

UNIVERSITY OF SOUTHAMPTON

Spatial networks: Growth models and dynamical processes

by

Garvin Haslett

Thesis for the degree of Doctor of Philosophy

in the
Faculty of Physical Sciences and Engineering
Electronics and Computer Science

September 2018

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by **Garvin Haslett**

Many inherently spatial systems have been represented using networks. This thesis contributes to the understanding of such networks by investigating the effect of imposing spatial constraints upon both the process network formation and dynamics that occur upon the network.

Degree heterogeneity is a feature of several real world networks. However, edge length is typically constrained in a spatial network, preventing the formation of the high degree nodes that are characteristic of degree heterogeneity. We instead constrain the network to be planar, producing networks that have a scale-free degree distribution. This model turns out to be a variant of random Apollonian growth and a one parameter family of models which incorporates the planar model alongside existing Apollonian models is proposed.

We identify the REDS model as a spatial model that does constrain edge length and exhibits a form of degree heterogeneity, albeit a weaker form than the scale-free distribution. REDS seeks to model social network formation by conceiving its nodes as agents who disburse a personal budget in order to maintain social bonds. We strengthen the model's plausibility by introducing uncertainty into the agents' budget expenditure decisions. The degree heterogeneity that was readily observed in the original model is now recovered only where decisions are subject to high levels of uncertainty.

An evolutionary game is a process that lends itself to simulation upon a spatial network. This is due to the fact that a spatially constrained population is more likely to exhibit network reciprocity known to result in increased levels of co-operation. We find those experiments within existing literature to be unsatisfactory in that network connectivity is assumed a priori. We address this issue by further extending the REDS model such that its nodes play prisoner's dilemma with their network neighbours. The budget with which agents form connections is now earned by accumulating payoffs from the dilemma game. This allows for a network topology that is now endogenous to the model. This model is further distinguished from prior coevolutionary models by its agents' ignorance of the details of their individual strategic interactions.

Contents

Declaration of Authorship	xiii
Acknowledgements	xv
1 Introduction	1
1.1 Structure	4
2 Literature Review	7
2.1 Spatially embedded scale-free network models	7
2.1.1 Link length penalisation	9
2.1.2 Embedding within a lattice	9
2.1.3 Space filling	9
2.2 Spatially embedded social networks	10
2.3 The evolution of co-operation	13
2.3.1 Overview	13
2.3.2 Static Networks	15
2.3.2.1 Cyclical dominance	18
2.3.3 Coevolution	19
2.3.4 Coevolution on lattices	20
2.3.5 General spatial models	21
3 Research aims and contributions	23
3.1 Planar growth	23
3.2 The stochastic REDS model	24
3.3 The evolutionary REDS model	26
4 Planar growth generates scale-free networks	29
4.1 The models and their degree distributions	30
4.1.1 Planar Growth, no planarity and no growth	30
4.1.2 Analysis of the degree distribution	31
4.1.3 Apollonian Planar Growth	33
4.1.4 Statistical test of the power law hypothesis	33
4.1.5 Robustness to variation of m	35
4.2 Analysis of Planar Growth	36
4.2.1 The small world property and assortativity	36
4.2.2 Angle distribution	37
4.3 Planarity relaxation	39
4.4 Comparison of APG with existing Apollonian growth	41

4.4.1	Area weighting and trisection	42
4.5	Summary	44
5	Spatial network growth based on stochastic agent decisions	47
5.1	The REDS model	49
5.1.1	Process	49
5.1.2	Results as reported by Antonioni et al.	50
5.1.3	Implementation of both variants	52
5.1.4	Summary	55
5.2	The stochastic REDS model	55
5.2.1	The stochastic REDS algorithm	56
5.2.2	Results	57
5.3	Stabilisation of the mean degree and the transition region	61
5.3.1	Defining the transition region	61
5.3.2	Characterising the transition region	62
5.3.3	Degree distributions for the stochastic model	64
5.4	Summary	67
6	Social dilemmas on spatial networks	69
6.1	Co-operation on the stochastic REDS model	70
6.1.1	The evolutionary REDS model	70
6.1.2	Results	71
6.1.2.1	Low strategy update temperature	72
6.1.2.1.1	High E	73
6.1.2.1.2	Low E	73
6.1.2.1.3	Intermediate E	73
6.1.2.2	High strategy update temperature	75
6.1.2.3	Low network update temperature	76
6.2	The coevolutionary REDS model	76
6.2.1	The model	76
6.2.2	Coevolutionary dynamics	77
6.2.2.1	Domination	78
6.2.2.2	Collapse	81
6.2.2.2.1	Dynamics prior to the collapse	81
6.2.2.2.2	Dynamics after the collapse	82
6.2.2.3	Cyclic	83
6.2.2.3.1	Rising co-operation	84
6.2.2.3.2	Transmission	84
6.2.2.3.3	Falling co-operation	86
6.3	Parameter variation	87
6.3.1	The effect of synergy and energy upon co-operation	88
6.3.2	The effect of update parameters upon co-operation	89
6.3.3	The effect of update parameters upon cyclic dynamics	92
6.4	Summary	94
7	Conclusions	97
Bibliography		103

List of Figures

4.1	A PG network with $m = 2$ at various stages of its growth. (a) $N = 0$ (b) $N = 50$ (c) $N = 100$ (d) $N = 250$ (e) $N = 500$	31
4.2	(a) Cumulative degree distributions for PG networks of order $N = 10^4$, APG networks of order 5×10^5 , PG-noplanarity networks of order $N = 10^4$ and PG-nogrowth networks of order $N = 2 \times 10^3$. All results averaged over 20 experiments with $m = 2$. The dashed line is the best fit for the APG experiment, a power law with exponent, $\alpha_{\text{APG}} = 2.77 \pm 0.01$. As with all exponents in this chapter, α_{APG} has been estimated using the method of Maximum Likelihood Estimators outlined in Clauset et al. Clauset et al. (2009). (b) Average maximum degree observed for the PG, PG-noplanarity and PG-nogrowth networks. The dashed line is the expected value of the maximum degree for a power law with exponent $\alpha_{m=2} = 2.83 \pm 0.01$, the estimated value of the exponent in the $n = 10^4, m = 2$ case. The dotted line is a plot of the expected maximum degree for a network with an exponential distribution.	32
4.3	(a) Mean characteristic path length for PG, PG-nogrowth and PG-noplanarity networks with varying order N . Note the logarithmic scaling on the x-axis. (b) Clustering for the same networks. Error bars in each image is one standard deviation.	36
4.4	(a) Visualisation of a 0.35×0.35 patch of a $N = 10^4, m = 2$ planar growth network. (b) Visualisation of a 0.01×0.01 patch of a $N = 10^4, m = 2$ PG-noplanarity network. (c) Visualisation of a 0.35×0.35 patch of a $N = 2 \times 10^3, m = 2$ PG-nogrowth network. (d), (e) and (f) the probability mass for the angle, ω , between successive clockwise edges at a node of the network immediately above. Patch sizes in (a), (b) and (c) have been chosen so as to best illustrate the network.	38
4.5	The normalised count of crossings, x_{norm} , observed in experiments with varying χ . Normalisation was observed by dividing x , the number of crossings, by 20,020, the number of edges. The value of $x_{\text{norm}} = 0$ when $\chi = 0$ cannot be represented on logarithmic axes and has been approximated by the value for $\chi = 0.01$	39
4.6	(a) Cumulative degree distributions for networks created using planar growth with a probability χ of accepting edge crossings. The dashed line is the power law with exponent $\alpha_{m=2}$, the best fit for the $\chi = 0.0$ experiment. (b) Average maximum degree observed in the same experiments. Dotted and dashed lines are the same references plotted in figure 4.2 and are fits for the $\chi = 0.0$ cases and $\chi = 1.0$ cases respectively.	40

4.7	(a) Average assortativity observed in PG networks with varying χ . (b) Average clustering observed in the same experiments. Each data point relates to twenty networks grown using $n = 10^4$, $m = 2$. Error bars represent one standard deviation.	41
4.8	A triangulation resulting from an Apollonian growth process represented as a ternary tree. Here a root node, 1, and one subsequent node, 2, have been added to the triangulation, resulting in five faces. This triangulation is represented by a ternary tree on the right hand side. The internal nodes of the tree correspond to the nodes of the triangulation and have been numbered as such. The leaves of the tree correspond to the faces and the colouring scheme indicates this.	42
4.9	Black dots are the degree distribution of a network of order $n = 10^5$ grown using Apollonian Planar Growth with $\beta = 100$. Faces were divided by trisecting in to three equal areas in this version of the model. Empty circles are the degree distribution of an Apollonian network of the same order. Dashed line is a plot the best fit of the exponent, $\alpha_{\beta=100} = 2.85 \pm 0.01$	44
5.1	Mean degree $\langle k \rangle$, clustering coefficient c and assortativity a for non-predictive REDS networks of order $n = 10^4$ and the maximum connection distance $R = 0.05$. These heatmaps are reproductions of those within Antonioni et al. (Antonioni et al., 2014).	51
5.2	Mean degree $\langle k \rangle$, clustering coefficient c and assortativity a for (a) re-implemented non-predictive REDS networks and (b) REDS networks. The value in each heatmap cell has been averaged over 10 instances of networks of order $n = 10^4$ and maximum connection distance, $R = 0.05$	53
5.3	Degree distributions averaged over 15 instances of networks created using the REDS algorithm with $S = 0.8$ and (a) $E = 0.03$, (c) $E = 0.06$ and (d) $E = 0.24$. Values have been normalised so that the area under the histogram is 1.0. Plot (b) is the reverse cumulative distribution of component sizes for the $E = 0.03$ networks. Individual scales have been used for each plot so as to best fit the panel.	54
5.4	Visualisation of two REDS networks. The square of side 0.2 with bottom left corner at (0.4,0.4) has been visualised for REDS networks with $S = 0.8$ and (a) $E = 0.03$, (b) $E = 0.06$	54
5.5	Dynamics of $\langle k \rangle$ over $t = 100$ epochs of 2.5×10^4 iterations for a stochastic REDS model with $T_n = 10^{-2}$, $S = 0.8$, $E = 0.24$	58
5.6	Mean degree $\langle k \rangle_t$, clustering coefficient c_t and assortativity a_t heatmaps for networks grown using the stochastic REDS algorithm for varying temperatures. (a) $T_n = 10^{-8}$ (b) $T_n = 10^{-4}$ (c) $T_n = 10^{-2}$. The value in each heatmap cell has been averaged over 10 instances of network. Each network is of order $n = 10^4$ with maximum connection distance, $R = 0.05$ and has been grown for 100 epochs of 2.5×10^4 iterations. Statistics for an individual network are calculated by averaging over the last 30 epochs of the process.	60
5.7	Heatmaps of $s_{\langle k \rangle}$ for networks grown using the stochastic REDS model for temperatures $T_n = 10^{-8}$, $T_n = 10^{-4}$ and $T_n = 10^{-2}$. Here, $s_{\langle k \rangle}$ is the standard deviation of the observations of $\langle k \rangle$ made over the last 30 epochs of the process. The value of $s_{\langle k \rangle}$ in each heatmap cell has been averaged 10 instances of the model.	62

5.8	Dynamics of $\langle k \rangle$ over $t = 4000$ epochs of 2.5×10^4 iterations for a stochastic REDS model with $T_n = 10^{-4}$, $S = 0.8$, $E = 0.102$	62
5.9	Plots of $\langle k \rangle_{500}$ for networks grown using the stochastic REDS algorithm for $S = 0.8$ and varying E . Temperatures are $T_n = 10^{-8}$ (red), $T_n = 10^{-4}$ (green) and $T_n = 10^{-2}$ (blue). Values of E have been chosen so as to cover the region of high $s_{\langle k \rangle}$ identified in figure 5.7. Error bars represent the standard error.	63
5.10	Degree distributions averaged over 15 instances of networks created using the stochastic REDS algorithm with $T_n = 10^{-2}$, $S = 0.8$ and (a) $E = 0.03$, (c) $E = 0.145$ and (d) $E = 0.24$. Values have been normalised so that the area under the histogram is 1.0. Plot (b) is the reverse cumulative distribution of component sizes for the $E = 0.03$ networks. Individual scales have been used for each plot so as to best fit the panel.	64
5.11	Visualisations of a stochastic REDS network with $S = 0.8$, $E = 0.145$, $T_n = 10^{-2}$. (a) The entire network. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$. (b) The square of side 0.2 with bottom left corner at (0.4, 0.4).	66
6.1	Plots of p_c vs b for the evolutionary REDS model with varying T_n , T_g and $E = 0.03$ (red), $E = 0.145$ (green) and $E = 0.24$ (blue). The parameters $n = 10^4$, $R = 0.05$, $S = 0.8$ are fixed. Each data point is the mean value of p_c over 15 instances of the model after 300 epochs of $t = 2 \times 10^6$ iterations. For each instance, p_c has been averaged over the last 50 epochs of the model run.	72
6.2	Visualisations of a evolutionary REDS network with $S = 0.8$, $E = 0.145$, $T_n = 10^{-2}$, $T_g = 0.2$ and (a) $b = 0.07$ (b) $b = 0.11$. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$	74
6.3	The quantities (i) $\langle k(c) \rangle$, (green) and (ii) $\langle k(d) \rangle$, (red) observed in the evolutionary REDS experiments with $E = 0.145$, $T_n = 10^{-2}$, $T_g = 0.2$ reported in figure 6.1a. The co-operator mean degree, $\langle k(c) \rangle$, is calculated for all co-operator nodes in the model and then normalised by the mean degree of all nodes in the model. Each data point is the the mean value of $\langle k(c) \rangle$ recorded for all the networks generated for the given value of b . A similar calculation, performed on defector nodes, gives $\langle k(d) \rangle$. The grey plot is the p_c curve for $E = 0.145$ from figure 6.1a.	75
6.4	Dynamics of p_c in three canonical scenarios for the coevolutionary REDS model. They are (a) domination, (b) collapse and (c) cyclic. Model parameters in the first two examples are $T_n = 10^{-2}$, $s = 0.8$, $e = 0.24$, $T_g = 0.2$, $w = 10^{-3}$, $b = 0.12$, (a) $l = 0.05$ and (b) $l = 0.5$. Model parameters in the final example are $T_n = 10^{-2}$, $s = 0.8$, $e = 0.03$, $T_g = 2.0$, $w = 10^{-1}$, $l = 0.95$, $b = 0.036$	78
6.5	Node strategies in the domination coevolutionary REDS example (see figure 6.4a). The square of side 0.2 with bottom left corner at (0.1,0.5) has been visualised after epochs (a) 30, (b) 70, (c) 110 and (d) 150. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$	79
6.6	Mean degree of nodes with the (a) co-operate and (b) defect strategy over the lifetime of the coevolutionary REDS domination example.	80

6.7	In (a) we plot the mean number of connections from a defector to (i) other defectors, $\langle k(d)_d \rangle$ (red), and (ii) co-operators $\langle k(d)_c \rangle$ (green). Similarly, in (b) we plot the mean number of connections from a co-operator to (i) other co-operators $\langle k(c)_c \rangle$ (green) and (ii) defectors $\langle k(c)_d \rangle$ (red).	82
6.8	Node strategies in a growing phase of the cyclic example. The entire network has been visualised after epochs (a) 54, (b) 55, (c) 56 and (d) 57. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$. The blue box in (b) highlights a region that is discussed in more detail in the text.	85
6.9	Visualisations of the cyclic example at epochs (a) 54.1, (b) 54.3 and (c) 54.5 within the square of side 0.2 with bottom left corner at (0.8, 0.4). This is the region highlighted by the blue square in figure 6.8b.	86
6.10	Node strategies in a collapsing phase of the cyclic example. The entire network has been visualised after epochs (a) 59, (b) 61, (c) 63 and (d) 65. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$	87
6.11	Plots of p_c vs b for the evolutionary REDS model with $l = 0.05, 0.5, 0.95$. For each l plots are for all combinations of $S = 0.0, 0.8$ and $E = 0.03, 0.24$ (see legend). Parameters $T_n = 10^{-2}, T_g = 0.2, w = 10^{-3}$ are fixed. For each data point, the standard error is below 0.1.	88
6.12	Plots of p_c vs b for the evolutionary REDS model with $w = 10^{-3}$ (blue), 10^{-1} (green) and 10^0 (red). Individual plots represent a $T_n \times T_g$ pair drawn from $T_n = 10^{-4}, 10^{-2}, 0.7$ and $T_g = 0.2, 0.5, 2.0$. Error bars represent the standard error.	90
6.13	Plots of p_n for each of the data points in figure 6.12. Here, is the proportion of experiments that exhibit cyclic dynamics. As in figure 6.12 the curves are colour coded $w = 10^{-1}$ (green) and 10^0 (red). Results for $w = 10^{-3}$ have not been plotted since $p_n = 0$ for all points with this timescale separation.	93

List of Tables

4.1	Estimated exponents of networks of order $n = 10^4$ with varying m . Each exponent, α , is calculated from a batch of twenty experiments that grow a network of order $n = 10^4$ using m as specified in the first row. Following Clauset et al. (2009) we use the standard error, σ , as our estimate of the uncertainty in power laws. For all other estimates of uncertainty in this chapter we use the standard deviation.	35
4.2	Log likelihood ratios of estimated power law distributions compared with other candidate distributions. The log likelihood ratio, \mathcal{R} , and their associated p -values, p , for fits of four alternative distributions compared with the fit of the power law distribution. Statistics were gathered for PG, planar growth with $m = 2$, $N = 10^4$, and APG, Apollonian planar growth with $N = 5 \times 10^5$. Positive values of \mathcal{R} indicate that the powerlaw hypothesis is the preferred model of the data, p is the significance value of the log likelihood ratio.	35
4.3	Estimated exponents of crossing probability networks. Each exponent α is estimated assuming a power law degree distribution for a batch of twenty experiments which grow an $N = 10^4$, $m = 2$ network with the crossing probability χ that is specified in the first row. Standard error, σ , is reported in the third row.	39
4.4	Variation of the degree distribution with area weighting. Estimated values of the exponent of the degree distribution, α , observed for APG networks of order 10^5 with varying area weighting exponent, β . The standard error, σ , in each case is 0.01.	43

Declaration of Authorship

I, Garvin Haslett, declare that this thesis entitled *Spatial embedding and network processes* and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Either none of this work has been published before submission, or parts of this work have been published as:
Planar growth generates scale free networks, Haslett et al., 2016
Planarity as a driver of spatial network structure, Haslett & Brede, 2015.

Signed:

Date:

Acknowledgements

Thanks most of all to my supervisors, Markus Brede and Seth Bullock. Markus in particular can be trusted for insight and clarity when pursuing the essence of an issue.

My girlfriend, Jan Eldekvist, has been an endless reservoir of support.

Alexandry Augustin and Sabin Roman both deserve credit for providing a enjoyable social atmosphere in the lab.

Lastly, thanks to Graeme Bragg for final proof reading of the introduction and conclusions.

Chapter 1

Introduction

The end of the last century saw the establishment and rapid growth of network science. This field has its roots in the mathematical discipline of graph theory and, when combined with increasing computing power, has offered insights into, e.g., neuroscience ([Bullmore and Sporns, 2009](#)), crime ([Duijn et al., 2014](#)), human language ([Stella and Brede, 2015](#)) and development economics ([Brummitt et al., 2017](#)).

A discernible sub-field of network science is that of spatial networks. In the most general formulation, spatial networks are those whose nodes are embedded in a metric space. In this thesis, we confine our interest to networks with nodes that have a position in \mathbb{R}^2 and distances between them are described by the Euclidean metric. Application domains in this regard include social networks ([Liljeros et al., 2001](#); [Jones and Handcock, 2003](#)), cities ([Levinson, 2012](#); [Rui et al., 2013](#); [Gudmundsson and Mohajeri, 2013](#)), electronic circuits ([Miralles et al., 2010](#); [Tan et al., 2014](#)), wireless networks ([Huson and Sen, 1995](#); [Lotker and Peleg, 2010](#)) and leaf venation ([Corson, 2010](#); [Katifori et al., 2010](#)).

In the examples cited, a graph is used as an intuitive representation of the system. Such representations can be allied with network modelling so as to illuminate both the formation of the graph and processes that take place upon its nodes. This thesis uses this approach to contribute to the understanding of spatial networks by introducing three models. The first two of these models speak to the issue of degree heterogeneity, i.e., the extent to which there are multiple scales for the number of connections that a node within the network may have. Of these two models, one of them lends itself to the investigation of social interaction and we subsequently deploy it to address questions in evolutionary theory.

In an influential review of the field, Barthélemy argued that the distinguishing feature of many spatial networks is that a cost is associated with the length of an edge ([Barthélemy, 2011](#)). Where a network model respects this constraint, its nodes are confined to forming connections within their immediate vicinity. Hence, in models with nodes uniformly

distributed upon the plane, hub nodes cannot form. Conversely, a central claim within network science is that many real world networks have a heterogeneous degree distribution (Caldarelli, 2007). Consequently, where spatial models have been used to describe networks with a heterogeneous degree distribution, the majority of studies have either supplied a power law as an input to the model, used an inhomogeneous distribution of nodes upon the plane or exploited topological features of the network (Hayashi, 2006). Furthermore, many spatial networks of interest are almost planar, e.g., a road network, and investigating this property in more depth has therefore been identified as a relevant goal for network science (Newman, 2010). Nonetheless, little, if any attention has been paid in the literature to the issue of planarity. Our first contribution speaks to both of these issues with a model, named planar growth (PG), that conserves planarity and results in a scale-free degree distribution, i.e., the degree distribution follows a power law and nodes therefore exhibit connectivity at all scales. We provide theoretical grounding for this contribution by reformulating planar growth as a process of repeated subdivision of the plane. This new process bears much similarity to Apollonian networks (Doye and Massen, 2005; Andrade Jr et al., 2005) and, hence, is named Apollonian Planar Growth (APG).

We have named social networks as an application domain with a relevant spatial embedding. Despite the social world being a topic with a considerable history in network science, there are few examples of spatial network models in which the nodes act as agents actively engaged in the task of forming social connections. Within this class, the REDS (Reach, Energy, Distance, Synergy) model (Antonioni et al., 2014; Iotti et al., 2017) is distinctive in that its nodes exhibit agency by evaluating their capacity, or lack thereof, to maintain the social connections they intend to make. The degree distribution produced by this model is notable in that it is bi-modal, i.e., we can identify two classes of nodes, some with a small number of connections and the remainder with a large number of connections. Thus, the REDS model exhibits degree heterogeneity across two scales. This aspect of REDS networks is of especial significance since degree heterogeneity is regarded as a defining feature of social networks (Amaral et al., 2000).

Agents making decisions in the REDS model do so with complete information with respect to the costs of edge formation. This can be seen by the fact that the agents neither inadvertently extend themselves beyond their inherent capacity to form social bonds nor subsequently re-evaluate their current bonds. By contrast, the reality of a social network is that agents form and relinquish connections in response to the opportunities and demands of the social world. Our second contribution incorporates this insight into the model by introducing uncertainty into an agent's decision to maintain a connection and to allow for the appraisal of existing connections. This new model, which we name the stochastic REDS model, is a dynamic one in which edges are perpetually updated. This contrasts with the original model which had a clearly defined stopping condition and so produced static networks.

A chief consequence of the stochastic REDS model is that the heterogeneous degree distribution is observed in more restricted circumstances. Whereas, the original model consistently produced networks with the bi-modal distribution within a region of its parameter space, the bi-modal degree distribution only obtains for the stochastic model when the uncertainty is high. Furthermore, where a bi-modal degree distribution is observed within a stochastic REDS network, high degree nodes form along the boundary of the space; a feature not observed in the REDS model.

Where a network is understood to represent interactions between social agents, the nature of those interactions are, in themselves, a topic of interest. These interactions are not always harmonious, to the extent that many social situations entail conflict between the immediate self interest of the individual and the overall welfare of the group (Van Lange et al., 2013). Such interactions are known as social dilemmas, the best known of which is the prisoner’s dilemma. Analytical insight into the prisoner’s dilemma and its evolutionary dynamics can be achieved via game theory and a mathematical model known as the replicator equation. Formalising the dilemma in this manner leads to the conclusion that within a rational population, agents will, under the pressure of natural selection, prioritise self interest ahead of the welfare of the group (Taylor and Jonker, 1978); a finding that is at odds with the numerous examples of co-operation observed in both the biological and social sciences (Wilkinson, 1984; Gomes et al., 2009; Fraser and Bugnyar, 2012).

One approach to resolving this contradiction is to note that the replicator equation assumes that its population is well mixed, i.e., its agents play the game with all other agents. When the population is instead structured such that co-operators are more likely to interact with each other, the pursuit of self interest is relinquished by some, or even all, of the agents. This phenomenon is known as network reciprocity (Nowak, 2006) and has a direct connection with spatial networks in that the structure of interacting co-operators can be realised in an intuitive manner when co-operative agents are co-located (Nowak et al., 1994).

co-operation can be further enhanced in network models by the introduction of coevolutionary mechanisms, i.e., the modification of network connections in tandem with agent strategy updates (Perc and Szolnoki, 2010). However, close reading of this literature reveals that the number of connections is an a priori trait of coevolutionary models, with the final arrangement attained by rewiring the initial network. This approach is counter to the everyday intuition that interactions ensue from those situations that demand them, as opposed to aligning with a fixed global capacity for mutual reaction. Thus, we highlight the need for coevolutionary models in which connectivity is established endogenously.

A further issue is that if a coevolutionary model is to promote the co-operative strategy, the network structure’s development needs to favour the formation of links between

co-operators. In order to achieve this outcome, agents in these models are attributed sufficient cognitive capacity so as to assess details of individual interactions, e.g., knowledge of a neighbour’s strategy. However, many living creatures, e.g. microbes, non-social animals, lack any such capacity. Furthermore, in human social interactions it is often unclear how credit for an outcome should be assigned. Thus, there is a need for co-evolutionary models with more limited cognitive abilities would better represent these situations.

In our final contribution, the coevolutionary REDS model, we address both these issues with a further extension of the stochastic REDS model. In the coevolutionary REDS model, payoffs from the evolutionary game are a resource that agents make use of when forming connections with others, thereby allowing for the endogenous growth of connectivity. Furthermore, since this model is based upon the stochastic REDS model, its agents make connections based upon their incomplete understanding of their capacity to form new connections. Since this capacity is determined by the resource earned by the agent from the evolutionary game, we argue that the coevolutionary REDS agents rely upon a summary of their interactions as opposed to awareness of the actions of individual neighbours.

A key element of any coevolutionary model is the time scale over which network connections are modified relative to the update of node strategy. We investigate the effect of varying this scale for the coevolutionary REDS and identify the regions of the model’s parameter space within which co-operation either dominates the model or collapses entirely. Between these two extremes we further identify regions where the dynamics result in the cyclical variation between cooperation and defection within the model.

1.1 Structure

The material which we have just discussed is divided into five chapters as follows:

Chapter 2: Literature Review. The second chapter reviews relevant the literature and is organised as it pertains to the individual contributions. The issues addressed by each model are placed in context and discussed in more detail. The literature review concludes, in section 3, with a summary of the research aims and motivates the questions addressed in light of the preceding discussion.

Chapter 3: Planar growth. Planar growth is a model in which edges are not allowed to cross. Additionally, this chapter demonstrates that planar growth model can be understood as a variant of random Apollonian growth.

Chapter 4: Stochastic REDS model. In the REDS model agents pay for connections and do so with complete information as to their budget and the cost of an edge. We

recreate this model and then, in the stochastic REDS model, modify the connection decisions so as to make them subject to uncertainty.

Chapter 5: Coevolutionary REDS model. The stochastic REDS model is extended so that its nodes play the prisoner's dilemma with their neighbours. The payoffs from the dilemma game are used as the budget with which the nodes in the stochastic REDS model form connections. Since the process of network formation occurs in tandem with updates to the strategy with which agents play the game we term this model the coevolutionary REDS model.

Chapter 6: Conclusions The final chapter summarises the contributions made by this thesis towards the modelling of spatial networks. Suggestions are made for future work.

Chapter 2

Literature Review

The introductory chapter identified three contributions that this thesis makes to the modelling of spatially embedded networks. These are a growth model that results in a scale free network, a model of uncertain agents forming a social network and a coevolutionary model in which uncertain agents earn the resources required to form a network from a social dilemma's payoffs. In this section we now review the literature relevant to each of these three contributions.

In what follows we use the term graph interchangeably with network. Formally, a graph, Γ , is defined by two sets; the first V , is the nodes. A connection between a pair of nodes, $i, j \in V$ is represented as $(i, j) \in E$, where E is the second set, the edges.

In section 2.1 we review existing spatial models that result in a scale-free network, highlighting that we can only find one example where the model uniformly distributes the nodes upon the plane. In section 2.2 we review examples of social network models that are spatially embedded. Section 2.3 introduces the evolution of co-operation in more detail before reviewing the literature, with a particular focus on the issues addressed in this thesis. Finally, section 3 presents our research aims in the context of this literature.

2.1 Spatially embedded scale-free network models

One of the most fundamental ways in which a network can be characterised is the degree distribution, which is the probability that a randomly chosen node i has k_i connections to other nodes and is typically written as $P(k)$. Many empirical studies have indicated that real world networks are characterised by a heterogeneous degree distribution (Amaral et al., 2000; Mislove et al., 2007; Agler et al., 2016), i.e., the distribution cannot be characterised by a single scale for the average degree of a node. A stronger version of this claim is that the degree distribution is of the form $P(k) \sim k^{-\alpha}$,

with the exponent in the range $2 < \alpha < 3$ (Caldarelli, 2007), the scale-free or power law distribution.

The most influential model to address this observation is that of Barabási and Albert (1999). Networks constructed by this model have been analytically demonstrated to exhibit power law degree distributions with exponent $\alpha = 3$ (Barabási et al., 1999). Barabási & Albert's (BA) model is a growth process which begins with a small number, m_0 , of nodes. Each step of the process adds a new node which forms $m \leq m_0$ connections with the existing nodes. The probability Π that a new node connects with existing node i is given by:

$$\Pi(k_i) = \frac{k_i}{\sum_{j \in V} k_j} \quad (2.1)$$

i.e., the new node attaches preferentially to nodes that already have a high degree.

This is an example of a Yule process (Yule et al., 1925), first advanced in 1925 to explain distributions in taxonomic groups. However, the Yule process is one among many mechanisms that can account for a power laws in nature (Newman, 2005). Consequently, within network science, there has been considerable effort to identify other procedures that can account for the power law degree distribution.

These efforts encounter difficulties when the network is spatially embedded. We illustrate the problem by considering the most fundamental spatial model, the random geometric graph (RGG) (Dall and Christensen, 2002). The RGG is generated by first placing the nodes on a unit square. Pairs of nodes are then connected if they lie within distance R of each other. The resulting degree distribution of the RGG is a Poisson distribution with mean degree $\langle k_{\text{RGG}} \rangle = n\pi R^2$. This outcome is due to the fact that nodes may only make connections within a specified radius and, as such, their maximum degree is limited by the density of nodes. Therefore, we do not see the formation of high degree hub nodes within the RGG. In fact, the procedure only results in a scale-free degree distribution when the nodes are themselves inhomogeneously distributed on the plane (Herrmann et al., 2003; Barnett et al., 2007; Bullock et al., 2010).

Thus, processes that generate spatially embedded, scale-free networks must relax the distance constraint to some extent. Three principal classes of mechanism have been identified which achieve this goal (Hayashi, 2006); (i) link length penalisation, (ii) embedding a scale-free network within a lattice and (iii) space filling. In the following subsections we briefly review them.

2.1.1 Link length penalisation

Two models make up this class; the modulated BA model and the geographical threshold graph. The modulated BA (Manna and Sen, 2002) is a spatially embedded growth model where nodes are sequentially added to the plane, by being placed uniformly at random within a square. The model extends the BA by modifying equation (2.1) so that the probability of node i receiving a connection from the newly added node follows $\Pi(k_i) \sim k_i l^\alpha$. Here, l is the Euclidean distance between the nodes and α is a parameter governing the influence of l on the connectivity. For $-1 < \alpha < 0$ the modulated BA has a degree distribution that is close to a power law. Where $\alpha < -1$, the distribution instead takes the form of the stretched exponential, $P(k) = a \exp(bk^\gamma)$

In the geographical threshold model (Masuda et al., 2005), nodes are again distributed uniformly at random and each node i is assigned a weight w_i . Connections between nodes are formed where the following condition is met:

$$(w_i + w_j)h(r_{ij}) > \theta \quad (2.2)$$

where h is an increasing function of r_{ij} , the distance between i and j , and θ is a threshold. In this model a scale-free network results when either the distribution of the weights or the function h follows a power law.

2.1.2 Embedding within a lattice

Models in this class consist of a lattice with periodic boundary conditions where each node i is assigned an intrinsic degree k_i (Rozenfeld et al., 2002; Ben-Avraham et al., 2003). The networks are constructed by selecting a node i at random and then connecting i to the k_i nearest neighbours on the lattice that (i) lie within the distance $r(k_i) = Ak_i^{1/d}$ and (ii) have not yet reached their intrinsic degree. The degree distribution in this case is precisely that which is assigned to the model during its initialisation and this is typically chosen to be the power law.

2.1.3 Space filling

The final class of spatial scale-free models recursively partition space by adding new nodes to the plane and then connecting them to the existing graph. The first of these models is the Apollonian network (hereafter DAN, the Deterministic Apollonian Network), which was discovered independently by both Doye and Massen (2005) and Andrade Jr et al. (2005). The model begins with three nodes placed upon the plane so as to form an equilateral triangle. In the first step a single node is placed in the middle of this

triangle and connected to the three initial nodes. The second step begins with a node being placed within each of the three triangles created by the first step. To complete the second step, each of these nodes are then connected to the vertices of the triangle that contains them. This process continues in a similar fashion at step t adding 3^{t-1} new nodes within each triangle and then 3^t edges so as to connect those nodes to the vertices of their containing triangle. Analytical treatment of this model reveals that its degree distribution is a power law with exponent $\alpha_{DAN} = 1 + \ln 3 / \ln 2 \approx 2.585$.

A stochastic variant of the DAN is the random Apollonian network (RAN) (Zhou et al., 2005). This model is initialised in the same way as the DAN and now, at each step, a single triangle is selected. Similarly to the DAN a new node is placed within this triangle and connected to its (the triangle's) vertices. The exponent of the degree distribution for the RAN is $\alpha_{RAN} = 3.0$.

Another Apollonian model of note is the evolutionary Apollonian network (EAN) (Zhang et al., 2006; Kolossváry et al., 2013). This model is also initialised in the same manner as the DAN. At step t each triangular face present in the EAN is, with probability q_t , trisected in the same manner as the RAN. This model produces triangulations of the plane with a power law degree distribution. Where $q_t \rightarrow 1$ the EAN approximates the DAN. Alternatively where $q_t \rightarrow 0$ and $\sum_{t=1}^{\infty} q_t \rightarrow \infty$ the EAN can be demonstrated to exhibit the same degree distribution and clustering coefficient as the RAN.

The final example in this class is the model of Mukherjee and Manna (2006). Here, new nodes are placed uniformly at random in the space. The new node is then connected to a random end of the nearest existing edge. The resulting degree distribution is a power law with exponent $\alpha = 3$.

2.2 Spatially embedded social networks

Social networks consist of nodes that represent people (or animals) and connections representing the social bonds between them. It can be credibly argued that this is the field of network science with the eldest pedigree, beginning with Moreno's sociograms in the 1930s (Moreno, 1934). Subsequently, social networks were a topic of extensive interest during network science's initial growth in the early 21st century. A consensus was formed during this period that social networks exhibit several statistical characteristics of interest. These were: high clustering, i.e., that a large number of triangles were present in the network; positive assortativity, i.e., that high (resp. low) degree nodes tend to link with other high (resp. low) degree nodes; community structure, i.e., that distinct cliques of highly interconnected nodes can be identified; the small world property, i.e., that the number of network hops between nodes remains small as network size grows and positively skewed degree distributions, i.e., that there exist a significant number of high

degree nodes in the network (Amaral et al., 2000; Newman, 2002; Jones and Handcock, 2003; Newman and Park, 2003; Toivonen et al., 2006).

Concomitant with research into the structure of these networks has been a modelling effort to gain insight into their formation. In many cases these models describe networks whose edges solely represent relationships between nodes, entirely neglecting the spatial separation between them. There are many reasons to follow this approach, for example, the aim may simply be to establish an efficiently reproducible network with many of the empirical features so as to allow for the investigation of dynamic processes upon that network (Toivonen et al., 2006). Alternatively, one might wish to understand how a psychological variable affects the network's structure, e.g., the agent's sense of identity (Qiao et al., 2014); in such cases ignoring the spatial embedding may be a desirable feature of the model so as not to confound the variable of interest. Other examples have concentrated upon online social networks, e.g., (Kumar et al., 2010), in which case geographical distance is considered to be less of a factor. Furthermore, as section 2.1 highlighted, difficulties that are specific to spatial models infringe upon attempts to model a network that has a heterogeneous degree distribution.

Nonetheless it remains the fact that face to face contact is a mainstay of social interaction. This assertion is supported by empirical studies of phone call data which indicate that people call those that are nearby to them more often (Lambotte et al., 2008; Mok et al., 2010; Phithakkitnukoon et al., 2012). If we consider frequency of calls as an indicator of the closeness of a social tie, these observations confirm the primacy of spatially constrained, person to person contact. Consequently, there is a need for models with an explicit spatial embedding. The literature that attends to this need is, at present, not large and we will review it in the remainder of this chapter.

One of network science's most fundamental models, the Watts-Strogatz (WS) model (Watts and Strogatz, 1998) has been interpreted as a social network model; the reason being that WS networks can exhibit the high clustering and short path length associated with social networks. The spatial component of the WS is of the most basic form, N nodes are arranged in a one dimensional space with periodic boundary conditions. Each node i is then connected to k nodes, chosen to be the $k/2$ nodes on either side of i . This initial construction results in a network that is highly clustered. To complete the model, each node i is considered in turn. Each edge ij that connects i to its $k/2$ neighbours is rewired with probability β to a node k . This node k is chosen uniformly at random from all the other nodes such that there are no self loops and the edge k does not duplicate an existing edge. For high β a network that is nearly random results. However, for low values of β the network's clustered structure is preserved while the rewiring introduces edges that act as shortcuts through the network. These shortcuts allow for the low β networks to exhibit short path length alongside the high clustering.

The spatial component of the WS model is simplistic, reflecting its specific research goal. A more sophisticated, early attempt to model the spatial embedding of social networks was that of [Boguñá et al. \(2004\)](#), where the probability of a connection between two individuals is inversely proportional to the distance between them. Boguñá et al. were able to show that their model exhibited high clustering, positive assortativity and community structure. [Wong et al. \(2006\)](#), in a related study, connected nodes with increased probability if the distance between the nodes is less than a specified threshold. This model created networks with short path lengths, skewed degree distributions and community structure.

Both Boguñá et al. and Wong et al. are considered examples of nodal attribute models (NAMs) ([Toivonen et al., 2009](#)). In a NAM each aspect of a node's social life is quantified by a real number. The node is then positioned within a linear space whose axes represent the aspects. These models are spatial in the sense that the node's attributes determine its location within a social space rather than a geographical one. These networks are thus appropriate for capturing homophily ([McPherson et al., 2001](#)), i.e., the tendency for social agents to link with agents which share similar characteristics. Both of these models are helpful in illuminating the influence of a spatial embedding upon the overall structure of the network. However, they have little to say about the decisions of individual social actors, by definition a key goal of any attempt to understand the social world. More promising in this regard are models which take an agent based view of the network's construction. We now turn to examples of these.

[zu Erbach-Schoenberg et al. \(2014\)](#) introduced a spatial model with weighted connections that represent the strength of the social interaction between the nodes. The model was a dynamical one where, at each time step, agents invite those network neighbours that lie within connection range R to a social gathering. Agents attending the same gathering had the strength of their connection increased; conversely, the strength of the connection depleted where agents did not meet at a subsequent time step. The resulting networks from this model exhibited high clustering, positive assortativity, community structure and short overall path length. Furthermore, these properties were demonstrated to be stable, even though the underlying structure of the network was being rewired in time.

[Antonioni and Tomassini \(2017\)](#) described a growing network model where nodes were placed, one at a time, uniformly at random, upon a unit square. With probability α a new node formed connections with the network following the preferential attachment rule, otherwise the node formed connections with the nearest Euclidean neighbours on the square. For α as low as 0.1, Antonioni and Tomassini demonstrated that this model produces networks with all the features associated with social networks, i.e., high clustering, positive assortativity, community structure, short overall path length and positively skewed degree distributions.

None of the models discussed in this section consider the capability of their nodes to form connections. This concept was central to the REDS model of Antonioni et al. ([Antonioni et al., 2014](#)); a social model in which friendships require effort to maintain. This model followed the RGG, in allowing its nodes to connect to those spatial neighbours within a specified distance. However, nodes in the REDS model were allocated an energy budget, E , with which connections were paid for. The cost of these connections was proportional to the distance between nodes. A second intuition underpinning the REDS model was that it is easier to maintain friendships with friends of friends. Another parameter named synergy, S , controlled the extent to which shared neighbours reduced edge costs. The process by which the REDS model constructs its network is that pairs of nodes are selected and, if they can both pay for it from their respective energy budgets, the edge between them is added.

It was found that this model produced three types of network, characterised by their degree distributions. High S, E regions of the parameter space recovered the RGG network. Low S, E regions of the parameter space produced low connectivity networks. Finally, between these two regions were networks that exhibited a form of degree heterogeneity in that their degree distribution was a bi-modal one, consisting of some nodes with high, RGG-like connectivity and others with low connectivity. Positive assortativity and high clustering was also recorded for both the high connectivity and bi-modal networks.

2.3 The evolution of co-operation

This section reviews literature pertaining to the evolution of co-operation. We first introduce the problem in a more thorough manner than was presented in the introduction in section [2.3.1](#). We then outline some key findings for networks that have a fixed structure in section [2.3.2](#). Following on from this we look at coevolution; for non-spatial networks in section [2.3.3](#) and for lattices in section [2.3.4](#). Finally, we review models that have a more general spatial structure than lattices in section [2.3.5](#).

2.3.1 Overview

In the introduction we discussed the possibility of formalising a social dilemma using game theory. In their most general form, games represent interactions between an arbitrary number of players, each of which has incomplete information about the game and can choose from one of an arbitrary number of strategies. However, within the literature on the evolution of co-operation, social dilemmas are commonly presented as symmetric, two player games, where the participants have complete information and, consequently, we concentrate on this form. Games of this sort can be represented by matrices with the following arrangement:

$$M = \begin{matrix} & \begin{matrix} C & D \end{matrix} \\ \begin{matrix} C \\ D \end{matrix} & \begin{pmatrix} R & S \\ T & P \end{pmatrix} \end{matrix} \quad (2.3)$$

Entries in the matrix represent the payoff for the first player. Rows of the matrix identify payoffs when the first player chooses to co-operate (C) or defect (D). Conversely, columns of the matrix identify payoffs based on the second player's strategy. So, e.g., if player 1 co-operates and player 2 defects, the payoffs will be S for player 1 and T for player 2. The naming convention for the entries in the matrix come from the terms reward R , sucker payoff S , temptation T and punishment P .

It is standard to scale the parameterisation of the game by setting $R = 1$, $P = 0$ and choosing $-1 \leq S \leq 1$ and $0 \leq T \leq 2$ (Perc and Szolnoki, 2010). Where S and T are chosen such that $T > R > P > S$ the matrix describes the prisoner's dilemma. If $R > T > P > S$ we recover the stag hunt game, a dilemma in which the incentive is for players to correlate their strategy. The snowdrift game is described by $T > R > S > P$ where, instead, the incentive is for players to anti-correlate their strategy. It is also possible to choose the parameters such that $R > T, S > P$, known as the harmony game since the incentive is for both players to co-operate.

In what follows we parameterise the temptation to defect as b . The reason for this separate parameter is that a common investigative strategy is to define a model and then test how well it supports co-operation as the temptation to defect increases. This experimental setup is achieved by setting $T = b$.

Within the prisoner's dilemma, when both players defect, neither player can improve their payoff by choosing co-operate instead. A combination of strategies that meets this criterion is known within game theory as a pure strategy Nash equilibrium. A crucial feature of PD is that the defect-defect strategy profile is the only pure strategy Nash equilibrium within the game, thus PD models a situation where a rational player will ultimately choose to defect, despite the significant mutual reward for both players should they choose to co-operate.

The formalisation of social dilemmas can be placed within the context of natural selection via the replicator equation (Taylor and Jonker, 1978):

$$\dot{s}_i = s_i[(W\mathbf{s})_i - \mathbf{s}^T W\mathbf{s}] \quad (2.4)$$

here \mathbf{s} is a vector whose elements s_i represent the frequency with which strategy i is played by the members of the population and W is the payoff matrix for a two player game. In this context \mathbf{s}^T should be interpreted as the transposition of the strategy vector.

The term $(W\mathbf{s})_i$ is the fitness of strategy i while $\mathbf{s}^T W \mathbf{s}$ is the average fitness in the population. Thus, the replicator equation embodies, in mathematical terms, the principle that individuals with strategies that reproduce more effectively spread throughout the population. Methods from dynamical systems can be used to demonstrate that the solutions of equation (2.4) are the Nash equilibria of W and these Nash equilibria are also the equation's fixed points (Hofbauer and Sigmund, 1998; Gintis, 2000). Thus, replicator dynamics for PD, i.e., where $W = M$, indicates that the defect strategy will be favoured by natural selection.

In spite of these pessimistic analytical results, examples of co-operation in nature are widely abundant. To give examples from different species, vampire bats share blood (Wilkinson, 1984), chimpanzees groom each other (Gomes et al., 2009) and ravens intervene in each others conflicts (Fraser and Bugnyar, 2012). The problem of the evolution of co-operation is to account, theoretically, for the observed abundance of co-operation in nature. Or, more succinctly, to answer the question: under what conditions can a co-operative strategy prevail within a Darwinian framework?

It is also the case that both the snowdrift and the stag hunt games also pose problems for the evolution of co-operation, see e.g., Hauert and Doebeli (2004). However, the results presented above underscore that the prisoner's dilemma is regarded as the strongest form of the problem. For this reason we will focus on the prisoner's dilemma within this thesis.

Numerous mechanisms have been proposed to resolve this problem including network reciprocity (Nowak and May, 1992), repeated interactions (Trivers, 1971; Axelrod and Hamilton, 1981), kin selection (Hamilton, 1964), reputation (Nowak and Sigmund, 1998), group selection (Traulsen and Nowak, 2006), memory (Wang et al., 2006) and teaching (Szolnoki and Perc, 2008). Among these, the introduction highlighted network reciprocity as having a spatial aspect. A brief synopsis of the argument is as follows: evolutionary populations are not well mixed, as is assumed within the replicator equation. Instead, agents within such populations interact with some individuals more than others with the consequence that co-operators can benefit where they play the game amongst themselves. Spatial co-location of agents can serve to promote this structure in their interactions.

It is here that the relevance of networks can be seen since the interaction structure can be described using a graph. We now review relevant literature in this regard.

2.3.2 Static Networks

We begin our survey of network reciprocity by considering simulations that are performed on static networks, i.e., networks with a fixed topology. This is a considerable area of research and, as such, cannot be thoroughly covered here. In this section we will

endeavour to cover key results. We refer the reader to [Roca et al. \(2009\)](#) and [Szabó and Fath \(2007\)](#) for two reviews that are widely used as starting points within the literature.

A fundamental assumption of the replicator equation is that the population in equation [\(2.4\)](#) is considered to be well mixed, i.e., individuals are assumed to play with all others [Roca et al. \(2009\)](#). This is in contrast with the social or biological worlds as commonly understood, where individuals interact in a more structured fashion. A seminal study addressing this issue was that of [Nowak and May \(1992\)](#); the key innovation of which was to simulate a finite population of agents placed within a lattice structure. These agents accumulate the payoffs from playing PD with their lattice neighbours. During each timestep, the model updates each agent's strategy to be that amongst its own and those of its network neighbours that has accumulated the largest payoff.

Nowak & May's model repeats this sequence of playing the PD game and updating the strategies of the nodes. It was found that the co-operative strategy not only survived but was the dominant strategy in simulations where the temptation to defect was $T = b < 1.8$. (By dominant, we mean that co-operation was the majority strategy on the network.) For $b > 2.0$ defection was the dominant strategy. In the range $1.8 < b < 2.0$, both strategies could coexist within the lattice.

The explanation for the dramatic improvement in the co-operators' fortunes observed within Nowak & May's model is that some co-operators now find themselves playing the game solely with other co-operators. Members of these co-operative clusters then outperform neighbouring defectors and ultimately convert them to the co-operate strategy. These dynamics were investigated in more detail by [Wang et al. \(2013\)](#) who identified two distinct phases, the enduring and the expansive. The enduring phase comes at the start of the model and is characterised by a fall in co-operation as defectors easily convert lone co-operators. At the end of this phase those co-operators that remain are grouped together in small clusters that can resist the defection strategy, i.e., the co-operators have successfully endured the defectors' advance. Furthermore these co-operators can now convert defectors in the immediate vicinity of their cluster, leading to an expansion in the number of co-operators. Wang et al. demonstrated that this phenomenon holds for different update mechanisms and lattice structures.

A subsequent critique of Nowak & May's model modified the strategy updating to be asynchronous, i.e., a single node was updated during a timestep rather than the entire network. Under these conditions co-operation was demonstrated to collapse in the model [\(Huberman and Glance, 1993\)](#). Subsequent research has demonstrated conditions under which synchronously updated models support high levels of co-operation [\(Nowak et al., 1994; Lindgren and Nordahl, 1994; Kirchkamp, 2000\)](#). Nonetheless, asynchronous updating now tends to be favoured and is used in all the articles we discuss from this point onwards.

Ease of implementation underscores the appeal of lattice models as the basis for an evolutionary simulation. However, they are less satisfactory when considered in terms of their capacity to capture more nuanced forms of population structure. In particular, it was noted in a prior section that a uniform degree distribution is an unsatisfactory description of a social network. Thus, Santos and Pacheco ([Santos and Pacheco, 2005](#)) investigate what happens where the network has a heterogeneous distribution by simulating PD for agents connected by a Barabási-Albert network. Their conclusion was that co-operation is supported for very high values of the temptation to defect. Specifically, they found that co-operation dominated the outcome for all $b \in [1, 2]$. An explanation of this phenomenon was offered by Santos et al. ([Santos et al., 2006a](#)). The essential idea being that where a co-operative hub connects to a defector hub, the co-operation strategy will ultimately prevail. The mechanisms underlying this phenomenon are twofold; firstly, highly connected defectors find their payoff augmented by the presence of co-operators in their immediate neighbourhood. Consequently, the defector hubs rapidly convert co-operators in their network neighbourhood, which, ironically, leads to a drop in the payoff for the hub. Secondly, high connectivity co-operators, while penalised by connections to defectors, nonetheless maintain enough connections to other co-operators so as to heighten their payoff and ultimately convert their neighbourhood to co-operators. Thus defector hubs are vulnerable to conversion while co-operator hubs are robust with respect to their strategy. Note that this account assumes that hubs are directly connected and Santos et al. report that, where these simulations are performed upon a heterogeneous network from which such connections are removed, the increased co-operation disappears.

Masuda introduces a participation cost, h into simulations upon Barabási-Albert networks and random regular graphs ([Masuda, 2007](#)). The participation cost represents the cost of maintaining a link and is introduced into the model by subtracting h from each of the four entries in M . Masuda then runs the evolutionary simulation upon a random regular graph and a scale-free network. In the case of the random regular graph, each node has the same degree and therefore each one accumulates the same total participation cost. Therefore results upon the random regular graph are unchanged by the introduction of h . By contrast, the heterogeneous degree of the scale-free networks entails that the high degree nodes accumulate more participation costs than the low degree nodes. Thus, for $h > 1$, the mechanisms we have just described as supporting co-operation on a scale-free network break down for the reason that high degree co-operators no longer derive a high overall payoff from their neighbours.

An important general principle, theoretically demonstrated by [Ohtsuki et al. \(2006\)](#), is that co-operation is favoured on networks where the mean degree is less than the cost benefit ratio of the co-operate strategy. The authors further conclude that the fewer the number of edges in the network the more likely it is that co-operation will be favoured in an evolutionary simulation upon that network. Numerical support for this proposition

comes from [Tang et al. \(2006\)](#) who studied how variation of average degree affected the uptake of co-operation on various network structures. Their results show that co-operation was favoured for low mean degree on Barabási-Albert, Watts-Strogatz and Erdős-Rényi models.

2.3.2.1 Cyclical dominance

We complete our overview of static networks by considering cyclical dominance. Briefly, these are models in which agents choose from, typically, three strategies. The strategies follow the form of the children's game rock, paper, scissors in that they are enumerated in a cycle and each strategy is inferior to its predecessor. It is this structure of the strategies that is termed cyclical dominance. There is a wide body of research in this area which has been reviewed in [Szolnoki et al. \(2014\)](#). A central theme in this work is that two of the possible outcomes allow for the survival of all the strategies. The first of these outcomes is that the model enters a steady state in which all three strategies coexist in constant proportion to each other. The steady state is maintained despite the fact that the nodes perpetually update their strategy and so this state is often termed a dynamic equilibrium. In the second outcome, known as global synchrony, the model transitions through a sequence of states in which each strategy briefly dominates the model without completely eradicating the others. Examination of this sequence of states reveals that it follows that of the cyclical dominance.

Results such as these speak to the evolution of co-operation in that they offer a possible mechanism by which a strategy, i.e., co-operation, that would normally be dominated by another, i.e., defection, can nonetheless prevail in an evolutionary scenario. A key point here is that cyclical dominance does not have to be explicitly imposed upon the strategies; rather, it can emerge in any evolutionary game where the competing strategies are three or more ([Szolnoki et al., 2014](#)). An example of this is Hauert & Szabó's voluntary participation model ([Hauert and Szabó, 2005](#)), where an agent has the third option of a loner strategy alongside the now familiar co-operation and defection. What the loner strategy means is that the agent chooses not to play PD and instead earns an intermediate payoff σ , such that $P < \sigma < R$. Agents that find themselves playing with loners are forced to adopt the loner strategy for that interaction.

In a well mixed population, the loner strategy dominates the model. However, note that the loner strategy is preferred by defectors that are connected to other defectors, thereby allowing loners to invade a population of defectors. In turn, co-operation will be preferred in a population of loners. Finally the cycle is completed by the fact that defection is favoured over co-operation. Hauert and Szabó found that when voluntary participation was played on lattices, the dynamic equilibrium of three coexisting strategies resulted. When the lattice structure was replaced with a network that allowed for long range connections, the globally synchronous state of perpetual transitions instead emerges.

Under certain dynamical conditions, cyclical interactions are possible even where the number of strategies is limited to two. Szolnoki et al. describe agents that modify their strategy with a probability W that is proportional to $P_i - P_j$, where P_i is the payoff that i gets from the evolutionary game and P_j is the payoff of j , one of i 's neighbours (Szolnoki et al., 2010). However, once i performs a strategy update the probability that it will do so again is capped for a fixed number of steps, H . By doing so this model allows for the existence of “young” C players, i.e., players that have held the co-operative strategy for less than H steps. These young co-operators tend to retain their strategy and, furthermore, can convert “old” D players. In this way Szolnoki et al. describe the following form of cyclic dominance: “old” $D \rightarrow$ “young” $C \rightarrow$ “old” $C \rightarrow$ “young” $D \rightarrow$ “old” D . Szolnoki et al. go on to show that co-operation can survive in this model at high b under this regime.

The appeal of this model is that additional strategies have not been explicitly specified; instead, a form of cyclic dominance arises from two strategies as a consequence of the dynamics. Nonetheless, it stands alone as the sole example of its class. It would therefore be of interest to identify further circumstances in which two strategy games can result in cyclic interactions.

2.3.3 Coevolution

A natural continuation from the static networks paradigm is instead to allow the network topology to evolve in tandem with the strategy update, a program of research known as coevolutionary models. An early example of this sort of model is that of Zimmermann et al. (2004) where nodes on an ER network were allocated either co-operate or defect strategy. In a similar fashion to Nowak and May, nodes on this network play the PD game, accumulate payoffs and then modify their strategy to mimic that of the neighbour with the highest payoff. The coevolutionary aspect of this model is that, in a third step, defectors with links to other defectors can, with probability p , rewire to another node at random. In doing so it was demonstrated that for p as low as 0.01 co-operation could dominate the network for a temptation to defect as high as $b = 2.0$, a noticeable improvement over Nowak and May. This result is instructive in that it demonstrates that a coevolutionary rule with a plausible motivation (in this case that defectors will be dissatisfied with links to other defectors) can boost the chances of co-operators.

In an influential review of coevolutionary models, Perc and Szolnoki (2010) identify timescale separation as an important factor that can affect the outcome of a coevolutionary model. Timescale separation identifies two different time scales, τ_s and τ_n , at which strategy and network topology updates occur. This idea is typically implemented by defining $w = \tau_s/\tau_n$, and at each iteration of the model updating the strategy of a randomly chosen node with probability $(1 + w)^{-1}$; otherwise the network topology is

modified instead. Thus, as $w \rightarrow 0$, we recover a static network model; conversely as $w \rightarrow \infty$, network topology is updated more rapidly than node strategy.

The most straightforward of this kind of example is that of [Santos et al. \(2006b\)](#). In this model co-operators connected to defectors rewire their connection with a probability that is proportional to the difference between payoffs that the two nodes, i.e., a rewiring event is more likely if the co-operator has a higher payoff. When such rewiring events occur, the co-operator links to one of the neighbours of the defector. Santos et al. varied the rate of node strategy updates relative to the rate of network topology updates. In doing so they found a critical value of time separation, w_c , above which co-operation always dominated the network.

[Fu et al.\(a\)](#) is a study in a similar vein; the innovation being that network nodes now have a reputation ([Fu et al., 2008](#)). By this, they mean that the agents know the past strategies of their nearest and next nearest network neighbours, those nodes that tend to co-operate are awarded a high reputation score. The initial network is one where all nodes have the same degree but are then randomly connected to each other. When chosen for a topology update, a node on this network drops the connection to the partner with the lowest reputation and, with probability p instead connects to the next nearest neighbour with the highest reputation. With probability $1 - p$ the node instead connects to a randomly chosen member of the population. [Fu et al.](#)'s chief finding is that the introduction of reputation significantly reduces w_c , the critical value of the timescale separation at which co-operation dominates the model.

It is not necessarily the case that coevolutionary models best support co-operation above some threshold in the timescale separation. For example, [Fu et al.\(b\)](#) ([Fu et al., 2009](#)) follow the model of Santos et al. above, the difference being is that only co-operators may rewire connections they have with defectors. Here, it is an optimal, rather than a critical value, of the time scale separation at which co-operation is best supported.

2.3.4 Coevolution on lattices

Thus far, this review has covered one example of a spatially embedded network, the lattice model of [Nowak and May \(1992\)](#). We now examine two issues that have been explored using lattice models: adaptive migration and reputation.

Within an adaptive migration model, agents move to more favourable regions of the network in order to boost their payoff. The foundational example here is [Helbing and Yu \(2009\)](#), whose model was a lattice containing more sites than agents. Each agent was given knowledge of the payoff that would ensue by moving to an empty site within a Moore neighbourhood of range M . The agents could then choose to move to a better performing site than their current location. Over time, co-operative agents in this model form clusters with defectors appearing on the borders. Once this happens a similar

dynamic to the Nowak & May's model takes hold, allowing for co-operation to convert defectors.

Further research into adaptive migration has established that models whose agents have only local knowledge also support co-operation (Jiang et al., 2010; Ichinose et al., 2013). In these models the agent has awareness of n_D , the number of defectors in its network neighbourhood. The agent then moves to a new location with probability n_D/k , where k is the agent's degree. These models indicate that, even with limited information, the co-operate strategy is well supported.

Khoo et al. (2016) apply reputation to spatially embedded nodes. In this model, the initial network is a lattice with connections between all pairs of nodes that are within Manhattan distance two of each other. As well as reputation, the nodes are assigned a preference to either rewire locally, i.e., to other nodes within a Manhattan distance of 3, or globally, i.e., to any other node in the model.

Khoo et al. demonstrated that, for low temptation to defect, co-operation dominated the model. The structure of the network was a small world network with high degree hubs. Further experiments revealed that an increase in the rate of rewiring increased the capacity of co-operation to overcome the temptation to defect as the co-operators could more easily rewire so as to become hub nodes.

2.3.5 General spatial models

We complete this section of the review by considering spatial models which allow for a more nuanced topology than lattice models.

Yang et al. (2011) play PD upon the Apollonian network. Since the Apollonian network is known to be scale-free it is perhaps unsurprising, in light of Santos and Pacheco's findings for the Barabási-Albert model, that co-operation dominates the model for high values of b . More specifically, they find that p_c , the proportion of co-operators in the equilibrium state, reliably reaches one of the two states $p_c = 0$ or $p_c = 1$. They then measure ϕ_c , the probability that the model reaches the all co-operators state ($p_c = 1$). Here, they find that when half the nodes in the model are initialised with the co-operation strategy, $\phi_c = 1.0$ when the temptation to defect reaches values as high as $b = 2.4$.

Buesser and Tomassini (2012) study the evolution of co-operation on three different spatial network models; the scale-free spatial network (SFSN), the Apollonian network and a modified RGG (mRGG). The SFSN is a form of intrinsic degree model where nodes are temporarily placed in a lattice. Each node is assigned a degree from a sequence k_1, k_2, \dots, k_n which follows the scale-free distribution. Edge (i, j) is added to node i by performing a random walk, starting at i , upon the lattice, until a node j , whose connectivity is at present less than k_j , has been located.

The construction of the mRGG is identical to that of the RGG except that each node is assigned a unique radius of connection R_i . If $R_i + R_j < D_{ij}$ then the two nodes are connected, here D_{ij} is again the Euclidean distance between the nodes i and j . Buesser & Tomassini then assign 1/16 of the nodes a radius of connection r and the remaining nodes a null radius, thereby inducing a network structure of a relatively small number of high degree hubs which connect to low degree, typically $k = 1$, nodes.

Buesser & Tomassini simulate 2 player evolutionary games on all three of these network structures and compare their results to those obtained from an RGG. co-operation, as measured by the fraction of co-operators in the equilibrium state, is found to be improved by all three of these network structures when compared to equivalent results on the RGG.

[Li et al. \(2013\)](#) investigate how the partner switching dynamics reported by Fu et al.(a) are affected when the network is spatially embedded. The initial network is a lattice and, as in the original study, nodes rewire from low reputation neighbours to higher reputation next nearest neighbours during the network update steps. However, in this model a node i may also rewire to any node j within within a specified Euclidean distance of i . The choice to rewire to a nearest network neighbour of the low reputation node or a spatial neighbour is governed by a probability. Li et al. found that co-operation was most favoured the rewiring was more likely to be directed towards the next nearest neighbours.

In a further experiment Li et al. awarded a reputation score $S_j^i(t)$ to each node j that can potentially be the target of a rewiring event at time t :

$$S_j^i(t) = R_j(t)/D_{ij} \quad (2.5)$$

here $R_j(t)$ is the reputation of node j at time t and D_{ij} is the Euclidean distance between i and j .

In this version of the model an agent then rewrites to the agent with the highest score. Under this condition, Li et al. found that co-operation dominates the model for $b \leq 2.0$ in scenarios where the network updating is rapid. However, perhaps unsurprisingly given equation (2.5), a high percentage of spatial neighbours were favoured.

Chapter 3

Research aims and contributions

In this chapter we motivate the questions that are addressed in the remainder of this thesis.

3.1 Planar growth

The discussion in section 2.1 revolved around the three mechanisms that have been identified to produce spatially embedded networks with a scale-free degree distributions, these were: (i) link length penalisation, (ii) embedding in a lattice and (iii) space filling. For class (i) two models were identified, the modulated BA and the geographical threshold model. In the modulated BA, the probability of an existing node receiving a new connection followed the form $\Pi(k_i) \sim k_i l^\alpha$. As such, this model is one whose connectivity is determined as much by its existing network topology as its spatial embedding. The geographical threshold model pre-assumed the power law as input, a feature that it shares with the models in class (ii).

In class (iii), several of these examples were of Apollonian growth processes. The underlying theme of these models is that they choose a face of the existing triangulation of the plane and then split it into three. Since it is the distributions of triangles in these models that determines where the new node is placed, it follows that these models are a further example of an inhomogeneous distribution of points upon the plane. Alternatively, one can interpret the Apollonian processes as ones that do not attribute any explicit point in space to their nodes, instead relying upon subtle aspects of the network topology to generate the power law.

Therefore, amongst all existing spatial models, that of [Mukherjee and Manna \(2006\)](#) results in a scale-free distribution when nodes are distributed uniformly at random on the plane and, further, this model does not assume any other input that has a power law form. Moreover, this model does not exploit network topology in the construction

process, indicating that it is purely spatial features of the model that contribute to the power law degree distribution. Since this is the only model that exhibits a scale-free degree distribution under these conditions, we contend that the discovery of further models with these properties is a significant finding within the discipline.

A further issue that we explore in this thesis is the extent to which the model we propose respects *planarity*: the property of a spatial network having edges that do not cross. While, for example, sexual contact networks may be embedded in space, they need not respect planarity. By contrast, the layout of a microchip must be planar since conductor lines may not cross without creating a junction. Transport networks tend to be nearly planar (a relatively small number of bridges and tunnels allow edges to cross without creating a junction vertex). Planarity is also a consideration in the construction of infrastructure such as wireless networks (Cairns et al., 2013). Despite the relevance of planarity considerations across a wide range of network domains, the role of planarity in network formation is an under-represented issue in the spatial networks literature (Newman, 2010).

Our first contribution is to present two related mechanisms, planar growth and Apollonian planar growth. The first, planar growth, constructs planar networks by the most basic mechanism we can think of. Specifically, where a new node is added to the network, it may not form a connection that crosses an existing one. In this way, we investigate the issue of planarity as it pertains to spatial networks and demonstrate that this model results in a network with a scale-free degree distribution. We then propose Apollonian planar growth as a means to explain why planar growth produces this result. This second mechanism is a reformulation of planar growth as an Apollonian growth process. It is well known that Apollonian networks are scale-free and, hence, the second model provides theoretical support for our findings.

3.2 The stochastic REDS model

The planar growth model reliably produces networks that have a heterogeneous degree distribution. Nonetheless, both planar growth and Mukherjee and Manna's model permit the addition of edges of any length. Thus, it remains a relevant endeavour to establish network models that maintain constraints upon edge length while exhibiting a heterogeneous degree distribution. In section 2.2 we discussed social networks, in particular noting that these are real world networks that exhibit degree heterogeneity while being made up of spatially constrained links. In light of the latter of these two properties we therefore reviewed models of spatially embedded social networks. Amongst these, we identified the REDS model (Antonioni et al., 2014) as being the sole example of a spatially embedded social network model that exhibited a form of degree heterogeneity.

The REDS model is one in which agents consider the merits of any potential connections. This idea is realised by assigning to each node an energy budget, E . Edges in the REDS model are only added if doing so does not cause either of their end nodes to exceed their budgets. We interpret this last feature as implying that nodes in the REDS model are partially rational agents; in the sense that they have complete knowledge of both their current budget and the cost of connections that they form. The idea that the model's agents evince certainty is further accentuated by the fact that once an edge is formed it becomes a permanent feature of the network. By contrast, the nature of real world social interaction is an exploratory process, in which connections are formed and then retained where they transpire to be beneficial, or at least sustainable, in the long run. In this light, we seek to further develop the REDS model so as to better reflect this social reality.

Our second contribution addresses this concern by alternatively formulating the REDS model so that its agents make connection decisions that are subject to uncertainty. As a result, these agents may form connections that exceed their energy budget and, therefore, we allow them to review and, where necessary, rescind an existing connection. In this regard our model bears some similarity to that of blinking networks (Belykh et al., 2005), where edges are added to a small world model only to be removed after a period of time, τ . We achieve this by making the decision to create, or destroy, an edge subject to a probability and for this reason we term our model the stochastic REDS model.

At the outset of this section we have discussed the significance of the heterogeneous degree distribution observed for the REDS model. More specifically a bi-modal degree distribution is observed within a small region of the REDS model's parameter space. We find that within the parameter space of the stochastic REDS model, the overall connectivity of the network in a stable state tends to match that of the original REDS model. However, we find that where the stochastic REDS model is parameterised with low uncertainty, a configuration that one would intuitively expect to closely follow the original REDS model, we do not observe the bi-modal degree distribution. Instead, the bi-modal degree distribution is now only observed in regimes where there is a high amount of uncertainty in the connection decisions.

Upon further inspection, we find a distinct spatial distribution of nodes in examples of the stochastic REDS model with a bi-modal degree distribution. Here we find that high degree nodes appear at the model's (spatial) boundaries while low degree nodes form a large cluster in the centre of the model. This is in contrast to the original REDS model, where patches of high degree nodes are separated by regions of low degree nodes.

In summary, our findings support a claim that the stochastic REDS model also exhibits a weak form of degree heterogeneity. However, the need for high uncertainty to attain this outcome and the spatial distribution of the nodes should be highlighted.

3.3 The evolutionary REDS model

Section 2.3 introduced the problem of the evolution of co-operation. Initially, we discussed networks with a static structure and the principle that clusters of connected co-operators can form in populations that are not well mixed. Connections within these clusters are mutually beneficial and the participants resist attempts to convert them to the defect strategy. This phenomenon of co-operators linking with each other is known as network reciprocity and can result in the entire population adopting the co-operate strategy (Nowak, 2006). Spatial networks are particularly relevant here, since co-location co-operators is a straightforward mechanism by which this structure of interaction can be achieved.

We then highlighted the further development of coevolutionary models (Perc and Szolnoki, 2010); the defining feature of which is that network structure evolves alongside changes to strategy. A key motivation for this development is that links that form in social situations typically change over time and so, the ongoing alteration of the connectivity is deemed to be more realistic. In these models the prospects for the co-operative strategy can be further boosted if the alterations to the network structure tend to favour greater interconnectivity between co-operators.

In all models discussed in section 2.3, the total number of edges is held at some fixed amount. Where topological modification occurs, it consists of rewiring undesirable connections. Implicit in this scheme is the assumption that edges are paid for from a shared budget and that connection costs are therefore invisible to the agent. *Prima facie*, two of the models discussed would seem to challenge this assumption, those of Masuda (2007) and Li et al. (2013). In Masuda's case a fixed cost was applied to each node. However, this was a static network, with the cost of the edge incorporated into the model by an adjustment to the payoff matrix. Ultimately the cost had no impact upon a node's capacity to form connections; a fact that was most apparent in the model's scale-free guise, where the total edges that a node could receive was, in principle, unlimited. Masuda's model then, serves to make explicit the underlying assumption of a shared budget. In Li et al. (2013), equation (2.5) penalises the formation of long edges. Nonetheless, connectivity within this model remains constant over its lifetime. Furthermore, short distances are merely a preference in this model; the longer edges can still be formed, and when they do so there is no additional cost for doing so. Thus, Masuda's model also conforms to the assumption of a shared budget.

To reiterate then, those coevolutionary network models that have been used to explore the evolution of co-operation have an overall connectivity that has been determined *a priori*; an assumption that implies that the model's agents have access to a shared budget with which they pay for their connections. This is an unrealistic scenario that can potentially be addressed by the stochastic REDS model, insofar as one of the model's

key features is that each of its agents maintain connections from a privately held budget. Since evolutionary games model strategic interactions in which agents maximise accumulation of a resource, it would be logical to derive the privately held budget from the payoff from the game. Such a model would extend the plausibility of coevolutionary network models in that connectivity can now arise in response to agent need and capacity, rather than being imposed upon the population.

Our final contribution therefore extends the stochastic REDS model so that its agents play the prisoner’s dilemma with their network neighbours. Doing so imbues the nodes with a state, i.e., the strategy with which they play the game. The agents are rewarded (or penalised) for their performance in the dilemma with more (or less) energy with which they can form edges in the underlying REDS network. In this way the state is intertwined with the network topology in the fashion of an adaptive network ([Gross and Sayama, 2009](#)) and we therefore name the model the coevolutionary REDS network. Since connectivity is determined in this fashion, there is no need to externally impose network structure.

A further issue with the coevolutionary models discussed in section [2.3](#) is that the agents are attributed with knowledge of their network neighbours. This knowledge takes diverse forms, e.g., the neighbourhood ([Santos et al., 2006b](#)), strategy ([Zimmermann et al., 2004](#)) or reputation ([Fu et al., 2008](#)) of a network neighbour. In the case of adaptive migration, agents are ascribed with knowledge of the potential payoff from moving to unoccupied slots of the lattice ([Helbing and Yu, 2009](#)). This knowledge is then applied during the network update phase of the model to guide the agent’s rewiring strategy. As a counterpoint to these assumptions, it is straightforward to imagine organisms that lack such capacity, consider microbes or many animals. Moreover, even human agents can find themselves in situations where assigning credit is difficult; consider salary negotiation, an activity so contentious that its outcome is almost universally confidential. In light of this discussion, a further merit of the coevolutionary REDS model is that it is populated by less cognitively endowed agents. Agents in the coevolutionary REDS network do not update their connectivity on the basis of any understanding of their neighbours’ state; instead they deploy what resources they have available at any given moment.

The coevolutionary REDS network exhibits three canonical forms of dynamics. One where co-operation dominates the model, a second where co-operation collapses entirely and a final example where co-operation rises and falls cyclically. The domination example is significant in that it demonstrates that co-operation can be promoted in a model where agents have no knowledge other than their overall reward from their strategic interactions with their neighbours. We perform a broad range of numerical experiments so as to establish the conditions that exhibit the varying canonical forms of the dynamics. We also note that the cyclic interactions arise from dynamical properties of the two

strategy game and, as such, bear some similarity to the model of [Szolnoki et al. \(2010\)](#) discussed in section [2.3.2.1](#).

Chapter 4

Planar growth generates scale-free networks

In section 3.1 we discussed models of spatial networks that result in a scale-free distribution. There, we identified [Mukherjee and Manna \(2006\)](#) as the sole example that results in a scale-free distribution when (a) nodes are distributed uniformly at random on the plane (b) does not assume an input to the model that has a power law form and (c) does not make use of network topology in the construction process. Furthermore, we identified planarity as an issue that has received little attention in the spatial networks literature.

In this chapter we present two related mechanisms; the first, planar growth (PG), seeks to directly address the impact of a planarity constraint on a network growth process. Briefly, PG incrementally builds a network by placing new nodes at random locations in space and connecting them to other nodes such that planarity is maintained. We introduce it in section 4.1 alongside two reference cases; one of which considers a network that grows in time but does not enforce planarity, while the other considers a network built over a static set of nodes through the addition of planarity-preserving edges. In contrast to the reference cases, PG results in a power law degree distribution and we present evidence to support this claim in section 4.1.2. Further investigation of PG is presented in sections 4.2 and 4.3, where we examine other key network measures and demonstrate the consequences of relaxing planarity, respectively.

The second mechanism is named Apollonian Planar Growth (APG) and is introduced in section 4.1.3 as a reformulation of PG as an Apollonian growth process. Consideration of APG as an object of study in its own right leads to further contributions. Firstly, the APG is inherently spatial; in contrast with the topological character of its precursors, the DAN and the RAN. Secondly, PG can be viewed as a generalisation of Apollonian growth processes to cases where m , the number of connections made when a node is added to the network, is less than 3. In section 4.4, we further develop APG as a single

parameter model, the variation of which tunes the exponent of the network's degree distribution. The deterministic Apollonian network (DAN) and the random Apollonian network (RAN) can then be seen as special cases of APG, with PG intermediate between them. Finally, we conclude this chapter in section 4.5 where we summarise our results.

4.1 The models and their degree distributions

We begin with a description of the models and an analysis of the degree distributions that they produce.

4.1.1 Planar Growth, no planarity and no growth

Planar Growth creates spatially embedded networks with $N + 10$ nodes and average degree $2m$ on a unit Euclidean square that has rigid boundary conditions. We wish to begin the process with a planar network that has nodes distributed uniformly on the plane. To do so ten nodes are placed uniformly at random upon the unit square with $m \times 10$ planar edges between them. As they are added, each node after the first is connected to an existing node; the edge being chosen so as not to violate planarity. Once all ten nodes have been placed, unconnected pairs are then chosen at random and an edge is chosen between them; again subject to the caveat that planarity is always maintained. We continue choosing node pairs until $m \times 10$ edges are added or until all possible node pairs have been tried. The resulting network is accepted irrespective of its final number of edges.

Tests of the procedure over 10,000 realisations show that for $m = 2$ the average degree of the initial network was $k_{ave} = 3.99$, while for $m = 3$ it was $k_{ave} = 4.2$. Despite the results for the $m = 3$ case we retain this method of initialisation since the number of nodes and edges of the initial network is statistically insignificant in comparison to the finished network and the method reliably initialises planar networks.

The algorithm now enters the growing phase where the following steps are repeated N times:

- (1) Place a new node, i , uniformly at random within the square.
- (2) Repeat m times:
 - (2a) Pick a node j where $j \neq i$.
 - (2b) If the straight line ij does not cross an existing edge then add ij otherwise go to 2a.

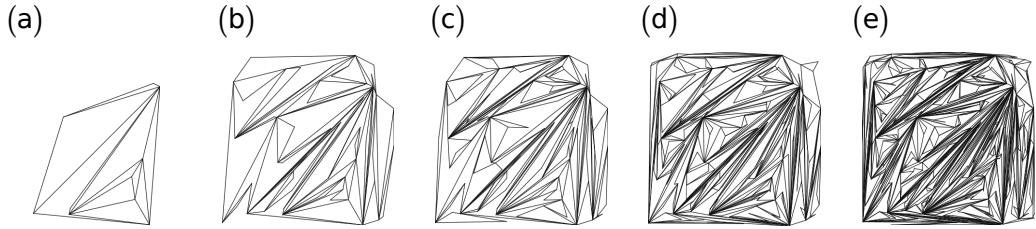


FIGURE 4.1: A PG network with $m = 2$ at various stages of its growth. (a) $N = 0$ (b) $N = 50$ (c) $N = 100$ (d) $N = 250$ (e) $N = 500$.

If step 2 cannot be completed because m valid nodes do not exist then remove node i and any associated edges and repeat step 1.

As reference cases for PG we consider two degenerate variants of the mechanism; one with no planarity constraint, PG-noplanarity, and one with no growth, PG-nogrowth. Choosing these two models allow us to determine what aspects of the PG model are responsible for the outcomes that we see. PG-noplanarity is very similar to PG except that edge connections are always allowed. This scenario is equivalent to the uniform attachment model originally introduced by [Barabási et al. \(1999\)](#), where it was shown to result in networks with an exponential degree distribution. In PG-nogrowth we create a static population of N nodes placed uniformly on the unit square. Pairs of nodes are picked at random and an edge is drawn between them, provided this new edge does not cross an existing one. We continue until $N \times m$ edges have been added.

4.1.2 Analysis of the degree distribution

Figure 4.1 is a series of visualisations of a PG network from its initialisation until it reaches 500 nodes. Qualitatively it seems that some nodes acquire a disproportionately high amount of connections hinting that the network has a skewed degree distribution. We proceed, in figure 4.2a, with a plot of the degree distribution for a planar growth experiment of order $N = 10^4$, along with a PG-noplanarity experiment of order $N = 10^4$ and a PG-nogrowth experiment of order $N = 2 \times 10^3$. A smaller value of N is reported for PG-nogrowth due to computational limits. Nonetheless the results show the degree distributions of both reference cases to be exponential while the PG experiment approximates a power law distribution.

To investigate finite size effects we plot, in figure 4.2b, how the maximum degree observed during these experiments varies with the size of the network. Following [Newman \(2003\)](#), we also plot the analytically derived relationship between $\langle k_{max} \rangle$, the mean maximum degree for networks with a power law degree distribution, and N ; $\langle k_{max} \rangle \sim N^{1/(\alpha-1)}$.

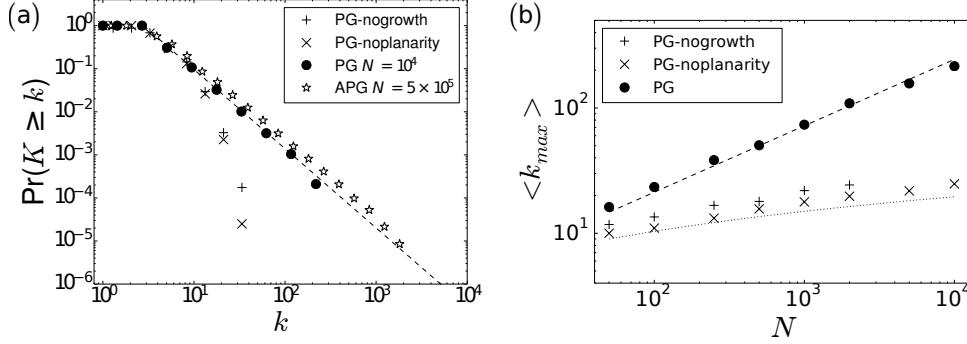


FIGURE 4.2: (a) Cumulative degree distributions for PG networks of order $N = 10^4$, APG networks of order 5×10^5 , PG-noplanarity networks of order $N = 10^4$ and PG-nogrowth networks of order $N = 2 \times 10^3$. All results averaged over 20 experiments with $m = 2$. The dashed line is the best fit for the APG experiment, a power law with exponent, $\alpha_{APG} = 2.77 \pm 0.01$. As with all exponents in this chapter, α_{APG} has been estimated using the method of Maximum Likelihood Estimators outlined in Clauset et al. [Claušet et al. \(2009\)](#). (b) Average maximum degree observed for the PG, PG-noplanarity and PG-nogrowth networks. The dashed line is the expected value of the maximum degree for a power law with exponent $\alpha_{m=2} = 2.83 \pm 0.01$, the estimated value of the exponent in the $n = 10^4, m = 2$ case. The dotted line is a plot of the expected maximum degree for a network with an exponential distribution.

We find it to be in good agreement with the observations, which provides strong support for the hypothesis of a power law distribution.

The expected value of the i^{th} member of a sequential ordering of the random variables of an exponential distribution with parameter λ is $E[X_i] = H_i/\lambda$, where H_i is the i^{th} harmonic number. Barabási & Albert found the degree distribution for the uniform attachment model to be $P(k) = e\lambda \exp(-\lambda k)$ with $\lambda = 1/m$. We therefore approximate the average maximum degree for a PG-noplanarity network of order N with $\langle k_{\max} \rangle \sim mH_N$. The plot of this curve also matches well with our empirical data supporting the claim that the degree distributions generated for both types of reference cases are exponential.

Considered as a whole, the evidence in this section suggests that the network produced by the planar growth process is scale-free. The necessary ingredients in order to produce this outcome are growth and the planarity conservation. When either of these aspects are removed we observe an exponential degree distribution. However, results discussed in this section are unsatisfactory in that the distribution has only been shown to hold over one order of magnitude. We will attend to this in the next section.

4.1.3 Apollonian Planar Growth

Zhou's original RAN algorithm ([Zhou et al., 2005](#)) starts with an equilateral triangle on the plane. Network construction proceeds by repeatedly choosing a face of the triangulation at random, placing a new node within it and connecting that node to the vertices of the face. Note that the probability of a node receiving a new edge is proportional to the number of triangles of which it is a vertex. This count of triangles is, in turn, equal to the degree. As such, the RAN is a form of linear preferential attachment; furthermore, its degree distribution can be analytically demonstrated to be a power law with exponent $\alpha_{\text{RAN}} = 3.0$ when the degrees of the three vertices of the external triangle are ignored.

Apollonian Planar Growth (APG) refines this algorithm by giving the nodes an explicit position on the face of the triangle. Which face is chosen to receive a new node is still random but now this probability is in proportion to the area of the face, i.e., face i is chosen with probability π_i defined by the following formula:

$$\pi_i(t) = \frac{a_i}{\sum_{j \in F_t} a_j}, \quad (4.1)$$

where a_i is the area of face i and F_t is the set of faces present in the simulation at step t .

The new node is then placed uniformly at random within triangle i and connected to its vertices. Clearly, this algorithm is equivalent to planar growth on a triangle with $m = 3$. It has the advantage that the triangulation can be represented as a ternary tree ([Albenque et al., 2008](#)), thereby allowing for more efficient implementation of the model. Thus, in figure 4.2a we present a plot of the degree distribution of an APG network of 5×10^5 nodes which shows the fit of the power law extending over two orders of magnitude on both axes with an estimated exponent of $\alpha_{\text{APG}} = 2.77 \pm 0.01$.

4.1.4 Statistical test of the power law hypothesis

In this section we have estimated several different exponents of assumed power law distributions using the method of Maximum Likelihood Estimation introduced by ([Clauset et al., 2009](#)). MLE can be used as a principled method to estimate the exponent but does not establish if a power law is an appropriate model to describe the data under consideration. To do so Clauset et al. describe two further steps; firstly, goodness of fit is quantified by a p -value calculated by bootstrapping from the estimated model and comparing using the Kolmogorov-Smirnov statistic. Secondly, the power law is compared with other candidate distributions via log likelihood ratios.

However, it has been demonstrated that p -value generated for the bootstrapping stage depends on the sample size of the data. An implication of this finding is that suitably large samples containing noise or imperfections will always fail the goodness of fit test. Further, results from the bootstrapping test may be inconclusive when comparing likelihoods of different models (Klaus et al., 2011; Alstott et al., 2014).

We therefore follow the approach recommended by Alstott et al. (2014) and use the second step as a means to identify the distribution that most appropriately describes the data. In table 4.2, we report the log likelihood ratios, \mathcal{R} , and associated p -values for two experiments, the PG network with $N = 10^4, m = 2$ and the APG network of order $N = 5 \times 10^5$. Following Clauset et al., the alternative distributions considered were the exponential:

$$P(k) = Ce^{-\lambda k} \quad (4.2)$$

the stretched exponential:

$$P(k) = Ck^{\beta-1}e^{-\lambda k^\beta} \quad (4.3)$$

powerlaw with cutoff:

$$P(k) = Ck^{-\alpha}e^{-\lambda k} \quad (4.4)$$

and lognormal:

$$P(k) = Ck^{-1}\exp\left[-\frac{(\ln k - \mu)^2}{2\sigma^2}\right] \quad (4.5)$$

where $\alpha, \beta, \lambda, \sigma$ & μ are the parameters to be estimated for the given distribution, C is a constant that is dependent on these parameters and k is the degree.

The power law model is favoured with high significance over the exponential and stretched exponential models for both PG and APG networks. The lognormal model is not found to be a significantly better fit than the power law model for both network models (indicated by the high p -values). The power law with cutoff model is found to be a significantly better fit than the power law model (and also a significantly better fit than the log normal model: with $\mathcal{R} = 3.0, p = 0.02$ for the APG network and $\mathcal{R} = 5.1, p \sim \mathcal{O}(10^{-8})$ for the PG network). This might be expected given equation 4.4 contains a further

TABLE 4.1: Estimated exponents of networks of order $n = 10^4$ with varying m . Each exponent, α , is calculated from a batch of twenty experiments that grow a network of order $n = 10^4$ using m as specified in the first row. Following [Clauset et al. \(2009\)](#) we use the standard error, σ , as our estimate of the uncertainty in power laws. For all other estimates of uncertainty in this chapter we use the standard deviation.

m	1	1.5	2	2.5	3
α	3.15	2.97	2.83	2.78	2.69
σ	0.03	0.02	0.01	0.02	0.01

TABLE 4.2: Log likelihood ratios of estimated power law distributions compared with other candidate distributions. The log likelihood ratio, \mathcal{R} , and their associated p -values, p , for fits of four alternative distributions compared with the fit of the power law distribution. Statistics were gathered for PG, planar growth with $m = 2$, $N = 10^4$, and APG, Apollonian planar growth with $N = 5 \times 10^5$. Positive values of \mathcal{R} indicate that the powerlaw hypothesis is the preferred model of the data, p is the significance value of the log likelihood ratio.

	exponential		stretched exp	
	\mathcal{R}	p	\mathcal{R}	p
PG	2.7×10^3	$\mathcal{O}(10^{-100})$	34	9×10^{-4}
APG	7.1×10^4	0	1.0×10^3	$\mathcal{O}(10^{-85})$

	lognormal		powerlaw with cutoff	
	\mathcal{R}	p	\mathcal{R}	p
PG	-3.4	0.10	-8.5	3.9×10^{-5}
APG	-2.3	0.26	-5.3	1.1×10^{-3}

parameter beyond the one in equation 4.2. Moreover, the estimated parameters for the functional form of the power law with cutoff suggest that the cutoff is not substantive. We observe that $\alpha = 2.76$, $\lambda = 7.80 \times 10^{-5}$ for the APG network and $\alpha = 2.77$, $\lambda = 0.0023$ for the PG network. We also note that the maximum degree observed for the APG network across all 20 experiments was $k_{max} = 5726$ while for the PG network it was $k_{max} = 271$. Both of these values are less than the corresponding λ^{-1} indicating that the magnitude of the cutoff does not significantly impact the power law.

4.1.5 Robustness to variation of m

We now vary m , the number of connections introduced with each new node, to determine if our observations are peculiar to the $m = 2, 3$ cases. Three is an upper bound on m , which can be established by consideration of Euler's formula for a planar graph, see discussion in reference [Barthélémy \(2011\)](#) for details. We therefore vary m between one and three. Non-integer values of m are attained by always attaching $\lfloor m \rfloor$ edges to a

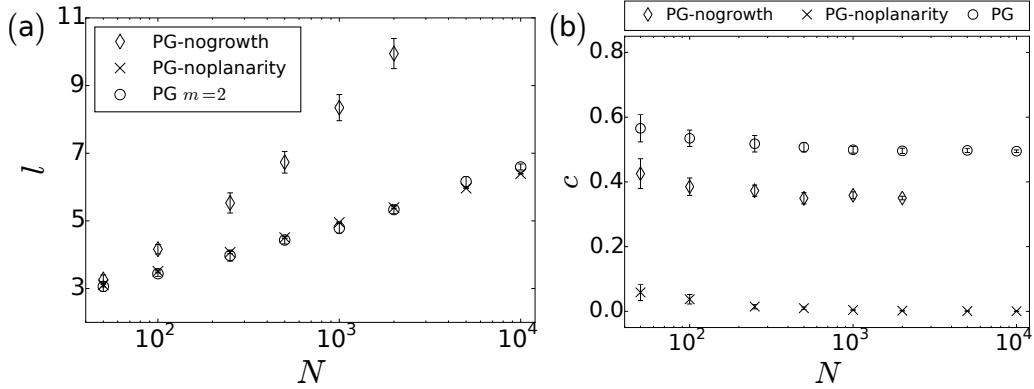


FIGURE 4.3: (a) Mean characteristic path length for PG, PG-nogrowth and PG-noplanarity networks with varying order N . Note the logarithmic scaling on the x-axis. (b) Clustering for the same networks. Error bars in each image is one standard deviation.

new node and then attaching a further node with probability $m - \lfloor m \rfloor$. The network size in these experiments was fixed at $n = 10^4$ and were observed to exhibit power laws. In table 4.1 we report the estimated exponents for these networks which decrease from $\alpha_{m=1} = 3.15 \pm 0.03$ to $\alpha_{m=3} = 2.69 \pm 0.01$. From this point of view PG can be thought of as a generalisation of APG, which strictly has $m = 3$, to any average degree less than three.

4.2 Analysis of Planar Growth

Having investigated the degree distribution we now take a look at other key indicators of global structure. We begin with the small world property and assortativity. Subsequently we examine how the planarity constraint affects the distribution of angles between edges.

4.2.1 The small world property and assortativity

We seek to determine if PG networks have the small world property; the defining characteristics of which are that the network's clustering coefficient (Watts and Strogatz, 1998), c , is high and the network's mean characteristic path length, l , scales with N as $l \sim \ln N$. Here, $l = \sum_{i,j \in V} d(i,j)/N(N-1)$ with V the set of vertices of the network and $d(i,j)$ the length of the shortest topological path between i and j . For random scale-free networks with $2 < \alpha < 3$ it is known that l scales with N as follows: $l \sim \ln \ln N$ (Cohen and Havlin, 2003). However the order of the networks, $N = 10^4$, does not permit the precision necessary to confirm if this is the case for PG. Instead, figure 4.3a, a plot of the observed l for PG networks with varying N and $m = 2$, indicates that l grows

logarithmically with N . Figure 4.3b shows that, for the same network, clustering is high for $m = 2$ and we therefore conclude that is a small world network.

Large values of c in this case are accounted for by the fact that when a node is added it will form connections with the end nodes of nearby edges. PG networks with $m \geq 2$ will therefore tend to form triangles with nearby edges. Furthermore, nearby edges deny a significant portion of the network to new nodes, thereby exacerbating this tendency. The short overall path length l , is a consequence of the high degree hubs in the network.

Now, we consider the assortativity coefficient, a , which we define, following Newman (Newman, 2002), as a correlation coefficient of the degrees at either ends of an edge, i.e.,

$$a = \frac{N^{-1} \sum_i j_i k_i - [N^{-1} \sum_i \frac{1}{2}(j_i + k_i)]^2}{N^{-1} \sum_i \frac{1}{2}(j_i^2 + k_i^2) - [N^{-1} \sum_i \frac{1}{2}(j_i + k_i)]^2} \quad (4.6)$$

where j_i, k_i are the degrees of the vertices at the end of the i^{th} edge.

Specifically, we investigate how a varies with m . For PG networks of order $N = 10^4$, it decreases from $a_{m=1} = -0.029 \pm 0.006$ to $a_{m=3} = -0.066 \pm 0.002$, i.e., the networks are mildly disassortative and this tendency increases as m increases. Plots (not presented) of this relationship show it to be roughly linear. These results are in line with the well known fact that random scale-free networks are disassortative (Park and Newman, 2003; Maslov et al., 2004). A partial explanation of this phenomenon that has been offered is that there is a limited number of possible edges that can lie between high degree hubs (Maslov et al., 2004). So, in general, a scale-free network must feature connections between high and low degree nodes.

4.2.2 Angle distribution

Visualisations of PG, PG-nogrowth and PG-noplanarity networks are shown in figure 4.4. A notable qualitative feature of the PG and PG-nogrowth plots are that edges emanating from the same node often closely bunch together. Combined with the observation of high clustering, this suggests that the space is characterised by triangles with at least one highly acute angle. By contrast the PG-noplanarity network looks markedly different to the naked eye with edges crossing each other freely.

Within city science, the distribution of angles between edges has been successfully employed to gain quantitative insight into road networks (Chan et al., 2011; Barthélémy et al., 2013). In a similar fashion we here consider those edges incident to a vertex in clockwise order and calculate the angle ω between subsequent pairs. The probability density of ω is presented as a series of histograms beneath the corresponding visualisations in figure 4.4. In the PG case three peaks are apparent; at zero, π and 2π

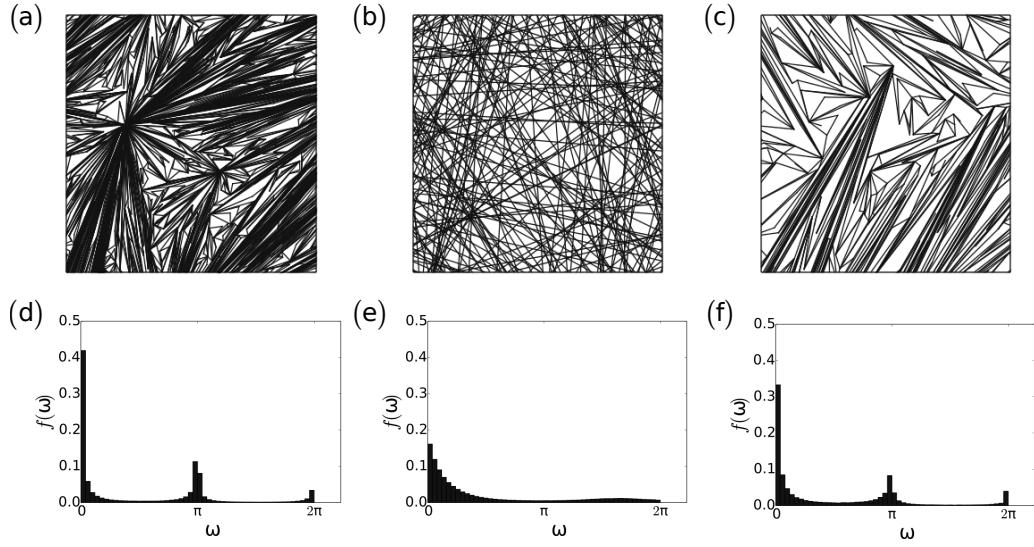


FIGURE 4.4: (a) Visualisation of a 0.35×0.35 patch of a $N = 10^4, m = 2$ planar growth network. (b) Visualisation of a 0.01×0.01 patch of a $N = 10^4, m = 2$ PG-noplanarity network. (c) Visualisation of a 0.35×0.35 patch of a $N = 2 \times 10^3, m = 2$ PG-nogrowth network. (d), (e) and (f) the probability mass for the angle, ω , between successive clockwise edges at a node of the network immediately above. Patch sizes in (a), (b) and (c) have been chosen so as to best illustrate the network.

radians. The peak at zero is the largest and indicates the high number of acute angles just described. The peak at 2π is evidence that in some cases the acute angle will be complemented by a large angle. The difference between the peaks at zero and 2π indicates that in many cases several acute angles will be recorded at a single node in a fan like structure. There will also be occasions when these fan like structures are formed next to pre-existing edges. When this happens the fan will spread towards the edge without crossing it, thereby resulting in two edges incident at the same node that form an almost straight line. It is this phenomenon that accounts for the peak at π . It should also be noted that the peak at π will be influenced by the boundary conditions; fans that appear next to the middle of a side will contribute to it.

The PG-noplanarity histogram shows a large proportion of small angles in a distribution that smoothly and rapidly tails off. There is a small bump at higher values of ω which is a consequence of the square's boundary, i.e., nodes at the corners will tend to have some $\omega > 3\pi/2$. Finally, the PG-nogrowth histogram is very similar to that of the PG networks, confirming that the angular structure is a consequence of the planarity constraint.

TABLE 4.3: Estimated exponents of crossing probability networks. Each exponent α is estimated assuming a power law degree distribution for a batch of twenty experiments which grow an $N = 10^4$, $m = 2$ network with the crossing probability χ that is specified in the first row. Standard error, σ , is reported in the third row.

χ	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
α	2.83	2.89	3.00	3.13	3.31	3.55	3.95	4.68	6.21	5.71	6.25
σ	0.01	0.01	0.02	0.01	0.02	0.04	0.03	0.08	0.32	0.07	0.05

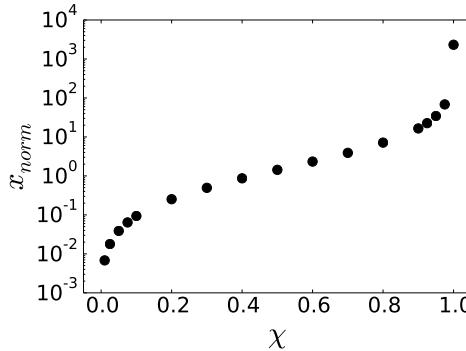


FIGURE 4.5: The normalised count of crossings, x_{norm} , observed in experiments with varying χ . Normalisation was observed by dividing x , the number of crossings, by 20,020, the number of edges. The value of $x_{norm} = 0$ when $\chi = 0$ cannot be represented on logarithmic axes and has been approximated by the value for $\chi = 0.01$.

4.3 Planarity relaxation

The contrast between the PG and the PG-noplanarity degree distributions is dramatic and we would like to investigate intermediate networks. To do so we introduce a new parameter; $\chi \in [0, 1]$, the crossing probability. This parameter is applied in step 2b of the PG algorithm where, each time we encounter a crossing, instead of rejecting the crossing outright, we allow it with probability χ . We grow networks with $N = 10^4$ and $m = 2$ while using a different value of χ in the range 0.0 and 1.0 for each experiment. Our first result, presented in figure 4.5, is a plot of the normalised number of crossings which shows that the number of crossings increases in a roughly exponential fashion between $\chi = 0.1$ and $\chi = 0.9$. Beyond $\chi = 0.9$ the number of crossings increases significantly in comparison to the previous regime.

The associated degree distributions are shown in figure 4.6a where we see a smooth transition from a power law to an exponential curve as χ increases from 0.0 to 1.0. Similarly, in figure 4.6b we present the average maximum degree where plots for low χ match the predicted maximum of a power law while increasing χ leads to curves that more closely match the exponential prediction. Taken together, this evidence shows a smooth transition from a heavy tailed to an exponential degree distribution as χ

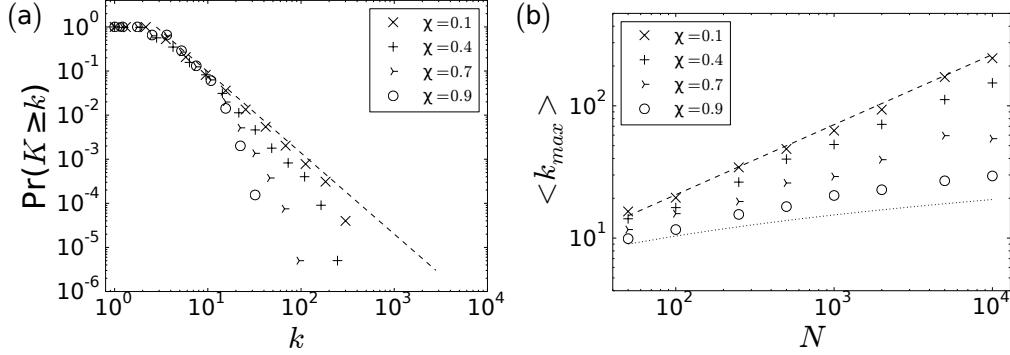


FIGURE 4.6: (a) Cumulative degree distributions for networks created using planar growth with a probability χ of accepting edge crossings. The dashed line is the power law with exponent $\alpha_{m=2}$, the best fit for the $\chi = 0.0$ experiment. (b) Average maximum degree observed in the same experiments. Dotted and dashed lines are the same references plotted in figure 4.2 and are fits for the $\chi = 0.0$ cases and $\chi = 1.0$ cases respectively.

increases. We also estimated exponents, assuming a power law distribution, and report the results in table 4.3, finding an increasing trend for the exponent with χ for networks with $\chi \leq 0.8$. However, from an examination of figure 4.6a, it is clear that networks for which $\chi > 0.7$ have a degree distribution that is exponential and we can therefore disregard exponents in this region of the parameter space.

Assortativity is plotted in figure 4.7a and again exhibits a smooth transition, this time from mild disassortativity to assortativity. To account for this change we first consider the strength of node i , $s_i = \sum_{j \in V(i)} w_{ij}$, (where $V(i)$ is the set of vertices connected to i and w_{ij} is the Euclidean distance between i and j). As discussed in section 4.2.1, preventing edge crossing has the consequence that new nodes are more likely to form connections with nodes that lie at the ends of nearby edges. Thus, where planarity is maintained, high strength nodes will be more likely to receive new connections since, by definition, strength is a measure of the total length of edges beginning at node i . These high strength nodes will also tend to be high degree nodes. Thus, for low χ , high degree nodes attract connections from new, low degree nodes thus accounting for the disassortativity. This tendency will be relaxed as χ increases leading to more assortative networks.

We consider the clustering of these networks in figure 4.7b noting a high $c_{\chi=0.0} = 0.49$ descending to a negligible value for $c_{\chi=1.0}$. High clustering occurs for $\chi = 0.0$ for the reasons outlined in section 4.2.1. On the other hand a new node connects freely to any existing node in the $\chi = 1.0$ case and, hence, this model displays no clustering, equivalent to the uniform attachment model.

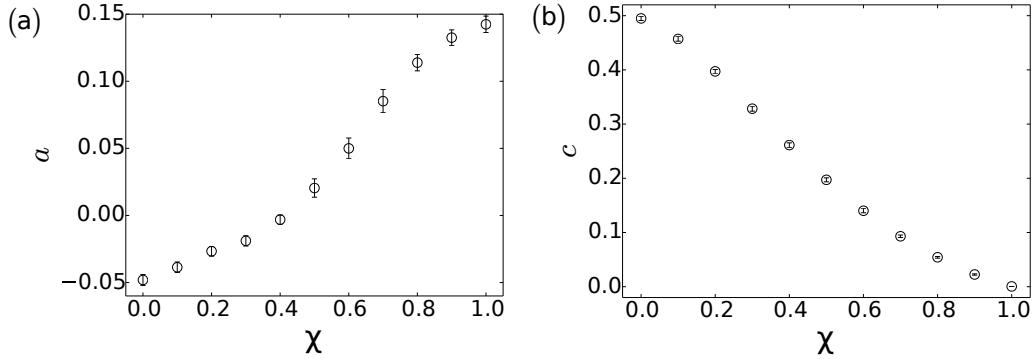


FIGURE 4.7: (a) Average assortativity observed in PG networks with varying χ . (b) Average clustering observed in the same experiments. Each data point relates to twenty networks grown using $n = 10^4$, $m = 2$. Error bars represent one standard deviation.

4.4 Comparison of APG with existing Apollonian growth

In section 4.1.3 we introduced APG as a refinement to the Random Apollonian Network noting that the exponent of its degree distribution was $\alpha_{APG} = 2.77 \pm 0.01$. We contrast this with the analytically derived exponents for the DAN, $\alpha_{DAN} = 2.585$, and the RAN, $\alpha_{RAN} = 3.0$. The APG's exponent lies between these two values and we contend that this is because APG can be thought of as a generalisation of the two existing models.

A triangulation created by any of the three Apollonian growth processes can be represented as a ternary tree where the internal nodes of the tree correspond to nodes of the network and leaves of the tree to the triangular faces (Albenque et al., 2008), see figure 4.8. In the case of the RAN, picking faces of the triangulation uniformly at random is equivalent to picking leaves of the tree uniformly at random. Therefore, the corresponding tree for the RAN will tend to grow in depth since leaves at the bottom of the tree will appear in greater abundance.

DAN constructions begins with K_4 embedded in \mathbb{R}^2 . Growth is an iterative process where, at each stage, a new node is placed within each of the graph's internal faces. Each of these new nodes is then connected to the vertices of its containing face resulting in three new faces. This recursive splitting of the triangle is repeated t times and the corresponding ternary tree has depth $t + 2$ and is both full and complete. Most importantly, the tree for the DAN is shallower than that of the RAN.

The consequence for the degree distributions of the triangulations is as follows: at depth t there are 3^t potential nodes of which $3 \times 2^{t-1}$ will connect to the triangulation's root node. The triangulation associated with the DAN is guaranteed to fill those locations that maximise the degree of the root node. Furthermore, note that the structure of the ternary tree is self-similar. As such any node within the DAN receives the maximum number of connections from its descendants on the tree. Meanwhile, degree in the RAN

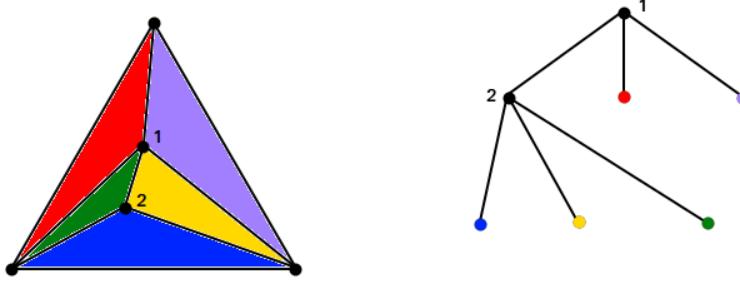


FIGURE 4.8: A triangulation resulting from an Apollonian growth process represented as a ternary tree. Here a root node, 1, and one subsequent node, 2, have been added to the triangulation, resulting in five faces. This triangulation is represented by a ternary tree on the right hand side. The internal nodes of the tree correspond to the nodes of the triangulation and have been numbered as such. The leaves of the tree correspond to the faces and the colouring scheme indicates this.

will be distributed more evenly since the new additions at greater depths will not have the same tendency to link to those at shallower levels of the tree. We therefore expect the DAN to exhibit a heavier tail in its degree distribution and this accounts for the fact that $\alpha_{DAN} < \alpha_{RAN}$.

In the case of APG there will be a tendency to place nodes within those faces with the greatest area. The intuition here is that the earlier a face is created the larger its area will be and therefore those faces that are at a shallow depth within the ternary tree will be favoured for selection. On the other hand, as more nodes are added, a greater proportion of the triangle's total area is covered by newer triangles at greater depth and these will come to be favoured over time. Therefore, nodes in the APG's ternary tree appear at a depth between those of the DAN and the RAN and this in turn explains why we see an intermediate exponent for the degree distribution.

4.4.1 Area weighting and trisection

We test the hypothesis of the previous section by varying the extent to which face area influences its selection within APG. To this end the equation 4.1 is modified as follows:

$$\pi_i(t) = \frac{a_i^\beta}{\sum_{j \in F_t} a_j^\beta}, \quad (4.7)$$

TABLE 4.4: Variation of the degree distribution with area weighting. Estimated values of the exponent of the degree distribution, α , observed for APG networks of order 10^5 with varying area weighting exponent, β . The standard error, σ , in each case is 0.01.

β	10^{-3}	10^{-2}	10^{-1}	10^0	10^1	10^2	∞
α	2.92	2.93	2.82	2.76	2.77	2.76	2.77

where β , the area weighting exponent, is a parameter controlling the influence of a given triangle's area. Clearly, as $\beta \rightarrow 0$, faces will be chosen at random and RAN will be recovered. Conversely, when $\beta \rightarrow \infty$ larger faces will be favoured. Since a triangle that is formed early in the process is more likely to be larger, we expect that faces corresponding to shallower levels of the ternary tree to be filled first. As such, we expect, in this case, to recover the DAN instead.

In table 4.4 we report the exponents for networks created using this variation of the APG and varying β . For values of $\beta < 1$ we see precisely the result predicted, as $\beta \rightarrow 0$, $\alpha \rightarrow 3$. On the other hand, for values of $\beta > 1$, a saturation effect has taken hold and the exponent remains around 2.76. This contradiction with the predicted behaviour occurs because we have assumed that shallow faces in the ternary tree will always have a greater area than deeper ones. Since nodes are placed randomly upon their containing triangle, this is not necessarily the case in Apollonian Planar Growth. Thus, nodes tend to appear deeper in the ternary tree than our initial hypothesis assumed.

In light of this reasoning we further modify the algorithm by placing each new node so that it exactly trisects its containing face; thereby guaranteeing a hierarchy of face sizes by depth within the tree. For $\beta < 1$ behaviour was again as expected; exponents were observed to decrease from 2.93 ± 0.01 to 2.68 ± 0.01 as β increased from 10^{-2} to $10^{-0.5}$. A further experiment with $\beta = 10^2$ gave an exponent of 2.85 ± 0.01 which, *prima facie*, suggests that the hypothesis is incorrect. However, it is apparent from figure 4.9 that the fit is not indicative of the degree distribution of the area weighted APG with trisection. This is because the Clauset et al. method is fitting the black dots in figure 4.9 and is therefore an inappropriate method for quantifying the exponent of power laws that exhibit this sort of discretisation.

To better understand the distribution as $\beta \rightarrow \infty$ we instead follow the formula for the degree distribution of a DAN presented in Andrade et al.'s original paper (Andrade Jr et al., 2005) and plot it on figure 4.9 alongside our own data. It is clear that the discretisation of the experiment closely matches that of the analytical calculation. In a further experiment we set $\beta = \infty$, i.e. the largest triangle was always chosen, and grew a network of 265,720 nodes, the order of an Apollonian network that has been iterated 11 times. In this case the analytical calculation exactly matches the experimental data,

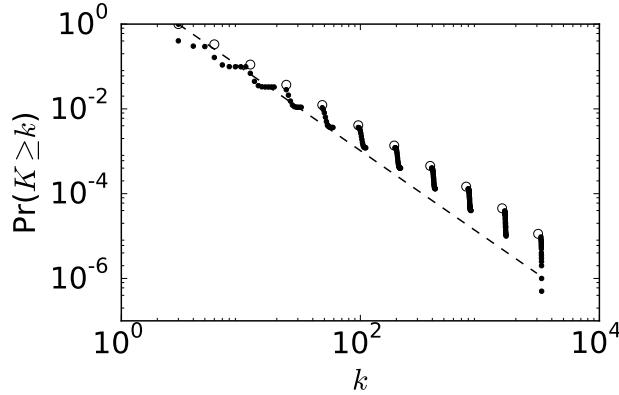


FIGURE 4.9: Black dots are the degree distribution of a network of order $n = 10^5$ grown using Apollonian Planar Growth with $\beta = 100$. Faces were divided by trisecting in to three equal areas in this version of the model. Empty circles are the degree distribution of an Apollonian network of the same order. Dashed line is a plot the best fit of the exponent, $\alpha_{\beta=100} = 2.85 \pm 0.01$.

confirming that the area weighted APG with trisection approximates the DAN as β increases.

To complete the analysis we considered networks with $\beta < 0$, results obtained indicate that an exponential distribution takes hold in this regime. In this regime, new nodes tend to appear within the model's smallest triangle, thereby creating new smallest face from the resulting trisection. Thus, new nodes will tend to congregate in the same region of the model. This contrasts with the $\beta \geq 0$ case where division of the largest face effectively lessens the probability of that region being selected in the next iteration of the process. Thus, the potential for nodes to be distributed over the entire face is a key feature in the onset of the power law degree distribution.

4.5 Summary

We have introduced planar growth as a model of spatial network formation in which a network is grown over time such that planarity is maintained. Resulting networks have been found to be scale-free, have the small world property and are mildly disassortative. It should be noted that PG attains the power law degree distribution with a uniform distribution of nodes in space. As far as we are aware this is only the second example, Mukherjee & Manna ([Mukherjee and Manna, 2006](#)) being the first, of a spatial growth process that attains this outcome under this constraint. The scale-free property is dependent on two aspects of the process; sequential growth and maintenance of planarity. Removal of either aspect results in a network with an exponential degree distribution.

To the extent that the planarity constraint is relaxed the degree distribution degrades from a power law to the exponential case. Similarly, smooth crossover was noted for the clustering and assortativity of these networks. We have discussed those spatial networks that are nearly planar and further note that Newman has articulated a desire for a quantification of the degree of planarity (Newman, 2010). We offer these results as an initial step towards resolving this question.

A refinement of the model, Apollonian Planar Growth, demonstrated a connection between planar growth and Apollonian networks. Weighting the area selection of triangles during Apollonian Planar Growth allowed us to easily recover the Random Apollonian Network while a further variation, trisection, was required in order to produce the Deterministic Apollonian Network. As such, weighted Apollonian Planar Growth with trisection acts as a framework that generalises the two existing Apollonian models.

Chapter 5

Spatial network growth based on stochastic agent decisions

In section 3.1 we discussed the REDS model (Antonioni et al., 2014), a spatially embedded growth process that models the formation of a social network. The distinctive feature of the REDS model is that its nodes are individually assigned a resource with which they maintain the connections that they make in the model. Networks formed by this model exhibit high clustering, positive assortativity and a form of degree heterogeneity, all of which are believed to be features of real world networks. The fact that a form of degree heterogeneity obtains in this model is especially satisfying since the length of edges in this model is constrained.

More precisely, the REDS model is a network growth process based upon the random geometric graph (RGG) (Dall and Christensen, 2002). The RGG is constructed by placing n nodes, uniformly at random, on the unit square and then connecting all pairs that are separated by a Euclidean distance that is no more than a specified radius, R . Nodes in the REDS model can form any edge described by the RGG. However, a node i pays for its edges from an energy budget, E_i , that is specific to i . The cost of an edge, (i, j) , is proportional to D_{ij} , the Euclidean distance between nodes i and j . Additionally, this cost decreases in proportion to the number of neighbours that i and j share. This effect is termed synergy and is governed by the parameter S .

Within the REDS model, an edge is created provided its addition to the network does not cause either end node to exceed its budget. Furthermore, once an edge appears, it is never removed. These aspects of the model imply that nodes in the REDS model are agents that have a high degree of certainty when forming connections with each other. This contrasts with everyday experience where social connections are often fleeting and their perpetuation is contingent upon the participants' unknown (at the time of meeting) capacity to maintain them over time.

Our main contribution in this chapter is instead to envision the nodes as agents whose decisions in this respect are instead subject to noise. The presence of noise in these decisions means that these agents may form connections that exceed their energy budget and they are therefore afforded the opportunity to discard connections. We achieve these goals by adding new edges between unconnected nodes in a probabilistic fashion, with the new edge more likely to be added if the energy budgets of its end nodes are not exceeded. Conversely, existing edges are more likely to be removed if either of the end nodes currently exceeds its energy budget. Since this new model incorporates additional stochastic elements we name it the stochastic REDS model.

We introduce noise into the stochastic REDS model via a parameter that we name the network update temperature, T_n . We grow networks using different values of this network update temperature and report key topological statistics. We find that, for low T_n , these statistics are nearly identical to those of the REDS model. That is, we find three distinct regions of the $S \times E$ parameter space; the first is a high connectivity region where mean degree, $\langle k \rangle$, matches that of the underlying RGG. The second is a low connectivity region. Finally, between the high and low connectivity regions there is a band of $S \times E$ space where we witness a sudden change in connectivity. At higher values of T_n , we find that this boundary region is shifted within $S \times E$ as higher energy and synergy is required in order to produce a high connectivity network.

We investigate dynamics for networks that lie within the transition between high and low connectivity. We find that, for low T_n , the network equilibrates to one of two states, either (i) a constant high connectivity value of $\langle k \rangle$ or (ii) a state where the mean degree does not settle but varies around a low value of $\langle k \rangle$. The probability with which a network in this transition region enters either of these two states is governed by E , the energy that is allocated to each node.

This is in contrast to networks within the transition regions of the REDS model, where resulting networks can exhibit an intermediate value of $\langle k \rangle$. These networks are hybrids of the high and low connectivity networks where parts of the network have RGG-like connectivity and the remainder is sparsely connected and, consequently, they exhibit a bi-modal degree distribution. However, we do find, for high T_n , scenarios where the stochastic REDS model creates networks whose connectivity lies between high and low values of $\langle k \rangle$.

In summary we find that the degree heterogeneity of the REDS model is not robust in the stochastic variant. The bi-modal distribution is only observed when there is high uncertainty in the stochastic variant of the model. We further examine the spatial distribution of the bi-modal distribution's high degree nodes and find that for the REDS model they (the high degree nodes) are located across the plane while for the stochastic REDS model they are located at the edges.

Our first step, in section 5.1, is to recreate the REDS model. In doing so we find a discrepancy between the original implementation and the algorithm as described in Antonioni et al. (2014). We implement both versions and find the results to be broadly similar. In section 5.2 we give the precise description of the stochastic REDS model and report network statistics for a range of temperatures. In section 5.3 we look more closely at the transition in mean degree.

5.1 The REDS model

Our chief aim in this section is to implement the REDS model. We begin by providing a precise description of the model in section 5.1.1. In section 5.1.2 we summarise those results of Antonioni et al. that are pertinent to the present study. Section 5.1.3 contains our own results. We recreate of REDS model and find that it matches Antonioni et al.'s findings. Further, we identify an error that was made with the original implementation and present results for a corrected version.

5.1.1 Process

In this chapter's introduction we outlined the construction of the Random Geometric Graph (RGG) (Dall and Christensen, 2002). The RGG is a fundamental example of a spatial network and directly inspires the REDS model. The RGG is formed by placing nodes uniformly at random upon a unit square and then connecting all nodes that lie within a specified distance of each other. We note here that the degree distribution of the RGG is a Poisson distribution with mean degree $\langle k_{\text{RGG}} \rangle = n\pi R^2$ (Dall and Christensen, 2002). Furthermore, the RGG has clustering and assortativity coefficients $c_{\text{RGG}} = a_{\text{RGG}} = 1 - \frac{3\sqrt{3}}{4\pi}$ (Antonioni and Tomassini, 2012) (Barnett et al., 2007).

The REDS model takes the RGG as its starting point but considers the nodes of the network to be social agents. Two intuitions about social interaction are embodied in the REDS model; (i) that agents have a limited capacity to maintain social bonds and (ii) that social bonds are more easily maintained between agents that share mutual acquaintances. Intuition (i) is realised within the model by allocating an individual budget, E_i , to each node i . An edge (i, j) incurs a cost c_{ij} which must be paid for from this budget. Intuition (ii) is realised by making c_{ij} inversely proportional to k_{ij} , the number of neighbours shared by i and j . These ideas are embodied in the following equation:

$$c_{ij} = \frac{D_{ij}}{1 + Sk_{ij}} \quad (5.1)$$

where D_{ij} is the distance between nodes i and j and $S \in [0, 1]$ is the synergy, a parameter that governs the extent to which shared neighbours reduce the cost of edge formation.

The REDS model is initialised in two steps; firstly we choose four parameters: the network order n , the maximum connection distance R , the synergy S and the per node energy budget, E . Secondly, each of the n nodes are placed uniformly at random within a unit square with periodic boundary conditions and allocated an initial energy $E_i = E$. The algorithm for growing the network then proceeds as follows:

1. A node i is picked uniformly at random. A second node is picked, uniformly at random, from all nodes j such that $D_{ij} < R$ and (i, j) does not exist.
2. An undirected edge (i, j) is added if $E_i \geq \sum_x c_{ix}^{+ij}$ and $E_j \geq \sum_x c_{jx}^{+ij}$. Here c_{ix}^{+ij} is the cost of the edge (i, x) in the network modified by the addition of edge (i, j) . The summation over x represents the summation over the neighbours of i in the first inequality and j in the second.
3. Steps 1 & 2 are repeated until there are no edges that satisfy the criteria laid down in step 2.

The aforementioned algorithm is the REDS model as described in [Antonioni et al. \(2014\)](#). During the preparation of this section of the thesis we communicated with one of the authors of Antonioni et al. and established that the results in the original paper were produced using an erroneous implementation of the REDS algorithm. Specifically, in step 2 of the algorithm, the implementation did not update the network before calculating the cost of an edge, i.e., the summations in step 2 were performed using c_{ix}, c_{jx} rather than $c_{ix}^{+ij}, c_{jx}^{+ij}$.

We interpret this variation of the algorithm as meaning that a node cannot easily calculate what the impact of adding the new edge will be upon its existing social network. For this reason we refer to the erroneous implementation as the non-predictive REDS model. Since the addition of a new edge (i, j) will reduce the cost of i and j 's shared edges, the non-predictive REDS model will disallow connections that would be valid under the REDS model.

5.1.2 Results as reported by Antonioni et al.

Antonioni et al.'s results can be broken into three main strands. Firstly, they reported the mean degree, $\langle k \rangle$, clustering, c , and assortativity, a , for networks with varying combinations of S and E . Secondly, they reviewed typical degree distributions, again for varying S and E . Finally, they examined community structure for two networks with differing values of S . Since we wish to establish if our recreation captures the essential features of the original network, a comparison with Antonioni et al.'s network statistics

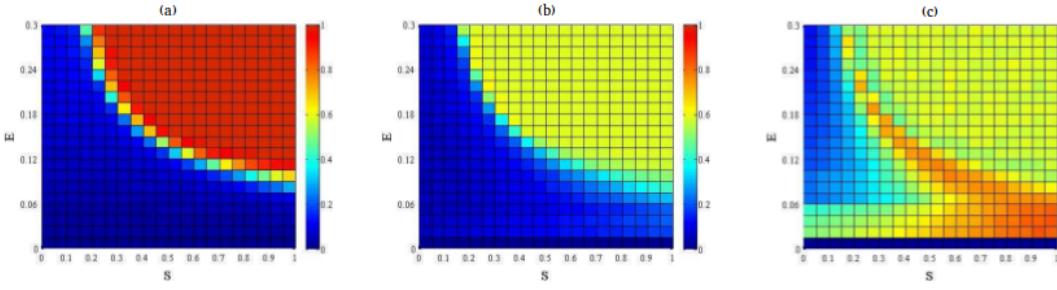


FIGURE 5.1: Mean degree $\langle k \rangle$, clustering coefficient c and assortativity a for non-predictive REDS networks of order $n = 10^4$ and the maximum connection distance $R = 0.05$. These heatmaps are reproductions of those within Antonioni et al. (Antonioni et al., 2014).

and reported degree distributions will be sufficient. Thus, we now review Antonioni et al.'s findings with respect to these issues.

Figure 5.1, is a series of three heatmaps that illustrate how the mean degree, $\langle k \rangle$, clustering, c , and assortativity, a , vary for networks grown by the non-predictive REDS model using different combinations of S and E . The heatmaps are taken from Antonioni et al.'s original work where networks of order $n = 10^4$ and maximum connection distance $R = 0.05$ were grown using the non-predictive REDS model. These choices of n and R effectively define the node density of these networks and determine the mean degree of the underlying RGG to be $\langle k_{\text{RGG}} \rangle = 78.54$. Thus, where Antonioni et al. report a value of $\langle k \rangle$ it is normalised by $\langle k_{\text{RGG}} \rangle$, a convention that we will follow in the remainder of this thesis.

The observations were made within the region of the $S \times E$ plane where $0.0 \leq S \leq 1.0$ and $0.0 \leq E \leq 0.3$. The upper bound on S is chosen to ensure that, for a node i within a fully connected clique, c_{ij} will always be at least the average distance between i and its neighbours. To justify E , we first observe that the mean distance between i and one of its neighbours is $2R/3$. Consequently, when $S = 0.0$, a node i of mean degree, $\langle k_{\text{RGG}} \rangle$, will require a budget of $E_i = 2n\pi R^3/3$ to form connections with all of its neighbours. Since, for $n = 10^4, R = 0.05$, this energy would be $E_i = 2.62$, the threshold $E = 0.3$ ensures that, in the $S = 0.0$ case, the non-predictive REDS model cannot form all of the connections required to form an RGG.

The most obvious feature in figure 5.1a is the division of the mean degree heatmap into two distinct regions. The top right of the heatmap is a region where $\langle k \rangle \sim 1.0$. Within this region the clustering is $c \sim 0.59$ and the assortativity is $a \sim 0.59$. These statistics are consistent with the analytical results for RGGs reported in section 5.1.1. In this region of $S \times E$ space, synergy allows more edges to be created than is possible in the $S = 0.0$ case. Thus, the networks in this region make almost all the possible connections that are available to them. For this reason we refer to this as the high connectivity region of $S \times E$ space.

Below and to the left of the high connectivity region S and E are not high enough to permit edge formation in any great quantity and, consequently, we instead observe a low connectivity region. Clustering is also low within this region. Assortativity within the low connectivity region is moderate for low values of E and high for a band of parameter values that lie just below the high connectivity region. The remainder of the low connectivity region has low, i.e., close to zero, assortativity.

Antonioni et al. also investigate the degree distributions networks drawn from different regions of the $S \times E$ heatmap. They assert that networks within the high connectivity region have the Poisson degree distribution with $\langle k \rangle \sim 1.0$. Networks within the low connectivity region have a degree distribution with a large peak at low values of k . They further identify a third region of the $S \times E$ heatmap in the boundary between the high and low connectivity regions. Here the networks exhibit a bi-modal degree distribution, in that the distribution exhibits two distinct peaks, a sharp peak at low k peak and a second, flatter peak at a higher value of k . Antonioni et al. argue that these two peaks indicate that these networks are composed of regions that are sparsely connected and other regions that have RGG like connectivity.

5.1.3 Implementation of both variants

In this section we recreate the non-predictive REDS model and implement the REDS model. We then compare Antonioni et al.'s results with our own. To allow for this comparison, we retain the parameters that govern network order, $n = 10^4$, and maximum connection distance, $R = 0.05$. Our implementations differ from the original specification in one respect; since real world social systems have defined geographic boundaries, we use rigid rather than continuous boundary conditions. In what follows, we find that observed phenomena closely match those reported by Antonioni et al. and for this reason we have persisted with the rigid boundary conditions in the remainder of this thesis.

The three heatmaps in figure 5.2a have been produced by following the non-predictive REDS algorithm. Since we do not have access to Antonioni et al.'s original results we instead rely on a qualitative comparison between figure 5.1 and figure 5.2a where we find a near exact match between the two sets of heatmaps. In particular we have reviewed values of mean degree, clustering and assortativity within the high connectivity region and find them to be consistent with the analytical values for an RGG, i.e., $\langle k \rangle \sim 1.0$ and $c, a \sim 0.59$.

The heatmaps in figure 5.2b are results for the REDS algorithm. Here, we again see plots that are similar to those in figure 5.1. The significant difference is that, for high S , the high connectivity region of the $\langle k \rangle$ plot extends further down the right hand side of the plot, i.e., high connectivity is observed for lower values of E . This difference is apparent

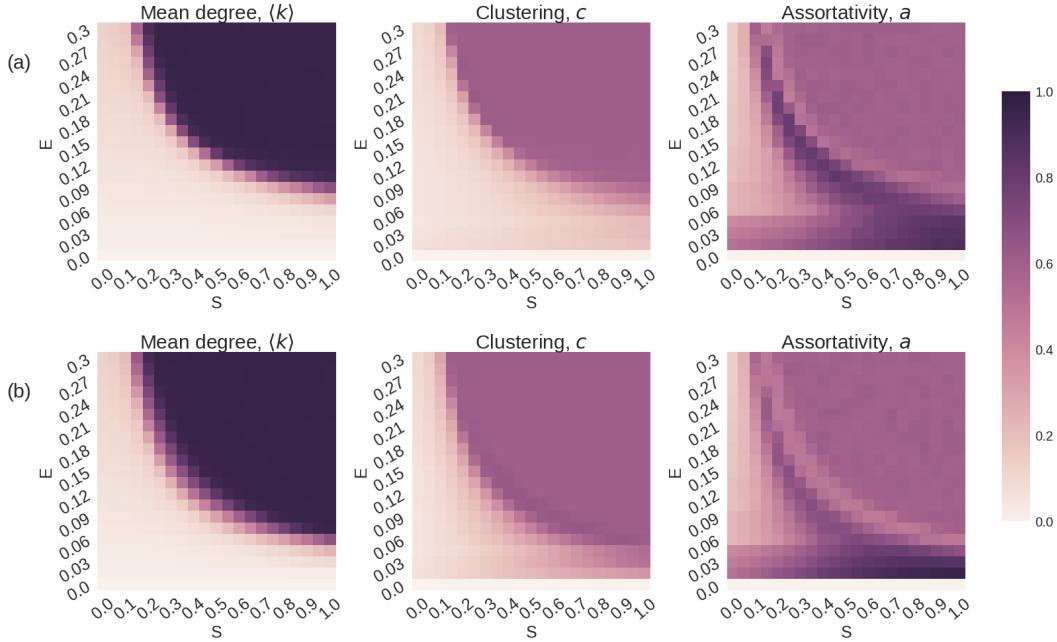


FIGURE 5.2: Mean degree $\langle k \rangle$, clustering coefficient c and assortativity a for (a) re-implemented non-predictive REDS networks and (b) REDS networks. The value in each heatmap cell has been averaged over 10 instances of networks of order $n = 10^4$ and maximum connection distance, $R = 0.05$.

in the clustering and assortativity plots. In both cases the region where $c, a \sim 0.59$ also extends further down the right hand side of the plot.

At the end of section 5.1.2 we discussed the three types of degree distribution that Antonioni et al. identified. Two of these distributions were characteristic of the high and low connectivity regions of $S \times E$ space, while the third was discovered in the boundary between these regions. In order to represent the equivalent regions of $S \times E$ space for the REDS model we fix $S = 0.8$ and choose $E = 0.03$ (low connectivity), $E = 0.06$ (boundary) and $E = 0.24$ (high connectivity). For each energy we generate 15 instances of the network.

Figure 5.3a is the degree distribution formed by compiling the degree sequences for all 15 $E = 0.03$ networks into a single, master degree sequence. The distribution peaks at $k = 1$ which rapidly decreases as k increases. The maximum degree observed across all the $E = 0.03$ networks is 27.

Figure 5.4a visualises a portion of a single $E = 0.03$ network. Here, patches of relatively high degree clusters are distinct amongst the predominately low degree nodes. The prevalence of low connectivity indicates that these networks may be fragmented, i.e., they consist of multiple components. We define g_c to be the number of components in a network find the mean number of components within the $E = 0.03$ networks to be $\langle g_c \rangle = 1331 \pm 8$.

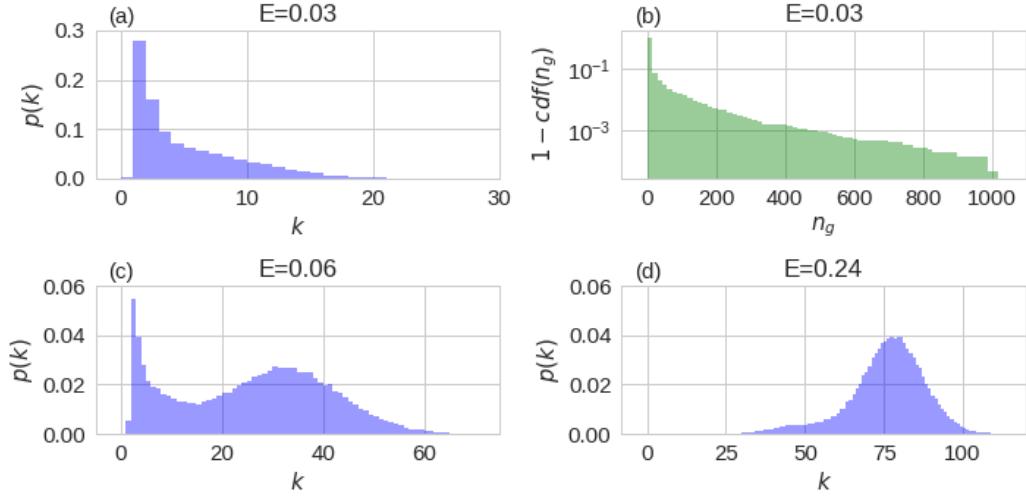


FIGURE 5.3: Degree distributions averaged over 15 instances of networks created using the REDS algorithm with $S = 0.8$ and (a) $E = 0.03$, (c) $E = 0.06$ and (d) $E = 0.24$. Values have been normalised so that the area under the histogram is 1.0. Plot (b) is the reverse cumulative distribution of component sizes for the $E = 0.03$ networks. Individual scales have been used for each plot so as to best fit the panel.

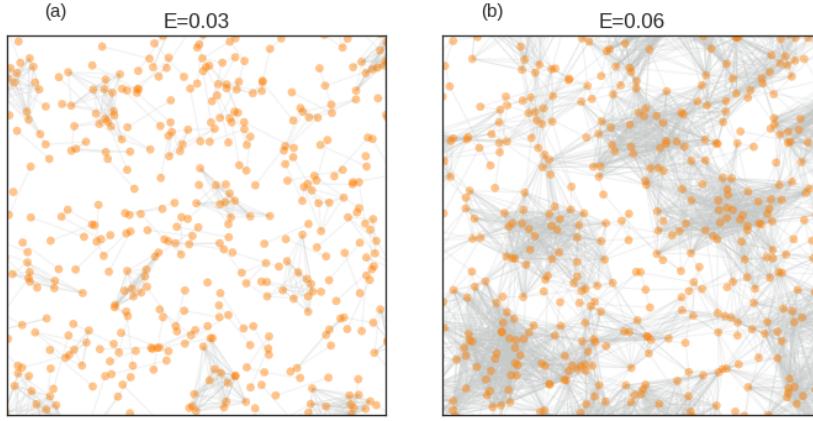


FIGURE 5.4: Visualisation of two REDS networks. The square of side 0.2 with bottom left corner at $(0.4, 0.4)$ has been visualised for REDS networks with $S = 0.8$ and (a) $E = 0.03$, (b) $E = 0.06$.

To analyse the order of these components, we define n_g as the number of nodes within a component and count n_g for each component. Figure 5.3b is a reversed cumulative distribution of n_g . Here we see that the majority of components are small, specifically 91% of these clusters are of order $n_g \leq 10$. However, it should be borne in mind that only 29% of the nodes lie within these $n_g \leq 10$ clusters.

In figure 5.4b, a visualisation of a $E = 0.06$ network, the space appears to be dominated by patches of high connectivity that are separated by boundaries of low connectivity. Figure 5.3c is the degree distributions for the $E = 0.06$ networks. Here, there is a sharp peak at low k and a second, smoother peak, centered around $k \sim 30$. In brief, this distribution and the visualisation confirm that $E = 0.06$ networks exhibit the bi-modal

distribution. Across all 15 of the $E = 0.06$ networks the largest component has order $n_g \geq 9990$. Thus, it seems sensible to say that these networks are a single component, albeit with a very small number of anomalies.

Finally, for $E = 0.24$ we find that, due to the density of edges, the resulting visualisation is unintelligible. However, we have produced the degree distribution in figure 5.3d. The distribution has an (unnormalised) mean degree of 75.2 and its shape is mostly similar to that of a Poisson distribution albeit with a noticeable right skew. These observations are consistent with a network that has reached a state that has formed most, but not all, of the connections required to form an RGG. In all cases these networks consist of a single connected component.

5.1.4 Summary

In this section we have investigated our implementation of the REDS model. Our first result was to implement the non-predictive REDS algorithm used by Antonioni et al. in their original work. We found good agreement between the mean degree, clustering and assortativity $S \times E$ heatmaps reported by Antonioni et al. and our implementation. On this basis we conclude that we have successfully recreated the original work.

For the remainder of the section we used the REDS model. We found the $S \times E$ heatmaps for this implementation to be similar to those for the erroneous implementation; the exception being that, where S is high, the high connectivity region extends towards lower values of E . Furthermore, the associated clustering and assortativity heatmaps also changed in a manner that was consistent with the extension of the high connectivity region. This expansion of the high connectivity region was to be expected since we argued that the effect of the error would be to reject edges that would otherwise be accepted when the effect of synergy was properly accounted for.

We examined degree distributions for three networks drawn from what we expected to be the distinct regions of the REDS model's $S \times E$ space. We found good agreement between the network structures reported by Antonioni et al. and our own results. This final result leads us to conclude that while the REDS model has an extended high connectivity region in comparison to the non-predictive REDS model, the character of the underlying networks remains the same.

5.2 The stochastic REDS model

In the previous section we reviewed and recreated the REDS model. A key theme of that growth process is that edge updates are deterministic; i.e. where a proposed edge is affordable it is added, otherwise it is not. In this chapter we instead investigate the

scenario where there is uncertainty in the process of edge formation with the consequence that nodes may overspend their budget. That being the case, we further require that nodes remove edges so as to counteract energy deficits. We implement this idea by introducing the stochastic addition (and removal) of edges and for this reason we name our new process the stochastic REDS model.

We precisely define this algorithm in section 5.2.1. In particular we introduce the network update temperature, T_n , a parameter which governs the amount of noise to which the stochastic edge update events are subject. In section 5.2.2 we generate networks using varying network update temperatures. We assess the stochastic REDS model by comparing the $S \times E$ heatmaps it produces with those produced by the REDS model. In particular, we find a near exact match between the REDS model heatmaps and those generated by the stochastic process at low network update temperatures. When this temperature is increased, the high connectivity region recedes towards the top of the heatmap and is replaced by an enlarged low connectivity region.

5.2.1 The stochastic REDS algorithm

We have begun this section by highlighting the deterministic nature of edge addition within the REDS model. Specifically, before edge (i, j) is added, the budgets of nodes i and j are checked and if either node cannot afford (i, j) then the edge is not added. We remind ourselves that the REDS model is intended to be a model of social network formation and that node i is considered to be an agent in this context. That being the case, we can interpret edge formation as a decision on the part of agent i . Furthermore, we can consider these decisions to be rational in the sense that i has complete knowledge of both its current expenditure on social connections and the cost of the new social connection.

In this section we instead consider agents whose decisions are subject to uncertainty. In particular these agents make a best guess as to how affordable a potential connection is and, as such, on occasion they will participate in connections that, temporarily, exceed their energy budget. Subsequent iterations of the model will allow agents to review their expenditure and remove edges that violate the energy constraint.

We implement this idea within the stochastic REDS model by using a Fermi function. A Fermi function is a probability distribution which has the following form:

$$f(\delta, T) = \frac{1}{1 + e^{-\delta/T}} \quad (5.2)$$

where $\delta = E - E_F$ is the difference between an energy state E and a ground state E_F . The temperature, T acts as noise parameter; as $T \rightarrow 0$ the Fermi function acts like a Heaviside function, alternatively, as $T \rightarrow \infty$, $f(\delta, T) \sim 0.5$ for all values of δ .

In the stochastic REDS model, the manner in which the Fermi function is applied depends on whether or not the edge under consideration is being proposed for addition or removal. When edge (i, j) is proposed for addition, we first calculate $\delta_i = E_i - \sum_x c_{ix}^{+ij}$. For a specified network update temperature, T_n , we say that node i accepts (i, j) if $r < f(\delta_i, T_n)$, where r is drawn uniformly from the range $[0, 1]$. Conversely, when (i, j) is proposed for removal, we first calculate $\delta_i = \sum_x c_{ix} - E_i$. We say that node i rejects (i, j) if $r < f(\delta_i, T_n)$.

We now define the network construction process as follows:

1. The process is initialised with n , S , E and R as for the REDS model above. In addition to these parameters we also choose a network update temperature T_n .
2. A node i is picked uniformly at random. A second node j is picked, uniformly at random, from all those nodes such that $D_{ij} < R$.
- 3a. If (i, j) does not exist, we **add** the edge if **both** i accepts (i, j) **and** j accepts (i, j) .
- 3b. Alternatively, if (i, j) exists we **remove** the edge if **either** i rejects (i, j) **or** j rejects (i, j) .
4. Steps 2 and 3 are then repeated for some specified number of iterations.

Informally, an edge is more likely to be added at step 3a if both nodes do not exceed their energy budget as a result of doing so. Nonetheless, it can occur, in step 3a that a node will accept an edge that violates this constraint or reject an edge that does not; possibilities that become more likely with increasing T_n . Where an existing edge is selected by the algorithm, step 3b allows for the end nodes to, in effect, reconsider that edge. However, step 3b is also prone to make errors with respect to the budget constraint. In this way we realise agents that are uncertain in their edge update decisions. The extent of this uncertainty is controlled by the network update temperature, T_n ; with low temperature modelling agents that have near complete information as to their expenditure and cost of connections. Conversely, high network update temperature models agents who are more uncertain in these matters.

5.2.2 Results

In this section we grow stochastic REDS networks. Where we grew networks using the REDS model above we relied upon the fact that the original model has a well defined end point, i.e., the process reaches a point where no further edges can be afforded. No such condition exists for the stochastic model. We remind ourselves that the characterisation of the REDS model's $S \times E$ heatmaps was determined primarily by the mean degree.

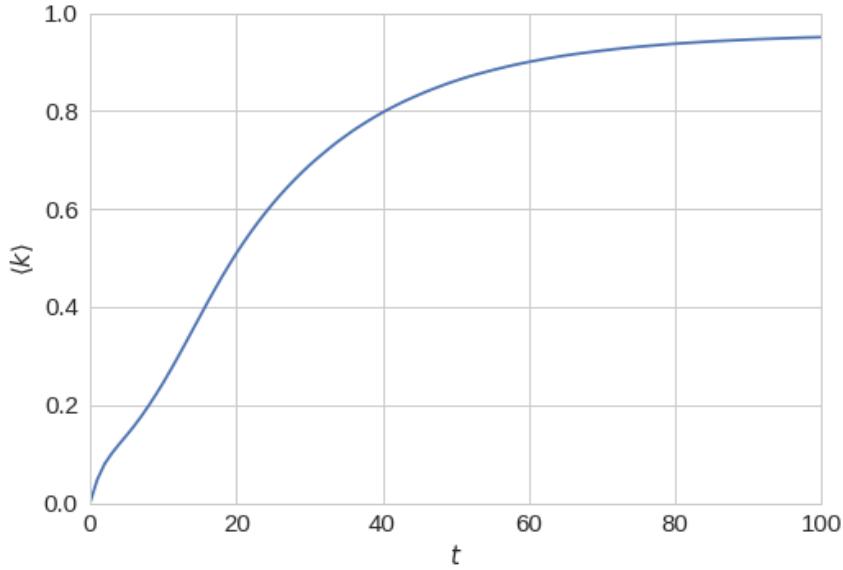


FIGURE 5.5: Dynamics of $\langle k \rangle$ over $t = 100$ epochs of 2.5×10^4 iterations for a stochastic REDS model with $T_n = 10^{-2}, S = 0.8, E = 0.24$.

Ultimately, we wish to compare the results of the stochastic model with those of the REDS. Therefore, in order to establish a stopping condition, we seek to determine when the value of $\langle k \rangle$ stabilises within the stochastic REDS model.

Since we intend to track the dynamics of the quantity $\langle k \rangle$ over the lifetime of the model we now clarify our definition of time for the model. We use the term epoch to mean a set number of iterations, where an iteration is the execution of steps 2 & 3 of the algorithm. In this chapter an epoch consists of 2.5×10^4 iterations. We also highlight here that we continue to fix network order and maximum connection at $n = 10^4$ and $R = 0.05$ to allow for comparison with previous results in this chapter.

As an initial experiment, we run the stochastic REDS model with $T_n = 10^{-2}, S = 0.8, E = 0.24$ and at the end of each epoch record $\langle k \rangle$. Figure 5.5 presents the dynamics of $\langle k \rangle$ for this network. In the plot, we see that the mean degree grows logarithmically until it reaches $\langle k \rangle = 0.95$ at epoch $t = 100$, i.e., this instance of the model has resulted in a high connectivity REDS network. Repeated experiments with parameterisations that result in a high connectivity network indicate that there is some growth in $\langle k \rangle$ beyond $t = 100$. However, the rate of increase of $\langle k \rangle$ remains gradual in this period.

We wish to assign a specific value of $\langle k \rangle$ to this network, however we do not simply use the value just reported for epoch 100. Instead we calculate $\langle k \rangle_{30}$, the mean of $\langle k \rangle$ over the last 30 epochs of the model, i.e.:

$$\langle k \rangle_{30} = \frac{\sum_{i=70}^{100} \langle k_i \rangle}{30} \quad (5.3)$$

where $\langle k_i \rangle$ is the mean degree at epoch t .

Doing so results in a value of $\langle k \rangle_{30} = 0.94$ for this network. Also of importance here is that the standard deviation of $\langle k \rangle_{30}$ is $s_{\langle k \rangle,30} = 0.01$, i.e., the value of $\langle k \rangle$ is stable in this region.

A similar experiment, this time with the energy reduced to $E = 0.03$, results in a network with low connectivity. For networks with these parameters, the mean degree rises to $\langle k \rangle \sim 0.03$ after $t = 50$ epochs and remains at this level until $t = 100$. More precisely, $\langle k \rangle_{30} = 0.03$ and $s_{\langle k \rangle,30} \sim 10^{-4}$ in this case.

In general we have found during initial, exploratory experiments, that these two examples give a good indication of the dynamics of the stochastic REDS model. In particular $t = 100$ epochs of 2.5×10^4 iterations is an adequate period of time for activity in the model, as measured by $s_{\langle k \rangle,30}$, to reach a low value. We will explore these matters in more depth in section 5.3. Before doing so, we will first explore the parameter space more thoroughly.

To do so, we again produce $S \times E$ heatmaps, this time for stochastic REDS networks. We have created three sets of $S \times E$ heatmaps, one for each network update temperature $T_n = 10^{-2}, 10^{-4}, 10^{-8}$. For an individual network we calculate $\langle k \rangle_{30}$ as described above. We also calculate clustering, c_{30} , and assortativity, a_{30} , in an analogous fashion, i.e., by averaging over the relevant observations for last 30 epochs of the model. Each cell in the heatmaps is then an average over 10 instances of the network.

In figure 5.6 we present $S \times E$ heatmaps generated from networks grown using the stochastic REDS model. We begin by reviewing the results for the lowest temperature, $T_n = 10^{-8}$, where we find the heatmaps in figure 5.6a to be near exact replicas of those presented in 5.2b., i.e, at this network update temperature, aggregate statistics for the stochastic model match those of the REDS model.

At network update temperature $T_n = 10^{-4}$ (figure 5.6b) we notice that the region of high $\langle k \rangle_t$ has receded slightly towards the top right corner of the heatmap. Similar phenomena can be observed in both the clustering heatmap and assortativity heatmaps. Note that just below and to the left of the high c_t region is a further band of intermediate clustering. In the assortativity heatmap, the band of low a_t that separates the region of moderate a_t in the top right corner from the band high a_t has disappeared.

Finally, for the highest temperature, $T_n = 10^{-2}$ (figure 5.6c), we see that the high $\langle k \rangle_t$, c_t regions have more definitively receded into the top right corner of the heatmap. In the

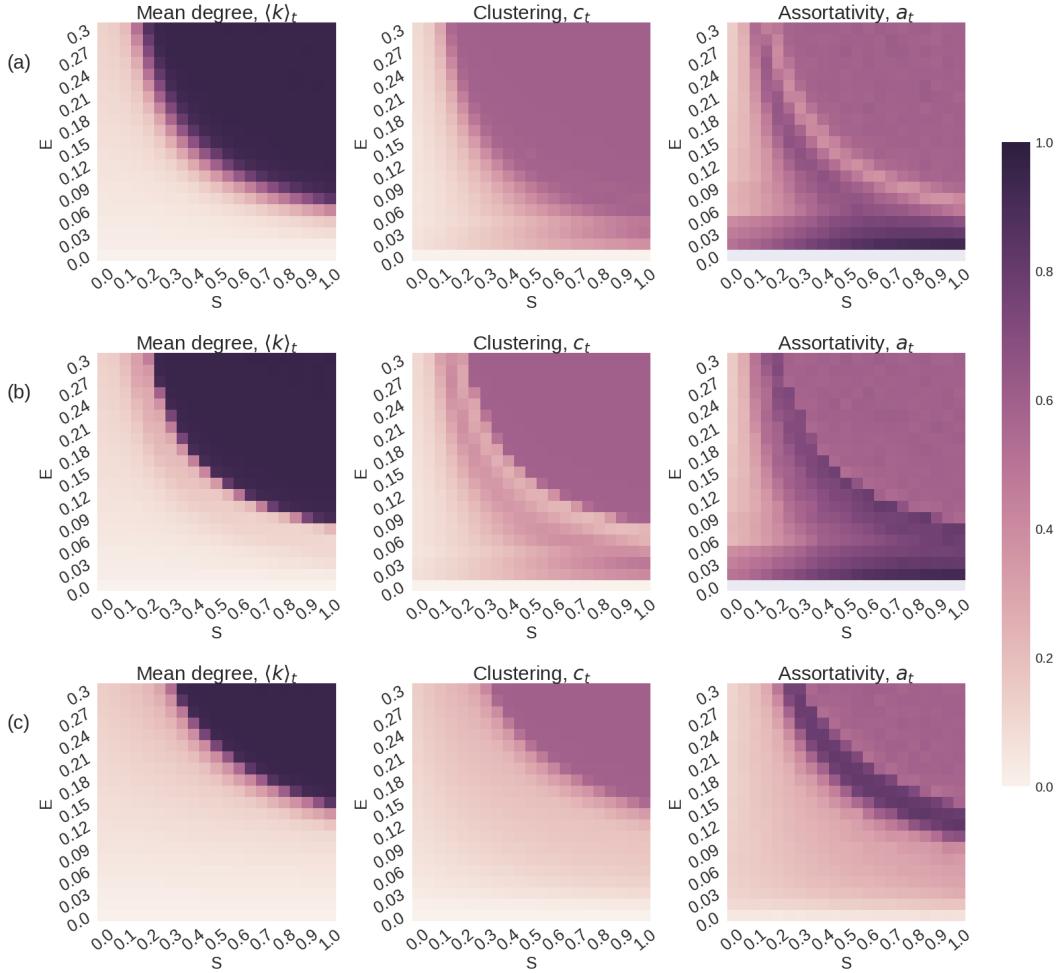


FIGURE 5.6: Mean degree $\langle k \rangle_t$, clustering coefficient c_t and assortativity a_t heatmaps for networks grown using the stochastic REDS algorithm for varying temperatures. (a) $T_n = 10^{-8}$ (b) $T_n = 10^{-4}$ (c) $T_n = 10^{-2}$. The value in each heatmap cell has been averaged over 10 instances of network. Each network is of order $n = 10^4$ with maximum connection distance, $R = 0.05$ and has been grown for 100 epochs of 2.5×10^4 iterations. Statistics for an individual network are calculated by averaging over the last 30 epochs of the process.

assortativity plot, the band of high a_t is noticeably thicker than that in the $T_n = 10^{-4}$ case.

We have also experimented with $T_n = 10^0$ and found the resulting heatmaps to be uniform, with low $\langle k \rangle$, c and a for all combinations of S and E . We find this result to be typical of all examples where we have raised temperatures beyond $T_n = 10^0$.

In order to interpret these results we consider the effect of the Fermi function invoked during step 3 of the algorithm. At low network update temperatures the Fermi function behaves more deterministically, adding edges in a manner that closely matches the original REDS model which accounts for the close resemblance between the heatmaps in figures 5.2b and 5.6a. As we increase network update temperature, edge addition and removal becomes more random and, consequently, the process cannot effectively exploit

synergy. Thus we find that the high $\langle k \rangle$ region recedes, particularly in regions of high S .

5.3 Stabilisation of the mean degree and the transition region

In the previous section we examined stochastic REDS networks for a range of different parameters. We argued there that stabilisation of the mean degree is an appropriate indicator that the process has reached an end state. To this end we discussed the quantity $s_{\langle k \rangle, 30}$, the standard deviation of the mean degree over the last 30 epochs of the model's execution. However, we only offered explicit values of $s_{\langle k \rangle, 30}$ in two example cases.

We address this shortcoming in section 5.3.1 where we report $s_{\langle k \rangle, 30}$ for the experiments discussed in figure 5.6 in section 5.3.1. We find that, in the majority of cases, the value of $s_{\langle k \rangle, 30}$ is low, i.e., the mean degree in the experiments has stabilised. Further, we find that those experiments for which $s_{\langle k \rangle, 30}$ is high are located upon the boundary between high and low connectivity stochastic REDS networks with $S \times E$ space.

This analysis allows for a more precise definition of the transition region. We use this definition in section 5.3.2 to specify a set of experiments that more thoroughly investigate the transition region for different network update temperatures. In doing so we find that it is only at $T_n = 10^{-2}$ that we observe networks with intermediate values of $\langle k \rangle$.

We examine degree distributions for stochastic REDS networks generated with $T_n = 10^{-2}$ in section 5.3.3. We find these distributions match the high, low and bi-modal distributions discussed in previous sections of this chapter.

5.3.1 Defining the transition region

In equation (5.3) we defined $\langle k \rangle_{30}$, the average value of the last 30 observations of the mean degree in an instance of the stochastic REDS model. We further introduced $s_{\langle k \rangle, 30}$, the standard deviation of $\langle k \rangle_{30}$. In figure 5.7 we present heatmaps of $s_{\langle k \rangle, 30}$ for the experiments from which figure 5.6 was generated. The value of an individual cell in figure 5.7 is calculated by first establishing $s_{\langle k \rangle, 30}$ for each of the ten experiments and then taking the average of these.

Within each of these plots we find a thin band of high $s_{\langle k \rangle, 30}$. It is noticeable that, for each temperature, this band of high $s_{\langle k \rangle, 30}$ is located at the border between the high and low connectivity regions of the heatmap. Closer examination of the three heatmaps indicates

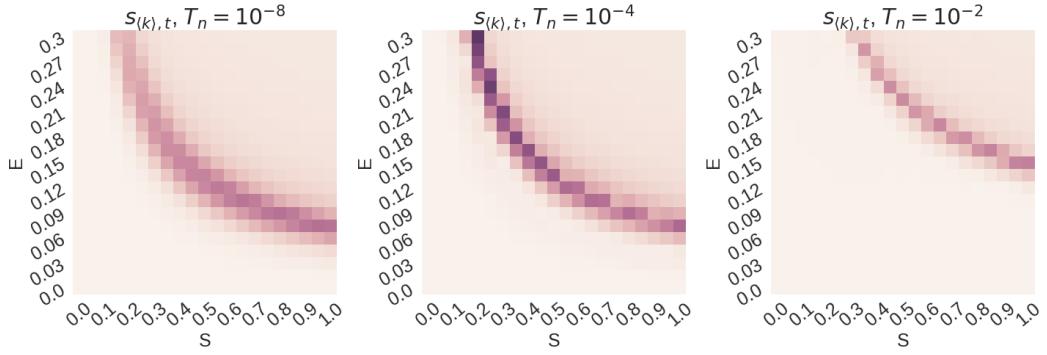


FIGURE 5.7: Heatmaps of $s_{\langle k \rangle}$ for networks grown using the stochastic REDS model for temperatures $T_n = 10^{-8}$, $T_n = 10^{-4}$ and $T_n = 10^{-2}$. Here, $s_{\langle k \rangle}$ is the standard deviation of the observations of $\langle k \rangle$ made over the last 30 epochs of the process. The value of $s_{\langle k \rangle}$ in each heatmap cell has been averaged 10 instances of the model.

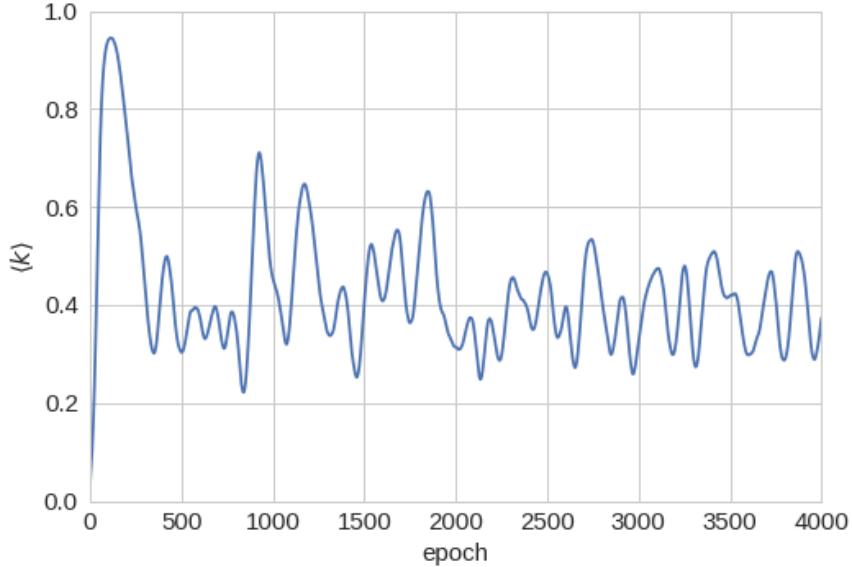


FIGURE 5.8: Dynamics of $\langle k \rangle$ over $t = 4000$ epochs of 2.5×10^4 iterations for a stochastic REDS model with $T_n = 10^{-4}$, $S = 0.8$, $E = 0.102$.

that within the low connectivity regions $s_{\langle k \rangle,30} \sim 10^{-4}$ while within the high connectivity regions $s_{\langle k \rangle,30} < 0.02$. We therefore define the transition region for the stochastic REDS model as being any region of the $S \times E$ heatmap for which $s_{\langle k \rangle,30} > 0.02$ after $t = 2.5 \times 10^6$ iterations of the algorithm.

5.3.2 Characterising the transition region

We now analyse the transition region more closely. To do so, we retain the three network update temperatures considered in section 5.2.2, i.e., $T_n = 10^{-2}, 10^{-4}, 10^{-8}$. We focus our analysis on a single axis of the $S \times E$ heatmap by fixing synergy at $S = 0.8$. For

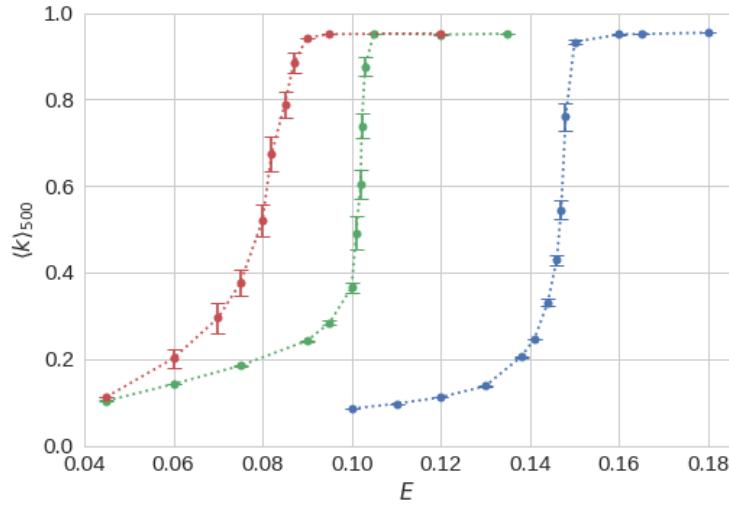


FIGURE 5.9: Plots of $\langle k \rangle_{500}$ for networks grown using the stochastic REDS algorithm for $S = 0.8$ and varying E . Temperatures are $T_n = 10^{-8}$ (red), $T_n = 10^{-4}$ (green) and $T_n = 10^{-2}$ (blue). Values of E have been chosen so as to cover the region of high $s_{\langle k \rangle}$ identified in figure 5.7. Error bars represent the standard error.

each network update temperature we vary the energy E so as to gather a set of data points that are within the transition region and lie along the $S = 0.8$ axis.

We define a stopping condition for the stochastic REDS model as follows: after completion of each step we calculate $s_{\langle k \rangle,500}$, the standard deviation of the mean degree over the last 500 epochs of the algorithm. Once $s_{\langle k \rangle,500} < 0.02$, we halt the algorithm's execution and record $\langle k \rangle_{500}$, the mean value of the last 500 observations of $\langle k \rangle$.

We find this stopping condition to be adequate in most of our experiments. However, we have found that for $T = 10^{-4}, E < 0.105$ and $T = 10^{-8}, E < 0.09$ the dynamics can follow the form of those presented in figure 5.8. Here we find that the value of $\langle k \rangle$ never truly settles with the consequence that the stopping condition is never reached, nonetheless the mean degree varies around a value of $\langle k \rangle \sim 0.4$ in the latter half of the plot. In light of these observations, we run each instance of the model for a maximum of 4000 epochs; in those cases where the stopping condition is not met we instead calculate $\langle k \rangle_{500}$ by averaging over the 500 epochs prior to epoch 4000.

We present plots of $\langle k \rangle_{500}$ in figure 5.9. For each data point, we have run sufficient instances of the model so as to reduce the standard error below 0.04. For all values of T_n , a high value of E produces a network with the high connectivity state. For the network update temperatures, $T_n = 10^{-2}, 10^{-4}$, we observe curves with a steep transition. For the highest network update temperature, $T_n = 10^{-8}$, this transition occurs between $E = 0.144$ and $E = 0.15$, while for $T_n = 10^{-4}$ the transition occurs between $E = 0.1$ and $E = 0.105$. At the lowest network update temperature, $T_n = 10^{-8}$, a more gradual transition occurs between $E = 0.07$ and $E = 0.09$.

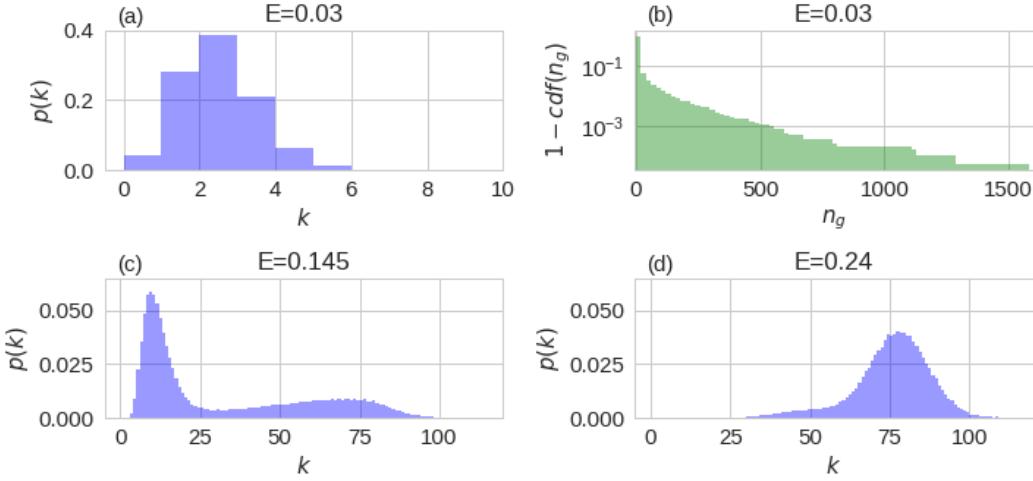


FIGURE 5.10: Degree distributions averaged over 15 instances of networks created using the stochastic REDS algorithm with $T_n = 10^{-2}$, $S = 0.8$ and (a) $E = 0.03$, (c) $E = 0.145$ and (d) $E = 0.24$. Values have been normalised so that the area under the histogram is 1.0. Plot (b) is the reverse cumulative distribution of component sizes for the $E = 0.03$ networks. Individual scales have been used for each plot so as to best fit the panel.

We have examined individual experiments within the transition regions the temperatures, $T_n = 10^{-4}, 10^{-8}$. In these two regions we find that the dynamics of the stochastic REDS model conforms to one of two scenarios. In the first scenario, the mean degree behaves as in figure 5.5, i.e., a logarithmic increase in the value of $\langle k \rangle$ resulting in a high connectivity network. The second scenario is that of figure 5.8, where the mean degree varies around a low value of $\langle k \rangle \sim 0.4$. Thus, intermediate points on the $T_n = 10^{-4}, 10^{-8}$ curves in figure 5.9 do not indicate networks that have reached an intermediate value of $\langle k \rangle_{500}$. Instead they represent the proportion of experiments that result in high connectivity.

By contrast, we find that for the high network update temperature, $T_n = 10^{-2}$, all data points are indicative of the value of $\langle k \rangle_{500}$ attained by the underlying networks. That is, at this network update temperature it is possible for the stochastic REDS model to result in a network that has a value of $\langle k \rangle$ that lies between the mean degrees associated with the high and low connectivity regions.

5.3.3 Degree distributions for the stochastic model

In section 5.1.3 we discussed the degree distributions of three instances of the stochastic REDS model. These distributions were a high connectivity distribution, a low connectivity distribution and a distribution that is a hybrid of high and low connectivity. In section 5.2.2 we concluded that, out of all the network update temperatures that were investigated, the stochastic REDS model with $T_n = 10^{-8}$ produced an $S \times E$ heatmap

that is very similar to that produced by the REDS model. However, we have just observed that, within the transition region, stochastic REDS networks with $T = 10^{-8}$ do not result in an intermediate value of $\langle k \rangle$ and therefore do not exhibit a bi-modal degree distribution. Thus, despite the near agreement between the heatmaps of figures 5.2a and 5.6a, the character of the transition between the high and low connectivity regions is subtly different in the stochastic, $T = 10^{-8}$, model when compared to the REDS model. A similar argument follows for the $T = 10^{-4}$ case.

However, for those experiments with a high network update temperature we did observe intermediate values of $\langle k \rangle$. In light of this outcome we seek to compare the degree distributions for stochastic, $T = 10^{-2}$, REDS networks with degree distributions for the REDS model. We consider stochastic REDS networks with $T_n = 10^{-2}, S = 0.8, E = 0.03, 0.145, 0.24$, thereby gathering representatives of the three regions of the high network update temperature $S \times E$ heatmap in figure 5.6c.

In the experiments presented in figure 5.9, the stochastic REDS model reaches the stopping condition after an average of $\langle t \rangle = 895$ epochs when $E = 0.145$. Further experiments indicate that, for $E = 0.03, 0.24$ cases, $\langle k \rangle$ stabilises after $t < 100$ epochs. To ensure stability in the mean degree, we run the stochastic REDS model for 100 epochs with $E = 0.03, 0.24$ and for 1000 epochs with $E = 0.145$.

In the same manner as section 5.1.3 we create 15 instances of the network. We then gather all 15 degree sequences into a master degree sequence. Finally we normalise the master degree sequence so that the area under the associated histogram has an area of one.

Figure 5.10a is the histogram for $E = 0.03$, the low energy case. We find that the degree distribution has a peak at $k = 2$ and that there are no nodes with $k > 8$. We plot the reversed cumulative distribution of n_g for these networks in figure 5.10b where we find that 89% of the components are of order $n_g \leq 10$ and 25% of nodes lie within these $n_g \leq 10$ components. The mean number of components that make up these networks is $\langle g_c \rangle = 1200 \pm 26$ (compare with $\langle g_c \rangle = 1331$ for the REDS model with $E = 0.03$). Furthermore, we note the longer tail in figure 5.10b in comparison to 5.3b.

In summary, the stochastic REDS, $E = 0.03$ networks have a similar structure to the REDS model with $E = 0.03$. However, the stochastic REDS networks tend to have lower overall degree than the REDS model. Furthermore, it appears that both these models produce similar amounts of components of order $n_g \leq 10$. However, the stochastic model produces fewer components with $n_g > 10$. The longer tail in figure 5.10b indicates that this is because components tend to be larger within the stochastic model.

In figure 5.10b we plot the degree histogram for the $E = 0.145$ networks. Here, the degree distribution also has a significant peak at a low value of k . There is also a second, flatter peak located around $k = 70$. Finally, figure 5.10c is the degree histogram

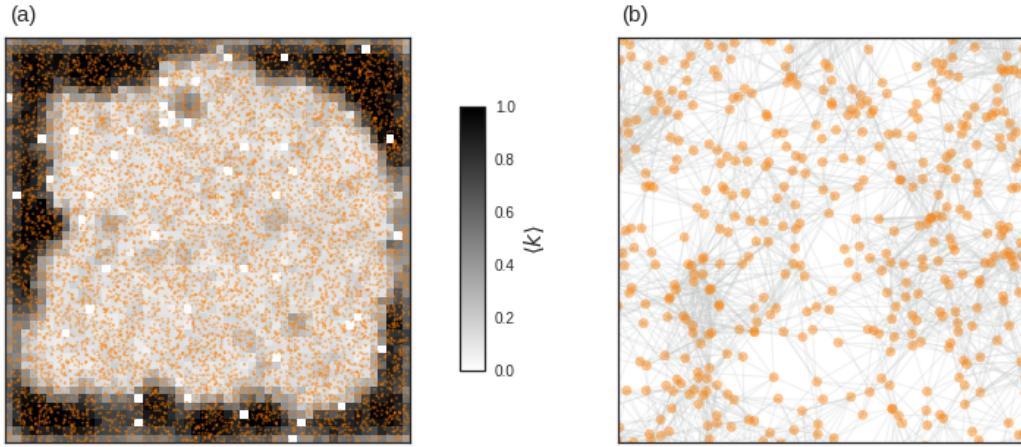


FIGURE 5.11: Visualisations of a stochastic REDS network with $S = 0.8, E = 0.145, T_n = 10^{-2}$. (a) The entire network. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$. (b) The square of side 0.2 with bottom left corner at (0.4, 0.4).

for the $E = 0.24$ networks. In this case the distribution is a large peak around a mean degree $\langle k_{E=0.24} \rangle = 75.1$. There is also a noticeable tail to the left of this peak. All the $E = 0.145, 0.24$ networks are a single component.

To complete our analysis of the $E = 0.145$ case we visualise a single instance of the network. The edge density of the network makes it difficult to ascertain the structure when the connections are plotted. Instead, we break the square down into a 50×50 grid and within each square add a greyscale hue to represent the mean connectivity. The resulting plot is presented in figure 5.11a. Here we see an obvious boundary effect where the RGG-like nodes, located along the edges of the square, form a ring around the less densely connected region. We capture a snapshot of this latter region in figure 5.11b where we visualise the square of side 0.2 and bottom left corner at (0.4, 0.4). This plot bears some similarity to the portion of the $S = 0.8, E = 0.06$ REDS network plotted in figure 5.4b in that there are patches of varying connectivity. However, it is clear that the nodes in figure 5.11b are not as well connected as those in figure 5.4b.

These results indicate that the structure of the stochastic REDS networks with $T_n = 10^{-2}, S = 0.8, E = 0.145, 0.24$ are consistent with the bi-modal and high connectivity regimes for the degree distribution that were reported by Antonioni et al. for the original REDS model and discussed for our own implementation of the REDS model in section 5.1.3. However, there is a notable boundary effect for the transition region example, i.e., the $T_n = 10^{-2}, S = 0.8, E = 0.145$ network. This boundary effect was not observed in the transition region REDS network, i.e., the $S = 0.8, E = 0.06$ case.

5.4 Summary

In this chapter we have studied the REDS model, a process which grows spatially embedded social networks. The key innovation of this model was to introduce a budget with which agents pay for connections and resulting networks exhibit many of the features of real world social networks; including, crucially, a weak form of degree heterogeneity. We began by recreating Antonioni et al.’s original research, the non-predictive REDS model. Following the original work, we summarised the parameter space of the non-predictive REDS model using $S \times E$ heatmaps in figure 5.2a and found them (the heatmaps) to be a near exact match of those presented in Antonioni et al.’s paper, indicating that we had successfully recovered those results.

A subsequent set of heatmaps, produced using the REDS model, were very similar to those for the non-predictive REDS model; the sole difference being an extended region of high connectivity at the bottom right of the mean degree heatmap. This difference is due to the fact that agents in the REDS model, by virtue of their capacity to properly calculate the cost of their connections, are better able to exploit synergy so as to achieve higher connectivity at high S , low E configurations.

The similarity between the REDS model and the non-predictive REDS model was further confirmed by the degree distributions presented in figure 5.3. Here it was demonstrated that the three classes of network structure reported by Antonioni et al. can be recovered by the REDS model and that these network structures are found in equivalent regions of the mean degree $S \times E$ heatmap. More specifically, both heatmaps featured three clearly identifiable regions; one of high, RGG-like connectivity, one of low connectivity and a third whose networks exhibit a bi-modal degree distribution consisting of both high and low connectivity nodes. Of these three distributions, the bi-modal is of the most significance since it is a form of degree heterogeneity.

We introduced an alternative formulation of the REDS model which allowed for the stochastic update of network edges, thereby simulating agents whose decisions to form connections are subject to uncertainty. The extent of the uncertainty within the stochastic REDS model was captured by the network update temperature, T_n , a parameter governing the amount of noise in the Fermi function. At the lowest network update temperature $T_n = 10^{-8}$, we observed $S \times E$ heatmaps whose structure was the near equivalent of those produced by the REDS model. Closer examination of the transition region of the mean degree heatmap for this temperature revealed that we do not observe networks with the bi-modal degree distribution. Instead, intermediate values in this mean degree heatmap were indicative of the proportion of networks that equilibrated to an RGG like state for the given S, E combination.

This result indicates that the bi-modal degree distribution is not a stable outcome for the stochastic REDS model as $T_n \rightarrow 0$. We remind ourselves that continual addition and

removal of edges is the essence of the stochastic REDS model whereas the REDS model halts once there are no more possibilities to add edges. Since this is the key difference between the two models it is this feature that must account for the difference between the two transition regions. A possible explanation is that continued addition of edges to an RGG-like region will result in over-budget nodes that will then have to relinquish edges. Where edges are removed from an RGG-like region the effect of synergy will be reduced, leading to more over-budget nodes. Assuming this brief hypothesis, the ongoing addition and removal of nodes will render high connectivity regions unstable.

When the network update temperature was raised to $T_n = 10^{-4}$, higher values of S and E were required so as to induce high connectivity networks (compare the dark area of the mean degree heatmap in figure 5.6b with that in figure 5.6a). In other words, agents require either more energy or the reduced edge cost that accrues from synergy in order to form the RGG-like network at this temperature. We interpret this result as meaning that as agents become more uncertain, they become less capable, as a population, of maximising the number of connections that they can make with their budget.

This pattern continued at the high network temperature $T_n = 10^{-2}$, where we found that the high connectivity region was described by even higher values of S and E , relative to the $T_n = 10^{-4}$ example. However, within the $T_n = 10^{-2}$ transition region we did find networks that exhibit a bi-modal degree distribution. Where we observed this distribution it was clearly influenced by the rigid boundary condition. To be specific, a high connectivity, RGG-like region along the boundary encircled a lower connectivity region in the centre of the square.

The appearance of the high connectivity nodes at the boundary of the square is most likely due to the fact that nodes at the boundary have fewer possibilities for making connections and therefore they more rapidly make connections with their near neighbours. These nodes can then exploit synergy to form clusters along the boundary. It is notable that we did not see this phenomenon for any REDS networks, indicating that the stochastic algorithm acts as a search mechanism to find this configuration. Furthermore, the $T_n = 10^{-8}, 10^{-4}$ temperatures do not exhibit this phenomenon either, indicating that high agent uncertainty is required to make this discovery.

Chapter 6

Social dilemmas on spatial networks

In section 3.3 we discussed the problem of the evolution of co-operation and the fact that spatial models readily produce interaction structures that promote co-operative behaviour via the phenomenon of network reciprocity. Our discussion focussed on two problems with coevolutionary models that have been applied to this problem; firstly, that the connectivity of the networks is exogenously determined and, secondly, that these models imbue their agents with significant cognitive ability and secondly .

In this chapter we define agents that, in contrast to the existing models, pay for links with a privately held budget as opposed to implicitly assuming a shared resource underpinning the model's connectivity. We achieve this by extending the stochastic REDS model so that its agents play the prisoner's dilemma game with their network neighbours. Taking this step begs the question as to where the agents acquire the energy with which they make connections and we therefore further integrate the stochastic REDS model with the evolutionary game by deriving the privately held budget from the payoff from the evolutionary game. In this new model, decisions to form connections with other agents are determined entirely by the stochastic mechanism outlined in the previous chapter. This approach has the further merit that the agents do not rely on any information other than their accumulated payoff from the game when forming connections. Thus, we can say that these agents rely on a reduced cognitive endowment in this regard.

In section 6.1 we run the Prisoner's Dilemma game on stochastic REDS networks. We name this model the evolutionary REDS model. In section 6.2 we define the coevolutionary REDS model and discuss three canonical forms of co-operation dynamics. In section 6.3 we more thoroughly investigate how well co-operation is supported across a broad range of the coevolutionary REDS model's parameter space. Finally, in section 6.4 we summarise our results.

6.1 Co-operation on the stochastic REDS model

In this section we introduce the evolutionary REDS model. The model is, simply, a static network generated by the stochastic REDS model upon which the prisoner's dilemma game is played. The intention is that the evolutionary REDS model will ultimately act as a reference case for the coevolutionary model that will follow in the next section. We name the current model the evolutionary REDS model. In section 6.1.1 we define the model and then present results in section 6.1.2.

6.1.1 The evolutionary REDS model

The evolutionary REDS model is initialised with the same set of parameters n, R, T_n, S, E that are used for the stochastic REDS model. In addition, we choose $T_g \in [0, \infty]$, a parameter that governs the extent of the stochasticity within the strategy update; a rule that will be introduced below. Having established the initial parameters we then grow the network using the stochastic REDS algorithm.

Within this model, nodes are players of an evolutionary game and therefore need to be allocated a beginning strategy. Half of the nodes are designated as co-operators, while the remaining half are designated as defectors. For the purposes of formalising the game, we denote the co-operation strategy by the two dimensional unit vector $s = [1, 0]^T$, and defection by $s = [0, 1]^T$. In this context T should be interpreted as a vector transpose operation.

A node i plays prisoner's dilemma (PD) with its neighbours on the network and earns an accumulated payoff P_i as follows:

$$P_i = \sum_{j \in \mathcal{N}_i} s_i^T M s_j \quad (6.1)$$

where \mathcal{N}_i is the set of neighbours of i and M is the payoff matrix for a game of PD.

We use a version of PD known as the donation game (Hilbe et al., 2013), so-called because the co-operator effectively donates b , the temptation to defect, to the defector:

$$M = \begin{matrix} & \begin{matrix} C & D \end{matrix} \\ \begin{matrix} C \\ D \end{matrix} & \begin{pmatrix} 1 & -b \\ 1+b & 0 \end{pmatrix} \end{matrix} \quad (6.2)$$

We use this formulation of PD so as to allow for comparability with the coevolutionary REDS model defined in section 6.2 and our motivation is outlined more fully there. For

now, we note that, from the point of view of co-operators, the donation game is a more exacting version of PD than that presented in equation (2.3).

An iteration of the evolutionary REDS model proceeds as follows: a pair of nodes is selected by first choosing a node i at random from all the nodes in the network. A second node, j , is chosen uniformly at random from i 's neighbours. For both nodes we calculate the payoffs, P_i and P_j . We then perform a strategy update by which we mean that we copy s_j to s_i with probability:

$$\phi(P_i, P_j, T_g) = \frac{1}{1 + e^{(P_i - P_j)/T_g}} \quad (6.3)$$

i.e., ϕ is a Fermi function where T_g is a temperature that influences the intensity of selection. This is a common choice of update function within the literature (Szabó and Töke, 1998; Blume, 2003; Traulsen et al., 2006), which we choose because T_g allows us to control the intensity of selection.

Finally, this process of selecting neighbouring nodes and stochastically updating their strategy continues until a specified number of iterations is completed.

To complete this section, we clarify the influence of T_g upon equation (6.3). When P_i exceeds P_j , node i will tend to retain its current strategy. As the difference between P_i and P_j increases, i retains its strategy with greater likelihood. Conversely, the extent to which P_j exceeds P_i governs the likelihood that i will instead tend to adopt j 's strategy. When T_g is low the update will almost always favour the strategy with the higher payoff. With increasing T_g , the function ϕ becomes more uniformly distributed between the two outcomes.

6.1.2 Results

In this section we examine the extent to which the co-operative strategy prevails within the evolutionary REDS model. To do so we define the quantity p_c , the proportion of the network's nodes that have the co-operative strategy. Having run the model extensively we find that the value of p_c tends to equilibrate rapidly, typically after 50 epochs of $t = 2 \times 10^6$ iterations. In this section we run the model for 300 epochs of $t = 2 \times 10^6$ iterations. Having done so we calculate p_c for an individual model run as being the mean value of p_c observed over the last 50 of these epochs. For each value of b we report the mean value of p_c observed over 15 instances of the model.

We seek to run the evolutionary REDS model on networks that are comparable to those discussed in section 5.3.3. Therefore, we fix the values $n = 10^4, R = 0.05, S = 0.8$ and choose the energy to be $E = 0.03, 0.145, 0.24$. To initialise the models, we run the stochastic REDS algorithm for 2.5×10^6 iterations in $E = 0.03, 0.24$ cases and for

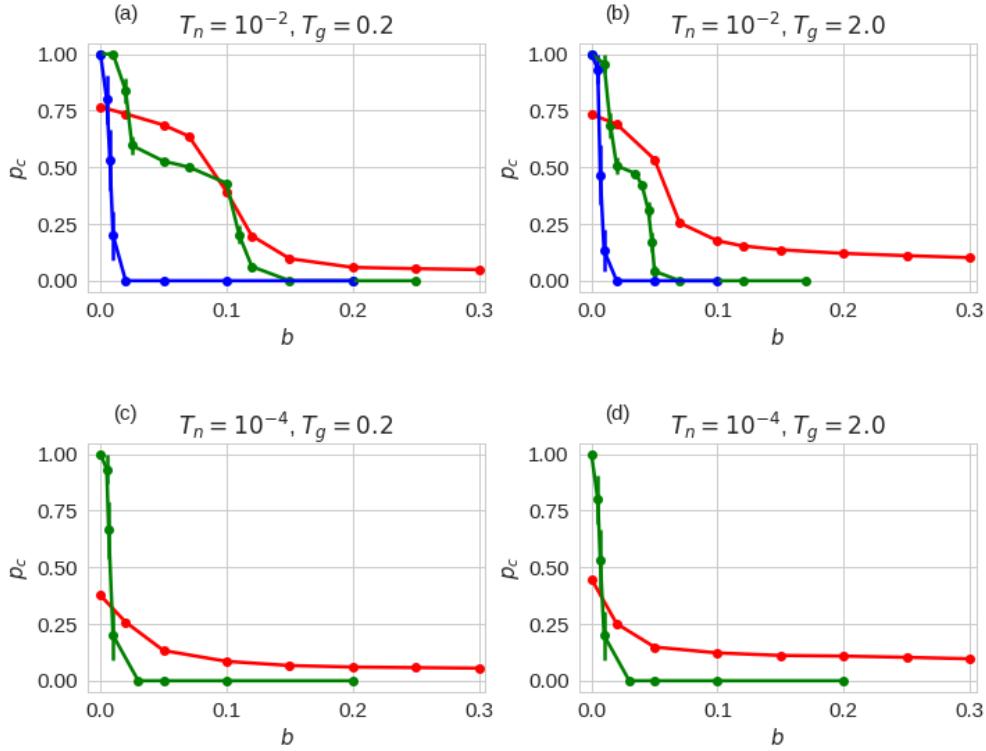


FIGURE 6.1: Plots of p_c vs b for the evolutionary REDS model with varying T_n , T_g and $E = 0.03$ (red), $E = 0.145$ (green) and $E = 0.24$ (blue). The parameters $n = 10^4$, $R = 0.05$, $S = 0.8$ are fixed. Each data point is the mean value of p_c over 15 instances of the model after 300 epochs of $t = 2 \times 10^6$ iterations. For each instance, p_c has been averaged over the last 50 epochs of the model run.

2.5×10^7 iterations in the $E = 0.145$ case, again to maintain fidelity with section 5.3.3. In order to choose the temptation to defect, we have found that the value of p_c does not vary beyond $b \geq 0.3$, therefore, we always choose $b \in [0.0, 0.3]$. The values of n , R and t quoted in this paragraph are used throughout this chapter and hence we do not repeat them beyond this point.

The last condition that we inherit from section 5.3.3 is the network update temperature $T_n = 10^{-2}$; subsequently, we shall also explore the temperature $T_n = 10^{-4}$. Finally, we shall also consider two values of the strategy update temperature, $T_g = 0.2, 2.0$.

6.1.2.1 Low strategy update temperature

Initially, we consider the low network update temperature $T_n = 10^{-2}$ and low strategy update temperature, $T_g = 0.2$. Figure 6.1(a) presents results for this combination. In the following three subsections we attend to each of the different energies.

6.1.2.1.1 High E

We begin with the high energy, $E = 0.24$, scenario. In this case, we have found that co-operation either completely dominates the model, i.e., $p_c = 1.0$, or that co-operation completely collapses, i.e., $p_c = 0.0$. When the temptation to defect $b = 0.0$, we find that co-operation always dominates the model while when $b \geq 0.05$ co-operation always collapses. Where $0 < b < 0.05$ the data point represents the proportion of runs where co-operation does dominate the model.

We recall from the previous chapter that the stochastic REDS model with $T_n = 10^{-2}$, $S = 0.8$, $E = 0.24$ is an example of an RGG like network. These results are therefore consistent with [Buesser and Tomassini \(2012\)](#) who found that, for RGGs with $n = 10^4$, $\langle k \rangle = 20$, the equilibrium value of co-operation in the PD game was $p_c = 0$ for all but very low values of b .

6.1.2.1.2 Low E

For the low energy, $E = 0.03$, scenario we find that at $b = 0.0$, co-operation equilibrates at $p_c = 0.77$. As we increase the temptation to defect we find that the level co-operation decreases but is significantly higher than in the $E = 0.24$ case. We reiterate Ohtsuki et al.'s conclusion, that co-operation is better supported on low degree networks than on high degree ones ([Ohtsuki et al., 2006](#)). It seems clear that this effect explains the improved performance of the $E = 0.03$ network over the $E = 0.24$ network in this case.

Nonetheless, a question remains as to why co-operation does not dominate in the $b = 0.0$ case. We can answer this by recalling from the previous chapter that $S = 0.8$, $E = 0.03$ stochastic REDS networks are composed of several unconnected components. Within some of these components, co-operation collapses entirely. Since, by definition, a component cannot be reached from another part of the network, there is no possibility for the co-operative strategy to re-enter the component at a later stage in the process. Thus co-operation never fully dominates the $b = 0$ case.

We further note that at the highest value of the temptation to defect that we have tested, $b = 0.3$, co-operation remains present in the $E = 0.03$ network at a value of $p_c = 0.05$. Conversely to the argument that we have just made, there are components of the network that are dominated by co-operators, even at this high temptation to defect, and so we observe a small fraction of co-operators at high b .

6.1.2.1.3 Intermediate E

For the intermediate energy, $E = 0.145$, we find that co-operation dominates the model for $b \leq 0.01$. As the temptation to defect increases we see a rapid drop in the proportion

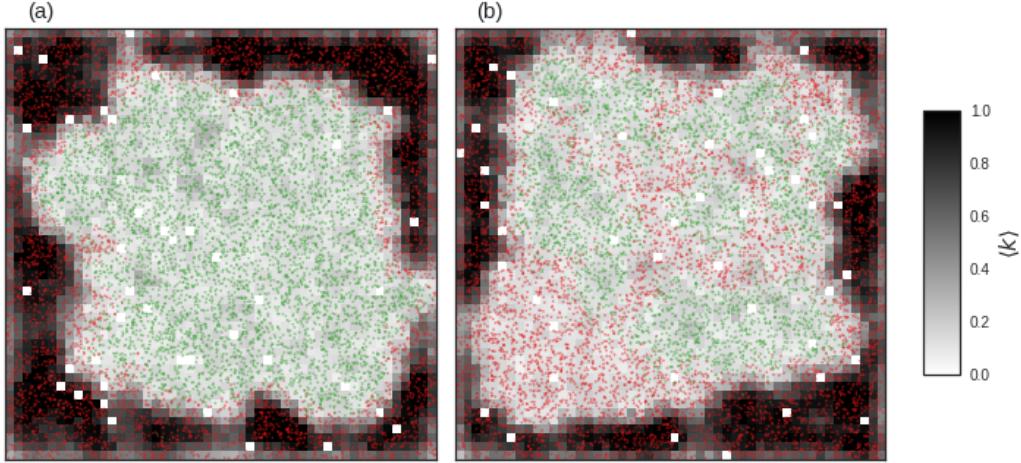


FIGURE 6.2: Visualisations of a evolutionary REDS network with $S = 0.8, E = 0.145, T_n = 10^{-2}, T_g = 0.2$ and (a) $b = 0.07$ (b) $b = 0.11$. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$.

of co-operation between $0.01 < b \leq 0.025$. In the range $0.025 < b \leq 0.1$, co-operation is relatively stable with values lying in the range $0.43 \leq p_c \leq 0.6$. Finally as the temptation to defect increases beyond $b > 0.1$ we see a second rapid collapse in the proportion of co-operation with $p_c = 0$ for all $b \geq 0.15$.

As a first step towards understanding the nature of these collapses we visualise two example $E = 0.145$ networks, one with $b = 0.07$ and the other with $b = 0.11$ in figure 6.2. These plots follow the same format as figure 5.11a, i.e., we use greyscale hue to indicate connectivity, and, in addition, we colour code nodes according to their strategy. We have chosen the value $b = 0.07$ so as to represent the stable, intermediate region and the value $b = 0.11$ to represent the second collapse.

We remind ourselves that, in section 5.3.3, we established that stochastic REDS networks with $S = 0.8, E = 0.145, T_n = 10^{-2}$ belong to the transition region of the $S \times E$ plane and that they consist of a regions of high, RGG like connectivity that bounds a lower connectivity, region. In figure 6.2a we see that defection dominates the high connectivity region while co-operation is the preferred strategy within the lower connectivity region. At the higher value of b , see figure 6.2b, defection is still dominant within the high connectivity region but now a significant number of defectors can be seen in the low connectivity region. Thus, the two separate collapses in the value of p_c observed along the $E = 0.145$ curve reflect the hybrid nature of the network. Reflecting what we have just established, co-operation in the RGG-like regions of the network collapses at low values of b . As we increase b beyond 0.1 we find that co-operation is progressively eliminated from the lower degree region as well.

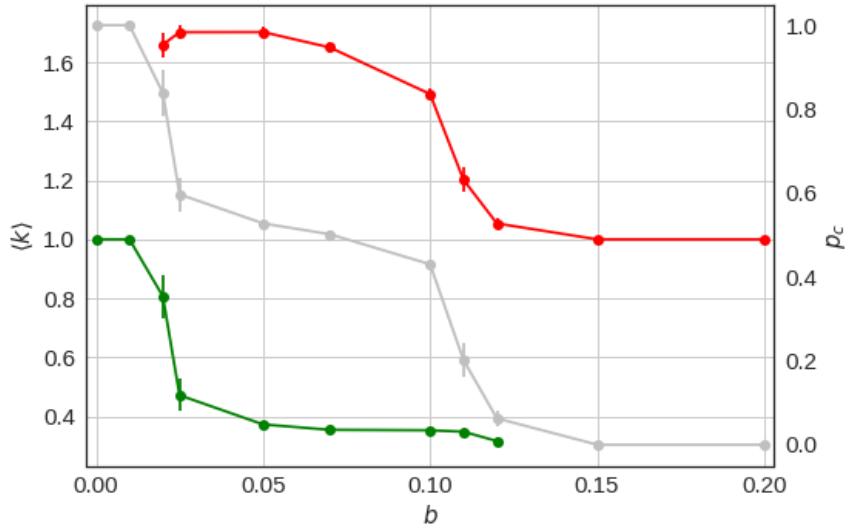


FIGURE 6.3: The quantities (i) $\langle k(c) \rangle$, (green) and (ii) $\langle k(d) \rangle$, (red) observed in the evolutionary REDS experiments with $E = 0.145, T_n = 10^{-2}, T_g = 0.2$ reported in figure 6.1a. The co-operator mean degree, $\langle k(c) \rangle$, is calculated for all co-operator nodes in the model and then normalised by the mean degree of all nodes in the model. Each data point is the the mean value of $\langle k(c) \rangle$ recorded for all the networks generated for the given value of b . A similar calculation, performed on defector nodes, gives $\langle k(d) \rangle$.

The grey plot is the p_c curve for $E = 0.145$ from figure 6.1a.

We confirm this account of the collapses by considering $\langle k(c) \rangle$ (resp. $\langle k(d) \rangle$), the mean degree of co-operator (resp. defector) nodes in the final epoch of the model. We normalise both $\langle k(c) \rangle$ and $\langle k(d) \rangle$ by the mean degree of the network in the final epoch. In figure 6.3 we plot the mean value of these quantities for the $E = 0.145$ experiments in figure 6.1a. For $b = 0.0, 0.01$ we find $\langle k(c) \rangle = 1.0$, since all the nodes are co-operators in these scenarios. In the $0.01 < b \leq 0.025$ range, the mean degree of co-operators rapidly drops to $k(c) = 0.47$, while the mean degree of defectors is around $k(d) \sim 1.7$, indicating that high degree nodes adopt the defect strategy for these low values of b . Some lower degree defectors appear as b increases and, consequently, in the range $0.025 \leq b \leq 0.1$ the mean degree of defectors then drops to $k(d) = 1.50$. Finally, all low degree nodes adopt the defect strategy as b is increased above 0.1 and we observe that the mean degree of defectors rapidly decreases to 1.0 in this region.

6.1.2.2 High strategy update temperature

In figure 6.1(b) we modify the parameterisation by raising the strategy update temperature to $T_g = 2.0$. The most obvious difference between this plot and figure 6.1(a) is that both of the $E = 0.03, 0.145$ curves have the same form as their $T_g = 0.2$ counterparts. However, for both curves, drops in the value of p_c occur at lower values of b . In

particular, bear in mind that the second collapse for the $E = 0.145$ networks was of the low degree co-operators.

6.1.2.3 Low network update temperature

Finally, in figures 6.1c & d, we present plots for a lower network update temperature, $T_n = 10^{-4}$ and strategy update temperatures $T_g = 0.2, 2.0$. The two plots are similar, co-operation has now completely collapsed for those $E = 0.145$ networks where $b \neq 0.0$. This is not surprising since $S = 0.8, E = 0.145$ networks lie well within the high connectivity region of the $S \times E$ heatmap for $T_n = 10^{-4}$. We have omitted the $S = 0.8, E = 0.24$ curve since the results are nearly identical to those of the $E = 0.145$ case. For $E = 0.03$, the curves in both plots are significantly lower than in the $T_n = 10^{-2}$ equivalents.

To summarise this section, we have found that co-operation is favoured within those regions of an evolutionary REDS network that exhibit low connectivity. This finding is in accord with previous claims within the literature. Increasing the strategy update temperature T_g undermines the capacity of these networks to support co-operation. Decreasing the network update temperature also decreases the capacity of the evolutionary REDS network to support co-operation. However, this observation follows from the fact that lowering T_n extends the high connectivity region of the $S \times E$ plane.

6.2 The coevolutionary REDS model

In this section we detail the model that is this chapter's main contribution. As discussed in the introduction, the coevolutionary REDS model is an adaptive network with endogenous connectivity whose nodes exhibit minimal cognitive ability. This minimal capability for cognition is realised in the process of edge formation which is exactly the one described for the REDS model in the preceding chapter, a process which is determined by an agent's capacity to maintain connections with its spatial neighbours.

We describe the coevolutionary REDS model in section 6.2.1 and then outline three canonical forms of its dynamics in section 6.2.2.

6.2.1 The model

We initialise the coevolutionary REDS model in the now familiar fashion; by placing n nodes, uniformly at random, upon a unit square. These nodes are assigned a default strategy in the same manner as the evolutionary REDS model.

Similarly to the stochastic REDS model, a node i has an energy budget, E_i , which is expended when the node forms an edge. We link the energy budget to the payoff from the game via the following equation:

$$E_i = lE + (1 - l)P_i \quad (6.4)$$

where E is a baseline allocation of energy and $l \in [0, 1]$ is a parameter which controls the relative amounts of baseline energy and payoff that compose the node's overall budget. Here, l can be thought of as an inherent capacity for nodes to form edges.

Where nodes play the PD game in the coevolutionary REDS model they accumulate the payoff described by equation (6.2). We choose this formulation of PD so that the amount of payoff available in equation (6.4) from a co-operate-defect edge is fixed at one unit, irrespective of the choice of b .

The coevolutionary REDS model is one with timescale separation, i.e., during an iteration of the model, either a topology or a strategy update occurs. These events occur asynchronously on two timescales; τ_n (network updating) and τ_s (strategy updating). We follow the scheme discussed in the literature and implement this concept by defining the ratio, $w = \tau_s/\tau_n$. At each iteration of the model a strategy update is performed with probability $(1+w)^{-1}$; otherwise a network update event takes place. As $w \rightarrow 0$, we recover a static network model; conversely as $w \rightarrow \infty$, network topology is updated more rapidly than node strategy. Where necessary, we refer to w as the timescale separation.

A network update event in the coevolutionary REDS model follows exactly the process of the stochastic REDS model, with the caveat that node i 's energy budget E_i is calculated using equation (6.4) above. A strategy update proceeds exactly as we have described for the evolutionary REDS model, i.e., we pick a pair of nodes, calculate their payoffs and modify the strategy subject to the Fermi function.

6.2.2 Coevolutionary dynamics

We have run the coevolutionary REDS model in a broad range of configurations and have identified three possible scenarios for the dynamics of p_c . In figure 6.4 we present an example plot of each case. The three scenarios are (i) domination, where all the nodes end up as co-operators, (ii) collapse, where all nodes end up as defectors and (iii) cyclic, where there is a pattern of rising and falling p_c with no fixed period. In the following sections we discuss an example of each of these scenarios.

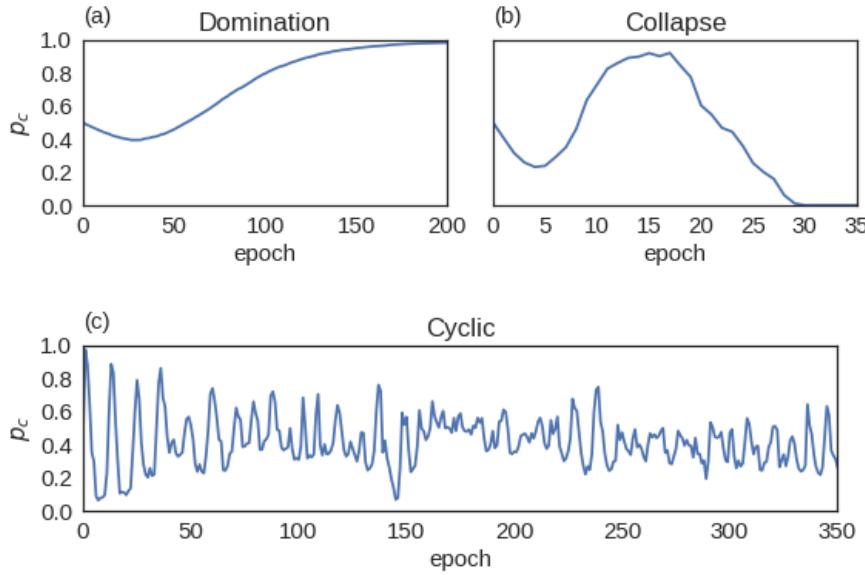


FIGURE 6.4: Dynamics of p_c in three canonical scenarios for the coevolutionary REDS model. They are (a) domination, (b) collapse and (c) cyclic. Model parameters in the first two examples are $T_n = 10^{-2}$, $s = 0.8$, $e = 0.24$, $T_g = 0.2$, $w = 10^{-3}$, $b = 0.12$, (a) $l = 0.05$ and (b) $l = 0.5$. Model parameters in the final example are $T_n = 10^{-2}$, $s = 0.8$, $e = 0.03$, $T_g = 2.0$, $w = 10^{-1}$, $l = 0.95$, $b = 0.036$.

6.2.2.1 Domination

Figure 6.4a plots the value of p_c over 200 epochs of a coevolutionary REDS model where $T_n = 10^{-2}$, $s = 0.8$, $e = 0.24$, $T_g = 0.2$, $w = 10^{-3}$, $l = 0.05$, $b = 0.12$. Initially we observe that p_c declines to a minimum by epoch 30. Subsequent to this the value of p_c rises so as to attain a value of 0.98 by epoch 200. Further experiments indicate that eventually all the nodes co-operate in models which exhibit similar dynamics to this one. For this reason we refer to this scenario as the domination example.

The initial decline in p_c can be explained by considering that, in the early development of the domination example's network, new edges connect nodes whose strategies are independent of each other. Furthermore, this new edge will connect two nodes that are otherwise unconnected. Therefore, where co-operator i connects with defector j , the resulting payoffs will be $P_i = -0.12$ and $P_j = 1.12$. Consequently, i will tend, by equation (6.3), to adopt the defect strategy and the payoffs become $P_i = P_j = 0$. Since neither i nor j now has access to much energy, edge (i, j) will tend to disappear during a network update. By this mechanism low degree defectors initially spread through the network and, after 30 epochs, the proportion of co-operators is $p_c = 0.39$ and mean degree is $\langle k \rangle = 0.01$. Figure 6.5a visualises a small portion of the network at this stage.

This early phase of the domination example's development also sees the creation of edges joining either pairs of defectors or pairs of co-operators. As discussed, edges between defector pairs tend to disappear. On the other hand, the payoffs are $P_i = P_j = 1$ in a

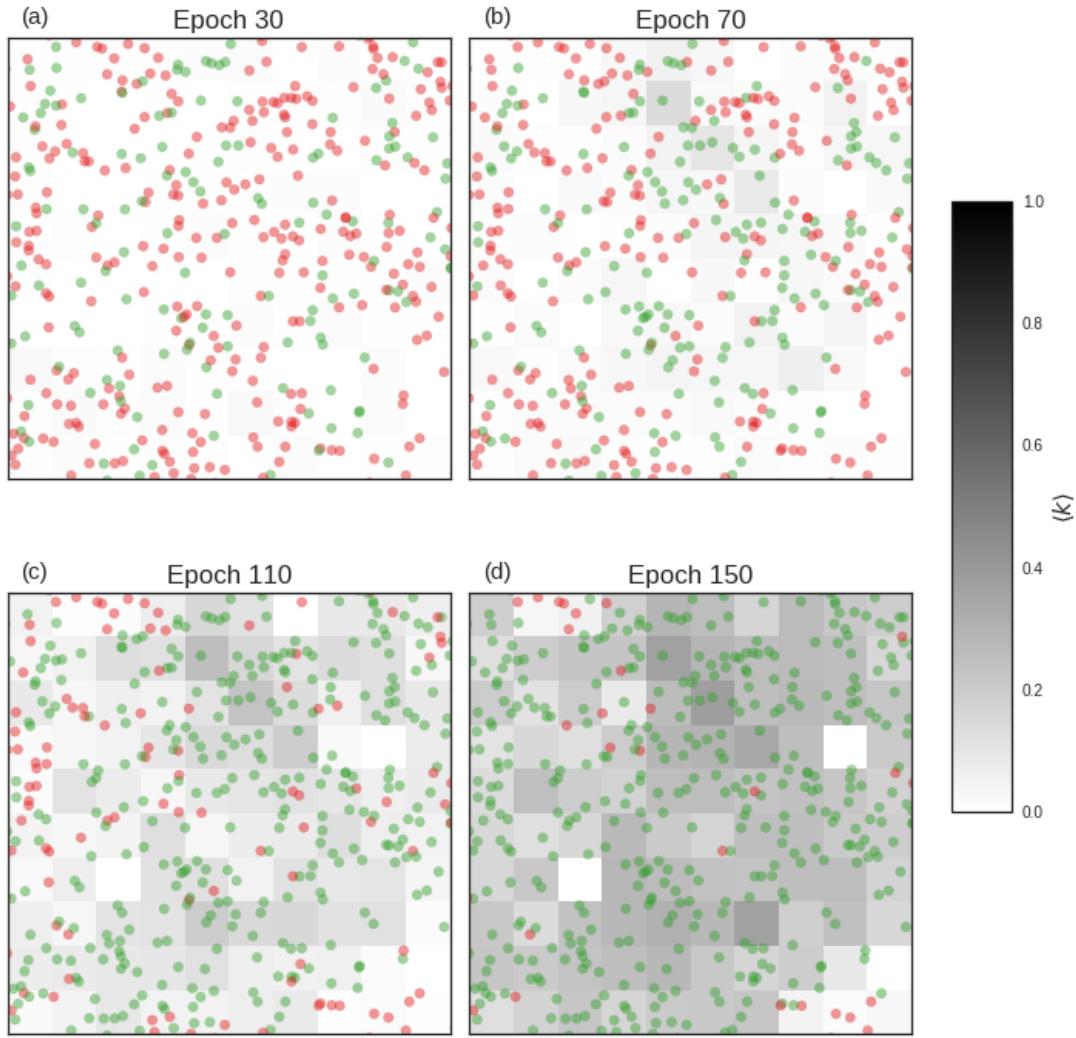


FIGURE 6.5: Node strategies in the domination coevolutionary REDS example (see figure 6.4a). The square of side 0.2 with bottom left corner at (0.1,0.5) has been visualised after epochs (a) 30, (b) 70, (c) 110 and (d) 150. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$.

co-operative pair, and so these edges will tend to perpetuate. Thus, clusters of connected co-operators form amongst the isolated defectors.

Furthermore, these clustered co-operators have an excess energy budget with which they can form connections with nearby defectors. Now, suppose one of these co-operators, i , now plays the game against a defector j . The payoff P_j is still boosted by the play of the defection strategy against co-operation. However, it also remains the fact that i plays the game with its multiple, co-operative neighbours; thus $P_i > P_j$ and j will be more likely to adopt a co-operative strategy. The outcome of these dynamics is visible in figure 6.5b which visualises the same portion of the network as figure 6.5a, only now after 70 epochs. The small pockets of co-operators increased their connectivity and in

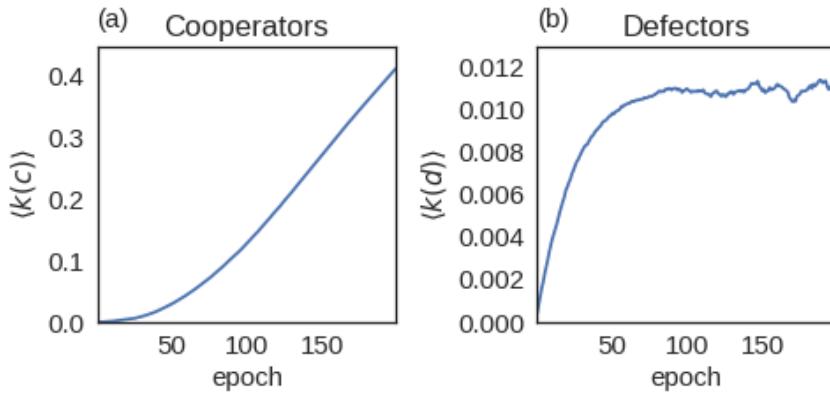


FIGURE 6.6: Mean degree of nodes with the (a) co-operate and (b) defect strategy over the lifetime of the coevolutionary REDS domination example.

some places are now converting nearby defectors. Considering the model as a whole, co-operation at epoch 70 is now $p_{c,70} = 0.6$.

Figures 6.5c (110 epochs) and 6.5d (150 epochs) demonstrate that this pattern continues for the remainder of the run. New regions of co-operation are established in the immediate vicinity of existing co-operators and co-operation increases in patches where there is a high proportion of co-operators. The statistics for co-operation are, at epoch 110, $p_{c,110} = 0.85$ and, at epoch 150, $p_{c,150} = 0.95$.

In summary, we have witnessed a process by which the co-operative strategy incrementally comes to dominate a small region of the model. The fundamental mechanism is that the co-operator nodes form local coalitions, thereby maximising their payoff. Members of a coalition can then readily ‘convert’ nearby defectors upon connection with them. We confirm that this explanation applies across the entirety of the domination example by examining the behaviour of $\langle k(c) \rangle$ and $\langle k(d) \rangle$. In this case, these quantities are again the mean degree of, respectively, co-operator nodes and defector nodes in the model. However, in contrast to section 6.1.2 we here normalise $\langle k(c) \rangle$ and $\langle k(d) \rangle$ by $\langle k_{RGG} \rangle$, the mean degree of the underlying RGG network.

Figure 6.6a is a plot of $\langle k(c) \rangle$, over the 200 epochs of the domination example. We find that, after a slow start, the mean degree of the co-operator nodes rises linearly to $\langle k(c) \rangle = 0.41$ by the end of the model run. Meanwhile, in figure 6.6b we see that the mean degree of the defector nodes rises logarithmically approaches a (noisy) maximum of just over 0.01 just after epoch 50.

In conclusion, these results are consistent with the process we have described. The logarithmic growth in defector connectivity occurs early in the simulation and is ultimately small in comparison to that achieved by the co-operators. The growth in the co-operator connectivity is slow initially since it takes time for co-operative clusters to form. However, once such clusters do appear they rapidly increase their connectivity.

6.2.2.2 Collapse

We now modify the parameterisation of the domination example so as to increase the ratio of baseline to payoff energy to $l = 0.5$. The dynamics of p_c for this example are presented in figure 6.4b where we see that co-operation drops to an initial minimum of $p_c = 0.22$ after 5 epochs before rallying to a maximum of $p_c = 0.92$ after 17 epochs. Following epoch 17, p_c falls steadily until it reaches zero at epoch 29. We refer to this scenario as the collapse example.

6.2.2.2.1 Dynamics prior to the collapse

Since l has now been increased from 0.05 to 0.5, node i 's energy budget, E_i , is, per equation (6.4), drawn more from the baseline energy allocation, E , than it is from the payoff, P_i . Therefore, relative to the domination example, edges form more easily in this example. In the initial phase of this example, defector/co-operator pairs appear more rapidly than in the domination example and the co-operators are converted to defectors. Hence, p_c reaches a local minimum more rapidly than in the domination example.

Increased baseline energy also supports the formation of defector clusters. Members of these clusters more effectively resist conversion to the co-operator strategy for the following two reasons: firstly, where defector i is selected for strategy update; its neighbour j is likely to be another defector, in which case the Fermi update rule dictates that i remains a defector. Secondly, should defector i be converted, any neighbouring defectors have their payoff boosted by the temptation to defect. Conversely, i now earns the sucker payoff for its interaction with these defectors. Over time, then, i will tend to revert back to the defect strategy.

Nonetheless, co-operative clusters still form in this scenario. However, since the defectors now resist conversion, we find defector co-operator pairings that persist. These defector nodes receive a high payoff from the PD game and can convert nearby co-operators. It is this phenomenon that is the trigger for the collapse in co-operation.

We quantify these statements by considering $\langle k(c)_d \rangle$, the mean number of connections that start from a co-operator node and terminate at a defector, and $\langle k(c)_c \rangle$, which is the mean number of connections that start from a co-operator node and terminate at another co-operator. To allow for comparison with the $\langle k(c) \rangle$ reported for the domination example, we normalise $\langle k(c)_d \rangle$ and $\langle k(c)_c \rangle$ by $\langle k_{RGG} \rangle$.¹

These statistics are plotted in figure 6.7a. Here we see that, up to epoch 17, $\langle k(c)_c \rangle$ rises linearly after a slow start; a similar pattern to that of $\langle k(c) \rangle$ in figure 6.3b. Meanwhile, $\langle k(c)_d \rangle$, is low during this period, demonstrating that the majority of connections that

¹The merit of this comparison is that $\langle k(c) \rangle = \langle k(c)_d \rangle + \langle k(c)_c \rangle$

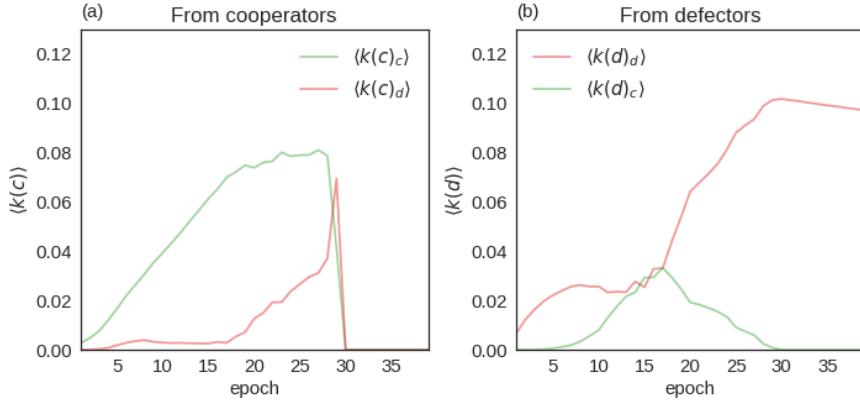


FIGURE 6.7: In (a) we plot the mean number of connections from a defector to (i) other defectors, $\langle k(d)_d \rangle$ (red), and (ii) co-operators $\langle k(d)_c \rangle$ (green). Similarly, in (b) we plot the mean number of connections from a co-operator to (i) other co-operators $\langle k(c)_c \rangle$ (green) and (ii) defectors $\langle k(c)_d \rangle$ (red).

start from co-operators connect to other co-operators. Thus, we conclude that, up to epoch 17, the co-operative nodes form clusters amongst themselves in much the same fashion as the domination example.

We similarly investigate the connectivity of defector nodes by considering $\langle k(d)_c \rangle$; start from a defector, end at a co-operator; and $\langle k(d)_d \rangle$; start from a defector, end at a defector. Figure 6.7b plots these two quantities. Up to epoch 5, $\langle k(d)_d \rangle$ rises to a value of 0.02 while $\langle k(d)_c \rangle$ remains negligible; confirming that defectors connect to and immediately convert co-operators in the early stages of this example. Note also that, in figure 6.3b, $\langle k(d) \rangle$ stabilised at around 0.01 for the domination example. Therefore, even at this early stage, defector nodes exhibit higher connectivity amongst themselves than in the domination example.

Beyond epoch 5 and up to epoch 16, the value of $\langle k(d)_d \rangle$ remains relatively stable at just over 0.02. Meanwhile, $\langle k(d)_c \rangle$ rises in this period, indicating that connections are being formed between defectors and co-operators. This second point is the crucial difference between the collapse and domination examples. We have found $\langle k(d)_c \rangle$ to be negligible in the domination example; reflecting the fact that defectors are near immediately converted after connecting to co-operators. Thus, we conclude that the accumulation of relatively high levels of defector to co-operator connections differentiates the collapse from the domination example and this is what triggers the fall in co-operation that begins at epoch 17.

6.2.2.2.2 Dynamics after the collapse

We now turn our attention to the nature of the collapse itself. In particular we wish to understand what drives the model to total collapse after epoch 17, instead of recovering

in the same fashion as observed after epoch 5. Returning to figure 6.7b we see that $\langle k(d)_c \rangle$ falls after epoch 17. This fall is not instantaneous indicating that in the early stages of the collapse the defectors press the advantage of a high number of outgoing connections to co-operators.

In the latter stages of the collapse $\langle k(d)_c \rangle$ falls quite significantly. However, in this period of the example $\langle k(c)_d \rangle$ rises noticeably. In words, individual defectors do not make multiple connections with co-operators as frequently as they did prior to the collapse. However, defector nodes are now better coordinated so as to connect to individual co-operators en masse. These co-operators receive multiple sucker payoffs from the PD game and so they convert to the defect strategy. It is this effect that leads to the completion of the collapse.

To summarise, the fundamental difference between the collapse and domination examples is the availability of baseline energy. In the domination example, baseline energy was in short supply, with the consequence that co-operators are better placed to exploit their mutually beneficial payoff via coalition formation. These coalitions go on to systematically convert nearby defectors. In the collapse example the introduction of baseline energy allows defectors to more easily form connections thereby forming their own coalitions. This leads to the situation that, in epoch 17, the following two ingredients combine to trigger the collapse: (i) defector clusters (ii) defector nodes with multiple connection to co-operators. Once the collapse is set in motion a third ingredient was added to this mix, that of co-operator nodes with multiple connections to defector nodes. It is this last component that fatally weakens the co-operator nodes, leading to their total collapse.

6.2.2.3 Cyclic

The dynamics of our final scenario are presented in figure 6.4c where the overall picture is that of rising co-operation that reaches a peak and then immediately drops. However, unlike the previous example, p_c never completely collapses; instead co-operation revives and the sequence begins again. This pattern of rising and falling p_c does not exhibit a constant period and we follow the convention from time series analysis and refer to this scenario as the cyclic example (Hyndman, 2011).

Initial investigations indicate that cyclic dynamics do not readily occur and so the parameters that induce the cyclic example differ significantly from those used for the domination and collapse examples; they are: $T_n = 10^{-2}$, $s = 0.8$, $e = 0.03$, $T_g = 2.0$, $w = 10^{-1}$, $l = 0.95$, $b = 0.036$. A consequence of modifying the parameters to this extent is that network connectivity and the cost benefit ratio of co-operation are markedly

different in this example when contrasted with those of the domination and collapse examples. As such, we cannot easily establish a basis for quantitative comparison between the three examples.

Nonetheless, visualisation of this time series indicates consistent patterns in its dynamics. Therefore, we will now give a qualitative description of an example of a rise and fall of p_c in the cyclic example's dynamics. For the remainder of this section we focus on epochs 54 to 65. In this period co-operation starts from a low of $p_c = 0.24$ and ascends to $p_c = 0.74$ at epoch 59. From this peak co-operation falls to $p_c = 0.24$ at epoch 65. We examine the rise and fall in their own sections. We also dedicate a section to demonstrating the phenomenon of transmission, i.e., a process by which a strategy can take root in a region of the square that is at some distance from an existing cluster of that strategy.

6.2.2.3.1 Rising co-operation

Figure 6.8 is a series of visualisations of the network between epochs 54 and 57, the period during which p_c is increasing. We colour code each node according to whether it is a co-operator (green) or a defector (red). Since the edge density is high in many portions of the network we do not render individual edges; instead we break the image into a 50×50 grid of squares of side 0.02 shade each square so as to represent its connectivity.

At epoch 54 there are two high connectivity clusters at the top left and bottom right of the network. These clusters are characterised by a core of co-operators and a boundary that consists of both co-operators and defectors. Both the co-operative core and the boundary have very high connectivity. The next epoch, 55, is almost exactly the same, with the exception that a small cluster of low connectivity co-operators appears on the middle right hand side of the network around the area marked by the blue square.

This small group of co-operators can now spread their strategy in a similar fashion to the domination strategy. Thus, we find in epochs 56 and 57 that co-operation spreads rapidly throughout the network. Also of note in epoch 57 is that a small group of co-operators is apparent in the bottom left of the square centred on the point (0.1, 0.3).

6.2.2.3.2 Transmission

We have observed the appearance of small groups of co-operators at some distance from existing communities of co-operators on two of the four epochs in the previous section. We gain some insight into how the community that was observed in epoch 55 appears by visualising the blue square from figure 6.8b at intermediate moments between epochs 54 and 55 in figure 6.9.

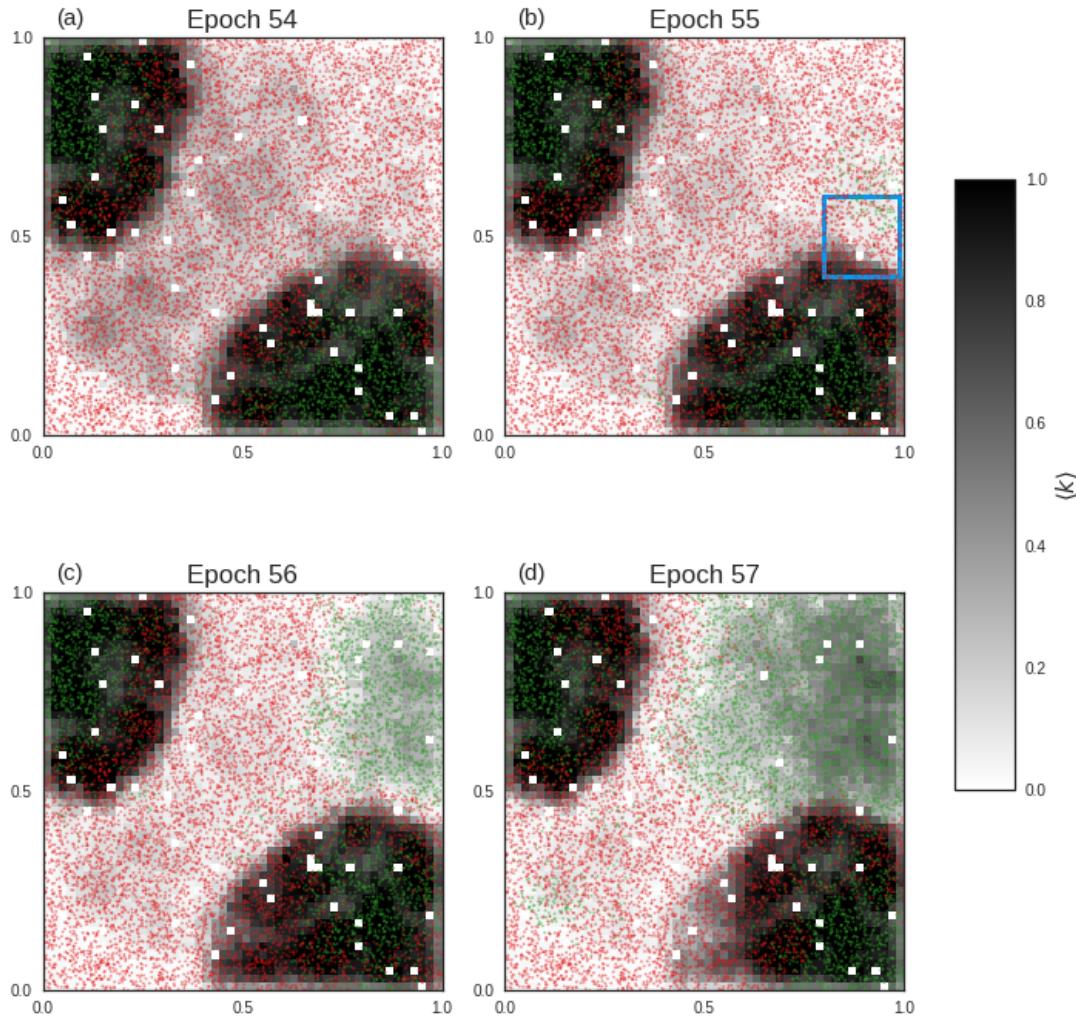


FIGURE 6.8: Node strategies in a growing phase of the cyclic example. The entire network has been visualised after epochs (a) 54, (b) 55, (c) 56 and (d) 57. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$. The blue box in (b) highlights a region that is discussed in more detail in the text.

At epoch 54.1, we find the border of the highly connected cluster at the bottom of the square. A few co-operator nodes lie beyond the cluster on the left of the image. This leads to the situation in epoch 54.3, where there are 7 co-operators within the top half of the square. None of these co-operators are within a distance of 0.05 of any other co-operator in the square, indicating that co-operation has arrived in this region via nodes that have subsequently converted back to the defect strategy. Finally, in epoch 54.5, the co-operators successfully form a cluster at the top of the square.

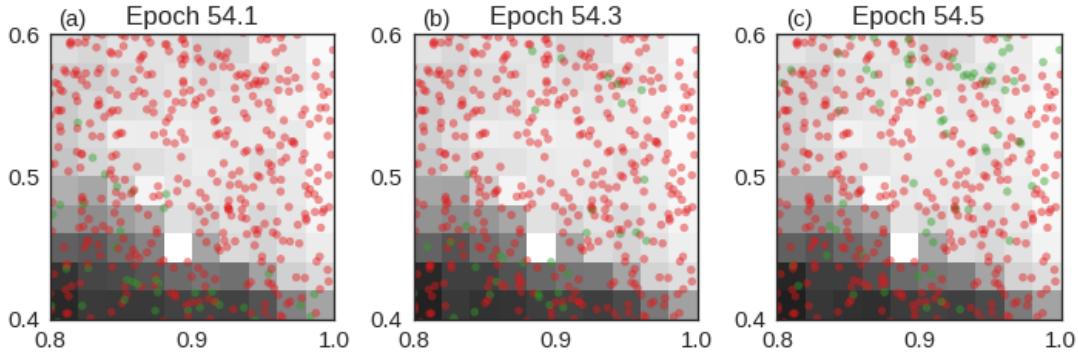


FIGURE 6.9: Visualisations of the cyclic example at epochs (a) 54.1, (b) 54.3 and (c) 54.5 within the square of side 0.2 with bottom left corner at (0.8, 0.4). This is the region highlighted by the blue square in figure 6.8b.

6.2.2.3.3 Falling co-operation

In figure 6.10 we present visualisations of the network between epochs 59 and 65, the period during which p_c is in decline. At epoch 59, the growth of the network reaches its apotheosis and the entire top right of the network is now a highly connected co-operative cluster. Furthermore, the small number of co-operators in the bottom left corner of the square has now grown significantly. Also noticeable, in the centre of the square, is a high connectivity patch consisting of both co-operator and defector nodes.

At epoch 61 we find that the defect strategy now begins to re-establish itself in two regions, around the points (0.5, 0.2) and (0.5, 0.75). This process continues in epoch 63 where the influence of the defect strategy is extended in two ways. Firstly nodes have been converted to the defect strategy across the entire top right of the plot. Secondly the defect strategy successfully invades the bottom right of the plot. The overall outcome of these developments is that a new highly connected cluster remains in the middle right of the square. Also notice that another highly connected cluster has formed in the bottom left of the square. Both of these new clusters share the features of the high connectivity clusters discussed previously, a co-operative core and a mixed strategy boundary.

Finally, in epoch 65, the cluster on the right of the square collapses and we reach a minimum of co-operation. In those areas where defection dominates, connectivity has now died away, leaving a region that can readily be converted by the co-operate strategy. This results in a network which has two high connectivity regions that can now seed the low connectivity region of defectors with the co-operate strategy, thereby beginning the sequence again.

This sequence of co-operation growth and collapse is underpinned by three key elements. Firstly, the high connectivity clusters allow for the preservation of the co-operate strategy. Secondly, transmission of the co-operation strategy means that high connectivity

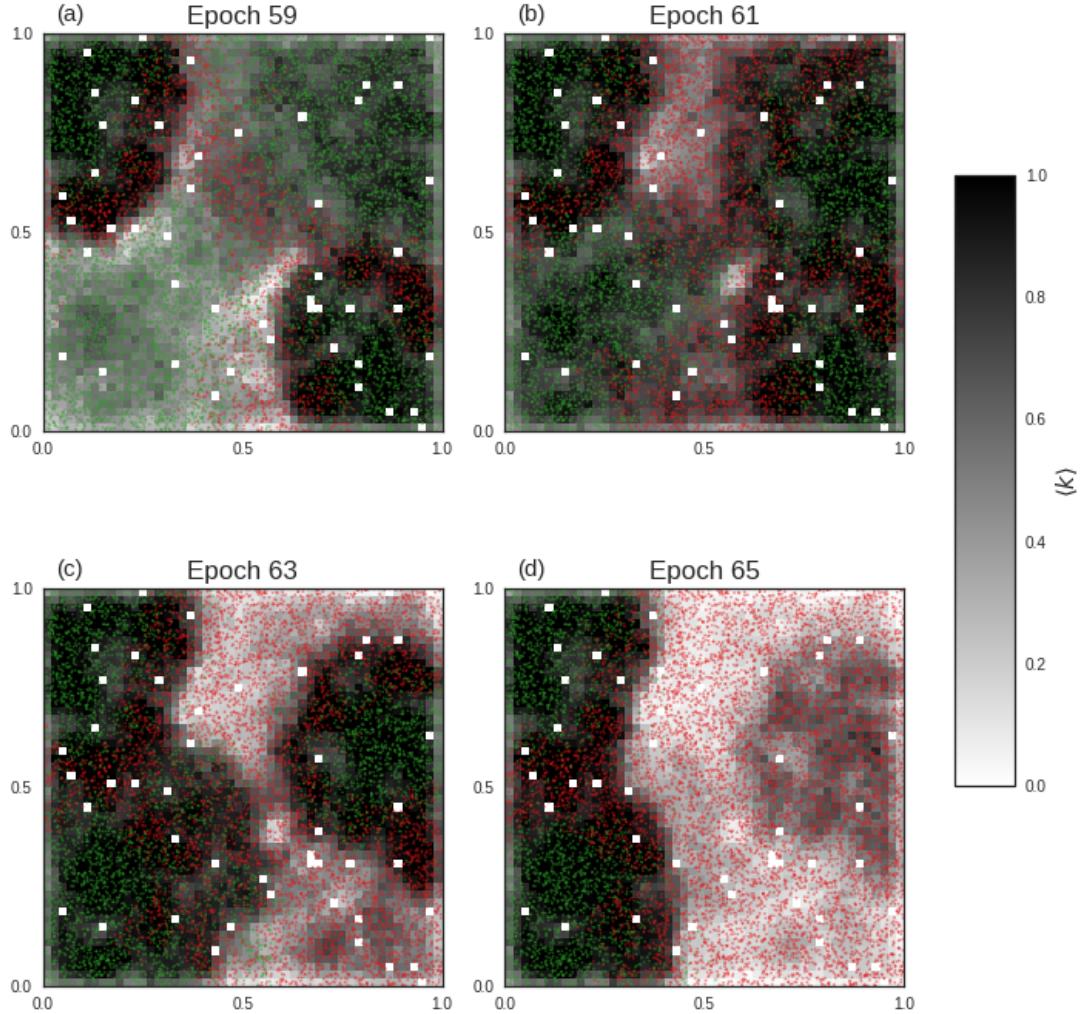


FIGURE 6.10: Node strategies in a collapsing phase of the cyclic example. The entire network has been visualised after epochs (a) 59, (b) 61, (c) 63 and (d) 65. Green dots are co-operators, red dots are defectors. The greyscale hue of the smaller squares represents $\langle k \rangle$ within that patch with darker hues representing higher $\langle k \rangle$.

clusters can propagate co-operation to regions of the network that are dominated by low connectivity defectors. Finally, while the high connectivity clusters appear to be robust over time, they are not infallible. Thus, defector nodes within the border regions can successfully initiate an invasion that leads to a collapse of co-operation within the high connectivity regions.

6.3 Parameter variation

In this section we more thoroughly assess the parameter space of the coevolutionary REDS model. Initially, we investigate the effect that variation of the baseline energy,

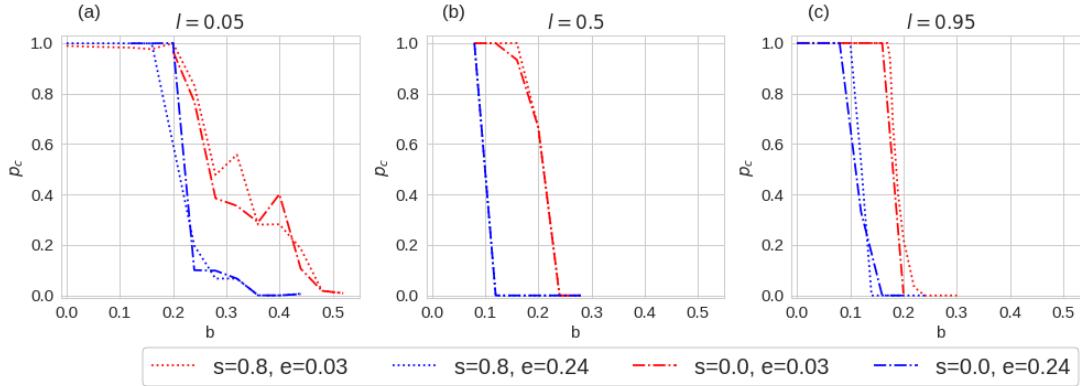


FIGURE 6.11: Plots of p_c vs b for the evolutionary REDS model with $l = 0.05, 0.5, 0.95$. For each l plots are for all combinations of $S = 0.0, 0.8$ and $E = 0.03, 0.24$ (see legend). Parameters $T_n = 10^{-2}, T_g = 0.2, w = 10^{-3}$ are fixed. For each data point, the standard error is below 0.1.

E , and synergy, S , have upon co-operation within the coevolutionary REDS model. In this regard, we remind ourselves that equation (6.4), introduced a weighting, l , between baseline energy and payoff. Energy is a node's inherent capacity to make connections while payoff is earned during the course of the simulation. We assess l alongside E and S in section 6.3.1, finding that low E and low l tend to promote co-operation while variation of S has little effect on the outcome.

In section 6.3.2 we turn our attention to the network update temperature T_n , the strategy update temperature, T_g and the timescale separation w which we, somewhat informally, refer to as the update parameters. We investigate the effect of the update parameters upon co-operation in section 6.3.2. We find that decreasing w , i.e., updating strategy more frequently, and lowering T_g , i.e., making the strategy update more deterministic, both promote co-operation.

6.3.1 The effect of synergy and energy upon co-operation

Our aim in this section is to investigate how synergy, S , and baseline energy, E , affect the level of co-operation in the coevolutionary REDS model. The parameter l governs the amount of baseline energy available to the nodes. Furthermore, the domination and collapse examples demonstrated that availability of baseline energy is a key driver of outcomes in the coevolutionary REDS model. Hence, it is appropriate to establish l 's impact upon co-operation at this stage. Our approach follows that of section 6.1.2, i.e., we run a given $l \times S \times E$ parameterisation of the coevolutionary REDS model against an increasing temptation to defect so as to establish the ranges of b in which co-operation persists.

In the domination and collapse examples, we used $l = 0.05, 0.5$, synergy $S = 0.8$ and baseline energy of $E = 0.24$. We broaden our exploration of $l \times S \times E$ space by choosing

$l = 0.05, 0.5, 0.95$, $S = 0.0, 0.8$ and $E = 0.03, 0.24$. Update parameters are fixed at network update temperature, $T_n = 10^{-2}$, strategy update temperature, $T_g = 0.2$, and timescale separation, $w = 10^{-3}$.

For an $l \times S \times E$ combination, the temptation to defect is varied between $b = 0.0$ and $b = 0.52$. We run sufficient instances of the model for each value of b so as to ensure that the standard error of p_c is less than 0.1. Irrespective of the standard error, at least 15 instances of each case is run. An instance of the model is run for 800 epochs. The value of p_c for an instance is calculated by taking the mean value of p_c observed over the last 50 epochs.

In figure 6.11 we present the resulting p_c vs. b plots for these experiments. It is clear that the low energy networks better support co-operation than the high energy ones, irrespective of the value of l , i.e., the red curves are further to the right than the blue curves in each plot. This finding is unsurprising given our discussion of the collapse and dominate examples. Quite simply, the more baseline energy that is present, the more easily the defection strategy can undermine the co-operators.

More surprising is how closely each $S = 0.0$ curve matches its $S = 0.8$ pair. In the domination example, co-operators locally formed coalitions which then connected to and converted nearby defectors, thereby extending the co-operative cluster. One might expect, since co-operators more readily form mutual connections, high synergy regimes to enhance this mechanism. This result, however, indicates that it is sufficient for co-operators to be spatially co-located in order to efficiently form coalitions.

When we compare results between different values of l it is clear that co-operation better prevails within the $l = 0.05$ networks (compare figure 6.11a with b & c). This result corroborates our analysis of the domination and collapse examples; recall that the sole change in the parameters between the two scenarios was an increase in l from 0.05 to 0.5. This is effectively an increase in baseline energy and results in the same outcome.

6.3.2 The effect of update parameters upon co-operation

In this section we turn our attention to how the update parameters affect co-operation in the coevolutionary REDS model. To motivate our choice of parameters we highlight an issue with the cyclic dynamics. Experimentation reveals that the cyclic dynamics are repeatedly observed when the model is initialised with the parameters discussed in section 6.2.2.3. However, small adjustments to these parameters result in dynamics that match those of either the domination or collapse examples. Therefore, alongside the key goal of investigating conditions under which co-operation prevails in the model, we also wish to establish in what regions of the parameter space cyclic dynamics are present.

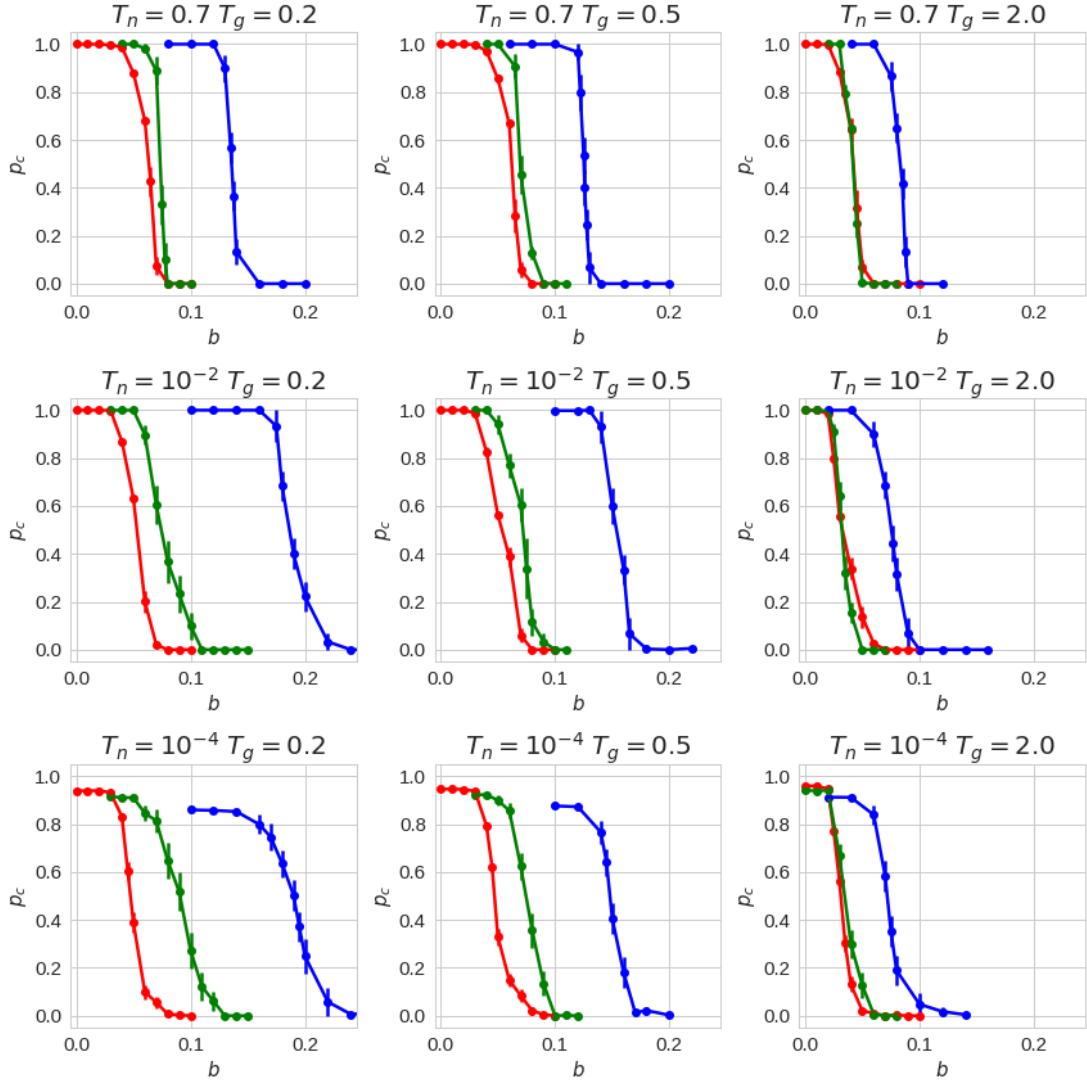


FIGURE 6.12: Plots of p_c vs b for the evolutionary REDS model with $w = 10^{-3}$ (blue), 10^{-1} (green) and 10^0 (red). Individual plots represent a $T_n \times T_g$ pair drawn from $T_n = 10^{-4}, 10^{-2}, 0.7$ and $T_g = 0.2, 0.5, 2.0$. Error bars represent the standard error.

In light of this latter goal, we fix $l = 0.95, S = 0.8, E = 0.03$, the same values that were used in the cyclic example. The three update parameters in the cyclic example were $T_n = 10^{-2}, T_g = 2.0, w = 10^{-1}$. We choose $T_n = 10^{-4}, 10^{-2}, 0.7$ for the network update temperature and $w = 10^{-3}, 10^{-1}, 10^0$ for the timescale separation so as to explore the immediate vicinity of the cyclic example. Finally, exploratory experiments indicate that lowering T_g is more conducive to co-operation, therefore we choose $T_g = 0.2, 0.5, 2.0$ as strategy update temperatures.

We follow the familiar pattern of running the coevolutionary REDS model for each $T_n \times T_g \times w$ combination and with varying values of b . We run each individual experiment for 400 epochs and calculate p_c the mean value of over the last 50 epochs. We again follow our convention of calculating p_c as being the mean value of p_c over the last 50 epochs of the model run. Where the standard deviation of p_c over the last 50 epochs is

greater than 0.01 we run the model for a further 400 epochs. We continue this process up to a maximum of 2400 epochs.

In the vast majority of cases, network dynamics follow the form of either the collapse or the domination example and do not equilibrate to intermediate values of p_c . There are two exceptions to this rule; firstly, some experiments result in cyclic dynamics. We will review this outcome in more detail shortly. Secondly, in some experiments with low network update temperature, $T_n = 10^{-4}$, the dynamics follow the domination example without co-operation reaching $p_c = 1.0$. This outcome occurs for the straightforward reason that the time required for co-operation to dominate these networks is considerably longer than 400 epochs.

Figure 6.12 presents the results of our experiments in a $T_g \times T_n$ grid of 9 plots, i.e., an individual plot represents a (T_g, T_n) pair. Within each plot there are three curves, one for each value of w . The most obvious feature of these plots is that, in all cases, the blue curves are the furthest to the right; indicating that the low timescale separation, $w = 10^{-3}$, is the most conducive to the perpetuation of the co-operative strategy.

Along the top row and rightmost column of the plots the red and green curves are nearly equivalent to each other. As we move towards the bottom left of the grid we find that, in general, the red curve ($w = 10^0$) is positioned somewhat to the left of the green curve ($w = 10^{-1}$). The interpretation of these observations is that increasing the timescale separation continues to adversely affect co-operation in the model, however this effect is only noticeable when both network and strategy update temperature is low.

In light of this analysis, we conclude that it is at low time separation that we best discern how variation of the update temperatures influences co-operation. Of the two we first address the strategy update temperature. The plots along the bottom row of figure 6.12 fix the network update temperature at $T_n = 10^{-4}$. As we scan from left to right along this row we see that the blue curve shifts leftwards within the plot, in other words co-operation is less well supported as strategy update temperature increases. The same pattern is apparent in the middle ($T_n = 10^{-2}$) and top ($T_n = 0.7$) rows.

We now shift our focus to the network update temperature. Along the top row of the plot, at the high network update temperature, the blue curves all exhibit near immediate transitions from $p_c = 1.0$ to $p_c = 0.0$. By comparison, the curves on the two lower rows are more gradual. The two lower rows are also similar in that their transitions occur for the similar values of b . The difference between the two cases is, as was discussed above, is that the transition begins from a value around $p_c = 0.87$ in the $T_n = 10^{-4}$, due to the longer time required to equilibrate the lower network update temperature model.

Considering the effect of the network update temperature on co-operation, it seems clear where $T_g = 0.2, 0.5$ co-operation is better supported by lowering the network update temperature. At the high value of the strategy update temperature, this pattern is

reversed, i.e., the blue curve in the $T_n = 0.7, T_g = 2.0$ plot is further to the left than in the other two $T_g = 2.0$ plots. However, we make two points in this regard; firstly, this effect is not especially strong and secondly we reiterate that co-operation is not well supported at the higher strategy update temperature.

We interpret these results as follows: firstly, lowering the network update temperature improves the level of co-operation in the model. This finding is consistent with the descriptions we gave of p_c dynamics in the domination and collapse examples. Since energy is most likely to be accumulated by mutual co-operators, low T_n ensures that these nodes are rewarded with an increased capacity to form edges. As T_n increases, so does the capacity for defectors to form edges.

Secondly, co-operation is best supported at the lowest value of the timescale separation considered, indicating that a more static network structure favours co-operation. It seems that this feature complements the improved co-operation for low T_n .

Finally, lowering strategy update temperature is also favourable to co-operation. For low T_g strategy update is more consistent and so co-operator coalitions will more reliably convert defectors. Conversely, increasing T_g will increase the possibility of defectors appearing randomly within co-operative clusters.

6.3.3 The effect of update parameters upon cyclic dynamics

We now discuss the presence of cyclic dynamics within our experiments. Our first task in this regard is to classify instances of the model with cyclic dynamics. We have seen that the cyclic scenario has an incessant sequence of peaks and troughs in its dynamics. By contrast, the domination and collapse dynamics settle to a constant value of p_c during the closing epochs of the model run. We therefore employ a simple method of peak counting so as to identify instances of the model which exhibit cyclic dynamics.

Our algorithm has four steps, the first of which is to smooth the series of values of p_c . Secondly, we identify the maxima and minima in the closing stages of the smoothed series. Thirdly, count those extrema that differ by a specified amount and finally we choose those examples where a sufficient number of peaks have been detected. The precise description of the algorithm is as follows:

1. Calculate the series $p_{c,5}$, which is, simply, the rolling average of p_c over a window of 5 epochs.
- 2a. Calculate, $dp_{c,5}$ the difference in $p_{c,5}$ between consecutive epochs.
- 2b. If $dp_{c,5} > 0$ (resp. $dp_{c,5} < 0$) at step t and $dp_{c,5} < 0$ (resp. $dp_{c,5} > 0$) at step $t + 1$ then mark epoch t as a maximum (resp. minimum).

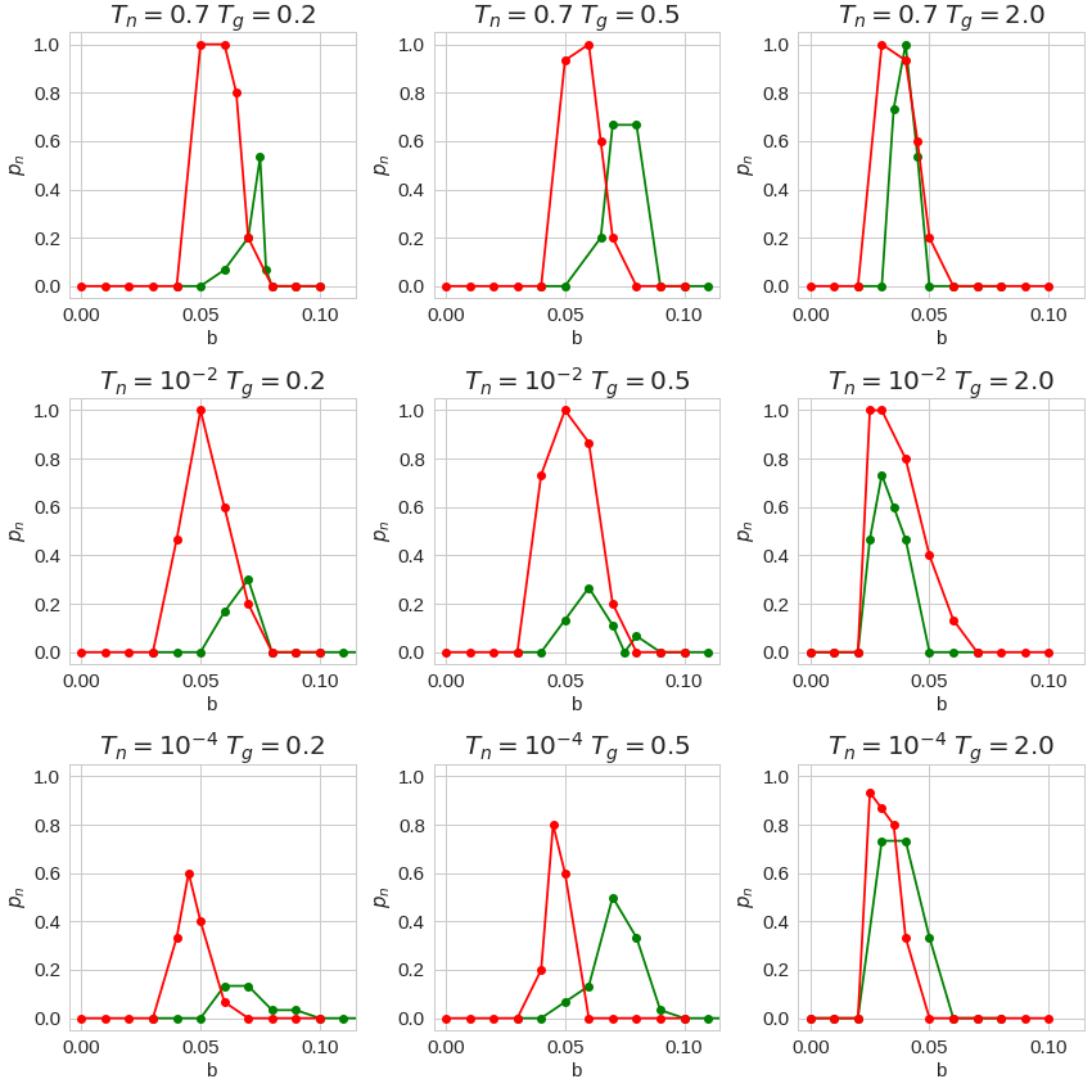


FIGURE 6.13: Plots of p_n for each of the data points in figure 6.12. Here, p_n is the proportion of experiments that exhibit cyclic dynamics. As in figure 6.12 the curves are colour coded $w = 10^{-1}$ (green) and 10^0 (red). Results for $w = 10^{-3}$ have not been plotted since $p_n = 0$ for all points with this timescale separation.

3. Within the last 50 epochs of $p_{c,5}$ count the number of instances where the absolute difference between consecutive extrema is greater than 0.05.
4. When this count is greater than 4 we classify the dynamics as being cyclic.

For each data point in figure 6.12 we use the foregoing algorithm to count the number of experiments with cyclic dynamics. We then calculate p_n , the proportion of experiments that have cyclic dynamics for that data point. In figure 6.13 we plot p_n for the $w = 10^{-1}, 10^0$ experiments encountered in figure 6.12. We have not plotted results for the $w = 10^{-3}$ experiments since $p_n = 0$ for all points with this timescale separation.

We highlight two general points about these plots; firstly, cyclic dynamics are present within the transitions of all the $w = 10^{-1}, 10^0$ curves. Secondly, at the higher timescale separation, $w = 10^0$ (red curves), cyclic dynamics are more prevalent. Examining the lower timescale separation, $w = 10^{-1}$ (green curves), more closely, we see that reducing either the network or strategy update temperature reduces the proportion of experiments that exhibit cyclic dynamics.

Returning to the red curves, we see that, along the bottom row of figure 6.12 decreasing the strategy update temperature reduces the incidences of cyclic dynamics. Raising the network update temperature to $T_n = 10^{-2}, 0.7$ maximises the the value of p_n , i.e., there is always a high prevalence of cyclic dynamics within the transition regions of the $w = 10^0$ experiments.

In summary, raising any of the three update parameters increases the likelihood that cyclic dynamics will be observed in the transition region of the associated p_c vs. b curve. This finding throws light on the transmission phenomenon discussed in section 6.2.2.3. Increased network update temperature implies increased independence of edge formation from a node's current energy balance. This allows for the appearance of temporary connections within the defector dominated, low connectivity regions of the network. These links allow for the possibility of transmission within the low connectivity region. Increased time separation has the consequence that these connections are long lived enough for their end nodes to be the subjects of strategy update events. Finally, increased strategy update temperature allows the co-operator strategy to overcome the advantage that defectors have in defector/co-operator pairs.

This account of the update parameters effect upon cyclic dynamics would appear to contradict the summary at the end of section 6.3.2. There, we argued there that low network connectivity favoured co-operators whereas the argument just given can be read as implying that it is, in fact, the case that increased connectivity now favours co-operators. However, we stress that the processes outlined in this section describe the formation of temporary pathways within the low connectivity region rather than a substantial increase in overall connectivity.

6.4 Summary

In this chapter we have extended the stochastic REDS model so that its nodes play prisoner's dilemma (PD) with each other. We began, in the evolutionary REDS model, by playing PD on static networks generated by the stochastic REDS model. In agreement with the predictions of Ohtsuki et al. (2006), two results in this set of experiments indicated that co-operation was best supported in regions of the network where connectivity was low. Firstly, higher levels of co-operation were found in the $E = 0.03$ evolutionary REDS networks than in the $E = 0.24$ case for all but the lowest values of the temptation

to defect. Decreased co-operation in the $E = 0.03$, low b experiments was due to the fact that these networks consist of multiple components. Some of these components were taken over by defectors early in the simulation and, hence, could not be reclaimed by co-operators. The second instance of low connectivity favouring co-operation was seen in the collapse of p_c for $E = 0.145$ evolutionary REDS networks. This collapse occurred in two stages as the temptation to defect increased; the first at low b for high connectivity nodes and the second at higher b for low connectivity nodes. In all our experiments with evolutionary REDS networks we found no examples where co-operation dominated the network for a temptation to defect $b > 0.01$.

We then introduced the coevolutionary REDS model, where the payoff from the PD game was used by the nodes as an energy source with which to form edges. Our key finding was that the capacity of these networks to support co-operation was improved by reducing the availability of baseline energy to the nodes, thereby maximising the advantage to co-operators from forming mutually beneficial links. In a further series of experiments we found that, in scenarios where the majority of energy available to a node was drawn from the payoff, i.e., where $l = 0.05$, co-operation dominated the network for $b \leq 0.2$. Thus, we conclude that the introduction of the coevolutionary rule significantly improves the uptake of the co-operative strategy when compared with the evolutionary REDS model.

Prior examples of coevolutionary models assume that their agents can distinguish between the separate interactions that they participate in. For example, in one of the fundamental coevolutionary models discussed section 2.3.3, that of [Zimmermann et al. \(2004\)](#), we find that the agents in the model are aware of the strategy of individual network neighbours. By contrast, agents in the coevolutionary REDS model can, at best, be described as assessing their accumulated payoff from playing the dilemma game with their neighbours. As such, our results demonstrate that co-operation can be supported in a model in which agents do not examine the details of specific interactions.

We found that co-operation was best supported in simulations where both the energy, E , and the parameter l were low. Where this was the case, the payoff, per equation 6.4, was the chief source of energy with which nodes could form connections. Under these conditions mutual co-operators can more easily use the higher payoff from the evolutionary game to form links. The spatially constrained nature of the model ensures that these co-operators then form links with nearby nodes, thereby driving the formation of locally co-operative clusters which can convert nearby defectors. When l or E was increased, defectors could make use of the baseline energy to form clusters amongst themselves, thereby resisting the co-operator's influence.

These dynamics underscore the importance of endogenous edge formation to the outcome of the coevolutionary REDS model. Such dynamics were also pertinent to the cyclic example. Here, we found that weakly connected defectors are vulnerable to invasion by

co-operators while, conversely, strongly connected defectors can resist invasion and even turn the tables on co-operators. That being the case, defector nodes can be thought of as being in one of two states; having either low or high degree. Furthermore, nodes within a strongly connected defector cluster will, over time, drop their connections thereby becoming a region of isolated individual defectors. From this perspective it can be said that the low degree defector state dominates the high degree defector state.

These three states; co-operators, high degree defectors and low degree defectors; bear comparison to the states in the model of voluntary participation presented in [Hauert and Szabó \(2005\)](#). We recall that in this model a loner strategy was introduced that could dominate defectors but was itself vulnerable to co-operators. Thus, in Hauert and Szabó's model loners dominate defectors, defectors dominate co-operators and co-operators dominate loners, an arrangement of the strategies known as cyclic domination. When simulated upon a network, it is possible for all three strategies of the voluntary participation model to persist. Where coexistence occurs in this model it is because, for any given strategy, there are always nodes with that strategy that have vulnerable neighbours with the consequence that the strategy can perpetuate itself.

In the coevolutionary REDS model low degree defectors dominate high degree defectors, high degree defectors dominate co-operators and co-operators dominate low degree defectors. Crucially, we found, in the cyclic example, that co-operators could exploit transmission so as to reach regions of low degree defectors, thereby closing the cycle between the three states. Consequently, the cyclic model supports the survival of all three states and, in this sense, we argue that it is analogous to the cyclic domination of Hauert and Szabó.

We finish this chapter by highlighting a connection between the cyclic example and the model of [Szolnoki et al. \(2010\)](#). In section 2.3.2.1 we outlined that Szolnoki et al.'s contribution was to define a model whose dynamics gave rise to cyclic interactions in a population of agents playing a two strategy game. In particular, these cyclic dynamics were induced by specifying a period after a strategy update during which a node has a reduced capacity to adopt a new strategy. The analogy between this example and the coevolutionary REDS model is that a high degree defector node in the cyclic example is prevented from updating to the co-operative strategy until such time as it relinquishes enough connections so as to render it vulnerable to conversion by a co-operator.

Chapter 7

Conclusions

This thesis has sought to contribute to the understanding of how spatially embedding a network affects both its structure and dynamics that occur upon it. We have realised this goal with the following three contributions: (i) a model that produces spatially embedded networks that exhibit a power law degree distribution, (ii) improving the plausibility of the REDS model so that its agents' decisions to form connections are subject to uncertainty and (iii) a coevolutionary network model in which the agents' capacity to form connections are constrained by a privately held budget.

The first of these contributions, the planar growth model, created a spatial network that respected the constraint that its edges do not cross. The most notable outcome of this model was that its networks displayed a heterogeneous degree distribution. Crucially, planar growth neither relied upon topological features of the existing network when forming connections with newly added nodes nor assumed a power law in any of its inputs. These details distinguish the new model in that, as far as we are aware, planar growth is only the second example with nodes uniformly distributed upon the plane that results in the power law degree distribution.

The importance of the planarity constraint was investigated in a variation of the model where edges were allowed to cross with differing probability. These experiments demonstrated that the degree distribution smoothly changes from the power law into an exponential distribution in line with increasing probability of edges crossing. Similarly smooth variations were noted for clustering and assortativity. Thus, planar growth constitutes an initial foray into the investigation of spatial networks that are almost planar, a topic that we identified as being under researched in section 3.1. Finally, for this model, we redefined planar growth as a random Apollonian process. We developed this idea sufficiently so as to incorporate planar growth and two existing Apollonian models within a single framework.

We intend that the results of the PG model will shed light upon those situations where a spatial network has a planar structure. Examples we have given in this thesis are roads,

circuits and wireless networks. One might object that these real-world examples exhibit a homogeneous degree distribution. However, we draw attention to those experiments in chapter 4 that investigated the relaxation of the planarity constraint. Further research in this direction has the potential to bridge the gap between our findings and empirically observed networks.

We then considered the REDS model, a process which simulates the formation of a social network. The distinguishing feature of this model was that its edges incurred a cost proportional to their length which agents, i.e., the nodes, paid for from a personal budget. The REDS model exhibits several statistics that are consistent with those considered to be characteristic of real world social networks. The most significant of these is a degree distribution that is heterogeneous in the sense that it is bi-modal, i.e., the distribution exhibits two scales, with several high degree nodes and the remainder of the nodes having low degree. Our innovation here was to view decisions to connect as being subject to uncertainty. Furthermore, edges were now subject to re-evaluation and possible removal, with the consequence that the new model was a dynamic version of the static original. We implemented these ideas in the stochastic REDS model by modelling decisions using a Fermi function. This function accepted a single parameter, T_n , representing the uncertainty in the decision.

When the stochastic REDS model was run with low uncertainty we found that we could no longer recover the bi-modal distribution; the network instead equilibrated to either high or low connectivity. This result indicated that high connectivity regions of the network were unstable in the face of the nodes' continual revision of their edges. By increasing the uncertainty we found networks that exhibited the bi-modal degree distribution. However, in this case, each of the two node types were located in specific regions of the square; high connectivity nodes appeared on its border and low connectivity nodes in the centre. This result indicated that the inhomogeneity in the space, i.e., the rigid boundary of the square, promoted the onset of the hybrid distribution and that increased uncertainty in the decision contributed to the stability of the high connectivity region along the border.

As was outlined in the literature review, social networks are amongst the most widely studied empirical models. Furthermore, a growing body of evidence confirms the intuition that spatial structure is inherent in real world social networks, even where those networks are enhanced by communications technology. Nonetheless, the REDS model is, as far as we are aware, the sole example of a spatially embedded social network model that exhibits a form of degree heterogeneity. We found the REDS model encouraging but voiced concerns as to the static nature of the finalised network. We offer the perpetual review of connections in the stochastic REDS model as a more plausible simulation of a social network. That being the case, the conditions under which the bi-modal distribution occurs in the stochastic REDS model is a new insight into the phenomenon of heterogeneous real-world social networks.

Our final contribution was the coevolutionary REDS model, in which we extended the stochastic REDS model so that its nodes played the prisoner’s dilemma against their network neighbours. The payoff from the evolutionary game was added to a node’s budget when forming edges in the underlying stochastic REDS model. By this mechanism connectivity within the coevolutionary REDS model was established endogenously, in contrast to existing co-evolutionary models which pre-assume overall connectivity.

We compared the coevolutionary REDS model with simulations in which nodes played prisoner’s dilemma upon static networks generated using the stochastic REDS model. When we did so we found that the coevolutionary model better supported co-operative behaviour within the population. These results follow a common pattern within the literature on the evolution of co-operation in that numerous coevolutionary models demonstrate increased support for co-operation relative to an appropriate static model. However, pre-existing coevolutionary models rely upon agents that discern some feature of each game they play, e.g., a neighbour’s strategy. By contrast the coevolutionary REDS model’s agents did not have the capacity to assess their individual strategic interactions, instead their knowledge consisted solely of the accumulated payoff across all the games that they played. Thus, the coevolutionary REDS model demonstrates the possibility for a population of agents with minimal cognitive capacity to favour the co-operative strategy.

It was also observed that the coevolutionary REDS model can exhibit cyclic dynamics. Two features ensured the onset of this phenomenon; firstly, defector nodes, depending on their connectivity, can be considered to be in one of two states. High connectivity defectors can successfully invade co-operators and, conversely, low connectivity defectors are liable to invasion by co-operators. The second feature was that appropriate values of the update parameters allowed for the appearance of temporary connections in the model that could transmit the co-operative strategy from co-operative clusters to regions of low connectivity defectors, thereby precipitating an invasion of the defectors by the co-operators.

We expect that these results will contribute to the development of evolutionary theory with respect to co-operation. To be clear, we do not discount that the cognitive capacity of biological agents contributes to the persistence of co-operative behaviour. However, our results indicate that, to some extent, co-operation may be supported ‘for free’ within a biological population. There is no reason to believe that this finding stands in opposition to existing work since the essential mechanism at the heart of evolutionary theory is one of incremental change. Thus, we argue that it may well be the case that cognitively enhanced agents build upon the spatial basis for co-operation that we have outlined.

Our findings in this regard offer the possibility of practical suggestions for promoting co-operative behaviour in real world scenarios. For instance, in order to promote greater

co-operation within an organisation, an obvious recommendation, in light of our findings, would be to encourage known co-operators to be located together. Our results indicate that, over time, such a group will spread their culture to the remainder of the hypothetical organisation.

Taken together, our contributions underscore the capacity of spatially embedded network models to enrich the field of network science. The formation of a scale-free degree distribution and increased support for co-operation within a coevolutionary model, revisit issues that have both generated sizeable bodies of research. A common theme of the planar growth and coevolutionary REDS models is that they offer more parsimonious accounts of the phenomena of interest by exploiting spatial features not previously explored within the relevant literature. As regards our other contribution, the original REDS model was already a spatial model. The incorporation of uncertainty led to an open-ended, dynamic model, as opposed to the fixed end point and static outcome of the original. Investigation of the dynamics of the stochastic REDS model revealed subtleties in the conditions required to induce a bi-modal degree distribution in the network; one of the key results for the original REDS model.

These contributions open up new avenues for research. We have analysed the effects of a planarity constraint on the structure of a growing network. It appears of interest to investigate planarity constraints in conjunction with other network formation mechanisms. For instance, an area in network science that has found much attention in the literature are questions of optimal design of network structures ([Colizza et al., 2004](#); [Gastner and Newman, 2006](#); [Brede, 2010](#)). It would be of interest to further investigate to what extent planarity restrictions can constrain such optimal network topologies.

Chapter 5 established clear differences between the transition regions of the REDS model and the stochastic REDS model. It would be desirable to further elucidate the nature of these differences. It would be particularly insightful to establish if the high connectivity region observed in the high T_n example is observed in a model that has periodic boundary conditions. A further open question is a precise description of the dynamics of edge creation and removal within the transition regions of the low T_n examples.

The literature review discussed a form of coevolutionary model known as adaptive migration. In these models, nodes were placed on a lattice and, instead of rewiring, moved to unoccupied spaces during the network update phase of the model. A typical component of adaptive migration models is that agents are given an indication of the payoff they will receive from the new lattice location, i.e., these are further examples of models whose agents have significant cognitive ability. It seems clear that the concept of adaptive migration could readily be applied to the coevolutionary REDS model by allowing nodes the capacity to move within the plane. Furthermore, it should be straightforward to do so in a manner which retains the reduced cognitive capacities which are integral to the contribution that the coevolutionary REDS model makes.

Finally, to unify the work in this thesis, it would be of interest to apply the planarity constraint to both the stochastic and the evolutionary REDS models. In the first case intuition tells us that planarity should enhance the degree heterogeneity seen in the model. In the second case it seems plausible that planarity may assist the formation of co-operative communities, thereby enhancing co-operation overall.

To sum up, we reiterate our assertion that spatial network models offer a promising way in which to deepen the insights of network science and offer the results in this thesis as proof of this claim.

Bibliography

Agler, M. T., Ruhe, J., Kroll, S., Morhenn, C., Kim, S.-T., Weigel, D., and Kemen, E. M. Microbial hub taxa link host and abiotic factors to plant microbiome variation. *PLoS biology*, 14(1):e1002352, 2016.

Albenque, M., Marckert, J.-F., et al. Some families of increasing planar maps. *Electronic Journal of Probability*, 13(56):1624–1671, 2008.

Alstott, J., Bullmore, E., and Plenz, D. powerlaw: a python package for analysis of heavy-tailed distributions. *PloS one*, 9(1):e85777, 2014.

Amaral, L. A. N., Scala, A., Barthélemy, M., and Stanley, H. E. Classes of small-world networks. *Proceedings of the National Academy of Sciences, USA*, 97(21):11149–11152, 2000.

Andrade Jr, J. S., Herrmann, H. J., Andrade, R. F., and da Silva, L. R. Apollonian networks: Simultaneously scale-free, small world, euclidean, space filling, and with matching graphs. *Physical Review Letters*, 94(1):018702, 2005.

Antonioni, A., Bullock, S., and Tomassini, M. Reds: an energy-constrained spatial social network model. In *Artificial Life 14: International Conference on the Synthesis and Simulation of Living Systems*. MIT Press, 2014.

Antonioni, A. and Tomassini, M. Degree correlations in random geometric graphs. *Physical Review E*, 86(3):037101, 2012.

Antonioni, A. and Tomassini, M. A growing social network model in geographical space. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(9):093403, 2017.

Axelrod, R. and Hamilton, W. D. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981.

Barabási, A.-L. and Albert, R. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

Barabási, A.-L., Albert, R., and Jeong, H. Mean-field theory for scale-free random networks. *Physica A: Statistical Mechanics and its Applications*, 272(1):173–187, 1999.

Barnett, L., Di Paolo, E., and Bullock, S. Spatially embedded random networks. *Physical Review E*, 76(5):056115, 2007.

Barthélemy, M. Spatial networks. *Physics Reports*, 499(1):1–101, 2011.

Barthélemy, M., Bordin, P., Berestycki, H., and Gribaudi, M. Self-organization versus top-down planning in the evolution of a city. *Scientific reports*, 3(2153), 2013.

Belykh, I., Belykh, V., and Hasler, M. Synchronization in complex networks with blinking interactions. In *Physics and Control, 2005. Proceedings. 2005 International Conference*, pages 86–91. IEEE, 2005.

Ben-Avraham, D., F Rozenfeld, A., Cohen, R., and Havlin, S. Geographical embedding of scale-free networks. *Physica A: Statistical Mechanics and its Applications*, 330(1):107–116, 2003.

Blume, L. E. How noise matters. *Games and Economic Behavior*, 44(2):251–271, 2003.

Boguñá, M., Pastor-Satorras, R., Díaz-Guilera, A., and Arenas, A. Models of social networks based on social distance attachment. *Physical review E*, 70(5):056122, 2004.

Brede, M. Coordinated and uncoordinated optimization of networks. *Physical Review E*, 81(6):066104, 2010.

Brummitt, C. D., Huremović, K., Pin, P., Bonds, M. H., and Vega-Redondo, F. Contagious disruptions and complexity traps in economic development. *Nature Human Behaviour*, 1(9):665, 2017.

Buesser, P. and Tomassini, M. Evolution of cooperation on spatially embedded networks. *Physical Review E*, 86(6):066107, 2012.

Bullmore, E. and Sporns, O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, 2009.

Bullock, S., Barnett, L., and Di Paolo, E. A. Spatial embedding and the structure of complex networks. *Complexity*, 16(2):20–28, 2010.

Cairns, D., Fayed, M. M., and Mouftah, H. T. Revisiting planarity in position-based routing for wireless networks. In Zheng, J., Mitton, N., Li, J., and Lorenz, P., editors, *Ad Hoc Networks*, pages 87–102. Springer, 2013.

Caldarelli, G. *Scale-free networks: complex webs in nature and technology*. Oxford University Press, 2007.

Chan, S. H. Y., Donner, R. V., and Lämmer, S. Urban road networks, spatial networks with universal geometric features? *The European Physical Journal B-Condensed Matter and Complex Systems*, 84(4):563–577, 2011.

Clauset, A., Shalizi, C. R., and Newman, M. E. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

Cohen, R. and Havlin, S. Scale-free networks are ultrasmall. *Physical Review Letters*, 90(5):058701, 2003.

Colizza, V., Banavar, J. R., Maritan, A., and Rinaldo, A. Network structures from selection principles. *Physical review letters*, 92(19):198701, 2004.

Corson, F. Fluctuations and redundancy in optimal transport networks. *Physical Review Letters*, 104(4):048703, 2010.

Dall, J. and Christensen, M. Random geometric graphs. *Physical Review E*, 66(1):016121, 2002.

Doye, J. P. and Massen, C. P. Self-similar disk packings as model spatial scale-free networks. *Physical Review E*, 71(1):016128, 2005.

Duijn, P. A., Kashirin, V., and Sloot, P. M. The relative ineffectiveness of criminal network disruption. *Scientific reports*, 4:4238, 2014.

Fraser, O. N. and Bugnyar, T. Reciprocity of agonistic support in ravens. *Animal Behaviour*, 83(1):171–177, 2012.

Fu, F., Hauert, C., Nowak, M. A., and Wang, L. Reputation-based partner choice promotes cooperation in social networks. *Physical Review E*, 78(2):026117, 2008.

Fu, F., Wu, T., and Wang, L. Partner switching stabilizes cooperation in coevolutionary prisoners dilemma. *Physical Review E*, 79(3):036101, 2009.

Gastner, M. T. and Newman, M. E. Shape and efficiency in spatial distribution networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(01):P01015, 2006.

Gintis, H. *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton university press, 2000.

Gomes, C. M., Mundry, R., and Boesch, C. Long-term reciprocation of grooming in wild west african chimpanzees. *Proceedings of the Royal Society of London B: Biological Sciences*, 276(1657):699–706, 2009.

Gross, T. and Sayama, H. *Adaptive Networks*. Springer, 2009.

Gudmundsson, A. and Mohajeri, N. Entropy and order in urban street networks. *Scientific reports*, 3, 2013.

Hamilton, W. D. The genetical evolution of social behaviour. ii. *Journal of theoretical biology*, 7(1):17–52, 1964.

Hauert, C. and Doebeli, M. Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature*, 428(6983):643, 2004.

Hauert, C. and Szabó, G. Game theory and physics. *American Journal of Physics*, 73(5):405–414, 2005.

Hayashi, Y. A review of recent studies of geographical scale-free networks. *Information and Media Technologies*, 1(2):1136–1145, 2006.

Helbing, D. and Yu, W. The outbreak of cooperation among success-driven individuals under noisy conditions. *Proc. Natl. Acad. Sci. USA*, 106:3680, 2009.

Herrmann, C., Barthélémy, M., and Provero, P. Connectivity distribution of spatial networks. *Physical Review E*, 68(2):026128, 2003.

Hilbe, C., Nowak, M. A., and Sigmund, K. Evolution of extortion in iterated prisoners dilemma games. *Proceedings of the National Academy of Sciences*, 110(17):6913–6918, 2013.

Hofbauer, J. and Sigmund, K. *Evolutionary games and population dynamics*. Cambridge university press, 1998.

Huberman, B. A. and Glance, N. S. Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences*, 90(16):7716–7718, 1993.

Huson, M. L. and Sen, A. Broadcast scheduling algorithms for radio networks. In *Military Communications Conference, 1995. MILCOM'95, Conference Record, IEEE*, volume 2, pages 647–651. IEEE, 1995.

Hyndman, R. J. Cyclic and seasonal time series. <https://robjhyndman.com/hyndtsight/cyclists/>, 2011. Accessed: 2018-02-7.

Ichinose, G., Saito, M., Sayama, H., and Wilson, D. S. Adaptive long-range migration promotes cooperation under tempting conditions. *Scientific reports*, 3:2509, 2013.

Iotti, B., Antonioni, A., Bullock, S., Darabos, C., Tomassini, M., and Giacobini, M. Infection dynamics on spatial small-world network models. *Physical Review E*, 96(5):052316, 2017.

Jiang, L.-L., Wang, W.-X., Lai, Y.-C., and Wang, B.-H. Role of adaptive migration in promoting cooperation in spatial games. *Physical Review E*, 81(3):036108, 2010.

Jones, J. H. and Handcock, M. S. An assessment of preferential attachment as a mechanism for human sexual network formation. *Proceedings of the Royal Society of London B: Biological Sciences*, 270(1520):1123–1128, 2003.

Katifori, E., Szöllősi, G. J., and Magnasco, M. O. Damage and fluctuations induce loops in optimal transport networks. *Physical review letters*, 104(4):048704, 2010.

Khoo, T., Fu, F., and Pauls, S. Coevolution of cooperation and partner rewiring range in spatial social networks. *Scientific reports*, 6, 2016.

Kirchkamp, O. Spatial evolution of automata in the prisoners dilemma. *Journal of Economic Behavior & Organization*, 43(2):239–262, 2000.

Klaus, A., Yu, S., and Plenz, D. Statistical analyses support power law distributions found in neuronal avalanches. *PloS one*, 6(5):e19779, 2011.

Kolossváry, I., Komjáthy, J., and Vágó, L. Degrees and distances in random and evolving Apollonian networks. *arXiv preprint arXiv:1310.3864*, 2013.

Kumar, R., Novak, J., and Tomkins, A. Structure and evolution of online social networks. In *Link mining: models, algorithms, and applications*, pages 337–357. Springer, 2010.

Lambiotte, R., Blondel, V. D., De Kerchove, C., Huens, E., Prieur, C., Smoreda, Z., and Van Dooren, P. Geographical dispersal of mobile communication networks. *Physica A: Statistical Mechanics and its Applications*, 387(21):5317–5325, 2008.

Levinson, D. Network structure and city size. *PloS one*, 7(1):e29721, 2012.

Li, Y., Min, Y., Zhu, X., and Cao, J. Partner switching promotes cooperation among myopic agents on a geographical plane. *Physical Review E*, 87(2):022823, 2013.

Liljeros, F., Edling, C. R., Amaral, L. A. N., Stanley, H. E., and Åberg, Y. The web of human sexual contacts. *Nature*, 411(6840):907–908, 2001.

Lindgren, K. and Nordahl, M. G. Evolutionary dynamics of spatial games. *Physica D: Nonlinear Phenomena*, 75(1-3):292–309, 1994.

Lotker, Z. and Peleg, D. Structure and algorithms in the sinr wireless model. *ACM SIGACT News*, 41(2):74–84, 2010.

Manna, S. S. and Sen, P. Modulated scale-free network in euclidean space. *Physical Review E*, 66(6):066114, 2002.

Maslov, S., Sneppen, K., and Zaliznyak, A. Detection of topological patterns in complex networks: correlation profile of the internet. *Physica A: Statistical Mechanics and its Applications*, 333:529–540, 2004.

Masuda, N. Participation costs dismiss the advantage of heterogeneous networks in evolution of cooperation. *Proceedings of the Royal Society of London B: Biological Sciences*, 274(1620):1815–1821, 2007.

Masuda, N., Miwa, H., and Konno, N. Geographical threshold graphs with small-world and scale-free properties. *Physical Review E*, 71(3):036108, 2005.

McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

Miralles, A., Comellas, F., Chen, L., and Zhang, Z. Planar unclustered scale-free graphs as models for technological and biological networks. *Physica A: Statistical Mechanics and its Applications*, 389(9):1955–1964, 2010.

Mislove, A., Marcon, M., Gummadi, K. P., Druschel, P., and Bhattacharjee, B. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42. ACM, 2007.

Mok, D., Wellman, B., and Carrasco, J. Does distance matter in the age of the internet? *Urban Studies*, 47(13):2747–2783, 2010.

Moreno, J. L. Who shall survive?: A new approach to the problem of human interrelations. *Nervous and mental disease monograph series*, 58, 1934.

Mukherjee, G. and Manna, S. Weighted scale-free networks in euclidean space using local selection rule. *Physical Review E*, 74(3):036111, 2006.

Newman, M. *Networks: an introduction*. Oxford University Press, 2010.

Newman, M. E. Assortative mixing in networks. *Physical Review Letters*, 89(20):208701, 2002.

Newman, M. E. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.

Newman, M. E. Power laws, pareto distributions and zipf’s law. *Contemporary physics*, 46(5):323–351, 2005.

Newman, M. E. and Park, J. Why social networks are different from other types of networks. *Physical Review E*, 68(3):036122, 2003.

Nowak, M. A. Five rules for the evolution of cooperation. *science*, 314(5805):1560–1563, 2006.

Nowak, M. A., Bonhoeffer, S., and May, R. M. Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Sciences*, 91(11):4877–4881, 1994.

Nowak, M. A. and May, R. M. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.

Nowak, M. A. and Sigmund, K. Evolution of indirect reciprocity by image scoring. *Nature*, 393(6685):573, 1998.

Ohtsuki, H., Hauert, C., Lieberman, E., and Nowak, M. A. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505, 2006.

Park, J. and Newman, M. E. Origin of degree correlations in the internet and other networks. *Physical Review E*, 68(2):026112, 2003.

Perc, M. and Szolnoki, A. Coevolutionary games a mini review. *BioSystems*, 99(2):109–125, 2010.

Phithakkitnukoon, S., Smoreda, Z., and Olivier, P. Socio-geography of human mobility: A study using longitudinal mobile phone data. *PLoS one*, 7(6):e39253, 2012.

Qiao, J., Huang, H.-Q., Li, G.-Y., and Fan, Y. Bridging the gap between different social networks. *Physica A: Statistical Mechanics and its Applications*, 410:535–549, 2014.

Roca, C. P., Cuesta, J. A., and Sánchez, A. Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. *Physics of life reviews*, 6(4):208–249, 2009.

Rozenfeld, A. F., Cohen, R., Ben-Avraham, D., and Havlin, S. Scale-free networks on lattices. *Physical review letters*, 89(21):218701, 2002.

Rui, Y., Ban, Y., Wang, J., and Haas, J. Exploring the patterns and evolution of self-organized urban street networks through modeling. *The European Physical Journal B*, 86(3):1–8, 2013.

Santos, F., Rodrigues, J., and Pacheco, J. Graph topology plays a determinant role in the evolution of cooperation. *Proceedings of the Royal Society of London B: Biological Sciences*, 273(1582):51–55, 2006a.

Santos, F. C. and Pacheco, J. M. Scale-free networks provide a unifying framework for the emergence of cooperation. *Physical Review Letters*, 95(9):098104, 2005.

Santos, F. C., Pacheco, J. M., and Lenaerts, T. Cooperation prevails when individuals adjust their social ties. *PLoS computational biology*, 2(10):e140, 2006b.

Stella, M. and Brede, M. Patterns in the english language: phonological networks, percolation and assembly models. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(5):P05006, 2015.

Szabó, G. and Fath, G. Evolutionary games on graphs. *Physics reports*, 446(4):97–216, 2007.

Szabó, G. and Tóke, C. Evolutionary prisoners dilemma game on a square lattice. *Physical Review E*, 58(1):69, 1998.

Szolnoki, A., Mobilia, M., Jiang, L.-L., Szczesny, B., Rucklidge, A. M., and Perc, M. Cyclic dominance in evolutionary games: a review. *Journal of the Royal Society Interface*, 11(100):20140735, 2014.

Szolnoki, A. and Perc, M. Coevolution of teaching activity promotes cooperation. *New Journal of Physics*, 10(4):043036, 2008.

Szolnoki, A., Wang, Z., Wang, J., and Zhu, X. Dynamically generated cyclic dominance in spatial prisoners dilemma games. *Physical Review E*, 82(3):036110, 2010.

Tan, H., Peng, M., Tse, C. K., and Wu, F. Global similarity tests of physical designs of circuits: A complex network approach. *Applied Mathematics and Computation*, 230: 96–103, 2014.

Tang, C.-L., Wang, W.-X., Wu, X., and Wang, B.-H. Effects of average degree on cooperation in networked evolutionary game. *The European Physical Journal B-Condensed Matter and Complex Systems*, 53(3):411–415, 2006.

Taylor, P. D. and Jonker, L. B. Evolutionary stable strategies and game dynamics. *Mathematical biosciences*, 40(1-2):145–156, 1978.

Toivonen, R., Kovanen, L., Kivelä, M., Onnela, J.-P., Saramäki, J., and Kaski, K. A comparative study of social network models: Network evolution models and nodal attribute models. *Social Networks*, 31(4):240–254, 2009.

Toivonen, R., Onnela, J.-P., Saramäki, J., Hyvönen, J., and Kaski, K. A model for social networks. *Physica A: Statistical Mechanics and its Applications*, 371(2):851–860, 2006.

Traulsen, A. and Nowak, M. A. Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952–10955, 2006.

Traulsen, A., Nowak, M. A., and Pacheco, J. M. Stochastic dynamics of invasion and fixation. *Physical Review E*, 74(1):011909, 2006.

Trivers, R. L. The evolution of reciprocal altruism. *The Quarterly review of biology*, 46 (1):35–57, 1971.

Van Lange, P. A., Joireman, J., Parks, C. D., and Van Dijk, E. The psychology of social dilemmas: A review. *Organizational Behavior and Human Decision Processes*, 120 (2):125–141, 2013.

Wang, W.-X., Ren, J., Chen, G., and Wang, B.-H. Memory-based snowdrift game on networks. *Physical Review E*, 74(5):056113, 2006.

Wang, Z., Kokubo, S., Tanimoto, J., Fukuda, E., and Shigaki, K. Insight into the so-called spatial reciprocity. *Physical Review E*, 88(4):042145, 2013.

Watts, D. J. and Strogatz, S. H. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.

Wilkinson, G. S. Reciprocal food sharing in the vampire bat. *Nature*, 308(5955):181, 1984.

Wong, L. H., Pattison, P., and Robins, G. A spatial model for social networks. *Physica A: Statistical Mechanics and its Applications*, 360(1):99–120, 2006.

Yang, D.-P., Lin, H., Wu, C.-X., and Shuai, J. Topological conditions of scale-free networks for cooperation to evolve. *arXiv preprint arXiv:1106.5386*, 2011.

Yule, G. U. et al. Ii.a mathematical theory of evolution, based on the conclusions of dr. jc willis, fr s. *Phil. Trans. R. Soc. Lond. B*, 213(402-410):21–87, 1925.

Zhang, Z., Rong, L., and Zhou, S. Evolving Apollonian networks with small-world scale-free topologies. *Physical Review E*, 74(4):046105, 2006.

Zhou, T., Yan, G., and Wang, B.-H. Maximal planar networks with large clustering coefficient and power-law degree distribution. *Physical Review E*, 71(4):046141, 2005.

Zimmermann, M. G., Eguíluz, V. M., and San Miguel, M. Coevolution of dynamical states and interactions in dynamic networks. *Physical Review E*, 69(6):065102, 2004.

zu Erbach-Schoenberg, E., Bullock, S., and Brailsford, S. A model of spatially constrained social network dynamics. *Social Science Computer Review*, 32(3):373–392, 2014.