we would need to keep track of the dynamics of the ODE, keeping track of how the initial conditions of the continuous-valued state variables affect the way in which the switches transition between states as time passes. In other words, we need to calculate the separatrices which split the state space up into various basins of attraction. A way to numerically calculate the extent of each basin of attraction is detailed in Appendix O. However, this separatrix problem, and the solution we propose is so complex that it was sensible to find a way to simplify the problem to obtain an approximation of the transition networks.

We developed the "racing switches" method to allow us to gain some insight into the dynamics of the system without having to work out what the basins of attraction look like for each focal point we identified. This method is able to identify all the possibilities for switch transitions for a given set of parameters and tells us where to look for basins of attraction to fully characterise the model's dynamics. In that sense, this method can be thought of as an essential first step before any attempts at basin of attraction calculations are made.

The method involves us comparing the behaviours of all three switches simultaneously, and inferring from this collected information which of the switches will undergo a transition from one state to the next. We described a way to carry out this process on the triple-switch model by constructing a

[8] Note, a simplex can be thought of as the simplest shape possible in any n-dimensional space. It is the shape that spans each available dimension with the minimum number of vertices. For example, in three dimensions the simplex shape is a tetrahedron.

"behaviour-state" table summarising the mapping between state C and behaviour b as shown in Figure 8.4. The full process is described in section 6.5.2.

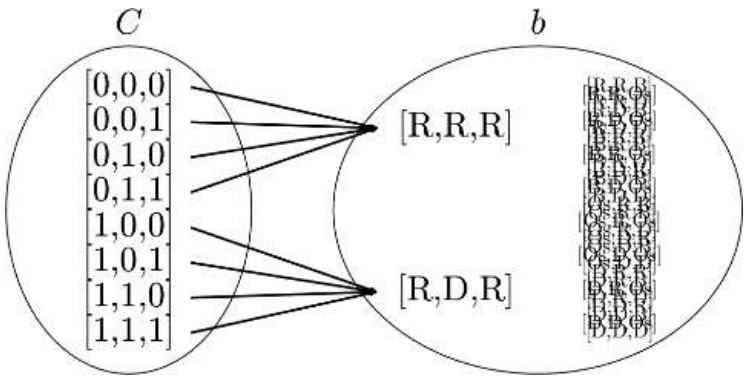


Figure 8.4 Example configuration-behaviour mapping

Once this "behaviour-state" table is produced, we can infer which switches must undergo transitions by considering them racing each other to their respective threshold values. Each possible result from our switch "races" is considered and presented as a "sibling" in the family of transition networks. From here, we can visualise these inferred transition networks in the form of the tree diagrams we discuss in section 6.6.

8.3.1. Requirements and limitations

This method requires that we are able to define the behaviour b (recovery, bistable, oscillation or damage) of each switch for every possible configuration C of the switches (i.e. the combination of off and on states of the switches). That's all the information required.

The limitations of this method are twofold:

1. It cannot identify the exact times at which each transition will take place, so the resulting networks of switch transitions cannot be associated with any exact timescale. We can only specify an approximate timescale based on the observed rate of convergence of ODE solutions to their focal points.
2. In instances where the relative timescales of two or more switches give rise of multiple possibilities for which switch undergoes a transition (we call these "multiple transition states" in section 6.5.2), this method simply shows us every possibility, and cannot give us any information as to which one should occur.
   This is in one sense a limitation, but also the strength of this technique. It is able to show us what we still don't know about the system we're studying, so that we can focus the more sophisticated methods in Appendix O on these areas if required.

8.3.2. Extension of this method to other hybrid dynamical systems

We could consider how this method might be applied to hybrid dynamical systems with more than three switches. The process would be exactly the same, except that the behaviour-state table and resulting transition table would have a few more rows and columns.

In instances where it is unclear which of a selection of the switches undergoes a transition, the method presents us with every possible outcome. As the number of switches in a hybrid dynamical system increases, the number of such possibilities also rapidly increases. In this thesis, we considered a system of three switches, which led to $3! = 6$ possibilities in such cases. This result is general, summarised as 8.5 below.

$$\text{Model with } n \text{ switches} \implies n! \text{ possibe switch transitions} \tag{8.5}$$

This in itself is an important result. It tells us that the more complicated a hybrid dynamical system becomes, the more possible ways there are that the system can behave due to minor differences in the rate at which the various switches change. As the number of switches increases, the more important it becomes to be able to define the relative rates at which the ODE dynamics bring about switch transitions. In biological systems, such questions of relative timescale are often hard to answer precisely or generally, as they may be subject to lots of unknown factors. This is a potential explanation for the unpredictability of chronic AD symptoms.

8.4. Calculating the basins of attraction for each transition networks

We considered how we might go about calculating the location of separatrices in the continuous-valued state space, as detailed in Appendix O. Before discussing this further, let us first note that the method is simply a proposal, and we have not realised and tested it in this thesis due to its evident complexity.

8.4.1. Requirements and limitations

Several problems emerge when one tries to actually carry out the method described.

1. Points sampled from the separatrix need to be organised into a coherent curve in the state space. This requires some computational geometry which may need more thought than we've given here.
2. Locations in the state space where the divergence of the ODE's vector field is high will be sampled less frequently than regions with lower divergence. This may lead to artifacts in the constructed separatrix curve.

8.4.2. Extension to other hybrid dynamical systems

The method we propose is applicable to any hybrid dynamical system, but the limitation listed above may become more acute if the ODE of the system is in any way complicated (containing cyclical dynamics for example).

We can imagine attempting this method on a hybrid dynamical system with more state variables and more switches. This would lead to us having to sample more points in the state space, and so the method would suffer a penalty in terms of computational time required. A model with N continuous-valued state variables would in general require $O(\ (N-1)^2\ )$ numerical integrations of the ODE to carry out this method.

8.5. Future work

There are a few avenues of investigation that were not adequately explored here, or that have become visible as this work has progressed, but which remain unexplored as our attentions were placed elsewhere. Broadly, these fall into two categories: firstly, those of the analytical form where we focus on fully elucidating how the mathematical model works and how we can derive information about its functions. Secondly, investigation of the biological/clinical implications of the model, including the effects of treatment and potential clinical applications of what we've learned here.

This section details each of them, from the easiest and best-defined tasks to the harder and more open-ended research questions. In each case, the model built up in this thesis provides a basis for study, and the techniques I have developed should prove helpful.

8.5.1. Specifying parameter ranges

In section 5.2.3 we found that there are four parameterisations of the triple-switch model, each of which induces a different combination of behaviours from the RK- and G-switch. We named these p1, p2, p3 and p4, and defined their values in table J.1. In this thesis, we used single set of values to represent each of these combinations of behaviours, rather than a range over which each behaviour combination is found.

We can work out which family a parameter set belongs to, by simply checking if conditions (8.6) and (8.7) are met (reproduced from section 5.2.3).

RK-switch bifurcation diagram condition: Not satisfied → {R, B, D} Satisfied → {R, B, Os, D}	
$P^{-}\left(1+\gamma_{PB}\bar{B}(1,*,*)\right)\bar{A}(1,*,*)-P^{+}\left(1+\gamma_{PB}\bar{B}(0,*,*)\right)\bar{A}(0,*,*)>0$	(8.6)

G-switch bifurcation diagram condition: Not satisfied → {R, Os, D} Satisfied → {R, B, D}	
$\frac{G_{\mathrm{on}}}{((2\cdot B^{-}\cdot\gamma_{BB}+1)^{2}-1)}-\frac{G_{\mathrm{off}}}{((2\cdot B^{+}\cdot\gamma_{BB}+1)^{2}-1)}>0$	(8.7)

The ease with which we can check these conditions should allow us to define a distribution of parameter values for which each of the four combinations of switch behaviours are produced. We can simply sample the parameter space randomly to record how these conditions were met at each sample.

Why is it better to define distributions rather than a single value?

1. All of our analysis after section 5.2.3 relies on checking just four points in the parameter space, and this limits our results to reflecting those locations. It is possible that had we settled on a slightly different set of parameter values for our experiments, we would be looking at a subtly different set of transition trees, and possibly end up drawing different conclusions as a result. It is important that we integrate out this source of ambiguity in order to improve our confidence in the inferences we draw from this model.
2. Specifying a range over which our parameter values tend to lead towards a particular connectivity of transition tree is the first step in connecting genotype to phenotype. It effectively means we can start to ask the question: given this genotype, what are the most likely mechanisms for chronic AD persistence in this patient? Answering this question could potentially inform treatment decisions, which was one of the aims of this project.

8.5.2. Calculating separatrices to find the extent of each basin of attraction.

Each transition network we draw is produced by a particular parameterisation of the triple-switch model and set of initial conditions for the latent variables P, B, A and D. Through the bifurcation analysis methods, we developed in section 5.2, we are able to define the regions of parameter space leading to each transition network, but we are not able to define the shape or size of the basin of attraction for the state variable P, B, A and D. We outline a method to calculate the shape and size of the basin of attraction in Appendix O. It is likely to be a difficult process, so why should it be pursued?

Obtaining the separatrices would allow us to calculate the volumes of the basins of attraction for each transition networks we propose rather than simply assuming that all the siblings in each family share

the state space out equally and thus have the same likelihood of occurring as a disease phenotype. While the assumptions we made in section 6.5.2 are reasonable, they are nevertheless likely to affect our predictions of the relative frequencies of the model dynamics. It's difficult to estimate how much these assumptions are affecting the results, so it would be nice to resolve the issue by working out what the results look like without having to make this strong assumption.

8.5.3. Checking the effects of common AD therapies

In section 7.4, we discussed the effect of three different types of therapy on our model of chronic AD. Our discussion was really just a thought experiment at this stage, since we do not yet have a rigorous set of inputs to the triple switch model to represent these therapies. A more thorough analysis would consider how emollient, corticosteroid and dupilumab therapies could affect our three switches and the mechanisms incorporated into the ODE and compare the resulting predictions about disease dynamics to clinical data on each of these therapies.

8.5.4. Stratification of AD patients based on their transition tree

We have established in this thesis that our mathematical model of chronic AD provides predictions of disease dynamics in the form of "transition trees". These trees provide a map of how the disease dynamics will evolve in the future based on the current symptoms. It should therefore be possible to use statistical methods to connect longitudinal data of symptoms to the portfolio of transition trees our model predicts. Once we have formed some ideas about the transition tree which best represents an individual's disease, we can start making predictions about how effective dupilumab therapy might be in their case.

8.6. Final remarks

Hybrid dynamical systems such as the triple-switch model potentially have applications beyond modelling AD. This type of model is useful in instances where mechanisms operating at drastically differing timescales interact. This is common when modelling biological systems, particularly when considering phenomena emerging from tissue-level interactions. Cells and tissues by their nature grow and move slowly relative to molecular-level phenomena like signalling, and biological systems usually consist of interactions between these differing time- and length-scale systems.

It seems likely that hybrid dynamical systems should be of utility in related atopic diseases such as asthma, which interact with similar components of the immune system and also operate on an epithelial tissue, should be particularly suited to a hybrid dynamical system.

In this thesis, we have presented a workflow of novel methods applicable to this class of models. We have shown that a range of complicated dynamics emerge, featuring single absorbing configurations where a single set of symptoms always emerges after the dynamics have ceased, periodic loops where symptoms repeat over time, multiutilities where multiple configurations of symptoms could persist for long periods, or even combinations of these. These complicated configurations correspond to the hard-to-predict disease dynamics of chronic AD. The methods presented in this thesis may be generalised to apply to other hybrid dynamical systems which may be developed for related atopic diseases.

In this chapter, we discussed the limitations of the novel methods we developed in this thesis. Specifically, the methods provide a useful approximation of the range of dynamics produced by hybrid dynamical systems, but do not provide us with a precise picture of the basins of attraction corresponding to each dynamic. Our methods may be extended to obtain information about basins of attraction using numerical integration techniques, which we discuss in Appendix O.

We summarise the following key findings, each associated with a testable prediction.

1. AD dynamics are strongly sensitive to the relative timescales of the three switches.

Even in the absence of more detailed information about basins of attraction, our novel methods can demonstrate the importance of the relative speeds of innate immune responses, calcium signalling and Th2 responses in determining the disease dynamics. Our methods show us how differing assumptions about the relative timescales of the three switches can lead to dramatically different "sibling" transition trees, featuring different absorbing states or loops and thus quite different experiences for a person with AD. We propose that this sensitivity to relative speeds of disease mechanisms can explain the observed complexity of AD dynamics. Two people whose disease may be represented by identical model parameters may in fact show notably different disease dynamics if the relative timescales of their three switches differ. One may, for instance, show a disease phenotype characterised by long periods of stable symptoms due to the presence of absorbing states in their transitions tree. Meanwhile others may experience relapsing-remitting symptoms due to the presence of an absorbing loop in their transition tree.

Testable prediction:

- Therapies which act to change the rate at which inflammation, itch or epidermal calcium signalling occurs (corresponding to our three switches) should be able to change the time course of chronic AD symptoms in a noticeable way, though not necessarily improve them. This corresponds to transitioning a patient from one transition tree to a "sibling" transition tree.

2. Sufficient observation of chronic AD time course may facilitate the identification of a patient's "transition tree".

We have established in this thesis that our mathematical model of chronic AD provides predictions of disease dynamics in the form of "transition trees". These trees provide a map of how the disease dynamics will evolve in the future based on the current symptoms. It should therefore be possible to use statistical methods to connect longitudinal data of symptoms to the portfolio of transition trees our model predicts. Once we have formed some ideas about the transition tree which best represents an individual's disease, we can start making predictions about their future disease dynamics, or how effective dupilumab therapy might be in their case.

One problem with this idea is that our transition trees (and indeed Markov chains) neglect to include information about the exact timing of any changes in switch configuration. Effectively, we know how the system will proceed along the branches of the transition trees, but we don't know how rapidly it will proceed through each transition. This presents us with a problem, because it prevents us from aligning our model's dynamics with experimental data. However, this can be overcome by using numerical integration to solve the hybrid dynamical system solutions.

Testable predictions:

- A hidden Markov model should provide a reasonable statistical model for inference of future AD dynamics. The Markov process in such a model would represent the transitions between different switch configurations, while the emission probabilities correspond to the measurement process such as providing a subjective score of a particular symptom. We predict that each patient's disease may have a different Markov process, but that once found, the absorbing configurations in the Markov process should be consistent for any given patient.
- We can recover a more precise view of the timescale of each transition by simulating the ODE itself, rather than relying on the focal point analysis we developed in this thesis. This would allow us to obtain a time course of our predictions for direct comparison with clinical time course data. This could facilitate better optimisation of this model predictions against experimental data.

It is also feasible that the hybrid dynamical system approach we have used to construct our model could in turn be applied to related diseases such as psoriasis, which is a skin complaint arising from a dysregulation of AMPs on the epidermis. This skin disease touches on many of the same physiological systems we have modelled here, and so it may be possible to obtain similar insights via this approach.

3. Dupilumab therapy can “rewire” the transition trees by interfering with the cross-talk between the RK- and T-switch. This can bring about a marked improvement in symptoms, but only in cases where the dynamics are driven by this particular cross-talk.

In Appendix P, we find that dupilumab therapy, uniquely amongst AD therapies due to its inhibitory on the cytokine signalling through IL-4 of Th2 cells, is able to increase the threshold value at which Th2 response is induced. This moving of the threshold values rewires the “transition tree” representing the disease dynamics. For some of the transition trees, this brings about an improvement in disease symptoms. Meanwhile, other transition trees are less responsive to dupilumab because the rewiring doesn’t affect the absorbing state or loop.

Testable prediction:

Application of Dupilumab should cause a change to the disease dynamics (by rewiring the transition networks). If we can identify a patient’s transition tree (possibly by applying the hidden Markov model described above), we predict that patients whose trees feature absorbing states with RK-T crosstalk should be the most responsive to dupilumab therapy.

References

[1] S. Kezic *et al.*, "Skin barrier in atopic dermatitis," *Front. Biosci. - Landmark*, vol. 19, no. 3, pp. 542–556, Jan. 2014.

[2] D. A. Lintzeri, N. Karimian, U. Blume-Peytavi, and J. Kottner, "Epidermal thickness in healthy humans: a systematic review and meta-analysis," *J. Eur. Acad. Dermatology Venereol.*, vol. 36, no. 8, pp. 1191–1200, Aug. 2022.

[3] J. A. Odhiambo *et al.*, "Global variations in prevalence of eczema symptoms in children from ISAAC Phase Three," *J. Allergy Clin. Immunol.*, vol. 124, no. 6, Dec. 2009.

[4] J. I. Silverberg and J. M. Hanifin, "Adult eczema prevalence and associations with asthma and other health and demographic factors: A US population–based study," *J. Allergy Clin. Immunol.*, vol. 132, no. 5, pp. 1132–1138, 2013.

[5] K. Abuabara, A. M. Yu, J.-P. Okhovat, I. E. Allen, and S. M. Langan, "The prevalence of atopic dermatitis beyond childhood: A systematic review and meta-analysis of longitudinal studies," *Allergy*, vol. 73, no. 3, pp. 696–704, Mar. 2018.

[6] I. A. Deckers, S. McLean, S. Linssen, M. Mommers, C. P. van Schayck, and A. Sheikh, "Investigating international time trends in the incidence and prevalence of atopic eczema 1990-2010: a systematic review of epidemiological studies," *PLoS One*, vol. 7, no. 7, p. e39803, 2012.

[7] M. I. Asher *et al.*, "Worldwide time trends in the prevalence of symptoms of asthma, allergic rhinoconjunctivitis, and eczema in childhood: ISAAC Phases One and Three repeat multicountry cross-sectional surveys," *Lancet*, vol. 368, no. 9537, pp. 733–743, 2006.

[8] M. R. Perkin, D. P. Strachan, H. C. Williams, C. T. Kennedy, J. Golding, and A. S. Team, "Natural history of atopic dermatitis and its relationship to serum total immunoglobulin E in a population-based birth cohort study," *Pediatr Allergy Immunol*, vol. 15, no. 3, pp. 221–229, 2004.

[9] D. Garmhausen *et al.*, "Characterization of different courses of atopic dermatitis in adolescent and adult patients," *Allergy*, vol. 68, no. 4, pp. 498–506, Apr. 2013.

[10] A. Lloyd-Lavery, L. Solman, D. J. C. Grindlay, N. K. Rogers, K. S. Thomas, and K. E. Harman, "What's new in atopic eczema? An analysis of systematic reviews published in 2016. Part 2: Epidemiology, aetiology and risk factors.," *Clin. Exp. Dermatol.*, vol. 44, no. 4, pp. 370–375, Jun. 2019.

[11] C. Flohr, S. G. O. Johansson, C. F. Wahlgren, and H. Williams, "How atopic is atopic dermatitis?," *J. Allergy Clin. Immunol.*, vol. 114, no. 1, pp. 150–158, Jul. 2004.

[12] P. M. Elias and M. Steinhoff, "'Outside-to-inside' (and now back to 'outside') pathogenic mechanisms in atopic dermatitis," *J Invest Dermatol*, vol. 128, no. 5, pp. 1067–1070, 2008.

[13] C. Flohr and J. Mann, "New insights into the epidemiology of childhood atopic dermatitis," *Allergy*, vol. 69, no. 1, pp. 3–16, 2014.

[14] A. D. Irvine, W. H. I. McLean, and D. Y. M. Leung, "Filaggrin Mutations Associated with Skin and Allergic Diseases," *N. Engl. J. Med.*, vol. 365, no. 14, pp. 1315–1327, Oct. 2011.

[15] L. Paternoster *et al.*, "Multi-ancestry genome-wide association study of 21,000 cases and 95,000 controls identifies new risk loci for atopic dermatitis," *Nat. Genet. 2015 4712*, vol. 47, no. 12, pp. 1449–1456, Oct. 2015.

[16] S. Weidinger *et al.*, “A genome-wide association study of atopic dermatitis identifies loci with overlapping effects on asthma and psoriasis,” *Hum. Mol. Genet.*, vol. 22, no. 23, pp. 4841–4856, Dec. 2013.

[17] S. Weidinger, L. A. Beck, T. Bieber, K. Kabashima, and A. D. Irvine, “Atopic dermatitis,” *Nat Rev Dis Prim.*, vol. 4, no. 1, p. 1, 2018.

[18] E. Domínguez-Hüttinger *et al.*, “Mathematical modeling of atopic dermatitis reveals ‘double-switch’ mechanisms underlying 4 common disease phenotypes,” *J. Allergy Clin. Immunol.*, vol. 139, no. 6, pp. 1861-1872.e7, Jun. 2017.

[19] E. Domínguez-Hüttinger, M. Ono, M. Barahona, and R. J. Tanaka, “Risk factor-dependent dynamics of atopic dermatitis: modelling multi-scale regulation of epithelium homeostasis,” *Interface Focus*, vol. 3, no. 2, p. 20120090, Apr. 2013.

[20] D. R. Hoffman, L. M. Kroll, A. Basehoar, B. Reece, C. T. Cunningham, and D. W. Koenig, “Immediate and extended effects of sodium lauryl sulphate exposure on stratum corneum natural moisturizing factor,” *Int. J. Cosmet. Sci.*, vol. 36, no. 1, pp. 93–101, Feb. 2014.

[21] M. D. Howell *et al.*, “Cytokine modulation of atopic dermatitis filaggrin skin expression,” *J Allergy Clin Immunol*, vol. 120, no. 1, pp. 150–155, Jul. 2007.

[22] Z. Zhang, T. B. Clarke, and J. N. Weiser, “Cellular effectors mediating Th17-dependent clearance of pneumococcal colonization in mice,” *J. Clin. Invest.*, vol. 119, no. 7, pp. 1899–1909, Jun. 2009.

[23] S. Celli, M. Day, A. J. Müller, C. Molina-Paris, G. Lythe, and P. Bousso, “How many dendritic cells are required to initiate a T-cell response?,” *Blood*, vol. 120, no. 19, pp. 3945–3948, Nov. 2012.

[24] M. Chen, L. Huang, Z. Shabier, and J. Wang, “Regulation of the lifespan in dendritic cell subsets,” *Mol. Immunol.*, vol. 44, no. 10, pp. 2558–2565, Apr. 2007.

[25] R. J. Tanaka, M. Ono, and H. A. Harrington, “Skin barrier homeostasis in atopic dermatitis: feedback regulation of kallikrein activity,” *PLoS One*, vol. 6, no. 5, p. e19895, 2011.

[26] T. Höfer, H. Nathansen, M. Löhning, A. Radbruch, and R. Heinrich, “GATA-3 transcriptional imprinting in Th2 lymphocytes: A mathematical model,” *Proc. Natl. Acad. Sci.*, vol. 99, no. 14, pp. 9364–9368, Jul. 2002.

[27] S. M. Langan, A. D. Irvine, and S. Weidinger, “Atopic dermatitis,” *Lancet*, vol. 396, no. 10247, pp. 345–360, Aug. 2020.

[28] J. A. Walker and A. N. J. McKenzie, “TH2 cell development and function,” *Nat. Rev. Immunol.*, vol. 18, no. 2, pp. 121–133, Feb. 2018.

[29] J. K. Gittler *et al.*, “Progressive activation of TH2/TH22 cytokines and selective epidermal proteins characterizes acute and chronic atopic dermatitis,” *J. Allergy Clin. Immunol.*, vol. 130, no. 6, pp. 1344–1354, Dec. 2012.

[30] M. F. Krummel, F. Bartumeus, and A. Gérard, “T cell migration, search strategies and mechanisms,” *Nat. Rev. Immunol.*, vol. 16, no. 3, pp. 193–201, Mar. 2016.

[31] D. L. Wallace *et al.*, “Direct Measurement of T Cell Subset Kinetics In Vivo in Elderly Men and Women,” *J. Immunol.*, vol. 173, no. 3, pp. 1787–1794, Aug. 2004.

[32] S. G. O. Johansson *et al.*, “Revised nomenclature for allergy for global use: Report of the Nomenclature Review Committee of the World Allergy Organization, October 2003,” *J. Allergy Clin. Immunol.*, vol. 113, no. 5, pp. 832–836, May 2004.

[33] Y. Yang, J. C. Ochando, J. S. Bromberg, and Y. Ding, “Identification of a distant T-bet enhancer responsive to IL-12/Stat4 and IFNγ/Stat1 signals,” *Blood*, vol. 110, no. 7, pp. 2494–2500, Oct. 2007.

[34] J. Zhu, “T helper 2 (Th2) cell differentiation, type 2 innate lymphoid cell (ILC2) development and regulation of interleukin-4 (IL-4) and IL-13 production,” *Cytokine*, vol. 75, no. 1, pp. 14–24, Sep. 2015.

[35] A. Y. Mitrophanov and E. A. Groisman, “Positive feedback in cellular control systems,” *BioEssays*, vol. 30, no. 6. Howard Hughes Medical Institute, pp. 542–555, Jun-2008.

[36] D. Angeli, J. E. Ferrell, and E. D. Sontag, “Detection of multistability, bifurcations, and hysteresis in a large class of biological positive-feedback systems,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 101, no. 7, pp. 1822–1827, Feb. 2004.

[37] T. Imai *et al.*, “Selective recruitment of CCR4-bearing Th2 cells toward antigen-presenting cells by the CC chemokines thymus and activation-regulated chemokine and macrophage-derived chemokine,” *Int. Immunol.*, vol. 11, no. 1, pp. 81–88, Jan. 1999.

[38] B. Yu *et al.*, “Differential regulation of thymus- and activation-regulated chemokine induced by IL-4, IL-13, TNF-α and IFN-γ in human keratinocyte and fibroblast,” *J. Dermatol. Sci.*, vol. 30, no. 1, pp. 29–36, Oct. 2002.

[39] G. Wirnsberger, D. Hebenstreit, G. Posselt, J. Horejs-Hoeck, and A. Duschl, “IL-4 induces expression of TARC/CCL17 via two STAT6 binding sites,” *Eur. J. Immunol.*, vol. 36, no. 7, pp. 1882–1891, Jul. 2006.

[40] S. A. Islam, M. F. Ling, J. Leung, W. G. Shreffler, and A. D. Luster, “Identification of human CCR8 as a CCL18 receptor,” *J. Exp. Med.*, vol. 210, no. 10, pp. 1889–1898, Sep. 2013.

[41] V. Kodelja, C. Müller, O. Politz, N. Hakij, C. E. Orfanos, and S. Goerdt, “Alternative Macrophage Activation-Associated CC-Chemokine-1, a Novel Structural Homologue of Macrophage Inflammatory Protein-1α with a Th2-Associated Expression Pattern,” *J. Immunol.*, vol. 160, no. 3, pp. 1411 LP – 1418, Feb. 1998.

[42] T. Hong, J. Xing, L. Li, and J. J. Tyson, “A simple theoretical framework for understanding heterogeneous differentiation of CD4+ T cells,” *BMC Syst. Biol.*, vol. 6, no. 1, p. 66, 2012.

[43] A. Yates, R. Callard, and J. Stark, “Combining cytokine signalling with T-bet and GATA-3 regulation in Th1 and Th2 differentiation: a model for cellular decision-making,” *J. Theor. Biol.*, vol. 231, no. 2, pp. 181–196, Nov. 2004.

[44] A. YATES, C. BERGMANN, J. LEO VAN HEMMEN, J. STARK, and R. CALLARD, “Cytokine-modulated Regulation of Helper T Cell Populations,” *J. Theor. Biol.*, vol. 206, no. 4, pp. 539–560, Oct. 2000.

[45] P. Allavena *et al.*, “IL-10 prevents the differentiation of monocytes to dendritic cells but promotes their maturation to macrophages,” *Eur. J. Immunol.*, vol. 28, no. 1, pp. 359–369, Jan. 1998.

[46] R. De Waal Malefyt *et al.*, “Interleukin 10 (il-10) and viral il-10 strongly reduce antigen-specific human t cell proliferation by diminishing the antigen-presenting capacity of monocytes via dowm’egulation of class h major histocompatibility complex expression,” *J. Exp. Med.*, vol. 174, no. 4, pp. 915–924, Oct. 1991.

[47] G. H. Fisher *et al.*, “Dominant interfering fas gene mutations impair apoptosis in a human autoimmune lymphoproliferative syndrome,” *Cell*, vol. 81, no. 6, pp. 935–946, Jun. 1995.

[48] R. Watanabe-Fukunaga, C. I. Brannan, N. G. Copeland, N. A. Jenkins, and S. Nagata, “Pillars article: lymphoproliferation disorder in mice explained by defects in Fas antigen that mediates apoptosis. 1992.,” *J. Immunol.*, vol. 189, no. 11, pp. 5101–4, Dec. 2012.

[49] S. Nagata and T. Suda, “Fas and Fas ligand: lpr and gld mutations,” *Immunol. Today*, vol. 16, no. 1, pp. 39–43, 1995.

[50] X. Zhang *et al.*, “Unequal Death in T Helper Cell (Th)1 and Th2 Effectors: Th1, but not Th2, Effectors Undergo Rapid Fas/FasL-mediated Apoptosis,” *J. Exp. Med.*, vol. 185, no. 10, pp. 1837–1849, May 1997.

[51] A. Maroof *et al.*, “Interleukin-4 can induce interleukin-4 production in dendritic cells,” *Immunology*, vol. 117, no. 2, pp. 271–279, Feb. 2006.

[52] S. L. Swain, A. D. Weinberg, M. English, and G. Huston, “IL-4 directs the development of Th2-like helper effectors,” *J. Immunol.*, vol. 145, no. 11, pp. 3796–3806, 1990.

[53] T. Huang *et al.*, “TGF-β signalling is mediated by two autonomously functioning TβRI:TβRII pairs,” *EMBO J.*, vol. 30, no. 7, pp. 1263–1276, Apr. 2011.

[54] O. Letourneur, J. F. Goetschy, M. Horisberger, and M. G. Grütter, “Ligand-Induced Dimerization of the Extracellular Domain of the TGF-β Receptor Type II,” *Biochem. Biophys. Res. Commun.*, vol. 224, no. 3, pp. 709–716, Jul. 1996.

[55] L. Ruiz *et al.*, “Unveiling the dimer/monomer propensities of Smad MH1-DNA complexes,” *Comput. Struct. Biotechnol. J.*, vol. 19, p. 632, Jan. 2021.

[56] M. A. Burke, B. F. Morel, T. B. Oriss, J. Bray, S. A. McCarthy, and P. A. Morel, “Modeling the proliferative response of T cells to IL-2 and IL-4,” *Cell. Immunol.*, vol. 178, no. 1, pp. 42–52, May 1997.

[57] A. W. Borkowski and R. L. Gallo, “The coordinated response of the physical and antimicrobial peptide barriers of the skin,” *J Invest Dermatol*, vol. 131, no. 2, pp. 285–287, 2011.

[58] M. E. Hibbing, C. Fuqua, M. R. Parsek, and S. B. Peterson, “Bacterial competition: surviving and thriving in the microbial jungle,” *Nat Rev Microbiol*, vol. 8, no. 1, pp. 15–25, 2010.

[59] R. L. Gallo and L. V Hooper, “Epithelial antimicrobial defence of the skin and intestine,” *Nat. Rev. Immunol.*, vol. 12, no. 7, pp. 503–516, 2012.

[60] E. Goleva, E. Berdyshev, and D. Y. M. Leung, “Epithelial barrier repair and prevention of allergy,” *J. Clin. Invest.*, vol. 129, no. 4, pp. 1463–1474, Apr. 2019.

[61] M. Suárez-Fariñas *et al.*, “Nonlesional atopic dermatitis skin is characterized by broad terminal differentiation defects and variable immune abnormalities.,” *J. Allergy Clin. Immunol.*, vol. 127, no. 4, p. 954, Apr. 2011.

[62] E. Guttman-Yassky *et al.*, “Broad defects in epidermal cornification in atopic dermatitis identified through genomic analysis,” *J. Allergy Clin. Immunol.*, vol. 124, no. 6, pp. 1235-1244.e58, Dec. 2009.

[63] F. Mascia, M. Denning, R. Kopan, and S. H. Yuspa, “The Black Box Illuminated: Signals and Signaling,” *J. Invest. Dermatol.*, vol. 132, no. 3, pp. 811–819, Mar. 2012.

[64] E. L. Simpson *et al.*, “Two Phase 3 Trials of Dupilumab versus Placebo in Atopic Dermatitis.,” *N. Engl. J. Med.*, vol. 375, no. 24, pp. 2335–2348, Dec. 2016.

[65] K. Kennedy, J. Heimall, and J. M. Spergel, “Advances in atopic dermatitis in 2017,” *J. Allergy*

Clin. Immunol., vol. 142, no. 6, pp. 1740–1747, Dec. 2018.

[66] E. H. Tham, E. Koh, J. E. A. Common, and I. Y. Hwang, "Biotherapeutic Approaches in Atopic Dermatitis," *Biotechnol. J.*, vol. 15, no. 10, p. 1900322, Oct. 2020.

[67] T. Nakatsuji *et al.*, "Development of a human skin commensal microbe for bacteriotherapy of atopic dermatitis and use in a phase 1 randomized clinical trial," *Nat. Med.*, vol. 27, no. 4, pp. 700–709, 2021.

[68] R. S. Moreci and T. Lechler, "Epidermal structure and differentiation," *Curr. Biol.*, vol. 30, no. 4, pp. R144–R149, Feb. 2020.

[69] V. A. Botchkarev, "The Molecular Revolution in Cutaneous Biology : Chromosomal Territories , Higher-Order Chromatin Remodeling , and the Control of Gene Expression in Keratinocytes," *J. Invest. Dermatol.*, vol. 137, no. 5, pp. e93–e99, 2016.

[70] M. Yi *et al.*, "TP63 links chromatin remodeling and enhancer reprogramming to epidermal differentiation and squamous cell carcinoma development," *Cell. Mol. Life Sci.*, vol. 77, no. 21, pp. 4325–4346, 2020.

[71] A. Smirnov *et al.*, "ZNF185 is a p63 target gene critical for epidermal differentiation and squamous cell carcinoma development," *Oncogene*, vol. 38, no. 10, pp. 1625–1638, 2019.

[72] K. Toufighi *et al.*, "Dissecting the Calcium-Induced Differentiation of Human Primary Keratinocytes Stem Cells by Integrative and Structural Network Analyses," vol. 11, no. 5, May 2015.

[73] N. A. Shah and C. A. Sarkar, "Robust Network Topologies for Generating Switch-Like Cellular Responses," *PLOS Comput. Biol.*, vol. 7, no. 6, p. e1002085, 2011.

[74] A. S. Borowiec, P. Delcourt, E. Dewailly, and G. Bidaux, "Optimal Differentiation of In Vitro Keratinocytes Requires Multifactorial External Control," *PLoS One*, vol. 8, no. 10, p. e77507, Oct. 2013.

[75] J. E. Dazard, J. Piette, N. Basset-Seguin, J. M. Blanchard, and A. Gandarillas, "Switch from p53 to MDM2 as differentiating human keratinocytes lose their proliferative potential and increase in cellular size," *Oncogene 2000 1933*, vol. 19, no. 33, pp. 3693–3705, Jul. 2000.

[76] E. Ogawa *et al.*, "p63/p51-induced Onset of Keratinocyte Differentiation via the c-Jun N-terminal Kinase Pathway Is Counteracted by Keratinocyte Growth Factor," *J. Biol. Chem.*, vol. 283, no. 49, pp. 34241–34249, Dec. 2008.

[77] A. Jadali and S. Ghazizadeh, "Protein Kinase D Is Implicated in the Reversible Commitment to Differentiation in Primary Cultures of Mouse Keratinocytes," *J. Biol. Chem.*, vol. 285, no. 30, pp. 23387–23397, Jul. 2010.

[78] M. D. Howell *et al.*, "Cytokine modulation of atopic dermatitis filaggrin skin expression," *J. Allergy Clin. Immunol.*, vol. 124, no. 3, pp. R7–R12, Sep. 2009.

[79] B. C. Nguyen *et al.*, "Cross-regulation between Notch and p63 in keratinocyte commitment to differentiation," *Genes Dev.*, vol. 20, no. 8, pp. 1028–1042, Apr. 2006.

[80] T. Sen, X. Chang, D. Sidransky, and A. Chatterjee, "Regulation of ΔNp63α by NFκB," *http://dx.doi.org/10.4161/cc.9.24.14093*, vol. 9, no. 24, pp. 4841–4847, Dec. 2010.

[81] A. W. Borkowski, K. Park, Y. Uchida, and R. L. Gallo, "Activation of TLR3 in Keratinocytes Increases Expression of Genes Involved in Formation of the Epidermis, Lipid Accumulation, and Epidermal Organelles," *J. Invest. Dermatol.*, vol. 133, no. 8, pp. 2031–2040, Aug. 2013.

[82] P. Duckney, H. K. Wong, J. Serrano, D. Yaradou, T. Oddos, and G. N. Stamatas, “The role of the skin barrier in modulating the effects of common skin microbial species on the inflammation, differentiation and proliferation status of epidermal keratinocytes,” *BMC Res. Notes*, vol. 6, no. 1, pp. 1–10, Nov. 2013.

[83] S. E. Lee, S. K. Jeong, and S. H. Lee, “Protease and Protease-Activated Receptor-2 Signaling in the Pathogenesis of Atopic Dermatitis,” *Yonsei Med. J.*, vol. 51, no. 6, p. 808, Nov. 2010.

[84] H. Dommisch *et al.*, “Protease-Activated Receptor 2 Mediates Human Beta-Defensin 2 and CC Chemokine Ligand 20 mRNA Expression in Response to Proteases Secreted by Porphyromonas gingivalis,” *Infect. Immun.*, vol. 75, no. 9, pp. 4326–4333, Sep. 2007.

[85] A. Abtin, L. Eckhart, R. Gläser, R. Gmeiner, M. Mildner, and E. Tschachler, “The Antimicrobial Heterodimer S100A8/S100A9 (Calprotectin) Is Upregulated by Bacterial Flagellin in Human Epidermal Keratinocytes,” *J. Invest. Dermatol.*, vol. 130, no. 10, pp. 2423–2430, Oct. 2010.

[86] I. Wanke *et al.*, “Skin Commensals Amplify the Innate Immune Response to Pathogens by Activation of Distinct Signaling Pathways,” *J. Invest. Dermatol.*, vol. 131, no. 2, pp. 382–390, Feb. 2011.

[87] P. M. Elias and J. S. Wakefield, “Mechanisms of abnormal lamellar body secretion and the dysfunctional skin barrier in patients with atopic dermatitis.,” *J. Allergy Clin. Immunol.*, vol. 134, no. 4, pp. 781-791.e1, Oct. 2014.

[88] J.-P. Hachem *et al.*, “Serine Protease Signaling of Epidermal Permeability Barrier Homeostasis,” *J. Invest. Dermatol.*, vol. 126, no. 9, pp. 2074–2086, Sep. 2006.

[89] Y. Hatano *et al.*, “NIH Public Access,” *J. Invest. Dermatol.*, vol. 129, no. 7, pp. 1824–1835, 2009.

[90] K. R. Feingold, “Thematic review series: skin lipids. The role of epidermal lipids in cutaneous permeability barrier homeostasis.,” *J. Lipid Res.*, vol. 48, no. 12, pp. 2531–46, Dec. 2007.

[91] P. M. Elias *et al.*, “Modulations in Epidermal Calcium Regulate the Expression of Differentiation-Specific Markers,” *J. Invest. Dermatol.*, vol. 119, no. 5, pp. 1128–1136, Nov. 2002.

[92] G. K. Menon, K. R. Feingold, and P. M. Elias, “Lamellar Body Secretory Response to Barrier Disruption,” *J. Invest. Dermatol.*, vol. 98, pp. 279–289, 1992.

[93] D. D. Bikle, Z. Xie, and C.-L. Tu, “Calcium regulation of keratinocyte differentiation,” *Expert Rev Endocrinol Metab*, vol. 7, no. 4, pp. 461–472, 2013.

[94] M. R. Mack and B. S. Kim, “The Itch–Scratch Cycle: A Neuroimmune Perspective,” *Trends Immunol.*, vol. 39, no. 12, pp. 980–991, Dec. 2018.

[95] K. M. Aberg *et al.*, “Co-Regulation and Interdependence of the Mammalian Epidermal Permeability and Antimicrobial Barriers,” *J. Invest. Dermatol.*, vol. 128, no. 4, pp. 917–925, Apr. 2008.

[96] J. Harder *et al.*, “Enhanced Expression and Secretion of Antimicrobial Peptides in Atopic Dermatitis and after Superficial Skin Injury,” *J. Invest. Dermatol.*, vol. 130, no. 5, pp. 1355–1364, May 2010.

[97] K. O. Kisich, C. W. Carspecken, S. Fiéve, M. Boguniewicz, and D. Y. M. Leung, “Defective killing of Staphylococcus aureus in atopic dermatitis is associated with reduced mobilization of human β-defensin-3,” *J. Allergy Clin. Immunol.*, vol. 122, no. 1, pp. 62–68, Jul. 2008.

[98] V. Soumelis *et al.*, “Human epithelial cells trigger dendritic cell–mediated allergic inflammation by producing TSLP,” *Nat. Immunol. 2002 37*, vol. 3, no. 7, pp. 673–680, Jun. 2002.

[99] J. Liu, X. Zhang, Y. Cheng, and X. Cao, “Dendritic cell migration in inflammation and immunity,” *Cell. Mol. Immunol. 2021 1811*, vol. 18, no. 11, pp. 2461–2471, Jul. 2021.

[100] T. Worbs, S. I. Hammerschmidt, and R. Förster, “Dendritic cell migration in health and disease,” *Nat. Rev. Immunol.*, vol. 17, no. 1, pp. 30–48, Jan. 2017.

[101] T. Ito *et al.*, “TSLP-activated dendritic cells induce an inflammatory T helper type 2 cell response through OX40 ligand,” *J. Exp. Med.*, vol. 202, no. 9, pp. 1213–1223, 2005.

[102] S. J. Galli, M. Tsai, and A. M. Piliponsky, “The development of allergic inflammation,” *Nature*, vol. 454, no. 7203, pp. 445–454, Jul. 2008.

[103] A. Elbe-Bürger *et al.*, “Overexpression of IL-4 Alters the Homeostasis in the Skin,” *J. Invest. Dermatol.*, vol. 118, no. 5, pp. 767–778, May 2002.

[104] P. Ahmad-Nejad *et al.*, “The Toll-like receptor 2 R753Q polymorphism defines a subgroup of patients with atopic dermatitis having severe phenotype [4],” *J. Allergy Clin. Immunol.*, vol. 113, no. 3, pp. 565–567, Mar. 2004.

[105] L. Y. McGirt and L. A. Beck, “Innate immune defects in atopic dermatitis,” *J Allergy Clin Immunol*, vol. 118, no. 1, pp. 202–208, 2006.

[106] M. L. Clausen, H. C. Slotved, K. A. Krogfelt, P. S. Andersen, and T. Agner, “In vivo expression of antimicrobial peptides in atopic dermatitis,” *Exp. Dermatol.*, vol. 25, no. 1, pp. 3–9, Jan. 2016.

[107] G. Tanaka, E. Domínguez-Hüttinger, P. Christodoulides, K. Aihara, and R. J. Tanaka, “Bifurcation analysis of a mathematical model of atopic dermatitis to determine patient-specific effects of treatments on dynamic phenotypes,” *J. Theor. Biol.*, vol. 448, pp. 66–79, Jul. 2018.

[108] R. Chopra, P. P. Vakharia, R. Sacotte, and J. I. Silverberg, “Efficacy of bleach baths in reducing severity of atopic dermatitis: A systematic review and meta-analysis,” *Ann. Allergy, Asthma Immunol.*, vol. 119, no. 5, pp. 435–440, Nov. 2017.

[109] D. C. Macallan *et al.*, “Rapid turnover of effector-memory CD4+ T cells in healthy humans,” *J. Exp. Med.*, vol. 200, no. 2, pp. 255–260, Jul. 2004.

[110] I. A. Rather, V. K. Bajpai, S. Kumar, J. Lim, W. K. Paek, and Y.-H. Park, “Probiotics and Atopic Dermatitis: An Overview,” vol. 7, no. 507, 2016.

[111] T. Nakatsuji *et al.*, “Antimicrobials from human skin commensal bacteria protect against Staphylococcus aureus and are deficient in atopic dermatitis,” *Sci. Transl. Med.*, vol. 9, no. 378, p. eaah4680, Feb. 2017.

[112] M. Baldry *et al.*, “Application of an agr-Specific Antivirulence Compound as Therapy for Staphylococcus aureus-Induced Inflammatory Skin Disease,” *J Infect Dis*, vol. 218, no. 6, pp. 1009–1013, 2018.

[113] R. G. Sanfelice, D. A. Copp, and P. Ñañez, “A toolbox for simulation of hybrid systems in matlab/simulink,” *HSCC 2013 - Proc. 16th Int. Conf. Hybrid Syst. Comput. Control. Part CPSWeek 2013*, vol. 1, pp. 101–106, 2012.

[114] M. R. Williams, T. Nakatsuji, J. A. Sanford, A. F. Vrbanac, and R. L. Gallo, “Staphylococcus aureus Induces Increased Serine Protease Activity in Keratinocytes,” *J. Invest. Dermatol.*, vol. 137, no. 2, pp. 377–384, Feb. 2017.

[115] K. Yamasaki *et al.*, “TLR2 Expression Is Increased in Rosacea and Stimulates Enhanced Serine Protease Production by Keratinocytes,” *J. Invest. Dermatol.*, vol. 131, no. 3, pp. 688–697, Mar. 2011.

[116] L. F. M. Ariens, D. S. Bakker, J. van der Schaft, F. M. Garritsen, J. L. Thijs, and M. S. de Bruin-Weller, “Dupilumab in atopic dermatitis: rationale, latest evidence and place in therapy,” *Ther. Adv. Chronic Dis.*, vol. 9, no. 9, pp. 159–170, 2018.

[117] S. McGregor, M. E. Farhangian, and S. R. Feldman, “Dupilumab for the treatment of atopic dermatitis: A clinical trial review,” *Expert Opin Biol Ther*, vol. 15, no. 11, pp. 1657–1660, 2015.

Appendix A Modifying the K-switch

In this appendix, we describe a modification to a model of kallikrein regulation in the epidermis [25]. This model is used in the double-switch model to justify the "sloped" K-switch [18]. The modification we make here has no effect on the rest of the modelling work in this thesis, but is nonetheless significant in that it provides a more accurate assumption for our model.

Introduction

Keratinocytes express a range of innate immune receptors such as protease activated receptor 2 (PAR2), which are constantly vigilant to common threats from potentially pathogenic bacteria and chemical irritants. When activated, PAR2 receptor induces innate immune responses which attract more sophisticated immune responses such as dendritic cells which can call on Th2 cells (i.e. the T-switch).

One of the effects of PAR2 activation is to increase the rate of production of active kallikreins (KLKs), which are extracellular enzymes which break down the corneodesmosomes which connect together the corneocytes in the uppermost layer of the skin barrier. The effect of this breaking is to increase the rate of shedding of the barrier, potentially ejecting the source of the environmental assault, albeit at the expense of some barrier integrity.

Given the crucial nature of maintaining barrier integrity, it comes as no surprise that KLK levels are regulated by a serine protease inhibitor molecule with the rather overwhelming name; lympho-epithelial kazal-type-related inhibitor (LEKTI). When it binds to activated KLK, it prevents it from cleaving the corneodesmosomes. We focus on this regulatory network in this model.

Two models for kallikrein regulation

The previously published model of PAR2 activation [25] demonstrated that their activation levels are described by a bistable switch with hysteresis with respect to the level of external stimulus, S, received by the skin. However, this model did not have a clear justification for the effect of external stimulus on the production of KLK5 precursors and LEKTI.

We modify the model and demonstrate that similar bistable behaviours emerge for the activation level of both KLK5 and PAR2. We redefine S as something that activates PAR2. We consider the following sources of PAR2 activation;

1. exogeneous sources such as allergens
2. bacteria such as $S. aureus$ which stimulate keratinocytes to release trypsin [114]
3. activated KLK enzyme

Previously published, "old" model

The previously published kallikrein model considers the regulation network in figure A.1Figure . The meanings of the letters are explained below.

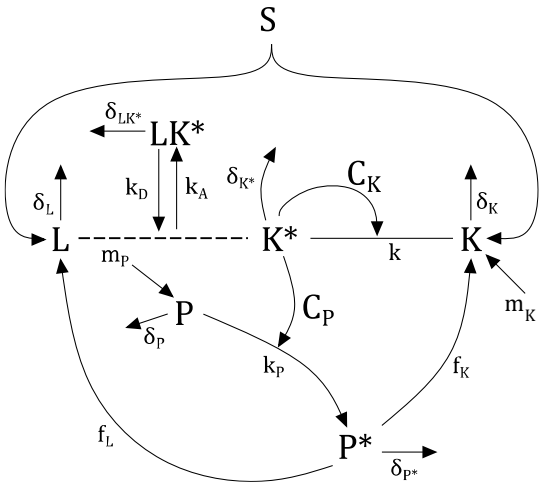


Figure A.1 Model schematic of the published kallikrein model

The model (figure A.1) is represented by the ODEs in equation A.1.

$$
\begin{aligned}
\frac{d[K]}{dt} &= -k\frac{[K^*][K]}{[K^*]+C_K} - \delta_K[K] + f_K[\mathrm{P}^*] + f_{KS}S + m_K, \\
\frac{d[L]}{dt} &= -k_A[K^*][L] + k_D[LK^*] - \delta_L[L] + t_L\left(f_L[P^*] + f_{LS}S + m_L\right), \\
\frac{d[K^*]}{dt} &= -k_A[K^*][L] + k_D[LK^*] + k\frac{[K^*][K]}{[K^*]+C_K} - \delta_{K^*}[K^*], \\
\frac{d[LK^*]}{dt} &= k_A[K^*][L] - [LK^*](\delta_{LK} + k_D), \\
\frac{d[P]}{dt} &= -k_P\frac{[K^*][P]}{[K^*]+C_P} - \delta_P[P] + m_P, \\
\frac{d[P^*]}{dt} &= k_P\frac{[K^*][P]}{[K^*]+C_P} - \delta_{P^*}[P^*],
\end{aligned}
\tag{A.1}
$$

where $[K]$, $[L]$, $[K^*]$, $[LK^*]$, $[P]$ and $[P^*]$ represent the concentrations of inactive KLK5 precursors, LEKTI, activated KLK, LEKTI- KLK complex, PAR2 receptors and activated PAR2 receptors, respectively. KLK5 precursors $[K]$ are activated by active KLK5s $[K^*]$, which activate PAR2 receptors irreversibly. LEKTI $[L]$ binds reversibly to active KLK5s $[K^*]$, forming LEKTI-activated KLK5 complex $[LK^*]$. PAR2 receptors $[P]$ are produced at a constant rate by keratinocytes. Production of LEKTI and KLK5 precursors is enhanced by activated PAR2 receptors $[P^*]$ and external stimulation.

The steady states, $[\overline{P^*}]$ and $[\overline{K^*}]$, both show a bistable switch with respect to external stimulus S (figure A.2). When S exceeds the external stimulus threshold S^+, both steady states converge to a high level, where the PAR2 activity $P^* > 0$ leading to increased expression of inflammatory genes. When S is below another threshold S^-, the PAR2 receptors are not activated, and so the skin does not become inflamed.

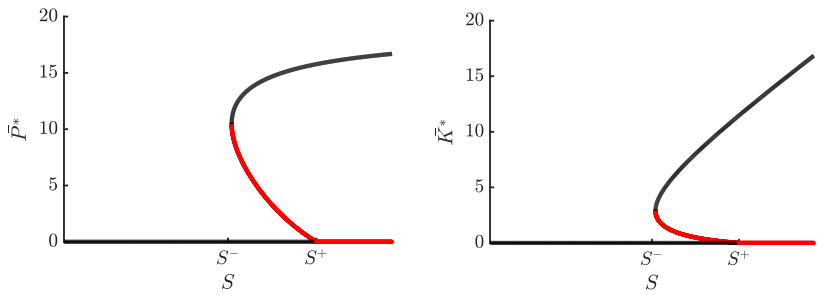


Figure A.2 Bifurcation diagrams of the published model of kallikrein regulation. The black and red lines correspond to stable and unstable steady states respectively. The parameter values are shown in Table .

New kallikrein model

In the double-switch model of AD [18], it is assumed that the PAR2 receptors (or innate immune receptors) are directly activated by external stimulus (or "environmental stressors" as called in the double-switch model) as represented in the schematic in figure A.3, left. However, in [25], the external stimulus S is assumed to enhance production of LEKTI and KLK5 precursors, rather than activation of P. Moreover, strong justification was lacking for the assumption that the rate of production of KLK precursors ($f_{KS}S$) and LEKTI ($f_{LS}S$) is proportional to the strength of external stimulation.

We therefore modified the original model to make it more compatible with the double-switch model, by redefining S as the entities that activate PAR2 and by assuming that PAR2 is activated by S through the action of exogeneous proteases or mast-cell derived tryptase released in response to pathogen recognition [83] (figure A.3, right).

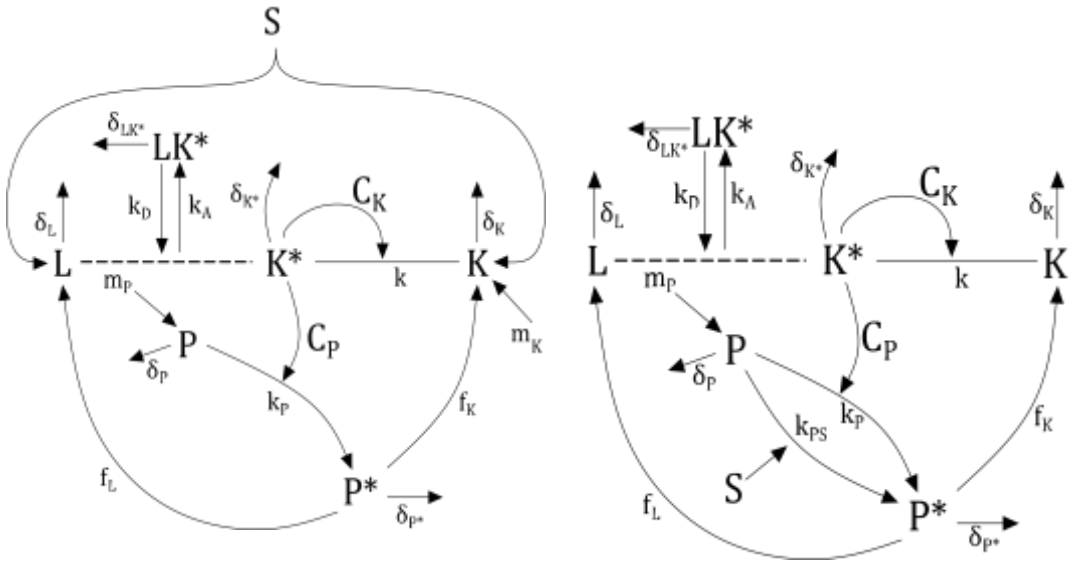


Figure A.3 Schematic of the published model (left). Schematic of the modified model of kallikrein activation (right). In the modified model (right), the external stimulus S activates PAR2 receptors P into activated PAR2 P^* at rate k_{PS}.

The modified part is highlighted in the modified ODEs of equation A.2, highlighted in green.

$$
\begin{aligned}
\frac{d[LK^*]}{dt} &= k_A[K^*][L] - [LK^*](\delta_{LK^*} + k_D), \\
\frac{d[L]}{dt} &= -k_A[K^*][L] + k_D[LK^*] - \delta_L[L] + t_L(f_L[P^*] + m_L), \\
\frac{d[K^*]}{dt} &= -k_A[K^*][L] + k_D[LK^*] + k\frac{[K^*][K]}{[K^*] + C_K} - \delta_{K^*}[K^*], \\
\frac{d[K]}{dt} &= -k\frac{[K^*][K]}{[K^*] + C_K} - \delta_K[K] + f_K[\mathrm{P}^*] + m_K, \\
\frac{d[P]}{dt} &= -k_P\frac{[K^*][P]}{[K^*] + C_P} - k_{PS}\frac{[S][P]}{[S] + 1} - \delta_P[P] + m_P, \\
\frac{d[P^*]}{dt} &= k_P\frac{[K^*][P]}{[K^*] + C_P} + k_{PS}\frac{[S][P]}{[S] + 1} - \delta_{P^*}[P^*].
\end{aligned}
\tag{A.2}
$$

For simplicity, we assume all protein species to decay at the same rate ($\delta_L = \delta_{LK^*} = \delta_{K^*} = \delta_K = \delta_{P^*} = \delta_P$) and some reaction rates are equal ($k_P = k_{PS} = k$).

The bifurcation diagrams for $[\overline{P^*}]$ and $[\overline{K^*}]$ (figure A.4) demonstrate a bistable switch with respect to S, as for the previous model.

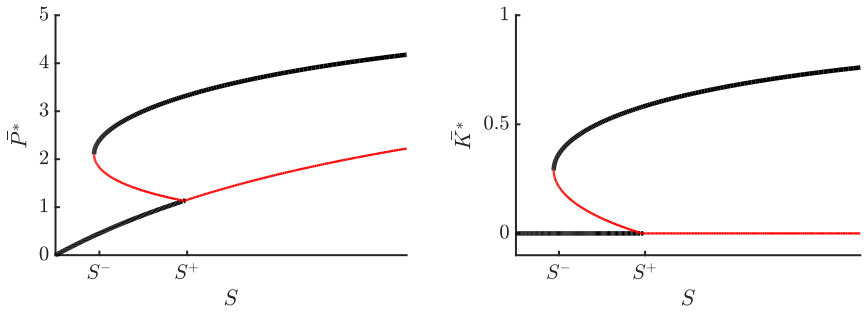


Figure A.4 Bifurcation diagrams for the healthy control. The black and red lines correspond to stable and unstable steady states respectively. The parameter values are shown in Table 1.

We also confirmed the similar results as for the previous model:

- Decreased production of LEKTI and increased pH lead to a decrease in S^-, making it harder to switch off inflammation (figure A.5).
- Increased expression of active KLK5 $\overline{K^*}$ for decreased production of LEKTI and increased pH (figure A.5 (b)).
- The inflammatory state becomes more persistent for increased strength of KLK5 feedback f_K (Figure A.6).

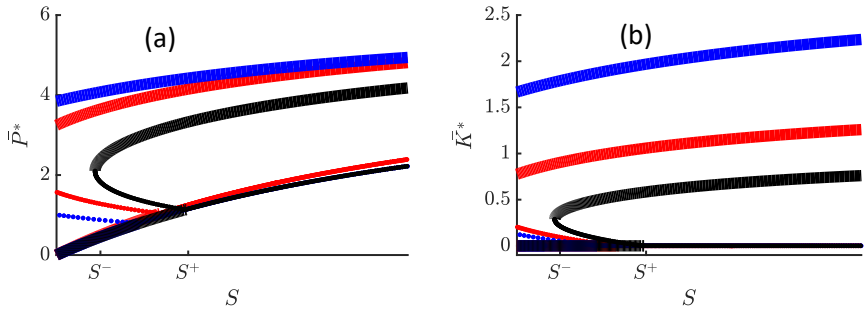


Figure A.5 Behaviours of the model with nominal "healthy" parameters from table 1 (black), high pH parameters (red) and LEKTI-deficient parameters (blue). Thick and thin lines represent stable and unstable steady states respectively.

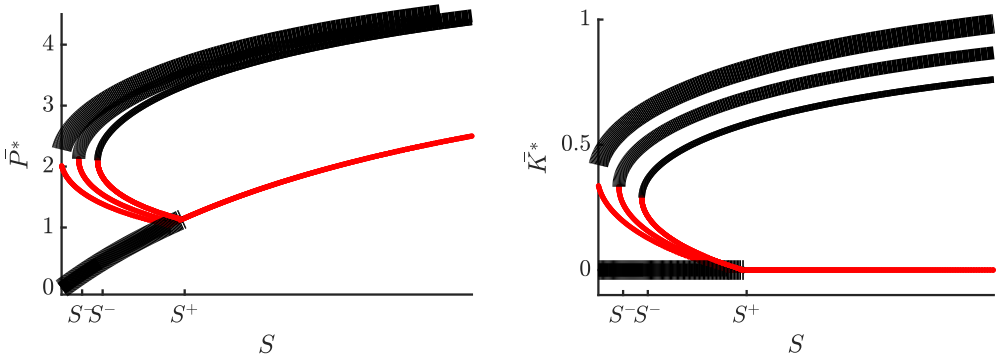


Figure A.6 Stronger positive feedback leads to more persistent inflammation. The strength of positive feedback corresponds to the thickness of each of the lines. Black and red lines correspond to stable and unstable steady states respectively.

K-switch conclusion

The revised model produces bistable behaviours (Figure (b)) but with an approximately linearly increasing $\overline{P^*}$ for $0 < S < S^+$, and an activated kallikrein level $\overline{K^*}$ remaining at zero. The approximately linear $\overline{P^*}$ response for low values of S results from the fact that the external stimulus S acts directly on PAR2 activation (P^*) in the modified model [83], rather than indirectly acting on them through KLK5 (K) and LEKTI (L) in the published model [25]. For low values of external stimulus ($S \ll 1$), the Hill function activation of the PAR2 receptor, $k_{PS}\frac{[S][P]}{[S]+1}$, is approximately proportional to the level of external stimulation S leading to the linear relationship where the Hill function may be approximated by $k_{PS}[S][P]$.

As S is increased, above 1, the steady-state values $\overline{P^*}$ and $\overline{K^*}$ each tend towards a finite limit. For these high values of external stimulus ($S \gg 1$), the Hill function becomes increasingly flat as the PAR2 receptors become saturated, with the Hill function approximated by $k_{PS}[P]$.

So, having seen these differences in the bifurcation diagrams for the new kallikrein model versus the old one, what does this mean for the "K-switch"? The old model produced a linearly increasing concentration of activated kallikreins, and we approximated this by a sloped upper branch of the K-switch definition $K_{\mathrm{on}} = Pm_{\mathrm{on}} - \beta$ (Figure left). In contrast, the new model results in a level of activated kallikreins which tends to a constant limit, so we approximate this by a "flat" K-switch definition, where K_{on} is just a constant (figure A.7 right).

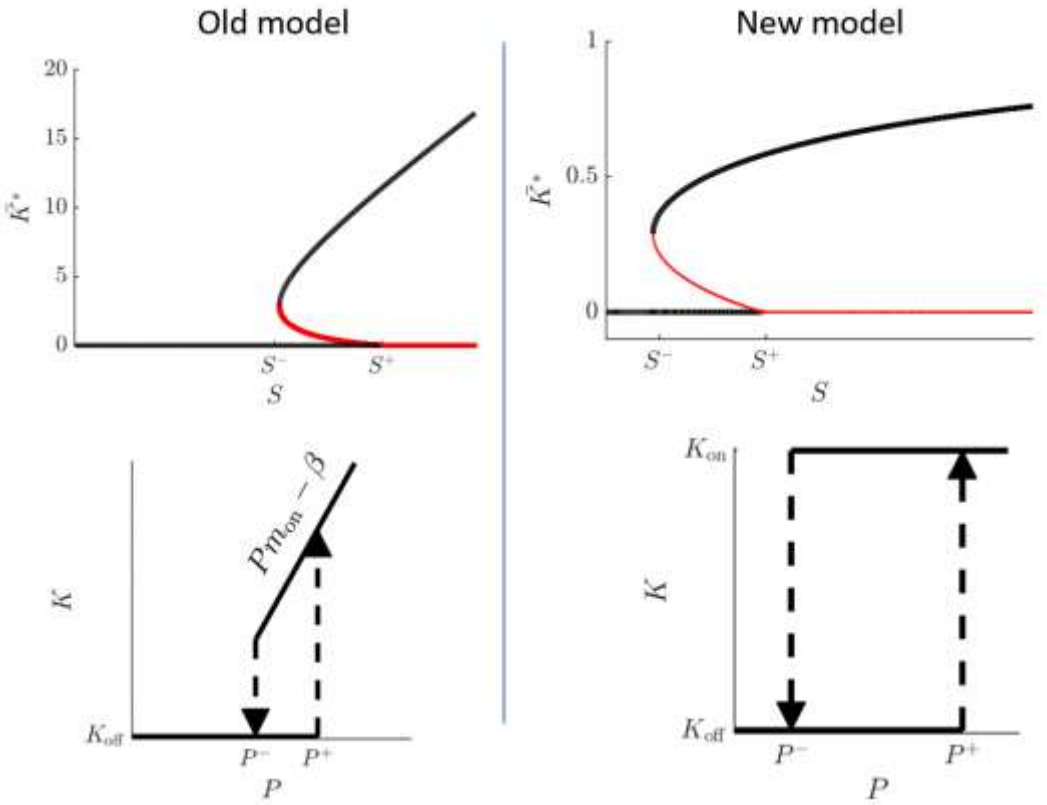


Figure A.7 Approximation of the kallikrein bifurcation diagram as a switch function. Left: the old kallikrein model, which was approximated by a sloped switch. Right: the new kallikrein model, which we approximate as a flat switch.

The previous double-switch model of AD [18] justified the bistable "R-switch" (which represents innate immune inflammatory responses) as being brought about by a bistable response of toll-like receptors (TLRs) to pathogen concentration in the skin. TLRs are receptors on keratinocytes which recognise danger-associated molecular patterns and pathogen-associated molecular patterns [115]. We assume that TLR activation and PAR2 activation occur at the same threshold values since pathogen-associated molecules and proteases can both originate from pathogens present on the skin. For this reason, the R- and K-switches still share the same threshold values P^- and P^+.

Table A.1 Parameters used to produce bistable KLK system.

Parameter	Interpretation	Value
k	Rate of proteolysis by active KLK5	1 (pH=4.5) 1.2 (pH=6.5)
k_A	Rate of LEKTI-KLK5 complex production	3.1 (pH=4.5) 3.6 (pH=6.5)
k_D	Rate of LEKTI-KLK5 complex degradation	1 (pH=4.5) 0.8 (pH=6.5)
κ_P	PAR2 activation rate	$\approx k$
δ_L	LEKTI degradation rate	1
δ_{LK}	LEKTI-KLK5* degradation rate	$\approx \delta_K$
δ_{K*}	KLK5* degradation rate	$\approx \delta_K$
δ_P	PAR2 degradation rate	0.5
δ_{P*}	PAR2* degradation rate	$\approx \delta_P$
C_K	Concentration of KLK5 required to activate inactive KLK at half the maximum rate	1
C_P	Concentration of KLK5 required to activate PAR2 receptors at half the maximum rate	1
m_L	Basal rate of production of LEKTI	1
m_K	Basal rate of production of KLK5 precursors	0
m_P	Basal rate of production of PAR2 receptors	10
f_L	Feedback strength from activated PAR2 receptors to LEKTI	5
f_K	Feedback strength from activated PAR2 receptors to KLK5 precursors	10
t_L	LEKTI production capability	1 (Healthy) 0.1 (LEKTI deficient)

Effect of flattening the K-switch on the model's bifurcation diagram

After we apply these changes to the model underlying the "K" part of the RK-switch, we need to consider the potential knock-on effect on the behaviours of the double-switch model. This process involves simply repeating our analysis from section 2.3.3 where we calculated the location of the BE bifurcation lines. In section 2.3, we had to make special consideration of the ODEs when the RK-switch is turned on because the K_{on} value becomes dependent on the state variable P. This P-dependence caused the equation for the focal point value to become quadratic (equation 2.7).

Now that we have modified the K-switch and concluded that K_{on} is just a constant, the process of calculating the focal point values of the double-switch model with the RK-switch on becomes somewhat easier. The equation for the focal point value may be summarised in a single form as equation A.3, making us of the inputs R and T to specify the state of the RK- and T-switch respectively. Note that this is exactly the same form as 2.5 except that it now generalises for all the combinations of switch configuration the model could take.

$$\bar{P}(R,T) = P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1+\gamma_{\text{B}}\bar{B}(R,T)\right)(\delta_{\text{P}}+\alpha_I R)}$$
$$\bar{B}(R,T) = \frac{\kappa_B}{\kappa_B + \delta_B K(1+\gamma_R \text{R})(1+\gamma_T T)} \quad \text{(A.3)}$$
$$\bar{D}(R,T) = \frac{\kappa_D \text{R}}{\delta_D}$$

Repeating the process of calculating the BE bifurcation lines, and once again plotting these in the κ_P, α_I plane, we see that the state of the T-switch still has not noticeable impact on the locations of the bifurcation lines (figure A.8).

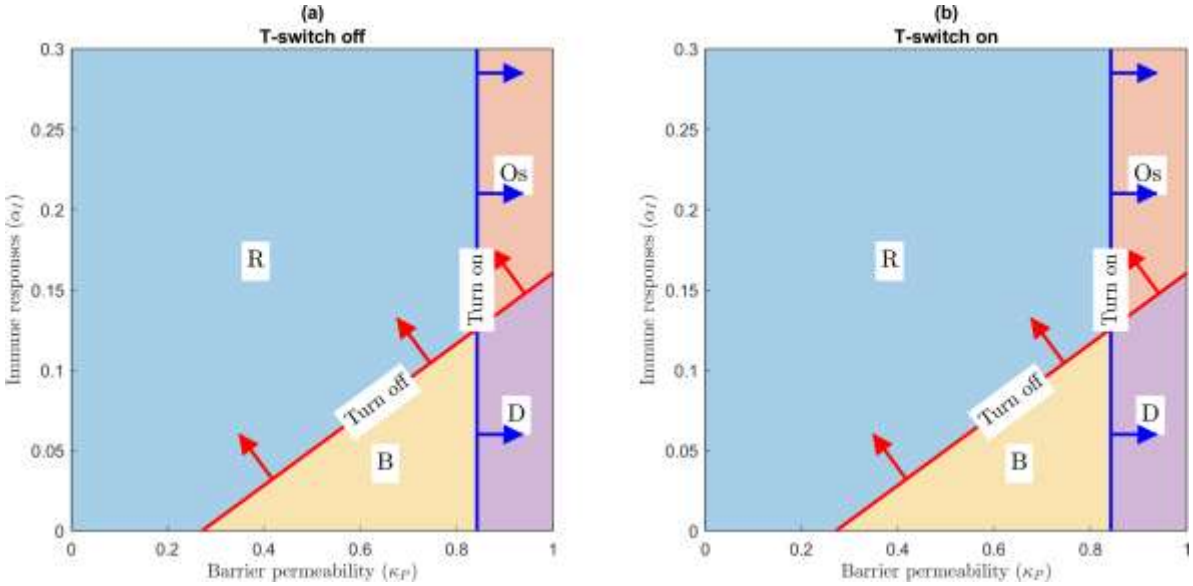


Figure A.8 Effect of flattening the K-switch on the bifurcation diagrams.

We can conclude that the modified assumptions made in the construction of the ODE model justifying the K-switch has had no noticeable impact on the macroscopic behaviour of the model for AD. The layout of the bifurcation diagram is the same, meaning that qualitative behaviours of the system remain unchanged relative to the previous model. The change we have made to the K-switch is nevertheless important in that it improves our confidence in our founding assumptions.

Conclusion

We've been looking for modifications we can make to the double-switch model that will help us describe chronic AD. Given the lack of impact on the double-switch model, we now know this change isn't it. However, the lack of impact on the double-switch AD model is not necessarily a bad thing.

We have improved the founding assumptions of the model for KLK regulation. This new model has improved agreement with the biological processes described in the literature, so we now have greater confidence in our model results. For this reason, use this modified KLK for the model of chronic AD developed in this thesis.

Appendix B Focal points and stability of the triple-switch model ODEs

In this appendix, we calculate the equilibrium values of the ODEs for the triple-switch model, and show that they are always stable.

The triple-switch ODEs are given by

$$\begin{aligned}
\frac{dP(t)}{dt} &= (P_{env} + P_{ext})\frac{\kappa_P}{1+\gamma_{PB}\cdot B} - \alpha_P \cdot A \cdot P - \delta_P \cdot P \\
\frac{dB(t)}{dt} &= \kappa_B \cdot \frac{1}{1+\gamma_{BB}\cdot B}\cdot\frac{1}{1+\gamma_{BR}\cdot R(P^-,P^+)} G(B^-,B^+) - \delta_B \cdot K(P^-,P^+)\cdot B - \delta_{Bext}(t)\cdot B \\
\frac{dA(t)}{dt} &= \kappa_A \cdot \frac{1}{1+\gamma_{AB}\cdot B}\cdot G(B^-,B^+) - \delta_A \cdot A \\
\frac{dD(t)}{dt} &= \kappa_D \cdot R(P^-,P^+) + D_{ext}(t) - \delta_D \cdot D
\end{aligned} \tag{B.1}$$

We assume that the external perturbation terms, $P_{ext}(t)$, $\delta_{Bext}(t)$ and $D_{ext}(t)$, remain constant, $P_{ext}(t) = \overline{P_{ext}}$, $\delta_{Bext}(t) = \overline{\delta_{Bext}}$ and $D_{ext}(t) = \overline{D_{ext}}$, and consider the switch states R, K, T and G constant. We replace the value of the switches R, K, T and G with the values for the switches off or on state later in our analysis to investigate how changing the switch state affects the focal point value to which our system will converge.

The focal points (described using an overbar, e.g., $\bar{P}$) are the values of the state variables P, B A and D at which the value of the derivative is zero. These are obtained by setting the left-hand side of the ODEs to 0,

$$\begin{aligned}
0 &= (P_{\mathrm{env}} + \overline{P_{\mathrm{ext}}})\cdot\frac{\kappa_{\mathrm{P}}}{1+\gamma_{\mathrm{PB}}\cdot B(t)} - \alpha_{\mathrm{P}}\cdot\bar{A}\cdot\bar{P} - \delta_{\mathrm{P}}\cdot\bar{P} \\
0 &= \kappa_{\mathrm{A}}\cdot\frac{1}{1+\gamma_{\mathrm{AB}}\cdot\bar{B}}\cdot G(B^-(R,T),B^+(R,T)) - \delta_{\mathrm{A}}\cdot\overline{\mathrm{A}} \\
0 &= \kappa_{\mathrm{B}}\cdot\frac{1}{1+\gamma_{\mathrm{BB}}\cdot\bar{B}}\cdot\frac{1}{1+\gamma_{\mathrm{BR}}\cdot R(P^-,P^+)}\cdot G(B^-,B^+) - \delta_{\mathrm{B}}\cdot K(P^-,P^+)\cdot\bar{B} - \overline{\delta_{Bext}}\cdot\bar{B} \\
0 &= \kappa_{\mathrm{D}}\cdot R(P^-,P^+) + \overline{D_{\mathrm{ext}}} - \delta_{\mathrm{D}}\cdot\overline{\mathrm{D}}
\end{aligned} \tag{B.2}$$

The third algebraic equation in B.2, corresponding to $\frac{dB}{dt} = 0$, may be rearranged to form a quadratic equation for $\bar{B}$,

$$0 = \gamma_{BB}\cdot\bar{B}^2 + \bar{B} - \frac{\kappa_B\cdot G}{(1+\gamma_{BR}\cdot R)\cdot\left(\delta_B\cdot K + \overline{\delta_{Bext}}\right)} \tag{B.3}$$

Solving this quadratic equation gives us two solutions. Since all of the parameter values are necessarily positive, this quadratic always has positive-valued coefficients for $\bar{B}^2$ and $\bar{B}$, and a negative constant. This means there is always one negative and one positive solution, which we denote $\bar{B}_{\mathrm{neg}}$ and $\bar{B}_{\mathrm{pos}}$ respectively as

$$\bar{B}_{neg} = \frac{-1 - \sqrt{1 + 4 \cdot \gamma_{BB} \cdot \frac{\kappa_B}{\left(\delta_B \cdot K + \overline{\delta_{Bext}}\right)} \cdot \frac{G}{(1 + \gamma_{BR} \cdot R)}}}{2 \cdot \gamma_{BB}} \tag{B.4}$$

$$\bar{B}_{pos} = \frac{-1 + \sqrt{1 + 4 \cdot \gamma_{BB} \cdot \frac{\kappa_B}{\left(\delta_B \cdot K + \overline{\delta_{Bext}}\right)} \cdot \frac{G}{(1 + \gamma_{BR} \cdot R)}}}{2 \cdot \gamma_{BB}} \tag{B.5}$$

Both solutions are always real since their discriminant

$$\left(1 + 4 \cdot \gamma_{BB} \cdot \frac{\kappa_B}{\left(\delta_B \cdot K + \overline{\delta_{Bext}}\right)} \cdot \frac{G}{(1 + \gamma_{BR} \cdot R)}\right) \tag{B.6}$$

is always positive for positive valued parameters.

Clearly, in our case, the negative-valued solution $\bar{B}_{neg}$ does not make physical sense since it isn't possible to have negative barrier integrity, so we consider only $\bar{B}_{pos}$ in our analysis. When we take this positive solution $\bar{B}_{pos}$, the values of $\bar{A}$ and $\bar{P}$ are also guaranteed to be positive. Further, the value of $\overline{D}$ is guaranteed to take a positive value regardless of the values of the other state variables.

We denote these positive focal point values with overbars $\bar{P}$, $\bar{A}$, $\bar{B}$ and $\overline{D}$, expressed in equation B.7.

$$\begin{aligned}
\bar{P} &= \frac{(P_{env} + \overline{P_{ext}}) \cdot \kappa_P}{(\alpha_P \cdot \bar{A} + \delta_P) \cdot (1 + \gamma_{PB} \cdot \bar{B})} \\
\bar{A} &= \frac{\kappa_A}{\delta_A} \cdot \frac{G}{1 + \gamma_{AB} \cdot \bar{B}} \\
\bar{B} &= \frac{-1 + \sqrt{1 + 4 \cdot \gamma_{BB} \cdot \frac{\kappa_B}{\left(\delta_B \cdot K + \overline{\delta_{Bext}}\right)} \cdot \frac{G}{(1 + \gamma_{BR} \cdot R)}}}{2 \cdot \gamma_{BB}} \\
\overline{D} &= \frac{\kappa_D \cdot R + \overline{D_{ext}}}{\delta_D}
\end{aligned} \tag{B.7}$$

We evaluate the stability of the positive-valued equilibrium point by calculating the Jacobian of the system

$$J = \begin{pmatrix}
-\delta_P - A\,\alpha_P & -P\,\alpha_P & -\frac{\gamma_{\mathrm{PB}}\,\kappa_P\,(P_{\mathrm{ex}}+P_{\mathrm{env}})}{(B\,\gamma_{\mathrm{PB}}+1)^2} & 0 \\
0 & -\delta_A & -\frac{G\,\gamma_{\mathrm{AB}}\,\kappa_A}{(B\,\gamma_{\mathrm{AB}}+1)^2} & 0 \\
0 & 0 & -\delta_{\mathrm{Bext}} - K\,\delta_B - \frac{G\,\gamma_{\mathrm{BB}}\,\kappa_B}{(B\,\gamma_{\mathrm{BB}}+1)^2\,(R\,\gamma_{\mathrm{BR}}+1)} & 0 \\
0 & 0 & 0 & -\delta_D
\end{pmatrix} \tag{B.8}$$

Since this Jacobian matrix is upper triangular, its eigenvalues are written on the leading diagonal. Evaluating these eigenvalues at the steady state point $(\bar{P}, \bar{A}, \bar{B}, \overline{D})$, we obtain

$$\begin{aligned}
\lambda_1 &= -\delta_P - \overline{A}\alpha_P & \lambda_2 &= -\delta_A \\
\lambda_3 &= -\delta_{\mathrm{Bext}} - K\delta_B - \frac{G\gamma_{BB}\kappa_B}{(\bar{B}\gamma_{BB} + 1)^2(R\gamma_{BR} + 1)} & \lambda_4 &= -\delta_D
\end{aligned} \tag{B.9}$$

Since all of the model's parameters are positive by assumption, eigenvalues λ_2 and λ_4 are clearly always negative. Further, since we are considering only the positive-valued solution of the steady state equations, both $\bar{B}$ and $\bar{A}$ are positive and so the eigenvalues λ_1 and λ_3 can also only take negative

values. Thus, the positive-valued equilibrium point is always stable, and since it is also the only equilibrium value in the admissible range for our state variables, it is also a local attractor.

We make use of the expression for the focal point value in equation B.7 in the subsequent analysis for the calculation of the boundary-equilibrium bifurcation lines in section 0.

Appendix C No bistability in the TGFβ feedback model using 1st-order cooperative inhibition

The ODE for non-dimerised TGFβ feedback (i.e. with TGFβ-mediated feedback of order-1 rather than order-2 as described in section 3.2.3) is written as C.1.

$$\frac{dT}{d\tau} = D\frac{T}{1+T} + \pi_2\frac{T^2}{1+T} - T \tag{C.1}$$

We usually write this with the common factor T pulled to the front (C.2) to make clear that $T = 0$ is always a steady state.

$$\frac{dT}{d\tau} = T\left(\frac{D}{1+T} + \pi_2\frac{T}{1+T} - 1\right) \tag{C.2}$$

The remaining, nonzero steady state solution is the solution for T of C.3.

$$1 - D\frac{1}{1+T} = \pi_2\frac{T}{1+T} \tag{C.3}$$

This nonzero steady state solution is C.4.

$$T = \frac{D-1}{1-\pi_2} \tag{C.4}$$

So, we have two steady state solutions to the ODE in equation C.1, and expressions for the values of each on.

We now assess the stability of these two solutions for different values of D. We can calculate the Jacobian of the ODE C.1, obtaining C.5.

$$\frac{\partial^2 T}{\partial\tau\partial T} = D\frac{1}{(1+T)^2} + \pi_2\frac{T^2+2T}{(1+T)^2} - 1 \tag{C.5}$$

We can evaluate the Jacobian when $T = 0$ to work out the stability of this steady state. The Jacobian simplifies to C.6.

$$\frac{\partial^2 T}{\partial\tau\partial T} = D - 1 \tag{C.6}$$

So, this steady state $T = 0$ is only stable for values of $D < 1$, where the Jacobian has a negative value.

We can evaluate the Jacobian when $T = \frac{D-1}{1-\pi_2}$, and we obtain C.7.

$$\frac{\partial^2 T}{\partial\tau\partial T} = \frac{(\pi_2 - 1)(D-1)}{D-\pi_2} \tag{C.7}$$

This tells us something rather interesting, which is that setting $k > 1$ can result in the $T = \frac{D-1}{1-\pi_2}$ steady state value becoming either unable or negative.

For $\pi_2 > 1$ and $1 < D < \pi_2$, the Jacobian C.7 is negative, but the focal point takes a negative value, which is a problem because negative population numbers of Th2 cells doesn't make sense.

If we set $\pi_2 > 1$ and $D > \pi_2$ simultaneously, the ODE would have no stable steady-states, and the dynamics would all simply increase without limit. All of the configurations for the system are summarised in Table 8.2.

Table 8.2 Configurations of the steady states of the ODE for various ranges of π_2 and D.

Ranges for π_2 and D (always positive by assumption)		Steady state value $T = 0$	Jacobian at $T = 0$ $D - 1$	Steady state value $T = \frac{D-1}{1-\pi_2}$	Jacobian at $T = \frac{D-1}{1-\pi_2}$ $\frac{(\pi_2 - 1)(D-1)}{D-\pi_2}$	Interpretation
$\pi_2 < 1$	$D < 1$	0	negative	negative	positive	One stable steady state at $T = 0$
$\pi_2 < 1$	$D > 1$	0	positive	positive	negative	One positive stable steady state at $T = \frac{D-1}{1-\pi_2}$
$\pi_2 > 1$	$D < 1$	0	negative	negative	positive	One stable steady state at $T = 0$
$\pi_2 > 1$	$1 < D < \pi_2$	0	positive	negative	negative	One stable steady state at $T = \frac{D-1}{1-\pi_2}$ but this solution is negative, so inadmissible.
$\pi_2 > 1$	$D > \pi_2$	0	positive	negative	positive	No steady states. ODE solutions increase without limit.

When $\pi_2 > 1$, the ODE behaves quite strangely, with no steady-state solutions when $D > \pi_2$, or having a negative steady-state otherwise. Such a badly-behaved ODE is of little use to us in biological modelling as we're primarily interested in observing the long-term behaviours of the system, and we know from biological knowledge that the level of Th2 cells does indeed reach a stable (and positive) equilibrium rather than climbing without limit to infinity. Flying off to infinity would clearly break our founding assumptions anyway. We also demand that the parameters of the ODE are positive valued since negative values would not make physical sense. For this reason, we limit our attention to the case where $0 < \pi_2 < 1$, so that our ODE remains well-behaved and makes physical sense.

Now we can set our attention on the value of D. Given the restriction on the range of π_2, we now require $D > 1$ for the focal point $T = \frac{D-1}{1-\pi_2}$ to take a positive value (negative values are not physically interpretable). Therefore, we restrict our attention to the range $D > 1$ for this steady state.

Taking these two sets of inequalities into account $0 < \pi_2 < 1$ and $D > 1$, we can see that the Jacobian for this steady state $T = \frac{D-1}{1-\pi_2}$ is always negative when the value of the steady state is positive, so this solution is always stable if positive. Any negative values this steady state could take up (which occur when $D < 1$) are unstable because the Jacobian takes a positive value.

Bringing together our knowledge of both steady states, we now know that they exchange stability in a transcritical bifurcation at $D = 1$ (illustrated in Figure 3.3). Crucially for our case, at no point can both steady states stable simultaneously, so it is not possible for the system to display bistability with the ODE in equation C.1.

Appendix D Reparameterization of the Th1/Th2 mutual inhibition model

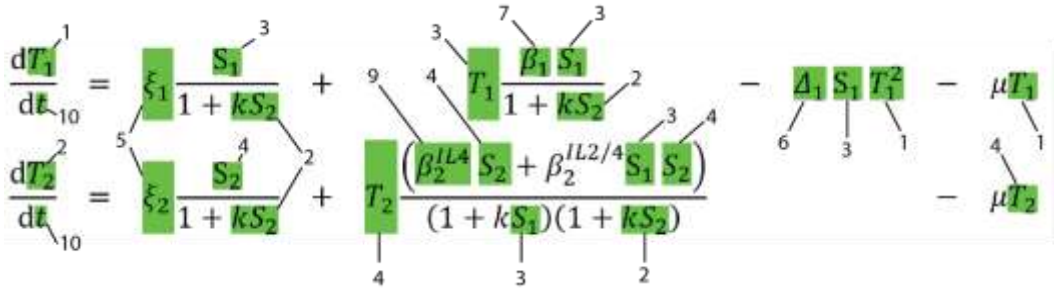
In this appendix, we nondimensionalise the Th1/Th2 mutual inhibition ODE in detailed steps using colour-coding to keep track of substitutions. The ODEs we work with are reproduced below

$$\begin{aligned}\frac{\mathrm{d}T_1}{\mathrm{d}t} &= \xi_1\frac{\mathrm{S}_1}{1+kS_2} + T_1\frac{\beta_1 S_1}{1+kS_2} - \Delta_1 \mathrm{S}_1 T_1^2 - \mu T_1\\ \frac{\mathrm{d}T_2}{\mathrm{d}t} &= \xi_2\frac{\mathrm{S}_2}{1+kS_2} + T_2\frac{\left(\beta_2^{IL4}S_2 + \beta_2^{IL2/4}\mathrm{S}_1S_2\right)}{(1+kS_1)(1+kS_2)} - \mu T_2\end{aligned} \tag{D.1}$$

We use the following scaling factors, numbered to help us keep track of where they're going.

①	②	③	④	⑤
$x_1 = k\alpha_1 T_1$	$x_2 = kS_2$	$S_1 = \frac{\alpha_1 T_1}{1+kS_2}$	$S_2 = \alpha_2 T_2$	$D_i = \xi_i\alpha_i/\mu$
⑥	⑦	⑧	⑨	⑩
$\delta_1 = \frac{\Delta_1}{\mu k^2\alpha_1}$	$\pi_1 = \beta_1/(\mu k)$	$\pi_2 = \beta_2^{IL2/4}/(\mu k^2)$	$\rho = \beta_2^{IL4}k/\beta_2^{IL2/4}$	$\tau = \mu t$

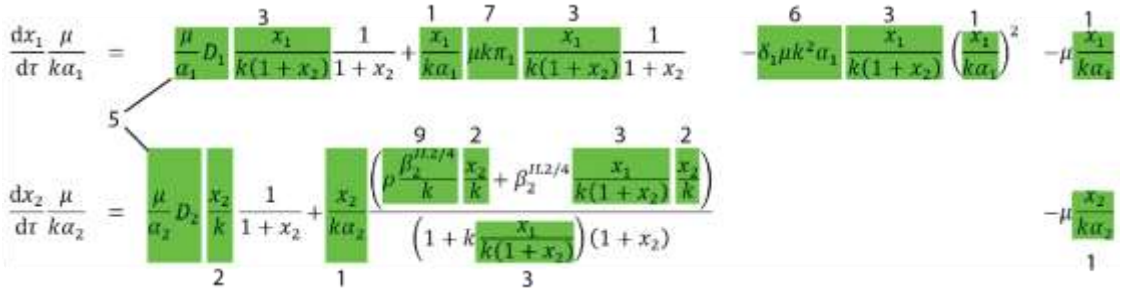
In equation D.2, we write this same ODE, but have added numbered labels to each term we rescale.



(D.2)

It is useful to note the following rearrangements of the scaling factors to help explain the substitutions made in this working. Where the substitution is not immediately obvious, we have labelled the appropriate term in D.3 with the following rearranged scaling factors.

①	②	③ combined with ① and ②		
$T_i = \frac{x_i}{k\alpha_i}$	$S_2 = \frac{x_2}{k}$	$S_1 = \frac{x_1}{k(1+x_2)}$		
⑤	⑥	⑦	⑧	⑨
$\frac{\mu}{\alpha_i}D_i = \xi_i$	$\Delta_1 = \delta_1\mu k^2\alpha_1$	$\beta_1 = \mu k\pi_1$	$\frac{\beta_2^{IL2/4}}{k} = \pi_2\mu k$	$\beta_2^{IL4} = \rho\frac{\beta_2^{IL2/4}}{k}$



(D.3)

with some further rearranging, we can expose the remaining factor $\frac{\beta_2^{IL2/4}}{k}$ in the x_2 ODE in D.4, using ⑧ to replace this factor with $\pi_2\mu k$ as shown in D.4 and D.5.

$$\begin{aligned}
\frac{\mathrm{d}x_1}{\mathrm{d}\tau}\frac{\mu}{k\alpha_1} &= \frac{\mu}{k\alpha_1}D_1\frac{x_1}{(1+x_2)^2} + \frac{\mu}{k\alpha_1}\pi_1\frac{x_1^2}{(1+x_2)^2} - \frac{\mu}{k\alpha_1}\delta_1\frac{x_1^3}{1+x_2} - \frac{\mu}{k\alpha_1}x_1, \\
\frac{\mathrm{d}x_2}{\mathrm{d}\tau}\frac{\mu}{k\alpha_2} &= \frac{\mu}{k\alpha_2}D_2\frac{x_2}{1+x_2} + \frac{x_2}{k\alpha_2}\frac{\left(\rho\frac{\beta_2^{IL2/4}}{k}\frac{x_2}{k} + \frac{\beta_2^{IL2/4}}{k}\frac{x_1}{(1+x_2)}\frac{x_2}{k}\right)}{\left(1+\frac{x_1}{1+x_2}\right)(1+x_2)} - \frac{\mu}{k\alpha_2}x_2
\end{aligned} \tag{D.4}$$

Substituting the scaling term ⑧, $\frac{\beta_2^{IL2/4}}{k} = \pi_2\mu k$, and grouping factors of x_2 we obtain D.5.

$$\begin{aligned}
\frac{\mathrm{d}x_1}{\mathrm{d}\tau}\frac{\mu}{k\alpha_1} &= \frac{\mu}{k\alpha_1}D_1\frac{x_1}{(1+x_2)^2} + \frac{\mu}{k\alpha_1}\pi_1\frac{x_1^2}{(1+x_2)^2} - \frac{\mu}{k\alpha_1}\delta_1\frac{x_1^3}{1+x_2} - \frac{\mu}{k\alpha_1}x_1, \\
\frac{\mathrm{d}x_2}{\mathrm{d}\tau}\frac{\mu}{k\alpha_2} &= \frac{\mu}{k\alpha_2}D_2\frac{x_2}{1+x_2} + \frac{1}{k\alpha_2}\frac{\pi_2\mu k\left(\rho\frac{x_2^2}{k} + \frac{x_1}{(1+x_2)}\frac{x_2^2}{k}\right)}{1+x_1+x_2} - \frac{\mu}{k\alpha_2}x_2
\end{aligned} \tag{D.5}$$

Finally, we can simplify the ODEs to the form used in section 3.2.4. We finally obtain D.6, which is what we have written in equation 3.9.

$$\begin{aligned}
\frac{\mathrm{d}x_1}{\mathrm{d}\tau} &= D_1\frac{x_1}{(1+x_2)^2} + \pi_1\frac{x_1^2}{(1+x_2)^2} - \delta_1\frac{x_1^3}{1+x_2} - x_1, \\
\frac{\mathrm{d}x_2}{\mathrm{d}\tau} &= D_2\frac{x_2}{1+x_2} + \pi_2\left(\rho + \frac{x_1}{1+x_2}\right)\frac{x_2^2}{1+x_1+x_2} - x_2,
\end{aligned} \tag{D.6}$$

Appendix E Calculating threshold values as a function of dupilumab concentration for the systemic IL4 feedback+Th1/Th2 mutual inhibition model

For the systemic IL4 feedback + Th1/Th2 mutual inhibition model discussed in section 3.2.4, the Th1/Th2 bistable region lies between the Th1-only focal point $(x_1^*, 0)$ and the mixed focal point (x_1^{**}, x_2^{**}) (yellow region of Figure 3.8 (b)). The upper threshold of the Th1/Th2 bistable region, D^+ (shown in Figure 3.8 (c)), occurs where the focal point $(x_1^*, 0)$ (pink region of Figure 3.8 (d)) becomes unstable. We can evaluate the value of D when this steady state becomes unstable by calculating the Jacobian of the ODE, from equation 3.9 (reproduced here as equation E.1 for ease of reading) and evaluating it at this focal point $(x_1^*, 0)$. In equation E.1, we substitute $D_1 = CD_2$ where we assume $C > 1$ such that we're in the $\frac{D_1}{D_2} > 1$ mode of behaviour.

$$\begin{aligned} \frac{dx_1}{d\tau} &= CD_2\frac{x_1}{(1+x_2)^2} + \pi_1\frac{x_1^2}{(1+x_2)^2} - \delta_1\frac{x_1^3}{1+x_2} - x_1, \\ \frac{dx_2}{d\tau} &= \frac{D_2}{1+[d]}\frac{x_2}{1+x_2} + \frac{\pi_2}{1+[d]}\left(\rho + \frac{x_1}{1+x_2}\right)\frac{x_2^2}{1+x_1+x_2} - x_{2,} \end{aligned} \tag{E.1}$$

Evaluating the Jacobian with at the "Th1 only" focal point $(x_1^*, 0)$, where $x_1 > 0$ and $x_2 = 0$, we obtain E.2.

$$J = \begin{pmatrix} CD_2 - 1 & 0 \\ \delta_1 x_1 - 2\pi x_1^2 - 2CD_2 x_1 & \frac{D_2}{1+[d]} - 1 \end{pmatrix} \tag{E.2}$$

Since the Jacobian is lower triangular, the eigenvalues of this Jacobian are just the diagonal elements in E.3 and E.4.

$$CD_2 - 1 \tag{E.3}$$

$$\frac{D_2}{1+[d]} - 1 \tag{E.4}$$

For this focal point to be stable, both eigenvalues must be negative so that we can write two corresponding inequalities in E.5 and E.6.

$$D_2 < \frac{1}{C} \equiv D_1 < 1 \tag{E.5}$$

$$D_2 < 1 + [d] \tag{E.6}$$

The first inequality E.5 simply states that $D_1 < 1$ is necessary (which means the yellow "Th1 only response" region in Figure 3.8 (b) has a boundary at $D_1 = 1$).

The other inequality E.6 tells us that when the external stimulation D_2 exceeds the value $1 + [d]$ the focal point with no Th2 response becomes unstable, and so the Th2 response switches on. Note the dependence on the dupilumab level $[d]$. This tells us that increasing the dupilumab concentration has a linear effect on the threshold value $D^+ = 1 + [d]$.

The lower threshold of the bistable region, D^-, is harder to calculate because we have no closed-form expression for the focal point when Th2 responses are activated. Therefore, we use numerical continuation techniques in MatCont to calculate the dependence of D^- on dupilumab concentration $[d]$. We continued from the limit point for D^-, allowing both the D_2 and $[d]$ parameters to vary simultaneously (since this is a codimension 2 bifurcation), which allows us to see how this threshold value moves through that plane. By this process, we obtain the red dashed line in Figure 3.15 (c).

Appendix F Calculating threshold values as a function of dupilumab concentration for the systemic IL4 feedback model

For the systemic IL4 feedback model with dimerization of TGFβ discussed in section 3.2.3, equation (3.7) reproduced below as F.1, the upper threshold D^+ (shown in Figure 3.15) occurs where the focal point $T = 0$ becomes unstable, making the system converge to the upper focal point $T = \frac{1-D}{\pi-1}$ if D is increased above D^+. We can evaluate this threshold value $D = D^+$ when this steady state becomes unstable by calculating the Jacobian of the ODE (reproduced below as equation F.2) and evaluating it at the focal point $T = 0$. In this way, we can deduce when the "no Th2 response" focal point becomes unstable and gives way to the "Th2 response" focal point.

Recall that the lowercase d represents the concentration of dupilumab, while the uppercase D represents the number of dendritic cells carrying an antigen.

$$\frac{dT}{dt} = \frac{D}{1+[d]}\frac{T}{1+T} + \frac{\pi}{1+[d]}\frac{T^2}{1+T^2} - T \tag{F.1}$$

The Jacobian, after partially differentiating with respect to T is equation F.2.

$$\frac{\partial^2 T}{\partial t \partial T} = \frac{D}{1+[d]}\frac{1}{(1+T^2)} + \frac{\pi}{1+[d]}\frac{2T}{(1+T^2)^2} - 1 \tag{F.2}$$

The Jacobian must be negative $\frac{\partial^2 T}{\partial t \partial T} < 0$ for the steady state $T = 0$ to be stable. Evaluating this Jacobian at $T = 0$, we obtain the expression F.3.

$$\frac{D}{1+[d]} - 1 < 0 \tag{F.3}$$

Therefore, the threshold value D^+ occurs at the value $D = 1 + [d]$, which causes the Jacobian to become positive, so we can write F.4.

$$D^+ = 1 + [d] \tag{F.4}$$

The lower threshold of the bistable region, D^-, is harder to calculate because we do not have a closed-form expression for the focal point when Th2 responses are activated. Therefore, we use numerical continuation methods using MatCont to calculate the dependence of D^- on dupilumab concentration $[d]$. To do this, we select the lower "limit point" on Figure 3.10, and carry out the continuation, allowing the parameters D and $[d]$ to vary (limit points are codimension-2 objects so two parameters may vary). By this process, we obtain the red dashed line in Figure 3.15 (a).

Appendix G Derivations of the boundary values for the stability regions in Figure 3.8

This appendix is concerned with deriving the regions of the D_1, D_2 plane in which the various steady state values of the "systemic IL4 feedback + Th1/Th2 mutual inhibition model" are stable. This model is defined by the ODEs in equation 3.9, but to save you flicking back and forth between pages too much, we reproduce the ODE in G.1 below.

$$\begin{aligned}\frac{\mathrm{d}x_1}{\mathrm{d}\tau} &= D_1\frac{x_1}{(1+x_2)^2}+\pi_1\frac{x_1^2}{(1+x_2)^2}-\delta\frac{x_1^3}{1+x_2}-x_1,\\ \frac{\mathrm{d}x_2}{\mathrm{d}\tau} &= D_2\frac{x_2}{1+x_2}+\pi_2\left(\rho+\frac{x_1}{1+x_2}\right)\frac{x_2^2}{1+x_1+x_2}-x_2,\end{aligned} \tag{G.1}$$

There are four steady state solutions to this system of ODEs, and we address each one in turn below. We make use of the Jacobian in each case, evaluating it at each steady state solution we tackle. It may be useful to see the Jacobian before evaluation as in G.2.

$$J=\begin{pmatrix}\dfrac{\partial^2 x_1}{\partial\tau\partial x_1} & \dfrac{\partial^2 x_1}{\partial\tau\partial x_2}\\ \dfrac{\partial^2 x_2}{\partial\tau\partial x_1} & \dfrac{\partial^2 x_2}{\partial\tau\partial x_2}\end{pmatrix} \tag{G.2}$$

$$\frac{\partial^2 x_1}{\partial\tau\partial x_1}=\frac{D_1}{(1+x_2)^2}+\frac{2\pi_1 x_1}{(1+x_2)^2}-\frac{3\delta x_1^2}{1+x_2}-1$$

$$\frac{\partial^2 x_1}{\partial\tau\partial x_2}=\frac{\delta_1 x_1^3}{(1+x_2)^2}-\frac{2D_1 x_1}{(1+x_2)^2}-\frac{2\pi x_1^2}{(1+x_2)^3}$$

$$\frac{\partial^2 x_2}{\partial\tau\partial x_1}=\frac{\pi_2 x_2^2}{(1+x_2)(1+x_1+x_2)}-\frac{\pi_2 x_2^2\left(\rho+\frac{x_1}{1+x_2}\right)}{(1+x_1+x_2)^2}$$

$$\frac{\partial^2 x_2}{\partial\tau\partial x_2}=\frac{D_2}{1+x_2}-\frac{D_2 x_2}{(1+x_2)^2}-\frac{\pi_2 x_2^2\left(\rho+\frac{x_1}{1+x_2}\right)}{(1+x_1+x_2)^2}+\frac{2\pi_2 x_2\left(\rho+\frac{x_1}{1+x_2}\right)}{1+x_1+x_2}-\frac{\pi_2 x_1 x_2^2}{(1+x_2)^2(1+x_1+x_2)}-1$$

Case 1. Completely off

Focal point $(\overline{x_1},\overline{x_2})=(0,0)$ for which the local Jacobian is diagonal and has eigenvalues G.3 and G.4.

$$\lambda_1=1-D_1 \tag{G.3}$$

$$\lambda_2=1-D_2 \tag{G.4}$$

The steady state $(0,0)$ is stable when these are negative, so when $D_1<1$ and $D_2<1$ (blue region "Off" in Figure 3.8 (a)).

Case 2. Th1 response only

Focal point at $(D_1,D_2)=(x_1^*,0)$, where x_1^* is the non-zero solution of the ODE for $\frac{dx_1}{d\tau}$ being equal to 0 when $x_2=0$. The ODE for $\frac{dx_1}{d\tau}$ forms the cubic equation G.5.

$$D_1x_1+\pi_1x_1^2-\delta x_1^3-x_1=0 \tag{G.5}$$

This cubic has solutions given by equation G.6 and $x_1^*=0$, the latter of which we can safely ignore since it is identical to 0 above.

$$x_1^* = \frac{\pi_1 \pm \sqrt{\pi_1^2 + 4\delta(D_1 - 1)}}{2\delta} \tag{G.6}$$

The larger of the two values G.6 can take is always positive while the other value is always negative, and we can safely neglect it. The positive solution is also real if the discriminant of G.6 is positive. This condition is written in G.7, rearranged for D_1

$$D_1 > 1 - \frac{\pi_1^2}{4\delta} \tag{G.7}$$

Evaluating the Jacobian with at the "Th1 only" focal point $(x_1^*, 0)$, where $x_1 > 0$ and $x_2 = 0$, we obtain G.8.

$$J = \begin{pmatrix} D_1 + 2\pi_1 x_1^* - 3\delta {x_1^*}^2 - 1 & \delta_1 {x_1^*}^3 - 2\pi_1 {x_1^*}^2 - 2D_1 x_1^* \\ 0 & D_2 - 1 \end{pmatrix} \tag{G.8}$$

Since this Jacobian is upper triangular, the eigenvalues at this positive focal point for x_1^* are on the diagonal as G.9 and G.10.

$$\lambda_1 = -2(D_1 - 1) - \frac{\pi_1\left(\pi_1 + \sqrt{\pi_1^2 + 4\delta(D_1 - 1)}\right)}{2\delta} \tag{G.9}$$

$$\lambda_2 = D_2 - 1 \tag{G.10}$$

The focal point $(x_1^*, 0)$ is stable only when both of its corresponding eigenvalues λ_1 and λ_2 are negative, as expressed in G.11 and G.12.

$$D_1 > 1 + \frac{\pi_1\left(\pi_1 + \sqrt{\pi_1^2 + 4\delta(D_1 - 1)}\right)}{4\delta} \tag{G.11}$$

$$D_2 < 1 \tag{G.12}$$

These two conditions define the extent of the yellow region "Th1 only" in Figure 3.8 (b)).

Case 3. Th2 response only

Focal point $(\overline{x_1}, \overline{x_2}) = (0, x_2^*)$, where x_2^* is the non-zero steady state solution of the ODE for x_2 when we set $x_1 = 0$.

$$\frac{\mathrm{d}x_2}{\mathrm{d}\tau} \quad = \quad D_2 x_2 + \pi_2 \rho x_2^2 - x_2$$

When $\frac{\mathrm{d}x_2}{\mathrm{d}\tau} = 0$, this has solution in equation G.13 and $x_2^* = 0$.

$$x_2^* = \frac{D_2 - 1}{1 - \pi_2 \rho} \tag{G.13}$$

We can safely ignore the $x_2^* = 0$ as this is just 0 above. This Jacobian evaluated at the G.13 equilibrium value becomes G.14.

$$J = \begin{pmatrix} -1 + D_1\left(\frac{1 - \rho\pi_2}{D_2 - \rho\pi_2}\right)^2 & 0 \\ \frac{\pi_2(D_2 - 1)^2(1 - \rho)}{(D_2 - \pi_2\rho)^2} & \frac{(D_2 - 1)(\rho\pi_2 - 1)}{D_2 - \rho\pi_2} \end{pmatrix} \tag{G.14}$$

This Jacobian is lower triangular, so the eigenvalues are on the leading diagonal, see G.15 and G.16.

$$\lambda_1 = -1 + D_1\left(\frac{1 - \rho\pi_2}{D_2 - \rho\pi_2}\right)^2 \tag{G.15}$$

$$\lambda_2 = \frac{(D_2 - 1)(\rho\pi_2 - 1)}{D_2 - \rho\pi_2} \tag{G.16}$$

Note that we only consider the case where $\rho\pi_2 < 1$, which causes the term $(\rho\pi_2 - 1)$ in G.15 to be negative. If we set $\rho\pi_2 \geq 1$, the Th2 cell number x_2^* has no stable positive steady states. This is biologically implausible since it would imply that immune responses composed solely of Th2 cells are impossible, however since patients with a Th2 response and no accompanying Th1 response are observed, we disregard this case.

The focal point G.13 is stable only if both eigenvalues G.15 and G.16 are negative. This requirement is summarised by conditions G.17 and G.18.

$$0 > \frac{(D_2 - 1)(\rho\pi_2 - 1)}{D_2 - \rho\pi_2} \tag{G.17}$$

$$0 > -1 + D_1\left(\frac{1 - \rho\pi_2}{D_2 - \rho\pi_2}\right)^2 \tag{G.18}$$

Given that we assume $\rho\pi_2 < 1$, and that D_2 must take a positive value, condition G.17 simplifies to G.19. Meanwhile, condition G.18 rearranges to G.20.

$$D_2 > 1 \tag{G.19}$$

$$D_1 < \left(\frac{D_2 - \rho\pi_2}{1 - \rho\pi_2}\right)^2 \tag{G.20}$$

Taken together, G.19 and G.20 define the extent of the "Th2 only" orange region in Figure 3.8 (c)). The quadratic form of G.20 defines the curved edge of this region.

Case 4. Mixed response, consisting non-zero numbers of both Th1 and Th2 cells

We finally consider the equilibrium point $(\overline{x_1}, \overline{x_2}) = (x_1^{**}, x_2^{**})$, marked by the green region in Figure 3.8 (d). At this equilibrium point, it becomes impossible to calculate closed-form expressions for the values x_1^{**} and x_2^{**}. This precludes the possibility of obtaining analytical expressions for the stability by evaluating the Jacobian around that fixed point.

With the analytical technique no longer an option, we resort to numerical continuation methods to work out what the border of this region looks like in the D_1, D_2 plane. The boundary of this region, which is the threshold value D^-, is determined numerically using MatCont. The resulting line of points provides us with the purple edge of the pink region in Figure 3.8 (d).

To identify the boundary of the region with MatCont, I selected the threshold value which defined the edge of this region, and carried out a numerical continuation, allowing the parameters D_1 and D_2 to vary. In MatCont, this threshold value appears as a "Limit point" and is identified as a "codimension-2" object, meaning that two parameters may be varied simultaneously while we track the location of the object in the state space.

Appendix H Derivation of T-switch behaviours

In this appendix, we generalise an analysis already presented in Domínguez-Hüttinger et al. [18], which only needed to consider the switching on of the T-switch because the T-switch was considered irreversible. Here, we use the same technique, but also incorporate the switching off of the new reversible T-switch in response to periodic input from the R-switch.

To determine the behaviour of the T-switch when the stimulation of innate immune receptors is intermittent, we apply a square save input $R(t)$ to the ODE and evaluate its solution $D(t)$. This solution $D(t)$, which represents the number of dendritic cells activating T-cells, causes the T-switch to turn on or off by crossing the thresholds D^- and D^+. This analysis extends a that found in [18], where this same problem was considered, but for an irreversible T-switch, which could only turn on, but not off. Here, we extend the analysis to consider the conditions under which the T-switch can turn off again once activated, allowing oscillating behaviour to be induced.

We begin by defining a squire wave input from the R-switch. The R-switch changes between the values R_{off} and R_{on}.We denote the rising edges of this square wave with capital letters $T_0, T_1, \ldots$ and the falling edges with lower-case letters $t_0, t_1, \ldots$ as shown in figure H.1.

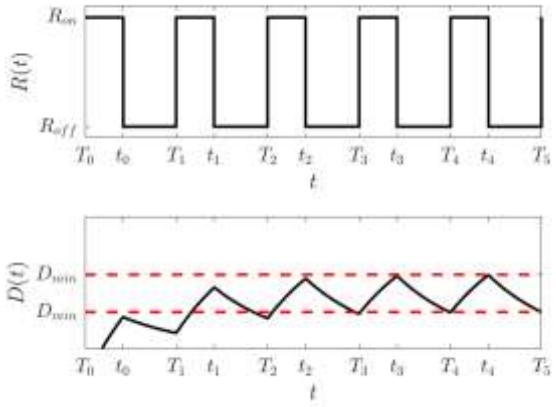


Figure H.1 Solution D(t) in response to square wave forcing function R(t). D_{min} is the minimum value of the solution D(t) after many periods of the square wave. D_{max} is the maximum.

The duty cycle of this square wave is

$$\begin{aligned} t_{on} &= t_k - T_k \\ t_{\text{off}} &= T_{k+1} - t_k \end{aligned} \tag{H.1}$$

Where $k = 0,1,2 \ldots$

The levels of the square wave are

$$R(t) = \begin{cases} R_{\text{on}} & \text{if} \quad T_k < t < t_k \\ R_{\text{off}} & \text{if} \quad T_{k+1} < t < t_k \end{cases} \tag{H.2}$$

The number of dendritic cells $D(t)$ is defined by the ODE in equation H.3 (which is just the D ODE from the full ODE system for the triple-switch model defined in equation 5.1). In this ODE, the population increases at rate κ_D when the innate immune receptors $R(t)$ are active, or superantigens stimulate mast cells directly $D_{\text{ext}}(t)$ and die at rate δ_D.

$$\frac{dD(t)}{dt} = \kappa_{\text{D}} \cdot (R(t) + D_{\text{ext}}(t)) - \delta_{\text{D}} \cdot D(t) \tag{H.3}$$

For simplicity, we assume the external perturbation external perturbation of DC number by superantigens $D_{\text{ext}}(t)$ to be constant over time $D_{\text{ext}}(t) = \overline{D_{\text{ext}}}$. Our use of this forcing function in this thesis is limited to merely considering the potential impacts in section 8.5.3 rather than a full mathematical analysis here.

The solution to the ODE can be expressed as

$$D(t) = e^{-\delta_D(t-t_0)}D(T_0) + \int_{T_0}^{t} e^{-\delta_D(t-\tau)}\,\kappa_D(R(\tau) + \overline{D_{\text{ext}}})d\tau \tag{H.4}$$

Where T_0 is the time at which we specify the initial number of dendritic cells $D(T_0)$.
Thus, the value of $D(t_0)$ after the first t_{on} duration (time between T_0 and t_0) is given by

$$\begin{aligned} D(t_0) &= D(T_0)\cdot e^{-\delta_D(t_1-T_0)} + \int_{T_0}^{t_0} \kappa_D(R_{\text{on}} + \overline{D_{\text{ext}}})e^{-\delta_D(t_1-\tau)}d\tau \\ D(t_0) &= D(T_0)\cdot e^{-\delta_D(t_0-T_0)} + \frac{\kappa_D}{\delta_D}(R_{\text{on}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D(t_0-T_0)}\right) \end{aligned} \tag{H.5}$$

And the value of $D(t)$ at the beginning of the next period of the square wave can be calculated by calculating this same integral solution again but setting the initial condition to take the final value $D(t_0)$ we just calculated, obtaining

$$\begin{aligned} D(T_1) &= D(t_0)\cdot e^{-\delta_D(T_1-t_0)} + \int_{t_0}^{T_1} \kappa_D(R_{\text{off}} + \overline{D_{\text{ext}}})e^{-\delta_D(T_1-\tau)}d\tau \\ D(T_1) &= D(t_0)\cdot e^{-\delta_D(T_1-t_0)} + \frac{\kappa_D}{\delta_D}(R_{\text{off}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D(T_1-t_0)}\right) \end{aligned} \tag{H.6}$$

Substituting $t_0 - T_0 = t_{\text{on}}$ and $T_1 - t_0 = t_{\text{off}}$, we can generalise these solutions for any $D(t_k)$ or $D(T_k)$, obtaining

$$\begin{aligned} D(t_k) &= D(T_k)\cdot e^{-\delta_D t_{\text{on}}} + \frac{\kappa_D}{\delta_D}(R_{\text{on}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D t_{\text{on}}}\right) \\ D(T_k) &= D(t_{k-1})\cdot e^{-\delta_D t_{\text{off}}} + \frac{\kappa_D}{\delta_D}(R_{\text{off}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D t_{\text{off}}}\right) \end{aligned} \tag{H.7}$$

Combining the equations above, we can write $D(T_k)$ as

$$D(T_k) = D(T_{k-1})\cdot e^{-\delta_D T} + \frac{\kappa_D}{\delta_D}\left((R_{\text{on}} + \overline{D_{\text{ext}}})\left(e^{-\delta_D t_{\text{off}}} - e^{-\delta_D T}\right) + (R_{\text{off}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D t_{\text{off}}}\right)\right) \tag{H.8}$$

This recurrence relation takes the form

$$a_k = r a_{k-1} + b \tag{H.9}$$

Where $k = 0,1,2\ldots$ is the index of the pulse in the sequence. We can write the terms of this recurrence relation for any value k in the sequence as

$$a_k = r^k a_0 + b\frac{1 - r^k}{1 - r} \tag{H.10}$$

We can now write an expression for the value $D(T_k)$ for any value $k = 0,1,2\ldots$ given the initial value $D(T_0)$ as

$$\begin{aligned} D(T_k) = e^{-\delta_D kT}D(T_0) &+ \frac{\kappa_D}{\delta_D}\Big((R_{\text{on}} + \overline{D_{\text{ext}}})\left(e^{-\delta_D t_{\text{off}}} - e^{-\delta_D T}\right) \\ &+ (R_{\text{off}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D t_{\text{off}}}\right)\Big)\frac{1 - e^{-\delta_D kT}}{1 - e^{-\delta_D T}} \end{aligned} \tag{H.11}$$

From this sequence, we can calculate the maximum values $D(t_k)$ from the $D(T_k)$ for $k = 1,2,\ldots$ using equation (H.7). We obtain a similar expression for $D(t_k)$ for any value $k = 0,1,2\ldots$. Note this looks very similar to (H.11), but the t_{off} and t_{on} have swapped places.

$$\begin{aligned} D(t_k) = e^{-\delta_D kT}D(T_0) &+ \frac{\kappa_D}{\delta_D}\Big((R_{\text{on}} + \overline{D_{\text{ext}}})\left(e^{-\delta_D t_{\text{on}}} - e^{-\delta_D T}\right) \\ &+ (R_{\text{off}} + \overline{D_{\text{ext}}})\left(1 - e^{-\delta_D t_{\text{on}}}\right)\Big)\frac{1 - e^{-\delta_D kT}}{1 - e^{-\delta_D T}} \end{aligned} \tag{H.12}$$

The structure of solutions (H.11) and (H.12) show us that the maximum and minimum values of the solution for D converge towards steady state values at an exponential rate, with each cycle of the R-switch reaching $e^{-\delta_D T}$ closer to the steady state value. We can evaluate these steady state minimum and maximum values from (H.11) and (H.12) respectively after many square wave cycles have passed by taking the limit as $k \to \infty$.

$$\begin{aligned} D_{\max} &= \lim_{k\to\infty} D(t_k) \\ D_{\min} &= \lim_{k\to\infty} D(T_k) \end{aligned} \tag{H.13}$$

Evaluating this limit, we obtain

$$\begin{aligned} D_{\max} &= \frac{\kappa_D}{\delta_D}\Big((R_{\text{off}} + \overline{D_{\text{ext}}})\big(e^{-\delta_D t_{\text{off}}} - e^{-\delta_D T}\big) + (R_{\text{on}} + \overline{D_{\text{ext}}})\big(1 - e^{-\delta_D t_{\text{off}}}\big)\Big)\frac{1}{1 - e^{-\delta_D T}} \\ D_{\min} &= \frac{\kappa_D}{\delta_D}\Big((R_{\text{on}} + \overline{D_{\text{ext}}})\big(e^{-\delta_D t_{\text{on}}} - e^{-\delta_D T}\big) + (R_{\text{off}} + \overline{D_{\text{ext}}})\big(1 - e^{-\delta_D t_{\text{on}}}\big)\Big)\frac{1}{1 - e^{-\delta_D T}} \end{aligned} \tag{H.14}$$

We are now in a position to write the inequalities describing when the T-switch must change state when the R-switch is oscillating. When the maximum value of the steady state oscillation is above the switching on threshold $D_{\max} > D^+$ we know the T-switch must turn on. Conversely, when the minimum value of the long-term oscillation lies below the switching off threshold $D_{\min} < D^-$ the T-switch must turn off. Thus, we can take the result from equation H.14 to write our transition inequalities as

$$\begin{aligned} &\text{Turn off} && \frac{\kappa_D}{\delta_D}\Big((R_{\text{on}} + \overline{D_{\text{ext}}})\big(e^{-\delta_D t_{\text{on}}} - e^{-\delta_D T}\big) + (R_{\text{off}} + \overline{D_{\text{ext}}})\big(1 - e^{-\delta_D t_{\text{on}}}\big)\Big)\frac{1}{1 - e^{-\delta_D T}} < D^- \\ &\text{Turn on} && \frac{\kappa_D}{\delta_D}\Big((R_{\text{off}} + \overline{D_{\text{ext}}})\big(e^{-\delta_D t_{\text{off}}} - e^{-\delta_D T}\big) + (R_{\text{on}} + \overline{D_{\text{ext}}})\big(1 - e^{-\delta_D t_{\text{off}}}\big)\Big)\frac{1}{1 - e^{-\delta_D T}} > D^+ \end{aligned} \tag{H.15}$$

The T-switch, like the other switches, can exhibit four possible behaviours: recovery, bistable, oscillation and damage, depending on which of the transition inequalities are satisfied. In this periodic R-switch case, we consider the extremes $D_{\min}$ and $D_{\max}$, so our transition inequalities become $D_{\max} > D^+$ to turn on the T-switch, and $D_{\min} < D^-$ to turn off. The four possible induced behaviours are summarised in table H.1.

Table H.1 Four possible behaviours emerge from the two conditions.

		Turn off condition $D_{\min} < D^-$	
		Satisfied	**No satisfied**
Turn on condition $D_{\max} > D^+$	**Satisfied**	Oscillation at the same frequency as the R-switch	Disease (switch remains on)
	Not satisfied	Recovery (switch remains off)	Bistable

Rearranging the inequalities in H.15, we can calculate the conditions on the duty cycle $t_{\text{on}}, t_{\text{off}}$ to turn on or turn off the T-switch. To turn on the T-switch, we require $D_{\max} > D^+$, which yields condition H.16.

$$t_{\text{off}} < \frac{1}{\delta_D} ln\left\{\frac{e^{-\delta_D t_{\text{on}}}\left(D^+ - \frac{\kappa_D}{\delta_D}(R_{\text{off}} + \overline{D_{\text{ext}}})\right) + \frac{\kappa_D}{\delta_D}(R_{\text{off}} - R_{\text{on}})}{D^+ - \frac{\kappa_D}{\delta_D}(R_{\text{on}} + \overline{D_{\text{ext}}})}\right\} \tag{H.16}$$

Conversely, to turn off the T-switch, we require $D_{\min} < D^-$ which yields the similar condition H.17.

$$t_{\text{on}} < \frac{1}{\delta_D} ln \left\{ \frac{e^{-\delta_D t_{\text{off}}} \left(D^- - \frac{\kappa_D}{\delta_D} (R_{\text{on}} + \overline{D_{\text{ext}}}) \right) + \frac{\kappa_D}{\delta_D} (R_{\text{on}} - R_{\text{off}})}{D^- - \frac{\kappa_D}{\delta_D} (R_{\text{off}} + \overline{D_{\text{ext}}})} \right\} \quad \text{(H.17)}$$

Given these inequalities H.16 and H.17, we can deduce that the T-switch will respond in one of four ways depending on the parameters of the R-switch function; t_{off}, t_{on}, R_{off} and R_{on}. Parameter regions corresponding to these four cases are shown in figure H.2, which we featured in our discussion in section 5.2.2.2 of this thesis.

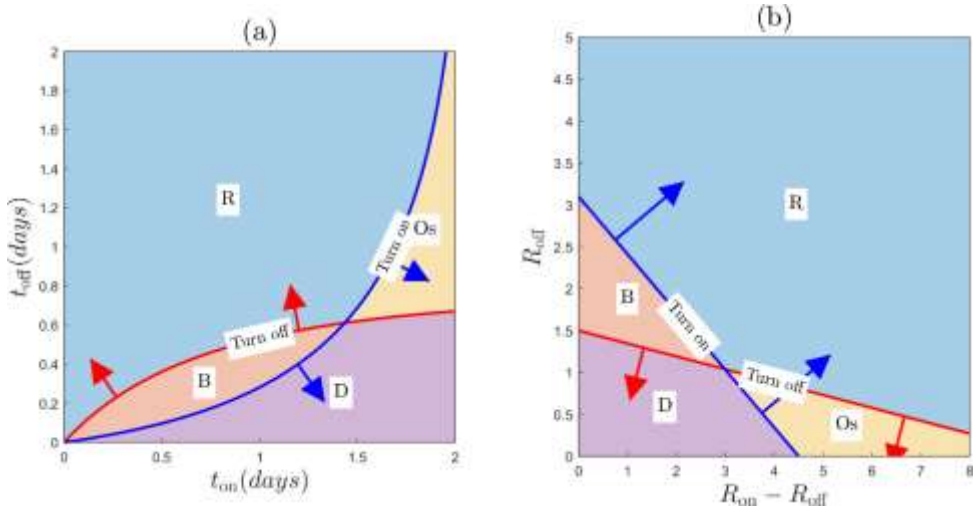


Figure H.2 Distribution of four distinct behaviours of the T-switch as we vary the duty cycle and amplitude of the R-switch.
(a) *t_{off} and t_{on} or* ***(b)*** *R_{off} and $R_{on} - R_{off}$ (chosen so that all points on the diagram satisfy $R_{on} - R_{off} > 0$). Parameters (unless varied on axes): $\kappa_D = 1, \delta_D = 1, t_{off} = 0.5$, $t_{on} = 0.5, R_{off} = 1, R_{on} = 5, D^- = 2, D^+ = 3.1$*

Finally, we visualise this family of periodic solutions in figure H.3 (overleaf). The regions of the figures where the background is white show where the T-switch is off, while the regions where the background is shaded pink show where the T-switch is on.

Note the induced behaviour is not dependent on the initial condition chosen for the $D(t)$ ODE, but in the case of the bistable behaviour, the initial condition does determine the final state (note state not behaviour) of the T-switch in the long-term. When the bistable behaviour is induced, and the initial condition $D(0)$ is set such that the T-switch is initially off, the T-switch remains off. Conversely, if $D(0)$ is selected such that the T-switch is initially on, the T-switch remain on.

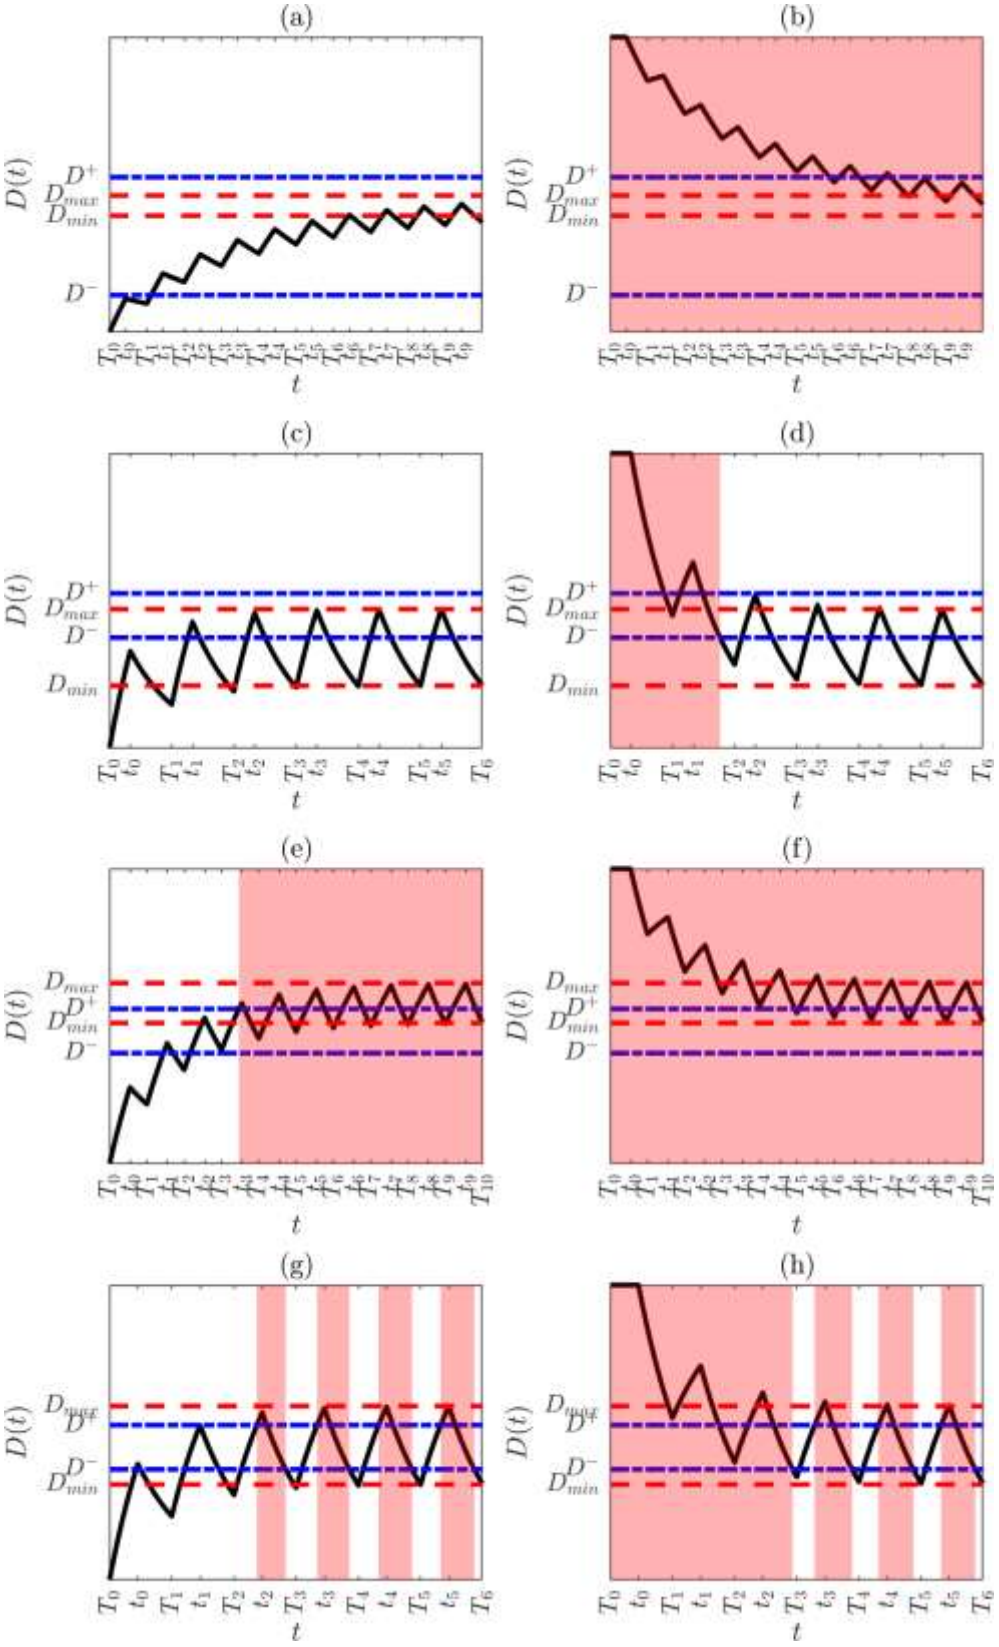


Figure H.3 Solutions $D(t)$ in response to oscillating R-switch, with duty cycles selected to induce each of the four possible behaviours of the T-switch (values selected from figure 10 in the body of this chapter). The white background shows times when the T-switch is off, while the pink shaded regions show times when the T-switch is on. (a), (c), (e) and (g) show the off, bistable, oscillation and on behaviours respectively, with the T-switch initially off. (b), (d), (f) and (h) show the off, bistable, oscillation and on behaviours respectively, with the T-switch initially on.

Appendix I Focal point values and BE bifurcation lines with the revisable T-switch

This appendix is concerned with calculating the boundary-equilibrium bifurcation lines for the double-switch model with the modified K- and T-switches.

First, we define the double-switch model by the ODEs in equation 2.1. We repeat them here as equation I.1 for convenience.

$$\begin{aligned}\frac{dP(t)}{dt} &= P_{\mathrm{env}}\frac{\kappa_{\mathrm{P}}}{1+\gamma_{\mathrm{B}}B(t)} - \alpha_{\mathrm{I}}R(t)P(t) - \delta_{\mathrm{P}}P(t)\\ \frac{dB(t)}{dt} &= \kappa_{\mathrm{B}}\frac{1}{1+\gamma_{\mathrm{R}}R(t)}\frac{1}{1+\gamma_{\mathrm{T}}T(t)}(1-B(t)) - \delta_{\mathrm{B}}K(t)B(t)\\ \frac{dD(t)}{dt} &= \kappa_{\mathrm{D}}R(t) - \delta_{\mathrm{D}}D(t)\end{aligned} \tag{I.1}$$

We modified the K-switch, in Appendix A with the result that it is now flat, so the RK-switch definition is now equation I.2

$$(R(t),K(t)) = \begin{cases}(R_{\mathrm{off}},K_{\mathrm{off}}) \to (R_{\mathrm{on}},K_{\mathrm{on}}) & \text{if } \dot{P}(t)>0 \text{ and } P(t)=P^{+},\\ (R_{\mathrm{on}},K_{\mathrm{on}}) \to (R_{\mathrm{off}},K_{\mathrm{off}}) & \text{if } \dot{P}(t)<0 \text{ and } P(t)=P^{-},\end{cases} \tag{I.2}$$

Where R_{on}, K_{on}, R_{off} and K_{off} are constants.

BE bifurcations under the IL4/TGFβ feedback model

If we take the "IL4/TGFβ negative feedback model" (section 3.2.3) for the Th2 cell number, the T-switch is defined by equation I.3. Note that the T-switch is now reversible, whereas the old definition was irreversible.

$$T(t) = \begin{cases}T_{\mathrm{off}} \to m\cdot(D(t)-D^{-})+T_{\mathrm{on}}, & \text{if } \dot{D}(t)>0 \text{ and } D(t)=D^{+}\\ m\cdot(D(t)-D^{-})+T_{\mathrm{on}} \to T_{\mathrm{off}}, & \text{if } \dot{D}(t)<0 \text{ and } D(t)=D^{-}\end{cases} \tag{I.3}$$

We wish to calculate the boundary-equilibrium bifurcation lines in the κ_P, α_I plane. To do this, we first calculate the focal point values for every possible configuration of the switch states, then we construct the inequalities defining where the focal point values cross the relevant switch thresholds.

Since there are two switches, each with two possible states, there are $2^2 = 4$ configurations for us to consider. These may be written in equations I.4, I.6, I.5 and I.7.

RK-switch off

T-switch off

$$\begin{aligned}\bar{P}(0,0) &= P_{\mathrm{env}}\frac{\kappa_{\mathrm{P}}}{\left(1+\gamma_{\mathrm{B}}\bar{B}(0,0)\right)(\delta_{\mathrm{P}}+\alpha_I R_{\mathrm{off}})}\\ \bar{B}(0,0) &= \frac{\kappa_B}{\kappa_B+\delta_B K_{\mathrm{off}}(1+\gamma_R R_{\mathrm{off}})(1+\gamma_T T_{\mathrm{off}})}\\ \bar{D}(0,0) &= \frac{\kappa_D R_{\mathrm{off}}}{\delta_D}\end{aligned} \tag{I.4}$$

RK-switch on

T-switch off

$$\begin{aligned}\bar{P}(1,0) &= P_{\mathrm{env}}\frac{\kappa_{\mathrm{P}}}{\left(1+\gamma_{\mathrm{B}}\bar{B}(1,0)\right)(\delta_{\mathrm{P}}+\alpha_I R_{\mathrm{on}})}\\ \bar{B}(1,0) &= \frac{\kappa_B}{\kappa_B+\delta_B K_{\mathrm{on}}(1+\gamma_R R_{\mathrm{on}})(1+\gamma_T T_{\mathrm{off}})}\\ \bar{D}(1,0) &= \frac{\kappa_D R_{\mathrm{on}}}{\delta_D}\end{aligned} \tag{I.5}$$

$$
\begin{aligned}
&\text{RK-switch off} && \bar{P}(0,1) = P_{\text{env}}\frac{\kappa_{\text{P}}}{\left(1+\gamma_{\text{B}}\bar{B}(0,1)\right)(\delta_{\text{P}}+\alpha_I R_{\text{off}})} \\
&\text{T-switch on} && \bar{B}(0,1) = \frac{\kappa_B}{\kappa_B+\delta_B K_{\text{off}}(1+\gamma_R R_{\text{off}})(1+\gamma_T(m\cdot(\bar{D}(0,1)-D^-)+T_{\text{on}}))} \\
& && \bar{D}(0,1) = \frac{\kappa_D R_{\text{off}}}{\delta_D} \\
&\text{RK-switch on} && \bar{P}(1,1) = P_{\text{env}}\frac{\kappa_{\text{P}}}{\left(1+\gamma_{\text{B}}\bar{B}(1,1)\right)(\delta_{\text{P}}+\alpha_I R_{\text{on}})} \\
&\text{T-switch on} && \bar{B}(1,1) = \frac{\kappa_B}{\kappa_B+\delta_B K_{\text{on}}(1+\gamma_R R_{\text{on}})(1+\gamma_T(m\cdot(\bar{D}(1,1)-D^-)+T_{\text{on}}))} \\
& && \bar{D}(1,1) = \frac{\kappa_D R_{\text{on}}}{\delta_D}
\end{aligned}
\qquad \text{(I.6)}, \text{(I.7)}
$$

Recall that we can define the switching conditions as in table I.1.

Table I.1 Switching conditions for the RK-switch.

State of the T-switch	**Switching conditions**	
	Region where the switch turns from **off** to **on** **"Turn on condition"**	Region where the switch turns from **on** to **off** **"Turn off condition"**
off	$\bar{P}(0,0) > P^+$	$\bar{P}(1,0) < P^-$
on	$\bar{P}(0,1) > P^+$	$\bar{P}(1,1) < P^-$

We can write expressions for the edges of these regions as in table I.2.

Table I.2 Edges of the bifurcation regions

State of the T-switch	**BE bifurcation line**	
	Edge of switching on region	Edge of switching off region
off	$\bar{P}(0,0) = P^+$	$\bar{P}(1,0) = P^-$
on	$\bar{P}(0,1) = P^+$	$\bar{P}(1,1) = P^-$

We can now write the switching conditions in terms of κ_P and α_I by rearranging the expressions for $\bar{P}$. We can evaluate the boundary-equilibrium bifurcation lines when the T-switch is off, and then again when the T-switch is on. These are written in equations I.8, I.10, I.9 and I.11.

BE-bifurcation lines when T-switch is off:

Edge of switching off region:
$$\alpha_I = \frac{P_{\text{env}}}{P^-\left(1+\gamma_B\bar{B}(1,0)\right)R_{\text{on}}}k_p - \frac{\delta_p}{R_{\text{on}}} \qquad \text{(I.8)}$$

Edge of switching on region:
$$\kappa_P = \frac{P^+\left(1+\gamma_{\text{B}}\bar{B}(0,0)\right)}{P_{\text{env}}} \qquad \text{(I.9)}$$

BE-bifurcation lines when T-switch is on:

Edge of switching off region:
$$\alpha_I = \frac{P_{\text{env}}}{P^-\left(1+\gamma_B\bar{B}(1,1)\right)R_{\text{on}}}k_p - \frac{\delta_p}{R_{\text{on}}} \qquad \text{(I.10)}$$

Edge of switching on region:
$$\kappa_P = \frac{P^+\left(1+\gamma_{\text{B}}\bar{B}(0,1)\right)}{P_{\text{env}}} \qquad \text{(I.11)}$$

BE bifurcations under the IL4 feedback + Th1/Th2 mutual inhibition model

If we were instead to consider the IL4 feedback + Th1/Th2 mutual inhibition model for the T-switch, we would take equation 3.16 as our definition of the T-switch. We repeat it here as I.12 for convenience.

$$T(t) = \begin{cases} T_{\text{off}} \to T_{\text{on}} & if \quad D(t) = D^{+} \; and \; \dot{D}(t) > 0 \\ T_{\text{on}} \to T_{\text{off}} & if \quad D(t) = D^{-} \; and \; \dot{D}(t) < 0 \\ m \cdot (D(t) - D^{++}) + T_{\text{on}} & if \quad D(t) > D^{++} \end{cases} \tag{I.12}$$

This form of T-switch is not hugely different from the IL4/TGFβ feedback model. The only difference is the flat section between D^{+} and D^{++}. Our analysis proceeds almost identically.

We once again calculate the focal point values for every possible configuration of the switch states and construct the inequalities defining where the focal point values cross the relevant switch thresholds.

The four sets of focal point values can now be written in equations I.13, I.14, I.15 and I.16.

RK-switch off, T-switch off:

$$\begin{aligned} \bar{P}(0,0) &= P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1 + \gamma_{\text{B}} \bar{B}(0,0)\right)(\delta_{\text{P}} + \alpha_I R_{\text{off}})} \\ \bar{B}(0,0) &= \frac{\kappa_B}{\kappa_B + \delta_B K_{\text{off}}(1 + \gamma_R R_{\text{off}})(1 + \gamma_T T_{\text{off}})} \\ \bar{D}(0,0) &= \frac{\kappa_D R_{\text{off}}}{\delta_D} \end{aligned} \tag{I.13}$$

RK-switch on, T-switch off:

$$\begin{aligned} \bar{P}(1,0) &= P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1 + \gamma_{\text{B}} \bar{B}(1,0)\right)(\delta_{\text{P}} + \alpha_I R_{\text{on}})} \\ \bar{B}(1,0) &= \frac{\kappa_B}{\kappa_B + \delta_B K_{\text{on}}(1 + \gamma_R R_{\text{on}})(1 + \gamma_T T_{\text{off}})} \\ \bar{D}(1,0) &= \frac{\kappa_D R_{\text{on}}}{\delta_D} \end{aligned} \tag{I.14}$$

RK-switch off, T-switch on:

$$\begin{aligned} \bar{P}(0,1) &= P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1 + \gamma_{\text{B}} \bar{B}(0,1)\right)(\delta_{\text{P}} + \alpha_I R_{\text{off}})} \\ \bar{B}(0,1) &= \begin{cases} \dfrac{\kappa_B}{\kappa_B + \delta_B K_{\text{off}}(1 + \gamma_R R_{\text{off}})(1 + \gamma_T T_{\text{on}})} & if \quad D^{-} < \bar{D}(0,1) < D^{++} \\ \dfrac{\kappa_B}{\kappa_B + \delta_B K_{\text{off}}(1 + \gamma_R R_{\text{off}})(1 + \gamma_T (m \cdot (\bar{D}(0,1) - D^{++}) + T_{\text{on}}))} & if \quad D^{++} < \bar{D}(0,1) \end{cases} \\ \bar{D}(0,1) &= \frac{\kappa_D R_{\text{off}}}{\delta_D} \end{aligned} \tag{I.15}$$

RK-switch on, T-switch on:

$$\begin{aligned} \bar{P}(1,1) &= P_{\text{env}} \frac{\kappa_{\text{P}}}{\left(1 + \gamma_{\text{B}} \bar{B}(1,1)\right)(\delta_{\text{P}} + \alpha_I R_{\text{on}})} \\ \bar{B}(1,1) &= \begin{cases} \dfrac{\kappa_B}{\kappa_B + \delta_B K_{\text{on}}(1 + \gamma_R R_{\text{on}})(1 + \gamma_T T_{\text{on}})} & if \quad D^{-} < \bar{D}(1,1) < D^{++} \\ \dfrac{\kappa_B}{\kappa_B + \delta_B K_{\text{on}}(1 + \gamma_R R_{\text{on}})(1 + \gamma_T (m \cdot (\bar{D}(1,1) - D^{++}) + T_{\text{on}}))} & if \quad D^{++} < \bar{D}(1,1) \end{cases} \\ \bar{D}(1,1) &= \frac{\kappa_D R_{\text{on}}}{\delta_D} \end{aligned} \tag{I.16}$$

The BE bifurcation line equations are defined in exactly the same way as in equations I.8, I.10, I.9 and I.11.

Appendix J Standard parameter sets

We make use of four standard parameter sets throughout our figures, each of which shows an interesting category of model output. These parameter sets are in table J.1.

Table J.1 Standard parameter sets developed in section 5.2.3.4

Parameter	P1 values	P2 values	P3 values	P4 values
P_{env}	1	1	0.7146	0.7146
$\overline{P_{ex}}$	0	0	0	0
$\overline{B_{ex}}$	0	0	0	0
κ_P	Range 0 - 35	Range 0 - 35	Range 0 - 35	Range 0 - 35
α_P	Range 0 - 35	Range 0 - 35	Range 0 - 35	Range 0 - 35
γ_{PB}	1	1	0.1371	0.1371
δ_P	1	1	2.7484	2.7484
κ_B	1	1	9.9311	9.9311
γ_{BB}	1	1	0.26097	0.26097
γ_{BR}	1	1	9.0002	9.0002
δ_B	1	1	3.0382	3.0382
κ_A	1	1	9.2564	9.2564
γ_{AB}	1	1	9.7497	9.7497
δ_A	1	1	1.0440	1.0440
κ_D	Range 0 - 1	Range 0 - 1	Range 0 - 1	Range 0 - 1
δ_D	Range 0 - 1	Range 0 - 1	Range 0 - 1	Range 0 – 1
R_{off}	0	0	0	0
R_{on}	1	1	1	1
K_{off}	1	1	1	1
K_{on}	2	2	2	2
P^-	0.2	0.2	0.2	0.2
P^+	0.8	0.8	0.8	0.8
G_{off}	2	8	2	2
G_{on}	1	0.2	1	1
$B^-(R_{\text{off}}, T_{\text{off}})$	0.2	0.45	0.2	0.45
$B^+(R_{\text{off}}, T_{\text{off}})$	0.8	0.55	0.8	0.55
$B^-(R_{\text{on}}, T_{\text{off}})$	0.3	0.45	0.3	0.55
$B^+(R_{\text{on}}, T_{\text{off}})$	0.9	0.55	0.9	0.65
$B^-(R_{\text{off}}, T_{\text{on}})$	0.1	0.45	0.1	0.35
$B^+(R_{\text{off}}, T_{\text{on}})$	0.7	0.55	0.7	0.45
$B^-(R_{\text{on}}, T_{\text{on}})$	0.15	0.45	0.15	0.4
$B^+(R_{\text{on}}, T_{\text{on}})$	0.75	0.55	0.75	0.5

Appendix K Behaviour-state dependence of the on the G- and RK-switch bifurcation diagrams

This appendix shows the results obtained by repeating the process described in section 5.2.4.1 but using all of the parameter sets P1, P2, P3 and P4. In section 5.2.4.1 we only showed the P1 results for the sake of brevity.

Bifurcation diagrams showing the behaviour of the G-switch under the P1 parameterisation (which were shown in section 5.2.4.1) are as follows.

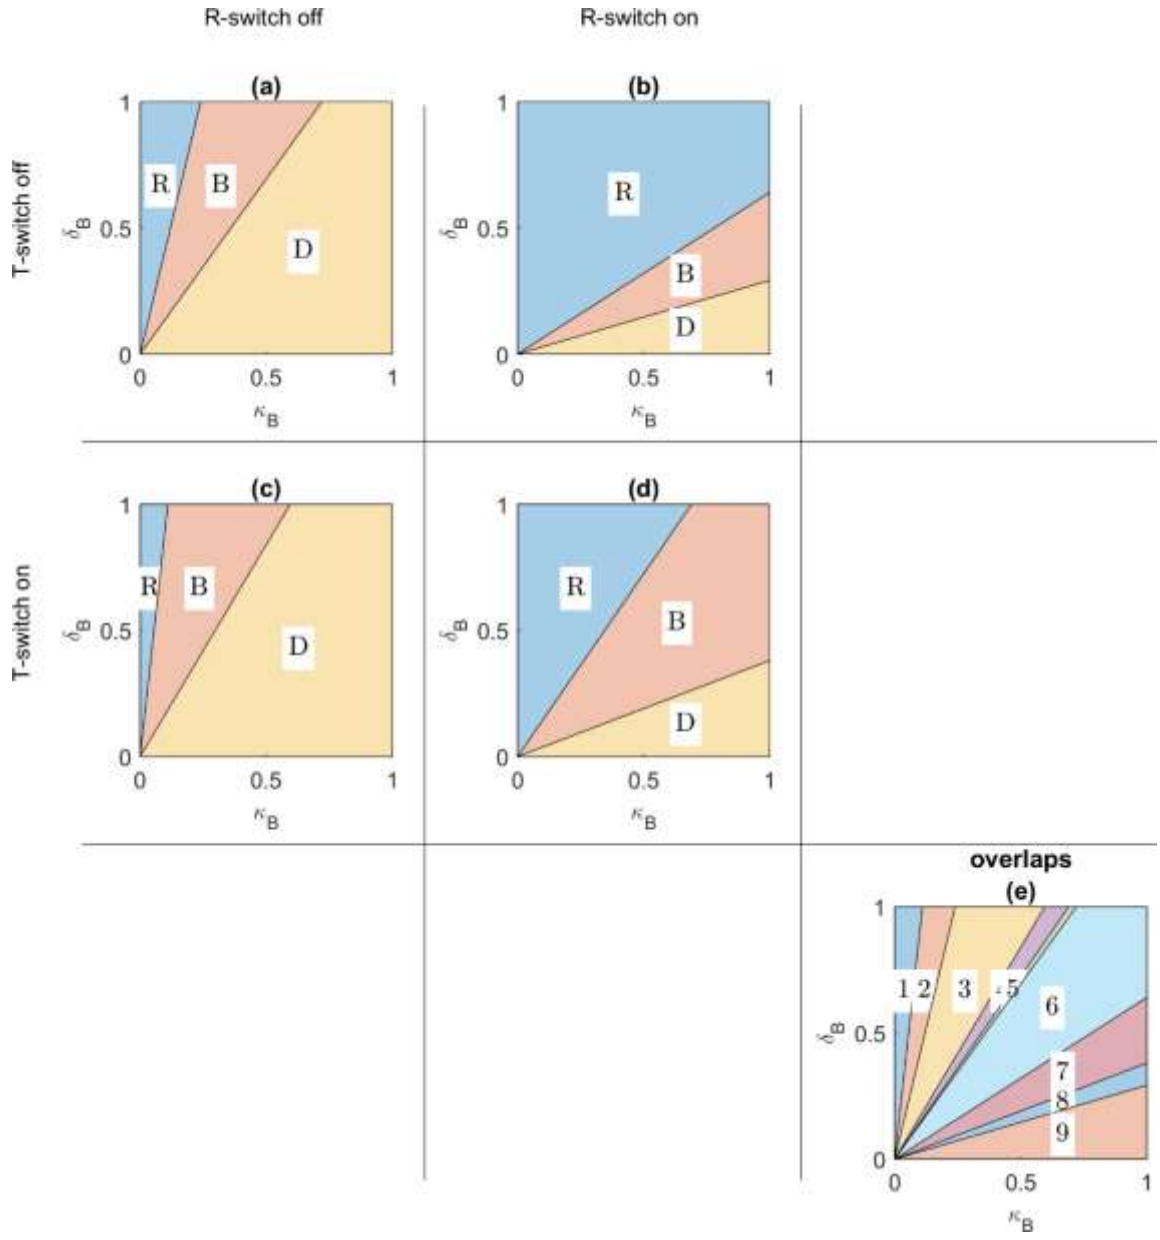


Figure K.1 G-switch bifurcation diagrams under the P1 parameterisation with the R-switch and T-switch in different states.
(Repeated figure from section 5.2.4)
(a) shows the bifurcation diagram with the R-switch off and T-switch off, (b) R-switch on and T-switch off, (c) R-switch off and T-switch on, (d) R-switch on and T-switch on, (e) shows the overlaps of all the other bifurcation diagrams (a) through (d).

The resulting diagram of overlaps between the behaviour regions figure K.1 (e) resulted in nine distinct behaviour-state relationships, listed in table K.1.

Table K.1 List of behaviour-state relationships under the P1 parameterisation.

Region number	RK-switch state	T-switch state	G-switch behaviour
1	0	0	R
1	1	0	R
1	0	1	R
1	1	1	R
2	0	0	R
2	1	0	R
2	0	1	B
2	1	1	R
3	0	0	B
3	1	0	R
3	0	1	B
3	1	1	R
4	0	0	B
4	1	0	R
4	0	1	D
4	1	1	R
5	0	0	B
5	1	0	R
5	0	1	D
5	1	1	B
6	0	0	D
6	1	0	R
6	0	1	D
6	1	1	B
7	0	0	D
7	1	0	B
7	0	1	D
7	1	1	B
8	0	0	D
8	1	0	B
8	0	1	D
8	1	1	D
9	0	0	D
9	1	0	D
9	0	1	D
9	1	1	D

When we take values for κ_B and δ_B from the centroid of each of these nine regions, we produce nine slightly different RK-switch bifurcation diagrams each can produce a slightly different number of

regions. For example, taking the κ_B and δ_B values from the centroid of region 1 in figure K.1, the RK-switch bifurcation diagrams loop like figure K.2

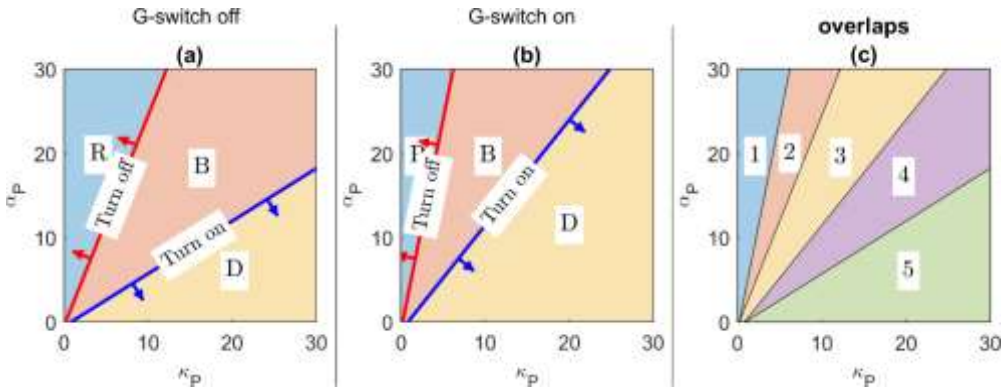


Figure K.2 RK-switch bifurcation diagram produced for the P1 parameterisation, taking values κ_B and δ_B from region 1 of the G-switch bifurcation diagram in figure K.1 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

Meanwhile, if we take values κ_B and δ_B from the centroid of region 9 of figure K.1, the RK-switch bifurcation diagrams loop like figure K.3 which is very slightly different.

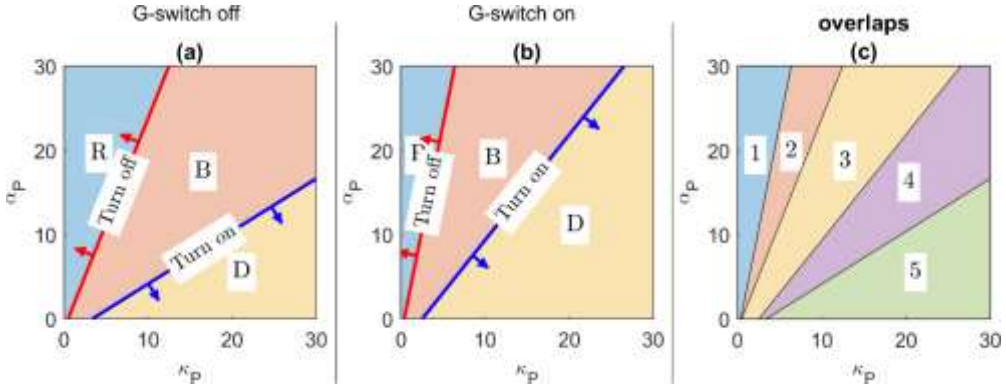


Figure K.3 RK-switch bifurcation diagram produced for the P1 parameterisation, taking values κ_B and δ_B from region 9 of the G-switch bifurcation diagram in figure K.1 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

In this example, the RK-switch bifurcation diagram produces nine “overlap” regions regardless of which regions of the G-switch bifurcation diagram we choose.

Bifurcation diagrams under the P2 parameterisation are as follows.

Under the P2 parameterisation, the G-switch has the Recovery, Oscillation and Damage behaviours available in the bifurcation diagram. The resulting overlaps from all the R- and T-switch state combinations leads to nine regions in figure K.4 (e).

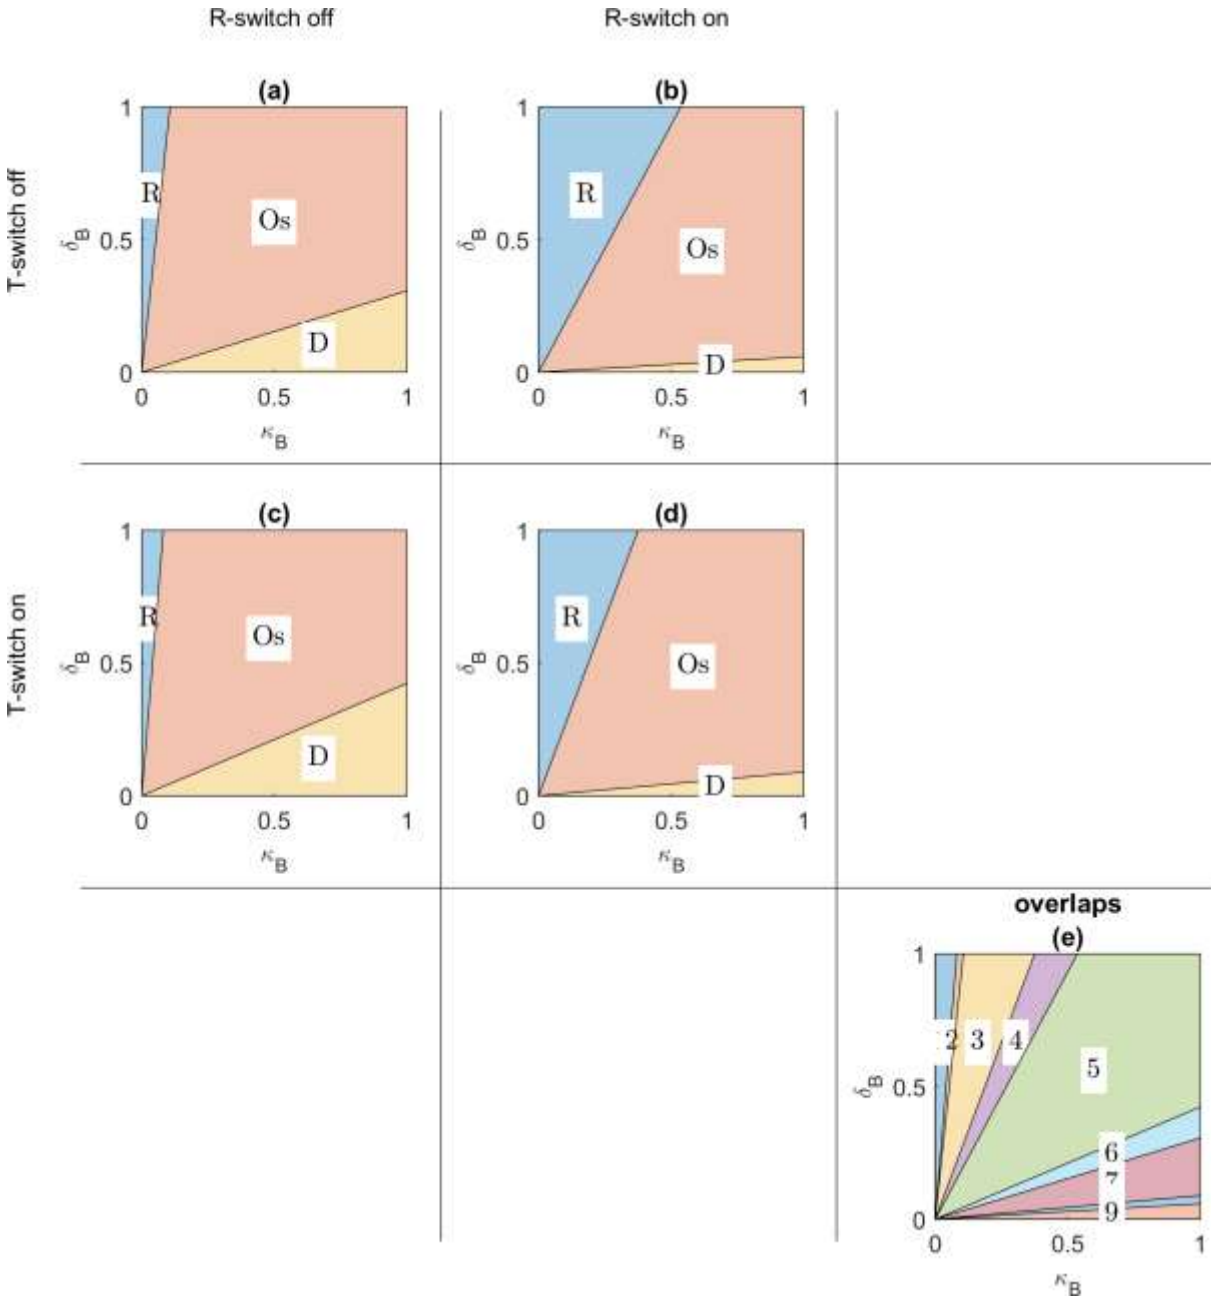


Figure K.4 G-switch bifurcation diagrams under the P2 parameterisation with the R-switch and T-switch in different states.
(a) shows the bifurcation diagram with the R-switch off and T-switch off
(b) R-switch on and T-switch off
(c) R-switch off and T-switch on
(d) R-switch on and T-switch on
(e) shows the overlaps of all the other bifurcation diagrams (a) through (d).

Each of the nine regions in figure K.4 (e) has a different behaviour-state relationship. All of them are listed in table K.2. Regions where the behaviour changes as a result of a change of R- or T-switch state are highlighted in green.

Table K.2 List of behaviour-state relationships under the P2 parameterisation.

Region number	RK-switch state	T-switch state	G-switch behaviour
1	0	0	R
1	1	0	R
1	0	1	R
1	1	1	R
2	0	0	R
2	1	0	R
2	0	1	Os
2	1	1	R
3	0	0	Os
3	1	0	R
3	0	1	Os
3	1	1	R
4	0	0	Os
4	1	0	R
4	0	1	Os
4	1	1	Os
5	0	0	Os
5	1	0	Os
5	0	1	Os
5	1	1	Os
6	0	0	Os
6	1	0	Os
6	0	1	D
6	1	1	Os
7	0	0	D
7	1	0	Os
7	0	1	D
7	1	1	Os
8	0	0	D
8	1	0	Os
8	0	1	D
8	1	1	D
9	0	0	D
9	1	0	D
9	0	1	D
9	1	1	D

In this p2 parameterisation case, the RK-switch bifurcation diagram always produces six overlap regions. The RK-switch bifurcation diagrams produced by taking κ_B and δ_B values from region 1 are shown in figure K.5.

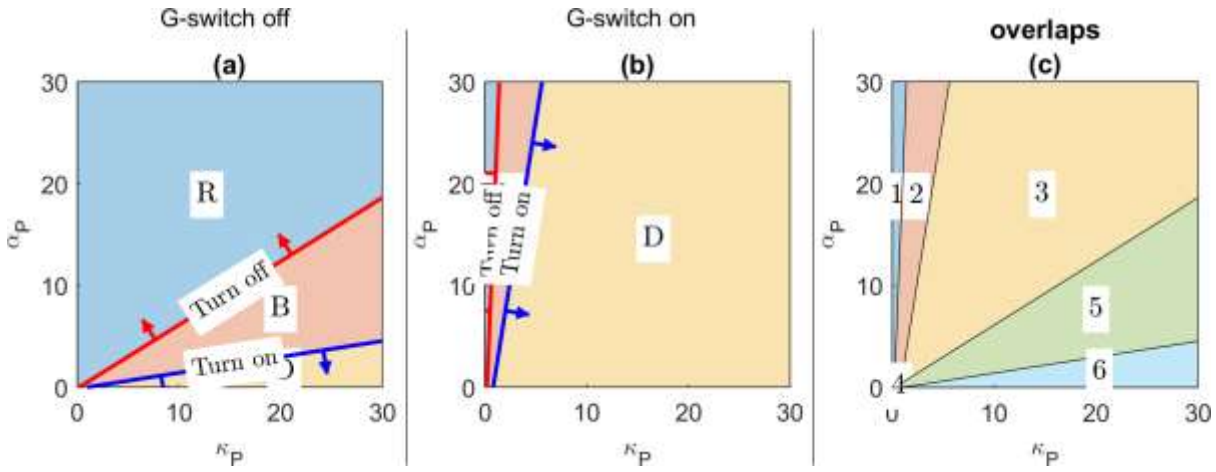


Figure K.5 RK-switch bifurcation diagram produced for the P2 parameterisation, taking values κ_B and δ_B from region 1 of the G-switch bifurcation diagram in figure K.4 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

Bifurcation diagrams under the P3 parameterisation are as follows.

Under the P3 parameterisation, the G-switch has the Recovery, Oscillation and Damage behaviours available in the bifurcation diagram. The resulting overlaps from all the R- and T-switch state combinations leads to nine regions in figure K.6 (e).

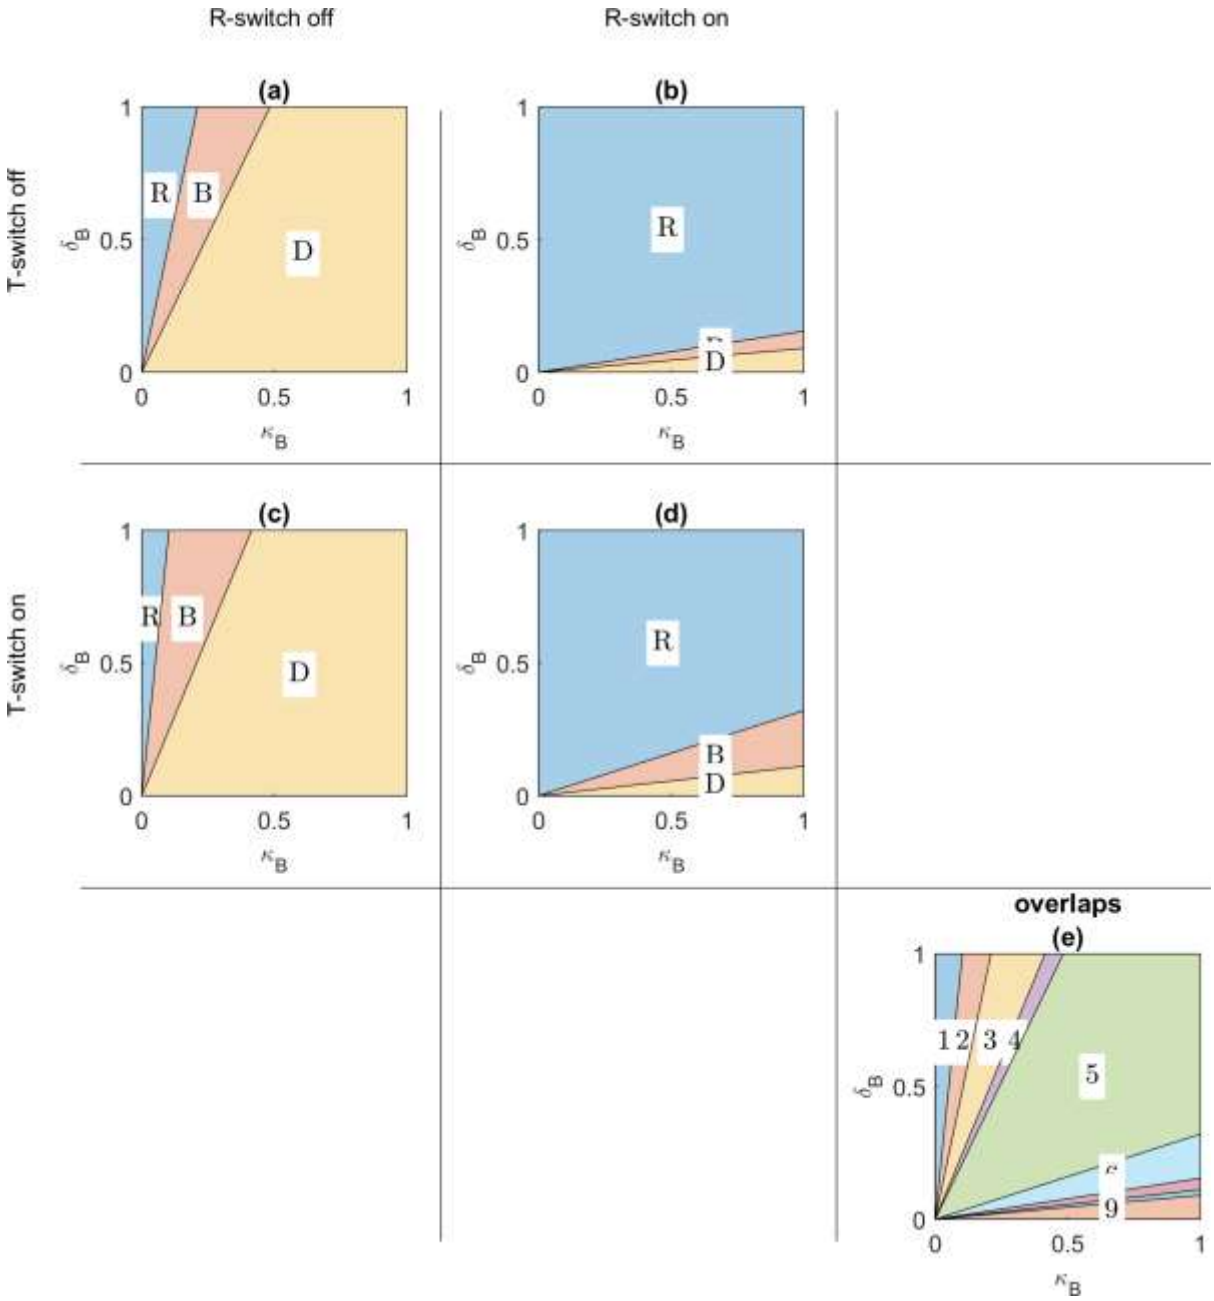


Figure K.6 G-switch bifurcation diagrams under the P3 parameterisation with the R-switch and T-switch in different states.
(a) shows the bifurcation diagram with the R-switch off and T-switch off
(b) R-switch on and T-switch off
(c) R-switch off and T-switch on
(d) R-switch on and T-switch on
(e) shows the overlaps of all the other bifurcation diagrams (a) through (d)

Each of the nine regions in figure K.6 (e) has a different behaviour-state relationship. All of them are listed in table K.3. Regions where the behaviour changes as a result of a change of R- or T-switch state are highlighted in green.

Table K.3 List of behaviour-state relationships under the P3 parameterisation.

Region number	RK-switch state	T-switch state	G-switch behaviour
1	0	0	R
1	1	0	R
1	0	1	R
1	1	1	R
2	0	0	R
2	1	0	R
2	0	1	B
2	1	1	R
3	0	0	B
3	1	0	R
3	0	1	B
3	1	1	R
4	0	0	B
4	1	0	R
4	0	1	D
4	1	1	R
5	0	0	D
5	1	0	R
5	0	1	D
5	1	1	R
6	0	0	D
6	1	0	R
6	0	1	D
6	1	1	B
7	0	0	D
7	1	0	B
7	0	1	D
7	1	1	B
8	0	0	D
8	1	0	B
8	0	1	D
8	1	1	D
9	0	0	D
9	1	0	D
9	0	1	D
9	1	1	D

In this p3 parameterisation case, the RK-switch bifurcation diagram produces a different number of overlap regions depending on the centroidal values of κ_B and δ_B we set. The RK-switch bifurcation diagrams produced by taking κ_B and δ_B values from region 1 are shown in figure K.7. In this case, five overlap regions are produced in figure K.7 (c).

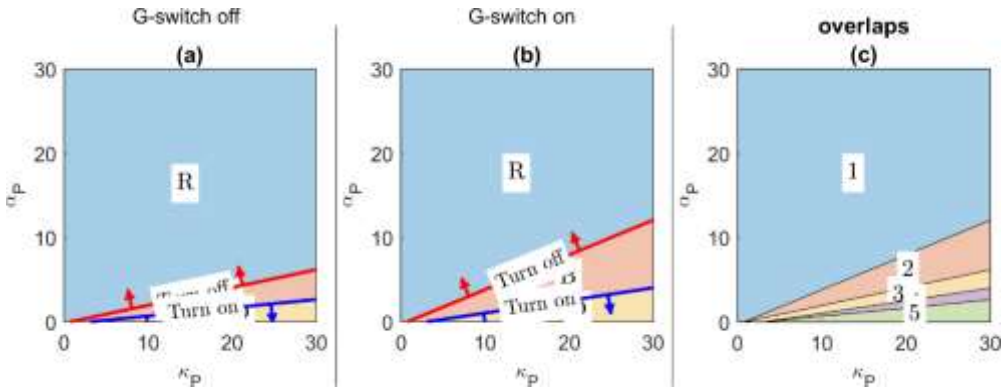


Figure K.7 RK-switch bifurcation diagram produced for the P3 parameterisation, taking values κ_B and δ_B from region 1 of the G-switch bifurcation diagram in figure K.6 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

If we instead take values for κ_B and δ_B from the centroid of region six of the bifurcation diagram in figure K.6 (e), we obtain eight overlap regions as shown in figure K.8.

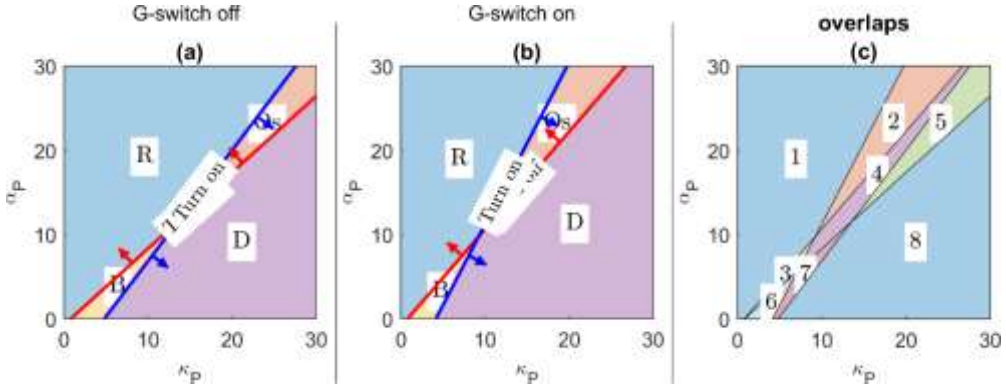


Figure K.8 RK-switch bifurcation diagram produced for the P3 parameterisation, taking values κ_B and δ_B from region 6 of the G-switch bifurcation diagram in figure K.6 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

We summarise the number of RK-switch bifurcation diagram "overlap" regions produced by each region from the figure K.6 (e) as follows in table K.4.

Table K.4 Number of overlaps found on the RK-switch bifurcation diagram for different choices of centroid values from figure K.6 (e)

P3 parameterisation	
Region number from figure K.6 (e)	Number of RK-switch "overlap" regions
1	5
2	5
3	7
4	7
5	8
6	8
7	7
8	6
9	5

Bifurcation diagrams under the P4 parameterisation are as follows.

Under the P4 parameterisation, the G-switch has the Recovery, Oscillation and Damage behaviours available in the bifurcation diagram. The resulting overlaps from all the R- and T-switch state combinations leads to nine regions in figure K.9 (e).

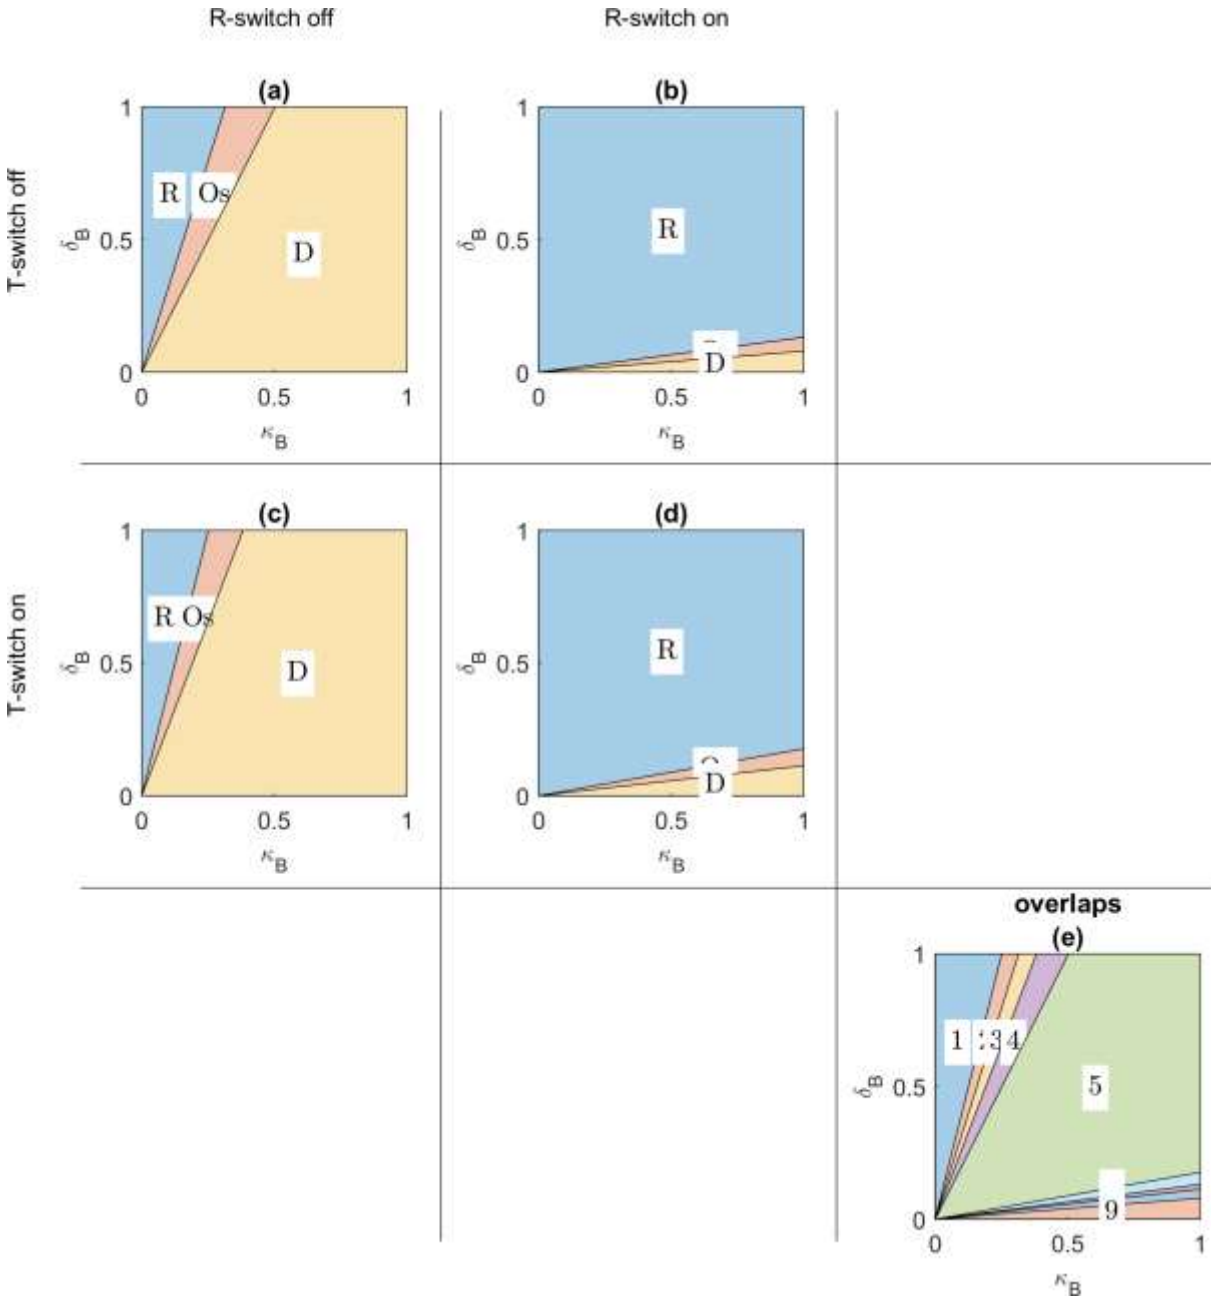


Figure K.9 G-switch bifurcation diagrams under the P4 parameterisation with the R-switch and T-switch in different states.
(a) shows the bifurcation diagram with the R-switch off and T-switch off
(b) R-switch on and T-switch off
(c) R-switch off and T-switch on
(d) R-switch on and T-switch on
(e) shows the overlaps of all the other bifurcation diagrams (a) through (d)

Each of the nine regions in figure K.9 (e) has a different behaviour-state relationship. All of them are listed in table K.5. Regions where the behaviour changes as a result of a change of R- or T-switch state are highlighted in green.

Table K.5 Behaviour state relationships of each of the nine regions produced for the G-switch bifurcation diagram under the P4 parameterisation.

Region number	RK-switch state	T-switch state	G-switch behaviour
1	0	0	R
1	1	0	R
1	0	1	R
1	1	1	R
2	0	0	R
2	1	0	R
2	0	1	Os
2	1	1	R
3	0	0	Os
3	1	0	R
3	0	1	Os
3	1	1	R
4	0	0	Os
4	1	0	R
4	0	1	D
4	1	1	R
5	0	0	D
5	1	0	R
5	0	1	D
5	1	1	R
6	0	0	D
6	1	0	R
6	0	1	D
6	1	1	Os
7	0	0	D
7	1	0	Os
7	0	1	D
7	1	1	Os
8	0	0	D
8	1	0	Os
8	0	1	D
8	1	1	D
9	0	0	D
9	1	0	D
9	0	1	D
9	1	1	D

In this p4 parameterisation case, the RK-switch bifurcation diagram produces a different number of overlap regions depending on the centroidal values of κ_B and δ_B we set. The RK-switch bifurcation diagrams produced by taking κ_B and δ_B values from region 1 are shown in figure K.10. In this case, five overlap regions are produced in figure K.10 (c).

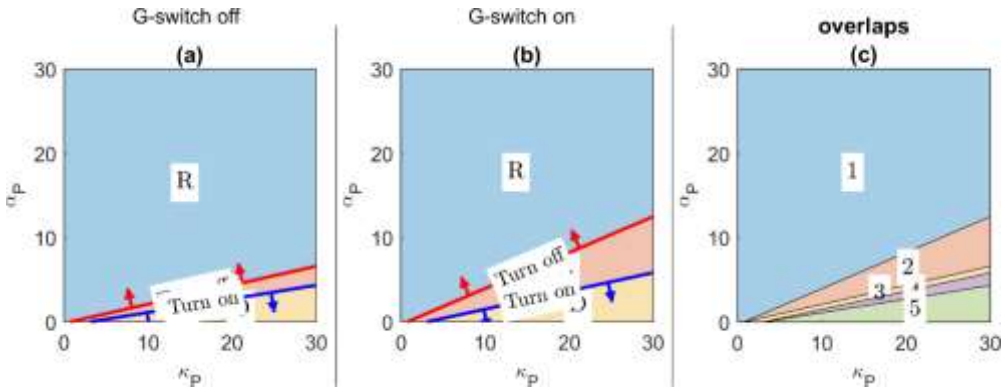


Figure K.10 RK-switch bifurcation diagram produced for the P3 parameterisation, taking values κ_B and δ_B from region 1 of the G-switch bifurcation diagram in figure K.9 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

If we instead take values for κ_B and δ_B from the centroid of region six of the bifurcation diagram in figure K.9 (e), we obtain eight overlap regions as shown in figure K.11.

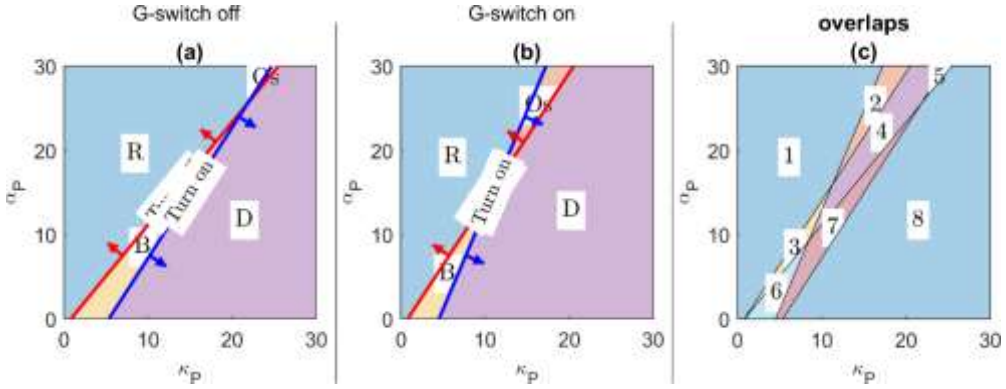


Figure K.11 RK-switch bifurcation diagram produced for the P3 parameterisation, taking values κ_B and δ_B from region 6 of the G-switch bifurcation diagram in figure K.9 (e). (a) with the G-switch off, (b) with the G-switch on and (c) the overlaps between the two.

We summarise the number of RK-switch bifurcation diagram "overlap" regions produced by each region from the figure K.9 (e) as follows in table K.4.

Table K.6 Number of overlaps found on the RK-switch bifurcation diagram for different choices of centroid values from figure K.6 (e)

P4 parameterisation	
Region number from figure K.6 (e)	Number of RK-switch "overlap" regions
1	5
2	6
3	7
4	7
5	9
6	8
7	7
8	6
9	5

Appendix L Sensitivity analyses

In this appendix, we show some further results of the sensitivity analysis for the RK- and G-switch bifurcation diagrams. Only one of these sensitivity analysis diagrams was included in Chapter 5 for the sake of brevity.

Sensitivity of the RK-switch behaviours

We evaluate the sensitivities of the bifurcation diagram for the R-switch when the four behaviours {off, bistable, oscillate, on}. The resulting parameter sensitivities are shown in figure L.1.

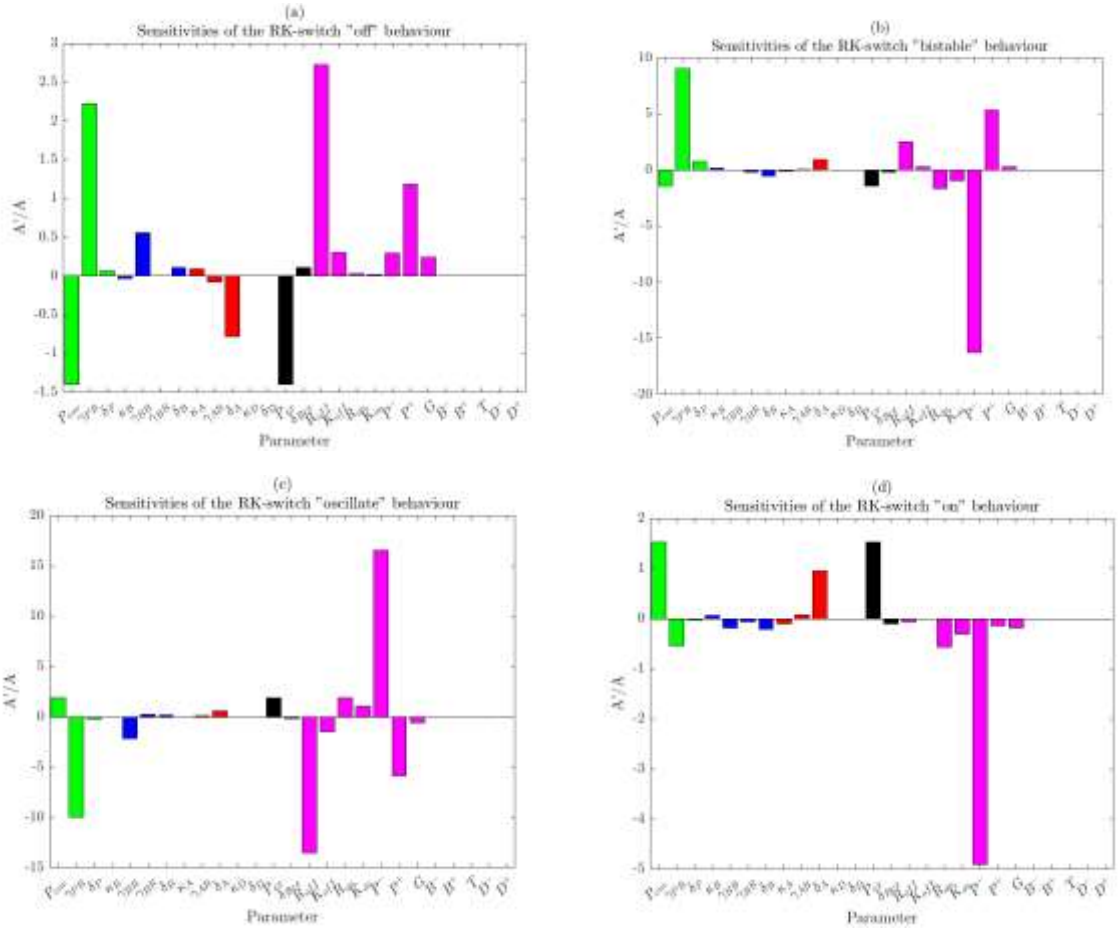


Figure L.1 Local sensitivities of the RK-switch for the "p4" parameterisation (parameter values in Appendix J).

- *(a) sensitivity of the "off" region,*
- *(b) sensitivity of the "bistable" region,*
- *(c) sensitivity of the "oscillation" region and*
- *(d) sensitivity of the "on" region.*

*Parameter bars are colour-coded as follows: **green** – pathogen-related parameters, **blue** – barrier-related, **red** – immunity-related, **black –** external perturbations, **purple** – switch-related parameters.*

These sensitivity analyses show us which parameters have no effect on the area of the behaviour regions. For example, the calculated sensitivity to the parameters κ_D and δ_D is zero. This is expected since these parameters appear only in the ODE for D, which plays no part in the activation of the RK-switch. In addition, B^-, B^+, D^- and D^+, which parameterise the G- and T-switches respectively, also

have no effect on the RK-switch regions. This highlights the parts of the system which are not directly coupled to the RK-switch behaviour.

The rest of the model has at least come effect on the areas of the RK-switch behaviour regions. Unsurprisingly, the RK-switch's threshold values P^- and P^+ are shown to be important, and the parameters directly relating to pathogen dynamics P_{env}, γ_{PB} come a close second.

Sensitivity of the G-switch behaviours

Next, we directly evaluate the local sensitivities of each of the behaviours of the G-switch when exhibiting the {off, bistable, on} phenotypes, for which only the "off", "bistable" and "on" behaviours are produced, and when exhibiting the {off, oscillation, on} phenotypes. With this information, we can detect which parts of the model most strongly affect the dependence of the G-switch on the κ_B and δ_B parameters. We also gain understanding of how much of the model affects the G-switch behaviour in comparison to the RK- and T-switch by seeing how many parameters are sensitive.

Figure L.2 shows the sensitivities of the G-switch behaviours under the P1 parameterisation. Since we are evaluating the sensitivities of the areas of these behaviour regions in the κ_B, δ_B plane, it does not make sense to evaluate the sensitivity to the κ_P or α_P parameters. They are therefore omitted from the list on the x-axis.

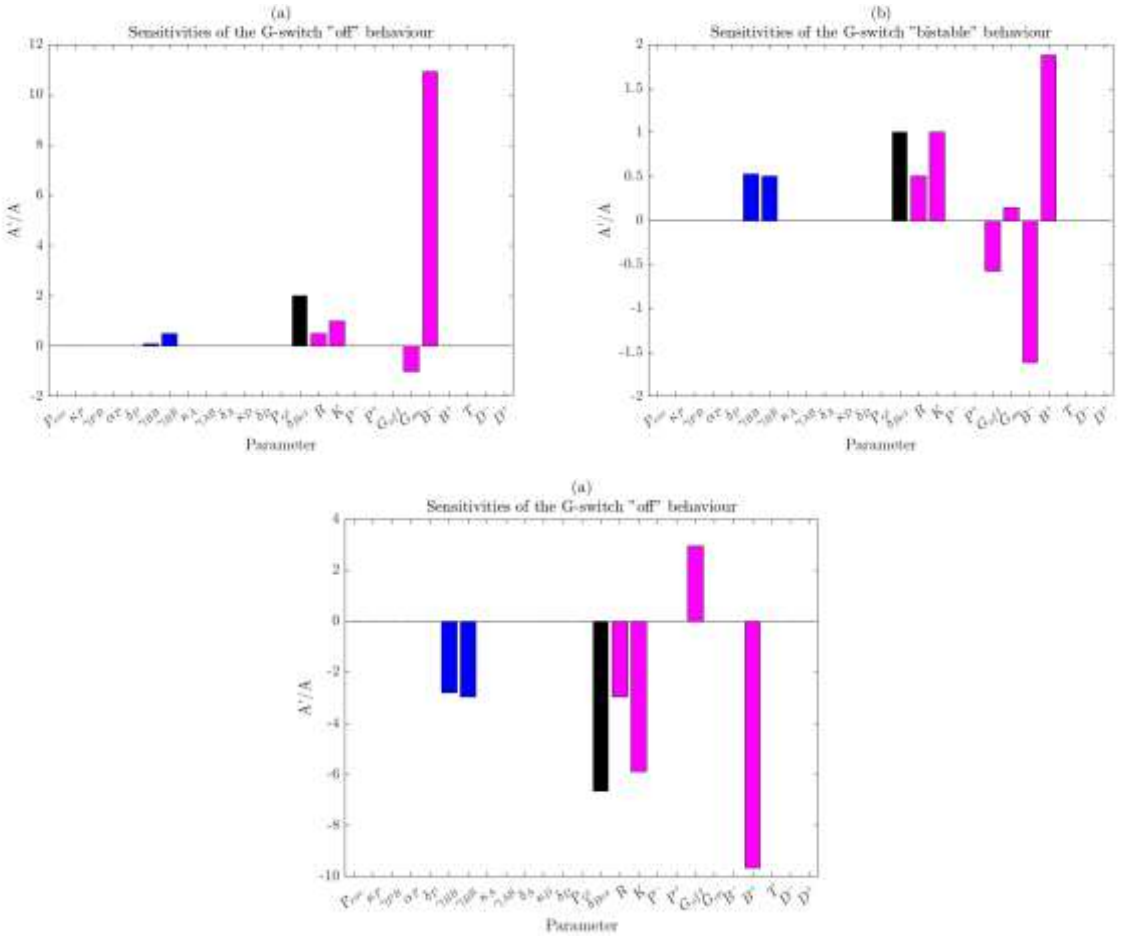


Figure L.2 Parameter sensitivities evaluated for the R, B and D behaviours of the G-switch

(a) sensitivity of the "off" region,

(b) sensitivity of the "bistable" region, and

(c) sensitivity of the "on" region.

Parameter bars are colour-coded as follows: ***green*** *– pathogen-related parameters,* ***blue*** *– barrier-related,* ***red*** *– immunity-related,* ***black –*** *external perturbations,* ***purple*** *– switch-related parameters.*

The G-switch has notably fewer parameters to which it is sensitive than is the case for the RK-switch, which makes sense since the BE bifurcation lines we derived for the G-switch contain fewer parameters than the RK-switch bifurcation lines. The parameters that are shown to be sensitive for the G-switch are all found in the condition we derived which determines the nature of the G-switch's bifurcation diagram (equation 5.39).

In this sensitivity analysis, we are assuming the RK-switch is on, setting $R = R_{\text{on}}$ and $K = K_{\text{on}}$, so the behaviours become sensitive to the parameter γ_{BR} . When we turn off the RK-switch, the parameter γ_{BR} no longer has any effect on the behaviours because it only ever multiplies by R in our ODEs, and when $R = R_{\text{off}} = 0$ this makes its effect becomes zero. This makes sense because the parameter represents the effect of innate immune responses on the barrier growth. In the absence of an innate immune response, it has no impact at all.

We also evaluate the local sensitivities for the P4 parameterisation where the {R, Os, D} behaviours are produced on the G-switch bifurcation diagram. These sensitivities are shown in figure L.3.

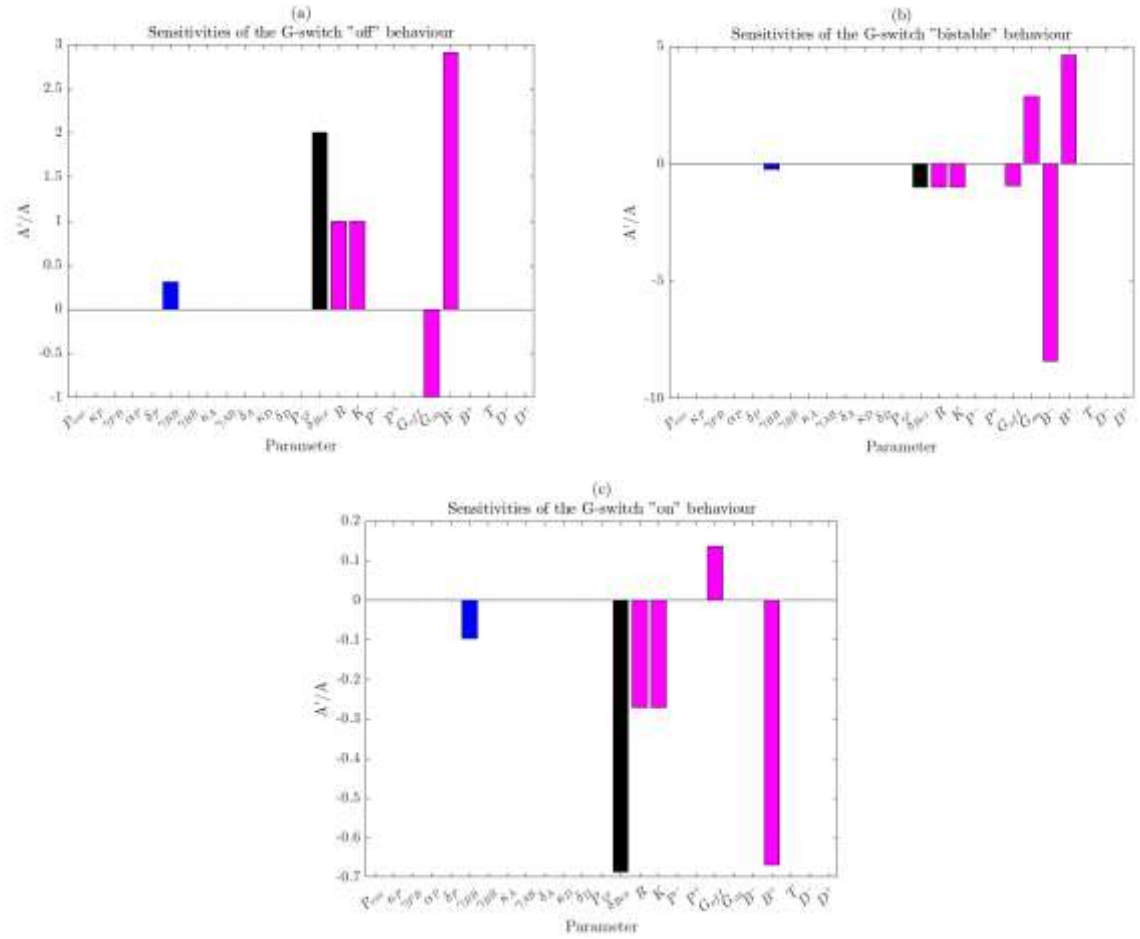


Figure L.3 Parameter sensitivities evaluated for the R, Os and D behaviours of the G-switch

(a) sensitivity of the "off" region,

(b) sensitivity of the "bistable" region, and

(c) sensitivity of the "on" region.

Parameter bars are colour-coded as follows: ***green*** *– pathogen-related parameters,* ***blue*** *– barrier-related,* ***red*** *– immunity-related,* ***black –*** *external perturbations,* ***purple*** *– switch-related parameters.*

The sensitivities of the G-switch in the case where {off, oscillate, on} behaviours are possible seems to show the same kind of profile of sensitivities, with the same few parameters affecting the area of these behaviours.

Appendix M Numerical integration of the triple-switch model

The most direct approach to working out the model's dynamics is simply to use numerical integration techniques. We can use a standard software such as MATLAB's ode45 to find numerical solutions to the ODE part of our model. The hybrid nature of the model presents us with a challenge in that we must identify the moments in time when the switch states transition, and then make sure the values of the switches in the ODE change value at that instant.

Here are the ODEs again in equation M.1.

$$\begin{aligned}
\frac{dP}{dt} &= (P_{env} + P_{ext})\frac{\kappa_P}{1+\gamma_{PB}\cdot B} - \alpha_P \cdot A \cdot P - \delta_P \cdot P \\
\frac{dB}{dt} &= \kappa_B \cdot \frac{1}{1+\gamma_{BB}\cdot B}\cdot\frac{1}{1+\gamma_{BR}\cdot R(P^-,P^+)} G(B^-,B^+) - \delta_B \cdot K(P^-,P^+)\cdot B - \delta_{Bext}(t)\cdot B \\
\frac{dA}{dt} &= \kappa_A \cdot \frac{1}{1+\gamma_{AB}\cdot B}\cdot G(B^-,B^+) - \delta_A \cdot A \\
\frac{dD}{dt} &= \kappa_D \cdot R(P^-,P^+) + D_{ext}(t) - \delta_D \cdot D
\end{aligned} \tag{M.1}$$

Transitions occur when the continuous-valued state variables cross a threshold value in a specific direction as defined in the switch definition, repeated below in equations M., M.3 and M.4.

$$\big(R(P^-,P^+),K(P^-,P^+)\big) = \begin{cases} (R_{\text{off}},K_{\text{off}}) \to (R_{\text{on}},K_{\text{on}}), & \text{if } \dot{P}(t)>0 \text{ and } P(t)=P^+ \\ (R_{\text{on}},K_{\text{on}}) \to (R_{\text{off}},K_{\text{off}}), & \text{if } \dot{P}(t)<0 \text{ and } P(t)=P^- \end{cases} \tag{M.2}$$

$$T(D^-,D^+) = \begin{cases} 0 \to T_{\text{on}} & \text{if} \quad D(t)=D^+ \text{ and } \dot{D}(t)>0 \\ T_{\text{on}} \to 0 & \text{if} \quad D(t)=D^- \text{ and } \dot{D}(t)<0 \end{cases} \tag{M.3}$$

$$G(B^-,B^+) = \begin{cases} G_{\text{off}} \to G_{\text{on}}, & \text{if } \dot{B}(t)>0 \text{ and } B(t)=B^+(R,T) \\ G_{\text{on}} \to G_{\text{off}}, & \text{if } \dot{B}(t)<0 \text{ and } B(t)=B^-(R,T) \end{cases} \tag{M.4}$$

Detection of switch transition events can be achieved using the "event location" feature, built into MATLAB ode solvers https://nl.mathworks.com/help/matlab/math/ode-event-location.html. Using this feature, we can specify the threshold values of the switches, and the direction which the threshold must be crossed to bring about a transition. When the event location feature is triggered, the ODE solver will stop and work out the time of event down to the tolerance we specify. It's important that we use this feature in our solution because otherwise the numerical ODE solver may just "step over" an event, evaluating points in time before and after the event, and not really be able to tell us at which instant in time the transition took place.

The process of integrating the hybrid dynamical system over time is summarised in figure M.1. A simplified version of our MATLAB code is placed at the end of this appendix, to help you get the idea of how the flow diagram in figure M.1 can be implemented.

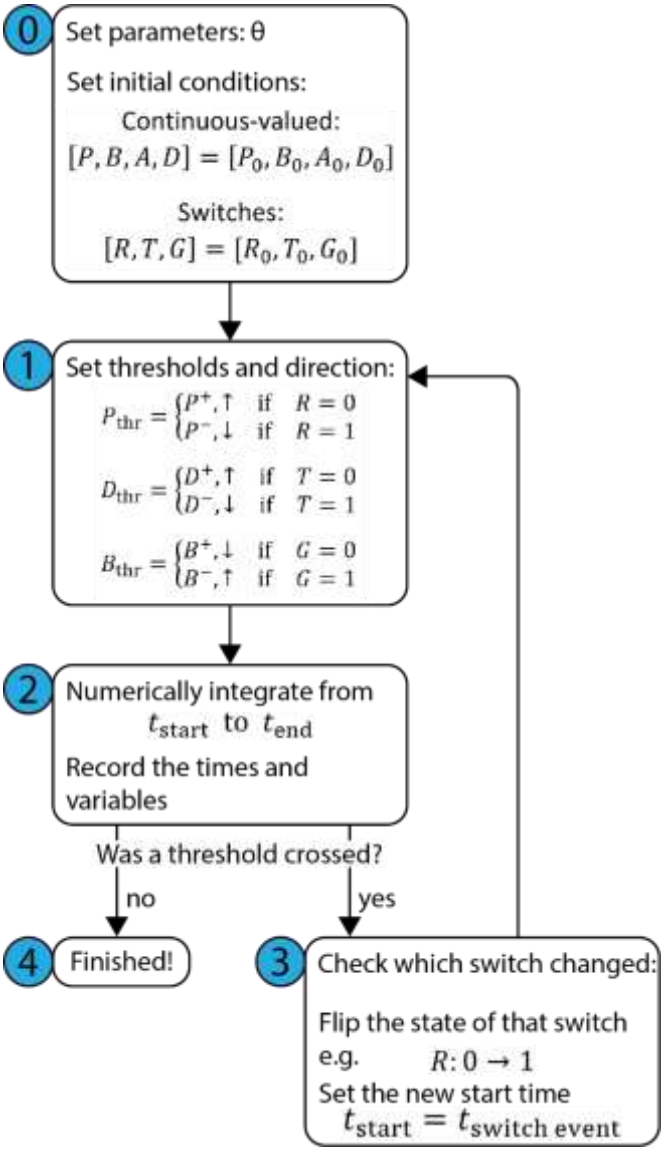


Figure M.1 Flow diagram explaining how to numerically integrate a hybrid dynamical system.

Simplified MATLAB code to numerically simulate the triple-switch model. Note it is spread across three separate functions. The “simulate” function calls the “diffEqn” and “eventsDetector” function, so make sure they’re accessible in the same folder.

```
function [tY, Y, tC, C] = simulate(tEnd, Y0, C0, parameters)
% simulate the triple switch AD model
% 0: set the initial conditions
Y = Y0; % continuous-valued state variables [P, B, A, D]
C = C0; % discrete-valued state variables [R, T, D]
tStart = 0; % we start integrating from time t=0
% record the initial switch configuration and start time
tC = 0;
C  = C0;

while(tStart < tEnd )
    % define the ODEs with these switch states and parameter values
    diffEqn   = @(t, y) diffEqn (t, y, C0, parameters);

    % 1: set the thresholds (task carried out by separate function)
    eventDetector = @(t, y) eventsDetector(t, y, C0, parameters);
    % tell MATLAB's solver about the event detector
    options = odeset('Events', eventDetector);

    % 2: numerically integrate the ODEs
    [ tOut, YOut, tEvent, ~, indexEvent ] = ode45(diffEqn, [tStart,tEnd], Y0,
options );
    % 2: append the continuous-valued state variables to our output vectors
    Y = [Y,YOut];
    tY = [tY,tOut];
    % 2: check to see if a switch threshold was crossed in the remaining
    % time to tEnd
    if ~isempty(tEvent) % 3: yes, a transition event occurred before tEnd
        % 3: flip the relevant switch state
        % The events are indexed from 1 as follows
        %  RK | T | G
        % ----+---+----
        %   1 | 2 | 3
        switch indexEvent
            case 1 % RK switch transition
                C0(1) = ~ C0(1);
            case 2 % T switch transition
                C0(2) = ~ C0(2);
            case 3 % G switch transition
                C0(3) = ~ C0(3);
        end
        % record the transition time, and the new switch states
        tC = vertcat(tC, tEvent ); % concatenate the switching time to the end
of the record.
        C  = vertcat(C, C0);

        % 3: set the next tStart time as the
        tStart = tEvent;
        % 3: set the initial conditions for the next integration
        Y0 = Yout(end,:);
    else % 4: no transition event
        return;
    end
end
end
```

```
function [value,isterminal,direction] = eventsFun (t, Y, C, parameters )
%% function to detect times at which switches undergo a transition
% state variables are Y = [P,B,A,D] in that order

%% work out the threshold value and expected direction of change for each
switch.
if C(1)==false % if RK-switch is off
    RK_switch_threshold = Y(1) - parameters.Pplus;
    RK_direction = 1; % stop only if going in positive direction
else % if RK-switch is on
    RK_switch_threshold = Y(1) - parameters.Pminus;
    RK_direction = -1; % stop only if going in negative direction
end

if C(2)==false % if T-switch is off
    T_switch_threshold = Y(4) - parameters.Dplus;
    T_direction = 1; % stop only if going in positive direction
else % if T-switch is on
    T_switch_threshold = Y(4) - parameters.Dminus;
    T_direction = -1; % stop only if going in negative direction
end

if C(3)==false % if G-switch is off
    G_switch_threshold = Y(2) - parameters.Bplus;
    G_direction = -1; % stop only if going in positive direction
else % if T-switch is on
    G_switch_threshold = Y(2) - parameters.Bminus;
    G_direction = 1; % stop only if going in negative direction
end

%% outputs
% value: the threshold value(s) for each switch in the current configuration
% MATLAB's event detection checks when the "value" crosses zero
value = [RK_switch_threshold, T_switch_threshold, G_switch_threshold];

% isterminal: takes the value 1 for each type of "event"
% setting to 1 makes the ODE solving stop so our "simulation" function gets
% a chance to record the results and set up the initial conditions for the
% next integration run
isterminal = [1,1,1];

% direction: which direction of crossing are we interested in detecting.
% 1 = increasing, -1 = decreasing, 0 = either direction.
direction = [RK_direction, T_direction, G_direction];
end
```

```
function dydt = diffEqn ( m, t, y, C, p )
%% diffEqn defines the ODEs of the triple-switch model
if C(1) == false % if RK-switch is off
    R = p.R_off;
    K = p.K_off;
else
    R = p.R_on;
    K = p.K_on;
end
if C(2) == false % if T-switch is off
    T = p.T_off;
else
    T = p.T_on;
end
if C(3) == false % if G-switch is off
    G = p.G_off;
else
    G = p.G_on;
end

P = y(1);
B = y(2);
A = y(3);
D = y(4);

dydt = [ (p.P_env + p.P_ex) * p.kappa_P / ( 1 + p.gamma_PB * B ) - p.alpha_P * A
* P - p.delta_P * P; % dP/dt
    p.kappa_B * G / ((1+p.gamma_BB*B)*(1+p.gamma_BR*R)) - p.delta_B * K * B; %
dB/dt
    p.kappa_A * G / (1+p.gamma_AB*B) - p.delta_A*A; % dA/dt
    p.kappa_D * R - p.delta_D * D ]; % = dD/dt
end
```

Appendix N Numerically checking the bifurcation diagrams

In section 5.2, we describe an analytical way to calculate the bifurcation diagrams of the triple-switch model. The work involved to get to a bifurcation diagram is fairly complex and there's a lot of room for human error in the various bits of MATLAB code we wrote to calculate and draw the diagrams. It's therefore important to have a way to double-check the method is working. The process of checking the bifurcation diagrams is described in this appendix.

The bifurcation diagrams we produced tell us what the behaviour of each switch is for a particular combination of parameter values and initial conditions of the other switches. For example, figure N.1 (reproduced from Chapter 5) shows the predicted behaviours of the G-switch for each configuration of the RK- and T-switch states and the nine regions of the plane within which the behaviour-state relationship is homogeneous.

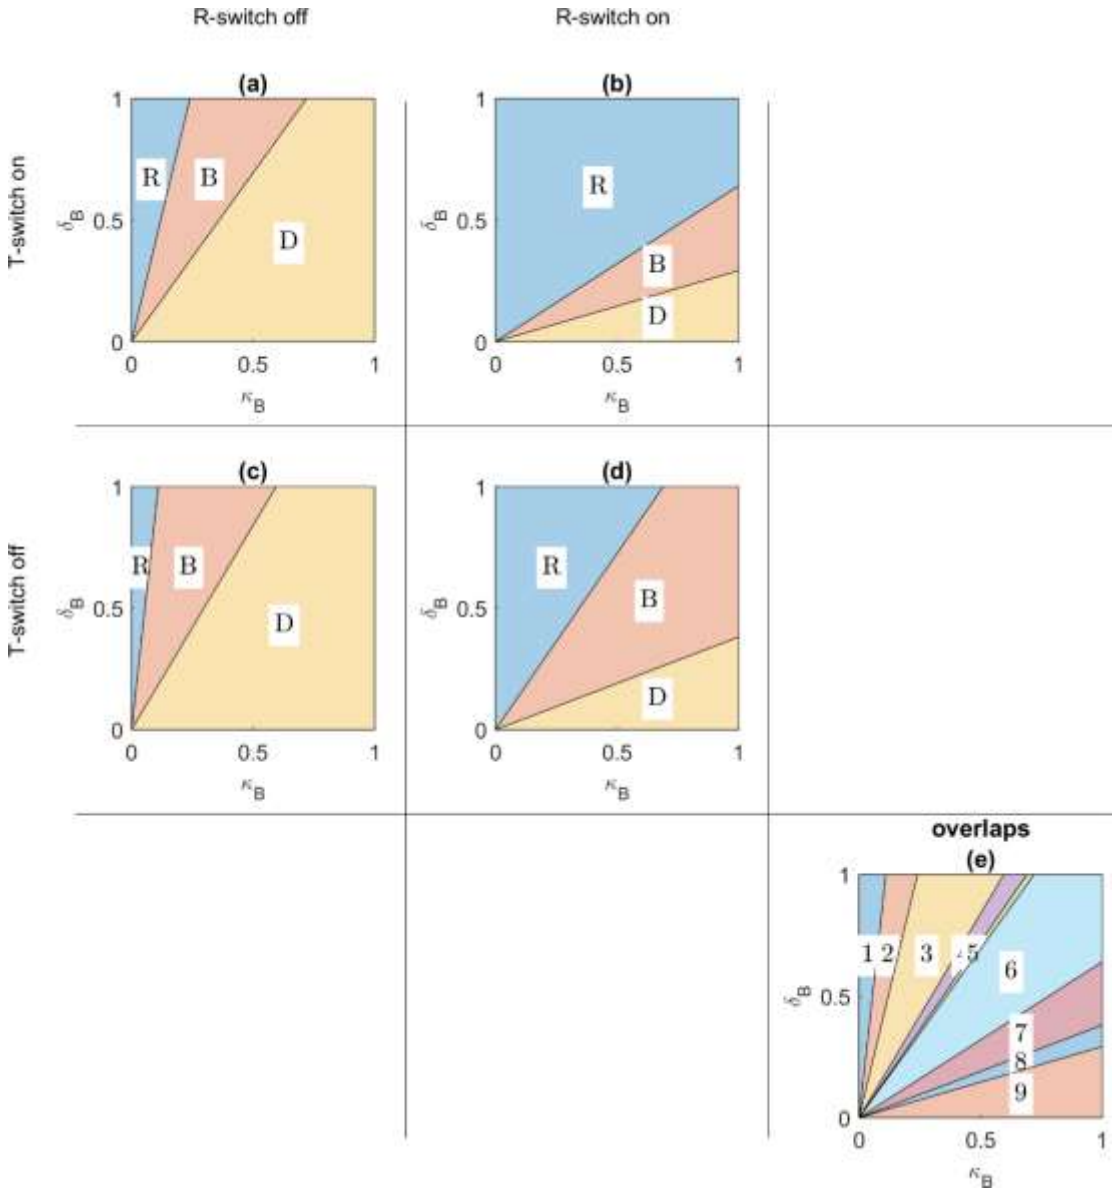


Figure N.1 Reproduction of Figure 5.16, showing how the RK-switch behaviour changes depending on the state of the T- and G-switch.

To check that the switch in question actually exhibits the predicted for each of the switch configurations specified in each region, we simulate the model directly using the technique described

in Appendix M. Once we have the numerically calculated trace of the model's dynamics at a particular location on the bifurcation diagram, we can proceed to work out if the switch's behaviour is as expected. Remember, there are only four possible behaviours: R (recovery), B (bistability), Os (oscillation) and D (damage), so we just have to check which of these four is produced.

Numerically checking a switch's behaviour requires at two numerical simulations, where we set the state of the switch being investigated to 0 and 1. After a long period of numerically integrated time, we check how the switch is behaving. For the R, B and D behaviours, it's very simple to check if they're occurring. If both simulations result in the switch being in the 0 (off) state, the behaviour is indeed R. If both simulations result in the 1 (on) state, the behaviour is D. If the simulation that began in the 0 state also ends in the 0 state, while the simulation that began the 1 state also ends in the 1 state, the behaviour is B.

The Os behaviour is a little tricker to look for because it is not stationary. We have to check how the switch is changing over time. We do this by checking the final few switch configuration the model goes through at the end of the simulation.

The steps described in table N.1 allow us to check the behaviour of a single location in the bifurcation parameter plane. Clearly, even if we find that single location to behave as expected, it doesn't mean we can expect the whole region to behave correctly. For this reason, we apply a mesh of points to every region in each bifurcation diagram, ensuring we check at least 10 locations in each polygonal region.

This method was useful in the development process of the bifurcation analysis described in section 5.2, and picked up a few errors in my code over the years. Every single bifurcation diagram found in this thesis has passed this numerical check process with at least one sample point within each region.

Table N.1 Steps to numerically check the switch behaviours.

Note: This algorithm checks the behaviour of one switch of the hybrid dynamical system. We denote this switch with a capital letter, while the other switches not under scrutiny are denoted by lowercase letters. In the numbered steps below, we write the example initial conditions for checking the behaviour of the RK-switch, however this same series of steps can be used for the T- and G-switches.

1. Set parameter values.

Take the parameter values θ which correspond to a particular point within the parameter region.

2. Set the initial conditions.

Set the switch configuration such that the other switches (i.e. the other two for which we're not checking the behaviour) match the state specified in the diagram.

For the switch we're investigating, we set one initial condition as off, and another initial condition as on. This allows us to see how the switch's dynamics evolve from each of these configurations.

Set the initial conditions of the continuous-valued variables (P_0, B_0, A_0, D_0) such that the initial conditions allow the switches to be in their specified states. For example, if the RK-switch is off, we cannot allow $P_0 > P^+$. A practical way to set the initial condition is to specify the "other" boundary; for example, if the RK-switch is off, set $P = P^-$ since this is in the admissible region by definition.

We end up with two sets of initial conditions. We denote the switch state from the two simulations with different superscripts.

For the RK-switch off configuration: $(P_0^0, B_0^0, A_0^0, D_0^0)$, $C_0^0 = (0, t, g)$
For the RK-switch on configuration: $(P_0^1, B_0^1, A_0^1, D_0^1)$, $C_0^1 = (1, t, g)$

3. Numerically integrate.

Numerically integrate from both initial conditions and record the switch configuration over time. The method and code to do this are described in Appendix M. Set the end time such as it is large relative to the timescale of this model so that we get to see a few switch transitions in our simulation.

We denote the nth switch state with a subscript.
From the simulation starting with the switch off, the sequence is written:

$$C_0^0, C_1^0, \dots, C_{n-2}^0, C_{n-1}^0, C_n^0$$

For the simulation starting with the switch on:

$$C_0^1, C_1^1, \dots, C_{m-2}^1, C_{m-1}^1, C_m^1$$

Note, in the notation above, we're assuming the two simulations produce n and m switch configuration, respectively.

4. Check the final three switch configurations.

For both numerical integrations, examine the final three switch configuration from each simulation $C_{n-2}^1, C_{n-1}^1, C_n^1$ and $C_{n-2}^1, C_{n-1}^1, C_n^1$

The following four situations tell us which behaviour was produced.

Results from the simulations		Behaviour found
$C_{n-2}^0 = C_{n-1}^0 = C_n^0 = 0$	$C_{m-2}^1 = C_{m-1}^1 = C_m^1 = 0$	Recovery
$C_{n-2}^0 = C_{n-1}^0 = C_n^0 = 0$	$C_{m-2}^1 = C_{m-1}^1 = C_m^1 = 1$	Bistable
$C_{n-2}^0 = C_{n-1}^0 = C_n^0 = 1$	$C_{m-2}^1 = C_{m-1}^1 = C_m^1 = 1$	Damage
$C_{n-2}^0 = C_n^0$ and $C_{n-1}^0 \neq C_n^0$	$C_{m-2}^0 = C_m^0$ and $C_{m-1}^0 \neq C_m^0$	Oscillation

Appendix O Calculating the shape of the basins of attraction

Note that this appendix describes a partially formed idea. The method proposed here for calculating the separatrices of the basins of attraction is intended as merely a starting point for further development.

We can calculate the shape of this separatrix in the state space by integrating backwards in time from the points of intersection of the relevant switch thresholds. In the case illustrated in Figure, we can take a set of points on the intersection line, where $B = B^+$ and $D = D^-$, then sample some values for the remaining continuous-valued state variables P and A, and backwards differentiate the ODE from this point using an ODE solver such as MATLAB's ode45.

For each starting point (P, B^+, A, D^-), the ODE solver will return a series of points on a line on the surface of the separatrix. Figure O.1 illustrates the separatrix marked as a curved black line with red chevrons to show the direction in which we calculate the solution (i.e. backwards, away from the focal point marked with a *). In this figure, we're showing a 2D slice of the 4D state space, and the separatrix is represented by a 1D line in this 2D slice. In general, in the triple-switch model the separatrix is actually a plane through the 4D continuous-valued state space.

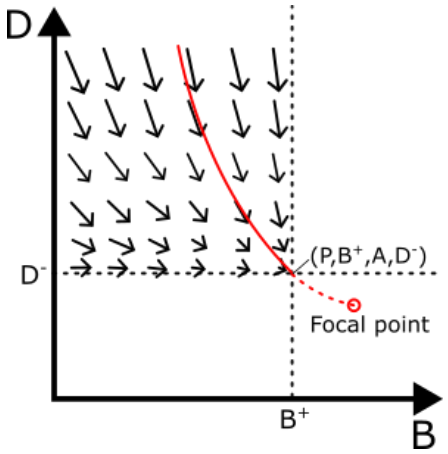


Figure O.1 projected slice of the separatrix, calculated by backwards differentiating from the point (P, B^+, A, D^-).

By repeatedly sampling the separatrix in this way for different values of P and A, we begin to build a picture of its shape in 4D space (though obviously we can't easily visualise this), and the relative sizes of the regions of the state space which lead to each possible transition can be straightforwardly calculated by taking the volume of the regions separated by the separatrix.

So how does this basin of attraction idea connect with the transition trees we've been drawing? We use the example we built up throughout Chapter 6, and reproduce its transition table here as table O.1 so you don't have to flip back to see it. This transition table showed us four switch state configurations which feature multiple outward transitions. These interesting switch state configurations are RK T G , RK T G , RK T G, and RK T G.

Table O.1 Transition table for the example problem used throughout Chapter 6. (reproduction of Table 6.12 to save you having to flip backwards to see it). The circled numbered correspond to those labelling the switch transition on figure O.2.

Row number	RK-switch state	T-switch state	G-switch state	**Predicted transition**		
				RK-switch	T-switch	G-switch
1	0	0	0	-	-	-
2	0	0	1	-	-	1→0
3	0	1	0	-	1→0	-
4	0	1	1	-	1→0 ④	1→0 ③
5	1	0	0	1→0 ①	0→1 ②	-
6	1	0	1	1→0 ⑦	0→1 ⑨	1→0 ⑧
7	1	1	0	1→0	-	-
8	1	1	1	1→0 ⑤	-	1→0 ⑥

This transition table produces the transition network illustrated in figure O.2. Around this transition table, we have sketched examples of the basins of attraction for each (the curves are intended for illustration only, they don't represent a calculated result).

For the switch configurations with just two possible outward transitions ,  and , the basins of attraction may be calculated by integrating backwards with respect to time from the point where the two relevant switch thresholds meet. We illustrate the resulting separatrix by a black line through the appropriate 2D slice of the continuous-valued state space. Each of the basins of attraction separated by these separatrices is labelled with a number corresponding to the transition marked in the transition network.

Transitions from  are marked as ① and ②.

Transitions from  are labelled as ③ and ④.

Transitions from  are labelled as ⑤ and ⑥.

The switch state configuration  has three possible outward transitions. Thus, there are in general three basins of attraction, one for each possible transition.

The imagined separatrix is illustrated at the bottom left of figure O.2, with the three basins of attraction corresponding to transitions ⑦, ⑧ and ⑨ labelled. These three basins form contiguous parts of a cube in the state space, bounded by the switch thresholds and the maximum values we have imposed in our illustration of the state space.

The separatrix separating these three regions resembles three two-dimensional planes, each protruding from a line of intersection of the switch thresholds. These three planes meet at a line which curves through the cube. This illustration simplifies the concept somewhat by omitting the state variable A. In reality for our model, the basins of attraction are actually 4D regions of the state space, and the separatrix a plane through the 4D continuous-valued state space, but clearly these can't be drawn on paper.

To calculate these separatrix planes, we can use the same approach as for the two-switch case, except that we now have many more start points for our backwards-in-time integrations. We integrate

backwards in time as before from the points of intersection of the relevant switch thresholds, except that we must consider the intersections of three switch threshold instead of two. We must consider every combination of intersections from these three thresholds, and also the intersection point of all three threshold planes. These intersections are summarised in table O.2.

Table O.2 Combinations of threshold values to be used as initial points for our backward-in-time integration.

$P = P^-$	$D = D^+$	$B = B^+$	**Notes**
✓	✓		Line of points
	✓	✓	Line of points
✓		✓	Line of points
✓	✓	✓	Single point

As we integrate backwards in time from the initial conditions in Table using a numerical solver such as ode45, we obtain a selection of points from the separatrix. It should be possible to infer the shape of the separatrix from these sampled points. In our cartoon example of figure O.2, the resulting basins of attraction are labelled ⑦, ⑧ and ⑨.

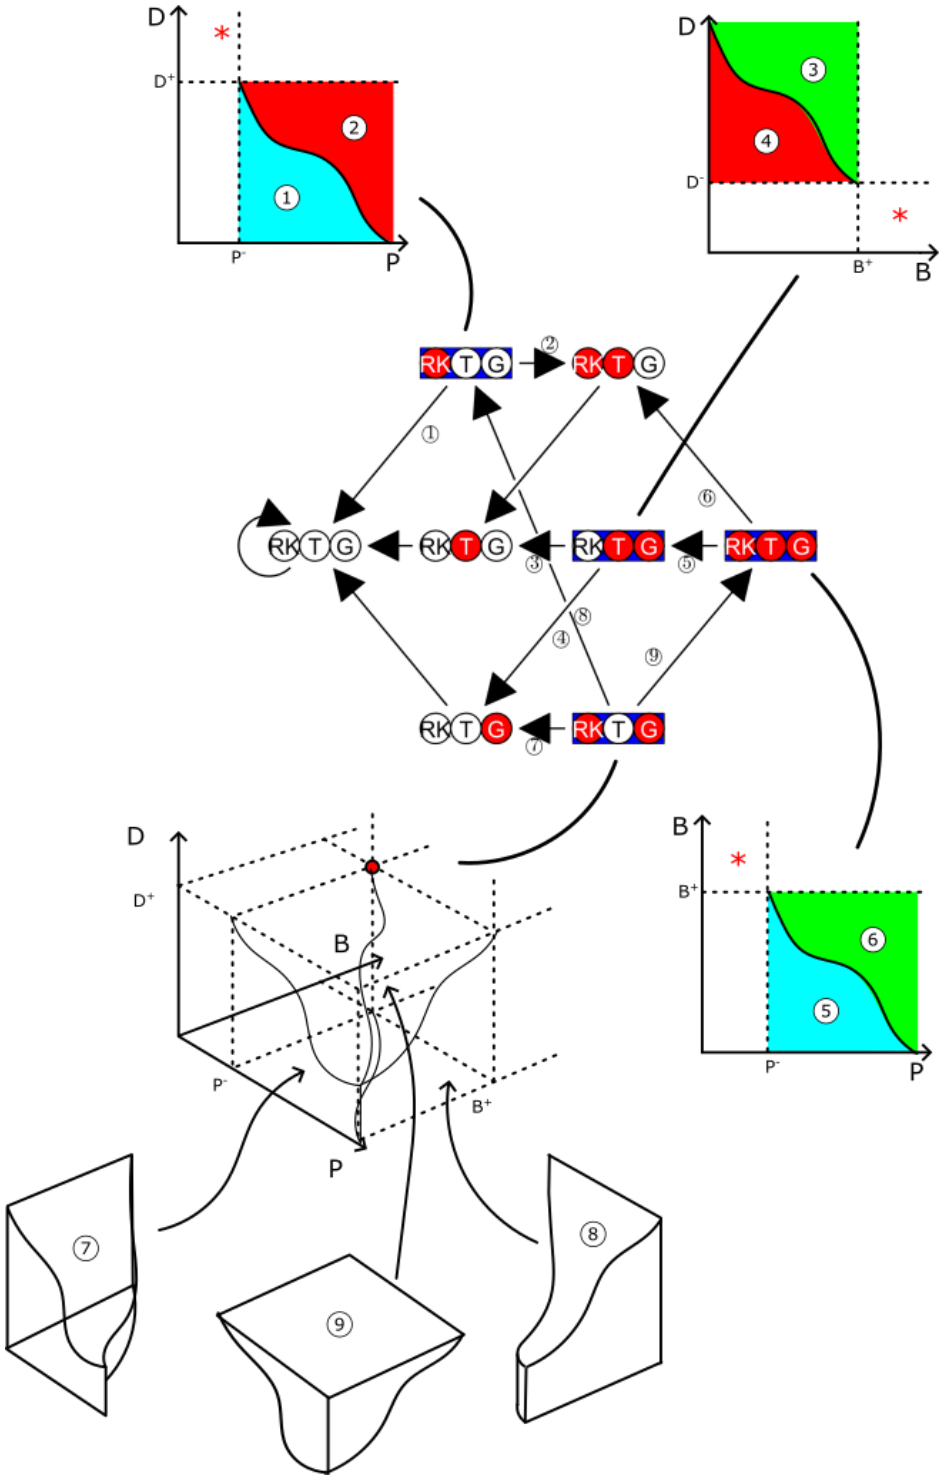


Figure O.2 Basins of attraction for each multiple-transition state in our example network.
Note, the basins of attraction are just sketches rather than numerically calculated results.

The method described above is only a half-formed idea at this point, and we have no practical implementation to in software. With some more thought and some skilled programming, it should in

principle be able to show us what the basins of attraction look like for any switch state transition chosen from a network.

What to do in cases where multiple switch transitions occur in sequence

It is common for race conditions to occur in a sequence such that the winner of one race condition undergoes the transition, leading the system to a new switch state which is itself also a race condition where multiple transitions could take place. We may wish to work out what the basin of attraction looks like for each combination of sequential transitions. In other words, what is the basin of attraction for each possible pathway through the network of transitions?

Thinking about this problem in terms of basins of attraction, it may be thought of as working out how the basins of attraction for a particular switch state configuration are connected to the basins of attraction for another switch state configuration. This allows us to work out the basin of attraction for a particular sequence of transitions through the network, such as those sequences proposed in Figure 6.16. This is essentially a geometrical problem and may be approached using the technique we propose above, applying it in a recursive manner.

First, we start at an absorbing configuration, calculating the basin of attraction in the (P,B,A,D) state space when the switch configurations are in the penultimate configuration leading to the absorbing configuration. Then we apply this same technique to each of the configurations leading to the absorbing configuration, and so on until we reach a switch state configuration with no incoming transition arrows

Let's illustrate this with an example. In figure O.3 (a), we show a particular path through the example network we used in figure O.2. This path passes through two switch configurations which have interesting basins of attraction. Firstly, we consider the transition from the penultimate switch state RK T G to the absorbing configuration RK T G. We've labelled this transition with a ①, and shaded the basin of attraction in the P,D plane in blue in figure O.3 (c).

Next, we apply the same method again to the transition from RK T G to the penultimate state RK T G. This transition is labelled with an ⑧, and the 3D shape of this basin in the (P,B,D) space is drawn in figure O.3 (b)[9], looking a little like a section of cliff-face on the UK coastline. Also shown in figure O.3 (b) is the shaded blue region corresponding to the basin of attraction from transition ①. This blue region is present on the back face of the "cliff" of shape ⑧. By using the curved edge of the blue region as the start points for backwards-in-time integrations of the ODE, we can see how the 3D basin of attraction for transition ⑧ gets split into a region where transition ① follows, and another where transition ② would follow. Effectively, we obtain a separatrix which separates the existing basin of attraction. It's a sort of basin within a basin.

Figure O.3 (d-f) shows the result: a 3D shape corresponding to transition ⑧ followed by transition ①. We've tried to draw this from a few different angles to show how the original cliff shape gets chopped up by the additional separatrix. This shape is labelled with a new symbol "8 & 1" in a circle, with both green and blue accent around it to denote that it represents transition ⑧ followed by transition ①.

This process could be extended to an arbitrarily long pathway, simply chopping up the basin of attraction further at each additional transition we consider. It's an iterative process, working backwards in time from the absorbing configuration to whatever initial state we might wish to

[9] The basin of attraction is a 4D shape, but we draw it here in 3D by neglecting the A dimension for visual clarity.

consider. The resulting basin of attraction shape is likely to be quite complicated, defined by lots of curved faces over which we have a selection of samples obtained through numerical procedures.

The geometrical computation required to turn this cloud of points into a clear image of the basin of attraction is in general quite complicated. Additionally, the method we propose to obtain the sample points by integrating the ODE backwards in time is not ideal, since regions of the state space where the vector field of the ODE in (P,B,A,D) space has high divergence will end up under-sampled as the solutions diverge away from these areas.

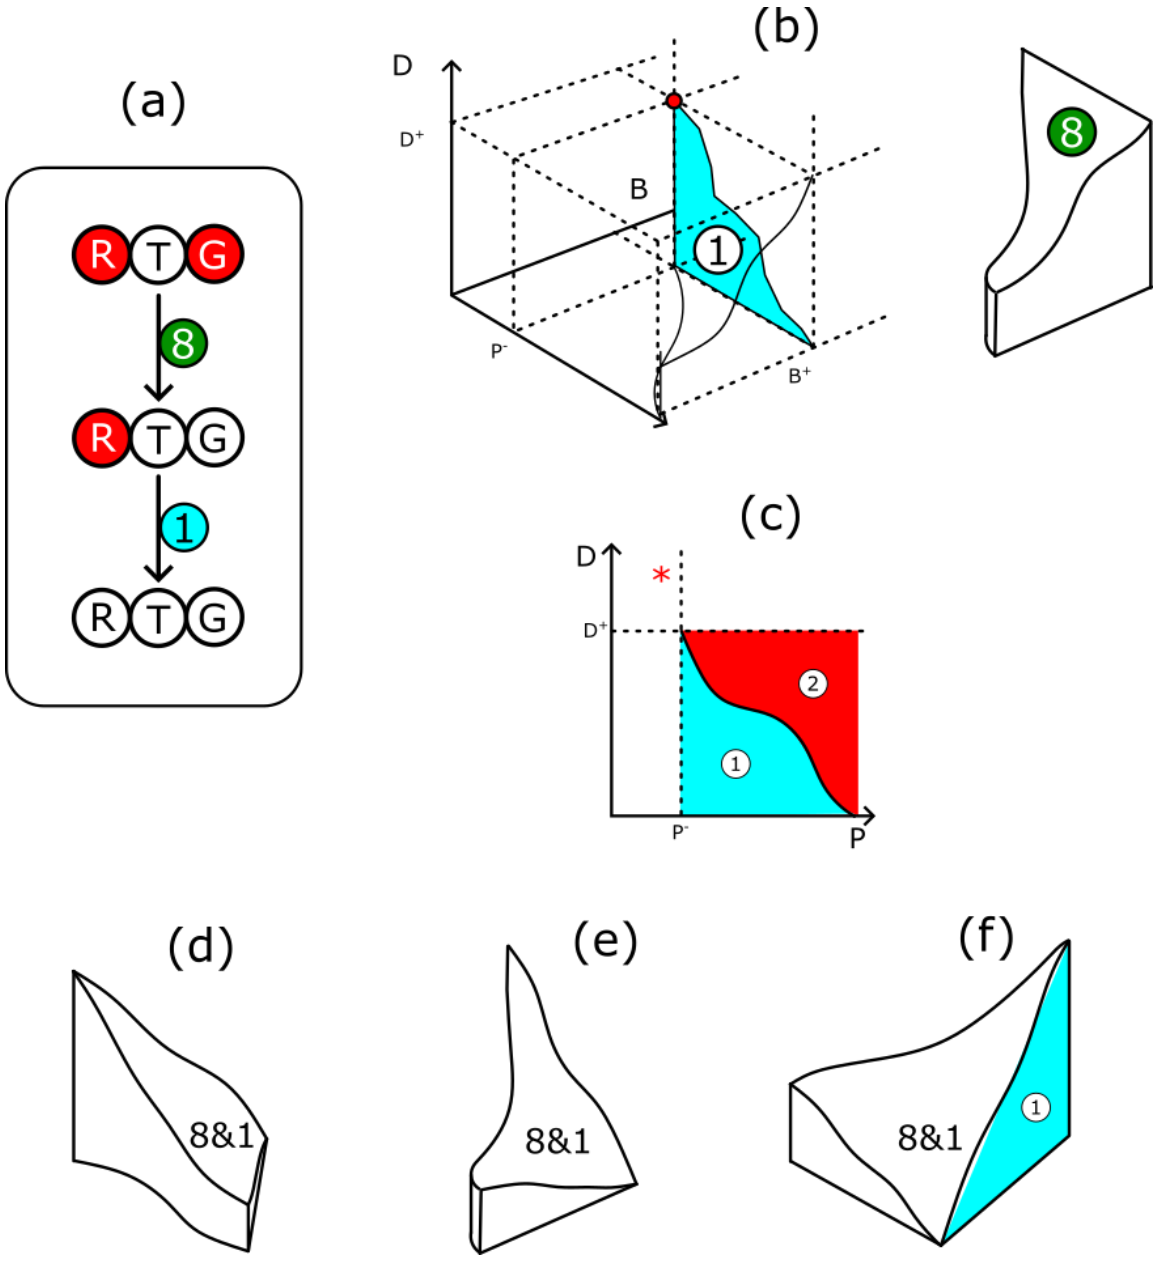


Figure O.3 Calculation of the basin of attraction for a path through the switch transition network

Appendix P Dupilumab therapy for treatment-resistant chronic AD

In cases of AD where symptoms are both persistent and resistant to the most commonly used therapies (emollients and corticosteroids), dupilumab therapy may be prescribed [116].

We showed in section 7.4 how existing treatments such as emollients and corticosteroids apply perturbations to the system from the forcing functions (we suggest $P_{ext}(t)$ and $\delta_{Bext}(t)$ in our ODEs), causing the squirrel to jump around between the branches of the transition network. However, the network itself remains unchanged in structure. In contrast, dupilumab acts on the adaptive immune response, acting on the T-switch itself rather than the ODE, and so is able to perturb the relationship between the switches that form the system as shown in figure P.1.

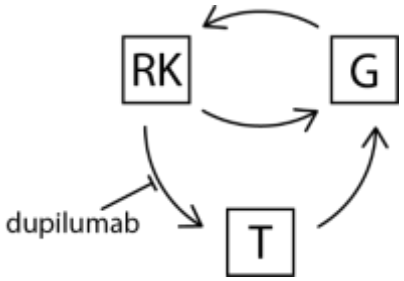


Figure P.1 Relationship between the three switches is interrupted by dupilumab.

Dupilumab is a competitive inhibitor for IL4 receptors, blocking both IL4 and IL13 signalling from Th2 cells [116]. The T-switch part of our triple-switch model considers regulation of the Th2 cell population, and IL4 feedback makes up one part of the feedback loops in that model. In section 3.3, we considered the effect of dupilumab therapy on this Th2 regulation model and found that it brings about an increase in the threshold value of DC activation required to bring about and sustain a Th2 response.

So, what does this increase in the threshold values mean for chronic AD according to our triple-switch model? Let us assume that the threshold value of DC activation level becomes increased beyond the maximum DC level possible in our triple-switch model, as defined in equation P.1. This would mean the dupilumab therapy makes it impossible to induce a Th2 response even if the innate immune response (RK-switch) is activated continuously.

Assumption: $$D^+ > \frac{R_{on}\kappa_D}{\delta_D} \quad \text{(P.1)}$$

This assumption means that all instances where the RK-switch is on would no longer lead to activation of the T-switch. Effectively, we have severed the arrow from RK- to T-switch in figure P.1 due to the inhibitory effect of dupilumab.

We can trace the impact of this assumption on our model results by recalculating what the switch transition trees look like under this new condition.

The result is a change to the structure of the transition trees because some of the switch transitions we previously predicted can no longer occur. We show below that it can change the absorbing configuration of the tree diagram, or even change the nature of the absorbing configuration(s), decreasing the period of the absorbing loop.

The following figures each show a family of transition trees, each sibling in the family shown stacked up vertically. There are two rows. The left row shows each sibling shown with dupilumab treatment applied following our assumption P.1. The right row shows that same sibling without treatment.

We've chosen a few interesting examples where the change brought about by dupilumab therapy is visible. This is by no means a rigorous analysis of our results but is intended as a good start point for potential future research.

Symptoms become less severe under dupilumab therapy.

In figure P.2, we show an example of a family of networks where dupilumab therapy has caused some of the absorbing loops to become shortened because the T-switch no longer becomes activated in the loop. We could interpret this as causing the long-term dynamics of chronic AD to no longer display a Th2-mediated inflammatory nature. Symptoms could be less severe under these conditions but retain their relapsing-remitting nature.

Figure P.3 shows a family of transition trees where the dupilumab treatment causes the model to lose its two absorbing configurations, and replaces them with a single, less severe absorbing configuration. This example demonstrates how the long-term dynamics of chronic AD can be made less severe by dupilumab therapy.

Disease dynamic develops relapsing-remitting nature due to dupilumab therapy.

In figure P.4, we show a family where the dupilumab treatment causes some more dramatic changes to the absorbing configurations and loops. Some members of this family respond to dupilumab treatment by forming absorbing loops where previously there were none. Others change the length of their loops or abolish an absorbing configuration. It is not clear in this case that dupilumab has had a positive effect, as the absorbing loops still contain some quite severe phenotypes (i.e. lots of the switches remain in the "on" state within the absorbing loops, even with dupilumab therapy).

Concluding remarks on Dupilumab therapy

The discussion in this appendix is intended as a start point for research wishing to identify how dupilumab therapy may interfere with the regulatory mechanisms which produce chronic AD.

Dupilumab is one drug from a whole family of new biologic therapies, many of which are still undergoing clinical trials. Dupilumab is just the first example of a successful biologic agent for application to AD, but other monoclonal antibody (mAb) therapies are at various stages of clinical trial [117] so it seems likely that there will be multiple mAb therapies applicable to AD in the near future.

A more rigorous study of the effects of dupilumab on the triple-switch model, and the effects of the various related mAb therapies may be able to shed light on their impact on the mechanisms giving rise to chronic AD, or even provide a method of stratification of patients according to their disease symptoms to ensure the appropriate mAb is applied in each case. It may be able to describe why some mAb therapies have been found to be ineffective for AD while dupilumab has proved effective.

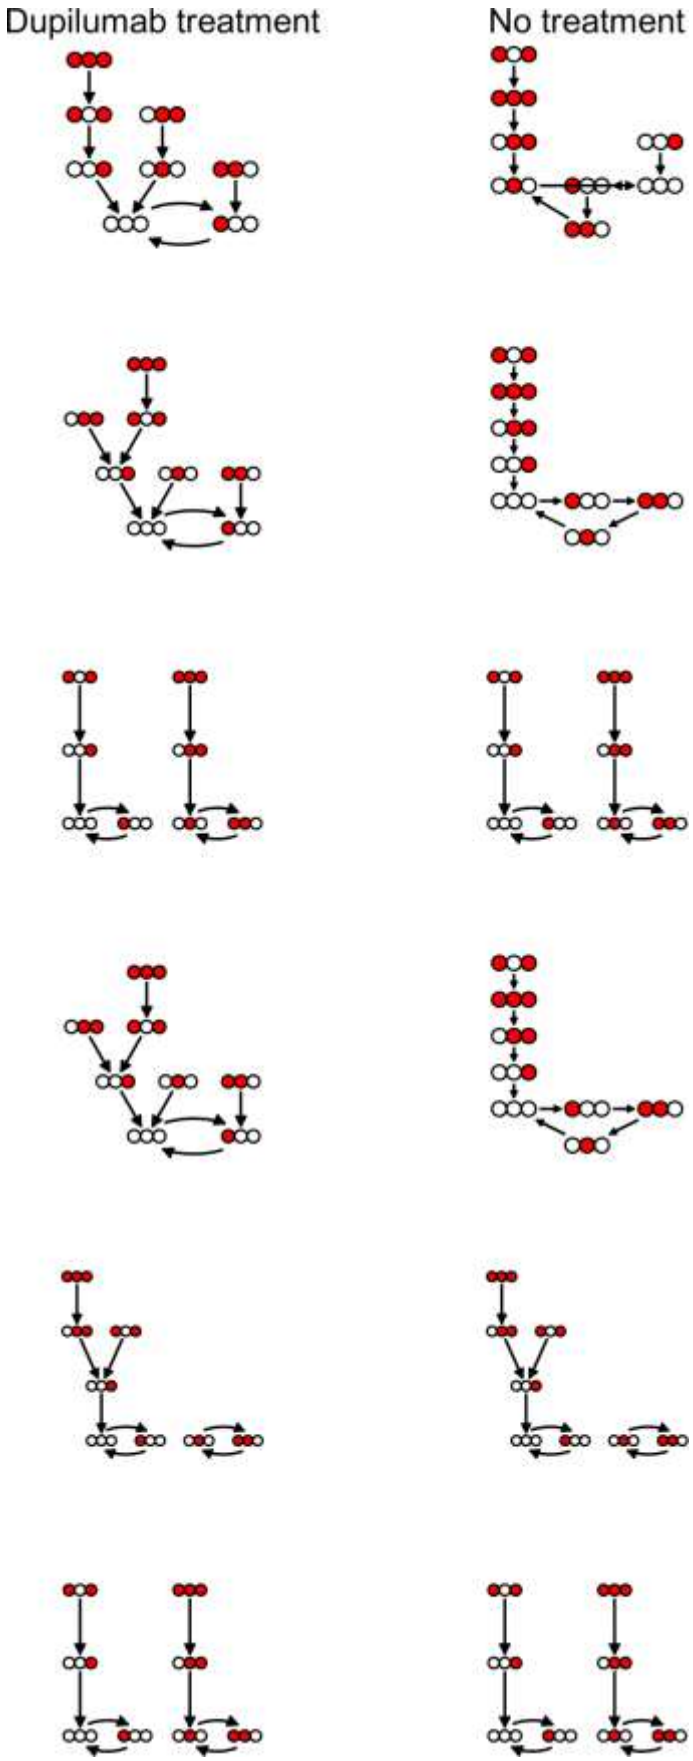


Figure P.2 Th2 inflammatory response is abolished from some of the absorbing configurations in this family. (Figure obtained from parameterisation p3, family 141).

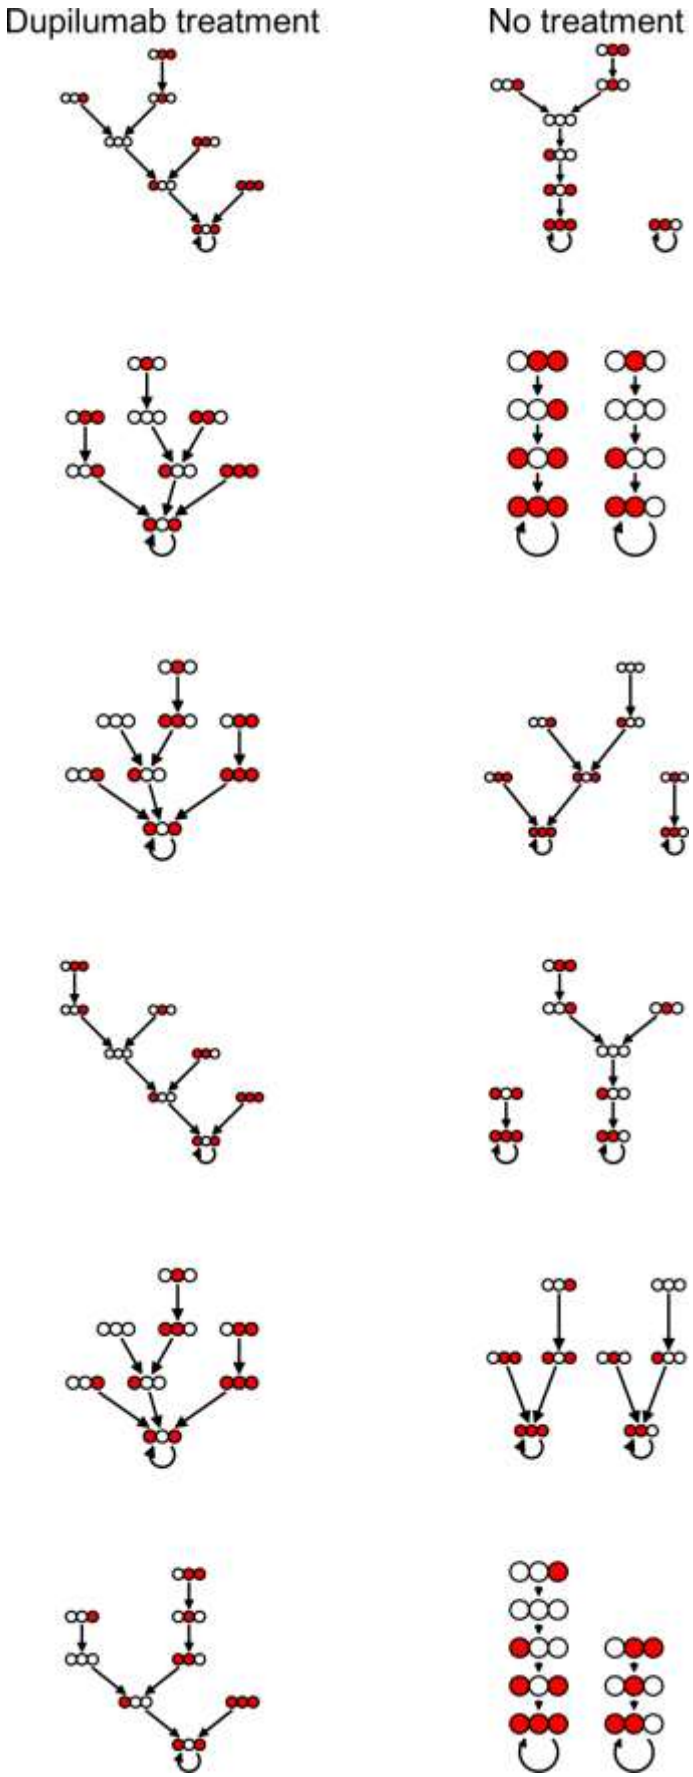


Figure P.3 bistability is abolished in favour of less inflamed absorbing configuration. (Figure obtained from parameterisation p3, family 139)

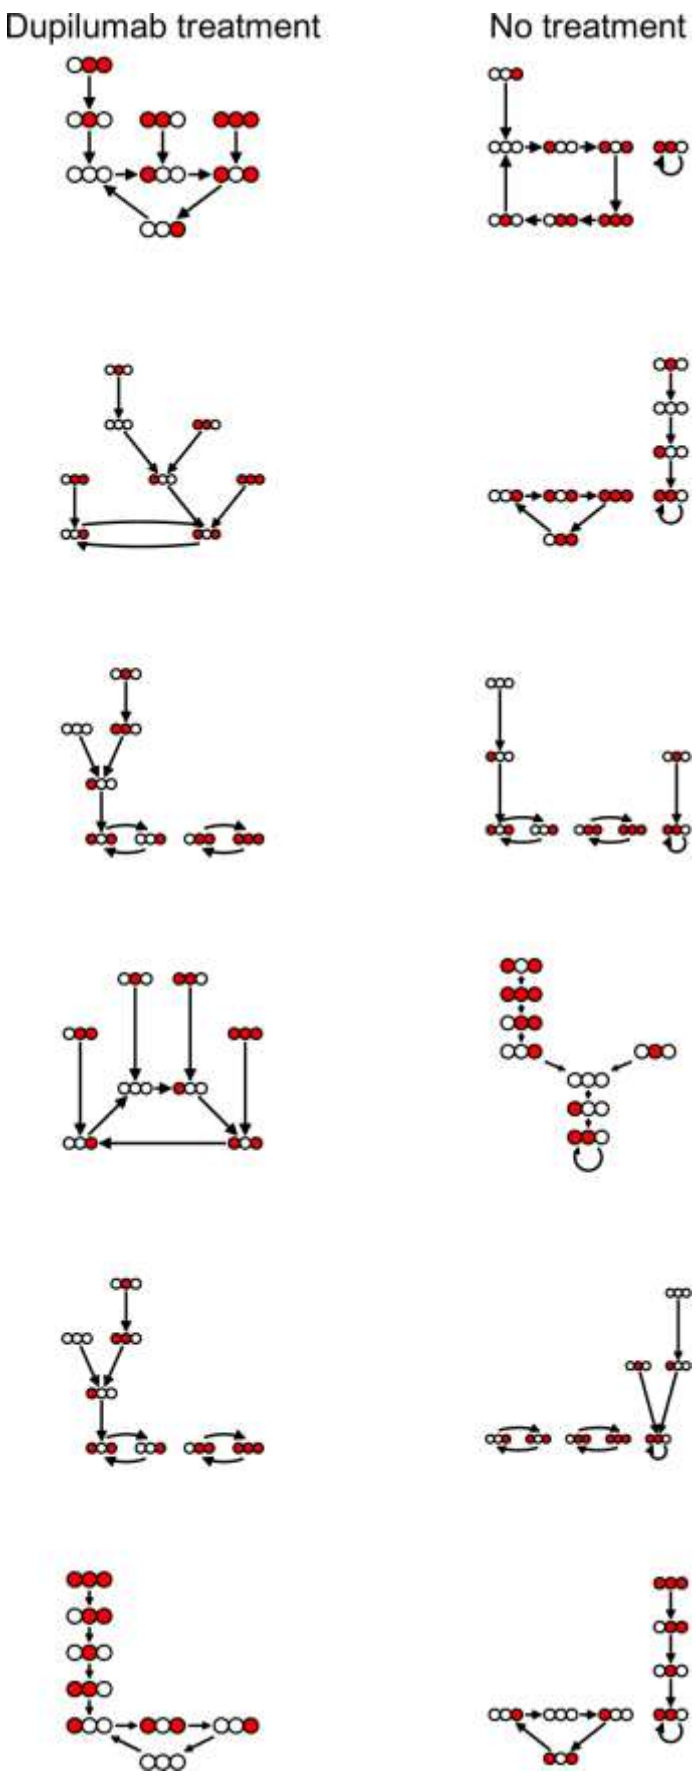


Figure P.4 Some trees develop new loops where previously a single absorbing configuration was present. (Figure obtained from parameterisation p3, family 136).